

Directed Altruism and Enforced Reciprocity in Social Networks*

Stephen Leider
University of Michigan

Markus M. Möbius
Harvard University and NBER

Tanya Rosenblat
Iowa State University

Quoc-Anh Do
Singapore Management University

June 22, 2009

Abstract

We conduct online field experiments in large real-world social networks in order to decompose prosocial giving into three components: (1) *baseline altruism* toward randomly selected strangers, (2) *directed altruism* that favors friends over random strangers, and (3) giving motivated by the *prospect of future interaction*. Directed altruism increases giving to friends by 52 percent relative to random strangers, while future interaction effects increase giving by an additional 24 percent when giving is socially efficient. This finding suggests that future interaction affects giving through a repeated game mechanism where agents can be rewarded for granting efficiency-enhancing favors. We also find that subjects with higher baseline altruism have friends with higher baseline altruism.

JEL Classification: C73, C91, D64

Keywords: modified dictator games, directed altruism, enforced reciprocity, social networks

*This paper replaces an earlier working paper version entitled “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, Tyler Williams and Kate Ambler 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).

1 Introduction

Real world social networks provide a natural laboratory to study prosocial behavior. Friends help each other frequently and often substantially. In the 1995 General Social Survey, 55 percent of Americans report they first approach close friends and family members when they need to borrow a large sum of money (Mobius and Szeidl 2007). In a 2007 UK YouGov survey, 48 percent of respondents reported lending, on average, \$1,800 to friends and relatives during the past 12 months (YouGov 2007). Similarly, close friends and relatives are the predominant source of informal insurance against risk in developing countries (Townsend 1994, Udry 1994).

Economists have explored two main explanations of prosocial behavior – selfless altruism and expectations of reciprocal favors in future exchanges. Since people tend to interact frequently with friends toward whom they have the strongest altruistic feelings, these mechanisms are difficult to distinguish empirically.¹ In this paper, we use two online field experiments in a real world social network to solve this identification problem. Our design allows us to selectively switch off the reciprocity mechanism in some treatments, and thereby separately measure the strength of the altruism and reciprocity mechanisms.

The ability of our methodology to distinguish between these two mechanisms has several applications. For example, many researchers have documented that as societies transition to market economies, market transactions replace bilateral reciprocal exchanges, leading to a decline in social capital². By using our diagnostic games to identify cultures where the social ties are either heavily reciprocity-based or heavily altruism-based, economists may be better able to explain differences in the social disruption caused by markets. Additionally, our finding (described below) that the reciprocal mechanism only applies to transactions that are efficient suggests a policy use of our techniques. Development programs like microfinance that use social ties to encourage lending and insurance may be most effective in communities where the social networks have a strong reciprocity component, because relationships based on trading efficiency-enhancing favors may help direct resources towards efficient uses. Thus, policymakers could use experiments like ours to identify areas where intervention may be most effective.

In our experiments, we distinguish three components of prosocial giving: (1) *baseline altruism* toward randomly selected strangers, (2) *directed altruism* that favors friends over random strangers, and (3) giving motivated by the *prospect of future interaction*. We begin by directly measuring the social network of Harvard undergraduates to identify, for each subject, socially close direct friends, less close friends-of-friends, and socially distant strangers. We then conduct a series of online games where subjects make unilateral allocation decisions for several types of *named* partners and

¹Economists and sociologists use both frequency of interaction (Marmaros and Sacerdote 2006) and intensity of altruism interchangeably to measure the strength of social connections (Granovetter 1974, Granovetter 1985, Marsden and Campbell. 1984, Mitchell 1987, Perlman and Fehr 1987, Mathews, White, Soper and von Bergen 1998).

²See for example Coleman (1993) on the industrial revolution in Western societies, Völker (1995) on the change from communism to capitalism in Eastern Germany, and Yellen (1990) on the increased participation of the Kalahari !Kung in markets.

one *nameless* partner (a randomly selected participant from the subject’s dormitory). The games are either modified dictator games, as in Andreoni and Miller (2002), or a new “helping game”. Subjects make multiple decisions but are paid for one decision selected at random. For some decisions, neither participant is informed which of the decision-maker’s choices was implemented. In these *anonymous* decisions, the difference in allocations between friends and strangers allows us to quantify the magnitude of directed altruism. In particular, we find that subjects send on average 52 percent more money to close friends than to strangers. We also demonstrate that giving to friends is strongly related to the decision-makers’ underlying baseline altruism toward nameless partners.

For other decisions, both the decision-maker and partner are informed which of the decision-maker’s choices was selected for payment. We use the difference between this *non-anonymous* and the anonymous treatment to measure the separate effect of future interaction on prosocial behavior. We find that the nature of the future interaction effect depends crucially on the social welfare effects of prosocial behavior. When giving *increases* joint surplus then subjects increase giving to friends (relative to strangers) by an additional 24 percent in the non-anonymous treatment. Thus, directed altruism is roughly twice as strong as future interaction effects in determining giving behavior. When giving *decreases* joint surplus, then subjects do not give more to friends in the non-anonymous treatment. In contrast, the directed altruism effect favors friends over strangers, both when prosocial behavior is socially efficient and when it is inefficient.

These differential effects of future interaction on prosocial behavior are well-explained by the theory of repeated games. Mobius and Szeidl (2007) develop a tractable theory for analyzing repeated games on social networks which we adapt to our setting and call the *enforced reciprocity model*. In that model, a decision-maker can safely grant favors (in the form of larger allocations) to partners when the relationship between them is stronger and more valuable, since the partner would rather repay the favor than damage the friendship. Granting favors, however, only benefits both the decision-maker and partner when giving increases social surplus. This conflicts with Benabou and Tirole’s (2006) model, where individuals give in order to signal being an altruistic type. The signaling model predicts that excess giving to friends should increase for all the games (and all exchange rates), since giving still distinguishes altruists from selfish types. The enforced reciprocity model also predicts that giving should increase if the decision-maker and partner share a greater number of common friends (controlling for social distance). We confirm this prediction in our data. This provides further evidence for the enforced reciprocity model and distinguishes it from related theories such as preference-based reciprocity (Rabin 1993, Dufwenberg and Kirchsteiger 2004). This result also highlights the influence of second-order links and network structure on economic decisions.

Lastly, we show that decision-makers who exhibit greater baseline altruism are treated more generously by their friends. However, we show that friends do not reward intrinsic kindness, but rather, that kind people tend to have friends who exhibit greater baseline altruism themselves.

Our paper builds on a rich experimental literature on other-regarding preferences and cooperation. Altruistic behavior toward strangers has been observed in a variety of laboratory contexts

(see Camerer (2003) for an extensive survey). Hoffman, McCabe and Smith (1996) suggests that a decrease in perceived “social distance” increases donations in dictator games. Our paper reduces social distance in two ways: (a) we distinguish between giving to friends versus giving to strangers; and (b) we either reveal or do not reveal which decision is implemented. Our online experimental design is a significant methodological advance because it provides a practical way to match subjects with their real-world friends.³ Lab experiments that do relax traditional anonymous matching of strangers typically reveal demographic characteristics such as gender or ethnicity (for example, see Fershtman and Gneezy (2001)). Very few lab experiments explicitly rely on subjects’ ongoing relationships with their friends as we do in our non-anonymous treatment. Instead repeated interactions are incorporated directly into the experimental design, increasing prosocial behavior (see Roth and Murnighan (1978) and Murnighan and Roth (1983) for early examples). A notable exception is the seminal work of Glaeser, Laibson, Scheinkman and Soutter (2000), who match subjects at various social distances in a trust game.⁴ In subsequent research, Goeree, McConnell, Mitchell, Tromp and Yariv (2008) 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 data from experiments with European university students). To the best of our knowledge, our design is the first to distinguish between directed altruism and future interaction effects in social networks.

The rest of the paper is organized as follows. Section 2 describes several relevant theories: enforced reciprocity, signaling, and preference-based reciprocity; and reviews their testable implications. Section 3 presents the experimental design. Section 4 summarizes the main features of the data. Our empirical results on directed altruism are presented in section 5. Section 6 analyzes decisions under non-anonymity. In section 7, we consider two benchmarks comparing the importance and magnitude of the directed altruism and future interaction effects. Section 8 shows that friends tend to have similar levels of baseline altruism. Section 9 concludes by discussing the implications of our results for the theoretical and empirical analysis of prosocial behavior.

2 Theoretical Framework

We briefly review the predictions of three different theories that explain decision-making when there is a prospect of future interaction. We refer the interested reader to the earlier NBER working paper for a detailed formal treatment of each theory (Leider, Mobius, Rosenblat and Do 2007). We assume throughout that there is a decision-maker, M and a partner, P , who are embedded in a social network. 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 of distance 2 (see Figure 1). The decision-maker’s allocation decisions

³Our experimental design fits the definition of an *artefactual field experiment* from Harrison and List (2004).

⁴See also Polzer, Neale and Glenn (1993) comparing allocations to friends and strangers in an ultimatum game.

in the anonymous treatment (neither decision-maker nor partner learn about the decision) and the non-anonymous treatment (both agents learn about the decision) are denoted with x_{MP} and \tilde{x}_{MP} , respectively (coded so that a larger x implies greater prosocial behavior).⁵

We assume, as a starting point, that decision-makers have altruistic preferences whose strength varies with social distance. We use a simple linear specification to capture decision-making in the anonymous treatment, that is a natural extension of existing preferences-based altruism models:⁶

$$x_{MP} = \alpha Z_{MP} + \gamma_1 D_{MP} + \gamma_M + \epsilon_{MP}, \quad (1)$$

We control for observable demographic characteristics of both individuals (Z_{MP}). The coefficient γ_M captures the decision-maker’s intrinsic generosity toward all partners (independent of social distance), which we refer to as her baseline altruism. The coefficient γ_1 determines how the decision-maker’s altruism varies with social distance, which we call directed altruism.

2.1 Enforced Reciprocity

Repeated games provide a natural framework to analyze the decision-maker’s allocation, \tilde{x}_{MP} , under the prospect of future interactions. However, repeated games typically admit many equilibria even in a two-person setting – this multiplicity problem is compounded for social networks where hundreds of agents often interact. Mobius and Szeidl (2007) provide a tractable framework for modeling repeated game effects in social networks. They assume that decision-maker and partner share a relationship that is “consumed” in the future and gives both of them utility V_{MP} . A decision-maker who is more prosocial under non-anonymity compared to anonymity, such that $\tilde{x}_{MP} - x_{MP} > 0$, grants a “favor” to the partner and can use the value of the relationship to the partner as “social collateral” to enforce repayment of the favor. If the partner refuses to return the favor their relationship breaks down and both agents lose the value of their friendship.

In our working paper, we formally show that in the unique subgame perfect equilibrium the decision-maker allocates a larger amount under non-anonymity compared to anonymity *only* when giving is *efficient*, and thus increases social surplus. In this case, both the decision-maker and the partner keep some of the surplus created by the favor. In contrast, when giving is inefficient a favor that has to be repaid makes both the decision-maker and the partner worse off.

We can also show that the difference in allocations between the non-anonymous and anonymous treatments, $\tilde{x}_{MP} - x_{MP}$, is increasing in the value of the relationship V_{MP} because the decision-maker is willing to grant, and the partner is willing to repay, larger favors. This motivates the

⁵Therefore, the decision-maker’s utility is decreasing in x and the partner’s utility is increasing in x .

⁶Andreoni (1990) model altruism as “warm glow”, while Fehr and Schmidt (1999), Bolton and Ockenfels (2000), and Charness and Rabin (2002) focus on preferences over payoff distributions.

following empirical model of allocation decisions under non-anonymity:

$$\tilde{x}_{MP} = \eta Z_{MP} + \theta x_{MP} + \phi V_{MP} + v_M + \epsilon_{MP}. \quad (2)$$

We include the decision-maker’s anonymous decision, x_{MP} , as a covariate, because our enforced reciprocity model predicts deviations from the anonymous benchmark. The model also predicts that enforced reciprocity and altruism are substitutes ($\theta < 1$): since the decision-maker will give the largest enforceable favor (when giving is efficient) and the partner will retain some of the surplus, the decision-maker’s marginal utility from giving an additional altruistic “gift” is smaller compared to the anonymous case (the partner is now “richer” than in the anonymous case).

We consider two proxies for the value of the relationship, V_{MP} . First, we expect that relationship value decreases with social distance because a decision-maker is less likely to interact with a socially distant partner in the future. Second, the enforced reciprocity model of Mobius and Szeidl (2007) suggests *maximum network flow* as an alternative measure. The maximum flow counts the number of distinct paths between decision-maker and partner (see Figure 1). Network flow captures structural features of the social network that are not captured by social distance alone. For example, having more friends in common increases network flow, while social distance is unchanged. Network flow formalizes a common intuition in the sociology literature that “dense” networks strengthen trust by facilitating informal arrangements (Coleman 1988, Coleman 1990). In Mobius and Szeidl’s (2007) enforced reciprocity model each common friend increases the “social collateral” between the decision-maker and the partner. The decision-maker is able to extract repayments for larger favors, since if the partner defaults he will lose his relationship with all of the common friends as well as the decision-maker. Throughout, we calculate network flow by only including links that are at most a distance $K = 2$ away from the decision-maker. Mobius and Szeidl (2007) justify this choice of a “circle of trust” of $K = 2$ by reviewing some of the existing empirical literature in economics and sociology, such as Granovetter (1974).

Table 1 summarizes the main predictions of the enforced reciprocity model (as well as the signaling and preference based-reciprocity models discussed below).

2.2 Signaling

In recent work, Benabou and Tirole (2006) propose a signaling model that provides an alternative theory for explaining greater generosity to friends under non-anonymity. In their framework, agents care about being perceived as an altruistic (rather than greedy) type so they act more generously when their actions can be observed. Moreover, it is reasonable to extend their model to assume that individuals care more about signaling generosity to friends than to strangers, because they are more likely to interact with friends in the future.

A key distinguishing prediction of the signaling model is that decision-makers should increase their allocation decision under non-anonymity compared to anonymity *both* when prosocial behavior

is efficient *and* when it is inefficient. Larger allocations are just as good (if not better) signals of generosity when giving is inefficient than when giving is efficient. Additionally, excess giving under non-anonymity is independent of the level of altruism in Benabou and Tirole’s (2006) model because signaling utility and distributional utility are additively separable. This effect contrasts to the enforced reciprocity model, where altruism and favors are substitutes.

2.3 Preference-Based Reciprocity

A third possible mechanism for future interaction effects is preference-based reciprocity: Dufwenberg and Kirchsteiger (2004) develop a psychological game theory model of sequential reciprocity, where an individual treats kindly (unkindly) those who have treated/will treat her kindly (unkindly). As in enforced reciprocity, the decision-maker increases her allocation between non-anonymity and anonymity only when giving is efficient, anticipating that with some probability, assumed to be decreasing in social distance, the partner may act to benefit the decision-maker.

Under preference-based reciprocity, however, the partner’s desire to return the decision-maker’s favor is intrinsic, rather than designed to preserve the relationship with the decision-maker or common friends. Therefore, unlike the enforced reciprocity model, preference-based reciprocity would not predict that the network flow measure independently correlates with increased generosity after having controlled for the frequency of future interaction (via social distance).

3 Experimental Design

Figure 2 presents the major features of our experimental design. In order to recruit more subjects and map a larger social network, all communication with subjects was conducted by email and all choices were submitted on a website (rather than in a lab) that subjects could access with a password through their own web browsers. We conducted two waves of the experiment. In each wave, we first used a novel task to elicit the social network truthfully. We then had subjects play an allocation game with other subjects in the network. In Wave 1, we used modified dictator games with varying exchange rates, and in Wave 2, we used a new helping game. In both waves, decision-makers first decided on allocations between themselves and other unnamed (“nameless”) partners, and then, a few days later, made several allocations between themselves and named partners (identified by real first and last name) at various social distances. For both waves, a single decision was randomly selected for payment and all players were informed of their earnings by email.

We conducted two major information treatments using a within-subject design in Wave 1 and a between-subject design in Wave 2. In these treatments we varied how much information was revealed to both parties at the end of the experiment. In the “anonymous” treatment, each subject was merely informed of their total payoff for the experiment - neither decision-makers nor partners found out which particular decision was implemented. In the “non-anonymous” treatment, both the

decision-maker and partner were told both their total payoff and which decision was implemented.

3.1 Design for Wave 1: Coordination Task and Dictator Games

Network Elicitation. To measure the social network, we used a *coordination task* to provide subjects incentives to truthfully report their friendships. Each subject listed her 10 best friends and the average amount of time per week she spends with each of them.⁷ Each time a listed friend also listed the subject, the subject was paid 50 cents, with probability 0.5 if their answers about time spent together disagreed, or with probability 0.75 if they agreed. We made the expected payoff (25 or 37.50 cents) large enough to give an incentive to list their friends truthfully, and small enough to discourage “gaming”. The randomization was included to limit disappointment if a subject was only named by a few people. To define the social network, we say that two subjects have a direct link if at least one of them named the other person. We call this type of social network the “OR-network”.⁸

Allocation game. After measuring the social network, we randomly assigned each subject the role of decision-maker or partner in the dictator games.⁹ Each decision-maker then played modified dictator games with a *nameless* partner randomly selected from the decision-maker’s dormitory.¹⁰ She was asked to make allocation decisions for both the anonymous and the non-anonymous conditions. The decision-maker allocated 50 tokens between herself and the recipient. However, the value of a token to the decision-maker and to the recipient differed between games. In one case a token was worth 10 cents to the decision-maker, and 30 cents to the recipient (“1:3”), while in the second case they were worth 20 cents to the decision-maker and 20 cents to the recipient (“1:1”), and in a third case they were worth 30 cents to the decision-maker and 10 cents to the recipient (“3:1”). We can therefore observe allocations when giving is efficient, neutral and inefficient. Each decision-maker made allocation choices for each of the three token-money exchange rates.

A few days later, decision-makers played in a second round, in which they were matched in random order with five different *named* partners listed using their real first and last 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 ($SD \geq 3$), and (5) a randomly selected student from the same dormitory who falls into none of the above categories.¹¹ Once again, the decision-maker was asked to make allocation decisions in both the anonymous and non-anonymous situations. The decision-maker made allocation decisions

⁷The choices were 0-30 minutes, 30 minutes to 1 hour, 1-2 hours, 2-4 hours, 4-8 hours or more than 8 hours.

⁸We find similar results using the “AND-network”, where a link exists only if both subjects name each other. The OR-network has desirable monotonicity properties: a subject with an above average number of actual friends will have an above average number of friends in the measured network even if the survey truncates his true network. This is not always true for the AND-network if truncation forces subjects to randomly select from a set of equally close friends.

⁹In the experimental instructions, we referred to two roles simply as player 1 and player 2.

¹⁰The decision-maker is told in the instructions that the recipient was selected from her dormitory.

¹¹Our selection algorithm used the “AND”-network definition for this step. Since social distance always (weakly) decreases when using the “OR”-social distance definition the number of observations for columns 1 to 4 in table 2 are not equal.

under the three different exchange rates (3:1, 1:1 and 1:3) for each partner.

To control for experimenter demand effects of presentation, we randomized whether subjects first made all the nameless decisions or first made all the named decisions. For each named partner in Wave 1, we randomized whether subjects first submitted their anonymous choice or their non-anonymous choice. We also randomized whether a subject first submitted all anonymous (non-anonymous) choices for all partners on a single screen, or whether a subject submitted both anonymous and non-anonymous choices on a separate screen for each partner. Finally, we randomized the order in which the three different exchange rates were listed on subjects' screens.

Note that each decision-maker made 6 decisions for each partner involving 3 different exchange rates under anonymous and non-anonymous treatments (36 decisions in total). This made it very difficult for a participant to identify which decision was chosen ex-post from her earnings.¹²

3.2 Design for Wave 2: Trivia Task and Helping Game

Network Elicitation. To focus on relationships with a high frequency of interaction, we developed the *trivia task* technique to elicit the social network for the second wave. Each subject listed 10 friends about whom they would answer questions. Over several weeks, several of the listed friends were randomly selected and each were sent an e-mail asking him to answer a multiple choice question about himself (e.g. "What time do you get up in the morning?"). The subject then received an e-mail directing her to a web page where she had 15 seconds to answer the same question about her friend. If the subject and her friend submitted identical answers, they both won a prize. The trivia task provides subjects with incentives to list friends that the subject spends time with frequently (and thus is more likely to know the friends' habits).

Allocation game. For the second wave we used a new allocation game to capture the kind of helping behavior that occurs in natural relationships, where an individual might face different opportunity costs to help a friend who needs a favor. In the helping game, each decision-maker was endowed with \$45, and each partner was endowed with \$0. 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. A random price between \$0 and \$30 was determined, and if her maximum willingness to pay was equal or greater than the random price the partner received \$30 and the random price was deducted from the decision-maker's endowment. Otherwise, the decision-maker's payoff equalled her endowment of \$45, and the partner's payoff equalled his endowment of \$0.

Effectively, the decision-maker revealed how much she valued a \$30 gain for the partner. As in the first design, subjects made two rounds of decisions: first they played with a nameless partner,

¹²While in principle the dictator could reveal his allocation to the partner after the experiment in the anonymous case, since the dictator was not told which decision was selected for payment, he would have had to make choices with unique payoffs so that he would know which partner to inform, and remember those choices several weeks later when payments were made. Post-experiment breaking of anonymity is also less of a concern in the Helping Game (described in the next section), since the partner could not verify which subject helped him, since all help yields the same payoff; nor could he verify the helper's cutoff, only that it was higher than the unknown price.

and then they played with four named partners. Also, we chose a between-subjects design for the anonymity manipulation: the decisions for the nameless partner were always anonymous, while in the second round decisions were either all anonymous or all non-anonymous. Every subject played both roles in the game, of a decision-maker and of a partner, which allowed us to test whether the baseline altruism of the recipient affects the decision-maker’s choice. To control for experimenter demand effects of presentation, we randomized the order of named partners.

4 Data Description

4.1 Dictator Games

In December, 2003, Harvard sophomores, juniors and seniors at two dormitories were recruited through posters, flyers, and mail invitation. Subjects who logged onto the website were asked to (1) list their best friends’ names using the coordination task and (2) fill in a basic demographic questionnaire. Subjects were required to name friends from the two participating dormitories. Subjects were paid their earnings from the coordination task, plus a flat payment of \$10 for completing the survey. They were also eligible to earn cash prizes in a raffle, adding \$3 (on average) in earnings. Experimental earnings were added to the student’s electronic cash-card.¹³

569 of the 806 students in those two dormitories (71 percent) participated in the social network survey. The survey generated 5690 one-way links. Of those, 2086 links were symmetric links where both subjects named each other.¹⁴ The resulting OR-network consists of a single connected component with 802 subjects, forming a comprehensive map of the social network.

The dictator game stage was conducted over a one-week period in May 2004. Half of all subjects who participated in the coordination stage were randomly selected to be decision-makers. Out of 284 eligible decision-makers invited, 193 participated in round 1 (decisions for nameless partners) and 181 of those participated in round 2 (decisions for named partners). The participants were representative of the coordination stage sample composition.

4.2 The Helping Game

Information on social networks was collected in December 2004 through an online trivia task advertised on the popular student social website `facebook.com`.¹⁵ 2,360 students signed up, generating were 12,782 links between participants out of 23,600 total links¹⁶ with 6,880 symmetric links. In total, 5,576 out of the 6,389 undergraduates at Harvard either participated in the trivia task or were named by a participant. Upperclassmen had higher participation rates, with only 34 percent

¹³These cards are widely used on campus as a cash substitute, and many off-campus merchants accept the cards.

¹⁴For 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.

¹⁵More than 90 percent of Harvard undergraduates were already members of `facebook.com` at that time.

¹⁶Subjects could also list non-participants in our experiment as friends.

of freshman responding, but with 45, 52, and 53 percent of sophomores, juniors, and seniors participating, respectively. 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 stage was conducted in May 2006¹⁷, over a one-week period with all juniors and seniors who had participated in the previous academic year’s trivia task. 776 subjects participated in the first part of the helping game (decisions for nameless partners), and 695 subjects completed the second part (decisions for named partners).

4.3 Summary Statistics

Table 2 shows the mean actions of decision-makers for the dictator and helping games, for both the anonymous and non-anonymous treatments. Two patterns are clear: in all games and in both treatments, the decision-makers’ generosity toward their partners decreased with social distance, and for any game and at any social distance, the decision-makers’ generosity is always higher under non-anonymity than under anonymity. Differences between treatments are 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’ allocation choices are significantly larger for partners who are direct friends than for partners at any other social distance.

In the dictator game with the 1:3 (efficient) exchange rate, the decision-maker passes an average of 19.19 tokens to a direct friend versus 12.20 tokens to a partner at social distance 4. With an (inefficient) exchange rate of 3:1, the decision-maker passes only 8.03 versus 6.15 tokens, respectively. In the non-anonymous treatment, for all social distances, the decision-maker passes about 4 to 5 more tokens when altruism is efficient and about 2 to 4 more tokens when altruism is inefficient. In the helping game’s anonymous treatment, the average willingness to pay is \$12.77 for a direct friend, and decreases to \$7.09 for a partner at social distance 4. Non-anonymity increases the cutoff by approximately \$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. The average level of giving to nameless partners in our dictator games is similar to the results reported in Andreoni and Miller (2002). Nameless partners are given 34.8%, 23.2% and 16.6% of the token endowment in the 1:3, 1:1, and 3:1 exchange rates, respectively, in our experiments, while nameless partners were given 32%, 24.3% and 20% of the token endowment in the equivalent treatments of Andreoni and Miller (2002).

We can interpret nameless decisions under anonymity as the decision-makers’ baseline or unconditional generosity, since she has no information about the partner. Our data replicates the

¹⁷During both waves of the experiment, the network information was also used for other experiments.

well-known finding of Andreoni and Miller (2002) and of Fisman, Kariv and Markovits (2007) that individuals are highly heterogeneous in their unconditional altruism. In particular, we 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.

5 Anonymity and Directed Altruism

In this section, we use the anonymous treatments to analyze how decision-makers' altruistic preferences vary with social distance. In section 2, we motivated the following specification for estimating the strength of directed altruism:

$$x_{MP} = \alpha Z + \gamma_1 D_{MP} + \gamma_M + \epsilon_{MP}. \quad (1)$$

Recall that x_{MP} is the decision-maker's action in the anonymous treatment. Since agents' actions are bounded below by zero and above by 50 in the dictator games, and by zero and 30 in the helping game, we use Tobit regressions to estimate equation (1). 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.¹⁸ 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 $SD = 1$) to $SD5$ (meaning social distance $SD = 5$). The omitted categories are $SD4$ for the dictator games and $SD5$ for the helping game.¹⁹ The estimated coefficient on $SD1$ in a dictator game, for example, should 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 toward 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 4.

We also estimate the specification with additional covariates, and report the results in the even-numbered columns of table 4. We include the decision-maker's action toward 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 baseline altruistic type by including the partner's nameless decision (since all subjects played both roles in the helping game). 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 they share a dormitory (helping game).²⁰

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

¹⁸Our results are very similar when we estimate equation 1 using standard random effects or fixed effects GLS.

¹⁹The social network that we used for the second experiment is much larger because it potentially involves all Harvard undergraduates. Therefore, the maximum social distance between subjects is higher.

²⁰Since the participation rate was lower for the helping game, the "same entryway" dummy variable is less useful.

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 the partner and the generosity displayed toward a nameless partner. Looking across all regression specifications, for both the dictator games and the helping game, each one unit increase in generosity toward a nameless partner is associated with a 0.56 to 1.40 unit increase in generosity toward a named player. Since the nameless decision and the named decisions were elicited one week apart, this continuity indicates a substantial degree of stability in the decision-makers’ preferences over time. Because the effect of “nameless altruism” on “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 heterogeneous trait that strongly influences the decision-maker’s action toward specific named partners.

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 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 absolute magnitude, which implies that decision-makers are making a greater relative sacrifice in the case of inefficient altruism.

Interestingly, neither the subjects’ gender, nor their geographic proximity, have a significant effect on generosity. However, the signs of the estimated gender coefficients for the decision-maker are consistent with 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.

6 Non-anonymity and Future Interaction Effects

We now examine how a decision-maker’s allocation changes under non-anonymity (i.e. both the decision-maker and partner are told which choice was selected for payment) and discuss which model – enforced reciprocity, signaling or preference-based reciprocity – best fits our results.

6.1 Graphical Analysis

In figure 3 we analyze the change in allocations (for the same named partner) between the non-anonymous treatment and the anonymous treatment. Since our helping game was a between-subject design, we can only perform this exercise for dictator decisions. We divide decision-makers into five groups depending on their generosity in the anonymous treatment (0 to 9, 10 to 19, etc.), where the most selfish subjects passed between 0 and 9 tokens. 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 group and by social distance (a proxy for relationship value V_{MP}).²¹

Decision-makers substantially increase their action from the anonymous treatment to the non-anonymous treatment, except the most generous under anonymity.²² This effect is strongest, up to ten extra tokens, when the decision-maker is selfish and when giving is efficient. The effect is half as large when giving is inefficient, where decision-makers pass at most 5 extra tokens.

The main insight from this graph is that the non-anonymity effect is large for close friends and declines with social distance when giving is efficient (and is also somewhat present when giving is neutral). However, when giving is inefficient, the decision-makers' additional contributions are quite small, and do *not* decrease with social distance for four out of the five groups. This result provides some preliminary evidence in support of the enforced reciprocity and preference-based reciprocity mechanisms and against the signaling mechanism.

The graph 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 2.1, we motivated the following specification for the non-anonymous condition:

$$\tilde{x}_{MP} = \eta Z + \theta x_{MP} + \phi V_{MP} + v_M + \epsilon_{MP}. \quad (2)$$

where \tilde{x}_{MP} is the decision-maker's action in the non-anonymous treatment when matched with a specific named partner P . We again use panel Tobit regressions to account for censoring, and 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.²³ The omitted social distance is $SD4$ for the dictator games and $SD5$ for the helping game. The estimated coefficient on $SD1$ in a dictator game, for example, should 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 academic class of the decision-maker and partner, because we expect the non-anonymity effect to be smaller for juniors and (especially) seniors, since the participants are less likely than sophomores to interact with each other in the future. On the right-hand side we also include the decision-maker's

²¹In the NBER working paper we use network flow as an alternative proxy with qualitatively identical results.

²²Even 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.

²³In one specification, we include both measures in the same regression.

action toward 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 toward 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-subject design because for no decision-maker/partner pair in the non-anonymous treatment, do we observe the decision-maker’s choice for that partner in the anonymous treatment. Therefore, we estimate it by running an auxiliary random-effects Tobit regression with data from the anonymous treatment, and include social distance dummies, and the same set of covariates Z (nameless decision, class dummies) as in our empirical specification of the non-anonymity mechanism.

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 5.

Result 3 *The observability of decisions by partners increases giving more for friends than for strangers. The differential effect is only induced when giving is efficient; therefore, the prospect of future interactions increases the surplus created.*

Our main finding is that, controlling for a decision-maker’s anonymous action, her response to non-anonymity increases with the strength of her relationship to the partner, but *only* if giving is efficient (i.e. surplus-increasing). This is true regardless of whether we proxy for the strength of a relationship using social distance or maximum network flow. Note that this result is consistent with both the enforced reciprocity and the preference-based reciprocity mechanisms, but not with the signaling mechanism (which predicts excess generosity toward friends for all exchange rates).

The magnitude of the non-anonymity effect is large, particularly 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 (statistically significant at the 1 percent level). 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.²⁴

Our results also provide some evidence for enforced, rather than preference-based, reciprocity.

Result 4 *The non-anonymity effect increases with maximum network flow.*

In table 5 we estimate a specification for each game including both social distance dummies *and* maximum network flow as covariates. For both the efficient dictator game and the helping game, we find that the coefficients on the social distance dummies decrease and become insignificant when we add network flow, while the coefficient on flow remains significant for the helping game. Additionally,

²⁴In contrast, when we tested for directed altruism, we did not find such a strong effect for friends of friends.

for the helping game, a likelihood ratio test rejects the specification excluding network flow in favor of the specification with both social distance and network flow ($p < 0.01$).²⁵

We also estimate a specification for the dictator game using network flow for allocations to *direct friends*, presented in table 3, to further demonstrate that network flow is capturing information about allocations not included in social distance (since here all partners have a social distance of one). Moreover, for direct friends we can include a control for the amount of time that the decision-maker spends with the partner, since this would be an aspect of the relationship possibly omitted by social distance that would still be consistent with the preference-based model of reciprocity.²⁶ We find that when giving is efficient, greater network flow increases the decision-maker’s generosity toward a direct friend under non-anonymity, even when we control for time spent together. Thus, network flow predicts generous allocations beyond the social distance information. Moreover, the estimated coefficient on time spent together is consistently insignificant and *negative*. We interpret these findings as evidence for the enforced reciprocity model.

Result 5 *The non-anonymity effect and directed altruism are substitutes. Altruistic individuals have a smaller change in allocations under non-anonymity than selfish individuals.*

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. The total difference between the non-anonymous allocation and the anonymous allocation will be smaller for a more altruistic subject than for a selfish subject. This conclusion is also consistent with the enforced reciprocity model.

7 Comparing Altruism to Enforced Reciprocity

Having established that both directed altruism and enforced reciprocity play a substantial role in determining allocation decisions when subjects are connected through a social network, we now attempt to benchmark the relative importance of each effect, compared to baseline altruism, for each condition we study. In table 6, we divide the estimated coefficients on social distance dummies (SD1 to SD3) from both the anonymous (i.e., measuring directed altruism) and the non-anonymous (i.e., measuring enforced reciprocity) allocations by the average generosity toward nameless partners in the anonymous treatment. In other words, we calculate what percent of the baseline altruism effect (relative to selfishness) the network effects represent. For the non-anonymous case, we also include the estimated effect of a one standard deviation increase in network flow. Directed altruism toward friends is equal to 52 percent of the average nameless generosity shown in the efficient dictator

²⁵For all three dictator games and the helping game, we cannot reject the specification with only network flow in favor of the specification with both network flow and social distance ($p > 0.40$ for all specifications).

²⁶However, the measure of time spent together is uncorrelated with our network flow measure for direct friends ($\rho = 0.03$). Thus, any predictive power of network flow appears not to relate to the frequency of interaction.

game, and is equal to 88 percent for the helping game. When altruism is inefficient, the directed altruism effect almost equals average nameless generosity (a 95 percent increase) - nearly double the magnitude of the efficient case (although, recall that the absolute increase in tokens is slightly smaller). Social distance, therefore, is as important a determinant of a subject’s generosity as her baseline altruism, particularly when giving is inefficient.

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. Friends 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.

Additionally, the non-anonymity effect is generally weaker for decision-makers who are juniors and seniors: the signs on the junior and senior dummies are consistently negative, although they 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 potentially shorter (lower) for upperclassmen, because they will live in the same dorm for a shorter time.

We also consider a different benchmark for the magnitude of directed altruism and enforced reciprocity: how large an increase in baseline altruism (i.e., how many standard deviations) would it take to equal the effect of social connections?²⁷ Therefore, in table 6, we also report the estimated coefficients on SD1, SD2, SD3 and network flow as percentages of the standard deviation of the distribution of nameless decisions. We find that greater social proximity to the partner in the anonymous dictator games moves the decision-maker’s generosity by at least 0.39 of a standard deviation, and by a maximum of 1.93 standard deviations in the anonymous helping game. Similarly, the effect of enforced reciprocity by this comparison represents an additional increase in generosity almost half as large as the altruism effect for the efficient dictator game and the helping game (a change of 0.14 and 0.69 of a standard deviation, respectively).

8 Correlation in Altruistic Preferences

Our design allows us to analyze whether subjects who are more altruistic also have friends who are more altruistic. We separate decision-makers into (approximate) quintiles based on their choices for nameless partners. Tables 7 and 8 present the resulting distributions of friends’ generosity.

Result 6 *Friends sort by baseline altruism. Subjects with a high level of baseline altruism have more friends with a high level of baseline altruism, while selfish subjects have more selfish friends.*

²⁷This measure is used in Andreoni and Miller (2002) and Fisman et al. (2007).

First, we find that altruistic and selfish subjects have the same number of friends.²⁸ 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, 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 25 percent more highly-altruistic friends than any other group; in the dictator game, the two most altruistic groups had over 20 percent more highly-altruistic friends than any other group. Moreover, a subject’s friends’ mean nameless allocation 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 $p < 0.026$). It seems that subjects either prefer to become friends with people who have similar social preferences, or their social preferences becomes more similar to each other after they initiate a friendship.

An important consequence of the correlation in friends’ baseline altruism is that it pays to be generous. For the anonymous treatment of the helping game, table 9 (column 1) displays the results of regressing the average allocation to partners from decisions made by direct friends on the partners’ baseline altruism (by quintile): partners with higher baseline altruism have substantially higher earnings. For example, direct friends of the most altruistic partners set the cutoff more than 5 dollars higher than the direct friends of the most selfish partners. Interestingly, this effect is entirely due to kinder partners having nicer friends, and is *not* due to generally kind partners being treated more nicely by their friends: we have already shown in our directed altruism regressions in table 4 that decision-makers do not treat generous partners better. Indeed, when we also control for the average baseline altruism of the decision-makers in table 9 (column 2), the partner’s baseline altruism no longer predicts her earnings from friends’ decisions.

9 Conclusion

We use large field experiments in real-world social networks to jointly measure two explanations of prosocial behavior: the prospect of future interaction and directed altruism. We find support for both mechanisms, although the future interaction effect is half as strong as the directed altruism effect. The mechanisms also have different welfare effects: directed altruism promotes equitable allocations, while enforced reciprocity promotes efficient allocations.

Our results suggest that it is difficult to empirically measure the importance of norms of reciprocity (based on dynamic game effects) without first measuring underlying directed altruism. For example, in empirical studies of informal lending and risk-sharing, it may be as important to mea-

²⁸The same holds for the number of people who list the subject as a friend.

sure directed altruism across social ties (using diagnostic games) as it is to observe cash and in-kind transfers. While both directed altruism and enforced reciprocity suggest that capital in small communities will more likely be directed toward close friends and relatives, the relative importance of altruism over reciprocity may suggest that resources are not being allocated efficiently to the most profitable projects. A better understanding of the strength and reach of altruistic ties within and across families can help policy makers target outside transfers more effectively.²⁹

Our result that friends cluster by baseline altruism raises another interesting question for future research: do our friends shape our social preferences (treatment effect), or do we seek out friends with similar social preferences (selection effect)? This could help explain to what extent the distribution of preferences, as observed in the lab by Andreoni and Miller (2002) and Fisman et al. (2007), is endogenous. In Leider, Mobius, Rosenblat and Do (2010), we provide preliminary evidence against selection based on altruistic type. When we elicit beliefs of partners about the generosity of their friends toward them and toward anonymous others, we find that partners expect to be treated better by their friends, but subjects do not incorporate their friends' specific baseline altruism into their beliefs. Therefore, it appears that while friends tend to cluster by their altruistic type, their beliefs about decision-makers' generosity are mostly driven by expectations of directed altruism. Future research can utilize carefully designed field experiments to measure the relative strength of selection versus social interaction in the formation and evolution of social preferences.

References


- Andreoni, James**, "Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving," *Economic Journal*, 1990, 100 (401), 464–477.
- and **John Miller**, "Giving According to GARP: An Experiment on the Consistency of Preferences for Altruism," *Econometrica*, March 2002, 70 (2), 737–753.
- and **Lise Vesterlund**, "Which is the Fair Sex? Gender Differences in Altruism," *Quarterly Journal of Economics*, February 2001, 116 (1), 293–312.
- Angelucci, Manuela, Giacomo De Giorgi, Marcos Rangel, and Imran Rasul**, "Insurance in the Extended Family," *Working paper, University of Arizona*, 2008.
- Benabou, Roland and Jean Tirole**, "Incentives and Prosocial Behavior," *American Economic Review*, December 2006, 96 (5), 1652–1678.
- Bolton, Gary and Axel Ockenfels**, "ERC: A Theory of Equity, Reciprocity and Competition," *American Economic Review*, March 2000, 90 (1), 166–193.

²⁹Angelucci, Giorgi, Rangel and Rasul (2008) show that random positive income shocks from the *Progres*a natural experiment are shared within rather than across extended families (consistent with a model of directed altruism).

- Brañas-Garza, Pablo, Ramon Cobo-Reyes, Maria Paz Espinosa, Natalia Jiménez, and Giovanni Ponti**, “Altruism in the (Social) Network,” DFAEII Working Paper, University of the Basque Country - Department of Foundations of Economic Analysis II 2006.
- Camerer, Colin**, *Behavioral Game Theory: Experiments in Strategic Interaction*, Princeton, NJ: Princeton University Press, 2003.
- Charness, Gary and Matthew Rabin**, “Understanding Social Preferences with Simple Tests,” *Quarterly Journal of Economics*, August 2002, 113 (3), 817–869.
- Coleman, James S.**, “Social Capital in the Creation of Human Capital,” *American Journal of Sociology*, 1988, 94 (1), 95–120.
- , *Foundations of Social Theory*, Cambridge, MA: Harvard University Press, 1990.
- , “The Rational Reconstruction of Society,” *American Sociological Review*, February 1993, 58 (1), 1–15.
- Dufwenberg, Martin and Georg Kirchsteiger**, “A Theory of Sequential Reciprocity,” *Games and Economic Behavior*, May 2004, 47 (2), 268–298.
- Fehr, Ernst and Klaus Schmidt**, “A Theory of Fairness, Competition and Cooperation,” *Quarterly Journal of Economics*, August 1999, 114 (3), 817–868.
- Fershtman, Chaim and Uri Gneezy**, “Discrimination in a Segmented Society: An Experimental Approach,” *Quarterly Journal of Economics*, February 2001, 116 (1), 351–377.
- Fisman, Ray, Shachar Kariv, and Daniel Markovits**, “Individual Preferences for Giving,” *American Economic Review*, December 2007, 97 (5), 1858–1876.
- Glaeser, Edward L., David I. Laibson, Jose A. Scheinkman, and Christine L. Soutter**, “Measuring Trust,” *Quarterly Journal of Economics*, August 2000, 115 (3), 811–846.
- Goeree, Jacob K., Margaret A. McConnell, Tiffany Mitchell, Tracey Tromp, and Leat Yariv**, “Linking and Giving Among Teenage Girls,” working paper, California Institute of Technology February 2008.
- Granovetter, Mark**, *Getting a Job*, Cambridge, MA: Harvard University Press, 1974.
- , “Economic Action and Social Structure: The Problem of Embeddedness,” *American Journal of Sociology*, November 1985, 91 (3), 481–510.
- Harrison, Glenn W. and John A. List**, “Field Experiments,” *Journal of Economic Literature*, December 2004, 42 (4), 1013–1059.
- Hoffman, Elizabeth, Kevin McCabe, and Vernon L. Smith**, “Social Distance and Other-Regarding Behavior in Dictator Games,” *American Economic Review*, 1996, 86(3), 653–60.
- Leider, Stephen, Markus M. Mobius, Tanya Rosenblat, and Quoc-Anh Do**, “What Do We Expect from Our Friends?,” *forthcoming in the Journal of the European Economic Association*, 2010, 8 (1).

- , — , **Tanya S. Rosenblat, and Quoc-Anh Do**, “Directed Altruism and Enforced Reciprocity in Social Networks,” working paper 13135, NBER May 2007.
- Marmaros, David and Bruce Sacerdote**, “How Do Friendships Form?,” *Quarterly Journal of Economics*, 2006, 121 (1), 79–119.
- Marsden, Peter V. and Karen E. Campbell.**, “Measuring tie-strength,” *Social Forces*, December 1984, 63 (2), 482–501.
- Mathews, K. M., M. C. White, B. Soper, and C. W. von Bergen**, “Association of indicators and predictors of tie-strength,” *Psychological Reports*, 1998, 83 (2), 1459–1469.
- Mitchell, J. Clyde**, “The components of strong ties among homeless women,” *Social Networks*, March 1987, 9 (1), 37–47.
- Mobius, Markus and Adam Szeidl**, “Trust and Social Collateral,” *NBER Working paper 13126*, 2007.
- Murningham, J. Keith and Alvin E. Roth**, “Expecting Continued Play in Prisoner’s Dilemma Games,” *Journal of Conflict Resolution*, 1983, 27 (2), 279–300.
- Perlman, Daniel and Beverly Fehr**, “The development of intimate relationship,” in D. Perlman and S. Duck, eds., *Intimate relationships*, Newbury Park, CA: Sage, 1987.
- Polzer, Jeffrey T., Margaret A. Neale, and Patrick O. Glenn**, “The Effects of Relationships and Justification in an Interdependent Allocation Task,” *Group Decision and Negotiation*, 1993, 2 (2), 135–148.
- Rabin, Matthew**, “Incorporating Fairness into Game Theory and Economics,” *American Economic Review*, December 1993, 83 (5), 1281–1302.
- Roth, Alvin E. and J. Keith Murningham**, “Equilibrium Behavior and Repeated Play of the Prisoners’ Dilemma,” *Journal of Mathematical Psychology*, 1978, 17 (2), 189–198.
- Townsend, Robert**, “Risk and Insurance in Village India,” *Econometrica*, May 1994, 62 (3), 539–591.
- Udry, Chris**, “Risk and Insurance in a Rural Credit Market: An Empirical Investigation in Northern Nigeria,” *Review of Economic Studies*, 1994, 61 (3), 495–526.
- Völker, Beate**, “Should auld Acquaintance be forgot...?” *Institutions of Communism, the Transition to Capitalism and Personal Networks: the Case of East Germany*, Amsterdam: Thela Thesis, 1995.
- Yellen, John E.**, “The Transformation of the Kalahari !Kung,” *Scientific American*, April 1990, pp. 96–105.
- YouGov**, “What’s Money Between Friends?,” PayPal UK Media Center December 2007.

Figure 1: Examples to illustrate difference between maximum network flow and social distance



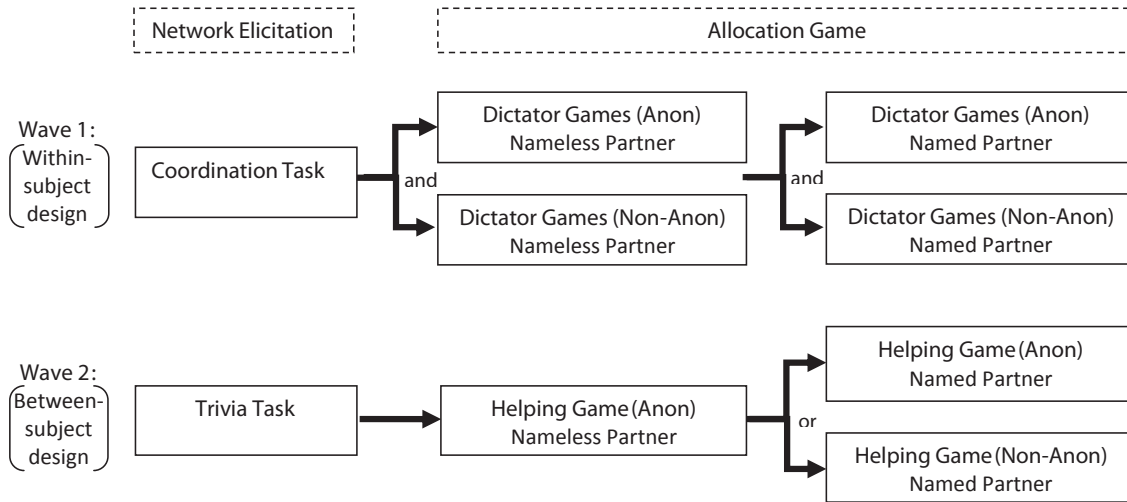
All links are assumed to have unit capacity. The examples illustrate the different features of the social network that are captured by social distance and maximum network flow, respectively: the addition of common friends will necessarily increase flow but can leave social distance unchanged.

Table 1: Testable predictions about decision-makers’ allocation decisions under non-anonymity versus anonymity

	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
Altruistic decision-makers are relatively less generous towards friends compared to strangers	Yes	No	Yes
Maximum network flow is a separate predictor of generosity beyond social distance.	Yes	No	No

We estimate the empirical model $\tilde{x}_{MP} = \eta Z + \theta x_{MP} + \phi V_{MP} + v_M + \epsilon_{MP}$ where \tilde{x}_{MP} and x_{MP} are the decision-maker’s actions under non-anonymity and anonymity, respectively, and V_{MP} describes the value of the relationship between decision-maker and partner (proxied by social distance).

Figure 2: Overview of experimental design (Waves 1 and 2)



In Wave 1, we randomized whether subjects first made all the nameless decisions or whether they first made all the named decisions. For each named partner in Wave 1, we randomized whether subjects first submitted their anonymous choice for that partner or their non-anonymous choice. We also randomized whether a subject first submitted all anonymous (non-anonymous) choices for all partners on a single screen, or whether a subject submitted both anonymous and non-anonymous choices on a separate screen for each partner. Finally, we randomized the order in which the three different exchange rates were listed on subjects’ screens.

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: 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 (1)	Dictator-1:1 (2)	Dictator-3:1 (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).

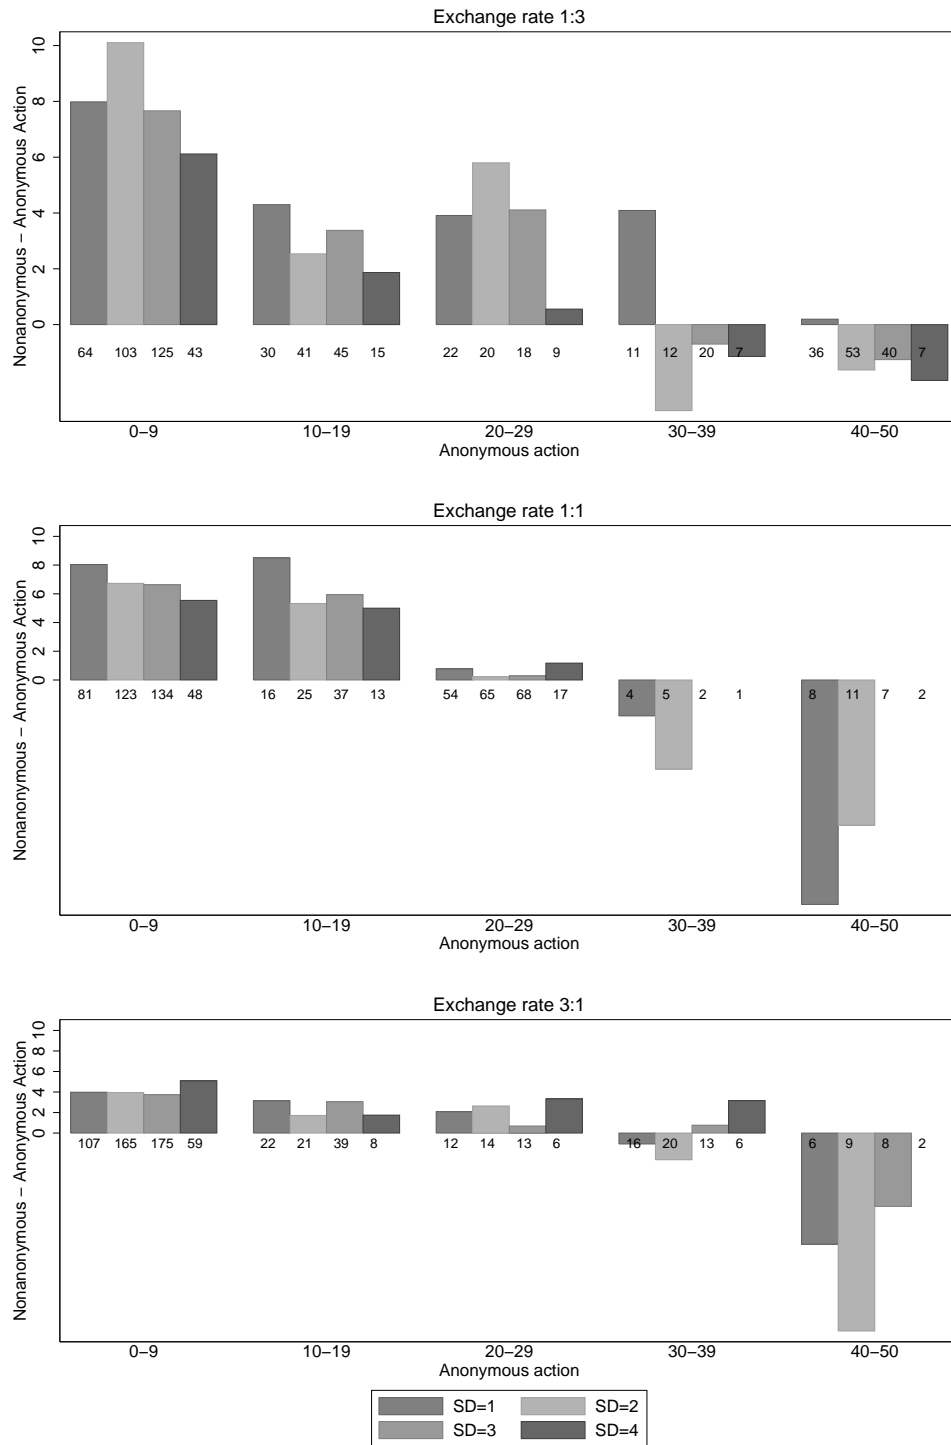
Table 4: Decision-makers' actions in the anonymous treatment (dictator and helping game) when paired with 5 partners at various social distances

	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)	-0.961 (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		-651 (1.335)		-0.24 (0.838)		-9.77 (1.165)		-5.23 (0.388)
Same entryway/house		0.732 (1.376)		-5.17 (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.

Figure 3: 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. The number of subjects in each group is presented below the bar.

Table 5: 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)	2.470 (1.175)*		2.498 (2.168)	-0.555 (1.556)		2.815 (3.018)	3.333 (4.618)		0.521 (4.665)
SD2	3.196 (1.346)*		1.591 (1.876)	1.445 (1.141)		1.461 (1.540)	0.689 (1.503)		2.604 (2.100)	1.214 (1.679)		0.193 (1.696)
SD3	0.513 (1.335)		0.451 (1.334)	1.030 (1.131)		1.031 (1.132)	-0.208 (1.487)		-0.148 (1.486)	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 (0.086)**
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	-0.957 (3.239)	-0.878 (3.237)	-0.936 (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)	-0.021 (0.937)	-0.148 (0.9)	-0.129 (0.933)
Partner is Junior	1.712 (1.058)	1.615 (1.054)	1.645 (1.058)	-0.758 (0.891)	-0.814 (0.891)	-0.757 (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)	-0.733 (0.817)	-0.815 (0.815)	-0.733 (0.817)	1.152 (1.096)	1.128 (1.089)	1.174 (1.095)	-0.227 (0.76)	-0.384 (0.51)	-0.330 (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)	-0.032 (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-subjects 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 6: Relative magnitudes of directed altruism and non-anonymity effects as percentages of all the decision-makers' average nameless action and as percentages of a standard deviation in nameless actions

Relative to:	Directed Altruism			Effect of Non-Anonymity			
	SD1	SD2	SD3	SD1	SD2	SD3	Network Flow
<i>Dictator Game (1:3)</i>							
Average	52	8	-8	24	18	3	19
Standard dev.	39	8	-4	14	10	2	11
<i>Dictator Game (1:1)</i>							
Average	52	16	3	21	12	9	10
Standard dev.	41	14	5	24	14	10	11
<i>Dictator Game (3:1)</i>							
Average	95	49	43	-7	8	-3	-6
Standard dev.	48	25	23	-4	4	-1	-3
<i>Helping Game</i>							
Average	88	36	12	35	13	19	30
Standard dev.	192	82	28	69	25	38	59

An "Average" row is calculated by dividing estimates for directed altruism (table 4) and the effect of non-anonymity (table 5) by the average nameless decision in the anonymous treatment (table 2, "Nameless" column). A "Standard deviation" row is determined by calculating the increase in nameless decision (measured in standard deviations) that would have an equal predicted increase in choice as the estimates for directed altruism and non-anonymity. 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 7: Correlation in baseline altruism among direct friends (dictator game (1:3))

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

Approximate quintiles based on choices for nameless partners (anonymous).

Table 8: Correlation in baseline altruism among direct friends (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

Approximate quintiles based on choices for nameless partners (anonymous).

Table 9: Regressing average allocation to partners from direct friends on baseline altruism 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.