Social Preferences and Payoff-Based Learning Explain Contributions in Repeated Public Goods Games

Chenna Reddy Cotla* and Ragan Petrie†

November 2019

Abstract

We conduct experiments to examine the interplay between social preferences and learning in repeated public goods games. In our design, we observe individual-level cooperation in a one-shot and repeated setting at both low and high costs to cooperate. Using behavior in the one-shot setting, we show that changes in the distribution of observed cooperation (free riding, conditional cooperation, and full cooperation) when prices change are consistent with a model of stable social preferences. In repeated settings, first-round behavior is best characterized by social preferences followed by simple payoff-based reinforcement learning in subsequent rounds. Predictions from this model are validated using a follow-up experiment and show there are benefits to subsidizing initial cooperation. By making the cost to cooperate low in the first interaction with others, a group can sustain a 27% higher level of cooperation later, compared to groups that faced a high initial cost.

JEL Codes: C63, C92, D83, H41

Keywords: public goods games, social preferences, learning, validation experiment

Acknowledgements: This research was funded by a National Science Foundation Dissertation Enhancement Grant (NSF#1324155).

*American Institutes for Research and Key Bank, chennacotla@gmail.com

†Texas A&M University, Department of Economics, and Melbourne Institute, rpetrie@tamu.edu

1 Introduction

There are benefits to cooperation in groups and the provision of public goods, yet individuals differ in their propensity to contribute. Some individuals are always helpful, others not at all, and some base their willingness to cooperate on how cooperative others are. When groups interact over time, the mix of individuals in a group can affect the level of public goods that end up getting provided. Individuals may bring a predisposition towards cooperation to a group interaction, but they may also modify their behavior or learn when they see the actions of others. How cooperative groups are then would be a combination of predispositions and learning. Understanding how this process works is important for fostering cooperative environments and outcomes. To gain some insight on this, we examine how cooperative predispositions and learning can characterize group cooperation across repeated interactions and what this implies for initial group behavior to sustain cooperation.

We use linear public goods experiments to examine the interplay of the propensity to cooperate and learning in groups. Experimental evidence from these type of games suggests that individuals vary in their predisposition towards cooperation and can be classified into a range of cooperative types (Fischbacher, Gächter, & Fehr, 2001; Fischbacher & Gächter, 2010; Kurzban & Houser, 2001, 2005; Burlando & Guala, 2005; Duffy & Ochs, 2009; Kocher, Cherry, Kroll, Netzer, & Sutter, 2008; Herrmann & Thöni, 2009; Muller, Sefton, Steinberg, & Vesterlund, 2008; Bardsley & Moffatt, 2007). Some individuals give everything (full cooperators), some give nothing (free riders), and some act as conditional cooperators. An individual's willingness to cooperate is not necessarily fixed and can be affected by the circumstance of the group interaction. For instance, cooperation tends to increase when the cost to do so decreases (Isaac, Walker, & Thomas, 1984; Cartwright & Lovett, 2014; Goeree, Holt, & Laury, 2002; Lugovskyy et al., 2017). The cooperative type of an individual is characterized as a function of fundamental other-regarding, or social, preferences (Fehr & Schmidt, 1999; Bolton & Ockenfels, 2000; Charness & Rabin, 2002; Cox, Friedman, & Sadiraj, 2008; Arifovic & Ledyard, 2012), and theory of social preferences provides guidance on how an individual's disposition towards cooperation might be affected by the parameters of the game.

In finitely repeated linear public goods games, contributions start around half the endowment and then steadily decline (Ledyard, 1995) and vary by cooperative type (Fischbacher & Gächter, 2010). A prominent explanation for these findings is that individuals start with random first-round contributions and employ directional learning based on their social preferences (Anderson, Goeree, & Holt, 2004; Wendel & Oppenheimer, 2010; Cooper & Stockman, 2002; Janssen & Ahn, 2006; Arifovic & Ledyard, 2012).[1] Full cooperators learn to contribute their full endowment, free riders learn to contribute nothing, and conditional cooperators learn to match their belief of what the other group members contribute on average. Based on the typical distribution of co-

[1] Learning and the parameters of the game are shown to be important to sustaining cooperation in Prisoner's Dilemma games (Embrey, Frechette, & Yutsel, 2018; Friedman & Oprea, 2012), a two-person version of a social dilemma.

operative types observed in public goods games, this behavioral theory can explain the decay in contributions and different levels of average contributions across cooperative types in a repeated game. It can also shed light on other findings from repeated public goods games including, as noted earlier, higher average contributions as the price of cooperation declines and the effects of group size on cooperation (Isaac & Walker, 1988; Isaac, Walker, & Williams, 1994; Diederich, Goeschl, & Waichman, 2016). Lower costs of cooperation and larger group sizes increase the probability of having conditional cooperators and full cooperators relative to free riders which in turn leads to higher average contributions. While this behavioral theory provides a plausible explanation of stylized findings, it is also possible that aggregate differences across types arise simply because initial contributions differ and individuals learn based on payoffs or type-specific equilibrium contributions. We compare these alternative explanations of behavior in repeated public goods games.

Two laboratory experiments are conducted to understand the roles of social preferences and learning in repeated linear public goods games. In the first experiment, our design allows us to compare competing behavioral specifications to explain contribution patterns in repeated games. These specifications propose different relationships between social preferences and learning that in theory can produce stylized differences in the average contributions across cooperative types. We examine whether social preferences can explain the observed changes in the cooperative disposition of participants when the cost to cooperate increases. This feature is crucial to validate the use of social preferences as a fundamental construct underlying cooperative dispositions. The second experiment serves as a validation exercise of the behavioral predictions derived from the data in the first experiment.

The experiments are based on Fischbacher and Gächter (2010). There are four treatments in a within-subject design, composed of a one-shot conditional contribution game (labeled the P-task) and a repeated game with partners matching (labeled the R-task) at two different costs to cooperate (i.e. marginal per capita return (MPCR) levels), one low and one high, that preserve the social dilemma. The one-shot game uses the strategy method to elicit an unconditional contribution and conditional contributions based on several possible average contributions of the other group members. Each one-shot game is then followed by a repeated game at the same MPCR level. In this game, for each participant along with a contribution decision in each round, we also elicited a belief about the average contribution of the other members of his group. The one-shot and repeated game sequence is then conducted again at a different MPCR level. The order of which MPCR is seen first is reversed across sessions to control for experience effects. Finally, across sessions, but not within sessions, we varied whether the repeated games were finite or indefinite.

We have a number of results. In our first experiment, we confirm that variations in the observed distribution of cooperative types when the cost to cooperate changes is consistent with stable social preferences. That is, as prices decline, individuals tend to shift to a more cooperative type in their behavior, but underlying preferences do not change. First-round choices in the repeated games are significantly different across cooperative types.

That is, choices are not random but are informed by social preferences. When we compare competing models of the relationship between social preferences and learning, the one that best fits the data is where choices in the first round are generated by social preferences and subsequent contributions are characterized by payoff-based reinforcement learning. Variation in first-round contributions and high inertia of reinforcement learning explain the differences in average contributions across cooperative types.

Our second experiment is designed to test the validity of the best-fit behavioral specification from the first experiment. An implication of this specification is that by changing the price of cooperation in the first round of a repeated game, relative to later rounds, one can affect average contributions in subsequent rounds. If individuals start with a randomly-chosen contribution and then learn equilibrium strategies, changing the first-round price of cooperation would not have any effect on later round contributions. Thus, the second experiment manipulates first-round price to test this. We find that lowering the price to cooperate in the first round leads to higher initial contributions, consistent with the findings from our first experiment. This is then followed in later rounds by a higher average level of contributions (27% higher), and higher earnings (33%), compared to the level attained by groups that faced a high price to cooperate in the first round. These findings lend support for our best-fit specification that social preferences drive initial choices followed by payoff-based reinforcement learning. Groups that can attain high levels of contributions initially will benefit over time as individuals base their subsequent contributions on what others did in the previous period.

Our results contribute to the ongoing discussion of the relative importance of social preferences and learning in explaining contributions in repeated public goods games. Learning in repeated public goods games has been modeled as individual level directional learning based on social preferences (Anderson et al., 2004; Wendel & Oppenheimer, 2010; Cooper & Stockman, 2002; Janssen & Ahn, 2006; Arifovic & Ledyard, 2012). Recent experimental results find that social preferences may be unnecessary to explain contributions in repeated games and that individuals learn based on payoffs in the game (Burton-Chellew & West, 2013; Burton-Chellew, Nax, & West, 2015). Our results suggest that the importance of social preferences may be overestimated in the former and underestimated in the latter. Social preferences matter to the extent that they determine the first-round contributions, thereafter, individual behavior is explained by pure payoff-based reinforcement learning.

In more recent work, Boosey, Isaac, Norton, and Stinn (2019) extend the findings of Fischbacher and Gächter (2010) to partner's matching and show a correlation between cooperative behavior in the strategy and repeated games, as we also find. They further show that the classification methods of Fischbacher and Gächter (2010) and Kurzban and Houser (2005) yield similar distribution of types. Ackermann and Murphy (2019) examine the interplay between preferences and beliefs and how these are affected by the behavior of others in repeated public goods games. They embed a strategy game in each round of the repeated game to model changes in social preferences. Our design allows us to model learning mediated by stable social preferences.

The paper proceeds as follows. In Section 2, we discuss the experimental design, and in Section 3, we present a descriptive analysis of the data and classification of cooperative types. Section 4 describes the social preference model and estimated type transitions, examines the repeated game data to disentangle the roles of social preferences and learning in repeated public goods games and presents the results from the validation experiment. We close with conclusions in Section 5.

2 Experimental Design

The experiment is based on the Fischbacher and Gächter (2010) public goods design. A participant is placed in a group and given an endowment of tokens. He decides how many tokens from the endowment to keep and how many to put into a public good account that benefits all members of the group. Each token put in the public good account pays a marginal per capita return (MPCR) to each group member at a rate < 1, and each token kept pays 1. For each decision in the experiment, a group is comprised of three participants, and the endowment is 20 tokens. Tokens convert to a monetary payoff at the end of the experiment at a rate of 20 tokens = \$1. Each participant is asked to complete four tasks in the following order: P1, R1, P2, R2.

In the $P - task$, the participant completes a conditional contribution table in which he decides how much to contribute to the public good account for 21 possible average contribution amounts of his group members (e.g. 0, 1, 2, ..., 20). He also makes an unconditional contribution decision by choosing how many of the 20 tokens to contribute to the public good. Payoffs from this task are determined as follows. Once all decisions are completed by the group members, two group members are chosen at random. The unconditional contributions of those two group members are averaged together and rounded up or down to the nearest integer k. The contribution of the third group member is determined by the amount specified in the conditional contribution table for the average contribution amount of k by the other two group members. The total amount then contributed to the public good account is the sum of the unconditional contributions of the first two group members and the conditional contribution of the third. Participant payoffs are based on this. Everyone knows these procedures before making their decisions.

In the $R - task$, participants make an unconditional decision of how many tokens to put in the public good account. This decision is repeated over several rounds, and the members of a group are fixed for all rounds in a partners-matching protocol. After deciding how much to contribute to the public good account, a participant is asked to state his belief of the average contribution of the other two group members in the current round. Participants are paid for the accuracy of this stated belief.[2] At the end of a round, a participant is informed of the exact contribution of each group member, the average

[2]Specifically, the individual is asked to guess the average contribution of the two other group members rounded to the nearest integer. If the guess is exactly equal to the rounded average contribution of the other group members, an individual is earns a bonus payment of three tokens (\$0.15). If the guess deviates by only one point, payment is 2 tokens, and if the guess deviates by two points, payment is 1 token. If the guess is off by three points, no tokens are paid. The financial incentive to elicit beliefs is small to avoid hedging.

contribution, and his payoff for that round.

In sessions with finitely repeated games, the number of rounds is fixed at seven. In the sessions with indefinitely repeated games, there is at least one round and after that the probability of a subsequent round is 0.85. The continuation probability of 0.85 yields, on average, seven rounds of play, thus the finitely and indefinitely repeated games are comparable in the expected number of rounds.

Each session has exactly 15 participants and thus 5 groups. Participants are randomly reshuffled across groups before each of the four tasks, and participants are aware of this before making decisions. Tasks P1 and R1 use the same MPCR, and tasks P2 and R2 also use the same MPCR. The MPCR used for P1 and R1 is different than that used for P2 and R2. There are two MPCR rates: 0.4 (Low) and 0.8 (High). In five sessions, the low MPCR is used for P1 and R1 and the high MPCR is used for P2 and R2. In five sessions, this is reversed, to control for order effects. We did not find any significant order effects for average contributions in the P-tasks or average contributions in the R-tasks (Two Sample T-test, $p > 0.10$ in both cases).

Participants know there are four tasks, and the instructions for each task are distributed and read out loud prior to the start of each task. Participant instructions for sessions with indefinitely repeated games are in Appendix 6.5. These include screen shots of the decision screens participants used to make decisions during the experiment. The instructions for the sessions with finitely repeated games are identical except that they state the repeated games last exactly seven rounds. To make sure participants understand the decisions they are asked to make and how to calculate payoffs, they are given a short set of questions to complete prior to the start of Task P1. Answers to the questions are given out loud and any remaining queries are addressed before starting Task P1. All decisions are made on a computer, privately and anonymously. Participants are paid their earnings for all four tasks and all rounds within a task. Earnings are paid privately and in cash at the end of the session.

The experiments were run at George Mason University during September and October of 2014. Ten sessions were run, with a total of 150 participants. Six of the ten sessions involved indefinitely repeated games and four sessions involved finitely repeated games. No one participated in more than one session. Participants were recruited via email from a pool of students who had all previously registered to receive invitations for experiments. Each experimental session lasted for approximately 1.5 hours. Average participant earnings were \$26.36 (s.d. \$8.67).

3 Cooperation and Classification of Types

We classified participants into cooperative types using the statistical classification algorithm of Kurzban and Houser (2005). This algorithm uses a linear conditional-contribution profile (LCP) to determine a given participant's type. The LCP is the result of an ordinary least squares regression of a participant's conditional contribution in the P-task on each of the 21 possible average contributions of the other group members and a constant. If the estimated LCP is strictly below half of the endowment for every possible aver-

Table 1: Distribution of cooperative types in strategy games computed using the LCP method and average contributions across types

	P1		P2		MPCR (0.4)		MPCR (0.8)	
Type	N	Cont.	N	Cont.	N	Cont.	N	Cont.
FR	36 (24%)	2.39	54 (36%)	2.47	54 (36%)	2.53	36 (24%)	2.30
CC	89 (59%)	9.66	69 (46%)	9.69	71(47%)	9.62	87 (58%)	9.72
FC	17 (11%)	18.10	20 (13%)	17.53	18 (12%)	17.65	19 (13%)	17.92
NC	8 (5%)	10.45	7 (5%)	9.80	7 (5%)	10.61	8 (5%)	9.74
Total	150 (100%)	8.92	150 (100%)	8.14	150 (100%)	8.08	150 (100%)	8.98

Note: FR is free rider, CC is conditional cooperator, FC is full cooperator. Types classified by LCP method.

age contribution, then the participant is classified as a Free Rider (FR). A participant is classified as a Full Cooperator (FC) if the LCP lies at or above half the endowment everywhere. If the LCP has a positive slope and lies both above and below half the endowment then he is a Conditional Cooperator (CC). Any participant who does not fall into one of these three categories is classified as a Noisy Contributor (NC).[3]

The distribution of types identified with the classification algorithm and their average contributions in the P-tasks are presented in Table 1. We present the type distributions and average contributions in P1 and P2 tasks and also by low and high MPCR. Consistent with previous studies, most participants are classified as conditional cooperators, and roughly one quarter are classified as free riders.[4] These two types account for more than 80% of participants in the combined P1 and P2 tasks and also when separated by the high and low MPCR treatments. Full cooperators and noisy contributors are less frequent. The distributions of types in the P1 and P2 tasks are not significantly different (Chi-squared test, $\chi^2(2) = 4.36; p = 0.11$).[5] From the low to high MPCR, as the cost to cooperate declines, there is a shift in the distribution, with the proportion of full cooperators rising. However, we do not find that the overall distribution under two different MPCR levels is significantly different (Chi-Squared Test, $\chi^2(2) = 4.06; p = 0.13$). While we do not find any aggregate distributional differences at the two MPCR levels, we evaluate how many participant-level type transitions are consistent with a social preference model in Section 4.1.

Looking at the contribution behavior of each type, Table 1 shows that in P1, the average contribution across all conditional contributions is 2.39 tokens for free riders, 9.66 for conditional cooperators, 18.10 for full contributors and 10.45 for noisy contributors. In P2, the average contribution is 2.47 for free riders, 9.69 for conditional cooperators, 17.53 for full contributors, and 9.80 for noisy contributors. In both P1 and P2, average

[3]In Kurzban and Houser (2005), participants are classified into only the first three groups only and three participants who did not fall into these categories are excluded.

[4]Previous studies include Fischbacher et al. (2001); Fischbacher and Gächter (2010); Kurzban and Houser (2001, 2005); Burlando and Guala (2005); Duffy and Ochs (2009); Kocher et al. (2008); Herrmann and Thöni (2009); Muller et al. (2008); Bardsley and Moffatt (2007)

[5]We excluded participants who either transitioned to or transitioned from being a noisy contributor in comparing the distributions.

Table 2: Contributions over rounds

	Contribution		
	(1)	(2)	(3)
Constant	11.79***	10.03***	9.59***
	(0.92)	(0.63)	(0.60)
Round	-0.17***	-0.079*	-0.057*
	(0.049)	(0.037)	(0.029)
Finitely Repeated		0.34	1.32
		(1.38)	(1.17)
Finitely Repeated × Round		0.025	-0.048
		(0.12)	(0.12)
highMPCR		2,87***	3.68***
		(0.72)	(1.05)
Finitely Repeated × highMPCR			-1.89
			(1.13)
Obs	2220	2220	2220
Sessions	10	10	10
R^2	0.014	0.051	0.055
F	11.92	20.42	17.93
$Prob > F$	0.0072	0.0002	0.0002

Notes: $*p < 0.1, **p < 0.05, ***p < 0.01$. Robust standard errors clustered at the session level are reported.

contributions across types are significantly different (Kruskal-Wallis tests, $p < 0.0001$). Similarly, we find that average contributions are significantly different across types in both Low MPCR and High MPCR treatments (Kruskal-Wallis tests, $p < 0.0001$).

In the repeated game task, four sessions involved finitely repeated games of seven rounds each, and six sessions used indefinitely repeated games. The probabilistic continuation rule produced rounds of the following length for R1 and R2 across these six sessions: {{6,7}, {3,25}, {7,12}, {5,2}, {9,5}, {6,5}}. While the finite and indefinitely repeated games had different lengths, we find no significant difference in contributions across these games.[6] This is illustrated in the regression results in Table 2. Column (1) shows the results of the amount contributed to the public good controlling for round, and Column (2) adds a high MPCR dummy, a finitely repeated game dummy and an interaction term with round number. Contributions decline over rounds in both finitely and indefinitely repeated games, but there is no significant difference in contribution behavior across game types. The high MPCR results in higher contributions in both the finitely and indefinitely repeated games as shown in Column (3) of Table 2, and this is consistent with previous research (Isaac et al., 1984; Lugovskyy et al., 2017).

[6] This was also found by Lugovskyy et al. (2017).

4 Results

We present the results in three stages. First, we examine if social preferences can organize the observed changes in the cooperative disposition of participants when the costs of cooperation change. Second, we assess the roles of social preferences and learning in explaining behavior in the repeated game environment. Third, we provide an empirical test of the relationship between social preferences and learning using a follow-up "validation" experiment.

4.1 Social Preferences and Cooperative Type Transitions

Our design allows us to classify an individual into a cooperative type when the cost to cooperate is high (MPCR=0.4) and low (MPCR=0.8). Transitioning from one type to another when the cost changes does not imply unstable preferences per se. It could be consistent with a model of stable social preferences.

We use the Arifovic and Ledyard (2012) model of social preferences in the context of public goods games to explore this. Consider a group of size N and an MPCR of M. Each individual $i \in \{1, 2, ...N\}$ is endowed with w. The payoff an individual i receives by contributing c^i when others in his group contribute on average o can be written as $\pi^i(c^i, o)) = w - c^i + M(c^i + (N-1)o)$. Similarly, the average payoff of the group can be written as $\bar{\pi}(c^i, o) = w - \bar{c} + MN\bar{c}$, where $\bar{c} = \frac{c^i+(N-1)o}{N}$. The utility derived by individual i is:

$$u^i(c^i, o) = \pi^i(c^i, o) + \beta^i \bar{\pi}(c^i, o) - \gamma^i max\{0, \bar{\pi}(c^i, o) - \pi^i(c^i, o)\} \tag{1}$$

Where $\beta^i \geq 0; \gamma^i \geq 0$ are social preference parameters. $\beta^i > 0$ implies that individual i has a preference for a higher average payoff to all agents in the group and thus higher welfare of group members. In other words, β^i characterizes an individual's altruistic preference. $\gamma^i > 0$ implies that individual i obtains a disutility when his/her payoff is smaller than the average payoff of the group, i.e. when $\bar{\pi}(c, o) > \pi^i(c, o)$. γ^i captures the discomfort individual i faces when being taken advantage of by the group. $\beta^i = 0$ and $\gamma^i = 0$ indicates that individual i is purely selfish.

In equilibrium, individual i would choose a contribution c^i as follows:

$$c^i = \begin{cases} 0 & \text{if } 0 \geq \left(M - \frac{1}{N}\right)\beta^i + M - 1 & \text{(Free Rider)} \\ \bar{c} & \text{if } \gamma^i\left(\frac{N-1}{N}\right) \geq \left(M - \frac{1}{N}\right)\beta^i + M - 1 \geq 0 & \text{(Conditional Cooperator)} \\ w & \text{if } \gamma^i\left(\frac{N-1}{N}\right) \leq \left(M - \frac{1}{N}\right)\beta^i + M - 1 & \text{(Full Cooperator)} \end{cases} \tag{2}$$

The social preference parameters (β^i, γ^i), along with the parameters of the public goods game (N, M), determine the cooperative type, i.e. if individual i behaves as a free rider, a conditional cooperator or a pure altruist (full cooperator).[7]

[7]The social preference models of Fehr and Schmidt (1999) and Charness and Rabin (2002) also

Table 3: Cooperative type transitions when costs to cooperate change

	FR(0.8)	CC(0.8)	FC(0.8)
FR(0.4)	$\beta \in [0, 0.43]$ $\beta \in [0, 8.57]$	$\beta \in [0, 8.57]$ $\beta \in [0.42, 1.40\gamma + 0.42]$	$\beta \in [0, 8.57]$ & $\beta \in [1.40\gamma + 0.42, \infty]$
CC(0.4)	$\beta \in [0, 0.42]$ $\beta \in [8.57, 9.43\gamma + 8.57]$	$\beta \in [8.57, 9.43\gamma + 8.57]$ $\beta \in [0.42, 0.42 + 1.40\gamma]$	$\beta \in [8.57, 9.43\gamma + 8.57]$ $\beta \in [1.40\gamma + 0.42, \infty]$
FC(0.4)	$\beta \in [0, 0.42]$ $\beta \in [9.43\gamma + 8.57, \infty]$	$\beta \in [0.42, 0.42 + 1.40\gamma]$ $\beta \in [9.43\gamma + 8.57, \infty]$	$\beta \in [9.43\gamma + 8.57, \infty]$ $\beta \in [1.40\gamma + 0.42, \infty]$

Note: FR is freerider, CC is conditional cooperator, FC is full cooperator. MPCR of 0.4 and 0.8 listed in parentheses. Range of parameter values predicted by model for each transition listed in cells. Cells in blue not predicted by model.

There are nine possible tranistions from one of the three types (FR, CC and FC) to another type when the cost to cooperate changes. While all transitions are possible, not all are predicted by the model. Table 3 shows all possible transitions as the MPCR changes from 0.4 to 0.8. The table includes the range of parameter values that an individual would need to have to be consistent with that type transition. The three cells in blue are type transitions and parameter ranges that are not predicted by the model. For example, the model predicts an individual could behave as a conditional cooperator when the MPCR is 0.4 and a full cooperator when the MPCR increases to 0.8, however, such an individual is not predicted to become a free rider.

Using the conditional choices in the P-task, we classify each participant into a type using the LCP method described in Section 2. This is done twice, once when the MPCR is 0.4 and once when the MPCR is 0.8. Thus, we observe each participant's type transition from low to high MPCR. Table 4 shows the distribution of the nine possible type transitions. We do not show the cases where the transition involves a noisy type.[8] Of the remaining 137 participants, there are 12 (8.7%) whose type transition is not consistent with the social preference model, and these are listed in blue in the table. This means that over 90% of participants respond to changes in the costs to cooperate in a way that is consistent with having stable social preferences.

In sum, the model of social preferences seems to be reasonably successful at organizing changes in cooperative tendencies. Almost all changes in observed behavior when the cost to cooperate increases can be rationalized by stable social preferences.

4.2 Disentangling the Roles of Social Preferences and Learning

We now turn to an examination of choices in repeated interactions with the same members of a group. Learning in repeated public goods games has been shown to be individual-

capture the effect of MPCR (M) on the cooperative type of an individual, and they suggest behavior is independent of group size. Thus, they do not pick up the variation in observed contributions across small and larger groups as identified in Isaac et al. (1994); Isaac and Walker (1988).

[8]Seven participants are classified as noisy in the low MPCR treatment, and 8 in the high MPCR treatment. There are 13 participants who are classified as noisy in at least one of the treatments.

Table 4: Number of participants by type transition

	FR(0.8)	CC(0.8)	FC(0.8)
FR(0.4)	29	20	2
CC(0.4)	5	59	5
FC(0.4)	2	5	10

Note: FR is freerider, CC is conditional co-operator, FC is full cooperator. MPCR of 0.4 and 0.8 listed in parentheses. Types classifed by LCP method. Cells in blue not predicted by model.

level directional learning based on social preferences (Anderson et al., 2004; Wendel & Oppenheimer, 2010; Cooper & Stockman, 2002; Janssen & Ahn, 2006; Arifovic & Ledyard, 2012). More recent experimental studies have argued that social preferences may not be needed to explain cooperation in repeated games and that individuals learn based on payoffs (Burton-Chellew & West, 2013; Burton-Chellew et al., 2015). Pure payoff-based learning, in which individuals learn to contribute zero, however, cannot reconcile the different average contributions across cooperative types as observed by Fischbacher and Gächter (2010).

We use the data generated from the experiments to test alternative theories of how social preferences and learning interact in public goods games. It is possible that aggregate differences across types arise simply because types start at different levels of initial contributions and use a learning rule to update their contributions in later rounds. This stands in contrast with the assumption in earlier literature that the aggregate differences across types arise because types start with random contributions and learn their equilibrium contributions over time (Anderson et al., 2004; Wendel & Oppenheimer, 2010; Cooper & Stockman, 2002; Janssen & Ahn, 2006; Arifovic & Ledyard, 2012). Since we observe the cooperative types of participants in the strategy task before they participate in repeated games, it is straightforward to test if contributions in the first round of a repeated game are purely random or align with types.

To begin, we compare the first-round contributions in repeated games and the unconditional contributions in strategy games and find they are not significantly different from each other (Paired Hotelling's T^2 test, $p = 0.118$). Across types, however, contributions are significantly different (Kruskal Wallis Tests: $p < 0.0001$ in all cases). Figure 1 presents the average contribution by cooperative type in the first round of a repeated game and the average unconditional contribution by cooperative type in the corresponding strategy game. Cooperative type information for P1 and R1 is computed using the conditional responses in P1 and the cooperative type information for P2 and R2 is computed using the conditional responses in P2. The figure shows a significant increasing trend in first-round contributions in the repeated games across free riders, conditional cooperators and full cooperators (R1: Jonckheere-Terpstra Test, $p < 0.0001$; R2: Jonckheere-Terpstra Test, $p < 0.0001$). The trend is also significant for unconditional contributions in the strategy games (P1: Jonckheere-Terpstra Test, $p < 0.0001$; P2: Jonckheere-Terpstra Test, $p < 0.0001$).

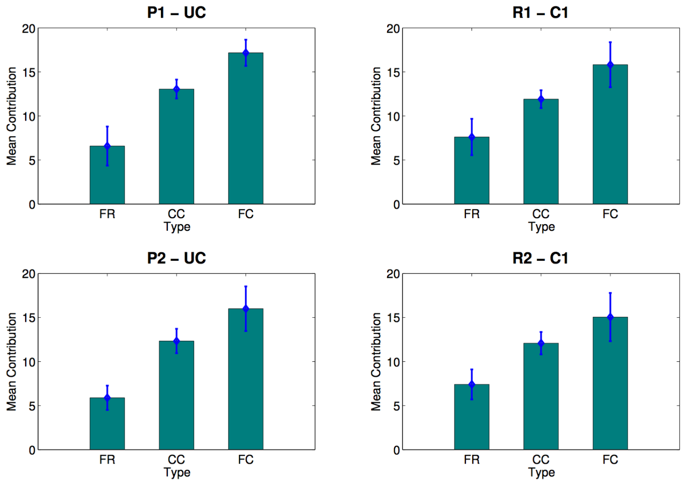

Figure 1: Unconditional contributions (UC) across types in P-tasks and the first-round contributions (C1) across types in R-tasks. Error bars represent 95% confidence intervals.

Thus, first-round contributions in the repeated game are similar to unconditional contributions in the corresponding strategy game, and they are informed by individual-level social preferences uncovered from the strategy game. In the first round of a repeated game, free riders tend to start with contributions that are significantly smaller on average (though not exactly zero), and full contributors start with significantly higher contributions on average (though not exactly full endowment). The remaining participants start around half of the endowment. In summary, participants make their first-round contributions based on their social preferences rather than starting with random contributions.

While it is clear that participants' first-round contributions are aligned with their social preferences, there are two potential explanations for how average contributions emerge across cooperative types in repeated public goods games. First, participants start with contribution levels that are aligned with their cooperative type and then learn based on payoffs. Second, participants start with contribution levels that are aligned with their cooperative type, then they learn their corresponding equilibrium contribution based on their social preferences. We use data from the repeated games and social preferences estimated from the strategy game to assess these two explanations. To be able to systematically disentangle the roles of social preferences and learning in repeated games, we consider three models that involve different assumptions about how learning and social preferences interact to produce contributions across cooperative types. These three models are summarized below:

- **RAND-UTIL**: In this model, individuals start with random contributions but learn based on utilities. In the learning phase, individuals use the utilities computed using their social preference parameters to learn their equilibrium contributions over the time. This is our baseline model (see Anderson et al. (2004); Wendel and Oppenheimer (2010); Cooper and Stockman (2002); Janssen and Ahn (2006); Arifovic and Ledyard (2012)).

- **SP-PAYOFF**: This model posits that individuals start their contributions according to their social preferences. Therefore, first-round contributions are in line with their cooperative type and modeled as a best response to their elicited beliefs in the first round given the distribution of social preferences estimated from P-tasks or strategy games. In the learning phase, individuals use simple payoffs in each round to update attractions of strategies.

- **SP-UTIL**: This model is identical to SP-PAYOFF except that learning is based on utilities computed using social preferences. Therefore, individuals learn their cooperative type specific equilibrium contribution over time.

Learning is modeled using Reinforcement Learning with Loss Aversion (REL). The REL model was proposed by Erev, Bereby-Meyer, and Roth (1999) to circumvent the problems faced by the reinforcement model of Erev and Roth (1998) in explaining behavior in games when a constant is added to all payoffs. It introduced two modifications to the original reinforcement model: sensitivity to payoff variability and insensitivity to payoff magnitude. The model assumes that strategies (contribution levels) have propensities and propensities are linked to the probabilities of choice using a choice rule. The learning algorithm of a given model dictates how the propensities of strategies are updated. Propensities of strategies are often referred to as attractions. There are two free parameters to be estimated in the REL model: the attraction sensitivity parameter λ and the strength of initial attractions $N(1)$. A full description of the REL model is provided in Appendix 6.1.

Social preference parameters and learning parameters are modeled as heterogeneous across individuals. The distribution of the learning parameters $(\lambda, N(1))$ is modeled as log-normal to capture heterogeneity across individuals. The distribution of social preference parameters (β, γ) is also modeled as log-normal distribution, since they cannot be negative, and uses choices in the P-tasks for estimation. In addition to social preference parameters, we also estimate a noise parameter ω which represents a random choice component of behavior.[9] ω is a probability between [0,1] and is modeled as a logistic-normal distribution in the estimation framework. The parameters are estimated using the simulated maximum likelihood method. A full description of the econometric framework used for the estimation is presented in Appendix 6.2.

The social preference parameter estimates from the strategy games are reported in Table 5. The median values of the social preference parameters are $(\beta = 10.23, \gamma = 11.45)$ indicating a median individual is a conditional cooperator (according to Equation 2). The median value of random choice parameter (ω) is 0.12 implying only 12% choices align with random selection of contribution levels. This indicates a good understanding of the decision context by participants.

The estimated distributions of (β, γ, ω) are used to compute the likelihood of first-round choices in the SP-UTIL and SP-PAYOFF models. Table 6 reports the parameter estimates of the REL model and the fit of RAND-UTIL, SP-PAYOFF and SP-UTIL. Models

[9] ω accounts for lack of understanding of the strategic environment or lack of attention that can lead to random choice behavior.

Table 5: Distributions of estimated social preferences (β, γ) and random choice propensity (ω) using data from strategy games.

	β	γ	ω
μ	2.33	2.44	-1.99
σ	3.09	1.79	3.34
N	150		
LL	-13625		

Notes: Parameter estimates are the mean and standard deviations of the underlying normal distribution before transformation. For example, $\beta \sim LN(2.32, 3.09)$.

Table 6: Fit statistics and estimated parameters for behavioral specifications

Model	Parameter Estimates		LL	AIC	BIC
SP-UTIL	$\lambda \sim LN(0.43, 0.28)$	$N(1) \sim LN(-2.66, 10.22)$	-5045	10098	10110
SP-PAYOFF	$\lambda \sim LN(0.31, 0.32)$	$N(1) \sim LN(-1.94, 4.99)$	-5049	10106	10118
RAND-UTIL	$\lambda \sim LN(0.36, 0.37)$	$N(1) \sim LN(-4.09, 9.02)$	-5100	10208	10220
RAND‡			-6759	13518	13518

Notes: ‡ RAND is the random choice model. According to this model any choice has a probability of $\frac{1}{21}$ of being chosen. $AIC = 2k - 2\hat{LL}$ and $BIC = kln(N) - 2\hat{LL}$, where k is the number of parameters of the model and N is the number of observations. All models involve estimating 4 parameters for the learning model using the data from 150 individuals.

whose first-round contributions are determined by social preferences achieve higher likelihood compared to the models that use random contributions in the first round. This aligns with our descriptive analysis earlier in this subsection. Table 7 reports two-sided p-values computed using the Vuong's test for each pair of the three models. SP-UTIL and SP-PAYOFF perform equally well in explaining contributions. That is, it does not seem to matter whether payoffs or utilities are used to model updating attractions in the REL learning model. Also of note is that they both outperform RAND-UTIL model.

These results highlight that initial contributions determined by social preferences combined with high inertia associated with REL learning model explain the aggregate differences in contributions across cooperative types. This result stands in contrast to the assumption that individuals start with random contributions and then move towards their equilibrium contributions based on learning with utilities derived from social preferences (Cooper & Stockman, 2002; Janssen & Ahn, 2006; Wendel & Oppenheimer, 2010; Arifovic & Ledyard, 2012).

Table 7: Two-sided Significances in Favor of Row Models Using the Vuong Test

	SP-PAYOFF	RAND-UTIL	RAND
SP-UTIL	0.667	0.024	0.000
SP-PAYOFF		0.037	0.000
RAND-UTIL			0.000

That there is no significant difference between the fit of SP-PAYOFF and SP-UTIL indicates that once initial-round contributions are determined by the social preferences of an individual, there is no additional information to be gained about the path of contributions using social preference based utilities to update the attractions of strategies rather than simple payoffs within the repeated game. In other words, social preferences matter insofar as they determine the first-round contributions and then individuals behave as if they are solely payoff-based learning.[10]

4.3 Validation Experiment: manipulating the cost to cooperate in the initial round influences contributions in subsequent rounds

Thus far, our analysis has uncovered three findings. First-round contributions reflect an individual's social preferences and later round contributions reflect payoff-based reinforcement learning. Choices mediated by social preferences imply that as the cost to cooperate declines contributions should increase. Changes in the classification of an individual's cooperative type are monotonic, in that as the cost to cooperate declines an individual should become no less cooperative.

In this subsection, we stress test a novel behavioral implication from these findings with data from an additional experiment. Taken together, our findings imply that lowering the cost to cooperate in the first round should increase subsequent contributions and these should reflect payoff-based reinforcement learning. The question we ask is, in the repeated games, could a one-time, initial lower cost to cooperate set the stage for higher cooperation over time, even if the subsequent cost to cooperate increases?

We answer this question with an additional experiment designed to assess the predictive validity of the behavioral results. We note that, if the first-round decisions of individuals are random and then they learn their equilibrium strategies, any change in the cost to cooperate in the first round would have no effect.

There are two treatments in this experiment, and details are presented in Table 8. The experiments were conducted at George Mason University between February and April of 2016. Ten sessions were conducted, and there were 160 participants. Participants made their decisions privately and anonymously. Five of the ten sessions specified a high cost to cooperate (e.g. a low first-round MPCR of 0.26) in the first round, and the remaining five sessions specified a low cost to cooperate (e.g. high first-round MPCR of 0.99). In the subsequent rounds 2-10, the cost to cooperate was identical (MPCR of 0.5) in both treatments. Each group had four members, and groups remained fixed for all ten rounds. Each experimental session lasted for approximately 1.5 hours. Average participant earnings were $22.27. Experimental instructions for the high first-round MPCR treatment are provided in Appendix 6.6.[11] These include screen shots of the decision

[10]It is interesting to note that a SP-PAYOFF model can also explain the restart effect found in the experiments of Andreoni (1988). When the repeated game is restarted, individuals appear to start again with initial contributions aligned with their social preferences and then learn based on payoffs.

[11]The instructions for the low first-round MPCR treatment are similar except that first-round MPCR

TABLE 8: VALIDATION EXPERIMENT TREATMENTS

Treatment	High first-round MPCR	Low first-round MPCR
First-round MPCR	0.99	0.26
$MPCR_{2-10}$	0.5	0.5
Group Size	4	4
No. of Rounds	10	10
Matching	Partners	Partners
No. of Sessions	5	5
Session sizes	{20, 20, 20, 16, 12}	{20, 16, 16, 8, 12}
No. of Subjects	88	72

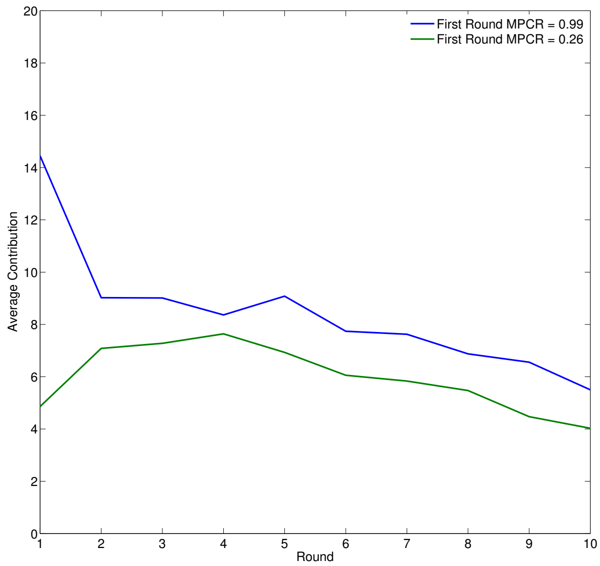

FIGURE 2: AVERAGE CONTRIBUTIONS IN THE TWO TREATMENTS OVER THE TEN ROUNDS OF THE PUBLIC GOODS GAME

screens participants used to make decisions during the experiment.

Average contributions over the ten rounds in both experimental treatments are shown in Figure 2. As expected, first-round contributions are much higher when the cost to cooperate is low (14.2 tokens) compared to when the cost to cooperate is high (5 tokens). We are interested in what happens to cooperation in the subsequent rounds when the cost is identical across treatments. Cooperation responds to the change in cost in the ways we would expect. As the cost increases, relative to the first round, cooperation declines, and vice versa. However, in the treatment where it was inexpensive to cooperate initially, the average contribution in rounds 2-10 is 7.75 tokens (s.d 4.11), and in the treatment where it was expensive to cooperate initially, the average contribution in rounds 2-10 is 6.09 tokens (s.d. 5.14). The difference between the two treatments is 1.66 tokens and

is 0.26. We note that the distribution of cooperative types derived from Task-P in the validation experiment with MPCR=0.5 is similar to that reported in Table 3 for MPCR=0.4. Since the prediction from the model is on repeated game behavior, we do not report the distribution.

significantly different from zero ($t(158) = 2.28, p = 0.024$). This implies groups that faced a low cost to cooperate in the first round managed to contribute 27% more to the public good in subsequent rounds than those that faced a high initial cost. This initial subsidy to cooperation also resulted in higher payoffs for group members. It was 3.82 times more expensive in round 1 to pay for cooperation in the low-cost treatment, but earnings were 33% higher per round in the subsequent rounds compared to the high-cost treatment.

These findings confirm that by just making cooperation less or more costly when a group first interacts with one another, subsequent cooperation can be significantly influenced. This suggests that it might be worth creating an environment in which it is easier to cooperate with one another when a team or group initially forms, as this sets the pathway for higher levels of cooperation as the group continues to interact in the future. These results validate that the behavioral specification obtained in this paper not only explains the observed behavioral regularities from the main experiment but also validates novel predictions in additional studies.

5 Conclusions

We characterize the decision to cooperate in a public goods experiment with a model of social preferences and payoff-based learning. The experimental design allows us to observe individual-level decisions in one-shot and repeated public goods games at different costs to cooperate. Using the one-shot game data, we show that the classification of participants into cooperative types at different costs to cooperation is consistent with a model of stable social preferences. Cooperative behavior in repeated games is most consistent with individuals making choices based on their social preferences in the first round and subsequently making choices based on payoff-based reinforcement learning in subsequent rounds.

We validate the predictive capacity of this behavioral specification using a follow-up experiment. We find that just by manipulating the price of cooperation in the first round, average contributions in later rounds are strongly affected. If the initial cost to cooperate with group members in a social dilemma is low, cooperation within the group in later interactions can be sustained at a higher level. This highlights the importance of how a group interacts initially. If the costs to cooperate are lowered the first time a group meets, they respond to this and cooperate more. This then sets the stage for future cooperation. Groups situated in an environment in which it was easy to cooperate with one another in the first interaction contributed 27% more to the public good for the remainder of the game, and made 33% higher earnings, than groups that faced a high initial cost to cooperate.

Our findings point to the importance of social preferences and learning in understanding cooperative behavior in groups. Individuals bring their innate willingness to cooperate to interactions with others but also respond to and learn from the behavior of those around them as they continue to interact. Importantly, our results suggest that how the initial stage is set for group interactions that would benefit from cooperation can be influential. Starting off on the right foot in terms of cooperating with others can set the group on a mutually beneficial pathway.

References

Ackermann, K. A., & Murphy, R. O. (2019). Explaining cooperative behavior in public goods games: How preferences and beliefs affect contribution levels. *Games*, *10*(1), 15.

Anderson, S. P., Goeree, J. K., & Holt, C. A. (2004). Noisy directional learning and the logit equilibrium. *The Scandinavian Journal of Economics*, *106*(3), 581–602.

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

Arifovic, J., & Ledyard, J. (2012). Individual evolutionary learning, other-regarding preferences, and the voluntary contributions mechanism. *Journal of Public Economics*, *96*(9), 808–823.

Bardsley, N., & Moffatt, P. (2007). The experimetrics of public goods: Inferring motivations from contributions. *Theory and Decision*, *62*(2), 161–193.

Bhat, C. R. (2001). Quasi-random maximum simulated likelihood estimation of the mixed multinomial logit model. *Transportation Research Part B: Methodological*, *35*(7), 677–693.

Bolton, G., & Ockenfels, A. (2000). ERC: A theory of equity, reciprocity, and competition. *American Economic Review*, 166–193.

Boosey, L., Isaac, R. M., Norton, D., & Stinn, J. (2019). Cooperation, contributor types, and control questions.

Burlando, R., & Guala, F. (2005). Heterogeneous agents in public goods experiments. *Experimental Economics*, *8*(1), 35–54.

Burton-Chellew, M. N., Nax, H. H., & West, S. A. (2015). Payoff-based learning explains the decline in cooperation in public goods games. *Proceedings of the Royal Society of London B: Biological Sciences*, *282*(1801), 20142678.

Burton-Chellew, M. N., & West, S. A. (2013). Prosocial preferences do not explain human cooperation in public-goods games. *Proceedings of the National Academy of Sciences*, *110*(1), 216–221.

Cartwright, E., & Lovett, D. (2014). Conditional cooperation and the marginal per capita return in public good games. *Games*, *5*(4), 234–256.

Charness, G., & Rabin, M. (2002). Understanding social preferences with simple tests. *Quarterly Journal of Economics*, 817–869.

Cooper, D. J., & Stockman, C. K. (2002). Fairness and learning: an experimental examination. *Games and Economic Behavior*, *41*(1), 26–45.

Cox, J. C., Friedman, D., & Sadiraj, V. (2008). Revealed altruism. *Econometrica*, *76*(1), 31–69.

Diederich, J., Goeschl, T., & Waichman, I. (2016). Group size and the (in) efficiency of pure public good provision. *European Economic Review*, *85*, 272–287.

Duffy, J., & Ochs, J. (2009). Cooperative behavior and the frequency of social interaction. *Games and Economic Behavior*, *66*(2), 785–812.

Embrey, M., Frechette, G., & Yutsel, S. (2018). Cooperation in the finitely repeated prisoner's dilemma. *Quarterly Journal of Economics*, *133*(1), 509–551.

Erev, I., Bereby-Meyer, Y., & Roth, A. E. (1999). The effect of adding a constant to all payoffs: experimental investigation, and implications for reinforcement learning models. *Journal of Economic Behavior & Organization*, *39*(1), 111–128.

Erev, I., & Roth, A. E. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review*, 848–881.

Fehr, E., & Schmidt, K. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, *114*(3), 817–868.

Fischbacher, U., & Gächter, S. (2010). Social preferences, beliefs, and the dynamics of free riding in public goods experiments. *American Economic Review*, *100*(1), 541-56.

Fischbacher, U., Gächter, S., & Fehr, E. (2001). Are people conditionally cooperative? evidence from a public goods experiment. *Economics Letters*, *71*(3), 397–404.

Friedman, D., & Oprea, R. (2012). A continuous dilemma. *American Economic Review*, *102*(1).

Goeree, J. K., Holt, C. A., & Laury, S. K. (2002). Private costs and public benefits: unraveling the effects of altruism and noisy behavior. *Journal of public Economics*, *83*(2), 255–276.

Herrmann, B., & Thöni, C. (2009, March). Measuring conditional cooperation: a replication study in russia. *Experimental Economics*, *12*(1), 87-92.

Isaac, R. M., & Walker, J. M. (1988). Group size effects in public goods provision: The voluntary contributions mechanism. *The Quarterly Journal of Economics*, 179–199.

Isaac, R. M., Walker, J. M., & Thomas, S. H. (1984). Divergent evidence on free riding: An experimental examination of possible explanations. *Public Choice*, *43*(2), 113–149.

Isaac, R. M., Walker, J. M., & Williams, A. W. (1994). Group size and the voluntary provision of public goods: experimental evidence utilizing large groups. *Journal of Public Economics*, *54*(1), 1–36.

Janssen, M. A., & Ahn, T.-K. (2006). Learning, signaling, and social preferences in public-good games. *Ecology and Society*, *11*(2), 21.

Kocher, M. G., Cherry, T., Kroll, S., Netzer, R. J., & Sutter, M. (2008). Conditional cooperation on three continents. *Economics Letters*, *101*(3), 175–178.

Kurzban, R., & Houser, D. (2001). Individual differences in cooperation in a circular public goods game. *European Journal of Personality*, *15*(S1), S37–S52.

Kurzban, R., & Houser, D. (2005). Experiments investigating cooperative types in humans: A complement to evolutionary theory and simulations. *Proceedings of the National Academy of Sciences*, *102*(5), 1803.

Ledyard, J. O. (1995). Public goods: some experimental results. In J. Kagel & A. Roth (Eds.), *Handbook of experimental economics.* Princeton, NJ: Princeton University Press.

Liu, Y.-H., & Mahmassani, H. S. (2000). Global maximum likelihood estimation procedure for multinomial probit (mnp) model parameters. *Transportation Research*

Part B: Methodological, *34*(5), 419–449.

Loomes, G. (2005). Modelling the stochastic component of behaviour in experiments: Some issues for the interpretation of data. *Experimental Economics*, *8*(4), 301–323.

Lugovskyy, V., Puzzello, D., Sorensen, A., Walker, J., Williams, A., & Hall, W. (2017). An experimental study of finitely and infinitely repeated linear public goods games. *Games and Economic Behavior*, *102*, 286–302.

Moffatt, P. G., & Peters, S. A. (2001). Testing for the presence of a tremble in economic experiments. *Experimental Economics*, *4*(3), 221–228.

Muller, L., Sefton, M., Steinberg, R., & Vesterlund, L. (2008). Strategic behavior and learning in repeated voluntary contribution experiments. *Journal of Economic Behavior & Organization*, *67*(3), 782–793.

Train, K. E. (2009). *Discrete choice methods with simulation.* Cambridge university press.

Wendel, S., & Oppenheimer, J. (2010). An agent-based analysis of context-dependent preferences. *Journal of Economic Psychology*, *31*(3), 269–284.

Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data.* MIT press.

6 Appendix

6.1 REL Model

We refer to strategy j available to individual i as s_i^j. There are m_i strategies available to individual i. The chosen strategy of i in round t is denoted as $s_i(t)$. $u_i(s_i^j, s_{-i}(t))$ is the payoff obtained by i in round t when it chooses s_i^j and others choose strategies given in $s_{-i}(t)$. In the payoff-based learning models, $u_i(s_i^j, s_{-i}(t))$ is simply the number of tokens received, whereas in the utility based learning models $u_i(s_i^j, s_{-i}(t))$ corresponds to the utility function in 1 which is based on social preferences. When the learning is based on tokens received individuals learn free riding equilibrium strategy and when it is based on utilities as specified in 1 different cooperative types are learning their equilibrium strategy as specified in 2. The attraction of strategy j of agent i in round t is denoted as $A_i^j(t)$. The three components of the REL model are as below:

- **Initial Attractions:** In the first round, attractions of all strategies are zero. Therefore, $A_i^j(1) = 0$. Thus in the first round, all choices have equal probabilities for being chosen. In the RAND-UTIL model, these initial attractions are used to choose a random initial contribution level. In SP-UTIL and SP-PAYOFF the first-round contributions are instead chosen based on social preferences and beliefs. The attractions in the learning model are used only from the second round onwards.

- **Attraction Updating:** Strategy j of agent i has attraction $A_i^j(t)$ in round t:

$$A_i^j(t) = \begin{cases} \left[\frac{A_i^j(t-1)[C_i^j(t-1)+N(1)]+u_i(s_i^j, s_{-i}(t-1))}{[C_i^j(t-1)+N(1)+1]}\right] & \text{if } s_i^j = s_i(t-1) \\ A_i^j(t-1) & \text{otherwise} \end{cases}$$

  where $C_i^j(t)$ is the number of times s_i^j has been chosen in the first t rounds and $N(1)$ is a free parameter that determines the strength of the initial attractions. A large $N(1)$ means that effect of actual payoffs in later rounds on attractions will be smaller. The attractions of unchosen strategies are not updated.

- **Stochastic Choice Rule:** Probability that agent i chooses strategy j in round t:

$$P_i^j(t) = \frac{e^{\frac{\lambda}{PV_i(t)} A_i^j(t)}}{\sum_{l=1}^{m_i} e^{\frac{\lambda}{PV_i(t)} A_l^j(t)}}$$

  where λ_i is a free parameter that determines the reinforcement sensitivity of the individual i. $PV_i(t)$ is the measure of payoff variability.

  The payoff variability is updated according to:

$$PV_i(t) = \frac{[PV_i(t-1)(t-1+m_iN(1)) + |u_i(s_i(t-1), s_{-i}(t-1)) - PA_i(t-1)|]}{[t+m_iN(1)]}$$

where $PA_i(t)$ is the accumulated payoff average in round t and m is the number of strategies. $PV_i(1) > 0$ is initialized as 1.

$PA_i(t)$ is calculated in a similar manner.

$$PA_i(t) = \frac{[PA_i(t-1)(t-1+m_iN_i(1)) + u_i(s_i(t-1), s_{-i}(t-1))]}{[t+m_iN_i(1)]}$$

We initialized $PA_i(1)$ as initial attraction of any strategy which is 0.

6.2 Econometric Framework

In this section, we formulate structural econometric models of discrete choice that can be estimated by maximum likelihood to estimate the social preference parameters and learning parameters allowing for heterogeneity. This is the appropriate approach when using the data generated in our experiments because choices are made on a discrete scale.

6.3 Estimation of Social Preferences from the Strategy Games

Based on the specification of the utility function that includes social preferences in Equation 1, we formulate structural econometric models of discrete choice that can be estimated by maximum likelihood. Let T_i is the number of decision situations an individual i has faced in the strategy games. Let $C_i = \{c_{it} | t \in \{1, 2, ..., T_i\}\}$ be the vector of observed contributions of the individual i and $O_i = \{o_{it} | t \in \{1, 2, ..., T_i\}\}$ be the vector of the average contributions of i's group members (excluding i itself). The average contributions of other group members are stated explicitly in the P-tasks.

We begin developing the econometric model by assuming the participants' decisions reflect maximization of utility function with social preferences specified in Equation 1. In the absence of any errors in decision making, for a given level of average contribution of others in his/her group, a participant chooses a contribution that maximizes his/her utility.

As a first step to allow for stochastic decision making, we add a standard extreme value distributed error term to the utility derived from each level of contribution. Assuming that these errors are independent of other parameters of the model and regressors, we obtain the logit probability of choosing a contribution level c_{it} as:

$$l_{it} = \frac{e^{U^{it}(c_{it}, o_{it})}}{\sum_{j=0}^{20} e^{U^{it}(j, o_{it})}} \tag{3}$$

The errors from a standard extreme value distribution that are added to the utility capture the idea that a participant's computation of subjective utility may be subject to some variability (Loomes, 2005). In addition to these errors, a number of experimental studies involving public goods games also found an evidence for the so-called "trembles" or "random choice errors" (Bardsley & Moffatt, 2007; Moffatt & Peters, 2001). These "trembles" account for a participant' failure to understand the decision problem or attention lapses during decision making. We model the propensity of a participant to chose randomly in any given task using a "trembling hand" parameter. A participant's tendency to make a random choice in a decision situation is given by a parameter ω_i. Since each participant is endowed with 20 tokens, the probability of choosing a given contribution level via random choice is $\frac{1}{21}$.

Then, for the participant i, the probability of observed level of contribution c_{it} for an average contribution o_{it} of his/her group members can be written as:

$$l_{it}(c_{it}, o_{it}, \beta_i, \gamma_i, \omega_i) = (1-\omega_i)\frac{e^{U^{it}(c_{it},o_{it})}}{\sum_{j=0}^{20} e^{U^{it}(j,o_{it})}} + \frac{\omega_i}{21} \tag{4}$$

A random coefficient model is employed to estimate the distribution of the individual-specific structural parameters β_i, γ_i and ω_i in the population. This has a better justification than doing a separate estimation for each participant since the number of observed choices will be rather small.

Since β, γ are constrained to be positive, we model them using log-normal distributions. To bound ω between 0 and 1, we model it as a logistic-normal distribution over [0 1]. For a concise notation define,

$$\eta_i = g_\eta(X_i^\eta \delta^\eta + \xi_i^\eta), \eta_i \in \{\beta_i, \gamma_i, \omega_i\} \tag{5}$$

η_i denotes one of the three individual specific parameters, X_i^η are $1 \times K^\eta$ vectors of regressors, δ^η are $K^\eta \times 1$ parameter vectors, and ξ_i^η are the unobserved heterogeneity components of the parameters. The first element of each X_i^η contains 1. The functions $g_\eta(.)$ impose theoretical restrictions on the individual specific parameters. For β, γ it is the exponential function ensuring that they are positive. For ω, it is the logistic distribution function ensuring that ω is always between zero and one. $g(X_i\delta + \xi_i)$ stands for a vector of the three functions.

We assume that $\xi_i = (\xi_i^\beta, \xi_i^\gamma, \xi_i^\omega)'$ follows a jointly normal distribution with a diagonal covariance matrix Σ independent of the regressors. The regressor matrix contains only ones in the estimations considered in this paper.

The likelihood contribution of participant i can be written as:

$$l_i = \int_{\mathbb{R}^3} \left[\prod_{t=1}^{T_i} l_{it}(c_{it}, o_{it}, g(X_i\delta + \xi)) \right] \phi(\xi) d\xi \tag{6}$$

where l_{it} is the probability given in Equation 4 and $\phi(.)$ denotes the density of multivariate normal ξ. The above integral does not have a closed form solution. We approximate it using $R = 1000$ Halton draws from ξ to obtain simulated likelihood (Train, 2009; Bhat, 2001). The simulated likelihood contribution of participant i is:

$$sl_i = \sum_{r=1}^{R} \frac{l_i(\xi_r)}{R} \tag{7}$$

The (simulated) log-likelihood is given by the sum of the logarithms of sl_i over all respondents in the sample. We maximized the log-likelihood function of entire sample using a two-step hybrid approach a multiple number of times as discussed in Liu and Mahmassani (2000) to avoid local maxima.[12] The variance-covariance matrix of the parameter estimates is computed using the *sandwich estimator* (Wooldridge, 2010). Standard errors are calculated using the sandwich estimator and treating all of each participant's choices as a single super-observation, that is, using degrees of freedom equal to the number of participants rather than the number of participants times the number of choices made. Standard errors for transformed parameters are calculated using the delta method.

6.4 Estimation of the Learning Model

Here, we describe in detail the estimation method used in the context of REL with parameter heterogeneity. The REL model has two structural parameters: the attraction sensitivity parameter λ_i and the parameter that defines the initial strength of attractions $N_i(1)$. We model both of them as random coefficients allowing for individual level heterogeneity. Using the notation in the previous subsection:

$$\eta_i = g_\eta(X_i^\eta \delta^\eta + \xi_i^\eta), \eta_i \in \{\lambda_i, N_i(1)\} \tag{8}$$

η_i denotes one of the two individual specific parameters, X^{η_i} are $1 \times K^\eta$ vectors of regressors, δ^η are $K^\eta \times 1$ parameter vectors, and ξ_i^η are the unobserved heterogeneity components of the parameters. The first element of each X_i^η contains 1. Since both $\lambda_i, N_i(1)$ are positive, we used the exponential function for g_η for both parameters. Assuming that $\xi_i = (\xi_i^\lambda, \xi_i^{N(1)})'$ follows a jointly normal distribution with a diagonal covariance matrix Σ independent of the regressors, the likelihood contribution of participant i can be written as:

$$l_i = \int_{\mathbb{R}^2} \left[\prod_{t=1}^{T_i} \left(\sum_{l=0}^{20} \mathbb{I}(c_{it}, l) P_i^l(t) \right) \right] \phi(\xi) d\xi \tag{9}$$

Where

$$P_i^l(t) = \frac{e^{\frac{\lambda_i}{PV_i(t)} A_i^l(t)}}{\sum_{k=0}^{20} e^{\frac{\lambda_i}{PV_i(t)} A_i^k(t)}} \tag{10}$$

is the probability of choosing contribution level l in round t. c_{it} is the observed contribution in t. $\mathbb{I}(c_{it}, l) = 1$ if $c_{it} = l$, 0 otherwise. T_i is the total number of rounds in the repeated games R1 and R2 that individual i has participated in (it should be noted that

[12] In the first step, we have employed a genetic algorithm to find parameters that maximize log-likelihood of the sample. Genetic algorithms are very effective in searching many peaks of likelihood function based on a rich "population" of solutions and thus reduce the probability of trapped into a local maximum. Since they do not require gradients to be computed, they are computationally very efficient for a global search of the parameters. In the second step, we used the solution of genetic algorithm as a starting point to Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm with numerical derivatives to maximize the log-likelihood function.

attractions, payoff variability, and accumulated payoff average in REL model will be reinitialized at the start of the R2). The integral in Equation 9 is computed using simulation and the total log-likelihood of the sample is computed as the sum of the logarithms of simulated individual level likelihoods of all respondents. We maximized the log-likelihood function of the entire sample using a two-step hybrid approach (described in the previous subsection) a multiple number of times to avoid local maxima. The variance-covariance matrix of the parameter estimates is computed using the *sandwich estimator*.

6.5 Experiment Instructions

Instructions

Welcome and thank you for participating in todays economic experiment. Please put away all your belongings and turn off your cell phones. You are not allowed to talk to any other participant during the experiment. If you have any questions, please raise your hand. We will come to you and answer your questions in private. The experiment will be run entirely on the computer and all interactions between yourself and others will take place via the computer terminal.

You have earned $5 just for showing up on time. This is yours to keep. In addition, depending upon the decisions you make, the decisions others make and random choice, you can earn more money. These instructions describe in detail the experiment and tasks you are asked to complete.

During the experiment, your earnings will be described in terms of tokens. At the end of the experiment, the total number of tokens you have earned will be converted to money at the following rate:

20 tokens = $1 (1 token = 5 cents)

The experiment consists of four tasks. You will receive instructions for each task prior to making decisions for that task. Your total earnings from the experiment will be the sum of your earnings in each task. At the end of the session, the total number of tokens earned across all four tasks will be converted to money and paid to you privately in cash, along with the $5 show-up fee.

Your Neighborhood

You will be placed in a network with 14 other participants as shown in the Figure below. The placement of participants in the network is random, and participants do not know who is connected to whom.

Each participant is placed at one position on the network, shown in the Figure below, and is connected to exactly two other participants. This placement and connection are fixed throughout each of the four tasks. You and the two other participants that are connected to you in the network define your neighborhood. In the Figure, for example, if you are placed in the position of the circle that is highlighted in pink then your neighbors are highlighted in yellow. In the network, there are three connected participants in each neighborhood and five neighborhoods.

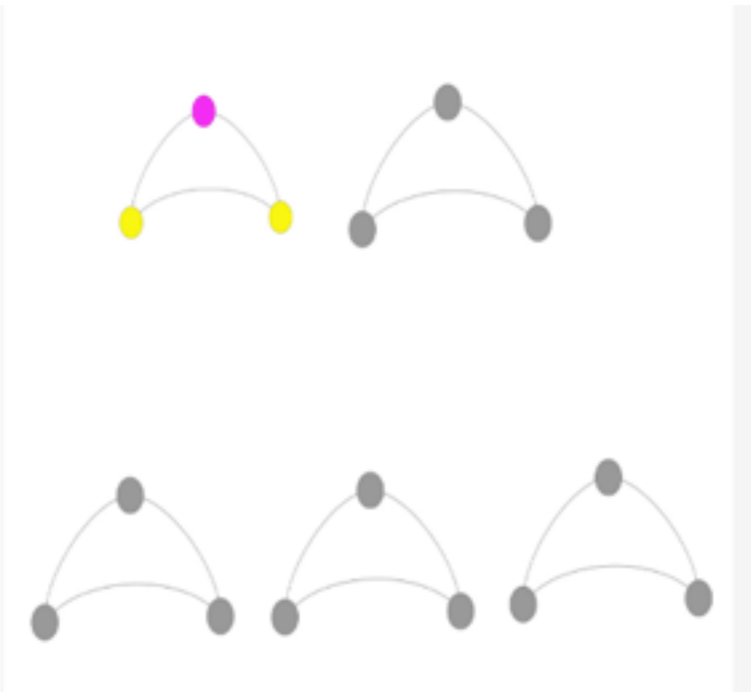

Figure 3: Groups in the experiment

The Decision Situation

Each participant is provided with 20 tokens and must decide how to allocate the tokens between a private account and a group project. You can choose to put none, all or some of your tokens into the group project. The tokens you choose not to contribute to the group project will remain in your private account. Everyone makes the same decision.

For each token put in the private account you earn exactly one token. You are the only one who earns tokens from your private account.

What you earn from the group project depends on the total number of tokens that you and your neighbors contribute to the group project. The more each member of the neighborhood contributes to the group project, the more each member earns. Remember that your neighborhood includes you and two other participants.

Your earnings from the group project are best explained by a number of examples.

Example 1: Suppose that you decided to contribute no tokens to the group project but the 2 other members of your neighborhood contribute a total of 36 tokens. Then your earnings from the group project would be 36 tokens $\times$ 0.4 = 14.4 tokens. Everyone else in your group would also earn 14.4 tokens.

Example 2: Suppose that you contribute 15 tokens to the group project and the 2 other members of your neighborhood invest a total of 36 tokens. This makes a group total of 51 tokens. Your earnings from the group project would be 51 tokens $\times$ 0.4 = 20.4 tokens. The other 2 members of the group would also earn 20.4 tokens.

Example 3: Suppose that you contribute 20 tokens in the group project but the other 2 members in your neighborhood invest nothing. Then you, and everyone else in the group, would earn from the group project 8 tokens (20 tokens $\times$ 0.4 = 8 tokens).

As you can see, every token contributed to the group project earns 0.4 tokens for every member of the neighborhood, not just the participant who puts it there. It does not matter who contributes tokens to the group project. Everyone will get a return from every token contributed therewhether they contributed tokens in the group project or not.

Your total earnings from the private account and group project will be:

Your total earnings = 20 - your tokens contributed to the group project + 0.4 $\times$ sum of tokens contributed to the group project by all members of your neighborhood

You will now complete some questions to make sure everyone understands how earnings are calculated.

Questions

Subject number: ______________

Please answer the following questions. These will help you understand how earnings are calculated. Your payoff is not affected by your answers to these questions.

Each token in the private account earns 1 token. Each token in the group project earns 0.4 tokens for each participant in the neighborhood.

1. Each participant has 20 tokens. Suppose you contribute 12 tokens to the group project and the other two participants contribute 18 tokens in total.
What are your earnings from your private account? ______________
What are your earnings from the group project? ______________
What are your neighbors earnings from the group project? ______________
What are your total earnings? ______________

2. Each participant has 20 tokens. Suppose you contribute 20 tokens to the group project and the other two participants contribute 38 tokens in total.
What are your earnings from your private account? ______________
What are your earnings from the group project? ______________
What are your neighbors earnings from the group project? ______________
What are your total earnings? ______________

3. Each participant has 20 tokens. Suppose you contribute 2 tokens to the group project and the other two participants contribute 38 tokens in total.
What are your earnings from your private account? ______________
What are your earnings from the group project? ______________
What are your neighbors earnings from the group project? ______________
What are your total earnings? ______________

Instructions for Task 1 – A-Task

In the A-task, you will make a choice for the decision situation described earlier. You will have 20 tokens and must decide how many to put into your private account and the group project. You will be randomly assigned to a 3-participant neighborhood.

Each participant has two decisions in this task: make an unconditional contribution and complete a contribution table. Details about these two decisions are as follows.

Unconditional Contribution: In this decision, you must decide how many of the 20 tokens you would like to put in the group project. You will make your decision on a screen such as the following.

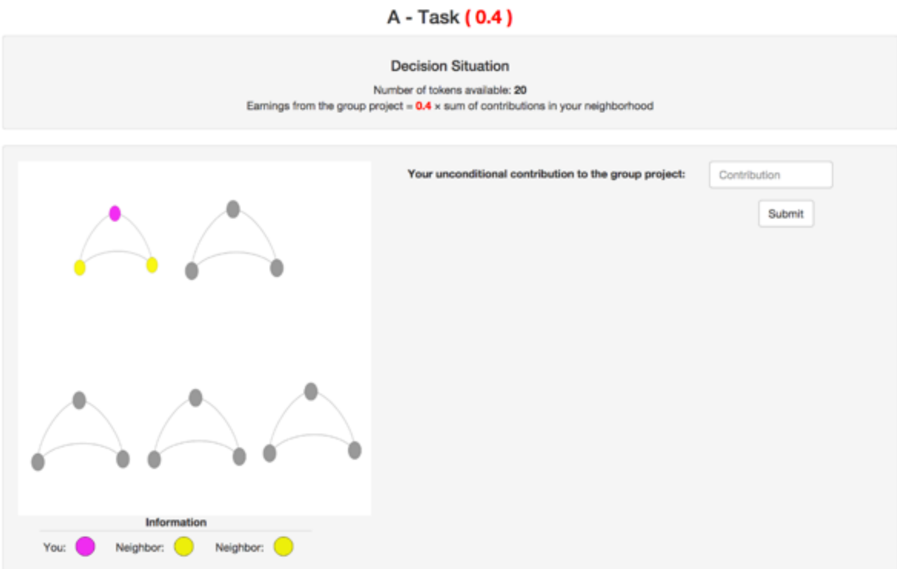

FIGURE 4: DECISION SCREEN FOR UNCONDITIONAL CONTRIBUTION CHOICE IN STRATEGY GAMES

Contribution Table: In this decision, you must decide how many tokens you would like contribute to the group project for each possible average contribution of your neighbors (e.g. 0, 1, 2,..., 20). That is, if your neighbors contributed 0 tokens on average, how much would you contribute? If they contributed 1 token on average, how much would you contribute? If they contributed 2 tokens on average, how much would you contribute? And so on, up to 20 tokens on average.

This means that in total you have to give 21 responses. You will make your decisions on a screen such as the following

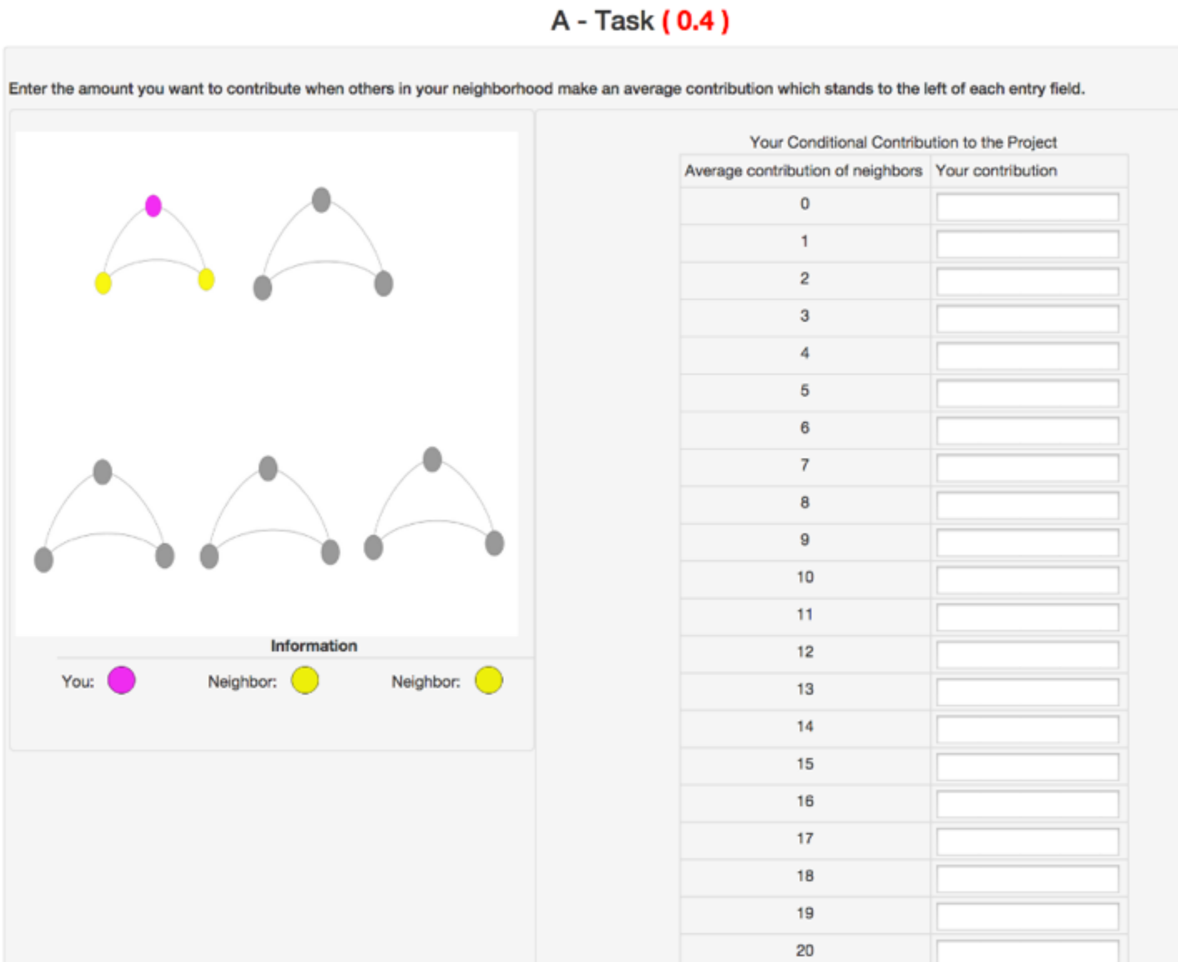

Figure 5: Decision screen for conditional contribution table in strategy games

Once each participant has made the unconditional decision and completed the contribution table, the computer will randomly determine if the unconditional contribution or the contribution table will be used to determine your earnings. In each neighborhood, one of the three participants is randomly chosen to have the contribution table count to calculate earnings. For the other two participants in the neighborhood, the unconditional contribution counts to calculate earnings. How is this done? If the participant is chosen to have his contribution table count for earnings, first the unconditional contributions to the group project of his neighbors are averaged and rounded to the nearest integer (e.g. 0,1,2,...,20). Then, the contribution table of the participant is used to determine how many tokens the participant contributes to the group project. The number of tokens contributed is the amount he specified for the average contribution of his neighbors.

So, if the average contribution of his neighbors is 16 tokens, and he specified 10 tokens if the average contribution of his neighbors is 16, the total contributed to the group project by everyone in the neighborhood would be 42 ($16\times2+10$) tokens.

You will not know if advance which decision, the unconditional contribution or the contribution table, will count to determine your earnings, so you should make each decision as though it will count for your earnings.

The following examples should help make this procedure clear.

Example 1: Suppose the contribution table was randomly chosen to count for you. This means that the decisions you made in the contribution table determine your earnings. For the other two neighbors their unconditional contributions determine their earnings. Suppose that the total contributions of the other two neighbors are 26 tokens, and the average contribution 13 tokens (26 tokens/2). In your contribution table, suppose you chose to contribute 4 tokens if the average contribution of your neighbors is 13, then your earnings for Task 1 would be: $20 - 4 + 0.4 \times (4 + 26) = 28$. If instead you chose to contribute 14 if the average contribution of neighbors is 13, your payoff would be: $20 - 14 + 0.4 \times (14 + 26) = 22$.

Example 2: Suppose the unconditional contribution was randomly chosen to count for your earnings. Also, suppose that the unconditional contribution of the neighbor who was not selected for the contribution table to count is 12. If your unconditional contribution is 20, then the average unconditional contribution is 16 tokens ($(20 + 12)/2$). If the neighbor selected to have his contribution table count chose 18 tokens if the average contribution of his neighbors is 16, then your earnings are: $20 - 20 + 0.4 \times (20 + 12 + 18) = 20$.

Are there any questions before we begin?

Instructions for Task 2 – B-Task

In the B-Task, you will be randomly assigned to a 3-participant neighborhood as described earlier. Your neighbors in Task 2 may be different from your neighbors in Task 1, however, you will remain with the same neighbors for all decisions you make in Task 2.

The B-Task lasts for several rounds. The number of rounds is randomly determined. In each round, you face the basic decision situation described at the beginning of the experiment. After each round, there is an 85% probability that there will be one more round. So, for instance, if you are in round 2, the probability there will be a third round is 85% and if you are in round 9, the probability there will be another round is also 85%. How this works is as follows. After each round, the computer will randomly draw a number between 1 and 100 (e.g. 1, 2, 3,..., 100), where each number is equally likely to be chosen. If the chosen number is 85 or lower, there will be another round. If the chosen number is 86 or above, there will be no additional rounds, and the task will end. You will know there is another round if you see the decision screen again and are asked to make a decision. If the task ends, you will get a message saying the task is done. You will not know ahead of time for how many rounds you will make decisions.

In each round, you will be given 20 tokens and must decide how many tokens you would like to contribute to the group project and how many you would like to put in your private account. You will receive earnings only from the group project that involves participants in your neighborhood. Your earnings from your contribution decision in a given round are determined as:

Your total earnings in a round = 20 - your tokens contributed to the group project + 0.4 × sum of tokens contributed to the group project by all members of your neighborhood

You will participate in the decision situation repeatedly with the same neighbors, until it is randomly determined that there are no more rounds.

You will make decisions on a screen such as the following:

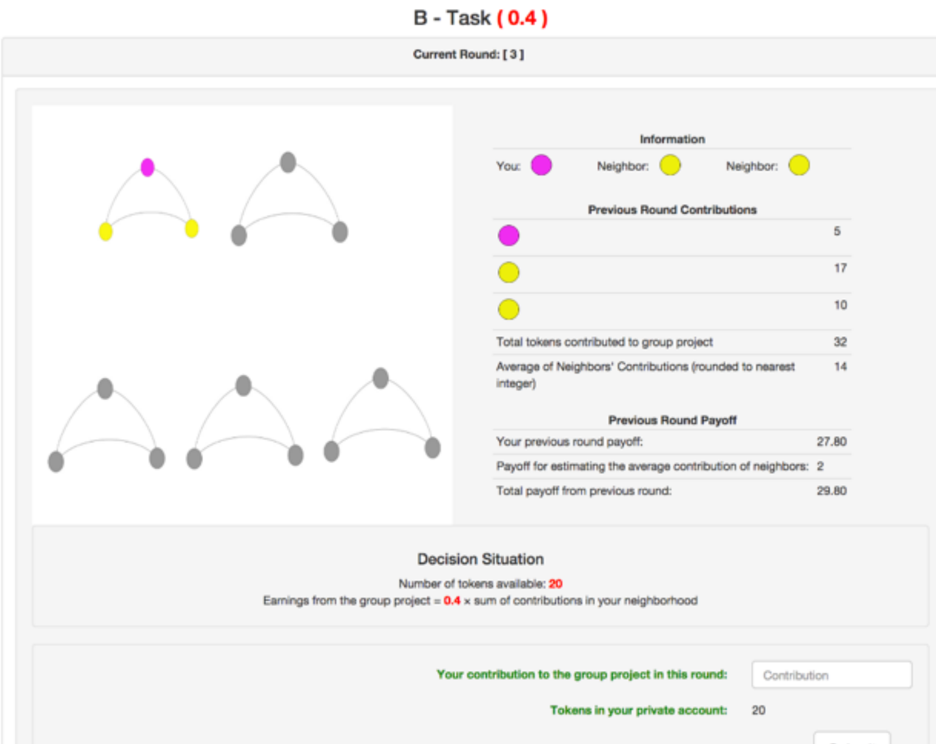

Figure 6: Decision Screen for contribution choice in repeated game

Here is an example to explain how earnings are calculated in each round:

Example 1: Suppose you chose to contribute 12 tokens and your neighbors chose to contribute 30 tokens in total. Your earnings in that round would be: 20 - 12 + 0.4 × (12 + 30) = 24.8 tokens.

In each round, after you decide how much to contribute to the group project, you will be asked to guess the average contribution to the project (rounded to the nearest integer) of your two neighbors. You will receive tokens for the accuracy of your estimate. If your guess is exactly equal to the average contribution of your neighbors you will receive 3 tokens in addition to your earnings for that round. If your guess was off by 1 token, you will get 2 additional tokens. If your guess was off by 2 tokens, you will get 1 additional token. And, if your guess was off by 3 or more tokens, you will get 0 additional tokens.

When everybody in your neighborhood has completed the two decisions, you will be shown each of their contributions, the total contributions to the group project, and the average contribution. You will only be informed of the contributions of those in your neighborhood. You will not be informed of contributions of participants in other neighborhoods. You will also be informed of your earnings for the current round.

Once all subjects in the experiment have completed the two decisions and are told their earnings and the contributions of their neighbors in the current round, the computer will randomly draw a number between 1 and 100 to see if everyone plays another round. If there is not another round, the task is done.

Are there any questions before we begin?

Instructions for Task 3 – A-Task

In Task 3, you will make two decisions again as you did in the Task 1 A-Task. The difference between this task and Task 1 is that for each token contributed to the group project you, and the other two neighbors, will get 0.8 tokens back.

Things to remember:

1. You will be randomly placed on the network at the beginning of the task and assigned to a 3-participant neighborhood. Your neighbors in this Task 3 may be different than in the previous two tasks.

2. You will make two decisions.

3. The first decision, the unconditional contribution, is how many of your 20 tokens you want to contribute to the group project.

4. The second decision, completing the contribution table, is how many of your 20 tokens you want to contribute to the group project for each possible average contribution of your neighbors (e.g. 0, 1, 2,..., 20).

5. Each token contributed to the group project will earn 0.8 tokens for each participant in the neighborhood.

6. One of the two decisions, the unconditional contribution or the contribution table, will be randomly chosen to determine earnings. You will not know ahead of time which decision will count.

Are there any questions before we begin?

Instructions for Task 4 − B-Task

In Task 4, you will make decisions again as you did in the Task 2 B-Task. The difference between this task and Task 2 is that for each token contributed to the group project you, and the other two neighbors, will get 0.8 tokens back.

Things to remember:

1. You will be randomly placed on the network at the beginning of the task and assigned to a 3-participant neighborhood. Your neighbors in Task 4 may be different than in the previous tasks, however, you will remain with the same neighbors for all rounds in this task.

2. You will face the same decision situation for several rounds. You must decide how many of your 20 tokens you want to contribute to the group project.

3. The number of rounds is randomly determined. After each round, there is an 85% probability that there will be one more round.

4. If there is another round, you will see a decision screen to make another choice. If there is not another round, you will be a message saying the task is over.

Are there any questions before we begin?

6.6 Validation Experiment Instructions

INSTRUCTIONS

Welcome and thank you for participating in todays economic experiment. Please put away all your belongings and turn off your cell phones. You are not allowed to talk to any other participant during the experiment. If you have any questions, please raise your hand. We will come to you and answer your questions in private. The experiment will be run entirely on the computer and all interactions between yourself and others will take place via the computer terminal.

You have earned \$5 just for showing up on time. This is yours to keep. In addition, depending upon the decisions you make, the decisions others make and random choice, you can earn more money. These instructions describe in detail the experiment and the tasks you are asked to complete.

During the experiment, your earnings will be described in terms of tokens. At the end of the experiment, the total number of tokens you have earned will be converted to money at the following rate:

20 tokens = \$1 (1 token = 5 cents)

The experiment consists of two tasks. You will receive instructions for each task prior to making decisions for that task. Your total earnings from the experiment will be the sum of your earnings in each task. At the end of the session, the total number of tokens earned across the two tasks will be converted to money and paid to you privately in cash, along with the \$5 show-up fee.

In each task, all participants will be randomly divided in groups of four members. Participants do not know who is in which group.

The Decision Situation

You will be a member of a group consisting of 4 people. Each group member is provided with 20 tokens and must decide how to allocate the tokens between a private account and a group project. You can choose to contribute none, all, or some of your tokens to the group project. The tokens you choose not to contribute to the group project will remain in your private account. Everyone makes the same decision.

For each token put in the private account you earn exactly one token. For example, if you put 20 tokens into your private account (and therefore do not contribute to the group project) your earnings from this account will be 20 tokens. If you put 6 tokens into your private account, your earnings from this account will be 6 tokens. You are the only one who earns tokens from your private account.

What you earn from the group project depends on the total number of tokens that you and the other members in your group contribute to the group project and a return rate. The return rate, denoted by M, specifies how much a token contributed to the group project returns to every member of the group. The return rate M will always be strictly between 0.25 and 1 in the experiment. For example, if the return rate is 0.5, then each token contributed to the group project returns 0.5 tokens to every member of the group. You will always be told the return rate before you make your contribution decision. The more each member of the group contributes to the group project, the more each member earns from the group project.

Your earnings from the group project are best explained by a number of examples.

Example 1: Suppose that you decided to contribute no tokens to the group project but the three other members in your group contribute a total of 40 tokens. This makes a total contribution of 40 tokens to the group project.

Suppose each token contributed to the group projects returns M = 0.3 tokens to each group member. Then your earnings from the group project would be 40 tokens × 0.3 = 12 tokens. Everyone else in your group would also earn 12 tokens from the group project.

Suppose instead that each token contributed to the group projects returns M = 0.5 tokens to each group member. Then your earnings from the group project would be 40 tokens × 0.5 = 20 tokens. Everyone else in your group would also earn 20 tokens from the group project.

Suppose instead that each token contributed to the group projects returns M = 0.8 tokens to each group member. Then your earnings from the group project would be 40 tokens × 0.8 = 32 tokens. Everyone else in your group would also earn 32 tokens from the group project.

Example 2: Suppose that you decided to contribute 10 tokens to the group project and the three other members in your group contribute a total of 40 tokens. This makes a total contribution of 50 tokens to the group project.

Suppose each token contributed to the group projects returns M = 0.3 tokens to each group member. Then your earnings from the group project would be 50 tokens × 0.3 = 15 tokens. Everyone else in your group would also earn 15 tokens from the group project.

Suppose instead that each token contributed to the group projects returns M = 0.5 tokens to each group member. Then your earnings from the group project would be 50 tokens $\times$ 0.5 = 25 tokens. Everyone else in your group would also earn 25 tokens from the group project.

Suppose instead that each token contributed to the group projects returns M = 0.8 tokens to each group member. Then your earnings from the group project would be 50 tokens $\times$ 0.8 = 40 tokens. Everyone else in your group would also earn 40 tokens from the group project.

As you can see, it does not matter who contributes tokens to the group project. Everyone will get a return from every token contributed-whether they contributed tokens in the group project or not.

Your total earnings will be the sum of the tokens you earn from your private account and the tokens you earn from the group project. Therefore, your total earnings will be:

Your total earnings = 20 - your tokens contributed to the group project + M $\times$ sum of tokens contributed to the group project by everybody in your group

You will now complete some questions to make sure everyone understands how earnings are calculated.

Questions

Subject number: _______________

Please answer the following questions. These will help you understand how earnings are calculated. Your payoff is not affected by your answers to these questions.

Each token in the private account earns 1 token.

1. Each participant has 20 tokens. Suppose you contribute 10 tokens to the group project and the other three other members in your group contribute 40 tokens in total. Each token in the group project earns M = 0.5 tokens for each participant in the group.

What are your earnings from your private account? _______________

What are your earnings from the group project? _______________

What are the earnings of your group members from the group project? _______________

What are your total earnings? _______________

2. Each participant has 20 tokens. Suppose you contribute 20 tokens to the group project and the other three members in your group contribute 30 tokens in total. Each token in the group project earns M = 0.8 tokens for each participant in the group.

What are your earnings from your private account? _______________

What are your earnings from the group project? _______________

What are the earnings of your group members from the group project? _______________

What are your total earnings? _______________

3. Each participant has 20 tokens. Suppose you contribute 12 tokens to the group project and the other three members in your group contribute 38 tokens in total. Each token in the group project earns M = 0.3 tokens for each participant in the group.

What are your earnings from your private account? _______________

What are your earnings from the group project? _______________

What are the earnings of your group members from the group project? _______________

What are your total earnings? _______________

Instructions for Task 1 - A-Task

In the A-task, you will make a choice for the decision situation described earlier. You will have 20 tokens and must decide how many to put into your private account and the group project. You will be randomly assigned to a 4-member group and will not know who is in your group.

Each token contributed to group project in this task returns M = 0.5 tokens to each group member.

Each participant has two decisions to make in this task: make an unconditional contribution and complete a contribution table.

Unconditional Contribution: In this decision, you must decide how many of the 20 tokens you would like to contribute to the group project. You will make your decision on a screen such as the following.

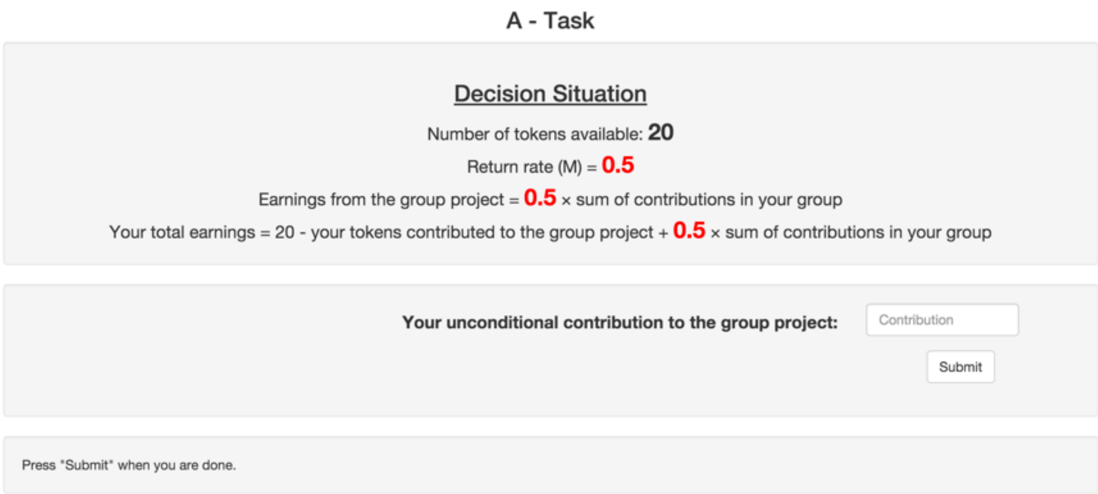

Figure 7: Decision screen for unconditional contribution choice in strategy games

Contribution Table: In this decision, you must decide how many tokens you would like contribute to the group project for each possible average contribution of the other group members (e.g. 0, 1, 2,..., 20). That is, if the other group members contributed 0 tokens on average, how much would you contribute? If they contributed 1 token on average, how much would you contribute? If they contributed 2 tokens on average, how much would you contribute? And so on, up to 20 tokens on average.

This means that in total you have to give 21 responses. You will make your decisions on a screen such as the following:

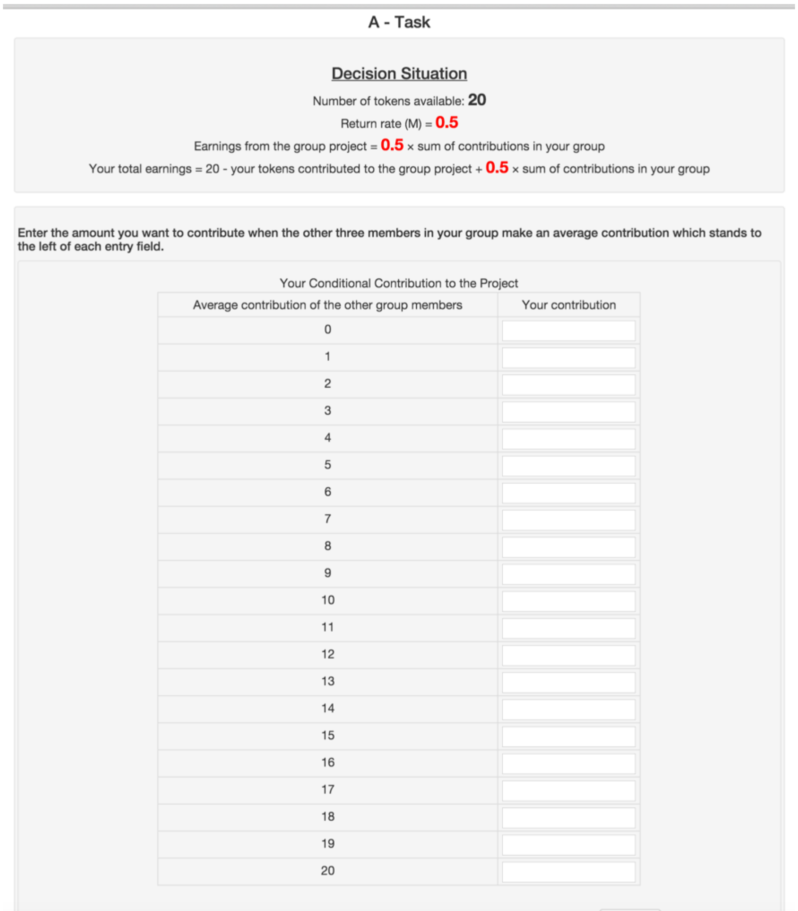

Figure 8: Decision screen for conditional contribution table in strategy games

Once each participant has made the unconditional decision and completed the contribution table, the computer will randomly determine if the unconditional contribution or the contribution table will be used to determine your earnings. In each group, one of the four participants is randomly chosen to have the contribution table count to calculate earnings. For the other three participants in the group, the unconditional contribution counts to calculate earnings. How is this done? If the participant is chosen to have his contribution table count for earnings, first the unconditional contributions to the group project of the three other members in his group are averaged and rounded to the nearest integer (e.g. 0,1,2, ,20). Then, the contribution table of the participant is used to determine how many tokens the participant contributes to the group project. The number of tokens contributed is the amount he specified for the average contribution of the other members in his group.

For example, if the average contribution of the other group members in his group is 16 tokens (48 tokens/3), and he specified 10 tokens if the average contribution of the other group members is 16, the total contributed to the group project by everyone in the group would be 58 (10 + 48) tokens.

You will not know in advance which decision, the unconditional contribution or the contribution table, will count to determine your earnings, so you should make each decision carefully as though it will count.

The following examples should help make this procedure clear.

Example 1: Suppose the contribution table was randomly chosen to count for you. This means that the decisions you made in the contribution table determine your earnings. For the other three members of your group their unconditional contributions determine their earnings. Suppose that the total contribution of the other three group members is 39 tokens, and the average contribution 13 tokens (39 tokens/3). In your contribution table, suppose you chose to contribute 11 tokens if the average contribution of the other group members is 13, then your earnings for Task 1 would be: $20 - 11 + 0.5 \times (11 + 39) = 34$. If instead you chose to contribute 1 if the average contribution of the other group members is 13, your payoff would be: $20 - 1 + 0.5 \times (1 + 39) = 39$.

Example 2: Suppose the unconditional contribution was randomly chosen to count for your earnings. Also, suppose that the unconditional contributions of the two group members who were not selected for the contribution table to count are 12, 18. If your unconditional contribution is 20, then the average unconditional contribution is 17 tokens ($(20 + 12 + 18)/3$). If the group member selected to have his contribution table count chose 10 tokens if the average contribution of his other group members is 17, then your earnings are: $20 - 20 + 0.5 \times (20 + 12 + 18 + 10) = 30$.

Are there any questions before we begin?

Instructions for Task 2 - B-Task

In the B-Task, you will be randomly assigned to a group of four as described earlier. Your group members in Task 2 may be different from your group members in Task 1, however, you will remain with the same group members for all decisions you make in Task 2.

The B-Task lasts for 10 rounds. In each round, you will be given 20 tokens and must decide how many tokens you would like to contribute to the group project and how many you would like to put in your private account.

Prior to making your contribution decision in each round, you will be informed of the return amount M ($0.25 < M < 1$) from each token contributed to the group project by your group. The return amount M could change across rounds. If each token contributed to the group project returns M tokens to each group member in a given round, your earnings from your contribution decision in that round are:

Your total earnings = 20 - your tokens contributed to the group project + M × sum of tokens contributed to the group project by everybody in your group

You will participate in the decision situation repeatedly with the same group members for all 10 rounds.

You will make decisions on a screen such as the following:

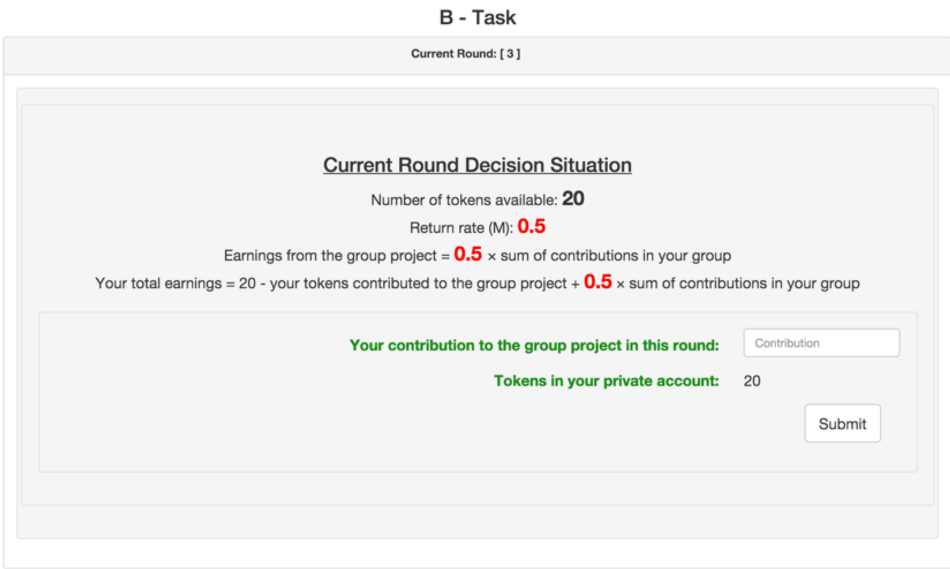

Figure 9: Decision Screen for contribution choice in repeated game

Here is an example to explain how earnings are calculated in each round:

Example 1: Suppose the return rate M = 0.5 tokens for a particular round. Suppose you chose to contribute 10 tokens and the other three group members chose to contribute 40 tokens in total in that round. Your earnings in that round would be: $20 - 10 + 0.5 \times (10 + 40) = 35$ tokens.

In each round, after you decide how much to contribute to the group project, you will be asked to guess the average contribution to the project (rounded to the nearest integer) of the three other members of your group. You will receive tokens for the accuracy of your estimate. If your guess is exactly equal to the average contribution of the other group members, you will receive 3 tokens in addition to your earnings for that round. If your guess was off by 1 token, you will get 2 additional tokens. If your guess was off by 2 tokens, you will get 1 additional token. And, if your guess was off by 3 or more tokens, you will get 0 additional tokens.

When everybody in your group has completed the two decisions in a given round, you will be shown the total contributions to the group project, the average contribution of the other three group members, and your total earnings in that round.

Once all subjects in the experiment have completed the two decisions and are told their earnings and the average contribution of their group members in the current round, the task proceeds to the next round. After the 10th round, the task is done and there will be no more rounds.

Are there any questions before we begin?