Trust and Reciprocity in 2-node and 3-node Networks

ac Alessandra Cassar and mr Mary Rigdon

University of San Francisco, University of Michigan

26. January 2008

Draft
acassar@usfca.edu
mrigdon@umich.edu

Trust and Reciprocity in 2-node and 3-node Networks

Alessandra Cassar — University of San Francisco
Mary Rigdon — University of Michigan*

January 26, 2008

Abstract

In this paper we focus on the interaction between exogenous network structure and bargaining behavior in a laboratory experiment. Our main question is how competition and cooperation interact in bargaining environments based on networked versions of the investment game. We focus on 3-node networked markets and vary the network structure to model competition *upstream*—multiple sellers paired with a monopsonistic buyer—and competition *downstream*—a monopolistic seller paired with multiple buyers. We describe two kinds of models of trust for such networked environments, *absolute* and *relativized* models, and use this structure to generate a general hypothesis about these environments: that information crowds in cooperation on the competitive side of the market. The experimental results support this hypothesis.

1 Introduction

It is well-known that bilateral bargaining under incomplete contracts is often cooperative and efficient even in environments where the equilibrium path favors no exchange at all. A buyer and seller manage to trade, often sharing the gains from exchange. But many environments have more structure to them than is represented by bilateral exchange: a principal may invest simlutaneously in multiple agents, and an agent may represent the interests of more than one principal. In each case, all of the connected parties are actors in the transaction and the structure of the network they find themselves in may have consequences for their choice behavior in the (thin) market in which they operate. The strategic

* The authors would like to thank the participants at the Economic Science Association meetings in Rome (June 2007) and Tucson (October 2007) and the Decision Consortium at the University of Michigan as well as Dan Friedman and Yan Chen for comments. For generous funding support, Cassar thanks the University of San Francisco and Rigdon thanks the University of Michigan.

behavior of real agents is thus often embedded in social networks, and so there is a need to study that behavior as it occurs as part of a complex system in which agents are linked to each other (Granovetter, 1985; Barabasi, 2002; Raub and Wessie, 1990).

There is a large literature, both theoretical and experimental, that focuses on strategic behavior in networks. The scope of issues is broad: the amount of cooperation and coordination achieved in various networked markets (Cassar, 2007; Cassar and Wydick, 2008); borrowing behavior in informal credit arrangements (Mobius and Szeidl, 2007); different patterns of contagion in financial crises (Cassar, 2001; Eisenberg, 1995; Allen and Gale, 2000); ultimatum bargaining outcomes (Corominas-Bosch, 2004; Fischbacher, *et al.*, 2003); employment and inequality in labor markets (Calvo-Armengol and Jackson, 2004); institutional efficiency (Deck and Johnson, 2004); social capital (Karlan, *et al.*, 2005); social learning (Gale and Kariv, 2003); provision of trust when contract enforcement is weak or nonexistent, and transmission of information about profitable trade opportunities (Rauch and Casella, 2001; Cassar, *et al.*, 2004).

In this paper we focus on the interaction between exogenous network structure and strategic behavior in a bargaining environment. Similar to some other research, our focus is on off-equilibrium but efficient cooperation. Unlike previous research, however, we focus on situations in which there are gains from the exchange between parties and in which equilibrium favors no exchange.

There are two ways that bargaining may be implemented on a network. The first, *bargaining over a network*, is familiar from the evolutionary game theory literature. Here individual interactions remain bilateral, but there is imposed structure on the population from which partners are drawn. Thus, agents on a ring bargain with their neighbor to their left and to their right, and agents on a lattice might bargain with the local Moore (8) neighbors or their von Neumann (4) neighbors. In each case, the individual interactions remain bilateral, but the population forces the pairing to be non-random, and this (in the context of an updating rule like the replicator dynamics) can drive the population toward efficient outcomes (Skyrms, 2004; Axelrod,*et al.*, 2002; Blume, 1995; Ellison, 1993; Nowak and May, 1993; Skyrms, 2004; Samuelson, *et al.*, 1998).

The other way of implementing bargaining on a network, what we call *networked bargaining*, changes the topology of the bargaining environment itself to model structured interactions for more than two agents. Here the interactions are no longer bilateral, and the network structure is part of the strategic environment. We focus on this type of interaction between network and bargaining, and in particular on networked bargaining in the investment game of Berg, Dickhaut, and McCabe (1995).

Our main question is how competition and cooperation interact in bargaining environments based on networked versions of the investment game. The level of competition in a market is likely to be a function of several things: the network structure the agents find themselves in, the information available to the parties in the exchange, and the agents likelihood of meeting exchange parties again. Thus, competition is a function, in part, of who the agents are (e.g., how thin the market is), their roles (e.g., buyer or seller), and also what each agent knows about the behavior of the other parties to the interaction. We focus on 3-node networked markets and vary the network structure to model competition *upstream*—multiple sellers paired with a monopsonistic buyer—and competition *downstream*—a monopolistic seller paired with multiple buyers. We use this structure to generate a general hypothesis about these environments: that information crowds in cooperation on the long-side of the market. Our experimental results generally support this hypothesis.

The next section characterizes two networked versions of the investment game. Section 3 describes the difference between absolute and relativized trust and develops a model of trust in the standard investment game using framework from Cox, Friedman, and Sadiraj (in press) and extends the model then to the networked game. Section 4 discusses two kinds of information flow in the network, partial and full information, and puts forward our main hypothesis about the relationship between information flow in the network and rates of cooperation among the agents. Section 5 discusses the experimental design and implementation and Section 6 contains the results. The final section concludes.

2 2-node and 3-node Networked Investment Games

The standard bilateral investment game is the limiting case of networked bargaining: it is a 2-node (dyadic) bargaining network, with one node occupied by Sender and the other by Receiver. Sender (S) is paired with Receiver (R), and each are endowed with M. In the first stage, S sends an amount X ($0 \leq X \leq M$) to R. That investment grows by some factor $r > 1$. In the second stage, R can return some amount Y of rX to S. The final payoff to S is $M - X + Y$ and the final payoff to R is $M + rX - Y$.

We use the 2-node case as a point of comparison, but focus on 3-node (triadic) networked investment games. In the networks we consider, each of the nodes is a different player—either a Sender or a Receiver—and the network is created through the interaction of the players. Moreover, we focus on *directed networks*, where each link represents a directional flow of investments or returns from one player to another player. There are then two unique types of triads: in one the induced competition

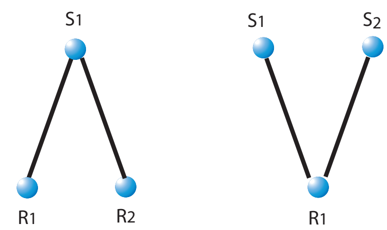

Figure 1: Triadic Networks

is on the Receiver's side (i.e., the buyer's) of the market; in the other the induced competition is on the Sender's side (i.e., the seller's).

The first triadic network structure adds an additional Receiver so that the Sender is paired with two Receivers, A and B (see Figure 1, left); we call this network structure [1S-2R]. S begins with M, and R_A and R_B each begin with M. In the first stage, S chooses an amount X_A of his endowment $(0 \leq X_A \leq M)$ to send to R_A and an amount X_B $(0 \leq X_B \leq M)$ to send to R_B, where $X_A + X_B \leq M$. The invested X_A and X_B are then multiplied by a growth factor, $r > 1$. In the second stage, R_A chooses some amount Y_A of rX_A, and R_B some amount Y_B of rX_B, to return to Sender. In this structure Sender is a monopolist and there is downstream competition between Receivers A and B.

The other triad reverses the asymmetry between Senders and Receivers: two Senders α and β are paired with one Receiver (see Figure 1, right); we call this network structure [2S-1R]. S_α and S_β each begin with M, and R begins with M. In the first stage, S_α chooses an amount X_α $(0 \leq X_\alpha \leq M)$ to send to R, and similarly S_β chooses an amount X_β $(0 \leq X_\beta \leq M)$. The invested X_α and X_β are then multiplied by a growth factor, $r > 1$. In the second stage, R then chooses some amount Y_α of rX_α to return to S_α and some amount Y_β of rX_β to S_β.[1] In this structure Receiver is a monopolist and there is upstream competition between Senders α and β.

The subgame-perfect equilibrium in the dyadic investment game has S opting out, investing nothing and all players earn only their initial endowments of M. Extending the standard backward induction argument to the triads: independent of whether the market power is concentrated in the hands of the Sender (downstream competition) or the Receiver (upstream competition), assuming that agents are Bayesian maximizers, Receivers will keep any investment. Under the standard common knowledge assumptions, Senders know this and (preferring more to less) invest nothing. Since the relevant

[1]Note that R return choices are bounded by the investment made by a particular sender, not by the aggregate investments: R can return some amount Y_i to S_i, where this amount returned is constrained by rX_i (i.e., the gains generated by S_i's investment) and not by $r(X_i + X_j)$.

inequalities between utilities are strict, the equilibrium is unique. Thus the equilibrium path in the triads also favors no exchange between Senders and Receivers.

3 Modeling Trust: Absolute and Relativized Trust

What distribution or outcome an agent prefers might depend not only on what the options for distributing the goods are when she makes her choice, but on how she arrived at that set of options. If preferences are understood purely classically, of course, this is problematic. But in the case of bargaining environments, there is ample experimental evidence of exactly this kind of dynamics of preferences. There are two relevant results. First, cooperative outcomes are chosen by second movers in investment and trust games at different rates (across subjects) depending on different histories of the game leading up to an information set (see, e.g., Croson and Buchan, 1999; Ortmann, *et al.*, 2000; Glaeser, *et al.*, 2000; Fehr, *et al.*, 2002; McCabe, *et al.*, 2002; Engle-Warnick and Slonim, 2005; Rigdon, 2007). With an interesting caveat, rates of return increase as a function of investment.[2] Second, if that history involves an opportunity cost for first movers (i.e., the difference between first movers' outside option and the payoff received if the second mover chooses defection is strictly positive), so that the cooperative move by a second mover can be deemed "reciprocal", there is comparably more cooperative play by second movers than if there is no such opportunity cost (McCabe, *et al.*, 2003).

Thus there is empirical support for the hypothesis that a Receiver in these environments chooses cooperatively only when Sender's action carries a signal of trust. Therefore by giving a more precise characterization of when a Sender action carries a signal of trust, we can generate more precise predictions for the conditions under which Receivers reciprocate, and test those predictions in the laboratory.

Trust in bilateral bargaining games like the investment game is, in general, measured by how much Sender invests. Given the empricial hypothesis that Receivers respond differentially to different levels of trust, it is then straightforward that what counts as trusting behavior (according to Receivers) is the investment level. But once we have non-trivial networked bargaining, there are multiplie measures that might be at work. In [1s-2r] the question is *trust in whom*? Does trust in Receiver A just equal the amount sent to A, X_A (an *absolute* measure), or does it also depend on the amount sent to the other Receiver, X_B (a *relativized* measure)? In [2s-1r] the question is *trust from whom*? Does trust

[2] The caveat is that there is a kink in returns for investment levels in the $(6, 8)$ interval in the standard baseline one-shot investment game.

from Sender α just equal the amount sent from α (absolute), X_α, or does it also depend on the amount sent from the other Sender, X_β (relativized)?

Our main concern is to examine whether and how changing the topology of investment games by imposing network structure on them can crowd-in efficient, off-equilibrium path play. Our initial hypothesis, to be refined below, is that it does. Our conjecture is crowding-in is due to two interacting factors: (1) what counts as a signal of trust in these networked games is some relativized trust measure; and (2) if the information conditions are right, that can get exploited to drive choice behavior toward more efficient play. Thus, for our purposes, it is not important to look to any one particular model of relativized trust. Rather, we can assume that *any* reasonable model of relativized trust says that it is some increasing function of differences in amount sent, and then test the class of all such models all at once. It is still useful to see some simple models of relativized trust, and so we provide some examples.

3.1 Trust in 2-node Networks

In the standard 2-node investment game, the distinction between absolute and relativized trust collapses: the amount sent by Sender (X) is considered a signal of trust. Intuitively, this is because as the investment X gets larger, so does the pie of gains from exchange that Receiver has to divide between them. So Receiver's space of options gets bigger and better while Sender's best-case options get bigger and better but also her *worst-case* option is worse than had she not sent anything at all.

To turn this intuition into a more precise characterization, we need to describe how agents' preferences can be conditional on the opportunities they face. The idea of "dynamic preferences"—that an agent's preferences over bundles might depend on how she arrives at the choice between bundles—is, of course, not new (see e.g. Sen (1997)). What is new is to use such a characterization of when an action carries a signal of trust to derive predictions for off-equilibrium cooperative behavior by Receivers in bargaining environments. Our characterization of trust in 2-node investment games uses some of the framework from a non-parametric model of preferences by Cox, Friedman, and Sadiraj (in press). But we depart from and extend the framework considerably to better suit our purposes, providing an interesting set of hypotheses for network environments.[3]

[3] The main point of departure is that our concerns are different: we want a framework for characterizing a set of hypotheses about networked bargaining, they want to account for well-known experimental data from simpler environments. Thus, they have no need for the ability to express the relativized trust measures we need. On the other hand, we have no special need for their assumptions and constraints about the changing preferences of Receivers: given our purposes, we are happy to take on board, as a general empirical result, that return behavior increases as investments go up. That is all any of our arguments require.

Begin with an n-player extensive form game. The basic idea is that agents can have very different preferences depending on what the set of feasible options is. Thus we conditon preferences on such sets. For a given game, let Π be the total set of possible payoff distributions. We will generally assume that this is a compact, convex subset of $\Re^2_+$. Since our interest is in 2-node and 3-node investment games, three types of distributions will be particularly relevant (where w is Sender's payoff and z is Receiver's payoff, and similarly for $z_{A/B}$ and $w_{\alpha/\beta}$):

- 2-node (baseline): $\pi = (w, z)$
- 3-node [1S-2R]: $\pi = (w, z_A, z_B)$
- 3-node [2S-1R]: $\pi = (w_\alpha, w_\beta, z)$

An *opportunity set* F is a non-empty subset of Π. Opportunity sets are thus just feasible budget sets.

Rather than preferences *simpliciter*, we want to talk about an agent's *preferences given an opportunity set.* Given an opportunity set F, i's preferences given F is a well-defined ordering over F that is convex and continuous. What we care about is that player i's preferences given F can be represented by a smooth utility function $u_i(\cdot|F)$ such that $\frac{\partial u_i(\cdot|F)}{\partial \pi_i} > 0$. Two possibilities are noteworthy here. First, it is possible that i (given F) cares about j's payoff: $\frac{\partial u_i(\cdot|F)}{\partial \pi_j} > 0$. Second, it is possible that $u_i(\pi|F) \neq u_i(\pi|F')$ even when $F \subseteq F'$. That is because the shape of i's preferences may well be affected by the set of opportunities she faces. It is convenient to identify what, given an opportunity set F, a maximizing agent i's best outcome is if i is an own-payoff maximizer; here we abuse notation (slightly, and to no harm) and write $\pi^*(i|F)$. This is the maximum feasible payoff to i in opportunity set F.

In the 2-node investment game, then, it is straightforward to say when one of Sender's actions is more trusting than another: just in case it determines a better budget space for Receiver and how much better is not dwarfed by how much better it is for Sender. That is:

Definition 1.

1. Opportunity set G is at least as generous as F, $G \trianglerighteq F$, iff:

   a) $\pi^*(R|G) \geq \pi^*(R|F)$; and

   b) $\pi^*(R|G) - \pi^*(R|F) \geq \pi^*(S|G) - \pi^*(S|F)$

2. G is more generous than F, $G \triangleright F$, iff $G \trianglerighteq F$ and $F \ntrianglerighteq G$

X carries at least as strong a signal of trust as X' if the opportunity set determined by X is at least as generous as the opportunity set generated by X'.

Thus, to say that one action carries a stronger signal of trust than another is to say that the first generates an opportunity set G that is more generous than the opportunity set F generated by the second. To say that G is more generous than F is just to say that a maximizing Receiver can do no worse in G than in F and that *how much better* R can do in G as compared to F isn't trumped by how much better Sender can do in G as compared to F. This second clause, notice, would not be satisfied if Sender faced no opportunity cost to investing.

Since opportunity sets are determined by Sender's action, it is sometimes useful to write them as a function of that action. In the investment game, Sender's choice is an investment bundle. In the 2-node case, it is a value for X; in the [1s-2r] case, it is a pair (X_A, X_B); and in the [2s-1r] case, each Sender makes an investment choice, X_α and X_β, and each thereby determines an opportunity set for the Receiver. Where c is an investment choice by a Sender, let $\Omega(c)$ be the opportunity set this choice determines.

Example. Consider two possible actions, or investment choices, by Sender in the 2-node investment game with endowment $M = 10$ for each player and a growth rate $r = 3$ (standard parameter values used in the investment game): $X = 2$ and $X' = 8$. Intuitively, the second action represents a stronger signal of trust than the first. This is confirmed by the model since $\Omega(X' = 8) \rhd \Omega(X = 2)$. To see this, let $F = \Omega(X = 2)$ and $G = \Omega(X' = 8)$. Then note that, for any chosen value for X, $\pi^*(R|\Omega(X)) = 3X + M$. Thus, $\pi^*(R|G) = 34 > \pi^*(R|F) = 16$. Hence a maximizing Receiver does better in G than in F. Second, note that $\pi^*(S|\Omega(X)) = M - X + 3X$—that is, the highest profit for Sender, given some initial choice for X, is to get all of the gains from exchange back. Now, $\pi^*(S|G) - \pi^*(S|F) = 26 - 14 = 12$, which is smaller than $\pi^*(R|G) - \pi^*(R|F) = 18$. Hence the size of the potential windfall to Receiver for being in G is greater than the potential windfall to Sender. Therefore, $G \rhd F$.

3.2 Short-side Trust: [1s-2r]

We now turn to 3-node networks in which the distinction between absolute and relativized models does not collapse. First, consider the case of [1s-2r]. Sender chooses an investment bundle (X_A, X_B) such that each X is between 0 and M and such that $X_A + X_B \leq M$. The question is how the bundle of investments is viewed by both Receiver A and how that same bundle is viewed by Receiver B. Since a single bundle could be quite trusting to one receiver and not to another, we have to measure trust-in-A (trust_A) and trust-in-B (trust_B).

An absolute model of trust says that Receiver i views X_i as a signal of trust, full-stop. Thus trust_i depends on comparing R_i's opportunities to S's, but only those of S's that are "local" to interaction with R_i. If this is the case, then X_{-i} does not matter. Such a model would predict no *no downstream competition*. We instead develop a relativized model: trust_i depends on comparing R_i's opportunities to S's "global" opportunities. Thus, since S's global opportunities depend in part on the link shared with R_{-i}, X_{-i} matters to the level of trust_i signalled by the investment bundle. Such a model will predict *downstream competition* between the two Receivers.

Since trust levels are tied to opportunity sets, there is no fact of the matter as to whether one opportunity set is "more generous than another" *simpliciter* in a 3-node network. So we generalize Definition 1, relativizing to a particular receiver. We do this in a way that links R_A's opportunities and X_B (and vice versa), thus making this a relativized model of trust.

Definition 2.

1. Opportunity set G is at least as generous to Receiver i as F, $G \trianglerighteq_i F$, iff:
   a) $\pi^*(R_i|G) \geq \pi^*(R_i|F)$; and
   b) $\pi^*(R_i|G) - \pi^*(R_i|F) \geq \pi^*(S|G) - \pi^*(S|F)$
2. G is more generous than F, $G \triangleright_i F$, iff $G \trianglerighteq_i F$ and $F \ntrianglerighteq_i G$

Let G_i be R_i's opportunity set determined by investment bundle $X = (X_A, X_B)$ and F_i be R_i's opportunity set determined by $X' = (X'_A, X'_B)$. X carries at least as strong a signal of trust_i as X' if the opportunity set determined by X is at least as generous as the opportunity set generated by X'.

So far, this looks like a mere notational variant of the earlier definition: we have swapped R_i for R. But, as we will see below, this is not really true. That is because the term involving Sender in Condition (b) depends on Receiver $-i$'s behavior when faced with opportunity sets G_{-i} and F_{-i}. Thus while it is true that payoffs to the receivers are independent (neither a function of the other), whether an opportunity set is more generous than another to R_i depends on the behavior (and so on the payoffs) of R_{-i}.

The following observation characterizes how differentially a bundle has to favor a Receiver in [1S-2R]—say, Receiver A—for that bundle to constitute a signal of trust$_A$.

Observation 1. Consider two (potential) investment choices by Sender, and let $G = \Omega(X_A, X_B)$ and $F = \Omega(X'_A, X'_B)$. Then $G \rhd_A F$ only if the difference in investments in A are at least $(r-1)$-fold greater than those in B.

Proof. By definition, $G \unrhd_A F$ only if:

Condition (b) $\pi^*(A|G) - \pi^*(A|F) \geq \pi^*(S|G) - \pi^*(S|F)$

The structure of the investment game gives us the following identities for R_A's payoffs (where r is the rate of growth):

$$\begin{aligned}\pi^*(A|G) &= M + rX_A \\ \pi^*(A|F) &= M + rX'_A\end{aligned} \tag{1}$$

Similarly for Sender's payoffs:

$$\begin{aligned}\pi^*(S|G) &= M - X_A - X_B + rX_A + rX_B \\ \pi^*(S|F) &= M - X'_A - X'_B + rX'_A + rX'_B\end{aligned} \tag{2}$$

The LHS of Conditon (b) can be simplified a bit:

$$M + rX_A - M + rX'_A = r(X_A - X'_A) \tag{3}$$

Similarly the RHS:

$$= (M - X_A - X_B + rX_A + rX_B) - (M - X'_A - X'_B + rX'_A + rX'_B) \tag{4}$$

can be simplified to:

$$(r-1)(X_A - X'_A) + (r-1)(X_B - X'_B) \tag{5}$$

So the requirement imposed by Condition (b) is that:

$$r(X_A - X'_A) \geq (r-1)(X_A - X'_A) + (r-1)(X_B - X'_B) \tag{6}$$

And that is equivalent to

$$X_A - X'_A \geq (r-1)(X_B - X'_B) \tag{7}$$

Which finally yields:

$$X_A - (r-1)X_B \geq X'_A - (r-1)X'_B \tag{8}$$

∎

Somewhat surprisingly, the comparative level of trust_A between two investment bundles X and X' requires a corresponding difference in investment level to B across those bundles. Notice that this states a necessary, and not a sufficient, condition on relative trust. Thus, the strength of signal of trust in A carried by a bundle X is in part a function of $X_A - X_B$, which is what we care about.

3.3 Long-side Trust: [2s-1r]

Next, consider the case of [2s-1r]. Here an investment bundle (X_α, X_β) in Receiver is determined by the joint actions of Sender α and Sender β. The question is how the bundle of investments is viewed by Receiver. Since there are two independent parts of that bundle (one from Sender α and one from Sender β), part of that bundle could be quite trusting and the other part not trusting. So we have to measure trust-from-α (trust^α) and trust-from-β (trust^β).

An absolute model of trust says that Receiver views X_i as a signal of trust^i. Thus trust^i depends on comparing S_i's opportunities to R's, but only those of R's that are "local" to interaction with S_i. If this is the case, then X_{-i} does not matter. Such a model would predict *no upstream competition*. We again develop a relativized model: trust^i depends on comparing S_i's opportunities to R's "global" opportunities. Thus, since R's global opportunities depend in part on the link shared with S_{-i}, X_{-i} matters to the level of trust^i signalled by the investment bundle. Such a model will predict *upstream competition* between the two Senders.

It is sufficient to say when investment bundles signal trust^α and trust^β. In the [1s-2r] network, what is relevant is how different investment bundles are viewed by Receivers, each from their own point of view. In [2s-1r], however, what is relevant is how a single investment bundle affects what the Receiver thinks about the two independent Senders: given the generosity of the space of opportunities R faces, we want to model α's contribution to that and β's contribution to that.

Given an investment bundle $X = (X_\alpha, X_\beta)$ that determines an opportunity set G (note that this is a set of triples, as in the [1s-2r] network) we will give a comparative assessment of when G's relative generosity is due to S_α as opposed to S_β. There are different ways of modeling this; we focus on a particularly simple version.

Notice that we can restrict G by eliminating one of its coordinates—say the coordinate for Sender α's profits; such a restriction is the *α-free slice of G*. Thus the α-free slice of G is the set of possible distributions in G to Receiver and Sender β. This allows us to evaluate relative generosity of G, an opportunity set that is a function in part of α's investment, while ignoring α's potential profits in G. Similarly, we can also take the *β-free slice of G*: the possible distributions in G involving Receiver and Sender α. Since G is a set of ordered triples, these slices are planes (sets of pairs) and are like opportunity sets from the 2-node investment game, except that Receiver's profits reflect that he is linked with two Senders. But they are (2-person) opportunity sets, and so can be compared for relative generosity. Receiver's range of profits are the same in these two slices, and so his maximum profit is the same. Thus if the slices differ in their generosity, it must be because one of the Sender's has a higher maximum profit in his slice than the other has in his; the former is less generous.

Definition 3. Let G be the opportunity set determined by investment bundle $X = (X_\alpha, X_\beta)$.

1. Restricted opportunity sets:

   a) α-free slice of G: $G_{\setminus\alpha} = \{(w_\beta, z) : (v, w_\beta, z) \in G \text{ for some } v\}$

   b) β-free slice of G: $G_{\setminus\beta} = \{(w_\alpha, z) : (w_\alpha, v, z) \in G \text{ for some } v\}$

2. Sender i is at least as generous in G as Sender j is iff $G_{\setminus i} \trianglerighteq G_{\setminus j}$

The level of trust^i signalled by investment bundle X is at least as great as trust^j iff Sender i is at least as generous in G as Sender j is.

Observation 2. Sender i is at least as generous in G as j is iff $X_i - X_j \geq 0$

Proof. Suppose Sender i is at least as generous in G as j is. Then $G_{\setminus i} \trianglerighteq G_{\setminus j}$. But this is so iff

(a) $\pi^*(R|G_{\setminus i}) \geq \pi^*(R|G_{\setminus j})$; and

(b) $\pi^*(R|G_{\setminus i}) - \pi^*(R|G_{\setminus j}) \geq \pi^*(j|G_{\setminus i}) - \pi^*(i|G_{\setminus j})$

But

$$\pi^*(R|G_{\setminus i}) = \pi^*(R|G_{\setminus j}) \tag{9}$$

so $G_{\setminus i} \trianglerighteq G_{\setminus j}$ iff (b) holds. That is, iff

$$0 \geq \pi^*(j|G_{\setminus i}) - \pi^*(i|G_{\setminus j}) \tag{10}$$

That is, iff $0 \geq 3X_j - 3X_i$ iff $X_i - X_j \geq 0$. ∎

4 Information

We now have characterizations for relativized trust both for networks with downstream competition ([1S-2R]) and for networks with upstream competition ([2S-1R]). Given the broad empirical generalization that return behavior in investment games increases with signals of trust by senders, we can test the effects of network structure on off-equilibrium but efficient behavior. But competition is a function of more than just network structure; it is also a function of how information about the bargaining environment flows throughout that network. We look at two kinds of information flow, full and partial.[4] Our main empirical hypothesis is about the relationship between the kind of information flow and the amount of cooperation achieved among the players:

Competition for Cooperation Hypothesis (CCH) Information crowds-in cooperation on the long-side of the market

Before turning to our experimental design and results, we will describe the CCH hypothesis in detail in this section by describing information flows in the network and how the level of information could be expected to impact the rates of cooperation.

Say that a network allows *full information flow* if: (i) every agent in the network knows every move made by every other agent at points higher in the game tree; and (ii) after the terminal nodes are reached, there is full-disclosure about all moves. This implies that in a [1S-2R] network with full information flow Receivers each know how much is invested in the other before deciding on a return and both learn how the other responds through full disclosure of all moves. That is:

- R_A knows the value of X_B before choosing Y_A
- R_A learns Y_B through full-disclosure
- R_B knows the value of X_A before choosing Y_B
- R_B learns Y_A through full-disclosure

Similarly, in a [2S-1R] network with full information flow, full-disclosure implies that Senders each learn about what the other invests following own investment choice and also learn how Receiver responds to the other Sender's investment:

- S_α learns the value of X_β and Y_β
- S_β learns the value of X_α and Y_α

[4]In the standard 2-node investment game there is no room for a difference between full and partial information flow in the sense defined below.

Say that a network allows for *partial information flow* if each agent *only knows* about (i.e., has full information about) the interactions in which that agent is a trader. Note that it follows immediately that agents on the short-side of the market know about all the interactions in that market. Moreover, in a [1S-2R] network with partial information Receivers do *not* know the amount Sender invests in the other and does *not* learn how the other Receiver responds:

- R_A does not know X_B
- R_A does not learn Y_B
- R_B does not know X_A
- R_B does not learn Y_A

Similarly, in the [2S-1R] network with partial information, is one in which Senders do not learn the amount the other invests or how the Receiver responds to that:

- S_α does not know X_β or Y_β
- S_β does not know X_α or Y_α

Thus CCH implies that in [1S-2R], with full information flow, there is competition on the long-side of the market between Receivers A and B for Sender's investment. That is because each Receiver knows the level of trust$_A$ and trust$_B$. So, given the broad empricial generalization about return behavior in investment games, we would expect more cooperative return behavior by each. But in the same network with only partial information flow, the signal of relativized trust is masked: neither Receiver A nor Receiver B has enough information to determine trust$_A$ and trust$_B$.[5] Thus, given the broad empirical generalization about return behavior in investment games, we would expect less cooperative return behavior by each.

Similarly, CCH implies that in [2S-1R], with full information flow, there is competition on the long-side of the market between Senders α and β for Receiver's attention. That is because, at the end of each stage, each Sender knows the level of trust$^\alpha$ and trust$^\beta$. Each also knows Receiver's reaction to those levels. So, if Receiver responds differentially to differences between trust$^\alpha$ and trust$^\beta$ (and given the broad empirical generalization about return behavior in investment games, this is to be expected), then we would expect an arms race in transfers between the Senders. But in the same network with only partial information flow, Receiver's differential responses to Sender α and Sender β cannot be traced to differences in trust$^\alpha$ and trust$^\beta$. Thus, Receiver cannot reveal his response-type to the long-side of the market and we would therefore expect less cooperative behavior on that side of

[5] Note that knowing only X_A, Receiver A cannot infer that $X_B = M - X_A$. S is under no obligation to invest all of M, and that is common knowledge.

the market.

Note that both of these implications of the CCH depend on exploiting relativized trust measures in the networks. We now turn to our experimental design and procedures; the results are discussed in Section 6.

5 Experimental Design and Implementation

The baseline condition is the standard 2-node investment game. We then cross type of triadic network {[1S-2R], [2S-1R]} with type of information flow {FULL INFO, PARTIAL INFO} to obtain our $1 + (2 \times 2)$ conditions. All treatments reported here use repeated interactions with random-matching.[6]

The sessions were run at the Learning and Experimental Economics Projects of Santa Cruz (UC-Santa Cruz) and the Robert B. Zajonc's Laboratory (Institute for Social Research, University of Michigan). The sessions were run May 2005 through July 2006. The only difference in the instructions for the treatment conditions was the description of either the network structure or the information available (see Appendices A and B for the [1S-2R] treatment).[7]

We use standard parameter values for initial endowments, $M = \$10$, and for the growth rate, $r = 3$. The experimental protocol for all treatments was single-blind. Subjects in each session were randomly assigned roles as either Sender or Receiver.

Subjects were required to complete a quiz regarding the interaction and completed calculations of payoffs prior to beginning, which the experimenter checked for accuracy. Subjects played a version of the game for 40 rounds with a known end-point, being randomly re-matched with a new counterpart(s) at the start of each round. Subjects who had previously participated in similar experiments were excluded from recruitment. Each session had 12 subjects and took less than 1 hour to complete. Three sessions of each treatment were conducted.

6 Results

Our experimental data support the Competition for Cooperation Hypothesis under both network structures: information crowds-in cooperation on the long-side of the markets. The flow of information

[6]The complete design for this project on networked bargaining has 3 matching treatments $\times$ 2 3-node network structures $\times$ 2 information conditions; in addition to the 2-node baseline.

[7]Computer interface screen shots for the [1S-2R] INFO treatment are available at http://www.umich.edu/~mrigdon/1s2rinfo.pdf.

in the environment is crucial for the competitive structure to be exploited for off-equilibrium path cooperation.

Figures 1 and 2 provide a visual display of the five-period moving averages of the results we report below. Figure 1 displays the differences in trust and trustworthiness over time within networks across information treatments; the dashed line corresponds to the PARTIAL INFO case and the full line corresponds to the INFO case. Figure 2 displays the differences in trust and trustworthiness over time within information treatment across networks; the dashed line corresponds to the [2S-1R] network and the full line corresponds to the [1S-2R] case. Tables 1–3 contain the results of random-effects Tobit regressions that provide direct tests of the theoretical predictions to be discussed in detail below. Tables 4–7 report detailed descriptive statistics on trust, trustworthiness, profit, and efficiency. Significant differences across the information treatment are reported with subscripts and superscripts to the right of the averages, where the subscripts report results using a non-parametric Mann-Whitney two-sided test and subscripts report results using a more conservative panel estimation. In the majority of cases, a result reporting a significant difference with the Mann-Whitney test is confirmed by the respective panel estimation (although the panel often has a larger p-value). The discussion below is based on the less conservative non-parametric results.

6.1 Dyads

The results in the 2-node treatment serves both as baseline and as a replication of earlier investment game experiments.

Return behavior by Receivers replicate the general empirical finding that higher generosity is reciprocated by higher amounts returned in dyadic exchange arrangements. In Table 1, we report the results from several random-effects Tobit regressions, where the dependent variable is the amount returned by Receiver.[8] The first regression has the amount sent by Sender as the only independent variable and we find that the Receiver returns approximately 45% of an extra dollar received. This result is robust once we add variables in the specification to capture the Receiver's homegrown altruism (amount returned by the same Receiver the previous period to a different Sender), an attempt by the Receiver to induce more generosity if the period before the Sender had sent a smaller amount (the amount received the period before from a different Sender), and time trends (round number and its square allowing for a nonlinear time trend). The results show that these additional regressors are all

[8]The censoring is between 0 and the amount X sent by the Sender.

highly significant. Tables 5–7 report that Senders on average send \$6.45 out of their \$10 endowment, and Receivers return about \$9 (or 39.43%).

6.2 1S-2R

Given CCH, we would expect to observe downstream competition in [1S-2R] with full information flow. That is because we would expect to find (i) evidence of relativized trust; and (ii) evidence that that relativized trust is exploited to crowd-in higher returns. It is a relativized measure of trust that Receiver behavior depends on, and it depends on it by being crowed-in toward off-equilibrium cooperative behavior. And, if CCH is false, then we would expect either no differential evidence of relativized trust or no differential evidence of crowding-in across information conditions.

First we examine the extent to which Receivers' choice behavior depends on trust_A and trust_B; that is, for example, whether Y_A depends on $X_A - X_B$. This, of course, is exactly what is predicted by a model of relativized trust: since trust_A depends on $X_A - X_B$, and given the general result that returns increase as level of trust increases, we would expect Y_A to increase as $X_A - X_B$ increases. Similarly for Receiver B: Y_B should increase as $X_B - X_A$ increases. But this dependence is predicted, given CCH, only under full information flow; under partial information flow, no such dependence is expected since neither Receiver knows $X_A - X_B$ and so neither knows $\text{trust}_{A/B}$.

In Table 2, we report the results from several random-effects Tobit regressions. The first two regressions estimate the effects separately by information condition, where the dependent variable is the amount returned by Receiver_i.[9] The two independent variables are amount sent to i and amount sent to the other Receiver. Under PARTIAL INFO, the amount returned depends only on the amount sent to Receiver_i and not on the amount sent to Receiver_{-i} ($p = 0.000$ and 0.599, respectively). Approximately 47% of each unit received is returned to the Sender. Under FULL INFO, it is also the case that the coefficient on the amount sent to Receiver_i is significant ($p = 0.000$); approximately 44% of each unit received is returned to Sender. Additionally, the coefficient on the amount sent to Receiver_{-i} is negative and significant ($p = 0.067$), indicating that Receiver i discounts the level of generosity of the Sender if she observes increasing generosity toward Receiver_{-i}. The result confirms the prediction that Receivers use a relativized version of trust to determine the level of reciprocity when the information is available to do so.

This result is robust. The third and fourth model specifications add a series of independent variables

[9] The censoring is between 0 and the amount X sent by the Sender.

to control for the effects of amounts sent to the Receiver by another Sender in the previous period, a Receiver's homegrown altruism, and time trends. There is strong evidence that all three have significant effects in both information conditions: Receivers adjust to the amount previously sent by returning more this period if they received a smaller amount the previous period, the Receivers exhibit homegrown altruism by returning more this period if they had returned more the previous period, and the time trend shows a non-linearity effect. Once controlling for these variables, under FULL INFO, the negative reaction to the amount sent by Sender to Receiver$_{-i}$ is still negative and weakly significant ($p = 0.10$).

The final regression in Table 2 has as the dependent variable the difference in amounts returned between Receivers A and B to the same Sender, $Y_A - Y_B$. The independent variables are the difference in the amounts sent, a dummy for the information treatment, an interaction term between information, the difference in the amounts sent, lags on the differences, and time trends. The results show that the difference between the amounts sent, $X_A - X_B$, is an highly significant factor: about 42% more is returned to the Sender for each unit sent to the Receiver in excess of what sent to Receiver$_{-i}$ ($p = .000$). Moreover, the interaction term of information treatment with the difference in amounts sent is highly significant, indicating that – when Receivers have information about what was sent to the other at the time they make their return decision – they strictly compare to determine how much to return ($p = 0.027$). This is clear evidence that Receivers utilize a notion of relativized trust.

By itself, the fact that Receiver behavior depends on a relativized measure of trust does not support CCH. For that we need to see, in addition, that return behavior is more cooperative when information about the bargaining environment is full. Then we will have evidence that (i) return behavior depends on comparative information about trust$_{A/B}$, and (ii) that the behavior on the long-side of the market is crowded-in. So, accordingly, we examine the extent to which information crowds-in cooperation by the Receivers. Here we see that, in fact, it does.

We can distinguish *strong* and *weak* reciprocity: a response (return amount) by a Receiver signals weak reciprocity if the amount returned is not less than the amount invested; it signals strong reciprocity if it is greater. We focus on the case of strong reciprocity since it provides the strongest test for CCH. Table 3 reports relative frequencies of (strong) reciprocity – i.e., the relative frequency of returns Y_i such that $Y_i > X_i$ – in the 2-node network and [1S-2R] under both information conditions. For the dyad, it is approximately 57% and is not significantly different from [1S-2R] under partial info ($p = 0.4485$). However, the fraction under FULL INFO is significantly higher than in the dyad, 63%, suggesting higher rates of (strong) reciprocity by Receivers A and B over their counterparts in the 2-node network ($p = 0.0180$).

6.3 2S-1R

Given CCH, we would expect to observe upstream competition in [2S-1R] with full information flow. That is because we would expect to find (i) evidence of relativized trust; and (ii) evidence that that relativized trust is exploited to crowd-in higher transfers. It is a relativized measure of trust that Senders' behavior depends on, and it depends on it by being crowed-in toward off-equilibrium cooperative behavior. And, if CCH is false, then we would expect either no differential evidence of relativized trust or no differential evidence of crowding-in across information conditions.

First, we look to see whether Receiver behavior depends on relative investments. That is, we see whether Y_α (Y_β) depends on both X_α and X_β instead of just X_α (X_β). We find evidence for this hypothesis.

In Table 4, we report the results from several random-effects Tobit regressions. The first two regressions estimate the effects separately by information condition, where the dependent variable is the amount returned by the Receiver to Sender$_i$; the two independent variables are amount sent by Sender$_i$ and the amount sent by Sender$_{-i}$.[10] Our results show that in both information conditions, as the amount sent by Sender$_i$ increases the amount returned to Sender i by the Receiver increases: 52% in PARTIAL INFO and 41% in FULL INFO of the gains from exchange are returned to the relevant Sender. This is consistent with the body of evidence from investment game experiments.

The more important result is that, under FULL INFO, Receivers continuously compare across Senders and send approximately *10% less* to Sender$_i$ if $X_i < X_{-i}$ ($p = 0.022$). Under PARTIAL INFO, on the other hand, when *no* information is available to Sender i regarding Receiver's return to Sender$_{-i}$, the Receiver rewards i slightly more – about 3% – the more she gets from the other Sender ($p = 0.014$). These results suggest of an intention by the Receiver to indicate to the stingier Sender that he has been punished and the more generous Sender has been rewarded.

The third and fourth model specifications in Table 4 add a series of independent variables to control for the effects of total amount sent to the Receiver in the previous period, a Receiver's homegrown altruism, and time trends. Under PARTIAL INFO, the Receiver returns about 55% of the gains from exchange and under FULL INFO returns a smaller percentage, approximately 40%. The positive and

[10] The censoring is between 0 and the amount X sent by Sender$_i$.

significant coefficients found for the lag amount returned confirm the homegrown preferences toward altruism already uncovered in the other networks. Interestingly, once controlling for these variables, under FULL INFO, the contemporaneous negative reaction to the amount sent by the other Sender is still highly significant, with the Receiver penalizing Sender$_i$ by returning 10% less for each additional dollar sent by Sender$_{-i}$ above what Sender$_i$ sent ($p = 0.000$).

The final regression we report in Table 4 has as the dependent variable the difference in amounts returned by the Receiver to Sender α and Sender β ($Y_\alpha - Y_\beta$). The difference between the amounts sent by the two Senders has a significant effect ($p = 0.075$). Additionally, the interaction effect of the information condition with the difference in the amount sent is highly significant ($p = 0.000$).

Under FULL INFO the Receiver discriminates based on Senders relative investment. Senders learn about this over time, triggering an arms race on high transfers. In 65% of the triads, $X_\alpha = X_\beta$. That is, Sender α and Sender β make *the same investment choice*, even though they make those choices simultaneously when they do not have access to the investment choice of each other. In these cases, the average amount sent is nearly the entire endowment, \$9.23. When everyone invests almost all of the endowment, there is little room to discriminate between them. However, in PARTIAL INFO, only 46% of the triads have equal investment levels and in these cases the average amount sent is much lower, \$6.20. This increase in competition occurs when the Senders receive information about what the other Sender sent and what the Receiver returned to the other. This information about today's interaction leaves its mark on tomorrow's behavior. The average amount sent is significantly higher—about 17% higher—under FULL INFO; otherwise the rate is the same as in the dyad (see Figure 1 and Table 5). As a result, trustworthiness is higher under FULL INFO in both absolute amount and in terms of proportion returned (see Tables 6A and 6B); the rates are not significantly different from the dyad under PARTIAL INFO.

These results confirm the CCH for information crowding-in cooperation on the long-side of the market. Each Sender knows the level of trust$^\alpha$ and trust$^\beta$, and each also knows Receiver's reaction to those levels. Consistent with the large body of investment game results, we see that Receiver reciprocates a trusting Sender in this network. Further, Receiver responds differentially to differences between trust$^\alpha$ and trust$^\beta$. Seeing this, the Senders compete for the attention of the Receiver: we have the predicted arms race in transfers. But in the same network with only partial information flow, Receiver's differential responses to Sender α and Sender β cannot be traced to differences in trust$^\alpha$ and trust$^\beta$. Thus, Receiver cannot reveal his response-type to the long-side of the market, there is no ensuing arms race, and we see less cooperative behavior on that side of the market.

6.4 Efficiency and Profits

In this last section on experimental results, we briefly discuss the impact of the observed behavior on two important market features, efficiency (i.e., degree of realized surplus) and profits (i.e., final distribution of that surplus). See Tables 6 and 7. In the baseline 2-node network, overall efficiency is only 82.2%; on average, Senders earn $11.60 and Receivers earn $21.30 (including the initial $10 endowment).[11] The [1S-2R] network is the most efficient at a rate of almost 94% and is independent of flow of information as the rates of trust are high in both information conditions; interestingly, Senders end up faring significantly better under FULL INFO than PARTIAL INFO, whereas Receivers, due to the competition, fare significantly worse under FULL INFO than under PARTIAL INFO.[12] The [2S-1R] network does not generate efficiency gains over the baseline with an overall efficiency of approximately 82%; however, FULL INFO is significantly more efficient than PARTIAL INFO due to higher rates of trust.[13] In terms of profits, the Senders and Receivers fare better than under the baseline network.

7 Conclusions

We utilize networked versions of the investment game to explore how competition and cooperation interact in more realistic bargaining environments, focusing on 3-node networked markets. The level of competition in these markets is a function of the particular network structure and the level of information available to the parties in the exchange.

Once we consider networked bargaining, we have to also consider whether standard behavioral measures of trust in these environments are adequate. Does amount sent still carry that signal, or do we need to know who sent it, or to whom they sent it? We described two plausible models of trust, absolute and relativized. We have given a simple picture of a relativized model, and used it to generate our main empirical hypothesis – that information crowds-in cooperation on the long-side of the market.

Our findings suggest two related conclusions which support the Competition for Cooperation Hypothesis. First, we find evidence that in networked bargaining, both in models of upstream and

[11] Efficiency is the ratio of the profits of the agents to the potential surplus available plus the endowment of the Receiver(s). For the baseline: $\frac{x+y}{40} \times 100$

[12] For the [1S-2R]: $\frac{x+y_A+y_B}{50} \times 100$

[13] For the [2S-1R]: $\frac{x_\alpha+x_\beta+y}{70} \times 100$

downstream competition, it is relativized and not merely absolute trust measures that are operative. Second, we find that when information flows fully in an environment, the network and relativized trust measure are exploited to crowd-in cooperation on the long-side (competitive) of the market. In our 3-node networks with downstream competition, this manifests itself in increasing returns. In our 3-node network with upstream competition, we instead see that Receivers – under full information only – treat different Senders differently, and that this triggers an arms race on investments on the competitive side of the market.

Table 1: Dyads - Random-effects Tobit Regressions

	Dependent Variable: **Amount Returned to S**†	
Number of observations:	n=646	n=581
Number of groups	18	18
Intercept	-1.088	-1.060
	(1.651)	(1.565)
Amt. Sent to R	0.451***	0.449***
	(0.029)	(0.030)
Lag Amt. Sent to R		-0.148***
		(0.035)
Lag Amt. Returned by R		0.315***
		(0.048)
Round number		0.162**
		(0.072)
Round number^2		-0.005***
		-0.002
Wald chi^2	236.70	282.92
Prob > chi^2	0.00	0.00

***99% significance, **95% significance, *90% significance.

†Censored lower limit= 0; upper limit = Amount Sent by S1

Table 2: Triads1S2R - Random-effects Tobit Regressions

	Dependent Variable:				
	Amount Returned to S†				**Diff. Returned**††
Information treatment:	**NoInfo**	**Info**	**NoInfo**	**Info**	
Number of observations:	n=829	n=774	n=710	n=630	n=489
Number of groups:	24	24	24	24	24
Intercept	-1.915**	-0.315	-1.223	-0.24	0.0597
	(0.858)	(0.713)	(0.996)	(0.903)	(0.786)
Amt. Sent to R Self	0.469***	0.442***	0.450***	0.492***	
	(0.020)	(0.022)	(0.021)	(0.025)	
Amt. Sent to R Other	0.011	-0.043*	-0.026	-0.042*	
	(0.020)	(0.024)	(0.021)	(0.026)	
Lag Amt. Sent to R Self			-0.102***	-0.106***	
			(0.029)	(0.032)	
Lag Amt. Sent to R Other			0.004	0.006	
			(0.020)	(0.025)	
Lag Amt. Returned by R Self			0.224***	0.208***	
			(0.045)	(0.048)	
Round number			0.077**	0.050	-0.075
			(0.037)	(0.048)	(0.062)
Round number^2			-0.003***	-0.003***	0.002
			(0.001)	-0.001	(0.002)
Difference Sent to R1 - R2					0.421***
					(0.024)
Information					0.252
					(0.854)
Info*Diff. Sent to R1 - R2					0.083**
					(0.038)
Lag Diff. Returned					0.130***
					(0.047)
Lag Diff. Sent					-0.057**
					(0.028)
Wald chi^2	985.33	726.84	1118.11	809.95	631.99
Prob > chi^2	0.00	0.00	0.00	0.00	0

***99% significance, **95% significance, *90% significance.

†Censored lower limit= 0; upper limit = Amount Sent by S to R Self

††Censored Regression. Lower limit=0 - Amount Sent to R2; upper limit = Amount Sent to R1

Table 3. Relative Frequencies of Strong Reciprocity: Y > X

	Dyad	Triad 1S2R	Triad 2S1R
Treatment			
Full Info	0.5764	**0.6333	***0.6771
	(0.4945, 720)	(0.4821, 960)	(0.4678, 960)
Partial Info		0.5948	***0.7010
		(0.4136, 960)	(0.4580, 960)

Notes: Standard deviations and number of observations are reported in parentheses
Superscripts *** report Mann-Whitney results.
$^{***}p\leq0.001$; $^{**}p\leq0.05$; $^{*}p\leq0.10$; $^{+}p\leq0.15$

Table 4: Triads 2S1R - Random-effects Tobit Regressions

	Dependent Variable: **Amount Returned to Same S†**				
Information treatment:	**NoInfo**	**Info**	**NoInfo**	**Info**	**Diff. Returned††**
Number of observations:	n=793	n=854	n=632	n=683	n=936
Number of groups:	12	12	12	12	24
Intercept	-3.994**	2.897**	-4.238**	0.873	0.296
	(1.769)	(1.431)	(1.700)	(1.368)	(0.579)
Amt. Sent by Same S	0.522***	0.413***	0.549***	0.403***	
	(0.017)	(0.031)	(0.198)	(0.030)	
Amt. Sent by Other S	0.034**	-0.101***	0.026	-0.117***	
	(0.014)	(0.022)	(0.016)	(0.022)	
Lag Amt. Total Sent			-0.060***	-0.029	
			(0.189)	(0.022)	
Lag Amt. Total Returned			-0.093***	0.203***	
			(0.025)	(0.024)	
Round number			0.093**	0.111	-0.022
			(0.054)	(0.074)	(0.061)
Round number^2			-0.002	-0.005***	0.00
			(0.01)	(0.002)	(0.001)
Difference Sent by S1 - S2					0.369***
					(0.017)
Information					-0.053
					(0.332)
Info*Diff. Sent by S1 - S2					0.136***
					(0.024)
Lag Diff. Sent					-0.036*
					(0.020)
Lag Diff. Returned					-0.046
					(0.037)
Wald chi^2	1014.38	198.94	884.97	348.20	1393.35
Prob > chi^2	0.00	0.00	0.00	0.00	0.00

***99% significance, **95% significance, *90% significance.

†Censored lower limit= 0; upper limit = Amount Sent by Same S

††Censored Regression. Lower limit=0 - Amount Sent by S2; upper limit = Amount Sent by S1

Table 5. Trust

	Dyad	Triad 1S2R		Triad 2S1R	
Treatment	Amount Sent	Tot. from S	Avg. to Each R	Tot. from S_{α}&S_{β}	Avg. from Each S
Full Info	6.45	${}^{***}_{**}8.43$	${}^{***}_{***}4.22$	${}^{***}_{***}14.55^{***}_{+}$	${}^{***}7.28^{***}_{+}$
	(3.44, 720)	(2.96, 480)	(3.04, 960)	(5.08, 480)	(3.46, 960)
Partial Info		${}^{***}_{***}8.72$	${}^{***}_{***}4.36$	${}^{***}_{***}12.46^{***}_{+}$	6.23^{***}_{+}
		(2.46, 480)	(2.76, 960)	(6.20, 480)	(3.82, 960)
Overall	6.45	${}^{***}_{***}8.58$	${}^{***}_{***}4.29$	${}^{***}_{***}13.51$	${}^{***}6.75$
	(3.44, 720)	(2.72, 960)	(2.90, 1920)	(5.76, 960)	(3.68, 1920)

Notes: Standard deviations and number of observations are reported in parentheses
Superscripts *** report Mann-Whitney results. Subscripts $_{***}$ report Panel results.
The **,*** at the left of the mean indicate signicant differences across network treatment with respect to the corresponding value for the dyad; at the right of the mean they indicate signicant differences within network across information treatment.
$^{***}p\leq0.01$; $^{**}p\leq0.05$; $^{*}p\leq0.10$; $^{+}p\leq0.15$

Table 6A. Trustworthiness - Absolute Amount Returned by Receiver

	Dyad		Triad 1S2R				Triad 2S1R			
Treatment	Returned	Returned (zeros incl.)	Tot. to S	Tot. to S (zeros incl.)	Avg. per R	Avg. per R (zeros incl.)	Tot. from R	Tot. from R (zeros incl.)	Avg. to each S	Avg. to each S (zeros incl.)
Full Info	8.98	8.06	${}^{***}_{*}11.19^{***}$	${}^{***}_{**}10.38$	${}^{***}_{+}6.43^{***}$	${}^{***}_{+}5.19$	${}^{***}_{***}21.85$	${}^{***}_{***}19.68^{***}$	${}^{***}11.06^{***}$	${}^{***}_{*}9.84^{***}$
	(7.50, 646)	(7.61, 720)	(5.24, 318)	(5.80, 480)	(4.82, 774)	(5.02, 960)	(12.72, 379)	(12.69, 480)	(7.18, 854)	(7.61, 960)
Partial Info			${}^{***}9.82^{***}$	${}^{***}_{*}9.76$	${}^{***}_{**}5.65^{***}$	${}^{***}_{**}4.88$	${}^{***}_{***}20.46$	${}^{***}_{***}16.09^{***}$	${}^{***}9.74^{***}$	8.05^{***}
			(4.87, 362)	(5.43, 480)	(4.67, 829)	(4.75, 960)	(13.35, 349)	(13.81, 480)	(7.34, 793)	(7.62, 960)
Overall	8.98	8.06	${}^{***}_{*}10.46$	${}^{***}_{***}10.07$	${}^{***}_{**}6.03$	${}^{***}_{**}5.03$	${}^{***}_{***}21.18$	${}^{***}_{***}17.88$	${}^{***}10.42$	${}^{***}_{*}8.94$
	(7.50, 646)	(7.61, 720)	(5.09, 680)	(5.62, 960)	(4.75, 1603)	(4.89, 1920)	(13.04, 728)	(13.38, 960)	(7.29, 1647)	(7.67, 1920)

Table 6B. Trustworthiness - Percentage Amount Returned by Receiver (of Amount Received from Sender)

	Dyad	Triad 1S2R			Triad 2S1R		
Treatment	% Returned	% Tot. to S	% Tot. to S (zeros incl.)	% Avg. per R	% Tot. from R	% Tot. from R (zeros incl.)	% Avg. to each S
Full Info	39.43	40.46^{***}_{***}	40.51^{**}_{**}	39.55^{**}	${}^{***}44.69$	${}^{***}45.52^{**}$	${}^{***}44.74^{+}$
	(27.35, 646)	(17.08, 318)	(17.45, 456)	(23.22, 774)	(23.36, 379)	(24.72, 475)	(25.45, 854)
Partial Info		${}^{**}35.47^{***}_{***}$	36.86^{**}_{**}	${}^{**}35.38^{**}$	${}^{***}42.68$	39.40^{**}	${}^{**}40.29^{+}$
		(15.65, 362)	(17.48, 467)	(21.79, 829)	(22.87, 349)	(23.67, 444)	(23.98, 793)
Overall	39.43	37.81	38.67	37.39	${}^{***}43.73$	${}^{***}42.56$	${}^{***}42.6$
	(27.35, 646)	(16.51, 680)	(17.56, 923)	(22.58, 1603)	(23.13, 728)	(24.40, 919)	(24.85, 1647)

Notes: Standard deviations and number of observations are reported in parentheses
Superscripts *** report Mann-Whitney results. Subscripts $_{***}$ report Panel results.
The **,*** at the left of the mean indicate signicant differences across network treatment with respect to the corresponding value for the dyad; at the right of the mean they indicate signicant differences within network across information treatment.
$^{***}p\leq0.01$; $^{**}p\leq0.05$; $^{*}p\leq0.10$; $^{+}p\leq0.15$

Table 7. Profit

	Dyad			Triad 1S2R			Triad 2S1R		
Treatment	S	R	R net	S	Avg. R	Avg. R net	Avg. S	R	R net by S
Full Info	11.61	21.29	11.29	${}^{***}11.94^{**}_{+}$	${}^{***}_{***}17.46^{***}$	${}^{***}_{***}7.46^{***}$	${}^{***}_{+}12.56^{*}$	${}^{***}_{***}33.98^{***}$	${}^{*}11.99^{***}$
	(6.06, 720)	(8.09, 720)	(8.09, 720)	(4.41, 480)	(6.00, 960)	(6.00, 960)	(6.13, 960)	(14.16, 480)	(8.29, 960)
Partial Info				11.05^{**}_{+}	${}^{***}_{***}18.19^{***}$	${}^{***}_{***}8.19^{***}$	${}^{**}11.81^{*}$	${}^{***}_{***}31.30^{***}$	${}^{*}10.65^{***}$
				(4.59, 480)	(5.77, 960)	(5.77, 960)	(5.51, 960)	(12.79, 480)	(7.89, 960)
Overall	11.61	21.29	11.29	${}^{*}11.49$	${}^{***}_{***}17.83$	${}^{***}_{***}7.83$	${}^{***}12.19$	${}^{***}_{***}32.64$	11.32
	(6.06, 720)	(8.09, 720)	(8.09, 720)	(4.52, 960)	(5.90, 1920)	(5.90, 1920)	(5.84, 1920)	(13.55, 960)	(8.12, 1920)

Notes: Standard deviations and number of observations are reported in parentheses
Superscripts *** report Mann-Whitney results. Subscripts $_{***}$ report Panel results.
The **,*** at the left of the mean indicate signicant differences across network treatment with respect to the corresponding value for the dyad; at the right of the mean they indicate signicant differences within network across information treatment.
$^{***}p\leq0.01$; $^{**}p\leq0.05$; $^{*}p\leq0.10$; $^{+}p\leq0.15$

Table 8. Efficiency - Realized Percentage of Theoretical Maximum Gains

Treatment	Dyad	Triad 1S2R	Triad 2S1R
Full Info	82.24%	$_{***}^{***}$93.73%	84.43%$^{***}_{+}$
	(17.22, 720)	(11.85, 480)	(14.52, 480)
Partial Info		$_{***}^{***}$94.87%	***78.47%$^{***}_{+}$
		(9.83, 480)	(17.71, 480)
Overall	82.24%	$_{***}^{***}$94.3%	*81.45%
	(17.22, 720)	(10.90, 960)	(16.45, 960)

Notes: Standard deviations and number of observations are reported in parentheses Superscripts *** report Mann-Whitney results. Subscripts $_{***}$ report Panel results.

The * at the left of the mean indicate signicant differences across network treatment with respect to the corresponding value for the dyad; at the right of the mean they indicate signicant differences within network across information treatment. $^{***}p\leq0.01$; $^{**}p\leq0.05$; $^{*}p\leq0.10$; $^{+}p\leq0.15$

Figure 1: —— Info - - - - NoInfo

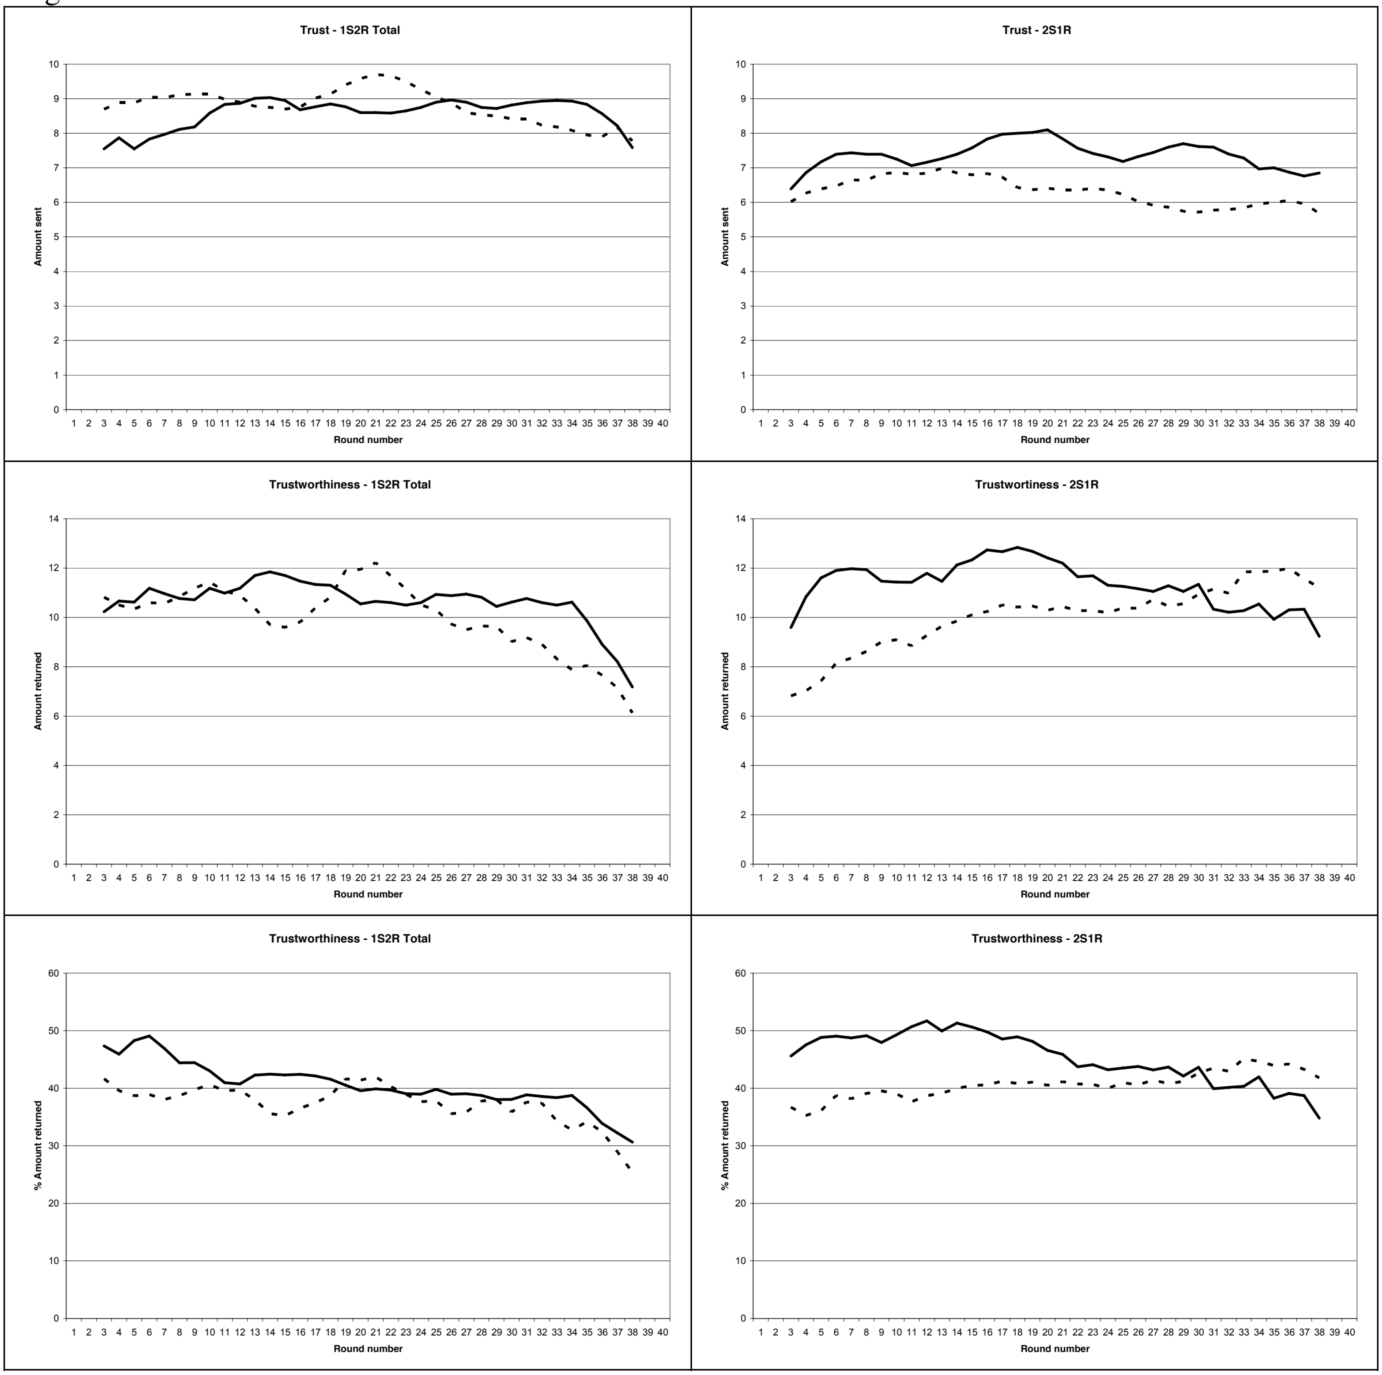

Figure 2: —— 1S1R – – – 1S2R Tot ----- 1S2R Each - - - - 2S1R Each

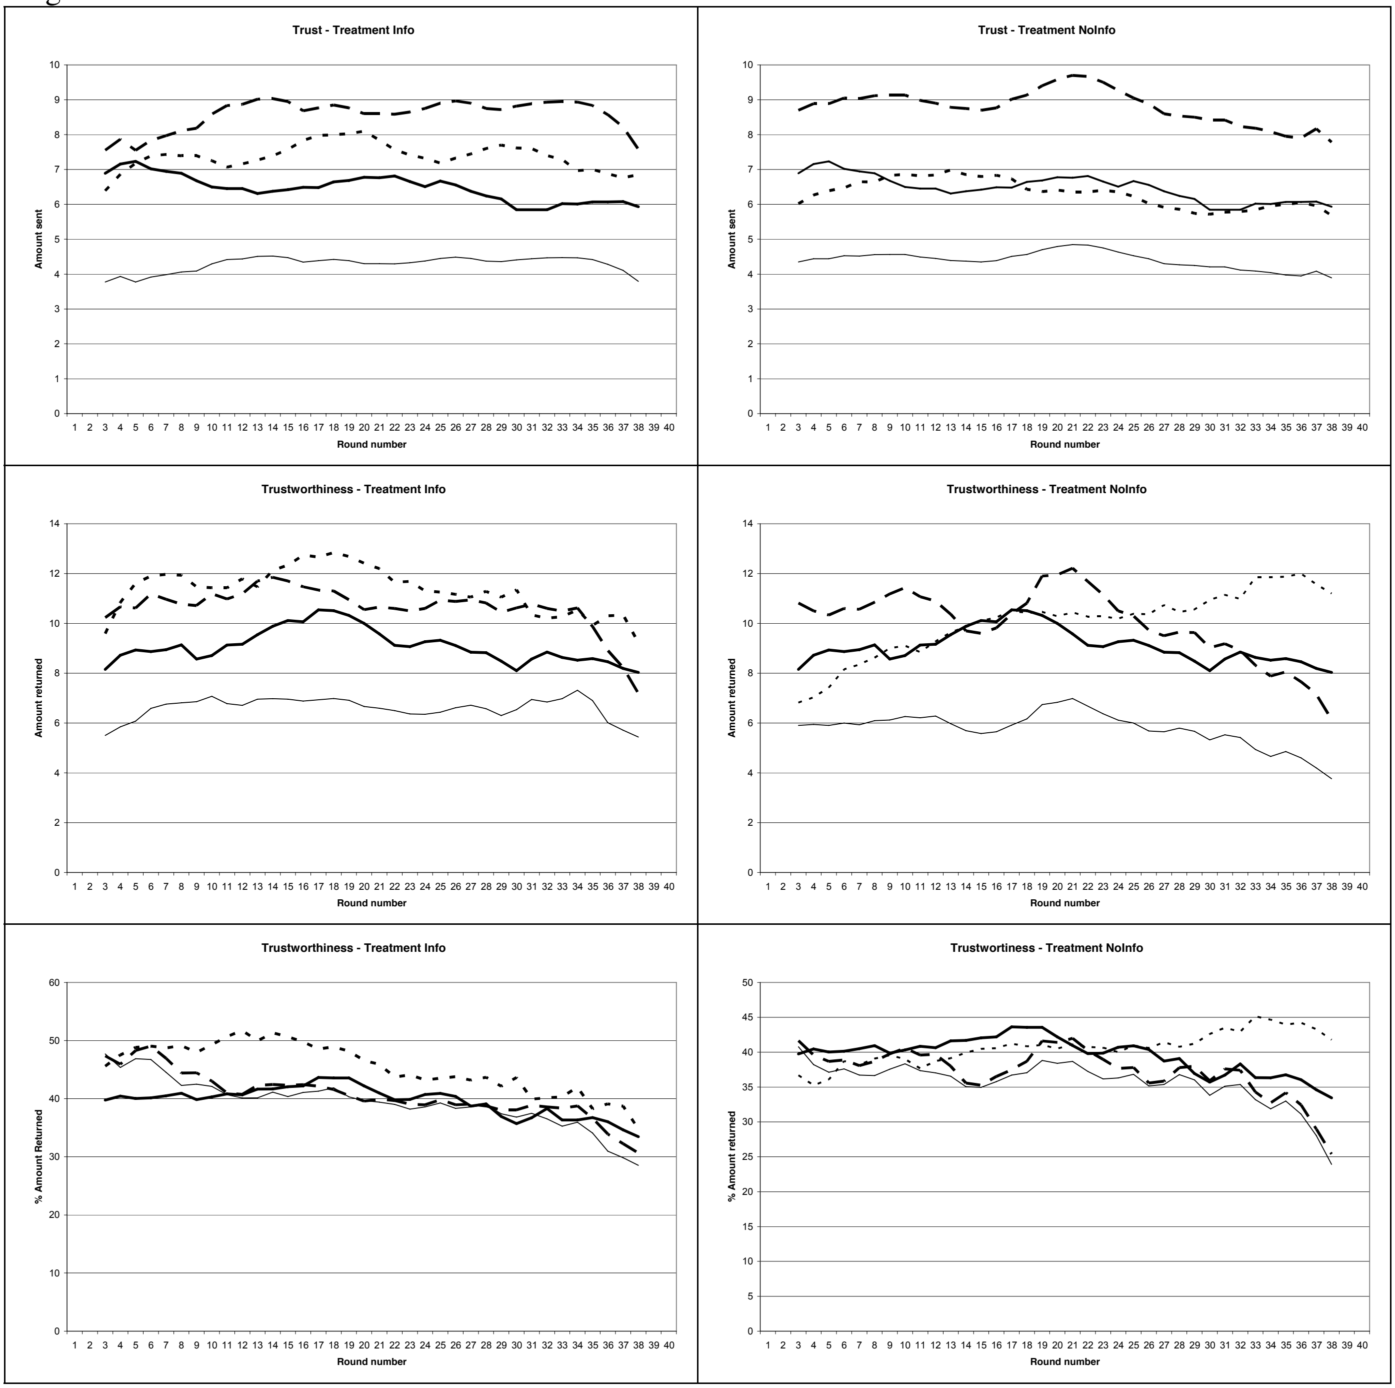

Bibliography

Allen, F. and D. Gale (2000). "Financial Contagion," *Journal of Political Economy*, **108**: 1–33.

Axelrod, R., R. Riolo and M. Cohen (2002). "Beyond Geography: Cooperation with Persistent Links in the Absence of Clustered Neighborhoods," *Personality and Social Psychology Review* **6**(4): 341–346.

Barabasi, A.L. (2002). *Linked: The New Science of Networks* (NY: Perseus).

Berg, J., J. Dickhaut, and K. McCabe (1995). "Trust, Reciprocity, and Social History," *Games and Economic Behavior* **10**: 122–142.

Blume, L. (1995). "The Statistical Mechanics of Best-Response Strategy Revision," *Games and Economic Behavior* **11**: 111–145.

Bohnet, I. and R. Zeckhauser (2004). "Trust, Risk, and Betrayal," *Journal of Economic Behavior and Organization* **55**(4): 464–484.

Calvo-Armengol, T., and M. Jackson (2004). "The Effects of Social Networks on Employment and Inequality," *American Economic Review* **94**(3): 426–454.

Cassar, A. (2007). "Cooperation and Coordination in Local, Random, and Small World Networks: Experimental Evidence," *Games and Economic Behavior* **58**(2): 209–230.

Cassar, A., D. Friedman, and P. H. Schneider (2007)."A Laboratory Investigation of Networked Markets," Working Paper, UCSC.

Cassar, A., and Duffy, N. (2001). "Contagion of Financial Crises under Local and Global Networks," in *Agent-Based Methods in Economics and Finance: Simulations*, F. Luna and A. Perrone (eds.) (Kluwer Academic Publishers).

Cassar, A. and M. Rigdon (2007). "Competition and Cooperation: A Laboratory Study of Trust and Reciprocity in Networks," Working Paper.

Cassar, A. and B. Wydick (2008). "Does Social Capital Matter? Evidence from a Five-Country Group Lending Experiment," Working Paper, University of San Francisco.

Charness, G., M. Corominas-Bosch and G. Frechette (2004). "Bargaining and Network Structure: An Experiment," *Journal of Economic Theory* **136**: 28–65.

Chaudhuri, A. S. Khan, A. Laksmiratan, A. Py and L. Shah (2003). "Trust and Trustworthiness in a Sequential Bargaining Game," *Journal of Behavioral Decision Making* **16**(5): 331–340.

Corbae, D. and J. Duffy (in press). "Experiments with Network Formation," *Games and Economic Behavior.*

Corominas-Bosch, M. (2004). "Bargaining in a Network of Buyers and Sellers," *Journal of Economic Theory* **115**: 35–77.

Cox, J. C., D. Friedman and V. Sadiraj (in press). "Revealed Altruism," *Econometrica.*

Croson, R. and N. Buchan (1999). "Gender and Culture: International Experimental Evidence from Trust Games," *American Economic Review, Papers and Proceedings* **89**(2): 386–391.

Deck, C. and C. Johnson (2004). "Link Bidding in a Laboratory Network," *Review of Economic Design* **8**(4): 359–72.

Eckel, C. and R. Wilson (2004). "Conditional Trust: Sex, Race, and Facial Expression in a Trust Game," Working Paper, Virginia Tech University.

Eisenberg, L.K. (1995). "Connectivity and Financial Network Shutdown," Santa Fe Institute Working Papers.

Ellison, G. (1993). "Learning, Local Interaction, and Coordination," *Econometrica* **61**: 1047–1071.

Engle-Warnick, J., and R. Slonim. (2005). "The Evolution of Strategies in a Trust Game," *Journal of Economic Behavior and Organization* **55**: 553–573.

Fehr, E., U. Fischbacher and S. Gächter (2002). "Strong Reciprocity, Human Cooperation, and the Enforcement of Social Norms," *Human Nature* **13**(1): 1–25.

Fischbacher, U., C. Fong, and E. Fehr (2003). "Fairness, Errors, and the Power of Competition," Institute for Empirical Research in Economics, University of Zürich, Working Paper.

Glaeser, E., D. Laibson, J. Scheinkman, and C. Soutter (2000). "Measuring Trust," *Quartery Journal of Economics* **115**: 811–46.

Granovetter, M. (1985). "Economic Action and Social Structure: The Problem of Embeddedness," *American Journal of Sociology* **91**: 481–510.

Jackson, M. O. and A. Watts (2001). "The Existence of Pairwise Stable Networks," *Seoul Journal of Economics* **14**(3): 299–321.

Kariv, S. and D. Gale (2003). "Bayesian Learning in Social Networks," *Games and Economic Behavior* **45**(2): 329–346.

Karlan, D., M. Mobius, and T. Rosenblat (2005). "Measuring Trust in Peruvian Shantytowns," Working Paper.

McCabe, K., M. Rigdon, and V. Smith (2003). "Positive Reciprocity and Intentions in Trust Games," *Journal of Economic Behavior and Organization* **52**(2): 267–275.

McCabe, K., M. Rigdon, and V. Smith (2002). "Cooperation in Single Play, Two-Person Extensive Form Games between Anonymously Matched Players," in *Experimental Business Research*, R. Zwick and A. Rapoport (eds.) (Boston, MA: Kluwer), pp.49–67.

Mobius, M. and A. Szeidl (2007). "Trust and Social Collateral," Working Paper.

Nowak, M. A. and R. M. May (1993). "The Spatial Dilemmas of Evolution," *International Journal of Bifurcation and Chaos* **3**(1): 35–78.

Nowak, M. A. and R. M. May (1992). “Evolutionary Games and Spatial Chaos,” *Nature*, **359**: 826–829.

Nowak, M. and K. Sigmund (1998). “Evolution of Indirect Reciprocity by Image Scoring,” *Nature* **393**: 573–577.

Ortmann, A., J. Fitzgerald, and C. Boeing (2000). “Trust, Reciprocity, and Social History: A Reexamination,” *Experimental Economics* **3**(1): 81–100.

Raub, W. and J. Wessie (1990). “Reputation and Efficiency in Social Interactions: An Example of Network Effects,” *American Journal of Sociology* **96**: 626–654.

Rauch, J.E., and A. Casella (2001). *Networks and Markets* (NY: Russell Sage Foundation).

Rigdon, M. (2007). “Trust and Reciprocity in Incentive Contracting,” Working Paper.

Samuelson, L., I. Eshel and A. Shaked (1998). “Altruists, Egoists and Hooligans in a Local Interaction Model,” *American Economic Review* **88**: 157–179.

Sen, A. (1997). “Maximization and the Act of Choice,” *Econometrica* **65**(4): 745–779.

Skyrms, B. (2004). *The Stag Hunt and the Evolution of Social Structure* (Cambridge: Cambridge University Press).

Watts, D. (1999). *Small World: The Dynamics of Networks between Order and Randomness* (Princeton, NJ: Princeton University Press).

A Instructions for One-shot [1s-2r], Info Treatment

You have been asked to participate in an experiment. Now that the experiment has started, we ask that you do not talk during the experiment. Please raise your hand if you have a question.

You will be randomly selected to either be a SENDER or a RECEIVER. If you are a SENDER you will be paired with two RECEIVERS, RECEIVER 1 and RECEIVER 2. If you are a RECEIVER you will be paired with one SENDER. You will not be told who these people are either during or after the experiment.

A round consists of two stages. At the beginning of the round, the SENDER begins with 10 points, and each RECEIVER begins with 10 points. In the first stage, the SENDER will have the opportunity to send any part of his/her 10 points (from 0, 1, ..., 10) to RECEIVER 1 and RECEIVER 2. Each point sent will be tripled. Note: The total amount sent to RECEIVER 1 *and* RECEIVER 2 should be less than or equal to 10 points prior to being tripled, or less than or equal to 30 points after being tripled. Each RECEIVER will see the number of points the SENDER sent to him/her as well as to the other RECEIVER.

In the second stage, both RECEIVER 1 and RECEIVER 2 have a decision to make: how many of the tripled points to send back to the SENDER and how many to keep. Once both of the RECEIVERS decide, the SENDER and the other RECEIVER will learn of the decision.

See the next page for a diagram representing the two stages.

You will participate in a total of ONE round.

Your earnings are the total number of points accumulated; points will be converted to U.S. dollars at the rate of 1 DOLLAR for 1 POINT.

Since your decisions are private, we ask that you do not tell anyone your decision either during, or after, the experiment.

Are there any questions before we begin?

Note: Please pay attention to your computer screen. Following the completion of each round, on the results page, there will be a “Finished” button. Be sure to click the button once you have viewed the results.

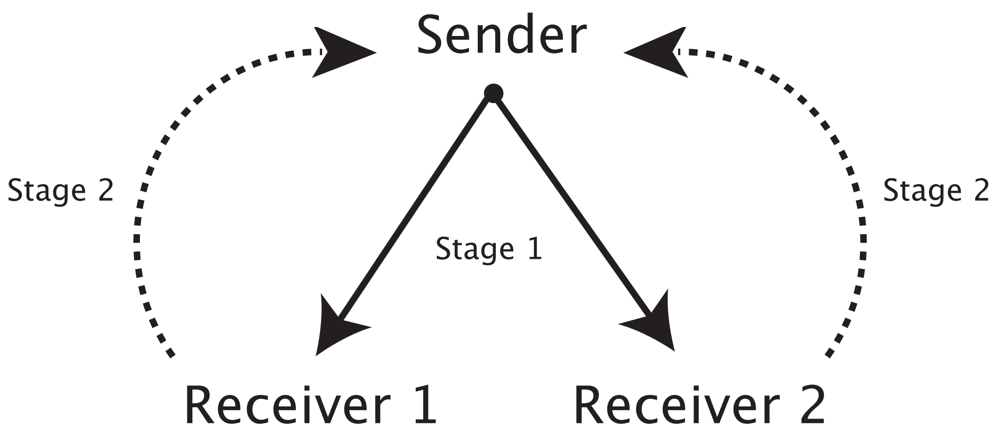
Sender
Stage 2
Stage 1
Stage 2
Receiver 1
Receiver 2

B Instructions for One-shot [1s-2r], No-Info Treatment

You have been asked to participate in an economics experiment. Now that the experiment has started, we ask that you do not talk during the experiment. Please raise your hand if you have a question.

You will be randomly selected to either be a SENDER or a RECEIVER. If you are a SENDER you will be paired with one RECEIVER. If you are a RECEIVER you will be paired with two SENDERS, SENDER 1 and SENDER 2. You will not be told who these people are either during or after the experiment. Each SENDER begins with 10 points, and the RECEIVER begins with 10 points.

A round consists of two stages. At the beginning of the round, the SENDER begins with 10 points, and each RECEIVER begins with 10 points. In the first stage, the SENDER will have the opportunity to send any part of his/her 10 points (from 0, 1, . . . , 10) to RECEIVER 1 and RECEIVER 2. Each point sent will be tripled. Note: The total amount sent to RECEIVER 1 *and* RECEIVER 2 should be less than or equal to 10 points prior to being tripled, or less than or equal to 30 points after being tripled.

In the second stage, both RECEIVER 1 and RECEIVER 2 have a decision to make: how many of the tripled points to send back to the SENDER and how many to keep. Once both of the RECEIVERS decide, the SENDER will learn of the decision.

See the next page for a diagram representing the two stages.

You will participate in a total of ONE round.

Your earnings are the total number of points accumulated; points will be converted to U.S. dollars at the rate of 1 DOLLAR for 1 POINT.

Since your decisions are private, we ask that you do not tell anyone your decision either during, or after, the experiment.

Are there any questions before we begin?

Note: Please pay attention to your computer screen. Following the completion of each round, on the results page, there will be a “Finished” button. Be sure to click the button once you have viewed the results.

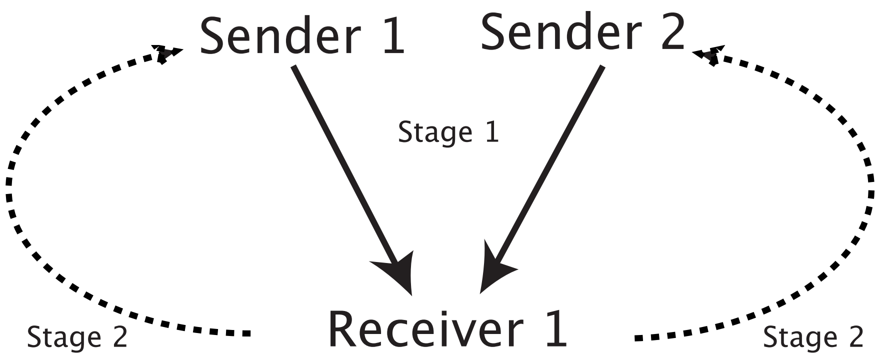
Sender 1
Sender 2
Stage 1
Receiver 1
Stage 2
Stage 2