

Leider, Stephen; Möbius, Markus M.; Rosenblat, Tanya; Do, Quoc-Anh

Working Paper

How much is a friend worth? Directed altruism and enforced reciprocity in social networks

Working Papers, No. 07-11

Provided in Cooperation with:

Federal Reserve Bank of Boston

Suggested Citation: Leider, Stephen; Möbius, Markus M.; Rosenblat, Tanya; Do, Quoc-Anh (2007) : How much is a friend worth? Directed altruism and enforced reciprocity in social networks, Working Papers, No. 07-11, Federal Reserve Bank of Boston, Boston, MA

This Version is available at:

<https://hdl.handle.net/10419/55563>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

How Much is a Friend Worth? Directed Altruism and Enforced Reciprocity in Social Networks

Stephen Leider, Markus M. Möbius, Tanya Rosenblat, and Quoc-Anh Do

Abstract:

We conduct field experiments in a large real-world social network to examine why decision-makers treat their friends more generously than strangers. Subjects are asked to divide a surplus between themselves and named partners at varying social distances, but only one of these decisions is implemented. We decompose altruistic preferences into baseline altruism towards strangers, and directed altruism towards friends. In order to separate the motives that are altruistic from the ones that anticipate a future interaction or repayment, we implement an *anonymous* treatment in which neither player is told at the end of the experiment which decision was selected for payment, and a *non-anonymous* treatment where both players are told the outcome. Moreover, in order to distinguish between different future interaction channels—including *signaling* one's propensity to be generous and *enforced reciprocity*, where the decision-maker grants the partner a favor because she expects it to be repaid in the future—the experiments include games where transfers both increase and decrease social surplus. We find that decision-makers vary widely in their baseline altruism, but pass at least 50 percent more surplus to friends as opposed to strangers when decision-making is anonymous. Under non-anonymity, transfers to friends increase by an extra 24 percent relative to strangers, but only in games where transfers *increase* social surplus. This effect increases with the density of the social network structure between both players. Our findings are well explained by enforced reciprocity, but not by signaling or preference-based reciprocity. We also find that partners' expectations are well attuned to directed altruism, but that they completely ignore the decision-makers' baseline altruism. Partners with higher baseline altruism have friends with higher baseline altruism and, therefore, are treated better by their friends.

JEL Classifications: C73, C91, D64

Stephen Leider and Quoc-Anh Do are Ph.D. candidates in economics at Harvard University. Their email addresses are sleider@hbs.edu and do@fas.harvard.edu, respectively. Markus Möbius is an associate professor of economics at Harvard University and a Faculty Research Fellow at the National Bureau of Economic Research. His email is mobius@fas.harvard.edu. Tanya Rosenblat is an assistant professor of economics at Wesleyan University and a visiting scholar at the Federal Reserve Bank of Boston. Her email address is trozenblat@wesleyan.edu.

This paper replaces an earlier working paper version titled "Social Capital in Social Networks." We are grateful to Ed Glaeser, Muriel Niederle, Adam Szeidl, and seminar participants at the 2004 Econometric Society Summer meetings, 2004 SITE, Max Planck Institute Ringberg Workshop on Social Networks, 2006 North American ESA meetings, Boston University, New York University, University of Connecticut, Texas A&M, Columbia University, Institute for Advanced Study, and University of Michigan for helpful comments. We are particularly indebted to Rachel Croson and Al Roth for extremely helpful conversations during the design stage of our experiments. Paul Niehaus and Raphael Schoenle were excellent research assistants. We thank Elizabeth Murry and Tyler Williams for very careful proofreading and numerous suggestions. Rosenblat thanks the Federal Reserve Bank of Boston for its hospitality. We are grateful to the Social Security Administration and the Center for Retirement Research at Boston College for financial support through a Sandell grant (Möbius, Rosenblat) and to the NSF (Leider).

This paper, which may be revised, is available on the web site of the Federal Reserve Bank of Boston at <http://www.bos.frb.org/economic/wp/index.htm>.

The views expressed in this paper are solely those of the authors and are not those of the Federal Reserve System, the Federal Reserve Bank of Boston.

This version: October 1, 2007

“Friendship is one mind in two bodies.” – Mencius (3rd century B.C.)

“Friendship is essentially a partnership.” – Aristotle (4th century B.C.)

1 Introduction

Our friends tend to treat us better than strangers do. In business settings, a significant share of team production therefore evolves around pre-existing social relationships. For example, 53 percent of all start-up firms in the United States have more than one owner, and 23 percent of these jointly owned start-ups have at least one pair of friends as co-owners (Davidsson and Reynolds 2005). Our objective in this paper is to answer the following questions: how much is it worth to have a friend rather than a complete stranger make payoff-relevant decisions, and what channels make a friend valuable?

While most research in experimental economics has focused on altruism between strangers, it is natural to assume that altruism’s strength varies with social distance. We will use the terms *baseline altruism* to denote a decision-maker’s altruism towards strangers and *directed altruism* to describe a decision-maker’s altruistic preferences towards socially close partners (usually friends). Increased altruistic motivations align the interests of a decision-maker and her partner in an economic relationship. Aside from directed altruism, friendships can be valuable because there is a greater prospect of future interaction with friends as compared to strangers. Economic theory suggests at least three mechanisms which induce the decision-maker to treat the partner more generously when there is a prospect of future interaction. First, the decision-maker can grant favors because she expects the partner to repay these in the future. We call this the *enforced reciprocity* mechanism. Second, the possibility of future interaction gives incentives for the decision-maker to *signal* her altruistic type to the partner (Benabou and Tirole 2006). Third, psychological game theory has modeled *preference-based reciprocity* where decision-makers behave generously because they expect the partner to behave kindly towards them in some future interaction, and because they derive utility from rewarding kind behavior (Rabin 1993, Dufwenberg and Kirchsteiger 2004).

We use field experiments to (a) separate the directed altruism from the future interaction channel, (b) determine the absolute and relative strength of both effects and (c) to isolate the mechanism behind the repeated interaction channel. Our experimental design has two stages. In the first stage we map a large real-world social network. In the second stage we invite a randomly selected sample of subjects from this network to play the role of a decision-maker who over a period of one to two weeks repeatedly divides surplus between herself and known partners (identified by the partner’s real first and last name) located at various social distances, plus a

“nameless” partner. In our first field experiment, decision-makers play modified dictator games with different exchange rates, such that in some games passing tokens to the partner increases social surplus and in other games it decreases social surplus. The second field experiment is a novel helping game where decision-makers have to choose how maximally willing they would be to pay to increase the partner’s payoff by a fixed amount. While each decision-maker makes multiple decisions, only one of these is randomly selected for payment at the end of the experiment. A decision-maker makes both an *anonymous* and a *non-anonymous* choice for each match. In the anonymous treatment, no player is told at the end of the experiment which of the matches was selected for payment, while in the non-anonymous treatment both players are informed about the selected decision-maker/partner match and the decision-maker’s action. The anonymous and non-anonymous treatments allow us to separate directed altruism from future interaction channels. By including games in which giving is both efficient and inefficient, we can distinguish the signaling mechanism from reciprocity mechanisms: if giving is efficient, decision-makers motivated by reciprocity should only transfer more surplus under non-anonymity to the socially close partner. In contrast, decision-makers who want to signal their generosity to the partner should always transfer more surplus, regardless if the setting is anonymous or not.

In the anonymous treatment, we find that a decision-maker’s action taken for a named partner is strongly determined by (a) her baseline altruism towards a nameless partner and (b) her directed altruism towards socially close partners. In the dictator game experiment, each token increase in generosity towards a nameless player is associated with approximately a one token increase in generosity towards any named player. This correlation is remarkable because decisions for named and nameless partners were taken one week apart. Controlling for their baseline altruism, decision-makers are substantially more generous towards friends: subjects pass at least 50 percent more surplus to friends as compared to strangers, when decision making is anonymous. This effect declines for indirect friends but can still be substantial for second-order friends. We do not find that standard demographic variables (such as the decision-maker’s or the partner’s gender) have significant effects. We also collect data on how generous partners expect(named) decision-makers at various social distances to behave. We find that partners’ expectations are well calibrated to directed altruism: their expectations decrease with greater social distance at approximately the same rate as actual giving. However, there is no evidence that a partner takes a decision-maker’s baseline altruism towards nameless players into account when forming expectations. This is true even if the partner and the decision-maker are socially close. Nevertheless, partners with high baseline altruism are treated substantially better by their friends than subjects with low baseline altruism because baseline altruism

towards nameless partners is positively correlated with altruism towards direct social neighbors. Therefore, it appears that while we seek out friends with similar altruistic preferences as ourselves, we base our beliefs about the decision-maker’s generosity mainly on social proximity or distance.

When we compare non-anonymous and anonymous decision making, we find that transfers to a friend versus a stranger are an *additional 24 percent higher only if* such transfers increase social surplus. This is true, for example, in a dictator game where tokens are worth more to the partner than to the decision-maker. Thus, an important outcome that distinguishes directed altruism from the effects of observability is that directed altruism leads to more equitable distributions of payoffs (that may diminish efficiency), while the partner’s ability to observe the decision-maker’s action pushes allocations towards social efficiency.

Just as in the case of directed altruism, the non-anonymity effect declines with social distance and is not affected by the demographic characteristics of the decision-maker or the partner. We also find that the effects of non-anonymity and directed altruism towards friends are substitutes: the more altruistic the decision-maker, the *lower* are the extra transfers under non-anonymity. Our experimental findings can be explained by a theory of enforced reciprocity, based on Möbius and Szeidl (2007) where, under the condition of non-anonymity, the decision-maker treats friends to extra “favors” because she expects these to be repaid in the future. In contrast, the signaling model of Benabou and Tirole (2006) predicts excess transfers to friends across all games, even when such transfers decrease social surplus. We cannot use the same test to distinguish the favors mechanism from preference-based reciprocity models because they make similar predictions. However, our enforced-reciprocity model also predicts an observed role for network structure: a measure of local network density – *maximum network flow* between decision-maker and partner – determines the decision-maker’s ability to make the partner repay a favor. Consistent with this model, we find that network flow predicts the decision-maker’s generosity under non-anonymity even after controlling for social distance. One might expect that in locally dense networks friends see each other more frequently and have more opportunities for reciprocation. As it turns out, our finding is not just an artifact of network flow being a proxy for the amount of time a decision-maker spends with a partner. For direct friends, we also collect data on the average amount of time spent together each week. For our data set, network flow is essentially uncorrelated with time spent together and it is network flow rather than time spent together that predicts generosity under non-anonymity. This finding suggests that our results are driven by enforced reciprocity rather than preference-based reciprocity.

Our paper relates to a rich experimental and theoretical literature on other-regarding pref-

erences and cooperation. Prosocial behavior of varying magnitudes has been observed in a variety of laboratory contexts (see Camerer (2003) for an extensive survey). Our directed altruism channel is a natural refinement of preference-based altruism as modeled by Andreoni (1990) in his “warm glow” model, or by Fehr and Schmidt (1999), Bolton and Ockenfels (2000), and Charness and Rabin (2002), all of whom focus on preferences over payoff distributions. In the lab, experiments on prosocial behavior with reduced anonymity have typically involved revealing the partner’s ethnic group or gender (for example, see Fershtman and Gneezy (2001)). In subsequent research, Goeree, McConnell, Mitchell, Tromp, and Yariv (2007) have adopted the anonymous treatment of our experimental design (using a standard dictator game) and also find strong evidence for directed altruism in a school network of teenage girls (also see Brañas-Garza, Cobo-Reyes, Paz Espinosa, Jiménez, and Ponti (2006) for experimental data with European university students). There are very few experimental papers that explicitly rely on the subjects’ ongoing relationships with their friends. A notable paper using this methodology is the seminal work of Glaeser, Laibson, Scheinkman, and Soutter (2000), who also explicitly match subjects at various social distances in a trust game. However, to the best of our knowledge, our design is the first within-subjects design which attempts to distinguish between directed altruism and future interaction effects.

In an important methodological advance, our two experiments were completely web-based. This ensured very high participation rates between 42 percent and 71 percent, which were crucial for generating a sufficient number of matches between direct friends during the course of the experiment. Our experiment also pioneers various techniques to conduct social network surveys online; we have already applied these techniques to a number of other projects and we hope that they will prove useful to other researchers.

The rest of the paper is organized as follows. Section 2 provides a simple theoretical framework and reviews the testable implications of various repeated game mechanisms such as enforced reciprocity, signaling, and preference-based reciprocity. The experimental design is described in section 3. In section 4 we summarize the main features of the data. Our empirical results on directed altruism are presented in section 5, and section 6 contains our comparison of decision making under anonymity and non-anonymity. Section 7 concludes by outlining how results presented in this paper can help build a broader theory of trust in social networks.

2 Theory

We start with a simple one-period model of directed altruism and then add a second period in which the partner learns and reacts to the decision-maker’s action from the first period.

2.1 Directed Altruism Channel

We assume that there is a decision-maker, M and a partner, P , who are embedded in a social network that consists of a set of nodes (people) and edges connecting social neighbors. Figure 2 shows an example of a decision-maker with four direct friends and two indirect friends at social distance 2 (“friend of a friend”) and social distance 3 (“friend of a friend of a friend”). We calculate the social distance D_{MP} between the decision-maker and the partner as the shortest path connecting them: for example, two direct friends have a social distance of 1, while a friend of a friend is distance 2 apart.

From some set $X = [0, x_{max}]$, the decision-maker chooses an action x that affects both the decision-maker’s payoff $\pi_M(x)$ as well as the partner’s payoff $\pi_P(x)$. We assume that the decision-maker’s payoff is decreasing in action x , while the partner’s payoff is increasing in x ; hence we say that the decision is “more generous” if the decision-maker chooses a higher x . The total social surplus generated by the decision-maker’s action is denoted as $S(x) = \pi_M(x) + \pi_P(x)$. We distinguish three important types of games.

Definition 1 *Giving is efficient (inefficient) in the game played between the decision-maker and partner if total surplus is increasing (decreasing) in the action x . Giving is neutral if $S(x)$ is constant.*

For example, in the classic dictator game, altruism is neutral. For modified dictator games, as studied in Andreoni and Miller (2002), where tokens are worth more to the partner than to the decision-maker, giving is efficient. When the tokens are worth less to the partner than to the decision-maker, giving is inefficient.

For the present purposes, we assume that the decision-maker’s choice x_{MP}^* remains unknown to the partner — this assumption will correspond to the anonymous treatment in our experiment. To model the action of a rational decision-maker with altruistic utility, we adapt the specification of Andreoni and Miller (2002) :

$$x_{MP}^* = \arg \max_x u(\pi_M(x), \pi_P(x); \gamma_{MP}). \quad (1)$$

We make the technical assumption that $u(\cdot)$ is quasi-supermodular, strictly increasing, concave in the decision-maker’s own payoff, and weakly increasing and concave in the partner’s payoff. We interpret γ_{MP} as the decision-maker’s altruism towards the partner and we assume that $\partial^2 u / \partial \pi_M \partial \gamma_{MP} < 0$ and $\partial^2 u / \partial \pi_P \partial \gamma_{MP} > 0$, such that the decision-maker’s choice x is always increasing with altruism. We also adopt a tie-breaking assumption: if the decision-maker is indifferent between two different actions, she will choose the action that gives the partner a

higher material payoff.

We allow altruism to depend on social distance:

$$\gamma_{MP} = \gamma_1 D_{MP} + \gamma_M. \tag{2}$$

We think of γ_M as defining the *baseline altruistic type* of the decision-maker M and of γ_1 as the strength of *directed altruism*.¹

In our experiments, we observe the decision-maker’s action x_{MP}^* for each decision-maker/partner match. From our baseline survey we know the social distance D_{MP} as well as basic demographic information such as the gender of both the decision-maker and the partner. We even have a proxy for the decision-maker’s baseline type, because we ask her to make a decision for a nameless partner. Since we ask each decision-maker for multiple decisions, we can use random effects to capture previously unobserved characteristics. We will estimate variants of the following empirical model:

$$x_{MP}^* = \alpha Z + \gamma_1 D_{MP} + \gamma_M + \epsilon_{MP}, \tag{3}$$

where Z are the characteristics of the decision-maker and the partner, and γ_M is a random effect for each decision-maker. We use social distance dummies to capture the directed altruism effect, whose magnitude we will compare to the average decision towards an anonymous player.

2.2 The Future Interaction Channel: Enforced Reciprocity

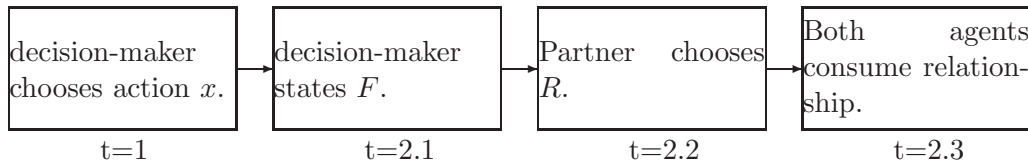
We capture the effects of future interactions between decision-makers and partners by adding a second period to our basic model in which the partner observes and reacts to the decision-maker’s action from the first period. This extension allows us to examine three mechanisms that provide incentives for the decision-maker to increase her action above her altruism-driven choice x_{MP}^* : (1) enforced reciprocity, (2) signaling, and (3) preference-based reciprocity. We start our analysis with a simple model of enforced reciprocity based on recent work by Möbius and Szeidl (2007), who provide a tractable model for capturing repeated game effects in social networks.

¹We do not take a position on whether altruism is the result of “warm glow”, as in Andreoni (1990) or arises from preferences over payoff distributions, as in Fehr and Schmidt (1999) or Charness and Rabin (2002).

2.2.1 Basic Enforced Reciprocity Model

We divide the second period into three stages, as shown in figure 1. In stage 2.1 the decision-maker can state a positive number F , which we interpret as “favor” for which the decision-maker requests repayment and which satisfies $0 \leq F \leq \pi_M(0) - \pi_M(x)$. Therefore, the decision-maker cannot claim favors that are larger than the payment loss from acting generously. In stage 2.2, the partner can either “repay” the favor and transfer $R = F$ to the decision-maker, or choose not to repay the favor and transfer $R = 0$. In stage 2.3 the decision-maker and the partner “consume” their friendship and both receive value V_{MP} from it. We interpret V_{MP} as the present value of the future relationship between the decision-maker and partner. Not repaying a favor implies that the relationship between the decision-maker and the partner breaks down, leading to no future utility for either party – otherwise each party receives the full value of the relationship.²

Figure 1: Timing of the actions of a decision-maker and partner in the enforced reciprocity model



We include the payoffs from both periods in the decision-maker’s altruistic utility function $u(\cdot)$, which we defined in the previous section:

$$u(\pi_M(x) + R + I(R, F)V_{MP}, \pi_P(x) - R + I(R, F)V_{MP}; \gamma_{MP}). \quad (4)$$

The indicator function $I(R, F)$ equals one if $R = F$ and is zero otherwise. For simplicity’s sake, we assume that the partner has selfish preferences:³

$$U_P(x, R) = \pi_P(x) - R + I(R, F)V_{MP}. \quad (5)$$

We focus on the unique subgame perfect equilibrium (SPE). It is easy to see that the

²We assume that friendships break down automatically if a favor is not repaid. Möbius and Szeidl (2007) give the borrower a choice between breaking off a relationship or continuing. In their model, breaking a relationship after non-payment is optimal because it signals to the lender that the borrower no longer values their relationship.

³Assuming altruistic preferences for the partner does not affect our qualitative results. However, if the partner is altruistic, two complications arise. First, he might repay favors that are greater than V_{MP} because he takes the utility loss of the decision-maker from breaking the relationship into account. Second, the partner might repay parts of large favors voluntarily.

partner's optimal action in any SPE is to exactly repay a favor by choosing $R = F$ when the value of the relationship with the decision-maker exceeds F , and to pay zero otherwise. Since a breakdown of the relationship hurts both players, the decision-maker will only do favors that are re-paid in equilibrium. Accordingly, the next theorem characterizes the decision-maker's new optimal action \tilde{x}_{MP}^* .

Theorem 1 (*Enforced Reciprocity Model*) (1) *When giving is efficient, the decision-maker takes action $\tilde{x}_{MP}^*(V_{MP})$, which is increasing in V_{MP} and satisfies $\tilde{x}_{MP}^*(0) = x_{MP}^*$. In the second period, the decision-maker requests that the partner repays a favor, $F = V_{MP}$. Both the decision-maker's utility and the partner's utility increase with V_{MP} , the relationship value.* (2) *When giving is inefficient, the decision-maker always chooses the same action, x_{MP}^* , that he chose under purely anonymous interaction with the partner.*

Proof: see appendix A

We might expect that the decision-maker always chooses her action \tilde{x}_{MP}^* such that, after accounting for the partner's repayment R , she has the same material payoff $\pi_M(x_{MP}^*)$ as under anonymous decision-making. However, the decision-maker will generally moderate her action, because her marginal utility derived from the partner's consumption decreases when the partner is made better off. This leads to the following supporting proposition:

Lemma 1 *Directed altruism and favors are substitutes in the sense that an altruistic decision-maker ($x_{MP}^* > 0$) chooses $\tilde{x}_{MP}^* < \bar{x}_{MP}$, where \bar{x}_{MP} is the solution to the equation:*

$$\pi_M(x_{MP}^*) - \pi_M(\bar{x}_{MP}) = V_{MP}. \quad (6)$$

A perfectly selfish decision-maker chooses $\tilde{x}_{MP}^ = \bar{x}_{MP}$.*

Proof: see appendix A

By way of illustration, in the case of a dictator game, where giving is efficient and increases the sum of the decision-maker's and partner's surplus, the above result implies that, compared to anonymity, under non-anonymity a perfectly selfish decision-maker will increase the number of tokens she passes more than will an altruistic decision-maker.

2.2.2 Relationship Value and Network Flow

Since we cannot directly observe V_{MP} , we proxy for it in our empirical analysis in two ways. First, we expect that relationship value decreases with greater social distance because in the

future a decision-maker is more likely to interact with a socially close partner. Therefore we will use social distance as a proxy for relationship value.

Our second proxy for the relationship value between the decision-maker and her partner is the new *maximum network flow* measure that was recently introduced by Möbius and Szeidl (2007). In their model, the network flow provides an upper limit for the amount of money that a borrower can informally obtain from a lender in the social network. Formally, the maximum network flow is defined as follows: all direct friendship links are assigned a relationship value of 1. A flow from partner to decision-maker is a set of transfers, F_{ij} , between any two direct neighbors i and j within the network such that (a) no transfer exceeds the value of a link and, (b) for any agent other than the partner and the decision-maker, the individual flows to and from that agent sum up to zero (flow preservation). The maximum network flow is the maximum value of the net flows to the decision-maker across all possible flows. Figure 3 provides a number of examples to illustrate the network flow concept and how it differs from the simple social distance measure.

The maximum network flow measure highlights information about the network structure that is not reflected in the “consumption value” of friendship captured by the simple social distance measure. Intuitively, the greater the number of distinct paths connecting the decision-maker with the partner, the higher the network flow. For example, sharing a large number of common friends will tend to increase network flow. Therefore, network flow formalizes a common intuition in the sociology literature that “dense” networks are important for building trust because these give agents the ability to engage in informal arrangements (Coleman 1988, Coleman 1990).

We can easily embed Möbius and Szeidl’s (2007) model in our framework by adding one more time period after the partner has decided whether to return the favor. If the partner fails to repay ($R = 0$), the decision-maker can propose an alternative set of transfers F_{ij} involving *all agents* in the social network. The net flow to the decision-maker, $\sum_i F_{iM}$, cannot exceed the initial request F . The alternative transfer arrangement needs to be accepted by all affected agents in the social network other than the partner. If any transfer F_{ij} between two agents is not made, the relationship between them breaks down. Following Möbius and Szeidl (2007) we say that a set of transfers is “side-deal proof” if the partner cannot propose an alternative set of transfers to a subset of agents in the network which that make her strictly better off and each agent in the network weakly better off.

Theorem 2 *The largest favor the decision-maker can request from the partner and secure through a side-deal proof transfer arrangement is equal to the maximum network flow.*

Proof: see Möbius and Szeidl (2007)

The transfer arrangement serves as an informal insurance policy for the decision-maker: if a partner acts selfishly, the decision-maker can still extract a payment F through common friends, who in turn “punish” the partner by extracting the favor through their direct dealings with the partner.

The basic network flow measure can include every member of the social network in the alternative transfer arrangement. Yet surely this assumption is unrealistic for large social networks. In his empirical analysis of job search networks Granovetter (1974) found that decision-makers predominantly utilize links that are, at most, a distance $K = 2$ to $K = 2.5$ away.⁴ We interpret K as the “circle of trust” that the decision-maker enjoys, and use $K = 2$ when we calculate network flow as a proxy of relationship value.

2.2.3 Testable Implications of the Enforced Reciprocity Model

We will estimate variants of the following empirical model:

$$\tilde{x}_{MP}^* = \eta Z + \theta x_{MP}^* + \phi V_{MP} + v_M + \epsilon_{MP}. \quad (7)$$

We include a random effect v_M for each decision-maker to control for unobserved heterogeneity in how the decision-maker responds to the prospect of future interaction. In table 1, we summarize the empirical predictions of the enforced reciprocity model. Theorem 1 implies that, controlling for directed altruism, decision-makers only treat friends more generously than strangers under non-anonymity ($\phi > 0$) when giving is efficient. The enforced reciprocity model also predicts that this effect will be more pronounced for more selfish agents ($\theta < 0$) and that the maximum network flow will be a better proxy for relationship value than is social distance.

2.3 The Future Interaction Channel: Signaling

In recent work, Benabou and Tirole (2006) analyzed a similar two-player model where the decision-maker is concerned about her reputation for generosity, and their theoretical model offers an alternative explanation for why decision-makers treat friends more generously than strangers under non-anonymity. In this framework, individuals care that others think they are an altruistic (rather than greedy) type; hence when their actions can be observed, they will act more generously in order to project a good impression.

⁴We define the distance of a link ij from an agent k as the average social distance of i to k and j to k .

2.3.1 Basic Signaling Model

In order to adapt Benabou and Tirole’s (2006) model to our framework, in this section, we assume that the decision-maker and partner play a modified dictator game in which tokens have value r_M to the decision-maker and value r_P to the partner. We also assume that the decision-maker has a quadratic utility function:

$$u(\pi_M(x), \pi_P(x); \gamma_{MP}) = \gamma_G(50 - x)r_M + \gamma_{MP}xr_P - k\frac{x^2}{2}. \quad (8)$$

The new parameter γ_G captures the decision-maker’s “greed”. The key feature of Benabou and Tirole’s (2006) model is that the preference parameters (γ_G, γ_{MP}) vary between individuals: while the decision-maker knows her own preferences, the partner can infer the decision-maker’s preferences by observing her action in the second period. Following Benabou and Tirole (2006), we focus on the case where (γ_G, γ_{MP}) are independently and normally distributed with means $(\bar{\gamma}_G, \bar{\gamma}(D_{MP}))$:

$$(\gamma_G, \gamma_{MP}) \sim N(\bar{\gamma}_G, \bar{\gamma}(D_{MP}), \sigma_G^2, \sigma_{MP}^2, \sigma_{G,MP} = 0). \quad (9)$$

Note, that we allow the partners’ mean beliefs about the decision-maker’s generosity to depend on social distance, but the precision of beliefs is independent of social distance. This assumption is validated in our empirical analysis, where we find that partners (correctly) expect decision-makers to whom they are socially close, to be kinder towards them. But, once we correct for social distance, we do not find that partners are any better at predicting the kindness of friends compared to predicting the kindness of strangers.

The decision-maker cares about being perceived as not greedy (low γ_G) by the partner in the second period, as well as being perceived as generous (high γ_{MP}). Her combined first and second period utility is therefore:

$$\underbrace{\gamma_G(50 - x)r_M + \gamma_{MP}xr_P - k\frac{x^2}{2}}_{\text{altruistic utility function}} - \underbrace{\mu_G E_P(\gamma_G|x, D_{MP}) + \mu_{MP} E_P(\gamma_{MP}|x, D_{MP})}_{\text{utility from reputation}}. \quad (10)$$

The parameters μ_G and μ_{MP} capture the intensity of the decision-maker’s concern for her reputation, as well as the probability that her partner will observe her action. We would expect that decision-makers care more about the beliefs of partners they are socially close to, and therefore that both μ_G and μ_{MP} decrease with greater social distance.⁵

Proposition 2 from Benabou and Tirole’s (2006) model then yields the optimal action for

⁵Our test for signaling works equally well if decision-makers care less about the beliefs of socially close partners.

the decision-maker:

$$\tilde{x}_{MP}^* = \underbrace{\frac{r_P \gamma_{MP} - r_M \gamma_G}{k}}_{x_{MP}^*} - \mu_G \chi(r_M, r_P) + \mu_{MP} \rho(r_M, r_P), \text{ where} \quad (11)$$

$\chi(r_M, r_P)$ and $\rho(r_M, r_P)$ are the standard normal signal extraction formulae:

$$\chi(r_M, r_P) = \frac{-r_M \sigma_G^2}{r_M^2 \sigma_G^2 + r_P^2 \sigma_{MP}^2} \quad \rho(r_M, r_P) = \frac{r_P \sigma_{MP}^2}{r_M^2 \sigma_G^2 + r_P^2 \sigma_{MP}^2}.$$

The difference in giving under non-anonymity and anonymity is therefore:

$$\tilde{x}_{MP}^* - x_{MP}^* = \frac{\mu_M r_M \sigma_G^2 + \mu_P r_P \sigma_{MP}^2}{r_M^2 \sigma_G^2 + r_P^2 \sigma_{MP}^2}. \quad (12)$$

From this expression, the main prediction of the signaling model follows:

Proposition 1 *In the signaling model, decision-makers always pass more (fewer) tokens to socially close versus socially distant partners under non-anonymity compared to anonymity if they care more (less) about their reputation with socially close agents.*

We also note that the effect of non-anonymity (meaning the excess giving under non-anonymous versus anonymous condition) is independent of the level of altruism. This effect contrasts to the enforced reciprocity model, where altruism and favors are substitutes.

2.3.2 Testable Implications of the Signaling Model

Table 1 summarizes the empirical predictions of the signaling model. Compared to the enforced reciprocity model, the main difference is that the signaling model predicts that decision-makers react similarly in games where giving is both efficient and inefficient.

2.4 The Future Interaction Channel: Preference-Based Reciprocity

The third channel through which future interaction can affect the decision-maker's action is through reciprocal incentives, under which an individual wishes to treat kindly (unkindly) those who have treated/will treat her kindly (unkindly). In other words, since observable actions create the possibility for a future exchange, the decision-maker will treat the partner more kindly because she anticipates the partner to behave kindly to her and wants to preemptively reciprocate that future kindness. (Dufwenberg and Kirchsteiger 2004).

2.4.1 Basic Preference-Based Reciprocity Model

We adapt Dufwenberg and Kirchsteiger’s (2004) model of sequential preference-based reciprocity. As in our enforced reciprocity model, we assume that the partner can make a payment R in the second period after observing the decision-maker’s action x . We assume that the partner can make this payment with probability $p(D_{MP})$, which is an increasing function of the social distance between herself and the partner. That is, we think of $p(D_{MP})$ simply as the probability that both agents will meet again in the future. The decision-maker has the same altruistic utility function as in our enforced reciprocity model, but also has “reciprocal incentives”:

$$\begin{aligned}
 (1 - p(D_{MP})) \cdot & \underbrace{u(\pi_M(x), \pi_P(x); \gamma_M)}_{\text{altruistic utility}} + & (13) \\
 + p(D_{MP}) \cdot & \underbrace{\left[u(\pi_M(x) + \hat{R}(x), \pi_P(x) - \hat{R}(x); \gamma_M) + \psi \kappa_M(x, \hat{R}(x)) \kappa_P(x, \hat{R}(x)) \right]}_{\text{altruistic utility } U_M(x, R) \text{ plus reciprocal incentives}}.
 \end{aligned}$$

The parameter ψ captures the strength of reciprocal incentives, which we assume to be constant across all decision-makers and partners. $\kappa_M(x, \hat{R}(x))$ represents the decision-maker’s kindness, which is a function of her action x and her belief about the partner’s future repayment $\hat{R}(x)$:

$$\kappa_M(x, \hat{R}(x)) = \max \left[0, \pi_P(x) - \hat{R}(x) - \pi_P(x_{MP}^*) \right]. \quad (14)$$

The decision-maker is kind if her action gives the partner *more* than his reference utility, which we define as his payoff under anonymous decision making.⁶ Our kindness function is also non-negative, guaranteeing the uniqueness of our equilibrium. $\kappa_P(x, \hat{R}(x))$ represents the decision-maker’s beliefs about the partner’s intended kindness:

$$\kappa_P(x, \hat{R}(x)) = \max \left[0, U_M(x, \hat{R}(x)) - U_M(x_{MP}^*, 0) \right]. \quad (15)$$

The decision-maker believes that the partner acts kindly if his action gives the decision-maker more than his payoff would have been under anonymous decision making.

⁶The choice of reference utility is a free parameter in all preference-based reciprocity models. For example, Dufwenberg and Kirchsteiger (2004) calculate the reference utility as the average of the highest payoff and the lowest payoff each of the two agents can secure for the other party (taking beliefs into account). Our particular choice simplifies our analysis considerably.

We similarly define the partner’s reciprocal utility function:

$$\pi_P(x) - R + \psi \kappa_P(x, R) \kappa_M(x, \hat{R}(x)). \quad (16)$$

As in our enforced reciprocity model, we assume that the partner has no intrinsic altruistic motivations. The partner’s intended kindness is $\kappa_P(x, R)$, and he believes that the decision-maker’s intended kindness is $\kappa_M(x, \hat{R}(x))$, where $\hat{R}(x)$ is the partner’s belief about the decision-maker’s belief $\hat{R}(x)$.

An equilibrium for the game consists of an action \tilde{x}_{MP}^* and a schedule $R(x)$ that taken together specify the partner’s optimal repayment for each of the decision-maker’s actions.

Proposition 2 *In the preference-based reciprocity model, there is a unique equilibrium with a repayment schedule that is increasing in x . When giving is inefficient, the decision-maker always makes the same choice $\tilde{x}_{MP}^* = x_{MP}^*$, as under the anonymous decision-making condition. When giving is efficient and agents’ reciprocal incentives, ψ , are sufficiently large, then the decision-maker chooses an action $\tilde{x}_{MP}^* > x_{MP}^*$ in equilibrium which increases with the probability of future interaction $p(D_{MP})$.*

Proof: see appendix B

When giving is inefficient, any giving beyond the anonymous level destroys social surplus, and hence at least one of the two players must receive a utility that is lower than under anonymous-giving. Hence, mutual positive kindness is not possible. Lowering x would allow the decision-maker to replicate any repayment schedule $R(x)$, leaving the partner’s payoff unchanged and increasing the decision-maker’s utility. But it would still give her lower utility than the anonymous allocation. However, when giving is efficient the decision-maker can “trade” an increase in x for a repayment $R(x)$, allowing both agents to behave kindly towards each other.

2.4.2 Testable Implications of Preference-Based Reciprocity

Table 1 summarizes the empirical predictions of the preference-based reciprocity model. Like the enforced reciprocity model it predicts extra generosity towards friends under non-anonymity only when giving is efficient. However, when we control for the frequency of future interaction (through social distance, for example), preference-based reciprocity differs from enforced reciprocity in that it does not predict that increased network flow has an independent effect on excess generosity.

3 Experimental Design

We used two web-based experiments to measure decision making in social networks: a series of modified *dictator games* with different exchange rates and a *helping game*. In all experiments, the decision-maker repeatedly chooses an action that determines her and the partner’s payoffs. The partner is either a nameless partner (randomly chosen from the network population), or a specific named player. Both experiments have two treatments: one where the decision-maker makes a choice anonymously, and one where the decision-maker makes his decision non-anonymously. One choice by each decision-maker is selected randomly at the end of the experiment, and both players are paid accordingly. If the selected decision is anonymous, the players are not told which decision was used to calculate their payments — otherwise they are told. The main difference between the dictator game experiments and the helping game experiments is that the anonymous and non-anonymous treatments are conducted within subjects, in the case of the dictator games, and between subjects in the case of the helping game. Moreover, in the dictator experiments, we also elicited partners’ beliefs about the expected generosity of various named decision-makers under anonymous and non-anonymous decision making scenarios. We elicited beliefs using an incentive-compatible mechanism which we describe below.

3.1 Measuring Friendship in Web-Based Experiments

Sociologists typically measure social networks by asking subjects about their 5 or 10 best friends. Since all of our experiments were web-based, we were concerned that the lack of interaction with a human surveyor would lead to more misreported friendships. We therefore developed two simple games that provide incentives for subjects to truthfully report their friendships.

For the dictator game, we used the *coordination game* technique. Each subject was told to list her 10 best friends and the average amount of time she spends with each of them per week (0-30 minutes, 30 minutes to 1 hour, 1-2 hours, 2-4 hours, 4-8 hours or more than 8 hours). The subject was paid some small amount (in our case, 50 cents) with 50 percent probability for each listed friend who also lists them. The probability increased to 75 percent if subjects also agreed on the amount of time they spend together each week. We made the expected payoff for each probability (25 or 37.50 cents) both large enough to give subjects an incentive to report their friends truthfully and small enough to discourage coordinated “gaming”. The randomization was included to limit disappointment if a subject was named by few people.

For the helping game, we developed the *trivia game* technique. Subjects also were asked to

list 10 friends in this game. Over the course of several weeks, a computer program randomly selected some of these subject-friend links and sent an e-mail message to the subject’s friend asking him to select the correct answer to a multiple choice question, such as what time he gets up in the morning. Once a subject’s friend had answered the question, the subject received an e-mail directing her to a web page where she had a 15 second time limit to answer the same multiple choice question about her friend. If the subject and her friend submitted identical answers, they both won a prize. The trivia game provides subjects with incentives to list friends with whom they spend a lot of time and with whose habits they are therefore familiar with.

The coordination game and the trivia game were run at Harvard University in two consecutive years (December, 2003, and December, 2004). This allowed us to directly compare both techniques by focusing on 167 subjects who participated in both studies.⁷ The 2003 experiment only allowed students to choose friends who lived in two neighboring dormitories which house about 17 percent of the undergraduate student population. On average, a subject listed 3.37 friends in 2004 whom they could have listed as a friend in 2003. Among this pool of friends, 64 percent were actually listed in 2003. Of all subjects, 34 percent listed all of their 2004 friends in 2003, and 77 percent listed at least half of them. This implies that over the course of the year subjects made about one new direct friend in their dormitory, which we consider a plausible outcome.

To define the social network, we say that two subjects have a direct link if one of them named the other person. We call this type of social network the “OR-network”.⁸

3.2 Dictator Game

After measuring the social network, we randomly assigned each subject the role of decision-maker or partner.⁹ Each decision-maker made several decisions over a period of several days. For each decision-maker, only one of her decisions was randomly selected at the end of the experiment, and she and her respective partner were paid accordingly and informed by e-mail about their earnings.

⁷569 subjects originally participated in the coordination game. Two thirds of the subjects were invited to the trivia game next year (except for seniors who left the university) and participation for the trivia game experiment was about 50 percent.

⁸We got very similar results when we used the “AND-network”, in which a link exists only if both subjects name each other. The OR-network definition has desirable monotonicity properties: a subject with an above average number of friends will have an above average number of friends in the measured network even when the network survey truncates his true network. This is not always true for the AND-network.

⁹In the experimental instructions, we referred to the decision-maker and the partner simply as player 1 and player 2.

At first, each decision-maker received an e-mail invitation to play modified dictator games with a *nameless* partner who was a randomly selected student in the decision-maker's dormitory (the player is made aware of this fact in the instructions). She was asked to make allocation decisions in two situations: in the anonymous situation (where neither decision-maker nor the partner learn each other's identity), and in the non-anonymous situation (where both players are informed about each other's identity by e-mail at the end of the experiment). In each situation, the decision-maker divided 50 tokens between herself and the partner three times under different exchange rates (3:1, 1:1, and 1:3). In the first decision each token was worth 30 cents to the partner and 10 cents to the decision-maker (meaning that giving is efficient). In the second decision, each token was worth 20 cents to both players (meaning that giving is neutral). In the third decision, each token was worth 30 cents to the decision-maker and 10 cents to the partner (meaning that giving is inefficient); as a result, for both subjects, the maximum earnings were \$15.

A few days after the first round of decisions, all decision-makers were invited by e-mail to participate in a second round, in which they are matched with five different *named* partners (we listed partners using their full first and last real names): (1) a direct friend (social distance $SD = 1$), (2) a friend of a friend ($SD = 2$), (3) a friend of a friend of a friend ($SD = 3$), (4) a student in the same staircase/floor who is at least distance 4 removed from the student, and (5) a randomly selected student from the same dormitory who falls into none of the above categories. Once again, the decision-maker was asked to make allocation decisions in both the anonymous and non-anonymous situations. In both situations, the decision-maker made the same three decisions as in the first round and allocated tokens at exchange rates of 1:3, 1:1, and 3:1, respectively.

Note that each decision-maker made 6 decisions for each partner involving 3 different exchange rates under anonymous and non-anonymous treatments. All together, she took 36 decisions, which made it very difficult to identify which decision became effective ex-post from one's earnings. In order to identify a match, a decision-maker would have needed to "code" each anonymous decision by making some unique allocation decision (such as passing 26 instead of 25 tokens), and then remember these separate decisions when the payments were made. We believe that few subjects in our experiment went to such lengths, since we also informed them in advance that payments would take two to three weeks to process.

We also measured the partners' beliefs of how many tokens decision-makers would pass in the anonymous treatment. Each partner in the network population got an e-mail invitation to participate in a single web-based experiment where the partner was asked to predict, out of 50 tokens, how many 5 different decision-makers (whose real names were presented to the partner)

would pass to the (named) partner under each of the three exchange rates (1:3, 1:1, and 3:1) in the anonymous treatment. The subject who played the role of the partner knew that *at most* 1 of these 15 decisions would be selected for payment. For each mispredicted token, 10 cents were subtracted from the partner's earnings. Therefore, the partner had incentives to report his *median* belief.

3.3 Helping Game

In the helping game, each decision-maker was endowed with \$45, and each partner was endowed with \$0. The experimenter secretly chose a random price between \$0 and \$30. The decision-maker was asked to report the maximum price that she would be willing to pay in order for the partner to receive a gain of \$30. If her maximum willingness to pay was below the price chosen by the experimenter, the partner got \$0 and the decision-maker kept her initial endowment.

Effectively, the decision-maker in the helping game revealed how much she values a \$30 gain for the partner. As in our dictator game design, subjects were invited to make two rounds of decisions: in the first round they played with a nameless partner, while in the second round they faced four named partners. Also, we chose a between-subjects design for the anonymity factor: the decisions for the nameless partner in the first round are always anonymous, while in the second round decisions are either all anonymous or all non-anonymous. Every subject played both roles in the game, as a decision-maker and as a partner. Of course only one decision in one role was selected for payment.

4 Data Description

All experiments were conducted with undergraduates at Harvard University who were at least at the sophomore level.

4.1 Dictator Games

In December, 2003, Harvard undergraduates at 2 (out of 12) upperclass dormitories were recruited through posters, flyers, and mail invitation and directed to a web site. Prospective subjects were asked to provide their e-mail addresses and were sent a password. Subjects without a valid e-mail address were excluded from the game. All future earnings from the experiment were transferred to the student's electronic cash-card account. These prepaid cards are widely used on campus as a cash substitute, and many off-campus merchants accept the cards as well.

Subjects who logged onto the website were asked to (1) report their best friends' names using the coordination game technique described in the previous section and (2) fill in a questionnaire asking basic demographic information. Subjects were restricted to naming friends from the two dormitories where our experiment was conducted. Subjects were paid their earnings from the coordination game, plus a flat payment of \$10 for completing the survey. Moreover, they were eligible to earn cash prizes in a raffle, which added, on average, another \$3 to their earnings.

Out of 806 students in those two dormitories, 569 (or 71 percent) participated in the social network survey. The survey netted 5690 one-way links. Of those, 2086 links were symmetric links where both subjects named each other.¹⁰ The resulting "OR"-network consists of a single connected component with 802 subjects. Fifty-one percent of subjects in the baseline survey were women; 49 percent were men. Thirty-one percent of the subjects were sophomores, 30 percent were juniors and 39 percent were seniors.

The dictator game experiment was conducted in May, 2004, over a one-week period. Half of all subjects who participated in the baseline study were randomly selected to be allocators. Out of 284 eligible decision-makers, 193 participated in round 1 and 181 participated in round 2. The participants were representative of the baseline sample composition: 58 percent were women, 28 percent were sophomores, 28 percent were juniors, and 44 percent were seniors.

4.2 The Helping Game

Information on social networks was collected through an online trivia game at the popular student website `facebook.com`, where students post their online profiles with biographic information as well as a list of their on-campus friends. More than 90 percent of Harvard undergraduates are members of `facebook.com`. As Ward (2004) notes, however, users often compile lists of over 100 friends, which contain many people with whom they maintain only weak social ties. The trivia game technique provides a particularly convenient method to identify the subset of strong friendships among facebook friends. In December, 2004, an invitation to the trivia game appeared for a four-week period on the home page of `facebook.com` after a member logged in. Over this period, 2,360 students completed the trivia game sign-up process. Upperclassmen had higher participation rates than freshmen, with only 34 percent of freshman responding, but 45 percent, 52 percent, and 53 percent of sophomores, juniors, and seniors participating, respectively.

¹⁰For symmetric links, the two subjects' assessment of the amount of time spent together in a typical week did not differ more than 1 category out of 5 in 80 percent of all cases (subjects could choose 0 to 30 minutes, 30 minutes to 1 hour, 1 hour to 2 hours, 2 hours to 4 hours, or more than 4 hours).

There were 12,782 links between participants out of 23,600 total links¹¹ and 6,880 of these links were symmetric. In total, 5,576 out of the 6,389 undergraduates at Harvard had either participated in the trivia game or had been named by a participant. The social “OR”-network of 5,576 individuals contains a single component (meaning all individuals are connected), with a mean path length of 4.2 between participants.

The helping game experiment was conducted in May, 2006, over a one-week period with all juniors and seniors who had participated in the previous year’s trivia game. In the first part of the helping game, 776 subjects participated, while 695 subjects completed the second part. Fifty-eight percent of participants were women. Forty-six percent were juniors and 54 percent were seniors.

4.3 Summary Statistics

Figure 4 and Table 2 show the mean actions of decision-makers for the dictator games and the helping game, for both the anonymous and non-anonymous treatments. Two major regularities are immediately apparent: in all games and in both treatments, the decision-makers’ generosity towards the partner decreased with social distance, and for any game and at any social distance the decision-makers’ generosity was always higher in the non-anonymous treatment than in the anonymous treatment. Differences between treatments were significant across all social distances in the dictator game, and for social distance 1 and 2 in the helping game. For both games and both treatments, decision-makers’ giving choices were significantly larger when the partners were their direct friends than when partners were at any other social distance.

In the dictator game, with the 1:3 exchange rate, the decision-maker passed about 19.19 tokens to a direct friend versus 12.20 tokens to a partner at social distance 4. With an exchange rate of 3:1, the decision-maker passed only 8.03 versus 6.15 tokens, respectively. In the non-anonymous treatment, for all social distances, the decision-maker passed about 4 to 5 more tokens when altruism was efficient, and about 2 to 4 more tokens when altruism was inefficient. In the helping game’s anonymous treatment, the average maximum willingness to pay \$12.77 to a direct friend decreased to \$7.09 for a partner at social distance 4. Non-anonymity increased the cutoff by about \$2 across social distances.

Curiously, in the anonymous treatment for all the dictator games and in the helping game, nameless partners are treated more generously than are friends of friends, despite the fact that the expected social distance of a randomly chosen partner is at least 3. In the non-anonymous treatment, on the other hand, the contributions to nameless partners closely track contributions to named partners at distance 3.

¹¹Subjects also listed non-participants in our experiment as friends.

We can interpret nameless decisions in the anonymous treatment as the decision-makers' baseline or unconditional generosity, since the decision-maker has no information about the partner. Our data replicates the well-known finding of Andreoni and Miller (2002) and Fisman, Kariv, and Markovits (2005) that individuals are highly heterogeneous in their unconditional altruism. In particular, we also find that many subjects are perfectly selfish: in the three dictator games, 28, 46, and 64 percent of subjects pass zero tokens, while in the helping game, 20 percent set a cutoff of zero dollars.

For the dictator games, we also collected the partners' beliefs about the decision-makers' actions, which we report in table 3 and figure 4. The partners' beliefs are reasonably accurate and anticipate the effect of greater social distance. Beliefs are most accurate when altruism is efficient. When altruism is inefficient, partners expect decision-makers to be somewhat more generous than they actually are.

5 Directed Altruism

In this section we use data from the anonymous treatments to analyze how decision-makers' altruistic preferences vary with the degree of social distance. For the dictator game experiment, we can also examine to what extent partners can predict decision-makers' preferences. Finally, we present evidence that subjects who behave more altruistically in general, as measured by their treatment of nameless partners, also tend to have more altruistic friends.

5.1 Decision-Makers' Actions

In section 2.1 we derived the following specification for estimating the strength of directed altruism:

$$x_{MP}^* = \alpha Z + \gamma_1 D_{MP} + \gamma_M + \epsilon_{MP}. \quad (3)$$

Recall that x_{MP}^* is the decision-maker's action in the anonymous treatment for one of the three dictator games or for the helping game. Since agents' actions are bounded below by zero and above 50, in the case of the dictator games, and by zero and 30, in the case of the helping game, we use Tobit regressions to estimate equation 3. We exploit the fact that we observe multiple actions for each decision-maker in the anonymous treatment, and control for unobserved heterogeneity in the decision-maker's baseline altruistic type γ_M by including random effects.¹² We control for the social distance, D_{MP} , between the decision-maker and the partner by including dummy variables $SD1$ (meaning a direct friend at social distance $SD = 1$)

¹²Our results are very similar when we estimate instead equation 3 using standard random effects or fixed effects GLS.

to $SD5$ (meaning $SD = 5$). The omitted categories are $SD4$ for the dictator games and $SD5$ for the helping game.¹³ The estimated coefficient on $SD1$ in a dictator game, for example, should therefore be interpreted as the number of extra tokens that the decision-maker passes to a direct friend compared to a distant partner in the anonymous treatment, while the estimated coefficient on $SD2$ captures directed altruism towards a friend of friend. The estimates of the Tobit regression for all of the dictator games and for the helping game are reported in the odd-numbered columns of table 5.

We also estimated the same specification with additional covariates, and report the results in the even-numbered columns of table 5. We included the decision-maker’s action towards a nameless player in the anonymous treatment as a proxy for the decision-maker’s baseline altruistic type γ_M . In the helping game, we can also control for the partner’s nameless decision because all subjects in the helping game played both the role of the decision-maker and the partner. Furthermore we added dummy variables for both players’ gender, their class (sophomores, juniors, or seniors), and whether they share a staircase (dictator game) or a dormitory (helping game) with the partner.¹⁴

Result 1 *Baseline altruism and directed altruism are correlated. Subjects who give more to nameless partners also give more to specific named partners.*

The two variables that consistently and strongly predict how generously a decision-maker treats a partner in her social network are the social distance from and the generosity displayed towards a nameless partner. Looking across all regression specifications, for both the dictator games and the helping game, each one unit increase in generosity towards a nameless partner is associated with a 0.56 to 1.40 unit increase in generosity towards a named player. Since the nameless decision and the named decisions were elicited one week apart, this continuity indicates a remarkable degree of stability in the decision-makers’ preferences over time. Because the pass-through from “nameless altruism” to “named altruism” is fairly close to 1, we view the nameless decision of a decision-maker as a useful proxy measure of her baseline altruism, a trait that strongly influences the decision-maker’s action towards specific named players.

Result 2 *Close social ties induce directed altruism. Allocations to friends are substantially higher than allocations to distant partners/strangers.*

Moreover, social distance also matters greatly: decision-makers are substantially more generous to direct friends than to partners located at greater social distance. Generosity decreases

¹³The social network that we used for the house experiment is much larger because it potentially involves all Harvard undergraduates. Therefore, the maximum social distance between subjects is higher.

¹⁴Since the participation rate was lower for the helping game, a dummy variable for sharing the same entryway is less useful.

quickly and monotonically with social distance, although the estimated coefficients on SD2 and SD3 are not significantly different from each other for all games. Given the three exchange rates in the dictator games, the distance coefficients are of similar magnitude, which implies that decision-makers are making a greater relative sacrifice in the case of inefficient altruism.

In order to assess the magnitude of directed altruism, in table 4 we compare the estimated coefficients on social distance dummies SD1 to SD3 with the average generosity displayed towards nameless partners in the anonymous treatment. Directed altruism towards friends is equal to 52 percent of the average nameless generosity shown in the efficient dictator game, and increases to 88 percent for the helping game. When altruism is inefficient, the directed altruism effects even exceed average nameless generosity. Social distance, therefore, is as important a predictor of a subject's generosity as is her baseline altruism.

There is another way to look at the magnitude of directed altruism: how much less selfish does a decision-maker become, relative to the fat-tailed population distribution of baseline altruism, when she makes decisions for socially close partners (Andreoni and Miller 2002, Fisman, Kariv, and Markovits 2005)? Therefore, in table 4, we also report the estimated coefficients on SD1, SD2 and SD3 as percentages of the standard deviation of the distribution of nameless decisions. We find that greater social proximity to the partner in the dictator games moves the decision-maker's generosity by at least 0.47 of a standard deviation, and by a maximum of 1.15 of a standard deviation in the helping game.

Interestingly, the decision-maker's and the partner's gender, and their geographic proximity, have no significant effect on generosity. However, the signs of the estimated gender coefficients of the decision-maker are consistent with the work of Andreoni and Vesterlund (2001), who found that men are more likely to exhibit social-surplus maximizing preferences: they are more generous in dictator games when giving is efficient and less generous when giving is inefficient. College juniors are somewhat more selfish than are sophomores and seniors; however, most of the coefficients on the class dummies are insignificant.

5.2 Partners' Beliefs

We now analyze to what extent partners are aware of the decision maker's baseline altruism and directed altruism. We simply take the same empirical specification from the previous section for directed altruism, equation 3, but estimate it using the partner's beliefs instead of the decision-makers' actions on the equation's left-hand side. We also specify random effects on the partner level (rather than on the decision-maker level), since our experiment provides us with multiple observations for each partner. The odd- and even-numbered columns in table 6 report our estimates with and without additional covariates. In contrast to table 5, we do not

list estimates for the helping game because we only asked for partners' beliefs in the dictator game experiment.

Result 3 *Partners' beliefs reflect directed altruism: subjects accurately predict that, on average, their friends will be more generous. However, they overestimate the generosity of friends of friends and do not predict individual differences in baseline altruism between friends.*

Our first finding is that partners are well aware of the decision-maker's directed altruism. The number of extra tokens that partners expect from their direct friends ($SD = 1$) is, on average, close to the actual number of extra tokens decision-makers pass to their direct friends. Partners believe that decision makers are slightly more altruistic than they actually are when giving is efficient (exchange rate 1:3), and that they are slightly less altruistic when giving is inefficient (exchange rate 3:1) than they actually are. Interestingly, partners expect *friends of friends* ($SD = 2$) to be significantly more generous than these decision-makers actually behave when giving is efficient or neutral. In the previous section, we showed that decision-makers do not treat friends of friends significantly more generously than strangers. Yet the corresponding comparison of estimates of partners' expectations show that partners expect friends of friends to be almost twice as generous as strangers when giving is efficient or neutral, and the estimates for friends of friends are also strongly statistically significant.

Again, none of the other demographic and geographic covariates matter except for the decision-maker's gender: partners expect male decision-makers to be significantly less generous when giving is neutral, and especially when giving is inefficient. Again, this result is consistent with Andreoni and Vesterlund's (2001) findings.

Surprisingly, partners' beliefs are completely unaffected by the decision-maker's baseline altruism. In contrast, in the previous section we found that each extra token a decision-maker passes to a nameless partner increases her contribution to a named partner by a comparable amount. One explanation for this discrepancy could be that partners are good at estimating social distance and have learned that decision-makers treat friends more generously than strangers, but they are unable to observe decision-maker's baseline altruism. One might expect that partners are better at observing the preferences of direct friends compared to socially distant agents. Therefore, we re-estimate our empirical model and include an interaction term between the decision-maker's nameless decision and the social distance dummy $SD1$. The results are reported in the odd-numbered columns of table 7 (without demographic and geographic covariates). We do not find any evidence that partners are any better in observing the preferences of a decision-maker who is socially close or distant; in fact, two out of the three estimates of the interaction term are negative.

For 204 out of the 563 beliefs matches between a specific partner and a decision-maker, we also have the decision-maker’s actual choice for this partner. For this subset, we estimate our empirical model again, but now use the decision-maker’s actual choice rather than her nameless decision as a proxy for her baseline altruism. The estimates are reported in the even-numbered columns of table 7. Again, neither the decision-maker’s actual choice nor the interaction between the actual choice and the social proximity to the decision-maker affect a partner’s expectations.

5.3 Correlation in Altruistic Preferences

We next examine whether more altruistic subjects also have more altruistic friends. We separate decision-makers into (approximate) quintiles based on their dictator game and helping game choices for nameless partners. Tables 8 and 9 present the resulting distribution of friends’ generosity for each quintile.

Result 4 *Friends sort by their baseline altruism. Subjects with a high level of baseline altruism tend to have more friends with a high level of baseline altruism, while selfish subjects tend to have more selfish friends.*

First, we find that altruistic and selfish subjects have the same number of friends. However, a subject’s baseline altruism is correlated with the baseline altruism of her friends (chi-square test: DG $p < 0.001$, HG $p < 0.01$). That is, more selfish subjects have a greater number of selfish friends, and fewer altruistic friends, while altruists have fewer selfish friends and a greater number of altruistic friends. In particular, the most altruistic quintile in the helping game has a 25 percent greater number of highly-altruistic friends than does any other group; in the dictator game, the two most altruistic groups had over 20 percent more highly-altruistic friends than did any other group. Moreover, a subject’s friends’ mean nameless choice increases with the subject’s baseline altruism. The most altruistic subjects have friends that are 25 percent more altruistic than the most selfish subjects in the dictator game, and 14 percent more altruistic in the helping game. Using t-tests, the 3rd, 4th, and 5th quintiles are significantly different from the 1st in the helping game, and the 4th and 5th are different from the 1st in the dictator game. A non-parametric equality-of-medians test rejects that the five quintiles are drawn from distributions with the same median (DG $p = 0.039$, HG $p < 0.026$). Hence, it seems that subjects tend to seek out and/or maintain friendships with others who have similar social preferences.

We confirm this finding in table 10 where we take all pairs of participating friends (including those not matched in our experiment), and regress each subject’s baseline altruism on the

average baseline altruism of all their friends. As expected, the baseline altruism of a subjects' friends is positively and significantly related to her own baseline altruism in the helping game. A 10 percent increase in the generosity of a subject's friends would increase the subject's generosity by 2 percent. In the dictator game we have many fewer observations (since we only observed nameless decisions for half the subjects); however, the relationship is directionally positive, as we expect. Moreover, the correlation of types is not driven by clustering by gender.

An important consequence of the correlation in friends' baseline altruism is that it pays to be generous. In the helping game anonymous treatment, table 11 (column 1) regresses the average allocation to partners from decisions made by direct friends on their own baseline altruism (by quintile): partners with higher baseline altruism have substantially higher earnings. For example, the direct friends of the most altruistic partners set the cutoff in the helping game more than 5 dollars higher than the direct friends of the most selfish partners. Interestingly, this effect is entirely due to the fact that kinder partners have nicer friends, but is *not* due to kinder partners being treated more nicely by their friends: we already showed in our directed altruism regressions in table 5 that decision-makers do not treat generous partners better. Indeed, when we also control for the average baseline altruism of decision-makers in table 11 (column 2), then the partner's baseline altruism no longer predicts her earnings from friends' decisions.

6 Non-anonymity vs. Anonymity

We now examine how a decision-maker adjusts her action for a named partner when she interacts with him non-anonymously (meaning that both the decision-maker and partner are told that this decision was selected for payment). We start with a simple graphical analysis, which allows us to preview the main results; we then confirm these findings by testing the empirical specification from section 2.2. We also discuss which of our three models – the enforced reciprocity, signaling and preference-based reciprocity models – best fit our results.

6.1 Graphical Analysis

In figures 5 and 6, we plot the extra tokens that a decision-maker passes to a specific partner in the non-anonymous treatment relative to the anonymous treatment. Since our helping game was a between-group design, we can only perform this exercise for dictator decisions. We divide decision-makers into five bins depending on how generously they treated their partner in the anonymous treatment. The most selfish decision-makers are those who passed between 0 and 9 tokens in the anonymous treatment. The other bins range from 10 to 19, 20 to 29, 30 to 39,

and 40 to 50. We then plot the average number of extra tokens passed in the non-anonymous treatment versus the anonymous treatment, $\tilde{x}_{MP}^* - x_{MP}^*$, by bin and by relationship value V_{MP} . For figure 5, we use social distance as a proxy for relationship value, and for figure 6, we use the maximum network flow, with circle of trust $K = 2$ as the proxy.¹⁵ Since in our dataset, network flow takes values ranging from 0 to 21, we need to group observations in order to obtain a meaningful plot; hence we define “strong relationships” as those with a network flow greater than the median value, 3, and “weak relationships” as those with a network flow less than or equal to the median value of 3.

Both figures show that decision-makers substantially increase their action in the non-anonymous compared to the anonymous treatment, unless they already behaved very generously in the anonymous treatment.¹⁶ This non-anonymity effect is strongest—equal to ten extra tokens sent to the same partner—when the decision-maker is selfish and when giving is efficient. The effect is less than half as large when giving is inefficient, when decision-makers pass at most 5 extra tokens in the non-anonymous treatment.

The main insight that we take from both graphs is that the non-anonymity effect declines with relationship strength (measured by either social distance or maximum network flow) when giving is efficient, and the effect somewhat declines when giving is neutral. However, when giving is inefficient, the decision-makers’ contributions do *not* decrease with social distance for four out of the five bins. This result provides some preliminary evidence in support of the enforced reciprocity and preference-based reciprocity mechanisms and against the signaling mechanism.

The two graphs also suggest that directed altruism and the non-anonymity effect are substitutes: controlling for the strength of a relationship (by fixing either social distance or maximum network flow), we find that the non-anonymity effect decreases monotonically in most cases, as decision-makers become more generous in the anonymous treatment.

6.2 Tobit Regressions

In section 2.2.3, we derived the following empirical specification for estimating the non-anonymity channel:

$$\tilde{x}_{MP}^* = \eta Z + \theta x_{MP}^* + \phi V_{MP} + v_M + \epsilon_{MP}. \quad (7)$$

Recall that \tilde{x}_{MP}^* is the decision-maker’s action in the non-anonymous treatment when matched with a specific named partner P in one of the three dictator games or in the helping game.

¹⁵Our results are qualitatively very similar for other values of K .

¹⁶Even in this case, the majority of decision-makers do not decrease their action – the negative averages result from a few decision-makers decreasing their contributions substantially in the non-anonymous treatment.

We again use Tobit regressions to account for censoring the left-hand side variable, and we also exploit the panel structure of our data again to control for unobserved heterogeneity in the decision-maker’s response to the non-anonymous treatment. We proxy for the strength of the decision-maker’s relationship with the partner, V_{MP} , by including either social distance dummies or the maximum network flow measure ($K = 2$).¹⁷ The omitted social distance categories are $SD4$ for the dictator games and $SD5$ for the helping game. The estimated coefficient on $SD1$ in a dictator game, for example, should therefore be interpreted as the number of extra tokens that the decision-maker passes to a direct friend under non-anonymity compared to the number of extra tokens that she passes to a stranger under non-anonymity. All of our regressions control for the class of the decision-maker and partner, because we expect the non-anonymity effect to be smaller for juniors and particularly for seniors, since they are less likely than sophomores to interact with the decision-maker in the future. On the right-hand side we also include the decision-maker’s action towards a nameless partner in the non-anonymous treatment as a proxy for the random effect v_M that captures heterogeneity in how decision-makers respond to non-anonymity.

Importantly, we control for the decision-maker’s intrinsic altruism towards the *same* partner P by including her decision in the anonymous treatment, x_{MP}^* , on the right-hand side of all of our regressions. This inclusion poses a problem for the helping game given its between-group design because for each decision-maker/partner match in the helping game’s non-anonymous treatment, we do not observe the decision-maker’s choice for that partner in the anonymous treatment. We therefore estimate it by running an auxiliary random-effects Tobit regression with data from the anonymous treatment, include social distance dummies, and the same set of covariates Z (nameless decision, class dummies) as in our empirical specification of the non-anonymity channel.

For each of the three dictator games and for the helping game, we estimate three variants of our empirical model. We first use only social distance to proxy for the strength of a decision-maker’s relationship to the partner, then use only maximum network flow, and finally use both measures in the same regression. All results are reported in table 12.

Result 5 *The observability of decisions by partners increases giving more for friends than for strangers. The effect is only induced when giving increases social surplus; therefore, it is efficiency-enhancing.*

Our main finding is that, controlling for a decision-maker’s action in the anonymous treatment, her response to non-anonymity decreases with the strength of her relationship to the partner,

¹⁷In one specification, we include both types measures in the same regression.

but *only* if giving is not inefficient. This is true regardless of whether we proxy for the strength of a relationship using social distance or maximum network flow. The magnitude of this effect is large and is most pronounced in the dictator game with exchange rate 1:3. In this game, a decision-maker increases her action by 4.18 tokens when her partner is a direct friend as opposed to a socially distant partner, a difference that is statistically significant at the 1 percent level. Moreover, the social closeness effect is smaller but still significantly different from zero at the 5 percent level for friends of friends (SD2) in the efficient dictator game.¹⁸

In order to compare the magnitude of this non-anonymity effect with the magnitude of directed altruism, we compare (table 4) the estimated coefficients on social distance dummies SD1 to SD3, as well as a one-standard deviation increase in network flow with average generosity towards nameless partners in the anonymous treatment. In the non-anonymous treatment, friends receive an extra transfer of surplus, equal to almost 24 percent of average nameless generosity, for the efficient dictator game and about 35 percent of nameless generosity for the helping game. Friend of friends receive an extra transfer of about 18 percent of nameless generosity in the efficient dictator game. We find a similar pattern but with slightly smaller magnitudes for the neutral dictator game. Taken all together, the effect of non-anonymity is about half as large as the directed altruism effect.¹⁹ Moreover, the non-anonymity effect is generally weaker for decision-makers who are juniors and seniors compared to sophomores: the signs on the junior and senior dummies are consistently negative, even though these are not always statistically significant. This finding is consistent with our model of enforced reciprocity, since the length of the future relationship (and thus its value) is shorter (lower) for upperclassmen.

6.3 Probit Regressions

An alternative way to compare the non-anonymous with the anonymous treatment in our dictator games is to examine when the decision-maker passes *strictly more* tokens in the non-anonymous treatment compared to the anonymous treatment. In about 50 percent of all matches between a decision-maker and some specific named partner, the decision-maker passes the same number of tokens in both treatments. As long as the decision-maker attaches zero value to the relationship with the partner, this phenomenon is consistent with all three of our theoretical models.

We therefore re-estimate our empirical model in equation 7, but replace the dependent

¹⁸In contrast, when we tested for directed altruism in the previous section we did not find such a strong effect for second-order friends.

¹⁹The relative comparison of the non-anonymity and directed altruism effects does not change if we instead normalize all estimates by using the standard deviation of the distribution of nameless decisions.

variable with a dummy variable equal to 1 only if the decision-maker passes strictly more tokens under non-anonymity. We use a random-effects probit estimation, and report the results in table 13, with all estimates interpretable as marginal effects. In short, the probit results are consistent with our Tobit estimates. When giving is efficient, decision-makers are about 91 percent more likely to increase their token allocation when interacting non-anonymously with direct friends, than when interacting with socially distant partners ($SD > 3$). Similarly, each unit increase in the maximum network flow raises the probability of passing more tokens by 6.5 percent. Both the social distance and network flow effects disappear when giving is neutral or inefficient.

Interestingly, when we use an analogous probit regression to test when the decision-maker passes *strictly fewer* tokens under non-anonymity, there is some evidence that the effect of social distance is reversed (see table 14): when giving is efficient, decision-makers are less likely to pass strictly fewer tokens to direct friends than to a stranger, but are more likely to do so when altruism is inefficient. However, the estimates behind these conclusions are only marginally significant.

6.4 Discussion

We can now review the testable implications that we derived for the enforced reciprocity, signaling, and preference-based reciprocity mechanisms outlined in section 2 (summarized in table 1). First of all, we found that when we control for the decision-maker’s intrinsic altruism towards a partner, under non-anonymity friends are only treated more generously than strangers if giving is efficient, or (to some extent) if giving is neutral, *but not if giving is inefficient*. This difference is consistent with both the enforced reciprocity and the preference-based reciprocity mechanisms, but not with the signaling mechanism, which would predict excess generosity towards friends under all exchange rates.

Result 6 *The non-anonymity effect increases with maximum network flow.*

In table 12 we estimate one Tobit specification for each of our four games that includes both social distance dummies *and* maximum network flow on the right-hand side. For both the efficient dictator game and the helping game, we find that the coefficients on the social distance dummies decreases and becomes insignificant when we add network flow, but that the coefficient on flow remains significant for the helping game. We find a similar pattern in table 13 using probit models.

It is possible that network flow does not proxy for the network’s ability to enforce repayment of favors, but rather provides a measure for the amount of time that the decision-maker spends

with the partner. For the dictator games, we collected information on the average amount of time subjects spend with their *direct friends* each week. We find that this measure is essentially uncorrelated with our network flow measure for direct friends (correlation coefficient is 0.03). In table 15 we re-estimate our empirical model for the non-anonymity channel and include both network flow and time spent together per week. We find that when giving is efficient, greater network flow increases the decision-maker’s generosity towards a direct friend under non-anonymity, even when we control for time spent together. Moreover, the estimated coefficient on time spent together is consistently insignificant and *negative*. We interpret these findings as evidence for the enforced reciprocity hypothesis.

Result 7 *The non-anonymity effect and directed altruism are substitutes.*

We also find that the estimated coefficients on the decision-maker’s anonymous action, x_{MP}^* , are always less than one, which implies that directed altruism and the decision-maker’s response to non-anonymity are indeed substitutes. This conclusion is also consistent with the enforced reciprocity mechanism.

7 Conclusion

Our purpose in this paper was to determine the value of having a friend rather than a stranger make payoff-relevant decisions. We identified two main effects. First of all, directed altruism makes a decision-maker, on average, at least half a standard deviation more generous relative to the distribution of the decision-makers’ baseline altruism. If giving increases social surplus and the partner can observe the decision-maker’s action, generosity increases, on average, by another quarter of a standard deviation. The latter effect is well explained by a model of enforced reciprocity, building on Möbius and Szeidl (2007) in an effort to provide a tractable repeated-game model in social networks.

Our findings are a first step towards a broader theory of *trust* in social networks. The bulk of the experimental and theoretical literature on trust has evolved around one-shot games played between strangers. In particular, the seminal work of Berg, Dickhaut, and McCabe (1995) introduced the investment or trust game. This existing literature basically asks why strangers trust each other (absolute trust). We ask a complementary question: why should we trust some decision-makers more than others (differential trust)? Social networks can be measured and therefore provide us with ample data to test and calibrate models of differential trust.

A natural next step in this research agenda is to ask whether partners do in fact choose “trustworthy” decision-makers. Karlan, Möbius, and Rosenblat (2006) examine this question in a naturalistic field experiment with micro-loans in Peruvian shantytowns. In their design, partners (“borrowers”) can ask different decision-makers (“lenders”) to provide them with a loan. The price (“interest rate”) associated with different lenders is exogenously randomized by the experimenter. The lending data is then analyzed for preference reversals: how much is a partner willing to pay to replace a socially distant with a socially close decision-maker? These estimates can then be compared to the value of friendship which we derived in this paper.

References

- ANDREONI, J. (1990): “Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving,” *Economic Journal*, 100(401), 464–477.
- ANDREONI, J., AND J. MILLER (2002): “Giving According to GARP: An Experiment on the Consistency of Preferences for Altruism,” *Econometrica*, 70, 737–753.
- ANDREONI, J., AND L. VESTERLUND (2001): “Which is the Fair Sex? Gender Differences in Altruism,” *Quarterly Journal of Economics*, 116, 293–312.
- BENABOU, R., AND J. TIROLE (2006): “Incentives and Prosocial Behavior,” *American Economic Review*, 96, 1652–1678.
- BERG, J., J. DICKHAUT, AND K. MCCABE (1995): “Trust, Reciprocity and Social History,” *Games and Economic Behavior*, 10, 122–142.
- BOLTON, G., AND A. OCKENFELS (2000): “ERC: A Theory of Equity, Reciprocity and Competition,” *American Economic Review*, 90(1), 166–193.
- BRAÑAS-GARZA, P., R. COBO-REYES, M. PAZ ESPINOSA, N. JIMÉNEZ, AND G. PONTI (2006): “Altruism in the (Social) Network,” DFAEII working paper, University of the Basque Country - Department of Foundations of Economic Analysis II.
- CAMERER, C. (2003): *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton, NJ: Princeton University Press.
- CHARNESS, G., AND M. RABIN (2002): “Understanding Social Preferences with Simple Tests,” *Quarterly Journal of Economics*, 113(3), 817–869.
- COLEMAN, J. S. (1988): “Social Capital in the Creation of Human Capital,” *American Journal of Sociology*, 94, 95–120.
- (1990): *Foundations of Social Theory*. Cambridge, MA: Harvard University Press.

- DAVIDSSON, P., AND P. REYNOLDS (2005): “Panel Study of Entrepreneurial Dynamics,” data set, University of Michigan.
- DUFWENBERG, M., AND G. KIRCHSTEIGER (2004): “A Theory of Sequential Reciprocity,” *Games and Economic Behavior*, 47(2), 268–298.
- FEHR, E., AND K. SCHMIDT (1999): “A Theory of Fairness, Competition and Cooperation,” *Quarterly Journal of Economics*, 114(3), 817–868.
- FERSHTMAN, C., AND U. GNEEZY (2001): “Discrimination in a Segmented Society: An Experimental Approach,” *Quarterly Journal of Economics*, 116, 351–377.
- FISMAN, R., S. KARIV, AND D. MARKOVITS (2005): “Individual Preferences for Giving,” Discussion paper, mimeo, UC Berkeley.
- GLAESER, E. L., D. I. LAIBSON, J. A. SCHEINKMAN, AND C. L. SOUTTER (2000): “Measuring Trust,” *Quarterly Journal of Economics*, 115(3), 811–846.
- GOEREE, J. K., M. A. MCCONNELL, T. MITCHELL, T. TROMP, AND L. YARIV (2007): “Linking and Giving Among Teenage Girls,” mimeo, California Institute of Technology.
- GRANOVETTER, M. (1974): *Getting a Job*. Cambridge, MA: Harvard University Press.
- KARLAN, D., M. MÖBIUS, AND T. ROSENBLAT (2006): “Measuring Trust in Peruvian Shantytowns,” Discussion paper, working paper.
- MÖBIUS, M. M., AND A. SZEIDL (2007): “Trust and Social Collateral,” working paper, Harvard University and UC Berkeley.
- RABIN, M. (1993): “Incorporating Fairness into Game Theory and Economics,” *American Economic Review*, 83(5), 1281–1302.
- WARD, B. A. (2004): “Distance and Social Capital: Can Distance Be Good?,” mimeo, Harvard University.

A Proof of Theorem 1

We assume that giving is efficient. We define \bar{x} as follows:

$$\pi_P(\bar{x}) - F = \pi_P(x^*). \quad (17)$$

The partner gets the same payoff from action \bar{x} and paying F as he gets from the anonymous action x^* . Because giving is efficient, the decision-maker’s payoff satisfies:

$$\pi_M(\bar{x}) + F > \pi_M(x^*). \quad (18)$$

Concavity and quasi-supermodularity of the utility function yield the following inequalities:

$$\begin{aligned} u_1(\pi_M(\bar{x}) + F, \pi_P(\bar{x}) - F) &< u_1(\pi_M(x^*), \pi_P(x^*)) \\ u_2(\pi_M(\bar{x}) + F, \pi_P(\bar{x}) - F) &> u_2(\pi_M(x^*), \pi_P(x^*)) \end{aligned} \quad (19)$$

Furthermore, we know that x^* maximizes the decision-maker's utility under anonymity:

$$u_1(\pi_M(x^*), \pi_P(x^*)) = u_2(\pi_M(x^*), \pi_P(x^*)). \quad (20)$$

By combining these inequalities, it follows that the marginal utility of the decision-maker's own consumption is lower than the marginal utility of the partner's consumption when the decision-maker takes action \bar{x} and declares F . Hence, the decision maker will optimally increase her action x . It follows that $\tilde{x}^* > x^*$. Using an analogous argument, we can show that $\pi_M(x^*) - \pi_M(\tilde{x}^*) < 0$.

This shows that the decision-maker's optimal action will give greater surplus to both players and hence will increase both of their utilities under non-anonymity.

B Proof of Proposition 2

To keep things simple, we assume $p = 1$ (our results easily extend to the general case). We start by solving the partner's problem of choosing the repayment R . The partner maximizes:

$$\pi_P(x) - R + \psi [U_M(x, R) - U_M(x^*, 0)]^+ \kappa_M(x, \hat{R}(x)). \quad (21)$$

Differentiating with respect to R and using the fact that in equilibrium $R = \hat{R}(x)$, we get the first-order condition for an interior solution:

$$-1 + \psi \frac{\partial U_M(x, R)}{\partial R} \kappa_M(x, R) = 0. \quad (22)$$

Lemma 2 *The unique equilibrium repayment schedule $R(x)$ satisfies $R(x) = 0$ for $x \leq x^*$ and is weakly increasing in x .*

Proof: By definition, we have $\kappa_M(x^*, R(x^*)) = 0$, such that the partner would always like to choose a lower R for any $R > 0$. The FOC for an interior solution can be rewritten as:

$$\psi \frac{\partial U_M(x, R)}{\partial R} = \frac{1}{\pi_P(x) - R - \pi_P(x^*)}. \quad (23)$$

The LHS of the equation is strictly decreasing in R and strictly increasing in x , since we assumed a quasi-supermodular utility function for the decision-maker. The RHS is strictly increasing in R and strictly decreasing in x . Therefore, there is a unique interior solution for R , and the solution is strictly increasing in x . ♠

The decision-maker's objective function is:

$$U_M(x, \hat{R}(x)) + \psi \kappa_M(x, \hat{R}(x)) [U_M(x, \hat{R}(x)) - U_M(x_{MP}^*, 0)]. \quad (24)$$

Consider her optimal decision when giving is inefficient. By choosing the same decision as under anonymity, she can guarantee herself a utility $U_M(x_{MP}^*, 0)$. Any other choice of x will

either set $\kappa_M = 0$ or $\kappa_P = 0$. Therefore, her utility will be $U_M(x, \hat{R}(x))$, which is lower than $U_M(x_{MP}^*, 0)$.

Next, consider the case where giving is efficient and fix some $x > x_{MP}^*$. From the partner's first-order condition, we can see that $R(x) \rightarrow \pi_P(x) - \pi_P(x_{MP}^*)$ as $\psi \rightarrow \infty$. Therefore, the decision-maker receives almost all of the partners's extra profits gained from the increase in her action above x_{MP}^* , as long as both agents are sufficiently reciprocal. The decision-maker will hence be strictly better off by choosing x instead of x_{MP}^* . This shows that the optimal action of the decision-maker is strictly greater than x_{MP}^* .

Figure 2: Illustration of decision-maker's social network with four direct friends ($SD = 1$) and two indirect friends at distance $SD = 2$ ("friend of a friend") and at distance $SD = 3$ ("friend of a friend of a friend")

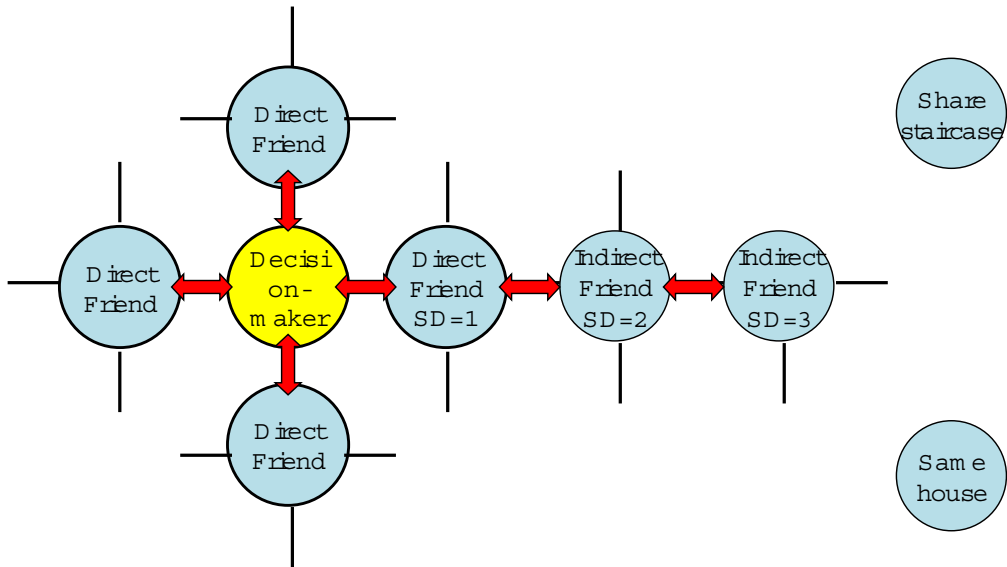
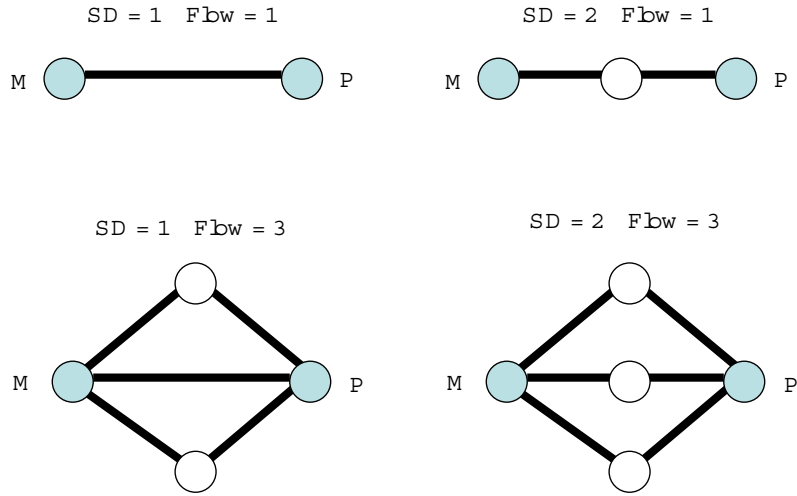


Figure 3: Examples to illustrate difference between maximum network flow and social distance



Network flow is calculated between decision-maker M and partner P under the assumption that all bilateral relationships have unit value. The examples illustrate the different information which is captured by social distance and maximum network flow, respectively: the addition of common friends will necessarily increase flow but can leave social distance unchanged.

Table 1: Testable predictions of enforced reciprocity, signaling, and preference-based reciprocity models when estimating the empirical model $\tilde{x}_{MP}^* = \eta Z + \theta x_{MP}^* + \phi V_{MP} + v_M + \epsilon_{MP}$

	Enforced Reciprocity	Signaling	Preference-based Reciprocity
Greater generosity towards friends ($\phi > 0$) when giving is <i>efficient</i>	Yes	Yes	Yes
Greater generosity towards friends ($\phi > 0$) when giving is <i>inefficient</i>	No	Yes	No
More altruistic decision-makers are relatively less generous towards friends compared to strangers under non-anonymity	Yes	No	Yes ²⁰
Maximum network flow is a better predictor of treating the partner generously under non-anonymity than social distance.	Yes	No	No

²⁰Assume a selfish and an altruistic decision-maker with the same future interaction probability $p(D_{MP})$. Compared to the selfish decision-maker, the altruist faces a repayment schedule with a flatter slope because the altruist's partner behaves less nicely by returning more tokens. At the same time, the marginal return to a returned token decreases – therefore the altruistic decision-maker will reduce his allocation of tokens.

Table 2: Summary statistics for decision-makers' actions in dictator and helping games

	Anonymous Treatment					
	SD=1	SD=2	SD=3	SD=4	SD=5	Nameless
Dictator Game	(N=206)	(N=286)	(N=312)	(N=97)	(N=4)	(N=193)
Ex. Rate 1:3	19.19 (19.64)	16.80 (19.30)	15.14 (18.79)	12.20 (15.47)	12.50 (25.00)	17.42 (18.21)
Ex. Rate 1:1	11.96 (13.53)	10.79 (12.68)	9.39 (11.89)	8.79 (10.25)	6.25 (12.50)	11.61 (12.83)
Ex. Rate 3:1	8.03 (13.55)	7.28 (12.88)	5.66 (11.10)	6.15 (10.72)	0.00 (0.00)	8.31 (13.23)
Helping Game	(N=876)	(N=149)	(N=73)	(N=181)	(N=78)	(N=776)
	12.77 (8.14)	8.97 (7.11)	7.14 (6.80)	7.68 (7.16)	7.09 (6.95)	9.52 (7.24)
	Non-anonymous Treatment					
	SD=1	SD=2	SD=3	SD=4	SD=5	Nameless
Dictator Game	(N=206)	(N=288)	(N=313)	(N=99)	(N=4)	(N=193)
Ex. Rate 1:3	24.32 (18.91)	21.67 (18.75)	19.79 (18.54)	14.80 (15.72)	37.50 (25.00)	19.87 (18.21)
Ex. Rate 1:1	16.33 (12.90)	14.62 (12.34)	13.99 (12.45)	12.16 (10.68)	18.75 (12.50)	13.98 (12.82)
Ex. Rate 3:1	10.52 (13.56)	9.88 (13.17)	9.18 (13.18)	10.15 (12.77)	0.00 (0.00)	9.62 (13.80)
Helping Game	(N=625)	(N=96)	(N=42)	(N=132)	(N=62)	
	14.54 (8.13)	11.28 (7.25)	9.26 (7.04)	8.83 (7.28)	8.11 (6.69)	

Table shows averages of number of passed tokens (dictator games) and average cutoffs (helping game) by social distance (OR-network). Standard deviations are in parentheses. Nameless refers to matches between the decision-maker and the partner where the identity of the partner is not known to the decision-maker.

Table 3: Summary statistics for partners' expectations in dictator games

	Anonymous Treatment				
	SD=1	SD=2	SD=3	SD=4	SD=5
Dictator Game	(N=262)	(N=371)	(N=401)	(N=140)	(N=2)
Ex. Rate 1:3	17.08 (15.84)	13.09 (14.22)	12.64 (14.84)	12.46 (12.83)	25.00 (14.14)
Ex. Rate 1:1	16.14 (12.06)	13.84 (11.77)	11.15 (11.30)	12.85 (11.82)	22.50 (3.54)
Ex. Rate 3:1	13.65 (14.49)	11.94 (13.86)	8.86 (12.68)	11.71 (14.34)	22.50 (3.54)

Table shows averages of number of expected tokens by social distance (OR-network). Standard deviations are in parenthesis.

Figure 4: Average number of tokens passed by the decision-maker/expected by the partner in dictator game (top/bottom) and average cutoff chosen in helping game (middle)

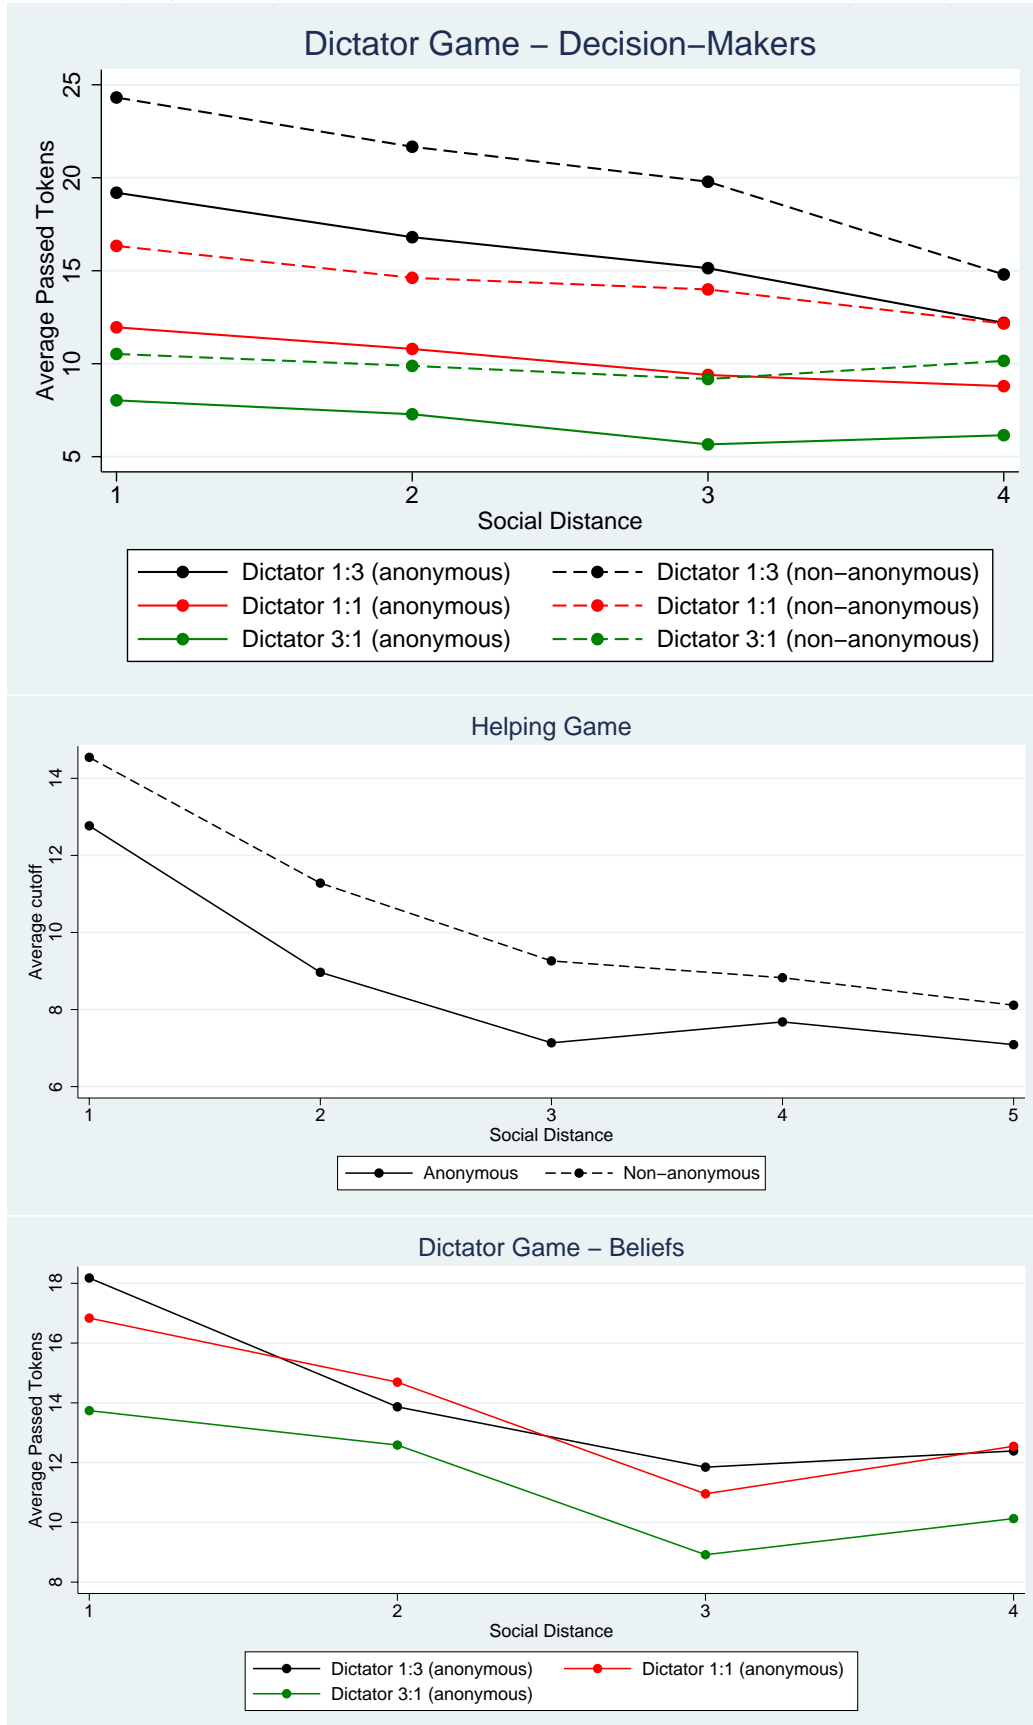


Table 4: Relative magnitudes of directed altruism and non-anonymity effects as percentages of all the decision-makers’ average nameless action and as percentages of standard deviation of nameless actions

Relative to:	Directed Altruism			Effect of Non-Anonymity			Network Flow
	SD1	SD2	SD3	SD1	SD2	SD3	
<i>Dictator Game (1:3)</i>							
Average	52	8	-8	24	18	3	19
Standard dev.	50	7	-7	23	18	3	18
<i>Dictator Game (1:1)</i>							
Average	52	16	3	21	12	9	10
Standard dev.	47	14	3	19	11	8	9
<i>Dictator Game (3:1)</i>							
Average	95	49	43	-7	8	-3	-6
Standard dev.	60	31	27	-4	5	-2	-4
<i>Helping Game</i>							
Average	88	36	12	35	13	19	30
Standard dev.	115	48	16	46	17	25	39

An “Average” row is calculated by dividing estimates for directed altruism (table 5) and the effect of non-anonymity (table 12) by the average nameless decision in the anonymous treatment (table 2, “Nameless” column). A “Standard deviation” row is calculated by dividing the same estimates by the standard deviation of the nameless decision in the anonymous treatment (table 2, “Nameless” column). For the “Network Flow” column we report the estimated effect of a one-standard deviation increase in network flow (equal to 10 units of network flow for “circle of trust” $K = 2$).

Table 5: Decision-makers' actions in the anonymous treatment (dictator and helping game) when paired with 5 partners at various social distances

	Dictator-1:3		Dictator-1:1		Dictator-3:1		Helping-Game	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
SD1	9.029 (2.331)**	9.915 (2.357)**	6.010 (1.388)**	6.244 (1.485)**	7.936 (1.935)**	8.838 (2.066)**	8.353 (0.769)**	8.045 (0.824)**
SD2	1.308 (2.304)	1.974 (2.331)	1.819 (1.365)	2.192 (1.458)	4.077 (1.886)*	4.623 (2.014)*	3.439 (0.898)**	3.337 (0.948)**
SD3	-1.340 (2.296)	-961 (2.304)	0.366 (1.361)	0.756 (1.443)	3.583 (1.887)†	4.337 (2.002)*	1.178 (1.073)	1.149 (1.177)
SD4							1.918 (0.885)*	1.451 (0.933)
Pass to Nameless (DM)		1.384 (0.136)**		1.186 (0.116)**		1.403 (0.164)**		0.564 (0.056)**
Pass to Nameless (P)								-0.039 (0.029)
Decision-maker is male		0.708 (4.547)		-2.833 (2.779)		-5.578 (4.052)		1.241 (0.822)
Partner is male		-0.651 (1.335)		-0.024 (0.838)		-0.977 (1.165)		-5.23 (0.388)
Same entryway/house		0.732 (1.376)		-0.517 (0.877)		0.381 (1.223)		0.574 (0.451)
Decision-maker is Junior		-16.356 (6.196)**		-5.507 (3.730)		-6.920 (5.365)		
Decision-maker is Senior		-10.614 (5.654)†		-5.181 (3.415)		-8.317 (4.917)†		0.475 (0.841)
Partner is Junior		0.965 (1.842)		0.802 (1.152)		1.663 (1.593)		
Partner is Senior		2.640 (1.651)		0.911 (1.046)		0.536 (1.459)		0.924 (0.467)*
Const.	4.326 (3.813)	-10.130 (5.680)†	-1.838 (2.286)	-9.253 (3.559)**	-18.845 (3.547)**	-18.679 (5.000)**	4.388 (0.84)**	-1.658 (1.175)
Obs.	901	836	901	836	901	836	1357	1193

Significance levels: † : 10% * : 5% ** : 1%
Standard errors are reported in parentheses. The dependent variable is the number of tokens passed by the decision-maker in the anonymous dictator games and the maximum cost the decision-maker is willing to pay in the helping game. Omitted distances are SD4 (dictator game) and SD5 (helping game). All specifications are estimated as Tobit regressions with decision-maker random effects. The coefficients on SD1 are significantly different from SD2 at the 5 percent level for all columns.

Table 6: Partners' expectations in the anonymous treatment of dictator game when predicting actions of five decision-makers at various social distance

	Dictator-1:3			Dictator-1:1		Dictator-3:1	
	(1)	(2)	(3)	(4)	(5)	(6)	
SD1	10.808 (1.697)**	10.970 (2.310)**	5.739 (1.399)**	6.001 (1.897)**	5.431 (2.020)**	6.291 (2.526)*	
SD2	4.774 (1.630)**	6.433 (2.230)**	3.072 (1.346)*	3.730 (1.833)*	2.777 (1.945)	4.091 (2.462)†	
SD3	1.675 (1.638)	2.550 (2.254)	-1.026 (1.352)	-0.872 (1.844)	-1.125 (1.961)	0.291 (2.468)	
Pass to Nameless (DM)	0.032 (0.035)	0.032 (0.035)	0.05 (0.043)	0.05 (0.043)	0.048 (0.057)	0.048 (0.057)	
Decision-maker is male	-1.720 (1.215)	-1.720 (1.215)	-2.377 (1.000)*	-2.377 (1.000)*	-4.311 (1.324)**	-4.311 (1.324)**	
Partner is male	3.971 (3.013)	3.971 (3.013)	-0.287 (2.079)	-0.287 (2.079)	-0.668 (2.989)	-0.668 (2.989)	
Same entryway/house	-1.842 (1.349)	-1.842 (1.349)	-0.922 (1.121)	-0.922 (1.121)	-2.344 (1.514)	-2.344 (1.514)	
Decision-maker is Junior	-0.981 (1.839)	-0.981 (1.839)	-1.398 (1.539)	-1.398 (1.539)	-1.284 (2.066)	-1.284 (2.066)	
Decision-maker is Senior	-0.099 (1.722)	-0.099 (1.722)	-1.103 (1.418)	-1.103 (1.418)	2.651 (1.911)	2.651 (1.911)	
Partner is Junior	-5.405 (4.062)	-5.405 (4.062)	-3.874 (2.811)	-3.874 (2.811)	-6.223 (4.026)	-6.223 (4.026)	
Partner is Senior	1.039 (3.589)	1.039 (3.589)	-0.985 (2.496)	-0.985 (2.496)	-4.775 (3.576)	-4.775 (3.576)	
Const.	6.148 (1.972)**	4.266 (3.798)	9.422 (1.484)**	11.933 (2.831)**	3.463 (2.204)	8.372 (3.958)*	
Obs.	855	577	855	577	855	577	

Significance levels: † : 10% * : 5% ** : 1%
Standard errors are reported in parentheses. The dependent variable is the number of tokens expected by the partner in the anonymous treatment of each dictator game. Omitted social distance is SD4. All specifications are estimated as Tobit regressions with partner random effects. The coefficients on SD1 are significantly different from SD2 at the 5 percent level for all columns.

Table 7: Accuracy of partners' beliefs

	Dictator-1:3		Dictator-1:1		Dictator-3:1	
	(1)	(2)	(3)	(4)	(5)	(6)
Pass to Nameless (DM)	0.059 (0.042)		0.071 (0.051)		0.013 (0.071)	
Pass to Nameless (DM) * SD1	-.071 (0.081)		-.001 (0.093)		0.119 (0.117)	
Pass to Partner		-.115 (0.084)		0.015 (0.079)		-.047 (0.126)
Pass to Partner * SD1		-.058 (0.131)		-.043 (0.14)		-.114 (0.205)
SD1	12.722 (2.717)**	9.535 (5.685) [†]	5.938 (2.235)**	2.701 (3.746)	6.001 (2.867)*	-3.993 (4.856)
SD2	6.538 (2.229)**	2.063 (4.555)	3.558 (1.829) [†]	-1.896 (3.090)	3.972 (2.520)	-8.441 (4.445) [†]
SD3	2.599 (2.266)	-3.194 (4.851)	-.827 (1.851)	-5.579 (3.187) [†]	-.600 (2.554)	-9.966 (4.588)*
Const.	3.192 (2.492)	11.769 (4.723)*	7.784 (1.943)**	12.288 (2.911)**	2.835 (2.726)	10.293 (4.084)*
Obs.	563	204	563	204	563	204

Significance levels: † : 10% * : 5% ** : 1%

Standard errors are reported in parentheses. The dependent variable is the number of tokens expected by the partner in the anonymous treatment for each dictator game. “Pass to Nameless (DM)” denotes the number of tokens the decision-maker passed to nameless partners and “Pass to Partner” indicates the actual generosity of the decision-maker towards the partner. Omitted social distance is SD4. All specifications are estimated as Tobit regressions with partner random effects.

Table 8: Correlation in baseline altruism among direct friends in the social network (dictator game with 1:3 exchange rate)

Nameless DG Choice	Percent of Subjects	Average # of Friends	Distribution of Friends' Types (percent)					Avg. Nameless DG Choice of Friends
			[0]	[1, 10]	[11, 15]	[16, 37]	[38, 50]	
[0]	35.23	16.75	38.65	11.04	16.56	13.19	20.55	17.02
[1, 10]	15.54	16.97	24.16	24.16	19.46	13.42	18.79	17.76
[11, 15]	13.99	17.44	37.50	20.14	9.72	10.42	22.22	17.19
[16, 37]	13.99	17.19	31.39	14.60	10.95	10.22	32.85	21.21
[38, 50]	21.24	17.83	28.39	11.86	13.56	19.07	27.12	21.39

Subjects are separated into approximate quintiles based on their dictator game (1:3) game choices for nameless partners (anonymous).

Table 9: Correlation in baseline altruism among direct friends in the social network (helping game)

Nameless HG Choice	Percent of Subjects	Average # of Friends	Distribution of Friends' Types (percent)					Avg. Nameless HG Choice of Friends
			[0]	[1, 5]	[6, 12]	[13, 15]	[16, 30]	
[0]	19.61	12.16	20.49	21.79	22.11	26.02	9.59	9.07
[1, 5]	20.39	12.36	19.88	20.77	16.02	32.34	10.98	9.66
[6, 12]	19.35	12.72	20.24	16.07	20.54	30.95	12.20	9.87
[13, 15]	30.13	12.27	15.84	21.58	20.59	31.49	10.50	9.83
[16, 30]	10.52	12.36	15.57	19.53	21.64	27.97	15.30	10.34

Subjects are separated into approximate quintiles based on their helping game choices for nameless partners (anonymous).

Table 10: Regressing subjects' baseline altruism on friends' average baseline altruism

	Dictator Game (1:3)		Helping-Game	
	(1)	(2)	(3)	(4)
Avg. Nameless Decision of Friends	0.266 (0.17)	0.267 (0.17)	0.154 (0.78)*	0.143 (0.78)†
Subject is Male		-0.688 (4.09)		-0.308 (0.60)
Percent Male of Subject's Friends		-0.370 (6.16)		-1.555 (0.96)
Const.	6.445 (3.82)†	6.929 (4.88)	7.115 (0.83)**	7.806 (0.90)**
Obs.	186	186	746	746

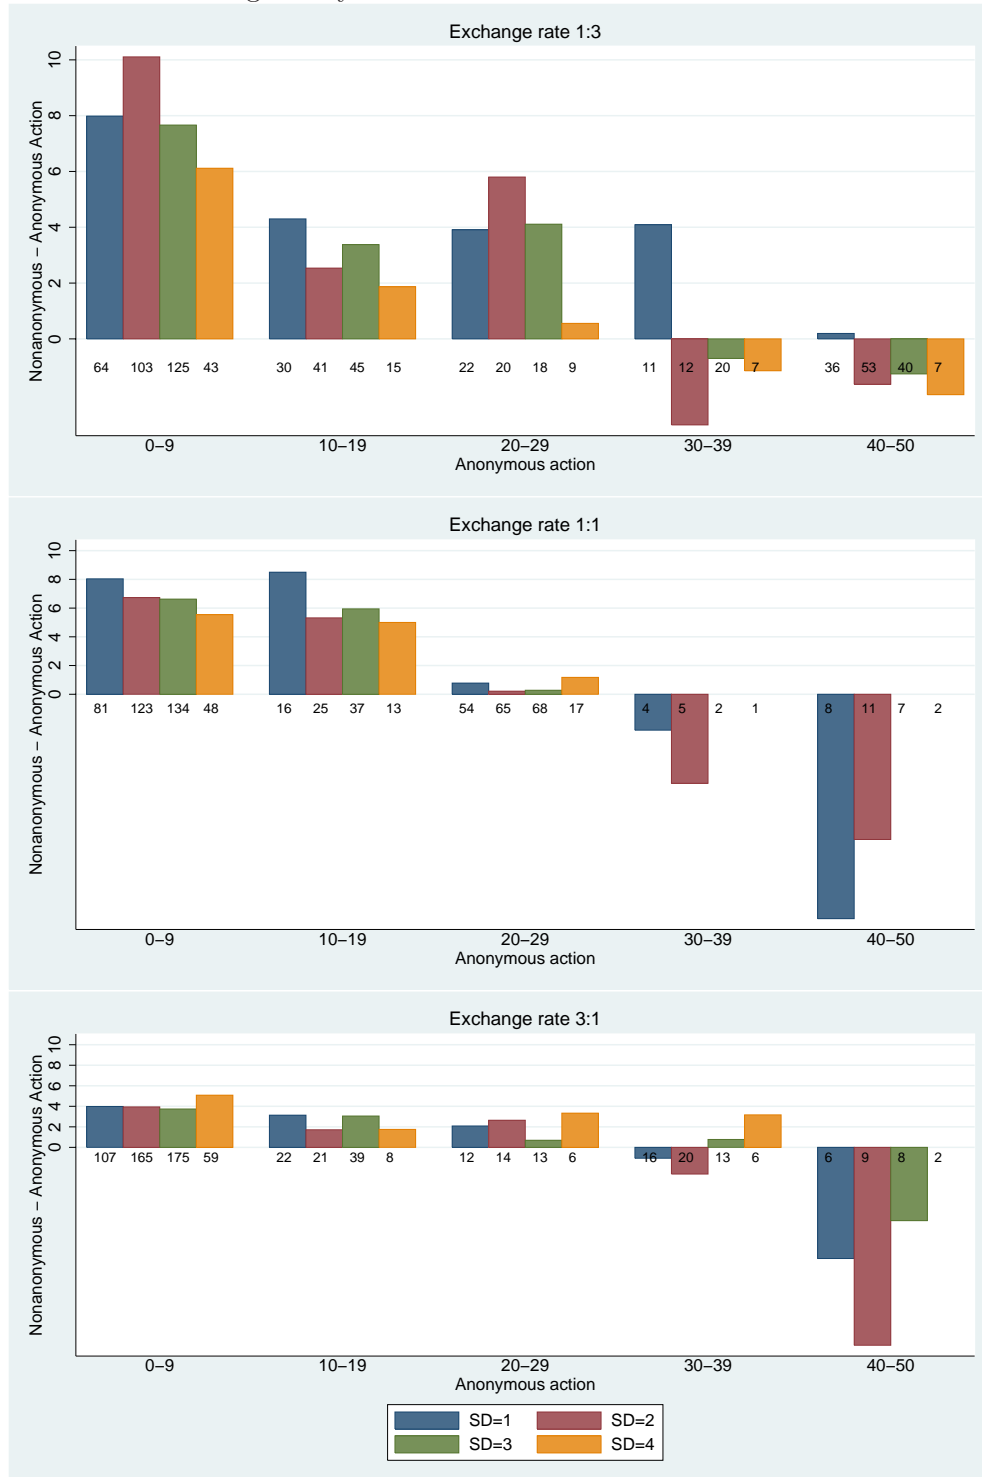
Standard errors are reported in parentheses. Tobit regressions of subjects' baseline altruism (measured by their nameless anonymous decision) on direct friends' average baseline altruism. All possible direct friends pairs are included.

Table 11: Regressing average allocation to partners whose direct friends made decisions for them (anonymous treatment only) on their own baseline altruism and the average baseline altruism of their friends (helping game)

	Helping-Game	
	(1)	(2)
Partner's Nameless Decision $\in [1, 5]$	1.048 (1.02)	0.357 (0.96)
Partner's Nameless Decision $\in [6, 12]$	3.074 (0.93)**	0.545 (0.91)
Partner's Nameless Decision $\in [13, 15]$	4.567 (0.92)**	0.521 (0.98)
Partner's Nameless Decision $\in [16, 30]$	5.275 (1.17)**	-0.384 (1.27)
Decision-Maker's Nameless Choice		-0.474 (0.054)**
Const.	9.747** (0.73)**	7.679** (0.73)**
Obs.	549	549

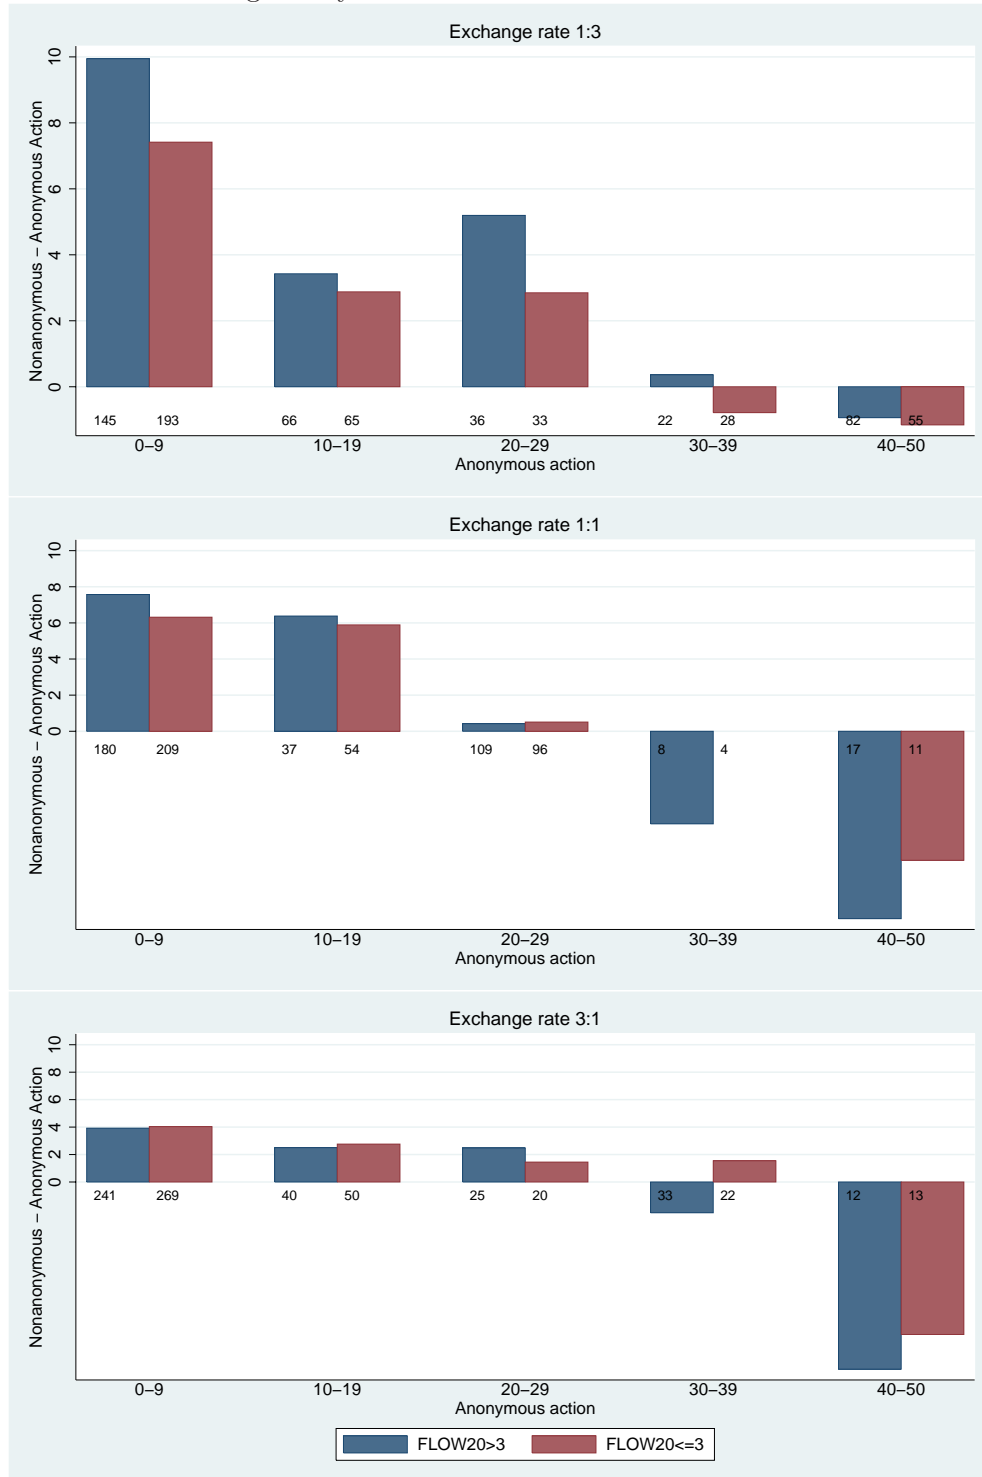
Standard errors are reported in parentheses. The dependent variable is a partner's average allocation in anonymous treatment from decisions made by friends.

Figure 5: Difference between number of passed tokens in the non-anonymous and anonymous treatments in the dictator game by social distance



For each decision-maker/partner pair, the difference between the number of tokens allocated in the non-anonymous and the anonymous treatments was calculated. Bars show average difference grouped by the decision-maker's contribution level in the anonymous treatment and by social distance.

Figure 6: Difference between number of passed tokens in the non-anonymous and anonymous treatments in the dictator game by network flow



For each decision-maker/partner pair, the difference between the number of tokens allocated in the non-anonymous and the anonymous treatments was calculated. Bars show average difference grouped by the decision-maker's contribution level in the anonymous treatment and by network flow for circle of trust $K = 2$ (median network flow equals 3).

Table 12: Comparing non-anonymous and anonymous treatments (dictator and helping games)

	Dictator-1:3			Dictator-1:1			Dictator-3:1			Helping-Game		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Anonymous Action	0.822 (0.047)**	0.81 (0.048)**	0.814 (0.048)**	0.508 (0.05)**	0.511 (0.05)**	0.508 (0.051)**	0.874 (0.064)**	0.874 (0.063)**	0.874 (0.064)**	0.582 (0.648)	0.676 (0.116)**	0.652 (0.644)
SD1	4.184 (1.381)**	1.373 (2.675)	1.373 (2.675)	2.470 (1.175)*	2.498 (2.168)	2.498 (2.168)	-555 (1.556)	2.815 (3.018)	2.815 (3.018)	3.333 (4.618)	3.333 (4.618)	0.521 (4.665)
SD2	3.196 (1.346)*	1.591 (1.876)	1.591 (1.876)	1.445 (1.141)	1.461 (1.540)	1.461 (1.540)	0.689 (1.503)	2.604 (2.100)	2.604 (2.100)	1.214 (1.679)	1.214 (1.679)	0.193 (1.696)
SD3	0.513 (1.335)	0.451 (1.334)	0.451 (1.334)	1.030 (1.131)	1.031 (1.132)	1.031 (1.132)	-208 (1.487)	-148 (1.486)	-148 (1.486)	1.809 (1.105)	1.809 (1.105)	1.850 (1.098)†
Network Flow		0.323 (0.073)**	0.234 (0.191)		0.115 (0.061)†	-0.002 (0.152)		-0.048 (0.082)	-0.278 (0.214)	0.281 (0.085)**	0.281 (0.085)**	0.281 (0.086)**
Pass to Nameless (DM)	0.544 (0.08)**	0.549 (0.08)**	0.547 (0.08)**	0.193 (0.055)**	0.192 (0.055)**	0.193 (0.055)**	-0.038 (0.074)	-0.037 (0.073)	-0.036 (0.074)	0.335 (0.382)	0.278 (0.092)**	0.294 (0.379)
Pass to Nameless (P)										0.007 (0.041)	0.018 (0.032)	0.017 (0.041)
Decision-maker is Junior	-1.701 (3.481)	-1.611 (3.479)	-1.651 (3.480)	-5.070 (2.641)†	-4.945 (2.634)†	-5.071 (2.641)†	-7.616 (3.503)*	-7.665 (3.489)*	-7.745 (3.503)*			
Decision-maker is Senior	-957 (3.239)	-878 (3.237)	-936 (3.237)	-5.214 (2.462)*	-5.129 (2.457)*	-5.214 (2.462)*	-4.262 (3.221)	-4.239 (3.209)	-4.328 (3.220)	-0.21 (0.937)	-148 (0.9)	-129 (0.933)
Partner is Junior	1.712 (1.058)	1.615 (1.054)	1.645 (1.058)	-7.758 (0.891)	-814 (0.891)	-757 (0.891)	0.33 (1.204)	0.366 (1.200)	0.391 (1.203)			
Partner is Senior	0.212 (0.97)	0.129 (0.962)	0.2 (0.969)	-733 (0.817)	-815 (0.815)	-733 (0.817)	1.152 (1.096)	1.128 (1.089)	1.174 (1.095)	-227 (0.76)	-384 (0.51)	-330 (0.755)
Const.	-4.608 (3.116)	-3.961 (2.943)	-4.516 (3.115)	5.447 (2.383)*	6.227 (2.221)**	5.446 (2.383)*	0.009 (3.165)	0.276 (2.970)	-0.32 (3.163)	0.944 (1.014)	1.321 (0.961)	0.952 (1.008)
Obs.	836	836	836	836	836	836	836	836	836	955	955	955

Significance levels: † : 10% * : 5% ** : 1%

Standard errors are reported in parentheses. The dependent variable is the number of tokens passed by the decision-maker to a specific partner in the non-anonymous dictator games and the maximum cost the decision-maker is willing to pay in the non-anonymous helping game when matched with a specific partner. All specifications are estimated as Tobit regressions with decision-maker random effects. “Anonymous Action” denotes the decision-maker’s action for the specific partner in the anonymous treatment. Because the helping game has a between-group design, we first predict the decision-maker’s action in the anonymous treatment by running an auxiliary Tobit regression with data from the anonymous treatment, controlling for social distance, nameless decision, and class dummies. Omitted social distance dummies are SD4 and SD5. Network flow is calculated for a circle of trust $K = 2$.

Table 13: Probit estimates (marginal effects reported) for decision-makers passing *strictly more* tokens in the non-anonymous versus anonymous treatments (dictator game only)

	Dictator-1:3			Dictator-1:1			Dictator-3:1		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Anonymous Action	-0.057 (0.011)**	-0.059 (0.011)**	-0.060 (0.011)**	-0.116 (0.017)**	-0.114 (0.017)**	-0.115 (0.017)**	-0.048 (0.015)**	-0.051 (0.015)**	-0.048 (0.015)**
SD1	0.908 (0.31)**		-0.088 (0.606)	0.085 (0.3)		0.439 (0.563)	-0.046 (0.298)		-0.037 (0.563)
SD2	0.373 (0.3)		-0.196 (0.426)	-0.068 (0.294)		0.134 (0.4)	0.113 (0.291)		0.117 (0.4)
SD3	0.152 (0.297)		0.148 (0.298)	-0.321 (0.29)		-0.318 (0.29)	-0.462 (0.288)		-0.462 (0.288)
Network Flow		0.065 (0.016)**	0.084 (0.044)†		0.019 (0.016)	-0.030 (0.04)		0.025 (0.016)	-0.0007 (0.04)
Pass to Nameless (DM)	0.046 (0.014)**	0.046 (0.014)**	0.047 (0.014)**	0.016 (0.014)	0.015 (0.014)	0.016 (0.014)	-0.016 (0.016)	-0.016 (0.016)	-0.016 (0.016)
Decision-maker is Junior	-0.778 (0.603)	-0.755 (0.605)	-0.761 (0.606)	-1.167 (0.652)†	-1.166 (0.645)†	-1.173 (0.654)†	-1.677 (0.784)*	-1.706 (0.778)*	-1.677 (0.784)*
Decision-maker is Senior	-1.095 (0.567)†	-1.084 (0.569)†	-1.076 (0.57)†	-1.721 (0.625)**	-1.718 (0.617)**	-1.728 (0.626)**	-1.036 (0.7)	-1.048 (0.693)	-1.036 (0.7)
Partner is Junior	0.357 (0.228)	0.331 (0.229)	0.332 (0.23)	0.011 (0.232)	-0.002 (0.229)	0.012 (0.232)	0.291 (0.248)	0.303 (0.243)	0.292 (0.249)
Partner is Senior	-0.021 (0.21)	-0.022 (0.21)	-0.034 (0.212)	-0.242 (0.213)	-0.221 (0.21)	-0.242 (0.214)	0.301 (0.222)	0.313 (0.216)	0.301 (0.222)
Const.	-0.532 (0.561)	-0.461 (0.511)	-0.513 (0.565)	1.134 (0.589)†	0.931 (0.538)†	1.126 (0.59)†	0.096 (0.662)	-0.118 (0.624)	0.095 (0.662)
Obs.	714	714	714	819	819	819	818	818	818

Significance levels: † : 10% * : 5% ** : 1%

Standard errors are reported in parentheses. The dependent variable is a dummy variable equal to 1 if the number of tokens passed to a specific partner in the non-anonymous treatment is *strictly greater* than the number of tokens passed to the same partner in the anonymous treatment, and zero otherwise. All specifications are estimated as Probit regressions with decision-maker random effects. “Anonymous Action” denotes the decision-maker’s action for the specific partner in the anonymous treatment. Omitted social distance dummies are SD4 and SD5. Network flow is calculated for a circle of trust $K = 2$. Data points where the decision-maker already passes the maximum number of tokens in the anonymous treatment are excluded.

Table 14: Probit estimates (marginal effects reported) for decision-makers passing *strictly fewer* tokens in the non-anonymous versus anonymous treatments (dictator game only)

	Dictator-1:3			Dictator-1:1			Dictator-3:1		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Anonymous Pass	0.054 (0.019)**	0.053 (0.019)**	0.053 (0.02)**	0.049 (0.017)**	0.045 (0.017)**	0.055 (0.02)**	0.013 (0.014)	0.013 (0.014)	0.014 (0.014)
SD1	-0.977 (0.55)†		-1.607 (0.976)†	0.279 (0.563)		-1.245 (0.984)	0.467 (0.533)		-0.754 (0.94)
SD2	-0.309 (0.493)		-0.667 (0.674)	0.852 (0.535)		-0.09 (0.711)	0.322 (0.529)		-0.399 (0.712)
SD3	-0.225 (0.47)		-0.225 (0.472)	0.358 (0.546)		0.381 (0.576)	0.073 (0.528)		0.071 (0.528)
Network Flow		-0.037 (0.028)	0.051 (0.064)		0.029 (0.024)	0.124 (0.063)†		0.044 (0.025)†	0.097 (0.061)
Pass to Nameless (DM)	-0.028 (0.02)	-0.027 (0.02)	-0.027 (0.02)	-0.037 (0.015)*	-0.036 (0.015)*	-0.040 (0.017)*	-0.0009 (0.016)	-0.0009 (0.016)	-0.0009 (0.015)
Decision-maker is Junior	-1.717 (0.925)†	-1.818 (0.922)*	-1.694 (0.932)†	0.475 (0.574)	0.492 (0.566)	0.656 (0.644)	-0.383 (0.634)	-0.345 (0.628)	-0.274 (0.626)
Decision-maker is Senior	-1.021 (0.674)	-1.081 (0.67)	-1.029 (0.68)	-0.094 (0.522)	-0.075 (0.516)	-0.107 (0.578)	-0.453 (0.526)	-0.465 (0.525)	-0.447 (0.521)
Partner is Junior	-0.030 (0.422)	0.057 (0.405)	-0.039 (0.426)	-0.570 (0.41)	-0.583 (0.401)	-0.669 (0.451)	-0.058 (0.417)	-0.033 (0.422)	-0.038 (0.423)
Partner is Senior	0.046 (0.39)	0.131 (0.365)	0.057 (0.392)	-0.280 (0.355)	-0.289 (0.342)	-0.305 (0.38)	-0.100 (0.354)	-0.066 (0.357)	-0.045 (0.359)
Const.	-2.616 (0.871)**	-2.817 (0.781)**	-2.634 (0.877)**	-2.598 (0.782)**	-2.213 (0.6)**	-2.849 (0.878)**	-1.845 (0.72)*	-1.864 (0.618)**	-1.894 (0.713)**
Obs.	496	496	496	432	432	432	294	294	294

Significance levels: † : 10% * : 5% ** : 1%

Standard errors are reported in parentheses. The dependent variable is a dummy variable which equals 1 if the number of tokens passed to a specific partner in the non-anonymous treatment is *strictly smaller* than the number of tokens passed to the same partner in the anonymous treatment, and zero otherwise. All specifications are estimated as Probit regressions with decision-maker random effects. “Anonymous Action” denotes the decision-maker’s action for the specific partner in the anonymous treatment. Omitted social distance dummies are SD4 and SD5. Network flow is calculated for a circle of trust $K = 2$. Data points where the decision-maker already passes the minimum number of tokens in the anonymous treatment are excluded.

Table 15: Effects of “average time spent per week” and network flow on decision-makers’ generosity towards direct friends under non-anonymity (dictator game only)

	Dictator-1:3	Dictator-1:1	Dictator-3:1
	(1)	(2)	(3)
Anonymous Pass	0.25 (0.063)**	0.238 (0.211)	0.474 (0.115)**
Network Flow	0.676 (0.306)*	-.024 (0.395)	0.155 (0.262)
Average Time Spent per Week	-.328 (0.368)	-.089 (0.488)	-.060 (0.323)
Const.	12.131 (3.640)**	13.946 (5.493)*	4.998 (3.286)
Obs.	206	206	206

Significance levels: † : 10% * : 5% ** : 1%

Standard errors are reported in parentheses. The dependent variable is the number of tokens passed by the decision-maker to a *direct friend* in the non-anonymous dictator games. All specifications are estimated as Tobit regressions with decision-maker random effects. “Anonymous Action” denotes the decision-maker’s action for the specific partner in the anonymous treatment. Network flow is calculated for a circle of trust $K = 2$. Average time spent per week is a categorical variable which takes the values 0 (less than half an hour per week), 1 (30 min to 1 hour), 2 (1 hour to 2 hours), 3 (2 hours to 4 hours), 4 (4 hours to 8 hours) and 5 (more than 8 hours a week).