Favor Trading in Public Good Provision

Sarah Jacobson[1]
Department of Economics
Williams College
24 Hopkins Hall Dr.
Williamstown, MA 01267
Sarah.A.Jacobson@williams.edu
Phone: 413-597-2476
Fax: 413-597-4045

Ragan Petrie
Interdisciplinary Center for Economic Science (ICES)
George Mason University
4400 University Drive, MSN 1B2
Fairfax, VA 22030
rpetrie1@gmu.edu

September 2013

Abstract

Favor trading is common. We do something nice for someone and they do something nice in return. Several motives might underlie such behavior, including altruism, strategic motives, and direct or indirect positive reciprocity. It is not yet well-understood how these fit together to affect behavior, how they interact in various institutional structures, and how they play out over time. We use a laboratory experiment to study the elements and dynamics of favor trading in a particular setting: the private provision of a public good. In our experiment, giving subjects the ability to practice targeted reciprocity by making a simple, low-cost change in information provision increases contributions to the public good by 14%. Subjects reward group members who have previously been generous to them and withhold rewards from ungenerous group members. Strategic concerns cannot explain all of this behavior, and it must be at least partly due to direct reciprocity. When someone cannot directly benefit from favor trading, he gives much less to the public good. People thus excluded from the "circle of reciprocity" provide a clean and strict test of indirect reciprocity. Contrary to previous studies in the literature, we do not observe indirect reciprocity.

Keywords: public goods, direct and indirect reciprocity, experiment, peer-to-peer fundraising

JEL Classifications: C92, H41, D01

[1] Corresponding author. James Cox, Vjollca Sadiraj, and participants in various workshops, conferences, and seminars (notably at Georgia State University) provided helpful comments. We thank the National Science Foundation (Award SES-0752754) and Georgia State University for funding this research.

"Do: Get excited. Engage your network, rally the masses and have fun doing it." – *charity:water advice on peer fundraising*

"If you don't give back no one will like you." – *Crowdrise tagline*

1. Introduction

Favor trading is ubiquitous. We lend tools to a neighbor and later ask him to watch our house while we are away. We recommend a friend for a job and then entreat him to give to our favorite charity or cause. Sometimes the exchange is less direct. We may support the school band of a colleague's child with no obvious opportunity for personal benefit in the near future. Such exchanges between non-kin people are common and may spring from motives such as strategic self-interest, reciprocal altruism, and indirect reciprocity. We know little about which of these factors drive behavior and how they interact to increase or reduce cooperative behavior. This is important to understand because, in institutions in which cooperation is pro-social, the ability to return favors may be an important determinant of efficiency.[1] In this paper, we study favor trading in the private provision of a public good. We manipulate the characteristics of the institution to study the importance of strategy, direct reciprocity, and indirect reciprocity.

We add to a substantial literature showing that people do contribute to public goods, but that cooperation is difficult to sustain (Ledyard, 1995; Chaudhuri, 2011). Reciprocal-type forces based on subjects' reputations may sustain giving and other pro-social cooperative behavior (e.g. Trivers, 1971; Fischbacher, Gächter, and Fehr 2001; Milinski et al. 2002; Gächter 2007), including through punishment and reward (e.g. Fehr and Gächter 2000; Andreoni et al. 2003; Masclet et al. 2003; Bochet et al. 2006; Houser et al. 2008; Almenberg et al. 2011). Social

[1] Social favor trading may drive behavior in many cooperative settings. The same dynamic may exist when time and effort are solicited, as in a volunteer advisory board. Klein (1990) argues that peer pressure played a role in the funding of turnpikes in early America. Jackson et al. (2012) show that favor trading can informally enforce contracts (including credit exchanges) within a social network using data from Indian villages, and La Ferrara (2003) discusses the role of intergenerational favor-trading in a similar context of informal credit in Ghana.

forces also may drive giving outside the lab; for example, DellaVigna et al. (2012) find evidence of "social pressure" in door-to-door fundraising.

The importance of a social element in giving can also be inferred from the fact that charities frequently encourage supporters to tap into their social networks. A supporter can trade favors with friends by asking them to contribute to his favorite cause with the expectation of being asked to support their causes in the future. Brick-and-mortar institutions like the Girl Scouts of the USA and the Susan G. Komen Race for the Cure have long used personal contact to fundraise through grassroots networks. Social media have made this process easier than ever with online tools like Facebook Causes and DonorPages. Favor trading in this environment is possible because different people favor different public goods and because institutions often make contribution information public.

When reciprocal giving supports public good provision, it may work in many ways.[2] It may be strategic, in which case a person gives only to garner a reward later and giving declines when future rewards are no longer possible. This motive was studied and compared to other-regarding preferences in Cabral et al. (2012), who find evidence of cooperation driven by strategic concerns but also find evidence of other-regarding preferences. Dreber et al. (2011) also study strategic cooperation as compared to cooperation based on social preferences, concluding that both are important, and Reuben and Suetens (2012a) find that much of the cooperation in their repeated Prisoner's Dilemma game is strategic. Reciprocal giving may be driven by direct reciprocity (i.e. reciprocal altruism), models of which include Rabin (1993) and Cox et al.

[2] As Sobel (2005) discusses, there are many terms to describe reciprocity. We will use phrases like "reciprocal act" to describe all kind acts that seem to respond to a kind act received in the past. We will describe these acts as "strategic" ("instrumental," per Sobel) if they are self-serving. We will use the phrases "direct reciprocity" and "indirect reciprocity" to describe *only* acts that are rooted in other-regarding preferences ("intrinsic," per Sobel).

(2008), in which case the giver's preferences are increasing in the payoff of a past benefactor.[3] Indirect reciprocity may play a role if an actor rewards a good deed done to a third party (e.g. Nowak and Sigmund 2005; Seinen and Schram 2006, Engelmann and Fischbacher 2009).[4]

Our experimental design identifies which institutional elements allow reciprocity to boost giving to a public good.[5] It also helps us determine the relative importance of direct reciprocity rooted in other regarding preferences, indirect reciprocity, and strategic self-interest.

We induce heterogeneous interests in "causes" by assigning asymmetric returns to the public good. In every round, one person (the "Stakeholder") gains more from the public good than do the others in the group (it is his "pet cause"). A group with a Stakeholder is similar to Olson's (1965) "privileged group," as studied in Reuben and Riedl (2009).[6] Our privileged group is even more effective because the Stakeholder position rotates from round to round through group members, creating opportunities for targeted reciprocal acts. To examine the importance of reciprocity, we compare behavior when subjects have full information regarding each other's actions (enabling conditional action) and when they do not have that information.

We also examine whether a group member who never has a "cause" will indirectly reciprocate by using his contribution to reward people who have been kind to others. This differs from many studies (e.g. Engelmann and Fischbacher, 2009; Seinen and Schram, 2006) that have

[3] Direct reciprocity or gift exchange has been studied extensively in a variety of settings (e.g. Fehr et al. 1993; Berg et al. 1995; Charness and Rabin 2002; Cox 2004; List 2006; Gneezy and List 2006).
[4] The types of indirect reciprocity most relevant to our setting include the "downstream reciprocity" of Nowak and Sigmund (2005), the "third party reward" of Almenberg et al. (2011), and the "social indirect reciprocity" of Stanca (2009).
[5] Non-experimental studies of peer solicitation could suffer from endogeneity. For example, Long (1976) found that the more "personal" a donor solicitation, the more contributions were solicited, but this analysis will likely overestimate this relationship because charities may focus personal solicitations on generous donors . However, Meer and Rosen (2011) find a similar result identified based on a solicitee's place in an alphabetized list.
[6] Isaac and Norton (2013) study a similar role, which they call an "agent of grace," and find that the presence of such a person may slightly boost group members' contributions.

found support for indirect reciprocity in that the people whose actions we study gain very little from a group norm of cooperation. In previous studies, subjects who indirectly reciprocate gain a large benefit if the group establishes a general norm of cooperation. We provide a stricter test.

We find that letting people engage in targeted reciprocity increases contributions by a modest but statistically significant 14.4%, although this effect may decline as rounds progress. More interestingly, we find strong evidence of reciprocal acts. Some may be strategic, but some are more consistent with direct reciprocity. We do not find evidence of indirect reciprocity. When a person's benefit from a norm of cooperation is indirect, he does not reward kindness with kindness.

The paper proceeds as follows. The next section explains the general experimental design. Section 3 describes three experimental treatments. In Section 4, we present results. We conclude in Section 5.

2. Experimental Design

The experiment is a linear public goods game with publicly revealed contributions and asymmetric payoffs. The design mimics a peer-to-peer fundraising network that allows favor trading. In each treatment, subjects are assigned into five-person groups. Each person i has an endowment of z tokens in each round t. He must choose how many tokens g_{it} to allocate to a public investment with some return to all group members and a private investment with return a to himself. In each round, one group member (with index $Stakeholder_t$) is the Stakeholder: he has a bigger stake in the public good because he gets a higher return c from tokens invested there as compared to the non-Stakeholder return b. The parameters are such that $b < a < c$: non-

Stakeholders maximize profit by not contributing and Stakeholders maximize profit by contributing. The Stakeholder position rotates through group members. Payoffs are:

$$\pi_{it} = \begin{cases} c\sum_j g_{jt} + a\left(z - g_{it}\right) & \text{if } Stakeholder_t = i \\ b\sum_j g_{jt} + a\left(z - g_{it}\right) & \text{if } Stakeholder_t \neq i \end{cases}$$

Parameters also ensure that $a < 4b + c$, so that the social optimum is achieved if everyone contributes fully. Since $c > a$, even a selfish Stakeholder always contributes to the public good. Since $b < a$, non-Stakeholders face a dilemma: they maximize profit by keeping all of their tokens, but this free riding is anti-social. The Nash equilibrium if all actors are selfish and have common knowledge of others' selfish preferences is that Stakeholders contribute fully and non-Stakeholders contribute nothing. Non-Stakeholder contributions are costly cooperation: they will be the focus of our analysis.

A selfish person is predicted to free ride when he is non-Stakeholder. An altruist or a person with warm glow preferences would contribute a positive amount if his marginal gain in utility from increasing others' payoffs is larger than his marginal utility loss from the reduction in his own payoff. Against this backdrop, the strategic and reciprocal forces described above will be at play, given sufficient information.

If a subject can only see a list or summary of his group members' contributions, he may exhibit conditional cooperation. However, the asymmetry provided by the Stakeholder position and a more complete information set (information on when each group member will be Stakeholder and a contribution history for each member) allow targeted reciprocal actions to occur. Subject *i* may interpret the contribution of another subject *j* when *i* is Stakeholder as a signal of *j*'s kindness since *j*'s contribution increased *i*'s earnings. Subject *i* may reciprocate by

contributing when *j* is Stakeholder. If another subject *k* is less generous during *i*'s Stakeholder stint, *i* may withhold contributions when *k* is Stakeholder.

Some reciprocal acts are motivated by other-regarding preferences—they reflect direct reciprocity. In this case, a person's willingness to pay to increase another person's payoff depends on the history he has observed. Other reciprocal acts occur when a strategic subject observes a history that changes his beliefs about the best strategic action. If subject *j* seems potentially generous and *i* is strategic, subject *i* may think that if he contributes when *j* is Stakeholder he will gain favor and earn future benefits when the contribution is reciprocated. The benefits of this strategic behavior are greatly reduced after a subject passes his last stint as Stakeholder, so even in the full-information Stakeholder setting, strategic motivations drop discontinuously in final periods. Thus, reciprocal giving in late rounds distinguishes other-regarding preferences from strategic giving.[7] A third reciprocal motive is indirect reciprocity: the drive, based in other-regarding preferences, to reward another subject who has previously been kind to other subjects rather than to himself.

Through these motives, information about contributions made by each group member and the timing of when each group member will be Stakeholder allows subjects to give reciprocally.[8] If information on Stakeholder timing and individual contribution histories is suppressed, targeted reciprocal acts (strategic or rooted in other-regarding preferences) are not possible. Thus, we can compare non-Stakeholder contributions across treatments to quantify the effects of allowing targeted reciprocity. Targeted reciprocity may increase total public good provision as subjects

[7] A misplaced application of a strategic reciprocal rule of thumb, i.e. cooperation caused by the mistaken belief that future rewards can be garnered, could also cause late-round reciprocation. Our study is not designed to examine this possibility. This idea is explored in Reuben and Suetens (2012b), who look at conditional cooperation that happens before end-period defection. This can be contrasted with the end-period cooperation we study.

[8] History could also affect current behavior through contagion: a person treated well (badly) in the past could react by behaving well (badly) solely because they have "caught" a good (bad) mood from their experience.

use contributions to reward each other. However, even if targeted reciprocal acts occur, it is not a foregone conclusion that provision will increase. For example, if initial contributions are low, reciprocators may respond by reducing contributions so that provision declines.

3. Experimental Treatments

We use three treatments: Private, Public, and Ineligible, described in detail below. All use endowment z = 20 tokens, private good return a = \$0.02, non-Stakeholder public good return b = \$0.01, and Stakeholder public good return c = \$0.03. For non-Stakeholders, the personal return from public good contribution is half the private good return; for Stakeholders it is 1.5 times that return. We use a within subject design. Subjects made decisions in all three treatments, in a different group for each treatment, with treatment order varied across sessions.[9]

The experiment is computerized and proceeds as follows. Subjects enter the lab and are given general instructions.[10] They are told that they will make decisions in three sets of multiple rounds, each with a different group, but they do not know the exact nature of the decisions they will make in each set (treatment) until the treatment-specific instructions are read. The first treatment begins with instructions that explain the information condition and the number of rounds for that treatment. The subjects play through all of the rounds for the treatment. When the first treatment is over, the groups are randomly re-matched. The second and third treatments proceed similarly, each with treatment-specific instructions read first. After all three treatments, subjects complete a questionnaire and are paid anonymously. Subjects' total earnings are the sum of their earnings in each treatment, which in turn are the sum of earnings in each round.

[9] We chose a within-subject design to investigate how an individual's behavior changes across treatments and to organize the data into subject "types." Such analysis did not yield interesting results and we do not present it.
[10] All instructions are available on the corresponding author's website.

In the software for each treatment, subjects see a decision screen and then, after making a decision, a review screen for each round. In both the decision and review screens, the central feature is the contribution table. This table contains a row for each round in the treatment. Columns contain information on the subject's contribution and the contributions of others in his group, the group's total contributions, and the subject's own earnings. Information is filled into this table after the decision stage of each round and remains visible for the rest of the treatment.

The Public treatment, which lasts ten rounds, follows the basic favor-trading public goods design outlined in the previous section. The Stakeholder position rotates through all five group members so everyone is Stakeholder twice. Contributions are publicly revealed and tracked individually, and Stakeholder assignments are common knowledge. Figure 1 shows the Public treatment decision screen (with simulated data). Each group member is randomly assigned a letter code and keeps the same letter code for all ten rounds. The contribution table shows in which rounds each subject will be the Stakeholder. The information contained in the contribution table allows subjects to reward each other for past generosity. For example, if subject *i* is A and *j* is B, *i* can see how much *j* contributed in Round 1 when *i* was the Stakeholder. Subject *i* can reward *j* with a large contribution when *j* is Stakeholder in Round 2, or *i* may withhold that reward if he deems *j*'s contribution stingy.

The Private treatment also lasts for ten rounds. As in the Public treatment, the Stakeholder position rotates through all group members so everyone is Stakeholder twice. However, the information environment differs from the Public treatment. Each subject still sees the disaggregated, individual contributions of his group members, but subjects are not assigned letter codes. It is no longer possible to track reputations. Figure 2 shows the review screen for the Private treatment. In each round's row, the contribution table reports the contributions of all

group members in a randomly-ordered list, re-shuffled for each round. Further, even if a subject thinks he can identify a group member as being worthy (or unworthy) of reward, he still does not know when that person will be Stakeholder. He only knows when he himself will be the Stakeholder, so he cannot target reciprocal acts toward any other subject.

Finally, the Ineligible treatment is similar to the Public treatment, with public reputations and public Stakeholder timing. However, one member ("the Bachelor") of each five-member group is ineligible to be the Stakeholder. The Bachelor is like the person in a social network who has no pet cause for which to fundraise. Because only four subjects are eligible to be Stakeholder, the Ineligible treatment lasts eight rounds so that each eligible subject is Stakeholder twice. The Bachelor is randomly chosen and remains the Bachelor for the entire treatment. In the screens for the Ineligible treatment (see Figure 3 for the review screen), the Bachelor is identified as the "Ineligible" person. The Stakeholder rotation skips the Bachelor: if person D is the Bachelor, the Stakeholder is A, then B, then C, then E, etc.

The main difference between the Public and Private treatments is that direct reciprocity, indirect reciprocity, and strategic reciprocity cannot motivate giving in the Private treatment. Subjects do not have the information they would need to respond to each other's actions. Unconditional altruism and general conditional cooperation, however, can affect giving in both treatments. Any difference between the Public and Private treatments must be due to some combination of targeted direct reciprocity, indirect reciprocity, and strategic giving. (Indirect reciprocity can only be cleanly isolated in the behavior of Bachelors in the Ineligible treatment, but it can still drive behavior of potential Stakeholders in the Public and Ineligible treatments.)

Within the Public treatment, we can see whether subjects give larger contributions when the current Stakeholder is someone who was previously generous. This is the favor trading we would expect to see in peer-to-peer fundraising or other non-lab settings. We can identify direct reciprocity as responsiveness of this type that does not disappear after the subject's last Stakeholder stint, since at that point strategic considerations have been eliminated.

The Ineligible treatment allows us to investigate indirect reciprocity. The Bachelor herself is not subject to direct reciprocity or strategic self-interest. If the Bachelor gives in a way that responds to the Stakeholder's past generosity, this is evidence of indirect reciprocity. Also, the presence of a Bachelor shrinks the "circle of reciprocity" from five people to four people. We study whether this change in the group dynamic affects non-Bachelor contributions. For example, others may be discouraged if Bachelor contributions are low.

The experiments were run in the Experimental Economics Center (E*x*CEN) at Georgia State University in six separate 20-subject sessions, for a total of 120 subjects. The software was written in z-Tree (Fischbacher, 2007), and the protocol was double anonymous.[11] Of the 120 subjects, 75 (62.5%) were female, and the average age was 21.8. The experiment lasted about 90 minutes, and average earnings were $24.33 (standard deviation $2.67).

4. Results

[11] Subjects could not identify which subjects they were interacting with, and the experimenters could not identify which subject made any set of decisions.

In our within subject design, each subject participated in all three treatments. The three treatments were run in all six possible orders, with each order run once. We do not observe effects of treatment order on variables of interest so we pool the data across sessions.[12,13]

Figure 4 shows contribution data across rounds. Stakeholder contributions in all treatments (the dashed lines) are close to the endowment. This is expected because the Stakeholder's return from the public good is greater than his return from the private good.[14] Non-Stakeholder contributions for each treatment are lower but positive in all rounds. These contributions show the downward trend usually seen in public goods games. Bachelor contributions in the Ineligible treatment are well below the other contribution trends and do not decline across the rounds.

Contributions by non-Stakeholders in the Private treatment compare well to previous research using linear public goods games with similar symmetric "prices of giving." Non-Stakeholder contributions start at 41% in round 1 and end at 21% in round 10, averaging 33%

[12] There is an order effect in that non-Stakeholder contributions are higher in the Private treatment if the Private treatment is before the Public treatment than if it is not first (43.44 as compared to 27.81 percent of endowment, Mann-Whitney test p=0.002). Because of this, if Private is first, there is no statistical difference between Private and Public non-Stakeholder contributions (paired Wilcoxon signed rank test p=0.368). The Private treatment may be more difficult to understand when presented first. This is supported by the fact that Stakeholder contributions are lower in the Private treatment if the Private treatment is first (93.19 as compared to 98.88 percent of endowment, Mann-Whitney test p=0.000). Our main results are unaffected by this order effect: we exclude those 40 subjects who were in the Private treatment first and all results still hold. Further evidence that our results are robust to these order effects come from parametric tests that control for order. To address any concerns that cross-treatment contamination affects our results, we also find the following if we use only the first treatment a subject experiences (thus losing a great deal of power): Public treatment non-Stakeholder contributions are no longer greater than those in the Private treatment; in Table 2 reciprocal acts are still statistically significant (Wilcoxon signed-rank p<0.01) when all rounds are considered for both the Public and Ineligible treatments but statistical significance is lost when only post-last Stakeholder stint rounds are considered; in Table 3 coefficients are similar but reciprocation loses statistical significance for post-last Stakeholder stint rounds in both Public and Ineligible treatments.

[13] In some treatment-order combinations, we see a small correlation between group members' non-Stakeholder contributions in an earlier treatment and a subject's non-Stakeholder contributions in later treatments. This could be caused by contagion, as discussed in footnote 8. However, this could merely reflect the fact that both of these variables should be correlated with a subject's inherent tendency to contribute. If contagion exists, the effect does not invalidate our results. Treatment effects should be either unaffected or diminished, and our reciprocity result is within-subject-within-treatment so should be unaffected.

[14] Contributions are not strictly 100%. This could be caused by subject error or myopic inequity aversion. Some (e.g., Saijo and Nakamura, 1995) have attributed such patterns to spite.

across all rounds. In the final round, 47% of non-Stakeholders make positive contributions. In Ledyard's (1995) survey of public goods game results, first period contributions range from 31%-68% of endowment, and final period contributions range from 9%-19% of endowment.

Results from asymmetric-return public goods games are difficult to compare because of differences in payoff structure, but the contribution levels in the Private treatment compare reasonably well to the literature. In a one-shot game with asymmetric returns, subjects with lower marginal per capita returns in somewhat similar treatments gave 20% on average in Goeree et al. (2002), 22.5% in Reuben and Riedl (2009), and 18% in Glöckner et al. (2011).

4.1 Treatment Effects

Recall that, in contrast to the Private treatments, the Public treatment opens the door to targeted direct reciprocity, indirect reciprocity, and strategic self-interest. We can test whether these forces increase provision of a public good simply by determining whether non-Stakeholder contributions are higher in the Public treatment than they are in the Private treatment.

Figure 4 shows that average non-Stakeholder contributions in the Public treatment always exceed those in the Private treatment. These differences are significant when pooled across rounds. As shown in Table 1, the average non-Stakeholder contribution is 14.4% larger in the Public (37.8% of endowment) than in the Private treatment (33.0%) and this difference is significant (paired Wilcoxon signed-rank test p=0.051).[15] However, on a round-by-round basis, the difference between Public and Private treatment contributions is only significant in rounds 1, 3, and 5, so the treatment-induced boost in cooperation is tenuous and may decline over time.

[15] Tests we report in the paper and tables use subject as unit of observation. If we use very strict tests in which comparisons are done at the group level (average group contributions), we find: in Table 1 the difference between Public and Private contributions is not significant (Mann-Whitney test p=0.471), Table 2 reciprocity results are essentially unchanged, and in Table 3 reciprocal action is still significant (p<0.01).

Thus, making a small low-cost change in the information structure of the game to allow targeted reciprocity increases cooperation by a modest but significant amount. This increase in contributions compares well to the effect of other public goods institutions that manipulate information on the contributions of others.[16] Targeted reciprocity need not increase cooperation, however: a reciprocator facing a largely uncooperative group may even contribute less overall when targeted reciprocity is possible than when it is not because he can discriminate against individual free-riders in the group.

We now turn to behavior in the Ineligible treatment. Table 1 shows that subjects who were randomly assigned to be Bachelors in the Ineligible treatment contribute significantly less (23.4% of endowment) in the Ineligible treatment than they did when these subjects were non-Stakeholders in the Private (29.9%) and Public (37.6%) treatments (paired Wilcoxon signed-rank test p=0.043 and p=0.020, respectively). Bachelors also give significantly less than non-Bachelor non-Stakeholders do in the Ineligible treatment (36.0% of endowment, Mann-Whitney test p=0.007). This is not true across all rounds, however; if Bachelor contributions are compared to the contributions of non-Bachelor non-Stakeholders in the rounds after the latter has had their last Stakeholder stint (those contributions averaging 21.9% of endowment) the levels are not significantly different (Mann-Whitney test p=0.794).

The low contributions of Bachelors in the Ineligible treatment, however, are not caused by idiosyncrasies of those subjects who were randomly assigned to be Bachelors. In the Private and Public treatments, the contributions of subjects who were assigned to be Bachelors in the

[16] Simply reporting disaggregated individual contributions, as in our Private treatment, rather than total contributions, increased contributions by 21% in Sell and Wilson (1991) but had no effect in Croson (2001). Revealing donors' contributions to each other increases contributions to external charities by 10% in Soetevent's (2005) study of church collections, but only for external causes and only temporarily. Andreoni and Petrie (2004) find an increase in giving of 59% when subjects see a photograph and contribution history for each group member, but no significant increase when only a photograph or only contribution history is revealed.

Ineligible treatment are not significantly different from the contributions of subjects who were not Bachelors in the Ineligible treatment.[17] The low contributions of Bachelors in the Ineligible treatment may be caused by several factors. First, Bachelors may feel a weakened urge to conform to a contribution norm because of their different role in the group. Second, their reduced earnings potential may render them less willing to trade off their payoff to benefit others. We explore Bachelors' behavior in more detail in the section on indirect reciprocity.

We might think that for non-Bachelors, shrinking the "circle of reciprocity" to four and adding a public good beneficiary outside that circle might dampen contributions. It does not. There is no difference between non-Bachelor non-Stakeholders' contributions in the Ineligible treatment (36.0% of endowment) and these subjects' behavior in the Private (33.8%) and Public (37.8%) treatments (paired Wilcoxon signed-rank test p=0.233 and p=0.410, respectively).

Finally, Figure 5 shows the cumulative distribution of non-Stakeholder contributions pooled across rounds of all sessions. In treatments that allow for targeted direct reciprocity (Public, non-Bachelors in Ineligible), the distribution lies to the right of the treatment in which the door to direct reciprocity is closed (Private). That is, there is a first-order shift in the distribution when direct reciprocity is possible. Subjects who were assigned to be Bachelors in the Ineligible treatment, however, give zero tokens more often in the Ineligible than the same subjects give in the Private (paired Wilcoxon signed-rank test p=0.007) and the Public (p=0.014) treatments. The distribution of contributions for all other treatments and roles lie to the right of

[17] Non-Stakeholder contributions in the Private treatment are not statistically different between subjects who were Bachelors in the Ineligible treatment and subjects who were not (Mann-Whitney test p=0.507). Similarly, in the Public treatment, subjects who were Bachelors in the Ineligible treatment do not give differently as compared to those who were non-Bachelors in the Ineligible treatment (Mann-Whitney test p=0.653). Subjects who were Bachelors in an early treatment may give less as non-Stakeholder in later treatments than do those who had been non-Bachelors, but this is only significant if the Ineligible treatment is first and if we compare only to the second treatment contributions (15.8% versus 37.9% of endowment, Mann-Whitney p=0.038).

that of Bachelors in the Ineligible treatment. Also, in the Public treatment, subjects give 20 tokens more often than in the Private (p=0.021) and Ineligible (p=0.098 for Bachelors, p=0.041 for non-Bachelors) treatments.

4.2 Direct Reciprocity and Other-Regarding Preferences

We next look at direct reciprocity, i.e. responsiveness to the current Stakeholder's past behavior. We first examine reciprocal acts in general, and then we isolate reciprocity rooted in other-regarding preferences. To do this, we use nonparametric within-subject tests of aggregate statistics and regression analysis. For each person, we examine whether, as non-Stakeholder, he gave more to the public good on average in rounds in which the current Stakeholder was previously generous to him as compared to rounds in which the current Stakeholder was previously ungenerous. Stakeholder past generosity is determined by the current Stakeholder's average contribution to the public good in rounds in which this subject was the Stakeholder.

To illustrate, assume that subject *i* is assigned position A and *j* is assigned B. Then *i* is Stakeholder in rounds 1 and 6, and *j* is Stakeholder in rounds 2 and 7. In round 2, *i* will remember how generous *j* was in round 1. "Stakeholder past generosity" will be *j*'s contribution in round 1. In round 7, when *j* is Stakeholder, our measure of "Stakeholder past generosity" for *i* would be the average of *j*'s contributions in rounds 1 and 6 when *i* was Stakeholder.

We look first at the effects of generosity using nonparametric tests. We define a "generosity threshold" $\bar{g}$ such that a group member whose cumulative average contribution in past rounds in which this subject had been Stakeholder is greater than this amount is called generous. For each subject, we calculate his average contribution across all rounds in which he

faced a Stakeholder whose past generosity met this threshold and his average contribution when he faced a Stakeholder whose past generosity did not. We tried many thresholds, including $\bar{g}$ = 6, 8, 10, 12, 15, and 19 tokens and an endogenous threshold calculated as the group's current cumulative average contribution. [18] All yielded the same qualitative and significant results. We report results from a threshold of $\bar{g} = 10$ tokens (50% of endowment) because it seems like a reasonable threshold for generosity that subjects may have in mind.

A subject shows reciprocal behavior if he gives more when facing a previously-generous Stakeholder than when facing a previously-ungenerous Stakeholder.[19] Table 2 presents averages of these measures in Panel I. In both treatments in which it is possible to attribute previous acts to a particular group member, generosity is rewarded: subjects give more to a previously-generous Stakeholder than to a previously un-generous Stakeholder (Public treatment 41.9% as compared to 27.1% of endowment, paired Wilcoxon signed-rank test p=0.000; Ineligible treatment 41.2% as compared to 23.9% of endowment, paired Wilcoxon signed-rank test p=0.000). As a placebo test, in the treatment in which previous generosity cannot be attributed to a particular individual, it is not rewarded (in the Private treatment, 27.6% was given to previously-generous Stakeholders as compared to 29.2% to previously-ungenerous Stakeholders, paired Wilcoxon signed-rank test p=0.773). We would not expect rewards in the Private treatment because subjects cannot tell who is Stakeholder or what the current Stakeholder did in

[18] To clarify further how the threshold works, suppose that a group contains only subjects *i*, *j*, and *k*. Subject *j* gave 15 tokens every time *i* was Stakeholder, and *k* always gave 2 tokens. Subject *i*'s average contribution to a generous Stakeholder is his average contribution when *j* was Stakeholder, and his average contribution to an ungenerous Stakeholder is his average contribution when *k* was Stakeholder.

[19] Readers who are concerned about the within-subjects design or intra-group correlation should note that this test relies on an individual subject treating members of his group (in a single treatment) in two different ways. Thus this result should be robust to either of those concerns.

the past. In the Public and Ineligible treatments, however, subjects give over 50% more to previously-generous Stakeholders than they give to previously-ungenerous Stakeholders.[20, 21]

Table 3 confirms these results with panel OLS fixed effects regressions. The first three columns show results for all rounds of each treatment. Non-Stakeholder contributions in each round g_{it} are regressed on characteristics of that round, including h_{ikt} (a summary of the current Stakeholder k's past generosity toward subject i: the current Stakeholder's cumulative average contributions when subject i was Stakeholder) and $\mathbf{X}it$ (other variables): $g_{it} = a + bh_{ikt} + CX_{it} + \varepsilon_{it}$. The coefficient b indicates reciprocal action (direct reciprocity or strategic self-interest). Group-level conditional cooperation could bias this coefficient upward, so we counteract that bias by including a control for group generosity in $\mathbf{X}it$: the group's cumulative average non-Stakeholder contribution in past rounds. For each subject in each round, this measure excludes his own past contributions and those of the current Stakeholder.[22] All regressions use individual subject fixed effects, and errors are clustered by group to address issues of contemporaneous correlation and correlation across rounds within the group.[23, 24] The coefficients on past generosity are significant and positive in the Ineligible and Public treatments (which are not significantly different from each other: p=0.535), but not in the Private treatment. Group average contributions do not significantly influence contributions in this treatment. The linear trend for round numbers is significant and negative only for one treatment (Ineligible). While one might

[20] Subjects could respond to past actions of the current Stakeholder and the current non-Stakeholders, although their contributions benefit the former three times as much as the latter. If subjects were responding to non-Stakeholders in this way, this would attenuate our within-subject measure of responsiveness to Stakeholder history.
[21] Misplaced reciprocal behavior could happen if subjects believe their behavior in an early treatment can affect outcomes in later treatments. Our finding of reciprocity is robust to this phenomenon because responsiveness is still significant in the Public treatment even when the Public treatment is the last treatment experienced.
[22] In the Ineligible treatment, this group measure also excludes data from the Bachelor (although the same results obtain if the Bachelor's data is included).
[23] Results change little when errors are clustered by session (of which there are 6) rather than group. Interactions in past treatments could cause inter-group correlation of the error term, but clustering by session removes that concern.
[24] If group level random effects are used (along with individual fixed effects) instead of clustering on group, results change little except that Public post-last Stakeholder stint reciprocity just loses significance (p=0.108).

expect a discontinuous dropoff in contributions after a subject's last Stakeholder stint, we only see this in the Private treatment.

Another way to look at direct reciprocation is to examine responsiveness, the difference between the amount given to a generous Stakeholder and the amount given to an ungenerous Stakeholder. Responsiveness is greater in both the Public (14.8% of endowment) and Ineligible (17.3% of endowment) treatments than in the Private (-1.6% of endowment) treatment (paired Wilcoxon signed-rank test p=0.000 and p=0.000, respectively). Responsiveness does not differ between the Public and Ineligible treatments (p=0.966). These results are confirmed in Table 4, which shows a stacked regression of responsiveness controlling for individual fixed effects and group-level correlation. Subjects are significantly more likely to give more to a "nice" Stakeholder in the Public and Ineligible treatments. These tests give us evidence of directly reciprocal behavior, although this behavior may be rooted in strategic self-interest.

We dispose of strategic concerns by looking for reciprocal giving after a person has passed his last Stakeholder stint. For example, again assume that subject *i* was Stakeholder in rounds 1 and 6 and *j* was Stakeholder in rounds 2 and 7. Assume no further rounds follow. If *j* was kind to *i* in rounds 1 and 6, will *i* reciprocate in round 7? If *i* is purely strategic, he has little to gain, so he should not contribute and therefore not reciprocate. This is also the logic that was used in one test of reciprocity rooted in other-regarding preferences in Cabral et al. (2012). It is notable in Panel II of Table 2 that responsiveness before the last Stakeholder stint is stronger than it is on average across all rounds; is it significant after the last Stakeholder stint?

We construct statistics of each subject's average contribution to previously generous and ungenerous Stakeholders after this subject's last Stakeholder stint. These results, in Panel III of

Table 2, show that subjects are not solely motivated by strategic concerns. Other regarding preferences are more strongly at play. Subjects in the Public treatment continue to discriminate between generous Stakeholders (to whom they give 30.0% of endowment) and ungenerous Stakeholders (16.9%) even when they have no strategic motive to do so. This difference is statistically significant (paired Wilcoxon signed-rank test p=0.026). In the Ineligible treatment, the sample size is reduced because it includes fewer subjects (only non-Bachelors) and fewer rounds. Post-last Stakeholder stint contributions to generous and ungenerous Stakeholders in the Ineligible treatment are compressed toward each other as compared to the Public treatment, but the amount of responsiveness in the Ineligible treatment is not statistically different from the Public treatment (Mann-Whitney test p=0.422). Discrimination in late rounds of the Ineligible treatment is not statistically significant, although the difference is in the same direction.

These results are confirmed in Table 3. The last three columns repeat the regressions in the first three columns but only use contributions in rounds after the last Stakeholder stint, when the subject has no strategic incentive to reward past good behavior. The coefficient on past generosity is significant and positive in the Public treatment, and the direction is the same in the Ineligible treatment, but it is not significant (although it is not statistically different from that in the Public treatment, p=0.563).

4.3 Indirect Reciprocity

We now examine evidence of indirect reciprocity. We define this type of positive reciprocity as the action of a disinterested party who rewards one agent for past generosity toward another agent. This kind of indirect reciprocity cannot be tested by looking at the responsiveness of a subject who has the opportunity to be Stakeholder because he is not

disinterested. He can benefit from fostering a norm of cooperation. The behavior of Bachelors in the Ineligible treatment gives us a clean test of indirect reciprocity. These subjects will never be Stakeholder, so they can never benefit directly from targeting generosity. Therefore, direct reciprocity and strategic self-interest cannot drive Bachelor giving. Bachelors can't even receive great benefits from conditional cooperation.[25]

We have shown that Bachelors give significantly less than non-Stakeholders in the Ineligible treatment and less than they themselves give when they were non-Stakeholders in other treatments. The subjects who were Bachelors in the Ineligible treatment also behave reciprocally when those subjects are part of the circle of reciprocity in another treatment: they have positive responsiveness in the Public treatment (giving 35.0% of endowment to previously generous Stakeholders and 22.3% of endowment to previously ungenerous Stakeholders, paired Wilcoxon signed-rank p=0.004). However, in the Ineligible treatment, Bachelors do not give more when facing previously generous Stakeholders. They give 21.9% of endowment when facing previously generous Stakeholders and 25.5% when facing previously ungenerous Stakeholders.[26] These levels are not significantly different (paired Wilcoxon signed-rank p=0.823). Therefore, indirect reciprocity does not seem to be a motivator in this setting.[27]

[25] Bachelors do benefit from a cooperative group, since their payoff is a function of group contributions; however, there is no way for a Bachelor to directly benefit by being nice to previously generous Stakeholder.

[26] For Bachelors, the current Stakeholder's past generosity is defined by how much the current Stakeholder gave on average as non-Stakeholder in all past rounds.

[27] We also tested whether contributions by non-Bachelor non-Stakeholders in the Ineligible treatment and by non-Stakeholders in the Public treatment were affected by past generosity to self and past generosity to others. Current Stakeholder past contribution to others is indeed significant in a regression; however, we caution against interpreting this as indirect reciprocity for two reasons. First, the correlation between Stakeholder past generosity to self and Stakeholder past generosity to others is high in all cases (although inclusion of this variable does not greatly affect the coefficient on Stakeholder past generosity to self). Second, and more importantly, a person who has been and will be Stakeholder may see someone else's generosity even in a round in which this subject is not Stakeholder as a signal of how this group member will behave when this subject IS Stakeholder, and thus he may have a great interest in promoting good behavior. A person in this position is far less disinterested than a Bachelor.

This result is intriguing because other studies have found evidence of indirect reciprocity, even in cases of positive rather than negative reciprocity (e.g., Engelmann and Fischbacher, 2009; Seinen and Schram, 2006). The re-matching structure of those experiments allows subjects to have a financial interest in the group's overall cooperation and propensity to reciprocate because they will eventually be direct beneficiaries. Subjects, therefore, are not as disinterested as our Bachelors. Both conditions are valid settings in which to examine cooperation. However, our results align more closely with the "bystander" interpretation of indirect reciprocity.

5. Conclusions

Favor trading is a natural element of social networks, and public good provision may benefit from tapping into it. In an experiment that allows different forms of reciprocity to be turned on and off, we explore the power of favor trading and the underlying reciprocal motivations through which institutions such as grassroots fundraising may work. In our setting, favor trading increases cooperation by a modest but significant 14.4%, although that boost may decay over time. However, we observe clear evidence of favor-trading in public good contributions, and some of this favor-trading is rooted in other-regarding preferences (is non-strategic).[28] Our setting in which subjects must come up with the idea of using reciprocity in this public goods game is somewhat richer than the classic trust or investment game (e.g., Berg et al., 1995) in which reciprocity has been studied in the past. Our results provide a step toward understanding how institutions outside of the lab can take advantage of these motives for social gain. In particular, our reciprocity findings are a complement to the result of DellaVigna et al. (2012) that social pressure is important in door-to-door fundraising.

[28] Cabral, Ozbay, and Schotter (2012) find that strategic forces are strong in their repeated veto game. As they note, while contributions do decline in final rounds (indicating that strategic concerns matter), they do not decline to zero, so other-regarding preferences play a role there as well.

Intriguingly, we find no evidence of indirect reciprocity in our Ineligible treatment. While most experiments studying indirect reciprocity allow agents to directly benefit from increased group cooperation, we strip our Bachelor of strong incentives to foster cooperative norms. When given the costly chance to reward a kind act without the possibility of future direct reciprocation, our Bachelors provide no such rewards. This also stands in contrast to third party punishment/reward games (e.g. Fehr and Fischbacher, 2004; Almenberg et al., 2011) in which disinterested parties will pay to punish bad behavior or reward good behavior. In those games, it is arguably costlier for the third party to take action. It is possible that one difference is that in those games the disinterested third party's attention is totally focused on the decision to punish or reward, while in ours such an action would occur through a change in contribution behavior.

How can it be that information conditions activate favor-trading and yet the increase in contributions we observe is so small? A reciprocator who has information that allows him to discriminate and target people's pet causes based on those peoples' past actions may find enough of them to be stingy that he does not increase, and perhaps even reduces, his overall contribution compared to the situation in which he does not have this information.

What is the true effect of reciprocal giving in non-lab social environments? It might increase contributions more than seen here; social relations last longer than lab relationships. However, it is possible that people either will contribute only to their own pet causes or will offset increased contributions to a friend's charity by reduced contributions to other charities. A charity may use social fundraising to increase its own revenues while crowding out contributions to other organizations. These phenomena are largely precluded in our experiment because the public good is linear. Field studies could answer this question, but most theories of other-regarding preferences should imply that if there is such crowd-out, it is incomplete.

References

Almenberg, Johan, Dreber, Anna, Apicella, Coren L. and Rand, David G., 2011, Third Party Reward and Punishment: Group Size, Efficiency and Public Goods. In: Nicholas M. Palmetti and Jennifer P. Russo (Ed.), *Psychology of Punishment*. Nova Publishing, pp. 73-92.

Andreoni, James, William Harbaugh, and Lise Vesterlund. 2003. "The Carrot or the Stick: Rewards, Punishments, and Cooperation." *American Economic Review*, 93:3, pp. 893-902.

Andreoni, James and Ragan Petrie. 2004. "Public goods experiments without confidentiality: A glimpse into fund-raising." *Journal of Public Economics*, 88:7-8, pp. 1605-23.

Berg, Joyce, John W. Dickhaut, and Kevin A. McCabe. 1995. "Trust, Reciprocity, and Social History." *Games and Economic Behavior*, 10:1, pp. 122-42.

Bochet, Olivier, Talbot Page, and Louis Putterman. 2006. "Communication and Punishment in Voluntary Contribution Experiments." *Journal of Economic Behavior and Organization*, 60:1, pp. 11-26.

Cabral, Luis, Ozbay, Erkut Y. and Schotter, Andrew, 2012. "Intrinsic and instrumental reciprocity: An experimental study." Working paper. New York University

Carpenter, Jeffrey and Erika Seki. 2011. "Do Social Preferences Increase Productivity? Field Experimental Evidence from Fishermen in Toyama Bay." *Economic Inquiry*, 49:2, pp. 612-630.

Charness, Gary and Matthew Rabin. 2002. "Understanding Social Preferences with Simple Tests." *Quarterly Journal of Economics*, 117:3, pp. 817-69.

Chaudhuri, Ananish, 2011. "Sustaining Cooperation in Laboratory Public Goods Experiments: A Selective Survey of the Literature." *Experimental Economics*, 14(1), 47-83.

Cox, James C. 2004. "How to Identify Trust and Reciprocity." *Games and Economic Behavior*, 46:2, pp. 260-81.

Cox, James C., Daniel Friedman, and Vjollca Sadiraj. 2008. "Revealed Altruism." *Econometrica*, 76:1, pp. 31-69.

Croson, Rachel T. A. 2001. "Feedback in Voluntary Contribution Mechanisms: An Experiment in Team Production," in *Research in experimental economics. Volume 8*. R. Mark Isaac ed: Amsterdam; New York and Tokyo: Elsevier Science, JAI, pp. 85-97.

DellaVigna, Stefano, List, John A. and Malmendier, Ulrike, 2012. "Testing for Altruism and Social Pressure in Charitable Giving." *Quarterly Journal of Economics*, 127(1), 1-56.

Dreber, Anna, Fudenberg, Drew and Rand, David G., 2013, "Who Cooperates in Repeated Games: The Role of Altruism, Inequity Aversion, and Demographics." Working paper.

Engelmann, Dirk and Urs Fischbacher. 2009. "Indirect reciprocity and strategic reputation building in an experimental helping game." *Games and Economic Behavior*, 67:2, pp. 399-407.

Fehr, Ernst and Fischbacher, Urs, 2004. Third-party punishment and social norms. *Evolution and Human Behavior*, 25:2, pp. 63-87.

Fehr, Ernst and Simon Gächter. 2000. "Cooperation and Punishment in Public Goods Experiments." *The American Economic Review*, 90:4, pp. 980-94.

Fehr, Ernst, George Kirchsteiger, and Arno Riedl. 1993. "Does Fairness Prevent Market Clearing? An Experimental Investigation." *Quarterly Journal of Economics*, 108:2, pp. 437-59.

Fischbacher, Urs. 2007. "z-Tree: Zurich Toolbox for Ready-Made Economic Experiments." *Experimental Economics*, 10:2, pp. 171-78.

Fischbacher, Urs, Simon Gachter, and Ernst Fehr. 2001. "Are People Conditionally Cooperative? Evidence from a Public Goods Experiment." *Economics Letters*, 71:3, pp. 397-404.

Gächter, Simon. 2007. "Conditional cooperation: Behavioral regularities from the lab and the field and their policy implications," in *Economics and Psychology. A Promising New Cross-Disciplinary Field*. Bruno S. Frey and Alois Stutzer eds. Boston, MA: The MIT Press.

Girl Scouts of the USA. 2011. "Girl Scout Cookies and Learning Life Skills." http://www.girlscouts.org/program/gs_cookies/, retrieved January 2011.

Glöckner, Andreas, Bernd Irlenbusch, Sebastian Kube, Andreas Nicklisch, and Hans T. Normann. 2011. " Leading with(out) Sacrifice? A Public-Goods Experiment with a Privileged Player." *Economic Inquiry*, 49:2, pp. 591-597.

Gneezy, Uri and John A. List. 2006. "Putting Behavioral Economics to Work: Testing for Gift Exchange in Labor Markets Using Field Experiments." *Econometrica*, 74:5, pp. 1365-84.

Goeree, Jacob, Charles Holt, and Susan Laury. 2002. "Private Costs and Public Benefits: Unraveling the Effects of Altruism and Noisy Behavior." *Journal of Public Economics*, 83:2, pp. 255-76.

Houser, Daniel, Erte Xiao, Kevin McCabe, and Vernon Smith. 2008. "When Punishment Fails: Research on Sanctions, Intentions and Non-cooperation." *Games and Economic Behavior*, 62:2, pp. 509-32.

Isaac, R Mark and Norton, Douglas A, 2013. "Endogenous institutions and the possibility of reverse crowding out." *Public Choice*, 156(1-2), 253-284.

Jackson, Matthew O., Rodriguez-Barraquer, Tomas and Tan, Xu, 2012. "Social Capital and Social Quilts: Network Patterns of Favor Exchange." *American Economic Review*, 102(5), 1857-1897.

Klein, Daniel B. 1990. "The Voluntary Provision of Public Goods? The Turnpike Companies of Early America." *Economic Inquiry*, 28:4, pp. 788-812.

La Ferrara, Eliana, 2003. "Kin Groups and Reciprocity: A Model of Credit Transactions in Ghana." *American Economic Review*, 93(5), 1730-1751.

Ledyard, John O. 1995. "Public goods: A survey of experimental research," in *The Handbook of Experimental Economics*. John H. Kagel and Alvin E. Roth eds. Princeton, NJ, USA: Princeton University Press.

List, John A. 2006. "The Behavioralist Meets the Market: Measuring Social Preferences and Reputation Effects in Actual Transactions." *Journal of Political Economy*, 114:1, pp. 1-37.

Long, Stephen H. 1976. "Social Pressure and Contributions to Health Charities." *Public Choice*, 28:Winter, pp. 55-66.

Masclet, David, Charles Noussair, Steven Tucker, and Marie-Claire Villeval. 2003. "Monetary and Nonmonetary Punishment in the Voluntary Contributions Mechanism." *American Economic Review*, 93:1, pp. 366-280.

Meer, Jonathan and Rosen, Harvey S., 2011. "The ABCs of charitable solicitation." *Journal of Public Economics*, 95(5–6), 363-371.

Milinski, Manfred, Dirk Semmann, and Hans-Jurgen Krambeck. 2002. "Reputation helps solve the 'tragedy of the commons'." *Nature*, 415:6870, pp. 424-26.

Nowak, Martin A. and Karl Sigmund. 2005. "Evolution of indirect reciprocity." *Nature*, 437:7063, pp. 1291-98.

Olson, Mancur., 1965, *The logic of collective action: public goods and the theory of collective action*. Harvard University Press, Cambridge, Mass.

Preston, Caroline. 2010. "Charities Must Find Multiple Ways to Persuade People of Different Generations to Give, Study Finds." *The Chronicle of Philanthropy*: Washington, DC.

Rabin, Matthew. 1993. "Incorporating Fairness into Game Theory and Economics." *American Economic Review*, 83:5, pp. 1281-302.

Reuben, Ernesto and Riedl, Arno, 2009. "Public Goods Provision and Sanctioning in Privileged Groups." *Journal of Conflict Resolution*, 53:1, pp. 72-93.

Reuben, Ernesto and Suetens, Sigrid, 2012a. "Revisiting strategic versus non-strategic cooperation." *Experimental Economics*, 15:1, pp. 24-43.

Reuben, Ernesto and Suetens, Sigrid, 2012b, "On the cognitive and motivational components of reciprocity." Working paper.

Saijo, Tatsuyoshi and Nakamura, Hideki, 1995. "The "Spite" Dilemma in Voluntary Contribution Mechanism Experiments." *The Journal of Conflict Resolution*, 39(3), 535-560.

Seinen, Ingrid and Arthur Schram. 2006. "Social Status and Group Norms: Indirect Reciprocity in a Repeated Helping Experiment." *European Economic Review*, 50:3, pp. 581-602.

Sell, Jane and Rick K. Wilson. 1991. "Levels of information and contributions to public goods." *Social Forces*, 70:1, pp. 107-24.

Sobel, Joel. 2005. "Interdependent Preferences and Reciprocity." *Journal of Economic Literature*, 43:2, pp. 392-436.

Soetevent, Adriaan R. 2005. "Anonymity in Giving in a Natural Context--A Field Experiment in 30 Churches." *Journal of Public Economics*, 89:11-12, pp. 2301-23.

Stanca, Luca. 2009. "Measuring indirect reciprocity: Whose back do we scratch?" *Journal of Economic Psychology*, 30:2, pp. 190-202.

Trivers, Robert L. 1971. "The evolution of reciprocal altruism." *The Quarterly Review of Biology*, 46:1, pp. 35-57.

Figures and Tables

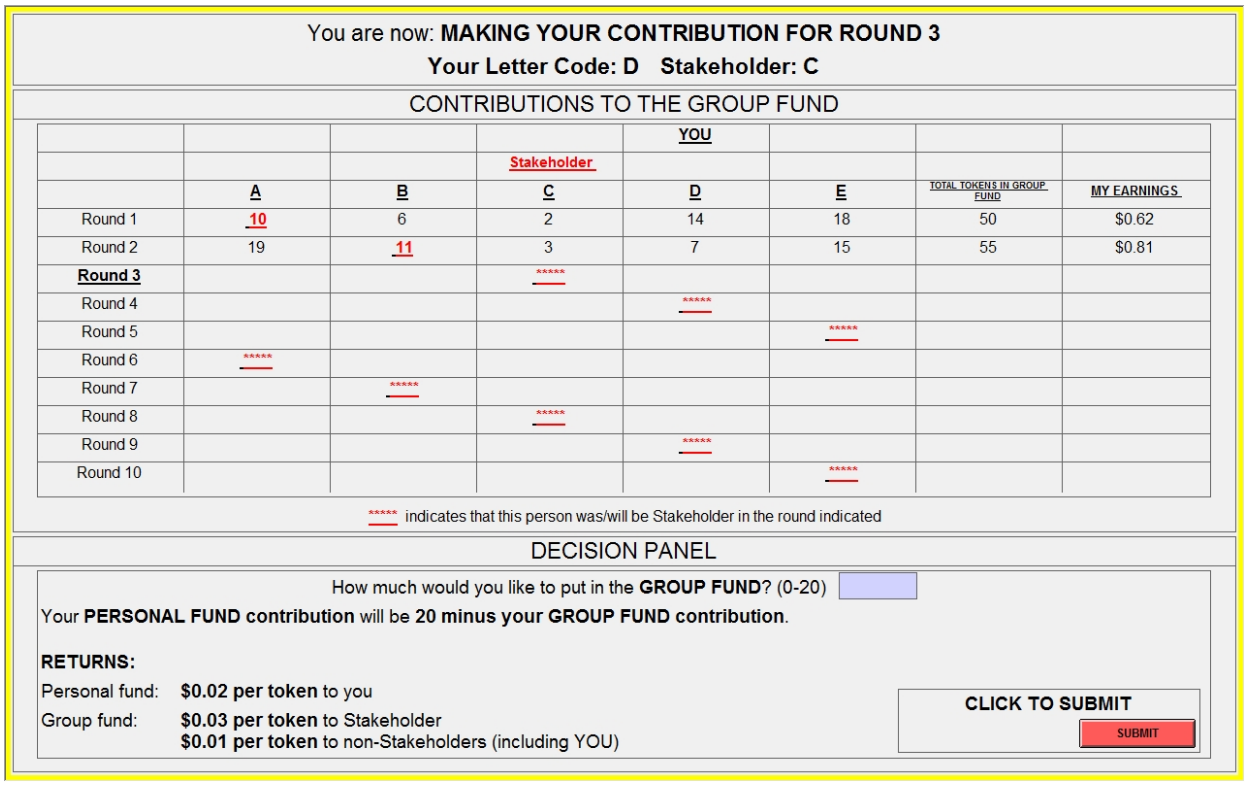

Figure 1. Public Treatment Decision Screen

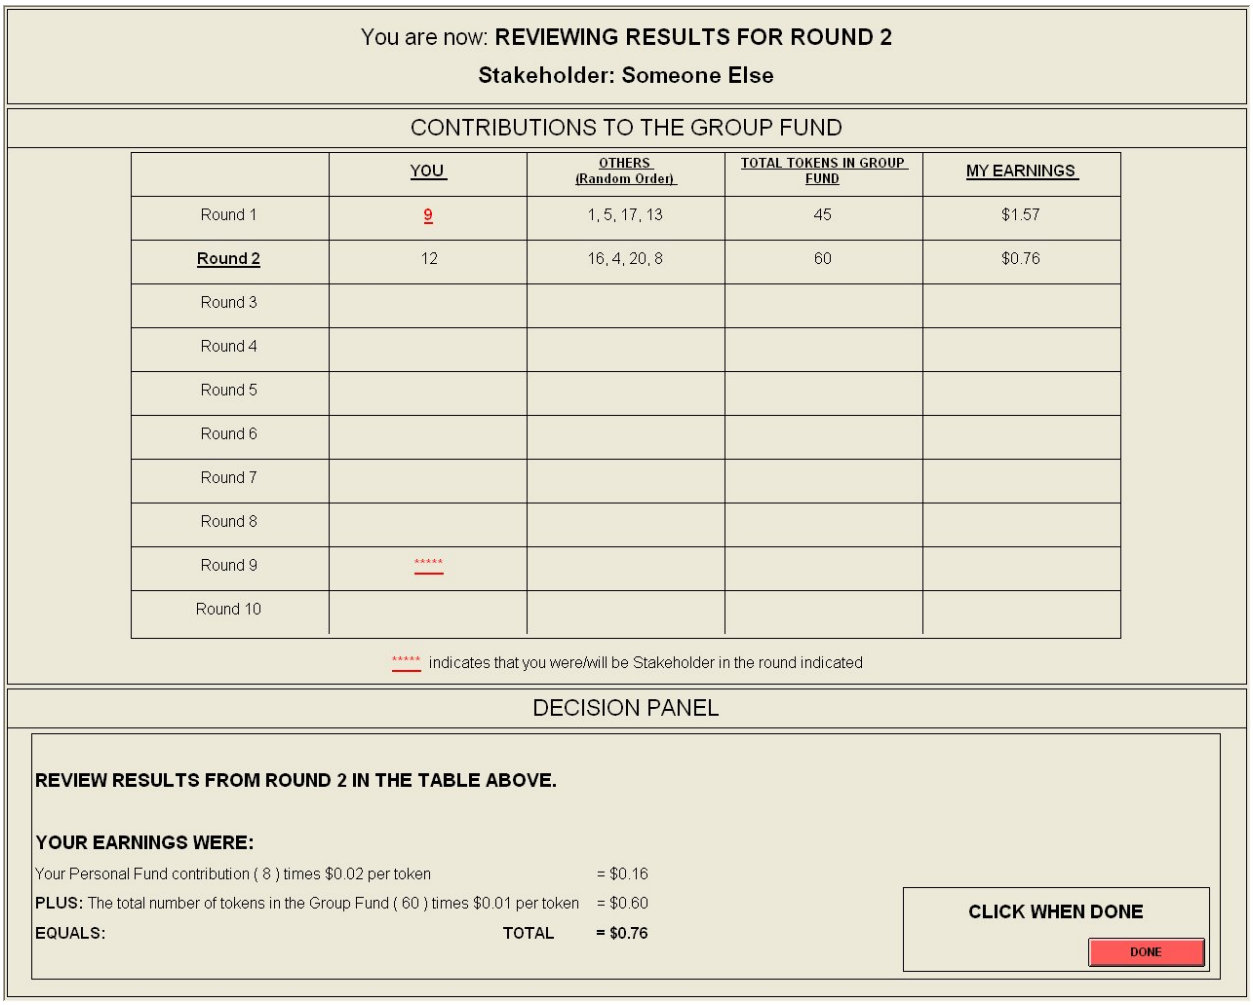

Figure 2. Private Treatment Review Screen

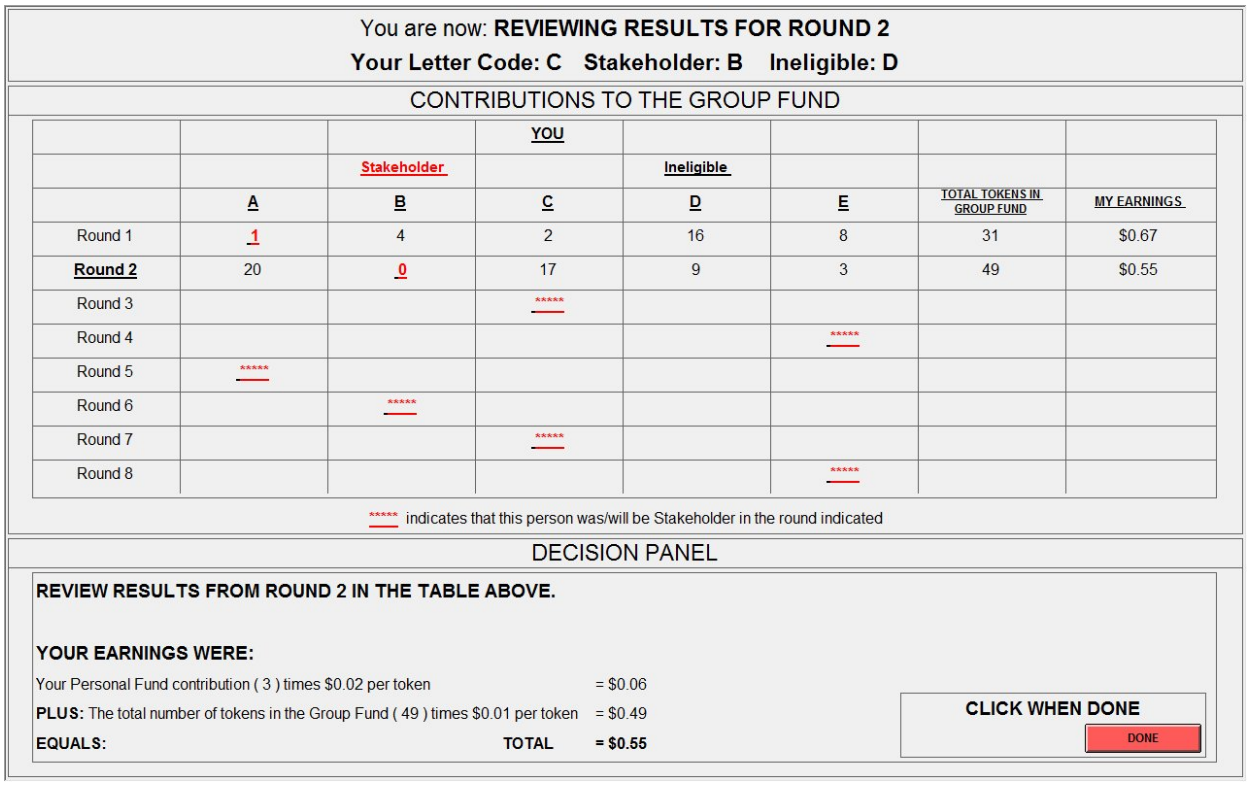

Figure 3. Ineligible Treatment Review Screen

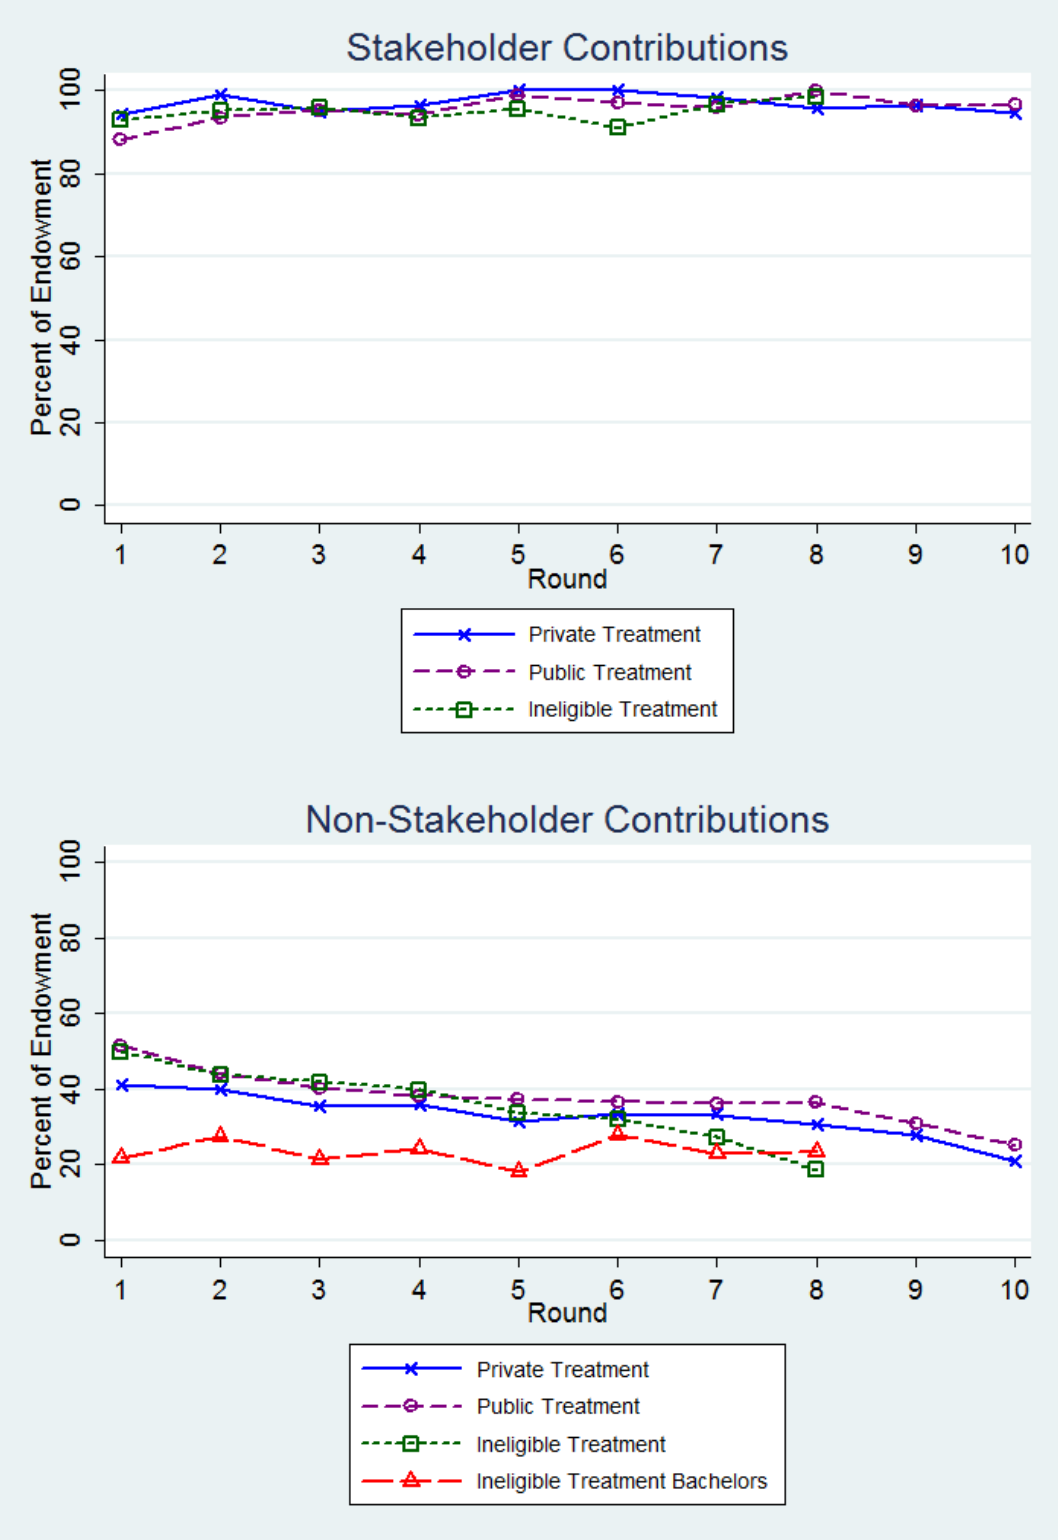

Note: Round numbers indicate round number within a treatment; data pooled across treatment orders

Figure 4. Contributions by Treatment and Role across Rounds (in Percent of Endowment)

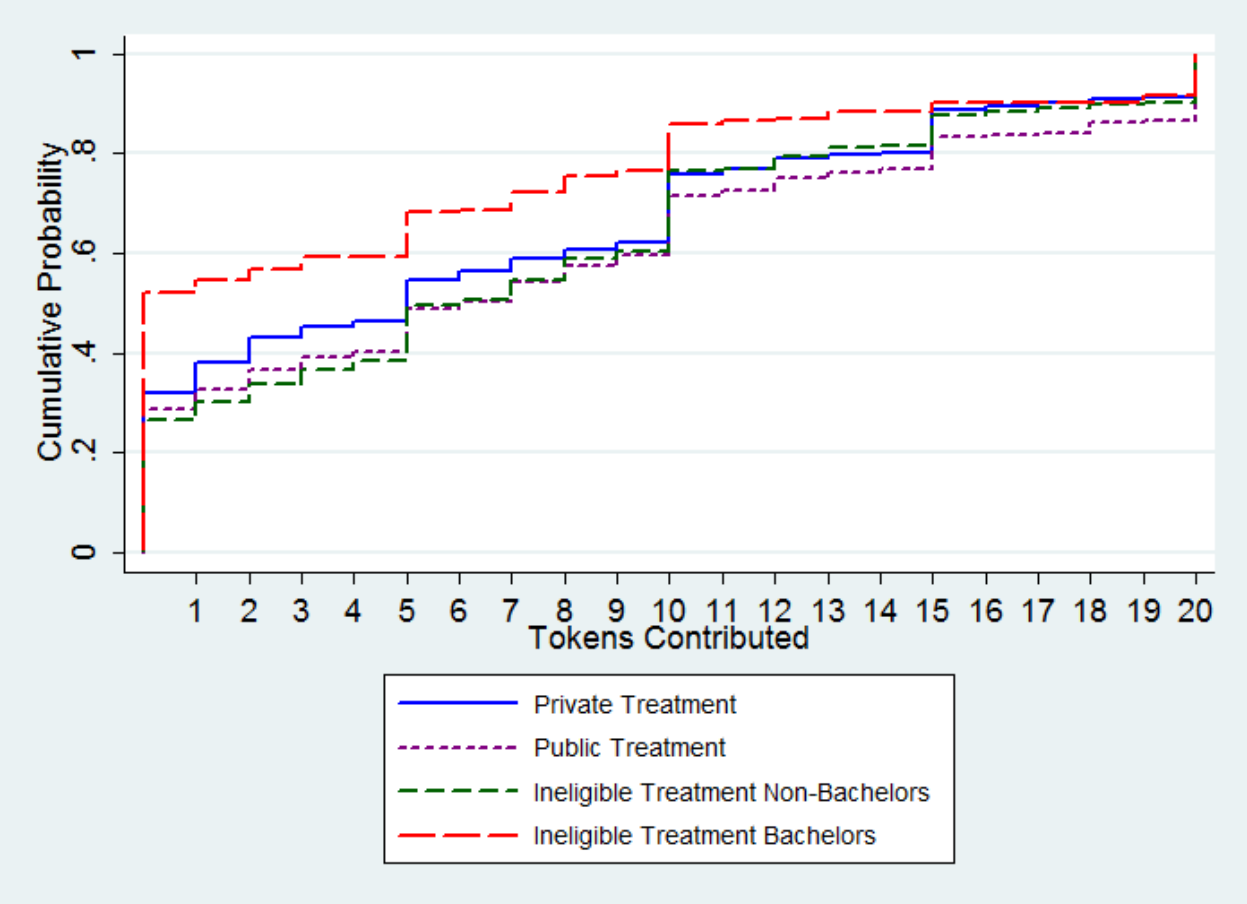

Note: unit of observation is an individual's contribution in a single round; data are pooled across rounds.

Figure 5. Cumulative Distribution of Non-Stakeholder Contribution Amounts (in Tokens)

Table 1. Average Contributions by Treatment and Role (Percent of Endowment)

	Stakeholder	Non-Stakeholder	Non-Stakeholders who had role of Bachelor in Ineligible treatment	Non-Stakeholders who did not have role of Bachelor in Ineligible treatment
Private Treatment	96.98 (10.30) N=120	33.02 (25.98) N=120	29.90 (26.89) N=24	33.80 (25.83) N=96
Public Treatment	95.63 (10.36) N=120	37.75 (25.93) N=120	37.60 (30.42) N=24	37.79 (24.86) N=24
Ineligible Treatment, Non-Bachelors	94.95 (11.74) N=96	35.96 (22.26) N=96	N/A	35.96 (22.26) N=96
Ineligible Treatment, Bachelors	N/A	23.39 (28.41) N=24	23.39 (28.41) N=24	N/A

Standard deviations in parentheses

Table 2. Non-Stakeholder Contributions by Subject Averaged across Rounds by Stakeholder's Past Generosity toward Subject (in Percent of Endowment)

	Private Treatment	Public Treatment	Ineligible Treatment (excl. Bachelors)
Panel I: All Rounds			
Generous stakeholder	27.57	41.92	41.22
	(26.62)	(32.06)	(29.11)
Ungenerous stakeholder	29.16	27.09	23.92
	(29.59)	(27.63)	(24.08)
N	82	95	75
p-value (paired Wilcoxon signed-rank test)	0.773	0.000	0.000
	Private Treatment	Public Treatment	Ineligible Treatment (excl. Bachelors)
Panel II: Only Pre Last Stakeholder Stint			
Generous stakeholder	35.02	42.44	46.29
	(27.59)	(32.27)	(32.40)
Ungenerous stakeholder	34.18	31.59	29.31
	(29.65)	(32.17)	(27.87)
N	53	91	58
p-value (paired Wilcoxon signed-rank test)	0.794	0.000	0.000
Panel III: Only Post Last Stakeholder Stint			
Generous stakeholder	14.60	30.00	26.60
	(25.70)	(30.99)	(31.98)
Ungenerous stakeholder	18.60	16.93	19.30
	(29.24)	(27.43)	(27.75)
N	31	41	25
p-value (paired Wilcoxon signed-rank test)	0.436	0.026	0.435

Standard deviations in parentheses. Note: Generous stakeholder is one who gave ≥ 10 on average in past rounds in which subject was a stakeholder. Ungenerous stakeholders gave < 10 in past rounds. Paired signed-rank test based on average contributions across all subjects within a group when facing a generous or non-generous stakeholder. Standard deviations are in parentheses. *N*'s are less than 120 (or 96, which is the maximum number for the Ineligible treatment) because some subjects did not face both a generous and an ungenerous Stakeholder, and/or because those subjects do not have more than two Post-Last Stakeholder stint rounds.

Table 3. OLS Fixed Effects Panel Regression of Non-Stakeholder Contribution (in Percent of Endowment) on Period-Level Covariates, Overall and Post-Last Stakeholder Stint

	Private	Public	Ineligible	Private Post-Last	Public Post-Last	Ineligible Post-Last
Stakeholder average past contributions to me	0.05	0.24***	0.20*	-0.02	0.21*	0.03
	(0.06)	(0.05)	(0.10)	(0.11)	(0.11)	(0.28)
Group average contributions	0.25	0.05	-0.12	0.39	-0.09	-0.14
	(0.16)	(0.17)	(0.19)	(0.44)	(0.32)	(0.55)
Round number	0.03	-1.69	-2.93**	-4.12**	-5.90***	-5.30**
	(0.65)	(1.22)	(1.33)	(1.83)	(1.96)	(2.17)
Post-last Stakeholder stint? (dummy)	-9.66**	-1.66	-5.34			
	(3.86)	(5.05)	(5.37)			
Constant	22.15***	34.65**	46.81***	48.39**	76.14***	67.07**
	(7.48)	(13.68)	(15.08)	(22.93)	(25.80)	(28.38)
Observations (rounds)	720	720	432	240	240	144
Number of subjects	120	120	96	96	96	72
R2 (within)	0.037	0.115	0.194	0.061	0.140	0.078
F	4.50	13.37	35.58	2.03	3.70	2.75

Robust standard errors in parentheses; errors are clustered on group (of which there are 24); individual fixed effects.
* significant at 10%; ** significant at 5%; *** significant at 1%

Table 4. Panel Stacked Regression of Average Contribution (in Percent of Endowment) on "Nice Dummy"

	Private Treatment	Public Treatment	Ineligible Treatment (excl. Bachelors)
"Nice dummy"	-1.59	14.83***	17.30***
	(2.72)	(2.68)	(2.72)
Constant	29.16***	27.09***	23.92***
	(1.36)	(1.34)	(1.36)
Subjects	82	95	75
Observations	164	190	150
R^2 (overall)	-0.001	0.058	0.096
F	0.34	30.54	40.37

Robust standard errors in parentheses; errors clustered on group (of which there are 24); individual fixed effects. There are two observations per subject: one to summarize the subject's average contribution when facing a Stakeholder who had been generous to him, and one for his average contribution when facing a previously-ungenerous Stakeholder. The "nice dummy" is the indicator that differentiates these two observations for each subject. * significant at 10%; ** significant at 5%; *** significant at 1%