Play it Again: Partner Choice, Reputation Building and Learning in Restarting, Finitely-Repeated Dilemma Games

by Kenju Kamei[1] and Louis Putterman[2]

[1] Department of Economics, Bowling Green State University, Bowling Green, OH 43403, USA. Email: kenju.kamei@gmail.com.

[2] Department of Economics, Brown University, 64 Waterman Street, Providence, RI 02912, USA. Email: Louis_Putterman@brown.edu.

Abstract:

Previous research has shown that opportunities for two-sided partner choice in finitely repeated social dilemma games can promote cooperation through a combination of sorting and opportunistic signaling, with late period defections by selfish players causing an end-game decline. How such experience would affect play of subsequent finitely-repeated games remains unclear. In each of six treatments that vary the cooperation premium and the informational basis for reputation formation, we let sets of subjects play sequences of finitely-repeated voluntary contribution games to study the competing forces of (a) learning about the benefits of reputation, and (b) learning about backward unraveling. We find, *inter alia*, that with a high cooperation premium and good information, investment in reputation grows across sets of finitely-repeated games.

Keywords: cooperation, reputation, voluntary contribution, public goods, sorting, endogenous grouping, group formation, experiment

JEL classification codes: C92, D74, D83, H41

1. Introduction

Situations in which cooperation is potentially beneficial are ubiquitous in economic and social life. With self-interest the most powerful motivator for individuals in most settings, temptations to cheat or free-ride, and responses including attention to others' reputations and attempts to defect in advance of one's partner, are equally ubiquitous. Instances of successful cooperation are nonetheless not difficult to find.

A key factor that facilitates cooperation in some domains is the potential to choose who one deals with and the resulting competition for trustworthy partners. When individuals have multiple potential partners or multiple groups they could join, incentives to earn a reputation for trustworthiness or cooperation may be strong. To be sure, in the most austere theoretical setting in which all agents are perfectly selfish and rational, there is no point in seeking out one partner rather than another, for in one-shot or finitely repeated interactions all are equally certain to "defect." But when enough individuals believe that some derive utility from helping, cooperating, or being trustworthy, or believe that enough others hold such a belief, the incentive to invest in a favorable reputation may support considerable amounts of cooperation. Presence of some individuals having actual preferences for cooperation is helpful although unnecessary in any one situation, but it may be critical to the long-term survival of beliefs that such types exist.

Laboratory experiments have shown that the ability to choose their partners in social dilemma interactions can be a powerful force encouraging more cooperative choices when informational conditions permit individuals to learn the past actions of prospective partners (Page, Putterman and Unel 2005; Ahn, Isaac and Salmon 2008; Ahn, Isaac and Salmon 2009; Coricelli, Fehr and Fellner, 2004; Bayer 2011; Riedl, Rohde and Strobel 2011; Charness and Yang, 2010; Wang, Suri and Watts, 2012; Riedl and Ule, 2013). While a number of factors may contribute to this effect, two motivational channels stand out in our view. First, many subjects appear to have an intrinsic preference for cooperating provided that their counterparts do the same.[1] Because each enjoys a material benefit from interacting with a more cooperative partner, partner choice permits cooperative subjects to team up and achieve higher payoffs through sustained cooperation. Second, selfish subjects, too, materially benefit from playing with

[1] A related but slightly different idea is that individuals show indirect reciprocity towards an interaction partner known to have acted pro-socially towards others in the past. See Nowak and Sigmund (1998), Seinan and Schram (2006) and Engelmann and Fischbacher (2009).

cooperators, and those among them who are sufficiently sophisticated realize that the best way to do this is to establish a reputation for cooperativeness themselves. In an environment with partner choice, such selfish players will mimic cooperativeness for so long as it is more profitable than free-riding—i.e., for as long as future opportunities to profit from the resulting reputation are sufficiently numerous. During such periods, sophisticated selfish mimickers would be indifferent about whether their partner of the moment has cooperative preferences or is simply another selfish mimicker of cooperation. In a finitely repeated sequence of interactions, those with a taste for cooperation and their mimickers may sometimes be distinguishable by differences in final period play only (Page *et al*., hereafter PPU).

Two questions regarding investment in reputation in finitely repeated endogenous grouping situations remain understudied, we think. First, in a given finitely repeated experiment, subjects may only gradually learn the benefit of investing in reputation. If reputation accumulates without a clearing of the slate, individuals who were initially too pessimistic or unaware of the potential of a cooperative strategy may be constrained by initial decisions they would like to revise, if given a fresh opportunity. Our first question, then, is whether the learning attained in earlier finitely-repeated super-games might lead to more investment in reputation given opportunities to make a fresh start.[2] This question has real life relevance because people often interact with different sets of individuals at different points in their lives, due to change of job, geographic location, and so forth. More broadly, what individuals' experiences teach them may affect others by the advice they give, for instance to the next generation.

[2] To be sure, the difficulty of altering one's reputation is greatest in environments like PPU and Coricelli *et al*. in which subjects carry with them their entire average past level of cooperation throughout the finitely-repeated series of interactions. In Ahn *et al*. (2008, 2009), Charness and Yang (2010) and Bayer (2011), subjects know only each prospective partner's behavior in the most recent few periods, so it is possible to create an entirely new reputation over time. Our first question thus applies in its strongest form to the first-mentioned designs, only. In all but one of the experiments mentioned, however, there is never a period after the very first one that an individual enters free of any reputation, whereas there will be such reputation-free restarts in our new design. The exception is Charness and Yang, whose subjects play one finitely repeated 15 period super-game with endogenous group formation in a "society" of 9 individuals, then play another such game in another society of 9 with new IDs and possibly new members. Their explanation recognizes the possibility of long-term reputation effects even with explicit reporting of recent periods' contributions only, as indicated by their statement "We chose to have a re-start, as an individual might get locked into a situation that is difficult to escape during a segment, but instead would receive a fresh start in the second segment." However, analyzing the effects of the restart is not a focus of their paper. We come back to it in footnote 41, below.

Second, if individuals move from one finitely-repeated super-game to another (the example of a geographic move may again be pertinent), the learning they take with them is likely to include experience of end-game type behaviors. Shortly before picking up and moving, one's co-worker, carpenter, or painter may cease to be reliable, and one's most recent favors or payments may suddenly stop being reciprocated. Upon entering a subsequent interaction that can also be expected to terminate at some point, individuals may then take greater protective actions, including trying to pre-empt partners' defections by attempting to defect one step ahead of them.

The two factors just discussed suggest that learning may point in two countervailing directions. On the one hand, learning may point towards greater initial cooperation, if cooperators have been seen to be rewarded in past play. On the other, it may point towards earlier defection, if end-game free-riding has been found to be common in past play.

The competition between learning the benefits of reputation-building and learning about end-game unraveling of cooperation has been studied by Selten and Stoecker (1986) and Andreoni and Miller (1993) in the context of playing prisoners' dilemmas with exogenously assigned counterparts. What their studies lack, with respect to our topic, is the dimension of partner choice, which past research suggests may substantially change behaviors.[3] Parallels to Andreoni and Miller's (hereafter AM) design and findings are discussed at length, below.

We implement experimental treatments to study the strength of incentives to invest in reputation for cooperativeness under conditions of mutual partner selection in a social dilemma framework using two-person voluntary contribution games as our elemental building block. The distinctive feature of our design is that we combine endogenous partner selection with the successive playing of multiple finitely repeated super-games between which there is no continuity of individual reputation. In line with past results, we expected that the initial optimism of some subjects regarding the presence of genuine cooperators might be sustained through the early periods of a super-game, but we also expected subjects to take note if defection became common in the final periods. Although the "backward unraveling" predictions of

[3] Subjects in two treatments of Hauk and Nagel (2001), which follow Andreoni and Miller in most respects, can refuse to play with the partner they are exogenously matched with. Unlike those in PPU, Ahn *et al*., and the other endogenous partner choice experiments mentioned above, Hauk and Nagel's subjects cannot compete for or seek out new partners of their own choice.

traditional economic models are rarely supported by the play of naïve subjects in the experimental lab, repeating a set of finitely repeated super-games (called phases, in the experiment) might lead to earlier and earlier appearance of widespread defections. But alongside this force making for potentially earlier unraveling of cooperation, recognition of the early period gains of cooperators by participants who had been less optimistic or less far-sighted in the initial super-game(s) could make for increasing initial cooperation. Studying these countervailing forces while obtaining additional observations of partner choice and further evidence of the presence of both cooperatively-oriented and strategically mimicking agents, is the goal of our study.

Because selfish subjects' incentives to engage in initial cooperative play depend on the gains from cooperation, our experiment varies across treatments the size of those gains—a potential 30% in our low gains versus 70% in our high gains treatments. And because the likelihood that greater current cooperation will lead to better future partner options depends on the conditions of information transmission, we vary the completeness of information through which reputations can be formed within a given finitely-repeated super-game, from a full (100%) report of each others' within-phase past average contribution, to reports based on a randomly chosen 50% of within-phase past contributions, to no (0%) information apart from potential recall of own partners' choices (enabled by within-phase fixity of subject IDs).

We find considerable evidence of subject investment in reputation and of preference for more cooperative partners. We find that returns from investing in a reputation for cooperativeness depend as predicted on our treatment variables and that, accordingly, there is substantially more cooperation as well as a tendency for cooperation to be higher in later phases in those treatments with higher returns from cooperation and better information. Cooperation levels evolve as subjects learn whether cooperation pays. Even when no information is provided about those with whom subjects have not interacted yet, some succeed in establishing fruitful partnerships. Indications of earlier unraveling of cooperation in later phases are also present, but fail to outweigh growing cooperation over four ten-period super-games when the cooperation premium is high and information good. Our evidence is also consistent with the presence of at least a small number of true conditional cooperators, whose presence helps sustain beliefs that partners might vary in type. Our data thus suggest that the beneficial effects of competition for partners can survive in a world in which people learn from playing multiple super-games as they transition from job to job, place to place, or among different spheres of interaction.

The remainder of our paper proceeds as follows. Section 2 provides a discussion of past literature which sets the stage for our study. In Section 3, we spell out our experimental design and discuss theoretical predictions. Section 4 presents our results and analysis. Section 5 concludes the paper.

2. Literature

Economists have devoted enormous attention to studying games such as the prisoners' dilemma, the voluntary contribution or public goods game, and the investment or trust game, because real world situations in which pairs or groups of individuals can mutually benefit by cooperating, but risk being attracted away from a jointly optimal outcome by the logic of self-interest, are widespread. At the macroeconomic level, the dilemma of providing public goods is viewed as essentially unsolvable without the coercive intervention of the state and its powers of taxation. But members of workgroups, partnerships, private voluntary organizations, and even those engaging in some kinds of exchange for which formal contracting and litigation are too costly or impracticable, have no corresponding remedy available and rely mainly on voluntary cooperation to achieve mutually beneficial outcomes.

When information is good and interactions are indefinitely ongoing, equilibria of cooperation are potentially available to uniformly selfish and rational individuals (e.g. Kandori 1992). But some interactions are repeated only a finite number of times. Consider a partnership or voluntary association aimed at producing a specific product, for instance the collaborators on an academic paper, or members of a community group which has decided to build a playground. Each might be induced to do her part in the initial stages by the reasonable expectation that others will likewise pitch in, but in the final stages some may begin to shirk if convinced that the others will make up for their declining effort.

A way out of a finitely repeated or one-shot social dilemma is sometimes provided by the presence of actors whose preferences aren't strictly selfish. Among the widely discussed possibilities are altruism (the others' outcomes receive unconditional positive weight in own utility), reciprocity (one gets utility from reciprocating the others' beneficial actions), inequality aversion (one prefers not to gain at the expense of others), and warm glow (one gets utility from "doing the right thing"). The potential benefits of such preferences both to the individuals concerned and to those with whom they will interact helps to explain substantial familial and societal investments in socializing children so as to imbue or strengthen these orientations. The mere recognition that others have or may have such preferences can at times spell the difference between another actor's cooperation or defection. Kreps, Milgrom, Roberts and Wilson (1982) famously demonstrated that the attachment of small probabilities to the counterpart having a non-selfish (or non-standard) preference can suffice to yield many periods of cooperation by rational, selfish agents playing a finitely repeated game.

The tit-for-tat strategy to which any non-standard actors who might exist are assumed to be committed, in Kreps *et al*., resembles the preference type called "conditional cooperator" in the literature on voluntary contribution mechanisms (hereafter VCM; see Fischbacher, Gächter and Fehr, 2001, Fischbacher and Gächter, 2010).[4] The tit-for tat player's actions are ones of reciprocation, except that a specific initial move is prescribed. In the VCM, in contrast, conditional cooperation has been thought of as allowing each decision including the first to be conditional on beliefs about what others will do. We understand a conditional cooperator to be an agent who prefers contributing as much to the public good as his counterparts,[5] but we can nonetheless imagine two equally conditionally cooperative individuals who contribute different amounts due to having different beliefs about what their counterparts are simultaneously contributing. To predict such an agent's play in a repeated VCM, it's accordingly important to know what information she has regarding counterparts' past behaviors and how she uses that information to form expectations of their next decisions.

Fischbacher *et al*. (2001) used a strategy-based protocol to elicit their initial evidence of conditional willingness to contribute. The predictive power of the decision schedule for

[4] The preference function of the conditional cooperator may be thought of as causing the material payoffs of a prisoners' dilemma to be have "assurance game" (Sen, 1967) or stag hunt game payoffs in utility (or psychological) terms. Alternatively, in our appendix we note that conditionally cooperative behaviors can also be rationalized by inequality averse preference along the lines of Fehr and Schmidt (1999).

[5] While Fischbacher and Gächter (2010) find that the average conditional cooperator less than fully matches her counterparts' contributions, we simplify discussion in the remainder of our paper by ignoring variation in the degree of completeness of reciprocity.

unconditional contributions is well demonstrated by Fischbacher and Gächter (2010). Strong evidence of conditional cooperation also comes from studies in which the experimenters sort subjects by contribution level without their knowledge and find that among high contributors, the rapid decay of contributions typical in unsorted groups is replaced by relatively sustained contributions (Gunnthorsdottir, Hauser and McCabe, 2007; Gächter and Thöni, 2005). This suggests that high contributors tend to reduce their contributions in the unsorted experiments not because they are learning to play the game, in a general sense, but because they find that their contributions are not adequately reciprocated.

An alternative to the assignment of subjects to groups of apparently like-minded type by an experimenter is to let them sort themselves if so inclined. One kind of endogenous sorting involves subject choice not of specific partners but of an institutional or parameter setting (e.g. Aimone, Iannacone and Makowsky, forthcoming).[6] Our focus, however, is on conditions in which subjects indicate preference for specific partners. For the VCM, such experiments begin with Ehrhart and Keser (1999). We focus initially, however, on PPU, one of the first to demonstrate a significant impact of endogenous partner choice on efficiency. Subjects in their endogenous regrouping treatment first played three periods of a standard VCM game in exogenously formed partner groups of four, then were shown the average contribution thus far by those in each of their session's four groups (hence 15 other individuals) and assigned as many of the latter as they wished to ranks indicating priority of preference for co-membership in their group for the following three period phase. For example, an individual could assign ranks 1 through 15 to other session participants, with 1 indicating the participant she most wanted to have in her next group of four. Subjects were informed that the computer would identify the four participants whose mutual sums of ranks were lowest, form a first group accordingly, then repeat the process until only four individuals remained and became a group by default. Groups were thus to be formed, and were reformed every three periods, by mutual preference. The procedure yielded results resembling those of the entirely different exogenous grouping procedure of

[6] In their "sacrifice" condition, Aimone *et al.*'s subjects each select a return from their private account from options 0.55, 0.60, …, 0.95 (equivalently, choose a penalty or sacrifice for private allocations ranging from 0.05 to 0.45) knowing that they will then play, with three others who chose the most similar levels of sacrifice, a four person linear VCM with *mpcr* = 0.4. This allows cooperators to signal intent to cooperate by choosing a large sacrifice and playing with like-minded others. The mechanism successfully raises cooperation despite leaving the pecuniary return from private allocation higher than that from allocation to the group account. Other examples of self-sorting by choice of institution or parameter include Gürerk, Irlenbusch and Rockenbach (2006) and Lazear, Malmendier and Weber (2012).

Gunnthorsdottir *et al*. in that contributors of similar amounts ended up grouped together and the typically rapid decay of contributions on the parts of higher contributors was not observed. However, because subjects in PPU, unlike those of Gunnthorsdottir *et al*., had full information about the grouping procedure to be followed, ones caring only about their own payoffs may still have ended up in relatively cooperative groups out of strategic calculation that they could earn more by investing in cooperative reputations.

Differences of design help to explain differences of outcome among the endogenous grouping experiments mentioned earlier. One dimension of differentiation is whether groups are of fixed size (four in PPU, two in Coricelli *et al.* [2004], Bayer [2011] and the present paper) or of variable size (Ehrhart and Keser [2009], Ahn *et al*. [2008, 2009], Charness and Yang [2010]), in which case variation with group size of the social and private returns from cooperation is often a focus in its own right. Another is whether both sides of each match have a say about whether they play together (PPU, Bayer, this paper, some treatments in Coricelli *et al*. and Ahn *et al*.) or individuals can join others' groups at will (Ehrhart and Keser, some treatments in Ahn *et al*.). Details of the matching process including its costliness to subjects also differ—for instance, a 2nd price auction is used in Coricelli *et al*., a Gale-Shapley stable marriage mechanism difficult to fully explain to participants in Bayer, majority or plurality voting in Ahn *et al*. and Charness and Yang, and a simple ranking and group assignment mechanism shared by PPU and the present paper, except that PPU required subjects to pay a small price to submit ranks so as to demonstrate subject beliefs that others might have persistent differences in type or at least strategy.[7] Charness and Yang's groups also have the power to expel members by vote, a feature shared with designs more focused on the expulsion opportunity including Cinyabuguma, Page and Putterman (2005) and Maier-Rigaud, Martinsson and Staffiero (2010).

Of particular importance here is the nature of the information subjects receive and by which they can form reputations. The prerequisite for investment in reputation, met by all of the experiments mentioned, is that subjects learn of at least the most recent contribution decisions of prospective partners. But the persistence of reputation differs, with subjects in experiments including Ahn *et al*. (2008, 2009), Bayer (2011), and Charness and Yang (2010) learning one

[7] Because PPU's subjects were free to submit any number of ranks, including none, traditional theory with common knowledge of own payoff maximizing type predicts that there will be no expenditure on ranking. To simplify our design somewhat, we eliminate this cost and require each subject to give a unique rank to each potential partner.

another's average contributions in the most recent periods only, those in PPU and Coricelli *et al*. (2004) seeing the full past average contribution (for the session as a whole, thus far) of each potential partner. In almost none of these environments—the exception being a restart for a second 15 periods in Charness and Yang—does a subject enter any period other than the first one free of reputation.

PPU note that individuals with a genuine preference for cooperating and high expectations of cooperation by others could have been highly ranked as partners, and thus have been grouped together with, strictly selfish (own payoff maximizing) participants who "mimicked" the cooperator types because they correctly anticipated that this would allow them to earn more. Conditional cooperators and those mimicking them might therefore be indistinguishable until the final period of play, and even then only a lower bound estimate of the number of conditional cooperators could be formed, since while every selfish mimicker of cooperation would by definition contribute zero in the last period, true conditional cooperators might also contribute zero if they assigned very low probability to the presence of other true conditional cooperators—indeed, of conditional cooperators having favorable estimates of their proportional representation and of one another's higher-order beliefs.

PPU carefully report last period behaviors because of their potential value as evidence of the presence of a 'non-payoff-maximizing' type. In groups composed of the mutually selected highest quarter of average contributors in their sessions, they find that 50% of subjects contributed their entire endowment (and more than 10% contributed other positive amounts) in the known last period. In groups of next-highest contributors, 43% gave their full endowment (and another 34% other positive amounts). Although their final contributions were on average lower, even groups the histories of whose members put them in the third quartile in their sessions had only 20% of subjects contributing zero in the final period, whereas in a baseline treatment without endogenous group formation almost 88% of subjects contributed zero in the last period. The evidence from the endogenously grouped subjects thus suggests that the subject pool contained somewhat more subjects with some form of taste for cooperation than mimickers of that type.

One thing all of the endogenous grouping papers have in common is that all demonstrate that given a choice of partners, nearly everyone prefers a high contributor to a low one. That subjects in many of the experiments are willing to incur monetary costs to be paired with or enter the groups of higher contributors suggests widespread belief that individuals are not all the same, and that an orientation or type can be signaled in a manner credible to the receivers despite self-interested incentives to free ride regardless of past behavior. Whether partner choice itself encourages greater cooperation seems to depend in part on whether one needs to "earn" one's partners by past behavior, as is true in treatments like PPU's or Ahn *et al*.'s restricted entry treatments in which the partners one seeks out have a say in the matter but not in treatments like Ehrhart and Keser's and Ahn *et al*.'s restricted exit treatments in which free-riders can freely join groups of relative cooperators.[8]

The idea that competition for better partners is a key factor behind cooperative outcomes in VCM experiments with partner choice including PPU and Bayer is also consistent with findings of a new experiment by Riedl and Ule (in process). In their design, subjects play a series of 60 prisoners' dilemma games with changing partners who lack fixed IDs and whose most recent decision to cooperate or defect only is made known to the counterpart. In one treatment, subjects have no choice but to play with the randomly assigned counterpart, in a second they are assigned a counterpart but can decide whether or not to play, and in a third, subjects are put into randomly formed triplets in which each is offered two potential counterparts each of whom they must rate as acceptable or not. In this last treatment, subjects can thus opt out of playing entirely or choose one of the two for interaction. By generating competition for partners, the third treatment leads to a cooperation rate rising to over 50%, whereas cooperation is declining and averages around 10% in the other two treatments.

While the evidence from past experiments suggests that the ability to choose interaction partners can indeed promote cooperation in finitely repeated social dilemmas both by allowing cooperators to sort and play with one another and by encouraging more selfish types to mimic cooperators, the two questions raised in our introduction remain to be studied. First, subjects in the regrouping treatments of PPU and Coricelli *et al*. accumulated their reputations over the entire course of their session. This means that an initially low contributor, whether a cooperator

[8] An exception might be cases in which when a self-regarding subject "captures" a partner who has no alternative for a series of interactions, they might both cooperate as the best remaining option—which might explain why one sided choice is effective in Hauk and Nagel (where one-sided opt out rather than partner choice proper gets good results). Whether a similar phenomenon accounts for Coricelli *et al*.'s better one- than two-sided result is unclear to us.

with initially pessimistic beliefs about others or a selfish player having similar beliefs or not yet recognizing the potential profitability of cooperating, could find it difficult and costly to build the reputation adequate for breaking into a group of already high contributors. Restarting the process might therefore lead some subjects to adopt more cooperative strategies from the outset, but this remains to be investigated, because most previous studies on partner choice lack treatments that start an entirely new finitely-repeated game after the first one ends.[9]

Second, even when high cooperation is achieved in most periods, there are usually end-game declines, the impact of which on play of other finitely repeated games is also understudied. In PPU, for example, despite the existence of very high contributions even in the last period in the most cooperative groups, the average contribution did drop considerably as the known last period approached—by about 30% in the final period itself. The present paper seeks to shed light on a question not answered there: If there had been an unannounced restart, or if the subjects had later participated in another experiment of the same design, would their observations of this end-game effect make them less trusting of others from an earlier point in the process? Could the cooperation-bolstering force implicit in the previous paragraph's reasoning survive such negative effects of the end-game on behavior?

While we know of no attempts to pursue these questions in a social dilemma environment involving partner choice, Andreoni and Miller (1993, hereafter AM) address similar questions regarding the potentially countervailing processes of increasing investment in reputation and progressive end-game anticipation in prisoners' dilemma games with exogenously assigned partners.[10] AM's subjects are informed from the outset that they will play twenty ten-period games. In the core treatment (see also Selten and Stoecker, 1986), each period of a given ten-period game is played with the same randomly chosen partner, which permits each member of the pair to build a reputation as a cooperator even if his intention is to defect before his counterpart does. The subsequent ten-period game is played with a new randomly-assigned partner, so reputation cannot carry over, but learning from previous games is possible. Much as in our discussion above and in the design we propose in the next section, this set-up makes it

[9] As mentioned, Charness and Yang (2010) are the exception, but the effects of the restart are not a significant focus of their analysis.

[10] In a design closely resembling AM's, Hauk and Nagel (2001) give subjects the option to choose not to play with a counterpart and to be assured a 0 payoff, which is preferable to the negative mutual defection and sucker's payoffs. But there is no possibility of seeking out partners based on their past histories.

possible for subjects to learn the benefits of early cooperation in a given finitely repeated super-game, while simultaneously building adeptness at forecasting counterparts' end-game defections. AM find that in their setting investment in reputation-building grows across super-games. While end-game defections initially become earlier on average, AM find them to settle into a relatively unchanging pattern after the first few finitely-repeated games. It is unclear *a priori* how the play of multiple finitely-repeated dilemma games with partner choice will compare to play by subjects in AM who cannot select their partners. On the one hand, competition for the most cooperative partners might lead to still greater reputation-building efforts; on the other hand, the possibility of switching to other partners mid-way through a super-game could in some cases encourage earlier defection under some informational conditions.[11] The greater menu of options in the VCM than the PD might also affect play. Thus, while AM provide a clear reference case for play of a social dilemma game without partner choice, investigating the parallel problem in the context of partner choice calls for new research.

3. Experimental Design and Predictions

3.1. Experimental Design

In our experiment, subjects belonging to sets of ten anonymous and randomly selected participants who remain together throughout their session, play four distinct 10-period sequences of VCM stage games, with no carry-over of reputation from one set to any other. Each sequence constitutes a finitely-repeated super-game in the terminology used above, and is called a "phase" in the instructions. We accordingly get to study whether end-game effects have future consequences, and our subjects are freer to change strategies, for example investing heavily in a cooperative reputation in a later phase although they failed to do so in an earlier one, or the reverse. The stage game group size is two, with potentially new partners assigned each period based on submission of rankings and pairing by mutual preference. Subject i's earnings in period t are given by:

[11] In our experimental treatments providing reputational information, subjects' behaviors are known to others during a super-game according to their past average contribution, with information not broken down into contributions of specific periods. A subject could accordingly contribute her full endowment for, say, three periods, then contribute nothing in anticipation of switching partners, and still have a high enough past average contribution to attract a somewhat cooperative new partner, if there is enough diversity of contributions within the relevant subject set.

$$E - C_{it} + mpcr \cdot \sum_{j=1}^{2} C_{jt}, \qquad (1)$$

where E is the uniform per-period endowment, which we set at 10 points, C_{it} is the contribution of subject i in period t, and *mpcr*, the marginal per capita return from allocations to a pair's joint account, can also be represented as F/n, with F being the factor by which returns under full cooperation exceed those under full free-riding and n being the group size, here 2.[12] We consider three treatments in which $F = 1.3$, dubbed "1.3 treatments," and three in which $F = 1.7$, dubbed "1.7 treatments," thus emphasizing by our terminology the potential 30% versus 70% gains from cooperation. Depending on treatment, then, *mpcr* equals either 0.65 or 0.85. At the end of each period, a subject is informed of her partner's contribution decisions and her own earnings. Under each F and *mpcr*, we vary the degree to which reputation stays with a subject during a phase, as explained presently.

Ranking procedure and information conditions

At the beginning of each period, each subject is given the opportunity to choose (or influence the assignment of) her partner through ranking. We adopt a simple procedure that requires each subject to rank five potential partners each period and that always assigns two subjects, i and j, to be one another's partners if j is most preferred by i among the candidates i is offered, and conversely. Subjects are offered a random subset of five of their ten set members as prospective partners each period thanks to a fresh random division of the set into two sub-sets of five.[13] During the ranking stage, each subject sees the Subject IDs (fixed for 10 period phase) of her 5 potential partners, and in four of the treatments, also information on their allocations to their joint accounts in past periods of the phase, and gives each a rank of 1, 2, …, or 5, with 1 indicating most and 5 least preferred counterpart. As explained to the subjects, the computer searches among the 25 possibilities for the pair with the lowest sum of ranks, breaking ties randomly, then searches for a next pairing of the remaining 4 subjects in each group using the

[12] One point exchanged for $0.045 (4.5 cents) at the end of the experiment. Hence, universal non-contribution would cause subjects to earn 40*10*$0.045 = $18 (plus a $5 show-up fee), universal full contributions would generate earnings of $23.40 (in 1.3 treatments) or $30.60 (in 1.7 treatments) plus show-up fee. In the event, overall earnings averaged $22.87, or with show-up fee $27.87, slightly past the mid-way point between average expected earnings with no and those with full cooperation.

[13] These divisions into sub-sets were adopted partly to speed ranking, since complete rankings could be decided on more quickly for five than for nine others. For our treatments with public reputation they also have the benefit of raising reputation's potential importance by making reliance on establishing a partnership with a given individual relatively infeasible.

same procedure, thereby assigning each set member to one of 5 pairs.[14] Subjects learn the ID number of the partner then assigned to them, and in the 100% (and 50%) treatments, are shown the past average contribution (of selected periods, in 50%) of that partner on the screen in which they indicate their own new contribution decision.

In addition to cooperation return *F*, we also vary across treatments the reputational information subjects have access to when they rank their prospective partners. In one treatment using each *F*, our "100%" treatments, subjects at each ranking stage see each of their prospective counterparts' average allocation to his or her past joint accounts in all periods of the phase thus far. In a second treatment with each *F*, called "50%," they see the average allocation in past periods of which each is randomly selected with probability 0.5, and are also shown how many past periods entered the averaging calculation. In the third treatment using each *F*, our "0%" treatments, subjects are shown no information on past allocations but can use subject IDs and recollection of their own interactions to inform their ranking.[15] The crossing of the two *F* values (1.3 and 1.7) with the three information conditions (100%, 50%, and 0%) yields six treatments, which we refer to as the 0%, 1.3 treatment, the 50%, 1.7 treatment, etc., as displayed in Table 1.

3.2. Predictions

If all participants care only about maximizing their own payoffs and if they have common knowledge that all are and know one another to be of that type, the prediction is the same for all six treatments: universal non-contribution to the joint accounts. With no reason to favor one

[14] The algorithm translating ranks into partner assignments is identical to that in PPU. We adopted this mutual ranking procedure for its simplicity and its potential to encourage contributions to the public accounts thanks to each subject's easy recognition of the desirability of playing with higher-contributing partners. The procedure avoids the "fleeing from free riders" feature of Ehrhart and Keser (1999), is easier to explain to subjects than the Gale-Shapley algorithm used by Bayer (2011), is arguably easier for subjects to understand than the auction device used by Coricelli *et al*., and is more symmetric and requires fewer steps than alternatives used by, for instance, Ahn *et al*. (2008) and Charness and Yang (2010). Whether the strategic complications from which it might theoretically suffer were problems in practice will be addressed below, where we'll find that by and large its operation was straightforward and quite effective.

[15] Subjects are required to submit ranks 1 through 5 even if entirely lacking relevant information, as is always the case in each phase's first period. Instructions in all treatments stated that subjects could take notes, presumably of especially valuable use in the 0% treatments. Note that while a subject, say ID 3, might wish to play again with another, say ID 6, based on favorable experience, random splitting into sub-groups of 5 each period produces a 4/9 chance that the desired partner is unavailable in any given period of the phase. In 50% treatments, the computer randomly determined at the end of each period whether each subject's allocation decision will be included in her later-displayed past average and informs her of the outcome. These decisions are not revisited; for example, if ID 3's period 2 decision is included in her past average record for the phase, as shown in period 3, her period 2 decision is also included in past averages shown in periods 4, 5, etc. Each subject and period's draw is independent.

prospective partner over another, ranks will be assigned randomly, and presence of ranking and pair assignment procedures will have no impact on play of the VCM stage games.

However, decades of experimental studies, including both experiments with opportunities for reputation building (e.g., AM) and ones giving subjects a say in partner choice (discussed above) indicate that the assumptions just applied are unlikely to yield good predictions. These studies suggest behaviors more consistent with models like that of Kreps *et al*. in which the decision-makers attach positive probability to the presence of some actors who do not simply maximize their own payoffs but in some conditions prefer altruistic or cooperative actions. In our context, we'll assume the 'non-standard' action to be reciprocating the expected contribution of the partner. Even if no actor is an actual conditional cooperator, it can be shown that if the payoff from mutual cooperation exceeds that from mutual free riding by a sufficient amount, if the benefit of one-time defection is not too great, and if sufficiently high probability is attached to others being conditionally cooperative, it can be selfishly rational to contribute to the joint account for a certain number of periods of finitely repeated play.

All of the above applies also to situations such as that of AM in which agents are matched for interaction with a sequence of different counterparts. How are the dynamics concerned affected by partner choice? Essentially, the potential to switch partners gives any individual wishing to encourage cooperation from her counterpart an additional source of leverage, since the individual can now implicitly threaten not simply to reduce her contributions if the partner is not sufficiently reciprocating, but to exit the relationship and to enter into one with a better partner. More importantly, whereas the recourse of withdrawing cooperation is likely to be the best response when faced with an uncooperative partner with whom a sequence of interactions must play itself out due to lack of other alternatives, the possibility of obtaining a better partner provides reason to continue to make large contributions so as to increase or maintain one's attractiveness in the market for partnerships. However, if contributions vary enough among the set of subjects in question, there might also be a temptation to "defect" on a thus-far cooperating partner mid-way through a phase. This is especially the case when the way reputation is displayed is based on averages, because the impact on past average contribution may be small enough not to rule out still-cooperative partnerships in remaining periods.[16]

In our Appendix, we provide a partial equilibrium model which shows that if all agents belong to one of two types—selfish payoff-maximizers, and conditional cooperators—then in any given phase (finitely-repeated super-game), there is a finite number of periods, k, over which the selfish players find it rational to contribute to their joint accounts at the same level as a conditional cooperator, after which they switch to full free riding. For the 100% history condition, we show k to be increasing in the proportion p of conditional cooperators and in the *mpcr*. For our finitely repeated super-games of 10 periods, it is self-evident that for any $0 < p < 1$ and any *mpcr* < 1, $k \leq 9$. If all players have identical beliefs and degrees of strategic sophistication, then all payoff-maximizers interacting at given *mpcr* and information condition will adopt identical strategies and therefore their behaviors and those of the conditional cooperators will be indistinguishable up to period k, thus providing no basis for choosing one partner over another.[17] In practice, however, individuals' contributions are likely to differ from one another, which may be explained *inter alia* by different beliefs about the value of p and different levels of strategic sophistication. This differentiation will lead to meaningful partner preferences and to a significant impact of the partner assignment mechanism, including the tendency for individuals having closely similar contribution profiles to be paired.[18]

Both a selfishly rational individual and a conditional cooperator, who also gets positive utility from own earnings, will prefer to interact with more cooperative partners. We thus expect to see a preference for more cooperative partners being expressed in the rankings, and we expect

---

[16] If almost all set members have been contributing their full endowments thus far, the impact of a one-time defection on one's reputation in a 100% treatment might effectively prevent one from obtaining good matches for the remainder of the phase. However, the impact can be more minor if contribution levels are relatively diffuse. Indeed, in the 50% treatments, there is a 50/50 chance that a given defection episode will be known only to the current partner.

[17] In period $k + 1$, the payoff-maximizers would at last distinguish themselves from the conditional cooperators, after which conditional cooperators will be able to favor one another as partners in period $k + 2$ and beyond (assuming additional periods remain, i.e. assuming $k \leq 8$). The Appendix can be viewed at http://www.brown.edu/Departments/Economics/Papers/2013/Appendix_2013_8.pdf.

[18] The basic version of our model considers a selfish player who calculates the optimal number of periods to mimic cooperation under the assumption that only actual conditional cooperators reciprocate their contributions. The alternative assumption that all rational selfish players will mimic cooperation for so long as this is profitable would predict a somewhat higher k. We believe that neither approach is fully realistic, because individuals differ in their degrees of strategic sophistication, causing the degree to which cooperation is mimicked to vary not only with own but also with beliefs regarding others' degrees of sophistication. Our model's qualitative conclusions should hold for a range of such adjustments.

the workings of our ranking mechanism to cause higher contributors to be matched with one another.[19] Since, in the 50% and 100% treatments, subjects see a sample of others' average contributions thus far in the phase at the beginning of each period (i.e., in the ranking stage, from $t = 2$ onwards), and since subjects in those treatments are shown in the same periods' contribution stages the (in 50% treatments, reported) past average contribution of the partner who gets assigned to them, they can infer whether they are lately playing with relatively high contributors, relatively low ones, or ones in between. With the help of such inferences, most subjects are likely to be able to infer, by the end of the first phase, that contributing more is a good way to increase one's chance to interact with a high-contributing partner. If the expected impact of the extra point one contributes to the joint account on the contribution of one's future partner—whose identity is to be determined under the influence of this decision—is high enough, it thus becomes profitable to contribute more, despite an *mpcr* < 1. Simplifying by considering only that return from this period's investment in reputation which is realized in the period immediately following, we see that to compensate for each extra point contributed, a subject would need to anticipate an extra contribution ∂C_j by the next partner per additional point of own contribution such that $mpcr * \partial C_j \geq (1 - mpcr)$. For $mpcr = 0.85$ (1.7 treatments), the requirement is $\partial C_{j,t+1}/\partial C_{i,t} \geq (1 - .85)/.85 \approx .176$, while for $mpcr = 0.65$ (1.3 treatments), it is $\partial C_{j,t+1}/\partial C_{i,t} \geq (1 - .65)/.65 \approx .538$. The much higher hurdle for the 1.3 treatments provides intuition for the prediction above that contributions are less sustainable with the lower *mpcr*.

The Appendix provides a more complete accounting, also considering benefits from potentially higher contributions by partners in periods beyond $t + 1$. Using this more comprehensive approach, we confirm that the k that optimizes earnings in 1.7 treatments (*mpcr* 0.85) is higher than that in 1.3 treatments (*mpcr* 0.65), usually considerably so, assuming the same known proportion p of conditional cooperators. Along with the assumption that participants' beliefs about p itself are unaffected by the returns from cooperation, this leads to our first prediction:

[19] To be sure, our ranking and grouping procedure can give rise to strategic issues, because once subjects have differentiated themselves with respect to contributions as a phase of the experiment begins to unfold, an individual who is only the second or third highest contributor in his sub-set of five has reason not to give his most preferred rank to the highest contributor in the sub-set containing his prospective partners. This is because while the highest contributor is his most preferred partner in an unconditional sense, the mutual ranking feature of the procedure may put her out of reach. As discussed in Section 4, we nonetheless find both a highly significant correlation between rank assignment and past average contribution of the subject being ranked, and we find that the mechanism is quite effective in pairing partners of similar contribution.

Prediction 1: Contributions will be more sustained in the 100%, 1.7 than in the 100%, 1.3 treatment.

The predictions in our Appendix are worked out under the assumption that any change in one's contribution in period *t* is fully reflected in one's known average past contribution at the beginning of period *t* + 1, and is thus most applicable to our 100% information treatments. What of the other information conditions? Consider first those on the opposite end of the information spectrum. In all of our treatments, subjects maintain fixed IDs during a phase, but in the 0% treatments only their counterpart of a given period learns their action, and there is no way for reputations to spread within the set. A subject might still attempt to build a within-phase reputation for cooperativeness with specific individuals, as do counterparts in AM, who found considerable, indeed growing, cooperation during the early periods of their finitely-repeated games. Such reputation building would be complicated for our subjects, however, by the random selection of the potential partner sub-set each period, since it renders a single pairing unlikely to be sustainable without interruptions even if both partners wish to maintain it. In view of this, we may expect to see at least some subjects attempting to build cooperative relationships with more than one other in a given phase.[20] And despite the difficulty of maintaining ongoing relationships in all treatments, it seems likely that given pairs of subjects will play more periods of a phase with one another in the 0% than in the 50% and 100% treatments, since in the 0% treatments prior play with an individual is the only way to ascertain cooperativeness.

While sparseness of information will cause more rank numbers to be assigned randomly in them, subjects will still preferentially rank whatever past partner had been most cooperative, and the converse for those who free rode, so the ranking and partner assignment mechanism can still be used to similar qualitative effect in the 0% treatments. With reputation so much more

[20] Some might suppose, based on Hauk and Nagel (2001), that the power to exit a given relationship is itself detrimental to cooperation. Those authors closely replicate AM's results in their "lock-in" treatment, while in two other treatments, their subjects are paired for 10 round super-games during each period of which they can choose an outside option of higher value than the mutual defect payoff in any given period. A preference to opt out is implemented if either one prefers it in their "mutual" treatment but only if both prefer it in the "unilateral" treatment. While the presence of exit options, especially "mutual" ones, reduces cooperation in Hauk and Nagel's "mutual" relative to their "lock in" treatment, we see no reason to expect that finding to extend to experiments involving actual partner choice without guaranteed opt-out payoff. This is because Hauk and Nagel's opt out treatments share with partner choice the possibility of avoiding a given partner, but not that of substituting a preferred one, which can create the possibility of competition for good reputation. What makes cooperation more difficult in our 0% treatments is not outside options *per se* but rather the impediments that exist both to building cooperation with any given partner and to building a reputation with which to compete for other partners.

difficult to build, however, we anticipate less overall cooperation in them. Finally, the factor leading us to predict more cooperation in 1.7 than in 1.3 treatments should generate a qualitatively parallel difference in the two 0% treatments.

Prediction 2: Contributions will be lower in 0% than in 100% information treatments. Letting C stand for predicted average contribution, we predict C(100%, 1.3) > C(0%, 1.7) > C(0%, 1.3).

Expectations for the 50% treatments lie somewhere in between those for the 0% and 100% treatments. A given contribution choice has only half the chance of preserving or improving one's outside reputation, if high, and harming it, if low, in a 50% as compared to a 100% treatment. The impression of a "mid-way" position between the other two information conditions seems likely to be misleading, however. Since there's no way to know in advance which choices will affect one's reputation with others besides the current partner, and since those choices that *are* recorded and reported will have about twice the weight on reputation as choices in the 100% treatment, investment in reputation seems likely to be more similar to that in the 100% than that in the 0% treatments. The effects of the differences in *mcpr* remain the same, and potentially strong enough that we cannot confidently predict the relationship between *C*(100%, 1.3) and *C*(50%, 1.7). What can be predicted is:

Prediction 3: C(100%, 1.7) > C(50%, 1.7) > C(0%, 1.7); C(100%, 1.3) > C(50%, 1.3) > C(0%, 1.3); and C(50%, 1.7) > C(50%, 1.3).

Between-phase dynamics

Thus far, we've considered change over time within a single phase only. By assuming a mixed population of conditional cooperators and payoff-maximizers, where $0 < p < 1$, our simple model in which payoff-maximizers contribute at the same level as conditional cooperators from periods 1 to *k*, thereafter contributing 0, implies that there will be a drop in average contribution in the final period or periods of the phase. Given the details of our design, subjects will be more fully informed of the *p* of their population as of the end of their first phase of play the smaller is *k* and the closer to 100% of others' histories they are shown.[21] Even if there is a well-known

[21] Subjects always receive feedback about their own counterpart's action at the end of each period, but learn the decisions of up to five others in the 50% and 100% information treatments only at the beginning of the next period. Hence, if end-game-like behavior appears in period 10 only, subjects' exposure to it in a single phase is very limited.

value of *p* for the universe of individuals from whom the subject set is drawn, *p* can vary randomly between randomly drawn sets of subjects, so it is likely that play of the second phase will differ from that of the first, at a minimum due to updating of beliefs about *p* (which leads to new solution values for *k*).

As pointed out above, however, actual subjects are likely to differ in their degree of strategic sophistication, as well as in their guesses regarding the degrees of such sophistication among fellow subjects, which means that there is room for learning from earlier experience, and that play in phase 2 is likely to differ from that in phase 1 for reasons additional to the revised estimates of *p*. As discussed in the previous sections, it is possible that some subjects will not have appreciated the potential benefits of establishing cooperative reputations at the beginning of the first phase, and since they are able to wipe their reputational slates clean only when one phase ends and another begins, any rise in their appraisal of those benefits may have an even bigger impact on their play at the outset of the next phase than in the remainder of the first one. It is also possible, especially under conditions which make benefits of cooperation small, that some subjects lower their appraisal of cooperation's benefits after the initial phase of play, leading to smaller contributions when the next phase begins. Without knowing the distributions of under- and over-estimates of the benefits of cooperation by payoff-maximizing subjects we cannot make confident predictions regarding these factors.

With regard to end-game learning, it is worth pointing out first that if subjects were all strategically sophisticated and either of conditional cooperator or payoff-maximizer type, with $0 < p < 1$, then there would be no theoretical reason to predict that observing the decline in contributions toward the end of the first phase would lead to earlier rather than to later "unraveling" of cooperation. Those payoff-maximizers whose expectations had proven over-optimistic might reduce their *k*'s, but just as many might have been overly pessimistic and thus have reason to revise those estimates of *p* and therefore their calculation of *k* upwards.[22] Once subjects achieve stable estimates of *p*, *k* should also remain fixed, paralleling the "settling into a stable pattern" dynamics found by AM.

[22] The Appendix, which simplifies by assuming that *p* is common knowledge, shows that the privately optimal *k* of a payoff-maximizer remains fixed over a range of *p* values because *k* is confined to the integers, so not all changes in *p* are associated with changes in *k*.

The presence of less sophisticated subjects whose understanding of the game improves with experience, however, might lead to a tendency for cooperation to "unravel" earlier with each phase. These subjects might not be able to solve for an optimal behavior assuming a distribution of types; rather, they might look for an "appropriate" behavior that seems not too far from the norm, while at the same time seeking to avoid being taken advantage of. Although never seeing the full distribution of last period behaviors, their impulse is to try to anticipate when others will reduce their contributions, and to do so themselves one period earlier.[23]

We end by noting that increases in reputation-building behaviors in early periods, and earlier end-game declines, are not mutually incompatible. Given some initially unsophisticated subjects having an ability to learn from experience, it seems likely that both some learning that it can pay to establish a cooperative reputation, and a tendency to try to stay a step ahead of others' end-game behaviors, may be present, at least when the conditions for reputation building are sufficiently favorable.

4. Results

12 experiment sessions, two for each treatment, each with 20 subjects, were conducted at a computer classroom at Brown University, from October 2012 through March 2013. Adding one under-populated session of 10 subjects, this makes for a total of 250 subjects.[24] The experiment was programmed using z-tree (Fischbacher 2007). Participants were undergraduates drawn from all subject areas, recruited through the BUSSEL (Brown University Social Science Experiment Lab) registration site.[25] All lacked prior experience in VCM experiments. Sessions

[23] Another factor contributing to earlier end-game declines could be uncertainty about their partners' plans on the parts of conditional cooperators, an uncertainty which is absent in a theoretical model with two discrete types and $k < 9$, but likely to be present in the messier real world with varying degrees of sophistication and varying beliefs not only about the distribution between our two assumed types but also about what behavioral types exist (there may be dozens of conceivable types the presence of which each individual estimates differently).

[24] By "under-populated," we refer to one session of the 50%, 1.7 treatment that had insufficient turnout and thus proceeded with only 10 subjects (one set). An extra session with two subject sets was therefore added for that treatment. Because the data from the under-populated session show no systematic differences from the others of its treatment, we keep all 5 set-level observations of it (Table 1).

[25] The university offers a wide range of science, engineering and mathematics, social science and humanities majors. Slightly under 17% of participants reported economics as their major or one of their major fields, almost identical to its representation in the undergraduate classes of the time. 56% of subjects were female, slightly above the 53% share of female undergraduates at the university.

lasted between ninety minutes and two hours. Instructions were neutrally framed and were read aloud by one of the experimenters as subjects read along.[26] Subjects then answered comprehension questions and were invited to ask questions which were answered by a member of the experiment team before the start of decision-making.

We can briefly preview the impacts of our treatments on contribution trends with the help of Figure 1, to which we will return in our more detailed analysis. We see that in all treatments, average contributions in each phase are positive but ultimately decreasing within the phase. At a given information level (0%, 50%, 100%), average contributions appear to be higher in the 1.7 than in the corresponding 1.3 treatment. Between information levels, average contributions tend to be highest at 100%, lowest at 0%, with the difference between both 100% and 50% information, on the one hand, and 0% information, on the other, usually more pronounced than that between 100% and 50%. Within-phase downward trends are noticeably attenuated in the first six to eight periods of a phase in the 50%, 1.7 and 100%, 1.7 treatments. The initial contribution of a phase tends to rise from phase to phase in all 50% and 100% treatments. Together, these last two results suggest that incentives to invest in a reputation for cooperativeness were increasingly effective in the treatments with an adequate basis for reputation formation, although their effectiveness within a phase was more sustained when returns to cooperation were greater (1.7 treatments). Finally, there are visual indications, especially in the 100%, 1.7 treatment, that end-game "unraveling" began earlier in the phase in late than in early phases.

a. Ranking, partner assignment, and payoff to reputation.

The desire to obtain cooperative partners can potentially explain high positive contributions in the early periods of a phase in our experiment. Our detailed analysis accordingly begins with a closer look at how the ranking and partner assignment mechanisms worked. First, we report estimates of individual random effects Tobit regressions in which the rank number assigned by each subject to each potential partner is predicted by the latter's past average contribution in the phase (100% treatments) or by the selected past average and the share of available past periods randomly selected for inclusion (50% treatments). The highly

[26] Full instructions are available at http://www.brown.edu/Departments/Economics/Papers/2013/Instructions_2013_8.

significant negative coefficients on all variables shown in Table 2 indicate that subjects, as expected, tended to give better (lower) rank numbers to those thus far reported to have contributed more on average to their joint accounts, and that in the 50% treatments subjects also showed a preference for counterparts information about whom was more complete.[27] For the 0% treatments, in which subjects could know the past contributions of others only for those periods of the phase in which they had been paired, we separately estimated regressions that show that there, too, subjects gave significantly better ranks to others based on their information about their contribution tendencies and in the 0%, 1.7 treatment, subjects also preferred those with whom they had interacted for more previous periods of the phase, other things being equal.[28]

Result 1: In the 50% or 100% treatments, at both factors 1.3 and 1.7, subjects were more likely to give better rank numbers to those who were reported to have contributed more in the past. In the 0% treatments, subjects were more likely to give better rank numbers to those perceived to be higher contributors based on direct past interactions. In the 50% treatments (the 0%, 1.7 treatment), subjects preferred to be matched with others whose history information was more complete (who had interacted more with the rank-giver).

In order to generate incentives to contribute more to one's joint account, the partner assignment mechanism should reliably assign high contributing subjects to interact with one another, preventing low contributors from accessing these preferred partners. The more highly correlated are the contribution levels of those paired together, then, the more effectively is the mechanism aiding incentive generation. To check for this functionality, we identified for each subject and period (except the first period of each phase, for which no information on past play

[27] To check whether the strategic concern of inability to compete for highest-contributing counterparts (see footnote 19) also influenced ranks given, we estimated probit regressions in which *not* giving one's best rank to the highest past contributor among available prospective counterparts is explained by one's own relative contribution standing (proxied by the ratio of own past average contribution to that of the highest contributor in the other subset of five [we used other rather than own subset because it is what the individual herself had information on]) and phase dummies. Lower relative standing does indeed significantly raise the likelihood of not giving the highest contributor one's best rank. See Appendix Table B.2. However, we were unable to pick up signs of such strategic voting using regression formats resembling Table 2. We therefore conclude that such exceptions to the rule of giving preferred rank to higher contributors did not interfere with our mechanism's operation as an effective generator of incentives to cooperate in order to acquire cooperative partners.

[28] Regression results for the 0% treatments are shown in Appendix Table B.3. For these regressions, we created a "perceived past contribution" variable that takes the value of each group member's average past contribution during interactions with oneself, if the two have interacted in the phase, or else the median value among those one has thus far interacted with, if the assessor has not played with the assessed individual. It is this "perceived past contribution" variable that obtains negative coefficients significant at the 1% level in the regressions mentioned.

was available) (a) the rank of each individual within their five-person matching subset based on their full contribution history in the phase thus far, and (b) in the 0% and 50% treatments, "perceived ranks" within matching subsets based on recorded history (50%) or own interaction history (0%) only. We then calculated Pearson's bivariate correlation coefficients for each period and treatment based on both the objective history and the perceived history approaches. In the 100% information treatments, for which the perceived and objective approaches are the same, correlations are positive and in every period significant at least at the 5% and usually at the 1% level. In the 50% treatments, correlations using either approach are almost always positive but are significant or marginally so in only 82% of periods, with most correlations that are significant being so also at the 1% level. In the 0% treatments, there is some successful matching, with at least marginally significant positive correlations using objective history (method (a)) in 14 of 36 cases for the 0%, 1.3 treatment and in 18 of 36 cases for the 0%, 1.7 treatment. Results are shown in Appendix Tables B.4 and B.5.

As a further check on the mechanism's success in pairing like contributors, we also estimated two regressions for each treatment, with individual random effects, in which the average past contribution of each subject *i* in each period *t* is dependent variable and the average past contribution of *i*'s period *t* partner is the only explanatory variable (except for a constant in one regression out of each pair). All coefficients are positive and significant at the 1% level, indicating that on average like was paired with like in all six treatments. Differences in the magnitudes of the coefficients nonetheless suggest differences in accuracy of pairings. In the treatments with factor 1.7, the magnitudes of the coefficients in the regressions for the 50%, 1.7 and the 100%, 1.7 treatments are similar and almost four times as large as the coefficients of the 0%, 1.7 treatment, indicating unsurprisingly that the ranking and pairing procedure achieved closer matches with respect to contribution tendencies when subjects had better information about their prospective partners. The estimated coefficients in the regression for the 100%, 1.3 treatment are slightly larger than (but not statistically significantly different from) those in the 100%, 1.7 treatment, but significantly larger than those for both the 50%, 1.3 and the 0%, 1.3 treatments, which do not significantly differ from each other. Thus, the tightest correlations between own and partner's past contributions, with coefficients around 0.5, are achieved when there is fuller information on which ranking can be based, with greater incentive to carefully rank perhaps contributing to the tighter correlation in the 50%, 1.7 than in the 50%, 1.3 treatment.[29]

Result 2: The ranking procedure sorted subjects and paired those with cooperative history with similarly cooperative subjects, those with less cooperative history with similarly less cooperative subjects. Its functioning was the most effective in the 100% treatments and least effective in the 0% treatments.

Ultimately, what matters to a mechanism's ability to encourage cooperation is not only the past behaviors of those you are matched with but how they behave when you interact with them. A subject's private incentive to contribute another point in period t depends on (i) the impact this will have on the contribution of the counterpart she is thereby able to secure in period $t + 1$, taking her contributions through period $t - 1$ as given, and on (ii) any corresponding impacts in periods $t + 2$ and beyond. In Table 3, we report individual fixed effect regression estimates of proximate impact (i), that is: the marginal impact of own period t contribution on counterpart's period $t + 1$ contribution, controlling (when possible) for own average contribution (in 50% treatments, reported average contribution) from period 1 to period $t - 1$ of the phase. The regressions are reported only for the 50% and 100% treatments, in which a reputation with third parties can be built via reporting of past contributions. We report separately estimates for period 1 of a phase, for which there are no prior recorded contributions, and for periods 2 through 7, in which the prior contribution average can be controlled for, leaving out the final three periods t (two periods $t + 1$), in which end-game effects seem most likely.[30]

All estimated coefficients suggest that contributing an additional point this period raises one's (possibly new) counterpart's contribution next period, and the estimates are significant at the 1% or 5% levels, with the exception of the regression for period 1 of the 100%, 1.3 treatment. Point estimates are higher for period 1, but the significance of the estimates for the later periods is perhaps more impressive, since, in view of the control, these are truly marginal impacts of contributing in period t, taking already-established reputation as given. Recall that it would be

[29] See Appendix Table B.6 parts I(a) and I(b), which report regressions and tests for significance of coefficient differences when all past contributions are considered. Alternative specifications in which only the information the prospective partners could take into account about one another are considered are shown in later parts of Table B.6 and show significant and in some cases somewhat higher correlations.
[30] That is, we include own contribution in period 7 and next counterpart's contribution in period 8 of each phase; thus, we drop only two of those periods that could have been included (i.e., for which a period $t + 1$ exists).

selfishly rational to contribute an additional point now at a sacrifice of (1 – *mpcr*) points if the impact on the next counterpart's contribution (recalling that the identity of the counterpart is itself not yet determined) is an increase of at least ((1 – *mpcr*)/*mpcr*) points. This implies that the value of $\partial C_{j,t+1}/\partial C_{i,t}$, which is estimated by the first coefficient in Table 3, must be at least (1 – 0.85)/0.85 ≈ 0.176 in the 1.7 treatments and at least (1 – 0.65)/0.65 ≈ 0.538 in the 1.3 treatments in order to spur a contribution from a strictly self-interested subject who myopically considers the next period only. The estimated coefficients in the regressions for the 1.3 treatments are not dramatically different from those for the 1.7 treatments, but since most are not far above the first cut-off value (.176), the required threshold is clearly met by the 1.7 treatments (columns (5) – (9)) only. This means that in the 1.7 treatment it was profitable to cooperate in early periods even if the benefit in the immediate next period only was considered, whereas the 1.3 treatments fall far from this standard.

Failure to meet the narrow criterion of profitability in terms of gain in period $t + 1$ does not in itself rule out that it was profitable for a far-sighted player to cooperate in some earlier periods of the 1.3 treatments. Such a non-myopic payoff-maximizer would ask whether foregoing 0.35 points now can be made up by the expected additional contributions of future partners—additional in the sense of not otherwise being forthcoming—that total more than 0.538 points over all remaining periods of the phase. To the estimate of the impact in period $t + 1$ given above, the far-sighted decision-maker would add impacts by way of the averaging of the period t contribution into the displayed past average contribution in periods $t + 2$, $t + 3$, etc., multiplied by the impact of each point of average contribution in later periods according to the second coefficient of the columns (2) or (6) estimate in Table 3. Supposing $t = 1$, for instance, she could calculate that to the 0.25 expected contributions by her partner in period $t + 1 = 2$, additional contributions of 0.93 in periods $t + 2 = 3$ through $t + 7 = 8$ are predicted by the column (2) estimate for the 50%, 1.3 treatment, for a total contribution of 1.18 points in periods 2 through 8 per point contributed in period 1. By this calculation, contributing one's full endowment in period 1 is clearly profitable.[31] However, the estimated total contributions by

[31] The calculation considers that if $t = 1$, each point contributed in t affects past average contribution (in expectation, for 50% treatments) by ½ in period 3, by 1/3 in period 4, by ¼ in period 5, and so on. Using only the coefficients in Table 3's estimates and assuming no further benefits in periods 9 and 10, the gain per point contributed in period 1 of the 50%, 1.3 treatment is thus **0.25** + **0.38** * (1) + **0.38** * (1/2) + **0.38** * (1/3) + **0.38** * (1/4) + **0.38** * (1/5) + **0.38** * (1/6) ≈ 1.18, where the two numbers (i.e., values) shown in bold typeface are those from our regression estimate.

future partners from own contributions diminish as the phase proceeds: to 0.70 per point contributed in period 2, and to less than the 0.538 threshold, or specifically 0.51 per point contributed in period 3, 0.38 in period 4, and 0.29, in period 5. Parallel calculations using the corresponding Table 3 regression coefficients show contributing to be selfishly rational with considerably larger margins of gain and for a larger share of the phase's periods in the 1.7 treatments.[32] Since subjects could not have perfectly estimated the impacts of their decisions (*inter alia* because they lacked the information based on multiple sets and sessions on which our regressions are estimated), a combination of myopia, uncertainty, and risk aversion would be sufficient to explain why contributions fall off rapidly in each period of the phase following the first, in the 1.3 treatments, whereas the much higher "profit margins" for early period contributions in 1.7 treatments support more sustained contributions within each phase.

Result 3: The benefit of contributing in terms of future partners' contributions exceeded the cost within a single period in the 1.7 but not in the 1.3 treatments. The total benefits of contributing exceeded the cost in the earlier but not the later periods of the 1.3 treatments when partners' contributions in up to seven future periods are considered, and total benefits exceeded cost by a larger margin and for a longer number of periods in the 1.7 than in the 1.3 treatments.

b. Contributions and trends within phases.

Our discussion of the effects of partner choice lead us naturally back to subjects' contribution decisions, which are averaged by treatment and period in Figure 1. In the present section, we discuss the levels and trends of contributions within given phases, focusing

[32] The estimated total contributions by future partners from own contributions of the 100%, 1.3 treatment is calculated similarly, and is 1.5 per point contributed in period 1, 0.89 points in period 2, 0.67 points in period 3, 0.52 points—less than the 0.538 threshold—in period 4, 0.41 points in period 5. The estimated total contributions by future partners from own contributions of the 50%, 1.7 treatment is greater than the 0.176 threshold, as discussed, and is 1.1 per point contributed in period 1, 0.63 points in period 2, 0.49 points in period 3, 0.39 points in period 4, 0.32 points in period 5, 0.26 points in period 6, and 0.21 points in period 7. Similarly, the estimated total contributions by future partners from own contributions of the 100%, 1.7 treatment is greater than the threshold, and is 1.9 per point contributed in period 1, 1.0 points in period 2, 0.75 points in period 3, 0.57 points in period 4, 0.44 points in period 5, 0.33 points in period 6, and 0.24 points in period 7. A caveat to this and the previous sentence is that the estimates in Table 3's even-numbered columns assume unchanging next-period and longer-term impacts within the period 2 – 8 range. Such estimates thus measure average impact, and do not rule out the possibility that the impact of own on partners' contributions was, for example, declining in the course of the phase, which could cause total benefit to dip below cost in an earlier period than that calculated (for example, before period 6 in the 50%, 1.3 treatment example).

especially on the first phase of the experiment. We then turn to how restarting the finitely-repeated game that constitutes a given phase altered behaviors.

The average contribution patterns in the initial phase reconfirm three of the most standard findings of the literature: (1) average contributions begin at around half of the endowment (with considerable variation among individuals), (2) the average contribution tends to decline with repetition, and (3) contributions are larger at a higher than at a lower *mpcr*.[33] Initial contributions cluster between 37 and 49% of endowment in the 1.3 treatments and between 63 and 77% of endowment in the 1.7 treatments. There are hints that contributions are a little higher on average with 50% or 100% reputation history than with 0%, but most differences are not significant.[34] The only noteworthy sign that the endogenous partner assignment and reputation dimensions of the experiment might be altering matters qualitatively is that the average contribution graph of the 100%, 1.7 treatment shows little sign of decline in contributions between periods 1 and 9, a pattern that remains for early periods in later phases of that treatment, and that is increasingly emulated in some phases more than in others by the 1.7, 50% treatment as well.

To obtain a more precise sense of the tendency of contributions to change within a phase, we estimate regressions that control for phase number with phase-specific dummy variables (on which we focus more in the next section's discussion) and that control for period-within-phase by treating period as semi-continuous (1, 2, …) and treating the sequence as repeating (periods 11, 21, and 31 of the session as a whole are the "period 1" of their phases, and so on). With most contributions being infra-marginal, we use least squares, but we control for the set of 10 subjects who interact only with one another by set fixed effects. Because the trends in the early and late periods of phases appear to be different for the 50%, 1.7 and 100%, 1.7 treatments, we estimate separate trends for periods 1 – 7 and for periods 8 – 10 of each phase. The results, shown in Table 4, indicate statistically significant downward contribution trends in every treatment and in both early and late periods of phases. As suspected, however, the rate of decline tends to be

[33] These three themes are already present, for example, in Davis and Holt's (1993) survey of the early literature. See also Ledyard (1995) and Zelmer (2003).

[34] Set-level Mann-Whitney tests indicate that the first contributions in phase 1 are not significantly different between any two treatments in the 1.7 treatments. The first contributions in phase 1, according to set-level Mann-Whitney tests, are not significantly different between the 100%, 1.3 and the 50%, 1.3 treatments or between the 50%, 1.3 and the 0%, 1.3 treatments, but, they are significantly different between the 100%, 1.3 and 0%, 1.3 treatment at the 5% level (p-value = 0.0421, two-sided).

considerably greater in periods 8 – 10 than in periods 1 – 7 in the two 1.7 treatments, with a slope almost three times steeper for the later than for the earlier periods in the 100%, 1.7 treatment, and with the slope differences between early and later periods being statistically significant for both treatments. Rates of decline also show significant between-treatment differences for periods 1 – 7 when comparing the 50%, 1.3 and 50%, 1.7 treatments, and in the comparison between the 100%, 1.3 and 100%, 1.7 treatments.[35]

Result 4: The average contribution followed a declining trend within phases in each of the six treatments. The rate of decline was significantly milder in periods 1 – 7 than in periods 8 – 10 in the 50%, 1.7 and the 100%, 1.7 treatments. Also, the rate of decline in periods 1 – 7 was significantly milder in the 50%, 1.7 than in the 50%, 1.3 treatment, and in the 100%, 1.7 than in the 100%, 1.3 treatment.

The estimates for the 0% treatments show the average rates of decline of contributions to be relatively mild in both periods 1 – 7 and periods 8 – 10 of each phase, due in part to the fact that the average contributions in earlier periods were already at low levels. The far greater difficulty of reputation-building in these treatments most likely explains those low levels, and their difference from the observed levels in the 50% and 100% treatments thus testifies to the power of incentives associated with partner choice to engender more cooperative behaviors. A close look at the data shows, however, that a few subjects were able to form successful partnerships with specific others for as many as 5 and in the limit up to 9 periods within a phase. Recalling that in the 0% treatments, one's actual partners are the only set members with whom one can build a reputation, it is striking that the average earnings of pairs, shown in Panel (1) of Appendix Table B.14, increase as the duration of pairing in a phase lengthens, in those treatments. Distinguishing between subjects' longest and shortest partnerships and performing Wilcoxon matched pair tests, we find that subjects earned significantly more on average during their longer than during their shorter partnerships in phases 2 to 4 in the 0%, 1.3 treatment and in all four phases in the 0%, 1.7 treatment.[36] Note that causality may run in both directions; that is,

[35] See the two-side F tests in Appendix Table B.8.

[36] For this analysis, we first, for each subject, identify both (a) the partners with whom he or she interacted the most in a given phase and (b) the partners interactions with whom were the fewest for that subject in that phase, with pairs playing three of more times being automatically excluded. Then, using the subset of subjects for whom the number of periods in (a) was at least four, and for whom a pair in (b) also existed, we performed Wilcoxon tests, comparing the average earnings in lengthy partnerships to those in short ones. Although the average amount earned tends to be

subjects would have given better ranks to partners who had been more cooperative, which helps lengthen duration, and subjects would have had reason to be more cooperative towards partners with whom they had hopes of playing for more future periods.[37]

Result 5: Some subject pairs managed to build successful partnerships in the 0% treatments, despite the random interruptions built into the design. 0% treatment subjects' long-duration partnerships were significantly more profitable than their partnerships having short durations.

c. Change across phases: Reputation building, and Unraveling of Cooperation.

The main innovation of our design is that it allows us to study what happens to the impact of potential competition for partners when more than one finitely repeated game is played. We speculated that subjects might on average learn to try even harder to invest in cooperative reputations in early periods of later phases, but that their experience of end game effects in early phases might lead to earlier "unraveling" of cooperation. Figure 1 appears to suggest that both effects were present to some degree in all four treatments with substantial possibility of investing in reputation, that is both 50% and both 100% treatments. Although contributions in the 1.3 treatments decay rapidly with repetition within phase, the average contribution in periods 11, 21 and 31 appear to be slightly higher than the average in period 1 for the 50%, 1.3 treatment, and the average appears to rise substantially from period 1 to period 11 and from period 11 to periods 21 and 31 (with the latter not much different from each other), in the 100%, 1.3 treatment. Despite their different within-phase trends, something similar applies to the initial periods of phases in both the 50%, 1.7 and the 100%, 1.7 treatments. Non-parametric tests looking at the first period of phase only find some of these differences to be statistically significant.[38]

---

larger for longer-lasting partnerships in the 50%, 1.7 and 100%, 1.7 treatments also, the corresponding Wilcoxon tests fail to find statistically significant differences based on partnership duration. See Appendix Table B.14.

[37] Appendix Table B.14 shows that for pairs of subjects who played large numbers of periods together, average earnings do not significantly differ between 0% and corresponding 50% or 100% treatments, suggesting that potential reputation formation to compete for other partners was not strictly necessary to cooperation. Among pairs playing fewer periods together, earnings are significantly higher in 50% and 100% treatments. We note that Bayer (2011) devotes considerable attention to the duration of partnerships and similarly finds that long-lasting partnerships show higher levels of cooperation. The fact that cooperation is significantly greater in Bayer's treatment encouraging longer partnerships by adding a bonus for continuing an existing match can be interpreted as evidence that a considerable part of the causality runs from duration to cooperation.

[38] Using as observations only their set-level averages, and thus having only four or five pairs of observations for each test, we find that the average contribution is higher in period 31 than in period 1, significant at the 10% level, in the 50%, 1.3 and 100%, 1.3 treatments according to a Wilcoxon signed ranks test. The corresponding differences are significant at the 5% level in the 50%, 1.7 treatment. Detailed results are found in Panel (2) of Appendix Table B.7.

Rather than focusing on first periods only, the phase dummy variables in Table 4's regressions show differences, if any, in the contribution level at which the linear within-phase trends begin. They indicate that contributions are significantly lower in the second than in the first phase in the two 0% treatments, and in the third than in the second phase in the 0%, 1.7 treatment, with little difference between third and fourth phase. The 50%, 1.3 treatment shows a small but significant decline in contributions in the second and third compared to the first phase, with a smaller and insignificant additional decline in the fourth. The two 100% treatments and the 50%, 1.7 treatment all show contributions to be increasing significantly from phase one to two and two to three, with little change to phase 4. Thus, the visual idea of a tendency for the early contribution trend to begin at a higher level in later than in earlier phases is supported statistically by the initial levels of the linear regression trends for three of our treatments, suggesting that at least some subjects took from their early experiences the idea that they would do well to invest more in the early periods of a phase than they had initially done. These findings are consistent with the predictions discussed in Section 3.2 and with the Table 3 regression results, which indicate that it was profitable to contribute in early periods of the 50%, 1.7 and 100%, 1.7 treatments especially.[39]

Result 6: The average contribution was significantly lower in phases 2 to 4 than in phase 1 in the 0% treatments and the 50%, 1.3 treatment. By contrast, it was significantly higher in phases 3 and 4 than in phase 1 in the 100% treatments and in the 50%, 1.7 treatment.

As mentioned, Figure 1 also gives the impression of earlier end-game unraveling of contributions in later than in earlier phases in those treatments (50%, 1.7 and 100%, 1.7) in which contributions are relatively sustained in the early periods of each phase. We use two approaches to investigate whether subjects were deciding to "free ride" earlier as the experiment progressed.

---

[39] In addition to the significant later phase dummy variables in Table 4, our conclusion that treatments with high benefit of cooperation and good informational basis for reputation formation see cooperation increase in later super-games is also supported by simpler regressions, one for each treatment, in which a phase variable (valued 1, 2, 3 or 4) and constant are the only factors explaining average contribution, with one observation per phase and set. In regressions using average contribution over all ten periods of each phase, phase obtains a positive coefficient in the regressions of the 100%, 1.7, 100%, 1.3 and 50%, 1.7 treatments and is significant at the 1% level for the latter two treatments. In corresponding regressions using contribution averaged over periods 1 to 7 of phase only, the phase variable's coefficient is also positive for the 50%, 1.3 treatment, and is significant at the 1% or 5% levels for the three treatments mentioned in the previous sentence. Results are shown in Appendix Table B.9.

First, for each subject in each treatment and phase, we identify the last period of the phase in which he or she chose $C_i > 0$. Table 5 shows both the mean and the median across subjects within each treatment of this last period of phase in which a positive contribution is made, phase by phase. For each treatment, the average and median period of last positive contribution becomes earlier as one moves from phase to phase, with the only exception occurring between phases 2 and 3 in the 50%, 1.7 treatment. There are interesting differences between treatments, however. In particular, the median subject still contributes a positive amount in the ninth period of most phases in the 1.7 treatments (including 0%, 1.7) whereas there are slightly earlier switches to complete free-riding in the 50%, 1.3 and 100%, 1.3 treatments and a much earlier switch—after period 4—in the 0%, 1.3 treatment.

Result 7: Median and especially average periods of last positive contribution are found to have come earlier as sets of subjects moved from phase to phase in all six treatments. Positive contributions were nevertheless the norm in all but the fourth phase's 9th periods in the 0% and 50%, 1.7treatments, and even in the fourth phase's 9th period in the 100%, 1.7 treatment..

Second, we look at the number of subjects making positive contributions in the final period of each phase. Barring confusion, only true conditional cooperators or possessors of other non-standard preferences should contribute anything in these periods, and the number contributing can be taken (as discussed above) as a lower bound on the number with such preferences. As discussed in Section 3, the existence of conditional cooperators, and subjects' first- and higher-order beliefs about their representation in the subject pool, may be critical to giving selfish subjects an incentive to invest in cooperative reputations. Figure 2 graphs the percentage of subjects contributing $C_i > 0$ in a phase's last period by treatment (each of six lettered sub-figures) and phase (bars within sub-figures). We see that within each treatment, the proportion making a positive contribution in the last period steadily declines from phase to phase, a trend that may be readily explained by decreasing optimism about others' last period contributions based on experience both within the current phase and in previous ones. Between treatments, the main difference is that the share contributing a positive amount in the initial period 10 is somewhat higher in each 1.7 than in each 1.3 treatment and that, among 1.7 treatments, the share appears higher in the 50%, 1.7 than the 0%, 1.7 and higher in the 100%, 1.7 than in the 50%, 1.7 treatment. In each case, higher last period contributions are consistent with higher estimates of the likelihood that counterparts are true conditional cooperators based on greater evidence of cooperation within the initial phase.[40,41]

How consistent are our data with others' measures of the prevalence of conditional cooperation or other non-standard preferences? To answer, we must first decide what inferences our data themselves support. It may be tempting to argue that subjects are initially naïve about the benefit of last period free-riding, learn that it is beneficial in the course of play, and that the 10 to 20% shares observed in period 40 are thus better indicators of true conditional cooperation than the observations from other periods. Since our instructions were clear and included tables showing own and partner's earnings for all possible contribution pairs, since subjects had to answer control questions, and since individuals of above-average intelligence are over-represented in our subject pool, however, we think it reasonable to attribute the greater part of the drop-off in final contributing to declining optimism about others' actions, a reassessment that subjects' experiences were certainly pointing towards. In the treatment and phase in which subjects had most reason to be optimistic about counterparts' contributions, phase 1 of the 100%, 1.7 treatment, about 58% of subjects contributed a positive amount in the phase's last period.[42] This does not differ markedly from estimates of the share of conditionally cooperative individuals in the subject pools of Fischbacher, Gächter and Fehr (2001), Keser and van Winden (2006), Herrmann and Thöni (2009), Kocher *et al.* (2010) and Kamei (2011). It also closely resembles the estimated share of cooperators based on the endogenous grouping treatment in PPU, 59%.

---

[40] Hypothetically, one could form inferences about p from Figure 2 and use them to compute optimal k for a payoff-maximizer as indicated by Appendix Table A.1. There is some rough alignment of behaviors with the resulting predictions. Recall, however, that the impact of behavior in the very last period of a phase is probably much smaller than it would be with fuller information, because each subject in our treatments saw feedback of his or her own partner's behavior in that period only.

[41] As mentioned earlier, Charness and Yang's experiment also includes a complete restart of a finitely repeated super-game with endogenous group formation, so comparing what they find with regard to changes in reputational investment and end-game behavior is of interest. Although not central to their analysis, their Figure 1 suggests that as in our 50%, 1.7 and 100%, 1.7 treatments, subjects in their two endogenous grouping treatments contributed more in the earlier periods of their second than those of their first super-game and that the end-game effect in the second super-game was more pronounced than that in the first. The more complex procedures for forming groups of quite variable size and the attention devoted to "expulsion" and "redemption" in Charness and Yang make the focus of the two papers rather different.

[42] Since subjects are randomly drawn from the same subject pool, it is unlikely that the true p differs much between the subjects in different treatments and likely that the better part of the differences apparent in Figure 2 are attributable to differences in treatment parameters, not in p.

5. Discussion and Conclusions

The desire to be a reputable partner in a world in which people get to decide who they interact with is a probable cause of much observed cooperation. Social dilemma experiments in which subjects in some conditions have a say over who they play with in a finitely repeated sequence of interactions have provided support for this intuition by demonstrating that cooperation rates are higher when subjects have than when they lack such say and when they have than when they lack the informational means to build a reputation with prospective interaction partners. Finite repetition has been used in these experiments because it allows one to rule out the potential cooperative equilibria of infinitely repeated games as alternative explanations for cooperation.

Analogous real life situations are sometimes better captured by infinite repetition, sometimes by finite repetition models. Even in interactions involving finite repetition, however, people may get to start over, for instance when they move or change careers. To see how experience of a finitely repeated social dilemma with partner choice might affect play in subsequent interactions of the same kind, we undertook to study how subjects' experiences would impact play in later finitely-repeated settings.

We found that when the benefit of mutually cooperating is high relative to the earnings foregone by not free riding, and when the informational conditions for acquiring a reputation within a population of potential partners are favorable, relatively high levels of cooperation (contributions averaging 70 to 90% of endowment, versus the typical start around 50% followed by immediate decline, when partner choice is absent) were sustained during most periods of our ten-period super-games. With respect to learning across super-games, we found that cooperation in early periods became stronger in later games. This suggests that subjects updated their beliefs about the returns to cooperating—a return from attracting more cooperative partners which we showed to be large and significant—in the direction of increased optimism. Even in treatments with a lower benefit from cooperation in which subjects displayed a much weaker tendency to sustain cooperation within a given super-game, first period contributions rose across games when there was adequate basis for reputation formation, which suggests learning that initial cooperation paid off—a fact supported by both theoretical and regression analyses.

Although early-in-phase cooperation grew over time especially in the high benefit treatments with good information, end-game associated declines began earlier in each phase than in the previous one, raising the possibility that with enough repetition, "learning the benefits of cooperation" might eventually be overwhelmed by "backward unraveling." Our 40 period, four (finitely repeated) super-game experiment may be too short to draw definite conclusions. Andreoni and Miller (1993) also saw end-game effects appear to strengthen during the first four or so of their twenty prisoners' dilemma super-games with fixed partners, but they concluded that within super-game cooperation thereafter settled into a stable pattern in which most subjects continued to build a cooperative reputation with their next partner in the early periods of their finitely repeated interactions. Given the greater complexity of our design, we could complete fewer games within the time span of a conventional decision experiment, so we are unable to tell whether a similar "settling down" would apply in our design.

Voluntary contribution games differ from prisoners' dilemmas in that intermediate degrees of cooperation or defection are possible, making the ways in which individuals can condition their cooperation on that expected from others more flexible. Competition for partners in our design but not AM's also makes direct inferences risky. Nevertheless, it is at least possible that if the non-standard "type" in our subject pool are conditional cooperators in the sense of the literature cited, they might well contribute less and less in each final period, as the average expected contributions of others continue to decline. Especially if conditional cooperators are biased towards contributing a little less than their partners (Fischbacher and Gächter, 2010), end-game contributions might eventually approach zero.

The above must remain only a conjecture, however. Our paper's contribution lies in demonstrating yet again that competition for partners can be a spur to cooperation, and in showing that both learning that cooperating pays and learning to anticipate end-game behaviors can coexist when a finitely repeated social dilemma game with partner choice is restarted. We show that competition for reputation can increase cooperation even among experienced subjects, but we find indications that the process might not be repeatable without limits.

Acknowledgments: We thank Yunan Ji and Philip Trammell for their excellent research assistance, and Cheewadhanaraks Matanaporn for help with experiment sessions. We also wish to thank Arno Riedl, Ralph Bayer and Dietmar Fehr from help on identifying relevant literature. This project was supported by a research grant from Murata Science Foundation.

References

Ahn, T.K., R. Mark Isaac, and Timothy Salmon, 2008, "Coming and going: Experiments on endogenous group sizes for excludable public goods," *Journal of Public Economics* 93(1-2): 336-51.

Ahn, T.K., R. Mark Isaac, and Timothy Salmon, 2009, "Endogenous Group Formation," *Journal of Public Economic Theory* 10(2): 171-94.

Aimone, Jason, Laurence Iannaccone and Michael Makowsky, forthcoming, "Endogenous Group Formation via Unproductive Costs," *Review of Economic Studies* (in press).

Andreoni, James and Miller, John, 1993, "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma: Experimental Evidence," *Economic Journal* 103(418): 570-85.

Bayer, Ralph, 2011, "Cooperation in Partnerships: The Role of Breakups and Reputation," The University of Adelaide School of Economics Research Paper No. 2011-22.

Charness, Gary and Chun-Lei Yang, 2010, "Exit, Exclusion, and Mergers: Endogenous Group Formation and Public Goods Provision," mimeo.

Cinyabuguma, Matthias, Talbot Page, Louis Putterman, 2005, "Cooperation under the threat of expulsion in a public goods experiment," *Journal of Public Economics* 89: 1421-35.

Coricelli, Giogio, Dietmar Fehr and Gerlinde Fellner, 2004, "Partner Selection in Public Goods Experiments," *Journal of Conflict Resolution* 48: 356-378.

Davis, Douglas and Charles Holt, 1993, *Experimental Economics*. Princeton: Princeton University Press.

Ehrhart, Karl-Martin and Claudia Keser, 1999, "Mobility and Cooperation: on the Run," CIRANO working paper 99s-24.

Engelmann, Dirk and Urs Fischbacher, 2009, "Indirect Reciprocity and Strategic Reputaiton Building in an Experimental Helping Game," *Games and Economic Behavior* 67: 399 – 407.

Fehr, Ernst and Klaus Schmidt, 1999, "A Theory of Fairness, Competition and Cooperation," *Quarterly Journal of Economics* 114: 817 – 868.

Fischbacher, Urs, Simon Gächter, and Ernst Fehr, 2001, "Are people conditionally cooperative? Evidence from a public goods experiment," *Economics Letters*, 71: 397-404.

Fischbacher, Urs and Simon Gächter, 2010, "Social Preferences, Beliefs, and the Dynamics of Free Riding in Public Goods Experiments," *American Economic Review* 100: 541-556.

Fischbacher, Urs, 2007, “z-Tree: Zurich Toolbox for Ready-made Economic Experiments,” *Experimental Economics* 10 (2): 171-8.

Gächter, Simon and Christian Thöni, 2005, “Social Learning and Voluntary Cooperation among Like-minded People,” *Journal of the European Economic Association* 3: 303-314.

Gunnthorsdottir, Anna, Daniel Houser and Kevin McCabe, 2007, “Disposition, history and contributions in public goods experiments,” *Journal of Economic Behavior & Organization* 62(2): 304-315.

Gürerk, Őzgür, Bernd Irlenbusch and Bettina Rockenbach, 2006, “The Competitive Advantage of Sanctioning Institutions,” *Science* 312(5770): 108-110.

Hauk, Esther and Rosemarie Nagel, 2001, “Choice of Partners in Multiple Two-Person Prisoner's Dilemma Games: An Experimental Study”, *Journal of Conflict Resolution* 45(6): 770-93.

Herrmann, Benedikt, and Christian Thöni, 2009, “Measuring conditional cooperation: a replication study in Russia,” *Experimental Economics*, 12: 87-92.

Kamei, Kenju, 2011, “From Locality to Continent: A Comment on the Generalization of an Experimental Study,” *Journal of Socio-Economics*, 41: 207-10.

Kandori, Michihiro, 1992, “Social Norms and Community Enforcement,” *Review of Economic Studies* 59: 63-80.

Keser, Claudia, and Frans A.A.M. van Winden, 2006, “Conditional Cooperation and Voluntary Contributions to a Public Good,” *Economics Bulletin* 3(8): 1-14.

Kreps, David, Paul Milgrom, John Roberts, and Robert Wilson, 1982, “Rational Cooperation in the Finitely Repeated Prisoners’ Dilemma,” *Journal of Economic Theory* 27: 245-52.

Lazear, Edward, Ulrike Malmendier and Roberto Weber, 2012, “Sorting in Experiments with Application to Social Preferences,” *American Economic Journal*: *Applied Economics* 4: 136-163.

Ledyard, John, 1995, Public goods: A survey of experimental research. In J. H. Kagel and A.E. Roth (eds.), The Handbook of Experimental Economics, 111-194, Princeton, NJ: Princeton University Press.

Maier-Rigaud, Frank, Peter Martinsson and Gianandrea Staffiero, 2010, “Ostracism and the Provision of a Public Good: Experimental Evidence,” *Journal of Economic Behavior and Organization* 73: 387-395.

Nowak, Martin and Karl Sigmund, 1998, “The Dynamics of Indirect Reciprocity,” *Journal of Theoretical Biology* 194 (4): 561-574.

Page, Talbot, Louis Putterman and Bulent Unel, 2005, “Voluntary Association in Public Goods Experiments: Reciprocity, Mimicry, and Efficiency,” *Economic Journal* 115(506): 1032-1053.

Riedl, Arno, Ingrid M.T. Rohde, Martin Strobel, 2011, “Efficient Coordination in Weakest-Link Games,” *IZA Discussion Paper* No. 6223.

Riedl, Arno and Aljaz Ule, 2013, “Competition for Partners Drives Cooperation among Strangers,” paper presented at the Fifth Biennial Social Dilemmas Conference, CalTech.

Seinen, Ingrid and Arthur Schram, 2006, “Social Status and Group Norms: Indirect Reciprocity in a Helping Experiment,” *European Economic Review* 50: 581-602.

Selten, Reinhard and Stoecker, Rolf, 1986, “End Behaviour in sequences of finite prisoner’s dilemma supergames: a learning theory approach,” *Journal of Economic Behavior and Organization* 7(1): 47-70.

Sen, Amartya, 1967, “Isolation, Assurance and the Social Rate of Discount,” Quarterly Journal of Economics 81: 112-24.

Wang, Jing, Siddharth Suri and Duncan Watts, 2012, “Cooperation and Assortativity with Dynamic Partner Updating,” *Proceedings of the National Academy of Sciences* 109(36): 14363-8.

Zelmer, Jennifer, 2003, “Linear Public Goods Experiments: A Meta-Analysis,” *Experimental Economics* 6: 299-310.

Table 1. Summary of treatments, and average contributions

Treatment	mpcr	The probability that contribution is recorded	Total # of sessions	Total number of sets	Total # of subjects	Average contributions
(a) Treatments with Factor 1.3						
0%, 1.3	.65	0%	2	4	40	1.66
50%, 1.3	.65	50%	2	4	40	3.57
100%, 1.3	.65	100%	2	4	40	3.85
(b) Treatments with Factor 1.7						
0%, 1.7	.85	0%	2	4	40	4.59
50%, 1.7	.85	50%	3	5	50	6.62
100%, 1.7	.85	100%	2	4	40	7.38
Experiment as a whole			13	21	210	

Figure 1. Average Contribution Period by Period

(a) Treatments with Factor of 1.3

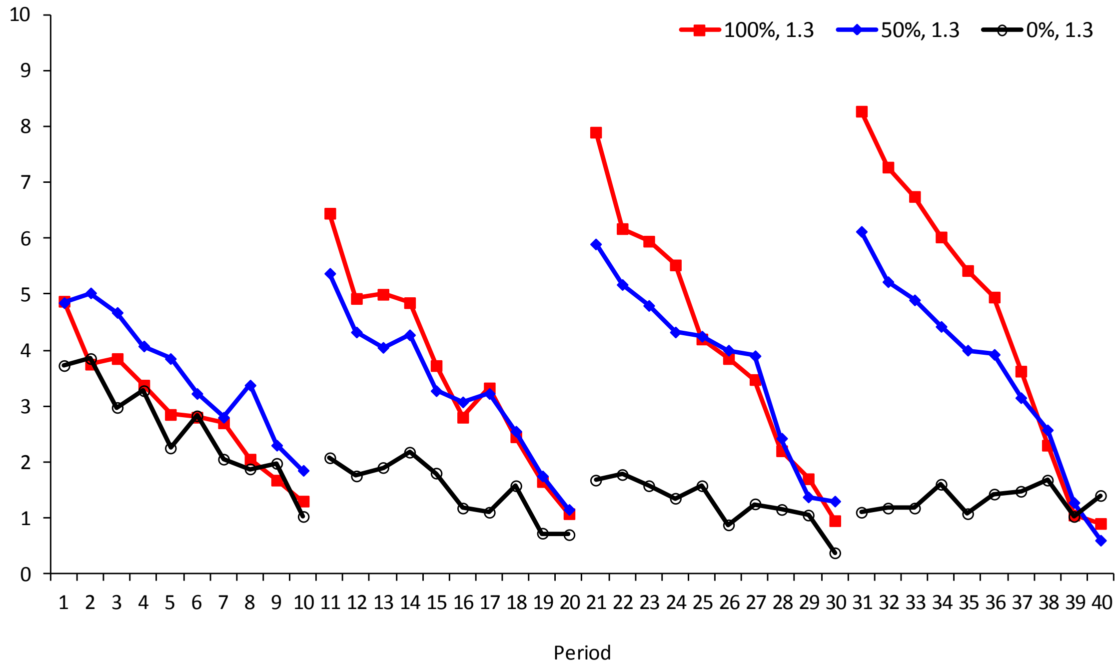

(b) Treatments with Factor of 1.7

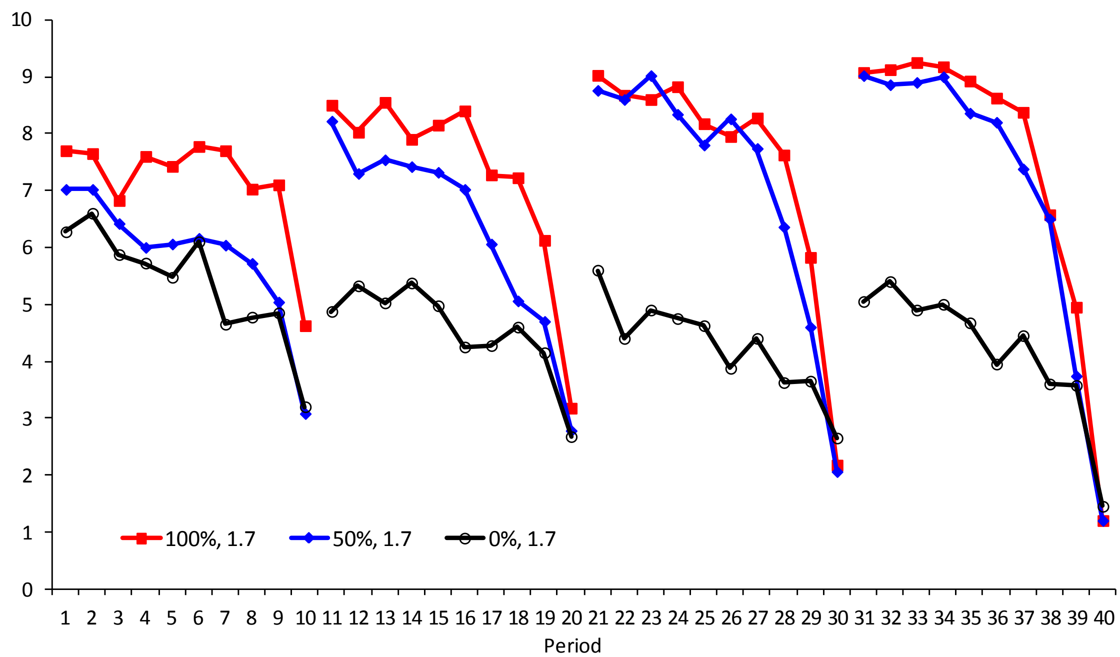

Notes: Non-parametric tests between different treatments, or across different phases, are found in Appendix Table B.7.

Table 2. Determinants of Ranking Decisions in the 50% and 100% info treatments

Dependent variable: Rank given to subject *j* in Period *t*.

	Group account efficiency			
	Factor of 1.3		Factor of 1.7	
Independent Variable	50% info (1)	100% info (2)	50% info (3)	100% info (4)
(a) subject *j*'s Average Previous Contribution	-0.42*** (0.0078)	-0.43*** (0.0089)	-0.44*** (0.0084)	-0.53*** (0.012)
(b) share of past periods for which information is included	-0.93*** (0.10)	----	-0.82*** (0.093)	----
Constant	5.45*** (0.14)	5.32*** (0.10)	6.82*** (0.10)	7.42*** (0.15)
# of Observations	6345	7200	8210	7200
Log Likelihood	-9508.9	-11452.1	-12886.4	-11587.3
Chi-squared	3235.5	2314.8	2869.4	1896.3
Prob > Chi-squared	0.000	0.000	0.000	0.000

Notes: Individual random effect Tobit regressions. Only observations whose variable (b) is greater than 0 are used. The numbers of left-(right-) censored observations are 1379(1236) in column (1), 1440(1440) in column (2), 1729(1540) in column (3) and 1440(1440) in column (4). Test results for the equality of the coefficient of variable (a) across treatments are found in Appendix Table B.1. Ex-post efficiency of the ranking procedure is measured by calculating the bivariate correlations between matched pairs' their past contribution decisions, whose results are found in Appendix Tables B.4 and B.5. *, **, and *** indicate significance at the .10 level, at the 0.05 level and at the .01 level, respectively.

Table 3. Period t contribution versus Period $t + 1$ partner's contribution

Dependent variable: Period $t + 1$ partner's contribution.

Independent Variable	50%, 1.3		100%, 1.3		50%, 1.7		100%, 1.7	
	$t = 1$	$t > 1$ & $t \leq 7$	$t = 1$	$t > 1$ & $t \leq 7$	$t = 1$	$t > 1$ & $t \leq 7$	$t = 1$	$t > 1$ & $t \leq 7$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Own contribution in period t	0.25** (0.066)	0.15** (0.044)	0.43 (0.20)	0.24*** (0.025)	0.37*** (0.039)	0.21** (0.060)	0.61*** (0.071)	0.24*** (0.045)
Own (recorded) past average contribution up to $t - 1$[1]	----	0.38** (0.10)	----	0.45*** (0.046)	----	0.29** (0.081)	----	0.54*** (0.047)
Periods within phase ($\in \{2, 3, \ldots, 7\}$)	----	-0.17** (0.040)	----	-0.33*** (0.057)	----	-0.28** (0.086)	----	-0.17* (0.054)
Phase 2 dummy {=1 if phase =2}	-0.83 (0.31)	-0.29 (0.14)	0.49 (0.71)	-0.052 (0.14)	-0.17 (0.55)	-0.012 (0.30)	-0.11 (0.41)	-0.027 (0.27)
Phase 3 dummy {=1 if phase =3}	-0.12 (0.45)	-0.021 (0.48)	1.11 (0.93)	-0.26 (0.18)	0.93* (0.42)	0.62 (0.39)	0.22 (0.76)	-0.038 (0.22)
Phase 4 dummy {=1 if phase =4}	-0.13 (0.71)	-0.29 (0.28)	2.05** (0.63)	-0.12 (0.30)	1.09** (0.26)	0.74 (0.53)	0.64 (0.57)	-0.089 (0.25)
Constant	3.79*** (0.66)	2.02** (0.47)	1.64 (0.94)	1.80*** (0.20)	4.39*** (0.49)	4.17** (0.92)	2.98** (0.62)	2.32** (0.42)
# of Observations	160	955[2]	160	960	200	1200	160	960
F	----	----	----	----	93.15	----	----	----
Prob > F	----	----	----	----	.0003	----	----	----
R-Squared	.1320	.3806	.3093	.4952	.1459	.2420	.2320	.3696

Notes: Individual fixed effects regression with robust standard errors clustered by setid. Only observations whose t is less than or equal to 7, and greater than 1 are used in columns (2), (4), (6) and (8).

[1] If a subject's contribution decisions have not been recorded yet, then the median of other group members whose contribution decisions have been recorded at least once is used.

[2] In one group, no contributions had yet been recorded as of period 22 in the 50%, 1.3 treatment; and thus the five observations of that group are excluded in this regression.

*, **, and *** indicate significance at the .10 level, at the 0.05 level and at the .01 level, respectively.

Table 4. Trends of Average Contributions by Treatment: Regression Analyses

Dependent variable: Set average contributions per period.

	Group account efficiency					
	Factor of 1.3			Factor of 1.7		
Independent Variable	0% info (1)	50% info (2)	100% info (3)	0% info (4)	50% info (5)	100% info (6)
(a) Phase2 dummy {1 if Phase 2; 0 otherwise}	-1.09*** (0.17)	-0.39** (0.19)	0.87*** (0.22)	-1.11*** (0.21)	0.60** (0.25)	0.19 (0.33)
(b) Phase3 dummy {1 if Phase 3; 0 otherwise}	-1.32*** (0.17)	-0.38** (0.19)	1.87*** (0.22)	-1.88*** (0.21)	1.81*** (0.25)	0.37*** (0.33)
(c) Phase4 dummy {1 if Phase 3; 0 otherwise}	-1.28*** (0.17)	-0.26 (0.19)	1.65*** (0.22)	-1.79*** (0.21)	1.89*** (0.25)	0.39*** (0.33)
(d) Period within phase1 {= 1, 2, …, 7}[#1]	-0.13*** (0.036)	-0.44*** (0.039)	-0.50*** (0.045)	-0.19*** (0.043)	-0.26*** (0.052)	-0.14** (0.069)
(e) Period within phase2 {= 8, 9, 10}[#2]	-0.13*** (0.022)	-0.46*** (0.023)	-0.63*** (0.027)	-0.21*** (0.026)	-0.47*** (0.031)	-0.41*** (0.041)
Constant	3.29*** (0.19)	5.59*** (0.21)	6.29*** (0.24)	5.22*** (0.23)	7.27*** (0.28)	8.67*** (0.37)
# of Observations	160	160	160	160	200	160
F	23.45	81.77	141.76	34.78	80.06	32.67
Prob > F	.0000	.0000	.0000	.0000	.0000	.0000
R-Squared	.3695	.4863	.6998	.4107	.5536	.3583
F test results						
(a) = (b)						
F	1.84	0.00	21.53***	13.98***	23.61***	0.31
p-value (two-sided)	.1774	.9468	.0000	.0003	.0000	.5794
(a) = (c)						
F	1.16	0.47	13.01***	10.76***	26.76***	0.35
p-value (two-sided)	.2826	.4962	.0004	.0013	.0000	.5537
(b) = (c)						
F	0.08	0.38	1.07	0.21	0.10	0.00
p-value (two-sided)	.7823	.5393	.3036	.6474	.7537	.9697
(d) = (e)						
F	0.02	0.58	16.07***	0.48	36.46***	31.93***
p-value (two-sided)	.8927	.4475	.0001	.4878	.0000	.0000

Notes: Set fixed effects linear regressions. The dependent variables are per-period set average contribution. [#1] The Period within phase 1 variable equals 0 if it is in period 8, 9 or 10. [#2] The Period within phase 2 variable equals 0 if it is in period 1, 2, ..., 6, or 7. Test results for the equality of the coefficient of each of variables (a) to variable (e) across treatments are found in Appendix Table B.8.

*, **, and *** indicate significance at the .10 level, at the 0.05 level and at the .01 level, respectively.

Table 5. End-effects Behavior by Phase and Treatment: The last period in which a subject contributed a positive amount

		Average				Median			
		Phase 1	Phase 2	Phase 3	Phase 4	Phase 1	Phase 2	Phase 3	Phase 4
Treatment	I. Factor of 1.3								
	0%, 1.3	7.05 (3.46)	5.55 (4.37)	4.50 (4.36)	4.15 (4.21)	8	8	4	4
	50%, 1.3	7.75 (3.12)	7.33 (2.81)	6.18 (3.47)	5.88 (3.15)	9	8	7	7
	100%, 1.3	7.63 (3.12)	7.53 (1.91)	6.83 (2.77)	6.60 (2.55)	9	8	8	7
	II. Factor of 1.7								
	0%, 1.7	7.38 (3.61)	6.83 (4.04)	6.63 (4.00)	5.88 (3.96)	9	9	9	7
	50%, 1.7	8.82 (1.77)	8.14 (2.46)	8.52 (1.36)	8.00 (1.67)	9	9	9	8
	100%, 1.7	9.35 (1.10)	8.85 (1.78)	8.83 (1.11)	8.48 (1.15)	10	9	9	9

Notes. Numbers in parenthesis are standard deviations. We use 0 for a subject's last period if the subject contributed nothing during the entire phase. Test results in comparing the average last periods between treatments are found in Appendix Table B.10. Parallel to this analysis, we also calculated the percentage of subjects that contributed nothing to their joint account by period and by treatment. The results are similar, and are omitted to conserve space.

Figure 2. The Percentage of Subjects that contributed positive amounts in the tenth period of each phase, by treatment

(a) Treatments with Factor of 1.3

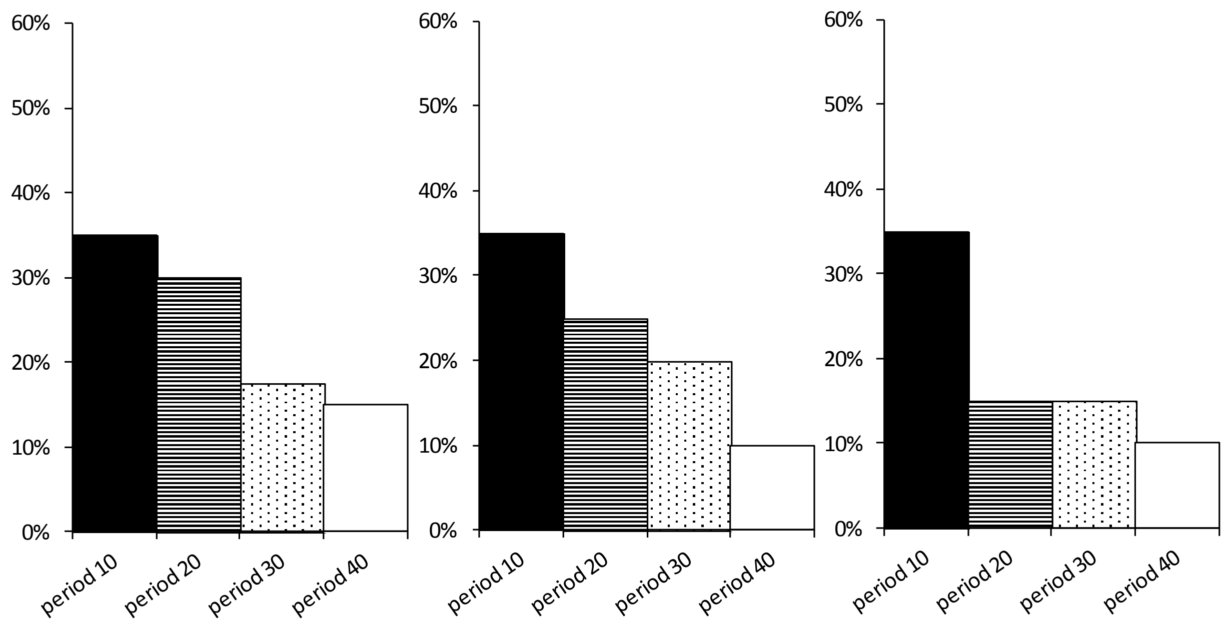

(a1) 0%, 1.3 (a2) 50%, 1.3 (a3) 100%, 1.3

(b) Treatments with Factor of 1.7

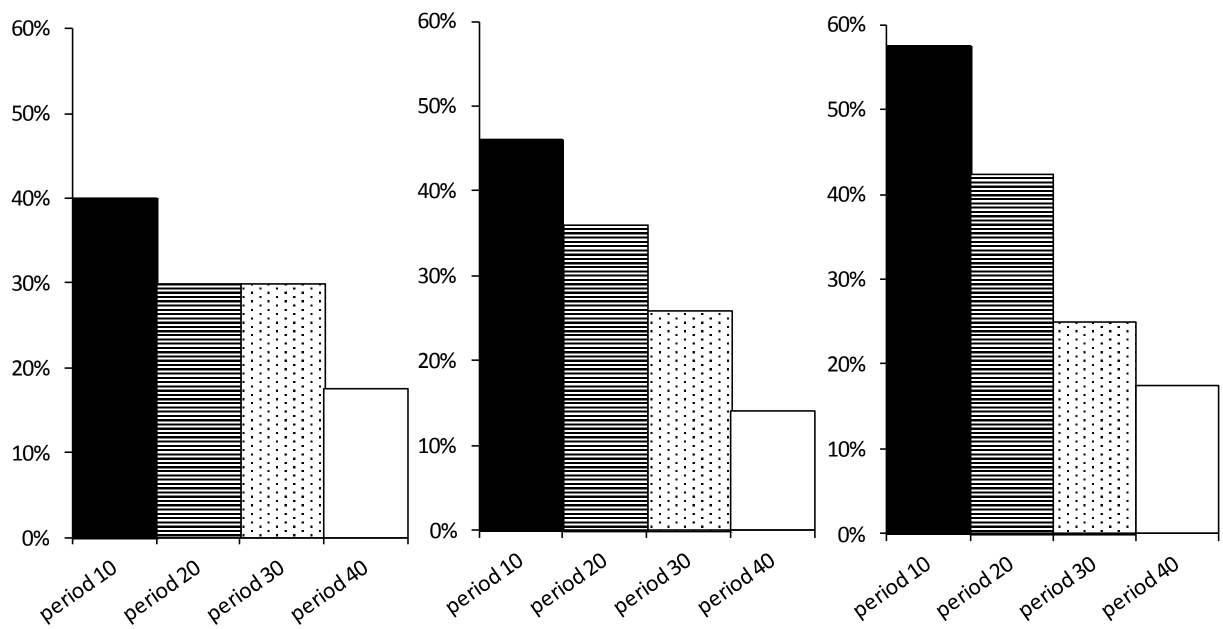

(b1) 0%, 1.7 (b2) 50%, 1.7 (b3) 100%, 1.7