

Favor Trading in Public Good Provision

Sarah Jacobson¹
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

June 2012

Abstract

Favor trading between non-kin people 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 reciprocity. It is not yet well-understood which of these 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. Charitable fundraising is often done through social networks, and peer-to-peer reciprocity may raise contribution levels. 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. We find that strategic concerns cannot explain all of this behavior, and that it must be at least partly due to direct reciprocity. When someone cannot 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, since they cannot benefit from a norm of cooperation. 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 have been attributed to motives such as strategic self-interest, reciprocal altruism, and indirect reciprocity. While any of these may be plausible explanations for reciprocal exchanges, we know little about which factors drive behavior and how they interact to increase or reduce cooperative behavior. This is important to understand because each motive has different implications for the success of reciprocal environments. In institutions where cooperation is pro-social, the ability to return favors and the information that enables such an action may be important determinants of the efficiency of outcomes.¹

In this paper, we study the motives behind favor trading and attempt to unpack them to better understand which are most salient. While favor trading occurs in many environments, we study the context of the private provision of a public good because favor trading in this setting may serve a special pro-social role. We use linear public goods experiments to explore this because they cleanly simulate a social environment in which cooperation increases efficiency.

¹ Social favor trading may drive behavior in many cooperative settings. The same dynamic may exist when time and effort are solicited, such as in a volunteer advisory board. Klein (1990) argues that peer pressure played a role in the funding of turnpikes in early America.

We systematically manipulate the characteristics of the institution to study the importance of strategy, direct reciprocity, and indirect reciprocity.

Our study can also be understood in the context of grassroots fundraising for charities. Charities frequently encourage supporters to tap into their social networks to increase donations. 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. Social media has made this process easier than ever: for example, charities use internet tools like Facebook Causes and DonorPages to increase donations. About one third of young people use email and social networking sites to inform friends of a charity (Preston, 2010). Brick-and-mortar institutions like the Girl Scouts of the USA use a combination of online solicitation and personal contact to fundraise through networks: Girl Scout troops send girls out into their communities every year to sell cookies in a \$700 million business to raise money for their Scouting activities (Girl Scouts of the USA, 2011).

The popularity of this social fundraising suggests that the method is fruitful. An essential element of this institution's success may be a social network's ability to leverage favor trading among peers. This favor trading is possible because of heterogeneous preferences for public goods and because the institution makes contribution information public. Our results, seen in this context, are important to both practitioners and researchers. We identify the role of reciprocity separately from the role of strategy as both work to boost contributions. We show that tapping into social networks is a low-cost and effective way to raise contributions to a public good.

To illustrate how favor trading in support of a public good may function, we present an example. Suppose Joe and Frank are coworkers. Joe's son is in Boy Scouts and Frank's daughter

is in Girl Scouts. Each organization has an annual fundraiser: Joe's Boy Scout sells popcorn in January and Frank's Girl Scout sells cookies in February. For each fundraiser, the children and their families solicit contributions from their social networks. Also, purchase (donation) decisions are fairly public. Each man has an interest in his own child's fundraiser because of the family's stake in the organization's success. As a result, Joe will buy popcorn and Frank will buy cookies. Also, Joe wants Frank to buy popcorn and Frank wants Joe to buy cookies. Even if Frank does not care about Boy Scouts (or popcorn), he may buy Joe's popcorn in the hope that Joe will reciprocate when Frank's cookie fundraiser comes around next month. These peers use the promise of future reciprocation to pressure each other to give to charity.

This reciprocal giving 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. It may be driven by direct reciprocity (i.e. reciprocal altruism), models of which include Rabin (1993) and Cox et al. (2008), in which case the giver's preferences are increasing in the payoff of a past benefactor. 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). Indirect reciprocity may play a role if an

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" (or "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).

actor rewards a good deed done to a third party (e.g. Nowak and Sigmund 2005; Seinen and Schram 2006, Engelmann and Fischbacher 2009).³

While there is ample evidence that people contribute to public goods, it is also clear that cooperation is difficult to sustain (Ledyard, 1995; Chaudhuri, 2011). Reciprocal-type forces based on subjects' reputations may enable 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). There is growing evidence that social forces also drive giving outside the lab; for example, DellaVigna et al. (2012) find evidence of “social pressure” in door-to-door fundraising.

Our experimental design differs from previous work in that we can identify which institutional elements are important for reciprocity to boost giving to a public good.⁴ We are also able to determine the relative importance of direct reciprocity rooted in other regarding preferences, indirect reciprocity, and strategic self-interest. We do not examine network formation and solicitation, but instead focus on the roles of social preferences and information in determining contribution behavior in a group.

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 (just as Frank cares about and thus has a stake in the Girl Scouts’

³ Of the many types of indirect reciprocity that have been studied, the 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).

⁴ 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. This was because a more personal solicitation allowed the fundraiser to exert peer pressure on the donor. However, the analysis will likely overestimate this relationship. Personal solicitations are more costly to perform, so charities may limit their most personal solicitations to donors known to be generous.

success). A group with a Stakeholder is similar to Olson's (1965) "privileged group," as studied in Reuben and Riedl (2009).⁵ We make the privileged group even more effective by making the Stakeholder position rotate through group members, creating opportunities for targeted reciprocal acts. The game is repeated and every person in the group has a turn at being the Stakeholder. To see the role of reciprocity, we compare behavior when subjects have full information regarding each other's actions (and thus can condition their giving on others' past acts) and when they do not have that information.

In a final treatment, we 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 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. Our treatment provides a stricter test.

We find that the effects of allowing people to engage in targeted reciprocity are significant. Average contributions increase by 14.4%. While some reciprocal acts may be strategic, we show that not all are. 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 small, he does not reward kindness with kindness. Taken together, these results suggest fundraising within a social network can have a significant impact on increasing donations, but only if there is an opportunity for favors to be returned.

⁵ Isaac and Norton (forthcoming) 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.

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 $Stake_t$) 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:

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

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 should

contribute fully and non-Stakeholders should contribute nothing. Non-Stakeholder contributions are costly cooperation: they will be the focus of our analysis.

A selfish person would be 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. Others' past behavior can help a subject decide when to make or withhold non-Stakeholder contributions. Going back to our illustrative example, a subject (Joe) may interpret the contribution of another subject (Frank) when Joe is Stakeholder as a signal of Frank's kindness since Frank's contribution increased Joe's earnings. Joe may reciprocate by contributing when Frank is Stakeholder. If another subject (Mary) is less generous during Joe's Stakeholder stint, Joe may withhold contributions when Mary is Stakeholder.

These reciprocal acts can be of two varieties. On the one hand, they may be motivated by other-regarding preferences—that is, they may reflect direct reciprocity. In this case, a person's willingness to pay to increase another person's payoff may depend on the history he has observed. On the other hand, if a subject is strategic, observed history may change his beliefs about the best strategic action. If Frank seems to be potentially generous and Joe is strategic, Joe

may think that if he contributes when Frank 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 may distinguish other-regarding preferences from strategic giving.⁶

In either of these ways, information about contributions made by each group member and the timing of when each group member will be Stakeholder allows subjects to give reciprocally.⁷ If the 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 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

⁶ 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.

⁷ 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.

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.

The experiment is computerized and proceeds as follows. Subjects enter the lab and are given general instructions.⁸ They are told that they will make decisions in three sets of multiple rounds with three different groups, 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.

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

⁸ All instructions are available on the corresponding author's website.

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 Joe is subject A and Frank is B, Joe can see how much Frank contributed in Round 1 when Joe was the Stakeholder. Joe can reward Frank with a large contribution when Frank is Stakeholder in Round 2, or Joe may withhold that reward if he deems Frank'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 may be Rita, Joe and Frank's officemate 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 still 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 indicated in the screen header and in the contribution table as the “Ineligible” person. The Stakeholder position 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 will be able to see whether subjects are responsive to past generosity. That is, we can see whether they give larger contributions when the current Stakeholder is someone who was previously generous.⁹ This is precisely 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 (Rita) 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 would be evidence of indirect

⁹ In each round, subjects may respond to the past behavior of both the current Stakeholder and the current non-Stakeholders. However, if we detect this kind of responsiveness with regard to only the current Stakeholder’s past actions, this is sufficient to demonstrate reciprocal giving.

reciprocity. Also, the presence of a Bachelor shrinks the “circle of reciprocity” from five people to four people. We can 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 (ExCEN) 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). The protocol was double anonymous (subjects could not identify which subjects they were interacting with, and the experimenters could not identify which subject made any set of decisions). 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

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.^{10,11}

¹⁰ 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 after the Public treatment. We believe this is due to subject error: the Private treatment is more difficult to understand on its own. This is supported by the fact that Stakeholder contributions are lower in the Private treatment if the Private treatment is first. To check that 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 (see Reviewer’s Appendix). These results tell the same story as results reported in this paper. In sum, these analyses give us confidence that our results are not caused by treatment order or cross-treatment contamination.

¹¹ 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 7. 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, it is small, and at any rate should not affect our results: our treatment effect should be either unaffected or diminished by contagion, and our reciprocity result is within-subject-within-treatment so should be unaffected.

Figure 4 shows contribution data across rounds.¹² 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.¹³ Non-Stakeholder contributions (the solid lines) 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 contributions in the other treatments 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% 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. 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). These results give us confidence that our overall results are in line with previous work, despite the asymmetry of payoffs.

¹² The round numbers used in Figure 4 indicate the round number within that treatment. Since data are pooled across orders, the first round of a given treatment (e.g. Private) is also pooled across all orders. This means that round one is behavior in the first round, even though that might not be the subjects first round overall (e.g. if Private was the second or third treatment).

¹³ Contributions are not strictly 100%. This could be caused by subject error or myopic inequity aversion.

4.1 Treatment Effects

Recall that the main difference between the Private and Public treatments is that the Public treatment opens the door to targeted direct reciprocity, indirect reciprocity, and strategic self-interest. Thus, we can test whether these forces increase the 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.¹⁴

Figure 4 shows that average non-Stakeholder contributions in the Public treatment exceed those in the Private treatment in all rounds. These differences are only statistically significant in a few rounds, but are significant when pooled across rounds. As shown in Figure 5, 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 significant¹⁵ (paired Wilcoxon signed-rank test $p=0.051$).¹⁶ Thus, by making a small low-cost change in the information structure of the game, which allows for targeted reciprocity to manifest itself, increases cooperation by a significant amount. This was not a foregone conclusion, because the size of an increase in giving driven by reciprocal forces will depend on the level of initial cooperation and the distribution of types within the population. Enabling reciprocal giving may even decrease contributions if initial contributions are too low to draw reciprocation, while in other settings (with other parameters or a different population) reciprocal giving may provide an even larger boost than we observe here.

¹⁴ Results in this and following sections are mostly based on non-parametric tests. All results in this section hold under parametric specifications as well, including pooled and panel regressions with individual fixed effects controlling for group average contributions, round number, and a dummy indicating whether the subject's last Stakeholder stint had passed. (Results of selected parametric tests are in the Reviewers' Appendix.)

¹⁵ As noted in footnote 8, if the Private treatment is before the Public treatment then Private treatment non-Stakeholder contributions are inflated so that for this order, the Public-Private treatment difference is not significant.

¹⁶ Note that because our design is within subjects, we are using paired statistical tests to take into account the non-independence across observations when comparing the behavior of the same subjects across treatments (within subjects). Each test uses individuals' behavior averaged across rounds within a treatment, thus removing within treatment correlation and giving us 120 observations for each test. When we compare behavior across subjects within a treatment, we use non-paired tests.

How does this increase in contributions compare to the effect of other public goods institutions that manipulate information on the contributions of others? 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 there is no significant increase when only a photograph or (most relevantly) only contribution history is revealed. In our setting, information provision is rendered more powerful by taking advantage of heterogeneous preferences for public goods to enable targeted reciprocity. The increase in contributions in our setting may be less costly on net than other institutional changes, such as punishment, even though some of these may yield greater increases.¹⁷

We now turn to behavior in the Ineligible treatment. Figure 6 shows that Bachelors contribute significantly less (23.4% of endowment) than they did in the Private (29.9%) or Public (37.6%) treatments (paired Wilcoxon signed-rank test $p=0.043$ and $p=0.020$, respectively). It also shows that Bachelors give significantly less than non-Bachelor non-Stakeholders do in the Ineligible treatment (36.0% of endowment, Mann-Whitney test $p=0.007$). However, if Bachelors' contributions are compared to non-Bachelor non-Stakeholder contributions after the non-Bachelors' last Stakeholder stint (21.9% of endowment) the levels are not significantly

¹⁷ Costly punishment, another tool of interest in public good provision, has generally been found to increase giving by a large amount but often decreases efficiency overall because of the cost of punishment. An important paper in this area is Fehr and Gächter (2000), and Bochet et al. (2006) contains a useful discussion of the literature. Nonmonetary "punishment" has also been studied, as the use of social disapproval in Masclet et al. (2003) and Carpenter and Seki (2011). However, while social disapproval increased contributions by 37% in the former, its effects in the latter were mixed and contributions actually decreased for some populations.

different (Mann-Whitney test $p=0.794$). Note that non-Bachelor contributions decline across rounds while Bachelor contributions do not.

The low contributions of Bachelors, however, are not caused by idiosyncrasies of those subjects who were randomly assigned to be Bachelors. Those who were randomly assigned to be a Bachelor in the Ineligible treatment contribute no differently in the Private and Public treatments as compared to those who were not Bachelors.¹⁸ 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.

Additionally, for non-Bachelors, shrinking the “circle of reciprocity” to four people and adding a public good beneficiary who is outside that circle does not significantly change average contributions. There is no difference between non-Bachelor non-Stakeholders' contributions in the Ineligible treatment (36.0% of endowment) and their 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 7 shows the distribution of non-Stakeholder contributions pooled across rounds of all sessions. All treatments show a peak at zero tokens, a possible peak at 6-10 tokens,

¹⁸ 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, Bachelors and non-Bachelors do not give differently (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 $p=0.038$).

a dip between 10 and 20, and a peak at 20 tokens. The only statistically significant differences across treatments are the following. Bachelors give zero tokens more often in the Ineligible than in the Private (paired Wilcoxon signed-rank test $p=0.007$) and the Public ($p=0.014$) treatments. 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 will first examine reciprocal acts that could either be strategic or rooted in other-regarding preferences, and then we will isolate reciprocity rooted in the latter. To do this, we use within-subject tests of aggregate statistics. 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 this point, let's return to our example from the Introduction. Suppose Joe is Stakeholder in rounds 1 and 6, and Frank is Stakeholder in rounds 2 and 7. In round 2, Joe will remember how generous Frank was in round 1. "Stakeholder past generosity" will be Frank's contribution in round 1. In round 7, when Frank is Stakeholder, our measure of "Stakeholder past generosity" for Joe would be the average of Frank's contributions in rounds 1 and 6 when Joe was Stakeholder.

We look at the effects of generosity using nonparametric tests. To do so, we define a “generosity threshold” such that contributions greater than this amount are called generous. For each subject, we calculate his average contribution when facing a Stakeholder whose past generosity meets this threshold and his average contribution across all rounds when facing a Stakeholder whose past generosity does not. We tried many thresholds, including 6, 8, 10, 12, 15, and 19 tokens and the group’s cumulative average contribution, with robust results. We report results from a threshold of 10 tokens (50% of endowment). To clarify further how this works, suppose that a group contains only Joe, Frank, and Mary. Frank gave 15 tokens every time Joe was Stakeholder, and Mary always gave 2 tokens. Joe’s average contribution to a generous Stakeholder is his average contribution when Frank was Stakeholder, and his average contribution to an ungenerous Stakeholder is his average contribution when Mary was Stakeholder.¹⁹

A subject displays reciprocal behavior if he gives more when facing a previously-generous Stakeholder than when facing a previously-ungenerous Stakeholder. Figure 8 presents averages of these measures. In all treatments in which it is possible to attribute previous generous and ungenerous 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 41.92% as compared to 27.09% of endowment, paired Wilcoxon signed-rank test $p=0.000$; Ineligible treatment 41.22% as compared to 23.92% 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

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.

treatment, 27.57% was given to previously-generous Stakeholders as compared to 29.16% to previously-ungenerous Stakeholders, paired Wilcoxon signed-rank test $p=0.773$).²⁰ We would not expect rewards to generous behavior in the Private treatment because subjects cannot tell who is Stakeholder or what the current Stakeholder did in 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.²¹

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, between treatments. 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 ($p=0.966$). This and the previous test 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 Joe was Stakeholder in rounds 1 and 6 and Frank was Stakeholder in rounds 2 and 7. Assume no further fundraising rounds follow. If Frank was kind to Joe in rounds 1 and 6, will Joe reciprocate in round 7? If Joe is purely strategic, he has little to gain, so he should not contribute and therefore not reciprocate.

²⁰ We find similar results in panel regressions of non-Stakeholder contribution regressed on the current Stakeholder's past contributions to the subject, group average past contributions, round, and a dummy for whether the subject is past his last opportunity to be Stakeholder.

²¹ Recall that 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 test of responsiveness to Stakeholder history.

This is also the logic that was used in one test of reciprocity rooted in other-regarding preferences in Cabral et al. (2012).

We test for non-strategic reciprocity by constructing statistics of each subject's average contribution to previously generous and ungenerous Stakeholders in rounds after this subject's last Stakeholder stint. These results, in Figure 9, show that subjects are not simply 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 after they have no strategic motive to do so.²² This difference is statistically significant (paired Wilcoxon signed-rank test $p=0.026$). In the Ineligible treatment, the sample size is reduced because only non-Bachelors are considered and the treatment has fewer rounds. Because of this reduced power, discrimination in late rounds is not statistically significant, although the difference is in the same direction.

4.3 Indirect Reciprocity

We now examine evidence of indirect reciprocity. We define this type of 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 receive the benefits of reciprocity. Therefore, direct reciprocity and strategic self-

²² Again, similar results hold in panel regressions, only using rounds after which the subject is no longer Stakeholder.

interest cannot drive Bachelor giving. Bachelors can't even receive great benefits from conditional cooperation.

We have shown that Bachelors give significantly less than non-Stakeholders in the Ineligible treatment and less than they themselves gave in other treatments. Bachelors also behave reciprocally when they are part of the circle of reciprocity, in that they have positive responsiveness in the Public treatment (giving 35.02% of endowment to previously generous Stakeholders and 22.28% of endowment to previously ungenerous Stakeholders, significantly different with paired Wilcoxon signed-rank $p=0.004$). However, Bachelors in the Ineligible treatment, when they are no longer part of the circle of reciprocity, 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.²³ 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.

This result is intriguing because other studies have found evidence of indirect reciprocity (e.g., Engelmann and Fischbacher, 2009; Seinen and Schram, 2006). The re-matching structure of those experiments, however, allows subjects to have a financial interest in the group's overall cooperation. Subjects are not wholly disinterested as our Bachelors are. 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*

²³ 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.

Favor trading is a natural element of social networks. Grassroots fundraising and other institutions can achieve pro-social ends by taking advantage of heterogeneous preferences and the desire to be nice to those who have been nice to you. In an experiment that allows different forms of reciprocity to be turned on and off, we explore the power of favor trading and the mechanisms through which institutions such as grassroots fundraising may work. In our setting, favor trading increases cooperation by 14.4%, an amount that is both statistically and economically significant. This result demonstrates that groups like Olson's (1965) "privileged groups," which contain members who have a dominant strategy of contributing, can be structured to further pro-social ends.

Our results also provide a new window into reciprocity. We find clear evidence of direct reciprocal giving that is non-strategic.²⁴ While some studies have found direct reciprocity in laboratory institutions like the investment game (e.g., Berg et al., 1995), others have criticized the abstraction of those institutions (e.g., List, 2006). The presence of direct reciprocity in our slightly richer setting may be a good step in understanding how institutions outside of the lab can take advantage of direct reciprocity 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.

Intriguingly, we find no evidence of indirect reciprocity in our Ineligible treatment. While most experiments that examine indirect reciprocity allow actors to directly benefit from an increased tendency to cooperate within the group, we strip our Bachelor of strong incentives to foster cooperative norms. When given the costly chance to reward a kind act without the

²⁴ 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.

possibility of future direct reciprocation, our Bachelors do not provide 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 have to occur through a change in contribution behavior.

Reciprocal behaviors, primarily driven by direct reciprocity rooted in other-regarding preferences, have the power to increase pro-social behavior and therefore efficiency. Favor trading can be harnessed in very low-cost ways through simple information provision. This is an informal mechanism we are all familiar with from our own lives. Charitable organizations also seem to recognize the power of reciprocal giving in social networks as they develop vehicles like Facebook Cause promotion and custom donor-created websites like DonorPages.

Will our finding of increased efficiency through reciprocal giving translate to non-lab social environments? Or will people tend to either contribute only to their own pet causes or to offset increased contributions to a friend's charity by reduced contributions to their own charities? The linear nature of the public goods in our experiment, while important in presenting a simple choice to subjects, clouds this generalization. Charities' interest in social fundraising does not necessarily imply that socially-elicited contributions to one organization do not crowd out contributions to other organizations. Field studies could provide an answer to this question, but it seems likely that most theories of other-regarding preferences would imply that if there is such crowd-out, it is incomplete.

Reviewers' Appendix: Parametric Tests of Reciprocity

In this Appendix, we use regression techniques to identify reciprocal behavior resulting from direct reciprocity and strategic self-interest. These results may be biased because of the endogeneity inherent in group dynamic behavior. We take care to limit the influence of this bias, but it is to some extent unavoidable.

We perform a panel regression for one treatment at a time, with fixed effects and errors clustered by group. 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 : here, the current Stakeholder's cumulative average contributions when subject i was Stakeholder) and \mathbf{X}_{it} (other variables):

$$g_{it} = a + bh_{ikt} + \mathbf{CX}_{it} + \varepsilon_{it}$$

If direct reciprocity or strategic self-interest is important, b (the coefficient on h_{ikt}) should be positive in both the Public and Ineligible treatments. In the Private treatment, b should be zero, because in each round, no-one knows who the current Stakeholder is or what that person has done in the past. 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} . Our control for group generosity is 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. 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).

The panel regression also includes in \mathbf{X}_{it} the current round number and an indicator for whether this subject has passed his last Stakeholder stint. If the coefficient on the round number is negative, cooperation decreases from round to round. The coefficient on the post-last-Stakeholder stint dummy can be interpreted as the importance of strategic giving. If a subject gives strategically in the rounds before his last Stakeholder stint, there should be a discontinuity in contributions that should be reflected in a large, negative coefficient on this dummy.

Results are shown in Table A-1. Directly reciprocal behavior is supported for the Public and Ineligible treatments. The coefficient on the post-last-Stakeholder stint dummy in the Public and Ineligible treatments is insignificant. This implies that strategic motives are not important. The same results obtain in an AR1 specification (results available upon request), except that in the Private treatment the post-last-Stakeholder stint dummy is no longer significant. In these results, as in the following Appendix results, the use of subject fixed effects effectively controls for order effects. Separate specifications without fixed effects and with order dummies instead (available upon request) tell the same story.

We can further test for the importance of other-regarding preferences by restricting our attention to the rounds after a subject has passed his final Stakeholder stint. See results in Table A-2. Due to the reduced population size, particularly for the Ineligible treatment, the power of the test is significantly reduced and the Stakeholder past contribution coefficient is not statistically significant for the Ineligible treatment. However, the coefficient on Stakeholder past contribution is positive for the Public treatment. This implies that reciprocal behavior is at least partly rooted in other-regarding preferences.

Finally, we build error clustering by group and fixed effects into a test using aggregate statistics, since this should eliminate the bias of our panel analysis while checking for robustness to intra-class correlation. We create two stacked observations per subject: one to represent the subject's average contribution in rounds in which the current Stakeholder was previously generous (using the 10 token, or 50%-of-endowment, threshold; similar results obtain for other thresholds) toward this subject, and one to represent his average contribution in rounds when the Stakeholder was ungenerous toward this subject. A "nice dummy" differentiates between the two observations for each subject. We perform a panel regression using these two observations per subject, in which we regress contributions on the "nice dummy." The same results obtain if we use group dummies, individual random effects, or individual fixed effects. In Table A-3 we show results of the individual fixed effects regression without group dummies. The "nice" dummy is significant and positive for the Public and Ineligible treatments, but not for the Private treatment. The result that reciprocal behavior continues after the last Stakeholder stint also persists in this specification.

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., 2011, "Who Cooperates in Repeated Games?" 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 Gächter, 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, Forthcoming. "Endogenous institutions and the possibility of reverse crowding out." *Public Choice*.

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

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.

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.
- 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

You are now: MAKING YOUR CONTRIBUTION FOR ROUND 3

Your Letter Code: D Stakeholder: C

CONTRIBUTIONS TO THE GROUP FUND

	<u>A</u>	<u>B</u>	<u>Stakeholder</u>	<u>YOU</u>	<u>E</u>	<u>TOTAL TOKENS IN GROUP FUND</u>	<u>MY EARNINGS</u>
Round 1	<u>10</u>	6	2	14	18	50	\$0.62
Round 2	19	<u>11</u>	3	7	15	55	\$0.81
Round 3			*****				
Round 4				*****			
Round 5					*****		
Round 6	*****						
Round 7		*****					
Round 8			*****				
Round 9				*****			
Round 10					*****		

***** indicates that this person was/will be Stakeholder in the round indicated

DECISION PANEL

How much would you like to put in the **GROUP FUND**? (0-20)

Your **PERSONAL FUND** contribution will be 20 minus your **GROUP FUND** contribution.

RETURNS:

Personal fund: **\$0.02 per token** to you

Group fund: **\$0.03 per token** to Stakeholder

\$0.01 per token to non-Stakeholders (including YOU)

CLICK TO SUBMIT

SUBMIT

Figure 1. Public Treatment Decision Screen

You are now: **REVIEWING RESULTS FOR ROUND 2**

Stakeholder: Someone Else

CONTRIBUTIONS TO THE GROUP FUND

	<u>YOU</u>	<u>OTHERS</u> (Random Order)	<u>TOTAL TOKENS IN GROUP FUND</u>	<u>MY EARNINGS</u>
Round 1	<u>9</u>	1, 5, 17, 13	45	\$1.57
Round 2	12	16, 4, 20, 8	60	\$0.76
Round 3				
Round 4				
Round 5				
Round 6				
Round 7				
Round 8				
Round 9	*****			
Round 10				

***** indicates that you were/will be Stakeholder in the round indicated

DECISION PANEL

REVIEW RESULTS FROM ROUND 2 IN THE TABLE ABOVE.

YOUR EARNINGS WERE:

Your Personal Fund contribution (8) times \$0.02 per token
= \$0.16

PLUS: The total number of tokens in the Group Fund (60) times \$0.01 per token
= \$0.60

EQUALS:
TOTAL = \$0.76

CLICK WHEN DONE

DONE

Figure 2. Private Treatment Review Screen

You are now: **REVIEWING RESULTS FOR ROUND 2**

Your Letter Code: C Stakeholder: B Ineligible: D

CONTRIBUTIONS TO THE GROUP FUND							
		<u>Stakeholder</u>	<u>YOU</u>	<u>Ineligible</u>			
	<u>A</u>	<u>B</u>	<u>C</u>	<u>D</u>	<u>E</u>	<u>TOTAL TOKENS IN GROUP FUND</u>	<u>MY EARNINGS</u>
Round 1	<u>1</u>	4	2	16	8	31	\$0.67
Round 2	20	<u>0</u>	17	9	3	49	\$0.55
Round 3			<u>*****</u>				
Round 4					<u>*****</u>		
Round 5	<u>*****</u>						
Round 6		<u>*****</u>					
Round 7			<u>*****</u>				
Round 8					<u>*****</u>		

***** indicates that this person was/will be Stakeholder in the round indicated

DECISION PANEL

REVIEW RESULTS FROM ROUND 2 IN THE TABLE ABOVE.

YOUR EARNINGS WERE:

Your Personal Fund contribution (3) times \$0.02 per token = \$0.06

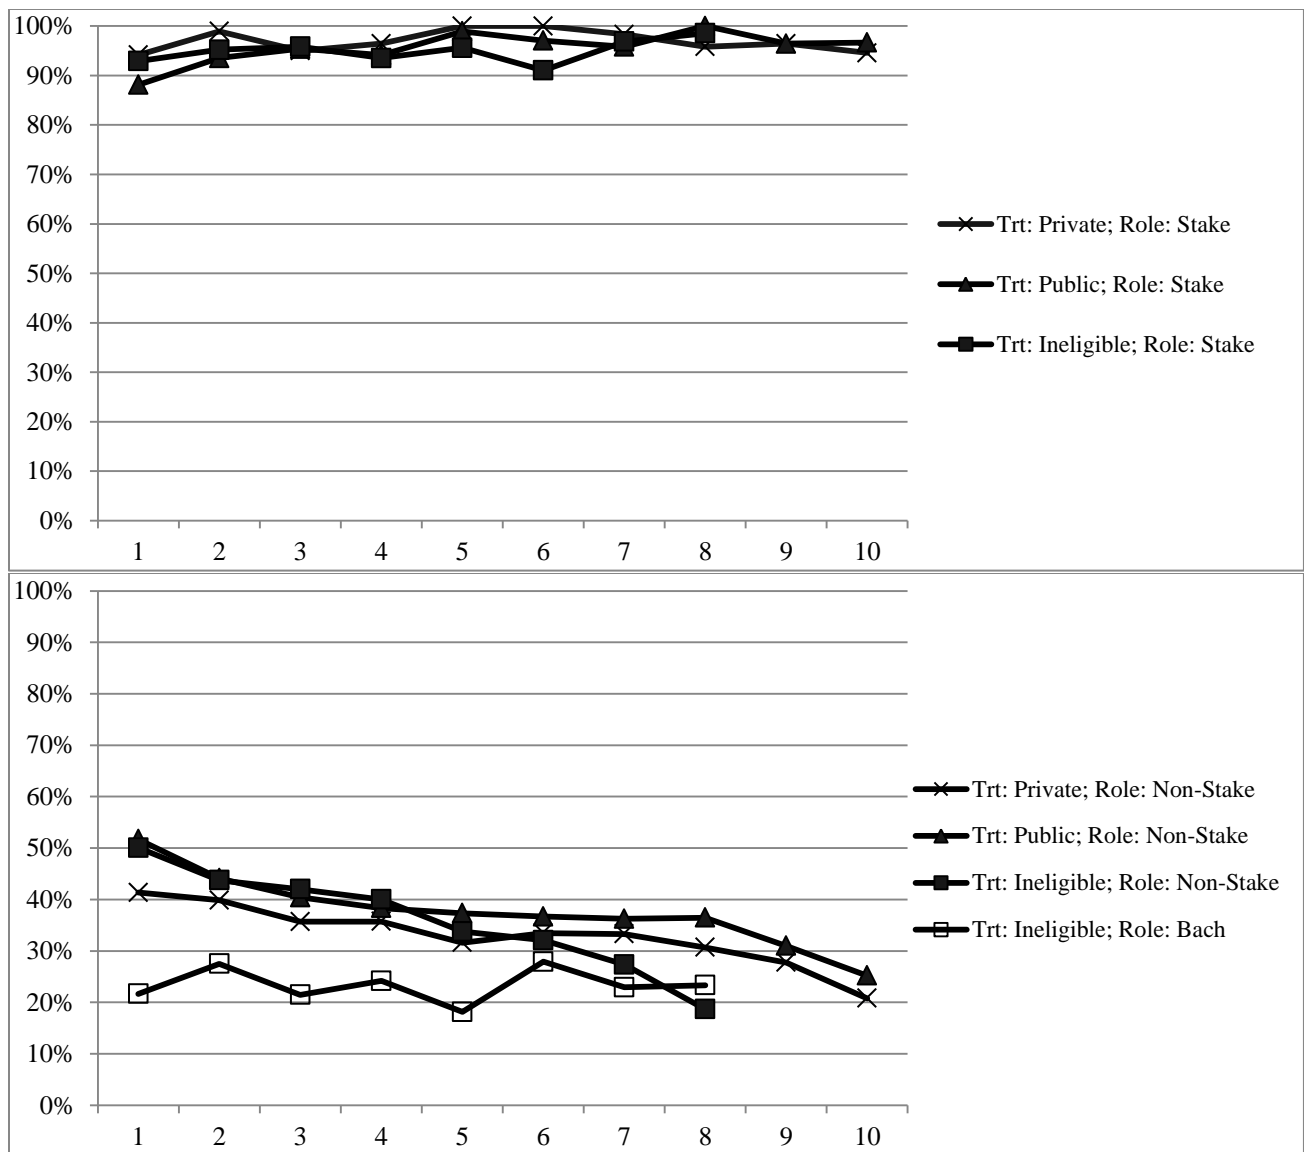
PLUS: The total number of tokens in the Group Fund (49) times \$0.01 per token = \$0.49

EQUALS: **TOTAL = \$0.55**

CLICK WHEN DONE

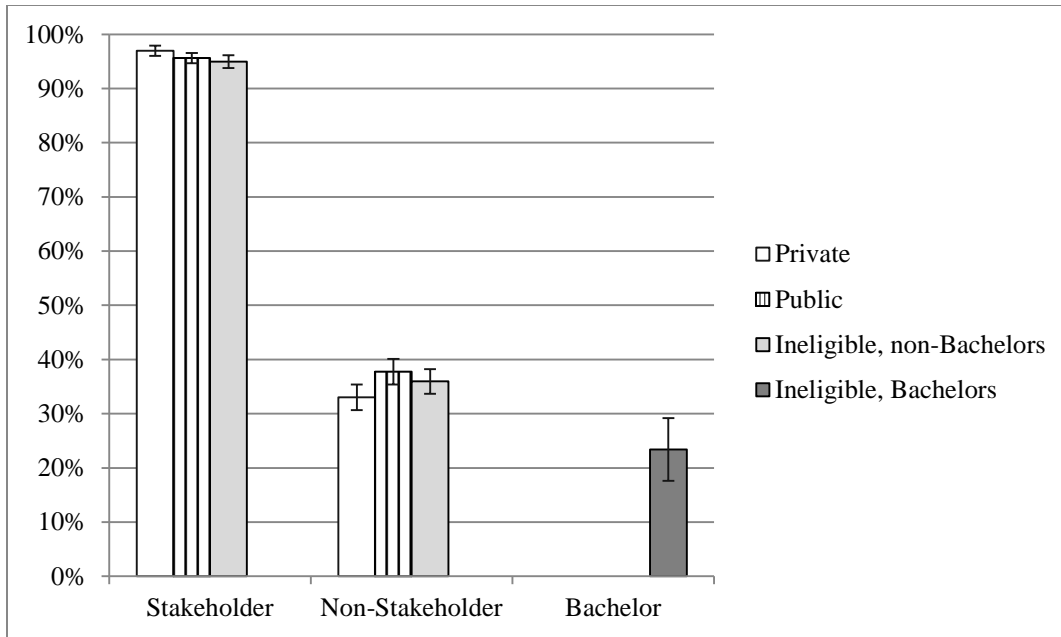
DONE

Figure 3. Ineligible Treatment Review Screen



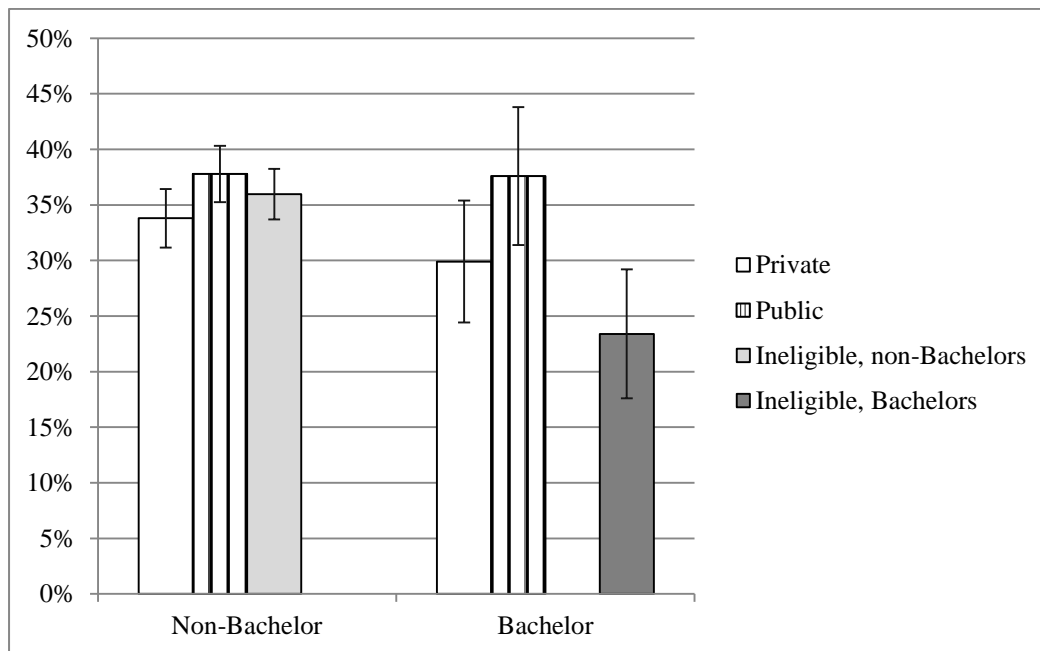
Note: Round numbers indicate round number within a treatment; data pooled across treatment orders

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



Note: Error bars indicate standard error of mean, using each subject mean as an observation. N=120 for Private and Public treatments; N=96 for Ineligible non-Bachelors; N=24 for Ineligible Bachelors.

Figure 5. Average Contributions by Treatment and Role (Percent of Endowment)



Note: Error bars indicate standard error of mean, using each subject mean as an observation. N=96 for non-Bachelors; N=24 for Bachelors.

Figure 6. Average Non-Stakeholder Contributions by Treatment and Bachelor Status (Percent of Endowment)

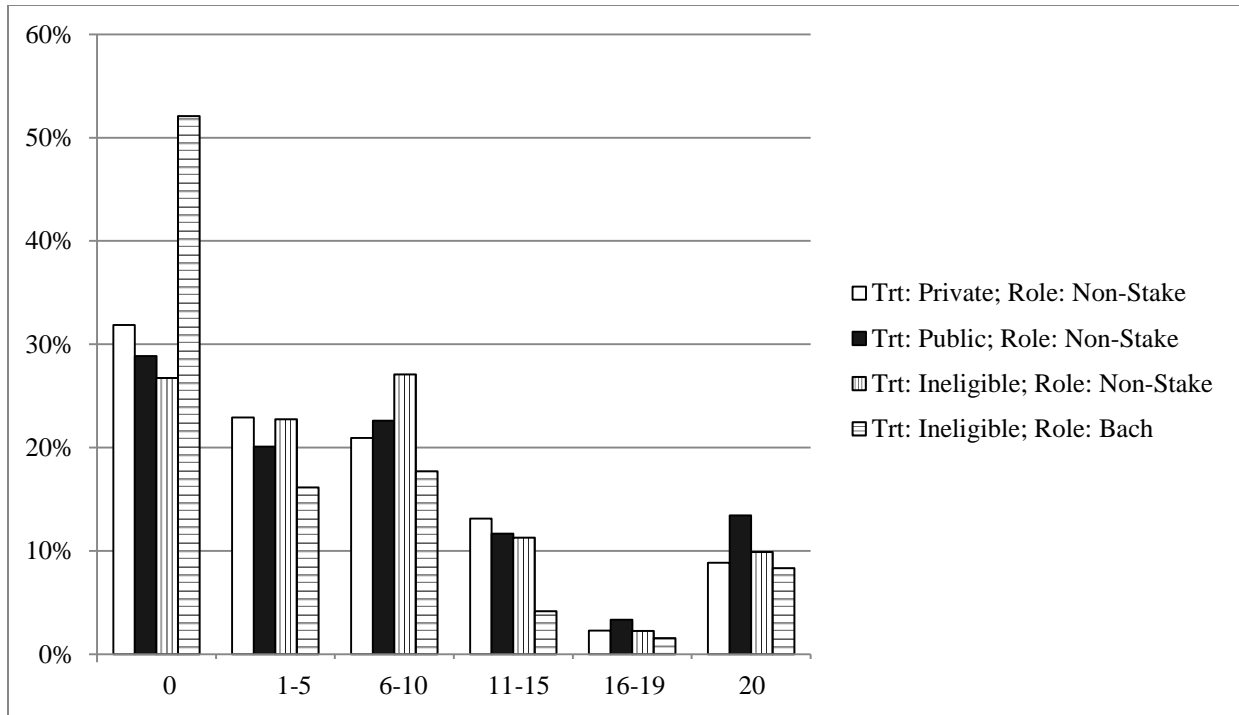
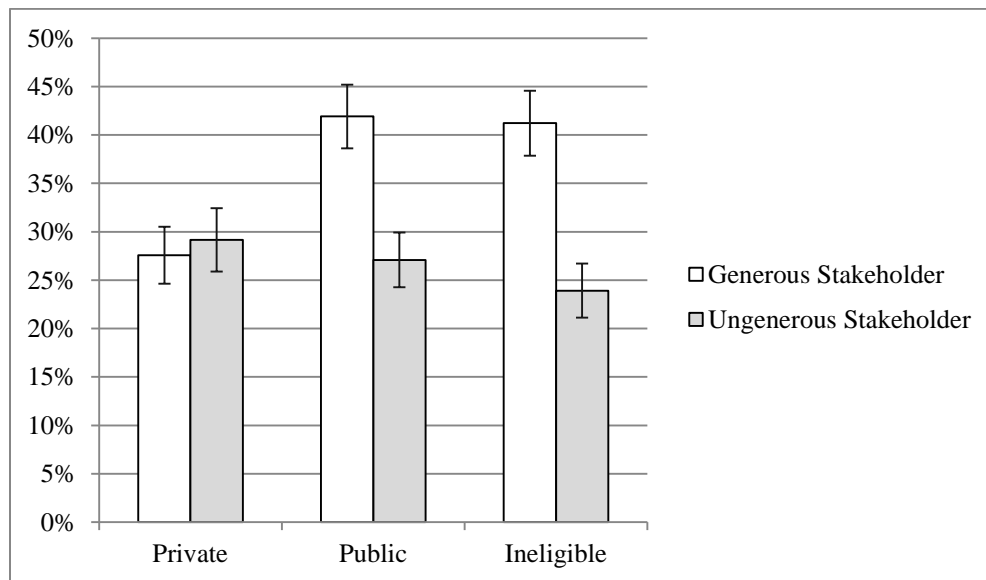
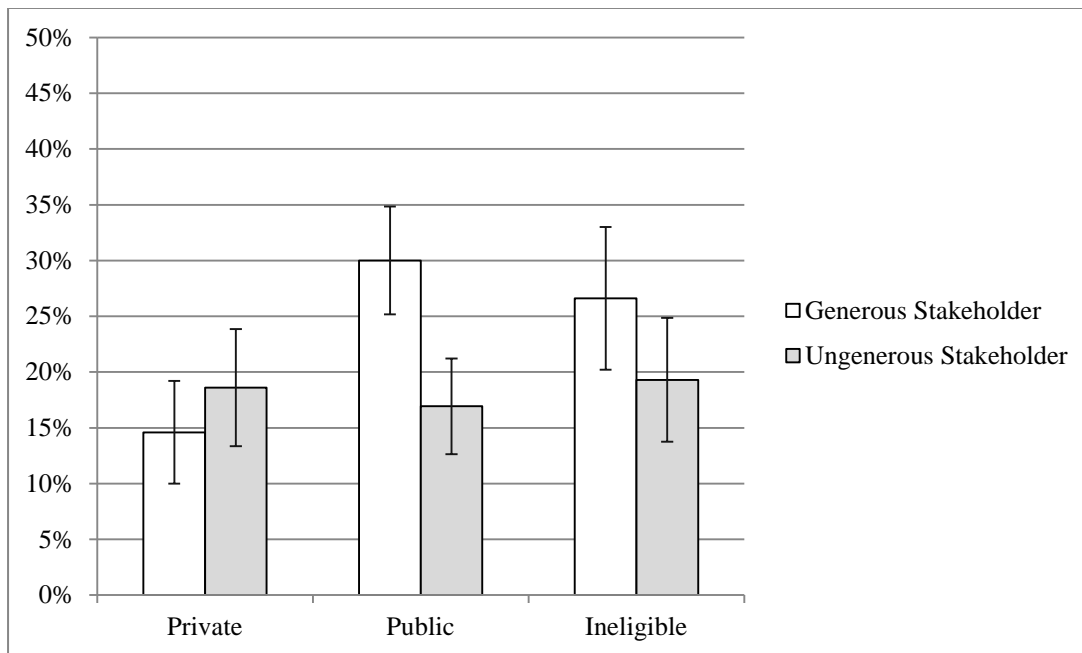


Figure 7. Distribution of Non-Stakeholder Contribution Amounts, Pooled across Rounds (in Tokens)



Note: Error bars indicate standard error of mean, using each subject mean as an observation. N 's are less than 120 because some subjects did not face both a generous and an ungenerous stakeholder. For this reason, Private $N=82$ (of 120), Public $N=95$ (of 120), Ineligible $N=75$ (of 96).

Figure 8. Evidence of Reciprocal Giving
Average Non-Stakeholder Contributions across All Rounds by Stakeholder's Past Generosity toward Subject (in Percent of Endowment)



Note: Error bars indicate standard error of mean, using each subject mean as an observation. *N*'s are less than 120 because 48 subjects had to be dropped from each treatment because there were less than two rounds remaining after their last Stakeholder stint; additionally, more subjects had to be dropped from each treatment if they did not face both a generous and an ungenerous Stakeholder after their last Stakeholder stint. Private *N*=31 (of 120), Public *N*=41 (of 120), Ineligible *N*=25 (of 96).

Figure 9. Evidence of Direct Reciprocity
Average Non-Stakeholder Contributions by Stakeholder's Past Generosity toward Subject, after Last Stakeholder Stint (in Percent of Endowment)

Table A-1. OLS Fixed Effects Panel Regression of Non-Stakeholder Contribution (in Percent of Endowment) on Period-Level Covariates

	Private Treatment	Public Treatment	Ineligible Treatment (excl. Bachelors)
Stakeholder average past contributions to me	0.05 (0.06)	0.24*** (0.05)	0.23** (0.09)
Group average contributions	0.25 (0.16)	0.05 (0.17)	-0.08 (0.27)
Round number	0.03 (0.65)	-1.69 (1.22)	-2.86** (1.34)
Post-last Stakeholder Stint? (dummy)	-9.66** (3.86)	-1.66 (5.05)	-4.85 (5.40)
Constant	22.15*** (7.48)	34.65** (13.68)	42.50** (16.18)
Observations (rounds)	720	720	432
Number of subjects	120	120	96
F	4.50	13.37	31.47
R ² (overall)	0.078	0.182	0.156

Robust standard errors in parentheses; errors are clustered on groups; individual fixed effects

* significant at 10%; ** significant at 5%; *** significant at 1%

Table A-2. OLS Fixed Effects Panel Regression of Non-Stakeholder Contribution (in Percent of Endowment) on Period-Level Covariates, Post-Last Stakeholder Stint

	Private Treatment	Public Treatment	Ineligible Treatment (excl. Bachelors)
Stakeholder average past contributions to me	-0.19 (0.11)	0.21* (0.11)	0.17 (0.22)
Group average contributions	0.39 (0.44)	-0.09 (0.32)	0.30 (0.60)
Round number	-4.12** (1.83)	-5.90*** (1.96)	-4.89 (2.26)
Constant	48.39** (22.93)	76.14*** (25.80)	42.89 (25.08)
Observations (rounds)	240	240	144
Number of subjects	96	96	72
F	2.03	3.70	1.78
R ² (overall)	0.036	0.106	0.083

Robust standard errors in parentheses; errors are clustered on groups; individual fixed effects

* significant at 10%; ** significant at 5%; *** significant at 1%

Table A-3. 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 (2.72)	14.83*** (2.68)	17.30*** (2.72)
Constant	29.16*** (1.36)	27.09*** (1.34)	23.92*** (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 groups; 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%