MOBILIZING INVESTMENT THROUGH SOCIAL NETWORKS: EVIDENCE FROM A LAB EXPERIMENT IN THE FIELD

EMILY BREZA†, ARUN G. CHANDRASEKHAR‡, AND HORACIO LARREGUY⋆

Abstract. In the absence of strong formal institutions, social networks play a significant role in contract enforcement, credit arrangements, and co-investment. We shed light on the relationship between network characteristics and investment decisions through a lab experiment in the field. We focus on the role for third parties to act as informal contract enforcers. Our protocol builds on a basic two-party trust game with a sender and receiver, to which we introduce a third-party to serve as either a monitor or punisher. The ex-ante benefits of a third party judge are ambiguous. On one hand, a third party may result in larger sender transfers due to her ability to punish. On the other hand, the punisher might act in a way to build reputation or may crowd-out intrinsic motivation. Importantly, these costs and benefits of a punisher might vary with her centrality in the network.

Our findings are consistent with both the role for the punisher to induce efficiency and to crowd out intrinsic motivation. They are also consistent with the effects of reputation-building by the punisher. Importantly, we find that very network-peripheral punishers are detrimental to efficiency, while network-central individuals may improve outcomes when given the technology to punish. We also show that these results cannot be explained by either the fact that the punisher also acts as a monitor, or by the punisher's characteristics such as elite status, educational attainment, caste, or proxies for wealth.

JEL Classification Codes: C92, L14

Keywords: Social norms, Social networks, Punishment, Third-party sanctions, Social preferences, Informal contracting

Date: May 2013.
We thank Matt Elliott, Ben Golub, Jacob Leshno, Daron Acemoglu, Abhijit Banerjee, Esther Duflo, and Matthew O. Jackson for detailed discussion. Manaswini Rao and Dhruv Parthasarathy provided excellent research assistance. Special thanks to Juan Pablo Xandri. We are grateful to the Russell Sage Foundation, the George and Obie Shultz Fund, and the MIT Center for International Studies for financial support. Breza and Chandrasekhar thank the National Science Foundation GRFP and Larreguy thanks the Bank of Spain and Caja Madrid Foundation. The social network data was gathered, in part by the authors, as part of a project with Abhijit Banerjee, Esther Duflo, and Matthew O. Jackson.
†Columbia Business School. Email: ebreza@columbia.edu.
‡Microsoft Research, New England and Stanford University. Email: arungc@stanford.edu.
⋆Massachusetts Institute of Technology. Email: larreguy@mit.edu.

1. Introduction

Sociologists and economists have repeatedly demonstrated the importance of social relationships in a variety of human interactions. Given that social interactions do matter, how can organizations, e.g., villages, firms or governments, harness existing social hierarchies to overcome inefficiencies in formal markets? In this paper, we identify social relationships within a society that permit maximal levels of cooperation. Specifically, by studying the behavior of pairs of participants in a simple sender-receiver investment game,[1] to which we add a third-party monitor or punisher, we mimic a co-investment or lending opportunity and shed light on how social networks affect the propensity for individuals to cooperate and enforce informal contracts.

The question of institutional and contract design is particularly important in developing countries. Without strong formal contracting institutions, social structures (networks) are frequently used to mediate economic and political interactions. This is especially true in rural settings where social hierarchies are particularly salient. Common examples of network-based economic relationships include social collateral in microfinance and ROSCAs, informal risk sharing arrangements, and increased prevalence of family firms. While these particular arrangements have been studied at length, (see Feigenberg et al., 2010, Kinnan and Townsend, 2010, and Bertrand and Schoar, 2006 for recent analyses of each) less is understood about the optimal contract structure given network characteristics as inputs. For example, which members of society serve best as third party monitors or punishers and lead to the most efficient outcomes?

We play games modeled after Berg et al. (1995) and Charness et al. (2008) with experimental subjects from 40 South Indian villages. Specifically, we ran each experimental session directly in each of 40 villages and drew the subject pool from the village population. We start with a two-person game where a *sender* (S) transfers money to a *receiver* (R). The transfer increases in size before it reaches the receiver, who then decides how much to return to the sender. In some treatments, we add a third-party punisher (T). Instead of anonymizing participants, the senders, receivers and third parties (where applicable) all sit together to play each game, and are thus able to identify each other before transfer and punishment decisions are

[1] The two-party game is also called a *trust game* in the literature. We view the game as one that naturally mimics co-investment (helping another community member invest in an opportunity that has arisen), for instance.

made. In one treatment the identifiable third party can levy costly punishments on the receivers (and hereafter we refer to such a third party as a *punisher*).

Because these are small villages, there is a high likelihood that any randomly chosen group of participants has non-trivial interaction outside the game, and we are precisely interested in how variation in the relative network position of the third party influences outcomes. In order to separate the mechanisms through which punishers affect experimental outcomes, we also include an experimental treatment with a third party who cannot punish (hereafter referred to as a *monitor*). In this treatment, the monitor merely observes the interaction between the sender and receiver. We combine the experimental results with household survey responses and village network data to determine how co-investment behavior is mediated by network and demographic characteristics.

We are mostly concerned with examining how the network position of a third party with a punishment technology affects the efficiency of the outcomes. Our experiment is designed to separate the effects of monitoring from punishment and to measure the causal effect of giving a punishment tool to different types of individuals in the village. The introduction of the third party may lead to ambiguous effects. The threat of punishment may induce the receiver to return more to the sender and, thus, the sender to transfer more to the receiver. However, the addition of a third party could also crowd out transfers from the sender to the receiver.[2]

We wish to study how the embedding of the sender, receiver and third party in the social network influences the efficiency of the interaction. For example, the eigenvector centrality of a node in a network can parametrize the extent to which an individual is important in an information transmission process. Nodes with higher centrality tend to both acquire more and propagate more information, and other nodes tend to have better information about the characteristics of central nodes (for empirical evidence see Alatas et al. (2010)). Viewing centrality as a reduced form measure of network-importance within each community, we ask what role the centrality of the third party plays in generating higher transfers. For instance, is it the case that when the third party is more central within the network we observe more efficient outcomes? We are also able to address how two categories of demographic characteristics of the judge influence outcomes: caste and whether an individual is a village elite.

[2]Crowd-out has been thoroughly documented in the experimental literature. Examples include Frey and Oberholzer-Gee (1997), Fehr and List (2004), and Fehr and Gächter (2002).

To unpack the role of the third party punisher, we develop a simple model of the three-party game. The model shows that giving the ability to punish to a peripheral judge should result in more inefficient outcomes in the laboratory game, relative to central judges. Essentially, if an individual (in this case the third party) can either be fair-minded/cooperative or uncooperative, then other villagers – who may potentially have repeated interactions with them in the future – would like to know the third party's type. The model uses the fact that individuals are likely to have better quality signals about high centrality judges as opposed to low centrality (or peripheral) judges to show that peripheral third party punishers should generate inefficiencies relative to central punishers. Namely, low centrality punishers will induce the sender to send lower initial transfers.

Empirically, we find that on average, adding a punisher (or a monitor) to a two-party investment game, neither increases nor decreases sender transfers. However, the absence of a level effect masks tremendous heterogeneity in the response to changes in the game structure and in the ability for individuals to achieve more efficient outcomes. We find that network characteristics do interact with the game in meaningful ways. Most notably, as our model would predict, sender transfers increase substantially when the third party is central in the social network and is given the ability to punish. These findings are not consistent with an alternative model where sender and receiver exploit the presence of a central monitor to signal that they are a cooperative type. Relative to the two-party game, adding a punisher who is peripheral in the social network is detrimental to efficiency.

We identify three additional ways in which network characteristics impact the game. First, we corroborate other papers in the literature and find that sender - receiver pairs of close social proximity are able to achieve better outcomes in the two-person game than pairs of socially distant individuals. Second, we show that social proximity can also interact with the punishment technology in negative ways: socially close sender - judge pairs result in lower sender transfers when the judge is given the ability to punish. This suggests that social closeness can be used to improve contracting outcomes, but it is important to consider the potential for collusion in institutional design. Third, in some games and specifications we find that the social importance of the sender is associated with lower sender transfers. While not robust to demographic controls, this third finding is possible support of the hypothesis that peripheral senders use the game as an opportunity to build reputation with the other players.

We show that our demographic characteristics of caste and elite status generally capture different dimensions of power within the village than our network measures. We define elites as *gram panchayat* members, self-help group officials, *anganwadi* teachers, doctors, school headmasters, or owners of the main village shop. Both high caste individuals and elites are afforded special status in their communities. However, in the data, the patterns of network centrality do not match (or even resemble) the patterns of demographic importance. Namely, having an elite or high caste individual as the third party does not significantly increase efficiency when given the ability to punish.

Note that our results differ from previous anonymous laboratory studies such as Charness et al. (2008), which finds that adding a punisher increases average sender transfers, and thus efficiency. We find that while adding an average punisher does not matter in our setting, the network characteristics of that punisher do matter. It is not surprising then that the same game played in an anonymous setting with strangers who have no chance of interacting in the future should produce different patterns when played with members of a tight-knit social network where individuals have a high probability of interacting in the future.

Our results take a step towards understanding how a community might enlist its own social fabric to overcome a lack of formal institutions. To our knowledge, no previous study has used high quality network data to analyze the play of investment games with third parties. Moreover, rural India is the type of setting where network effects should matter most for economic outcomes. Our results also highlight how social connections might have first-order effects when transplanting contracting institutions that work in the (anonymous) lab to the field.

Relevant Literature. Our baseline game builds from the literature started by the Berg et al. (1995) trust game. While the Nash equilibrium has zero transfers for anonymous partners, the authors find that senders make positive transfers and some receivers fully reciprocate. However, senders who transfer tend to lose money on average. Charness et al. (2008) add a third-party punisher and find that senders transfer more and receivers reciprocate to a greater degree than in the case without a punisher. Initial transfers are 60% higher when a punisher is present, significantly increasing total payoffs.

While most experimental games are played with anonymous interactions, a smaller subset of the social preferences literature examines how the outcomes of experimental games change as the level of familiarity between agents is manipulated within the experiment (e.g., Hoffman et al., 1996; Bohnet and Frey, 1999; Burnham, 2003; Charness and Gneezy, 2008).[3] Recently, researchers have begun to combine experimental games with existing network structures. Goeree et al. (2010) elicit peer networks of middle school students and run dictator games with the students and find that dictator offers can largely be explained by inverse social distance. Two experiments to distinguish between motives for giving in a dictator game, namely the extent to which transfers can be explained by altruism as opposed to reciprocity, have looked at networks. Using networks of Harvard students in online dictator games Leider et al. (2009) find social distance effects while studying members of a village in Paraguay, Ligon and Schecter (2012) find effects based on the number of links a household has. The closest paper to ours is Glaeser et al. (2000), who study the investment game with Harvard students. They elicit data on number of months that sender has known the receiver, number of friends they have in common, and demographic characteristics. They do not find any significant effects of similarity of demographic characteristics (e.g., same/different nationality/ethnicity of S and R) on behavior. They find that senders send more and receivers return a greater share if the pair has known each other longer or have more friends in common. Finally, they find that the sender differentially benefits from social status – having more friends makes the sender send less and the receiver send more.

Meanwhile, the network economics literature has developed a rich language to characterize the importance of an individual in the network via measure such as

[3] Hoffman et al. (1996), Bohnet and Frey (1999), Burnham (2003), and Charness and Gneezy (2008) randomly give dictators fairness priming, information prompts, pictures of the receiver, or allow the dictator to see the receiver. All find allocations made by the dictator to increase under these treatments. Bohnet and Frey (1999) also add a treatment where both players visually identify each other and find that dictators are far more likely to split the surplus according to the "fair" 50-50 allocation rule. While these papers give importance evidence that familiarity (or a humanizing force) affects experimental outcomes, they fall short of being able to explain how realistic social dynamics interact with each participant's strategic behavior. In the developing country context, Fehr et al. (2008) and Hoff et al. (2009) use similar games to investigate how caste affects trust and co-investment in India.

eigenvector centrality.[4] This centrality measures typically reflect a node's importance in transmission; more important nodes may be able to better punish others through reputational or social capital channels. Jackson (2008) provides a detailed discussion of the concepts. Empirical network papers employing eigenvector centrality include Hochberg et al. (2007), Banerjee et al. (2011), and Schechter et al. (2011).

Our paper makes several contributions to these literatures. First, we are exactly interested in studying a non-anonymous environment, wherein individuals have deep, preexisting relationships that influence the way they behave. Individuals are called often to interact with community members, sometimes in unanticipated circumstances: e.g., serving on a committee, PTAs, co-investing in a public good. Understanding how the variation in the network position of the actors influences the outcomes of these interactions is important and requires randomly matching individuals as well as obtaining detailed data on the underlying social network of the community. We are in the unique position of having detailed social network data in 40 villages so that we can analyze those nuanced network features suggested by theory to play a role in social interactions. Second, unlike dictator games or two-party interactions, we are specifically interested in questions of institutional design and the role for outside parties to monitor or mediate economic decisions. Our experimental treatments with third parties allow us to ask which network properties must a third party authority possess to generate efficient (or inefficient) behavior? Third, as rural village networks mediate most economic transactions in developing countries (and potentially substitute for more formal institutions or credit markets), it is crucial to understand how barriers to joint investment can be overcome in exactly these types of settings.

Structure of the Paper. The remainder of the paper is organized as follows. In section 2, we describe the experimental subjects, network and survey data sources and the experimental design. Section 3 provides a simple model of third-party punishment. In section 4 we present the results, and section 5 concludes.

[4]Degree is the number of neighbors a node has, eigenvector centrality is a recursively defined measure which defines the centrality of a node as proportional to the sum of its neighbors' centralities, and betweenness centrality computes the share of shortest paths between all pairs of nodes that pass through the node whose centrality we are measuring.

2. Data and Experimental Design

2.1. **Setting.** Our experiment was conducted in 40 villages in Karnataka, India which range from a 1.5 to 3 hour's drive from Bangalore. We chose these villages as we had access to village census demographics as well as unique social network data set, previously collected in part by the authors. The data is described in detail in Banerjee et al. (2011) and Jackson et al. (2010).

The network represents social connections between individuals in a village with twelve dimensions of possible links, including relatives, friends, creditors, debtors, advisors, and religious company. We work with an undirected and unweighted network, taking the union across these dimensions, following Banerjee et al. (2011) and Chandrasekhar et al. (2011). As such, we have extremely detailed data on social linkages, not only between our experimental participants but also about the embedding of the individuals in the social fabric at large.

Moreover, the survey data includes information about caste and elite status. Here local leader or elite is someone who is a *gram panchayat* member, self-help group official, *anganwadi* teacher, doctor, school headmaster, or the owner of the main village shop. Finally, the survey data contains educational attainment as well as proxies for household wealth. We use information such as house size, electrification, building materials, and toilet amenities to construct a ranking within each village.

2.2. **Experiment.** Each participant played five to six total rounds of three experimental treatments.[5] Players were randomly assigned one of three roles in each round: sender (S) with endowment Rs. 60, receiver (R) with endowment Rs. 60, and third party (T) with endowment Rs. 100. A total of 14 surveyors moderated the experiments, each overseeing only one group of two or three participants at a time.

The baseline game (T1) is a two-player investment game with no third-party monitor or judge. The surveyors select two participants at random and assign them to roles of S and R. S can then make a transfer to R, which then triples in size. Finally, R decides how much of his or her wealth from the game to return to S. This transfer does not grow when sent by R. Ending balances are then recorded by the surveyors.

[5]We also attempted, albeit unsuccessfully, to implement a fourth treatment involving having (S, R) pairs interact in anticipation of a T who is not from their village by using a cellular phone.

In the other two treatments, we add third parties who can either monitor or punish. Three players are randomly selected and given roles of S, R, and T. S and R then make the same transfer decisions as in T1. In T2, T watches the transfers take place, but does not take any additional action within the game. In T3, T observes the transfers, and further, has the option to spend his or her own resources to levy a monetary punishment on R. For every Rs. 1 spent by J, R loses Rs. 4.

Each participant played either five or six randomly ordered rounds of the experimental games, including two rounds each of T2 and T3. Half of participants played T1 once and the other half twice. Out of the 5 or 6 total rounds played, participants were each given their ending wealth values for one randomly-chosen round. Participants were given a fixed participation fee of Rs. 20 in addition to their earnings from the game. The average payoff from participating in the experiment was approximately Rs. 110, or approximately half to three-fourths of a daily agricultural wage.

In many settings individuals can choose the individuals with whom they interact. However, there are also plenty of real-world occasions when people find themselves needing to work together with a potentially diverse group of community members. Examples include PTA boards at schools, new business ventures, or new clients. Randomly matching individuals to pairs and triples of individuals in the lab we try to mimic some aspects of the latter types of relationships. Furthermore, in some cases, it is possible to choose the enforcement or governance structure of a group of individuals. Thus, our goal is to detect who is the best at filling the role of the leader or punisher. As such, randomly forming groups of individuals is instructive.

2.3. **Norms.** An investment game has a scaling parameter α. Suppose that the sender transfers τ_S rupees to the receiver. The transfer grows by a factor of α before reaching the receiver. Finally the receiver transfers τ_R back to the sender.

In order to choose α, we proceeded as follows. We designed our experiments to be able to separate between different possible response strategies receivers may be using. There may be natural focal points for how players choose to divide resources among themselves, which may also act as reference points that third parties use when deciding whether or not to punish. In a behavioral economics survey paper, Rabin (1998) discusses several sharing norms. We consider five natural possibilities for how the sender and receiver may choose to share resources:

(1) Keep the Entire Transfer: $\tau_R = 0$.

(2) Keep the Surplus: $\tau_R = \tau_S$.
(3) Split the Transfer: $\tau_R = \frac{\alpha \tau_S}{2}$.
(4) Share the Pie: $\tau_R = \frac{(\alpha+1)}{2}\tau_S$.[6]
(5) Return the Full Surplus: $\tau_R = \alpha \tau_S$.

The choice of $\alpha = 3$ allows us to distinguish between all five cases. Thus, the norms as a fraction of the amount that reaches the receiver are (1) $\frac{\tau_R}{\alpha \tau_S} = 0$, (2) $\frac{\tau_R}{\alpha \tau_S} = \frac{1}{\alpha} = \frac{1}{3}$, (3) $\frac{\tau_R}{\alpha \tau_S} = \frac{1}{2}$, (4) $\frac{\tau_R}{\alpha \tau_S} = \frac{\alpha+1}{2\alpha} = \frac{2}{3}$, (5) $\frac{\tau_R}{\alpha \tau_S} = 1$.[7]

We look empirically for evidence that receivers play (some of) these norms in Section 4.1, and our model in Section 3.2 incorporates the idea that there is a finite set of norms that receivers may play.

2.4. **Descriptive Statistics.** Table 1 presents the descriptive statistics. In each village, 24 individuals between the ages of 18 and 45 were randomly invited to our experiment. All together, 960 individuals participated.[8] The average age is 30 with a standard deviation of 8.1 years. 61% of the participants are female, and the average education level is 8.14 with a standard deviation of 4.31.[9] About 63% of the participants are general or "otherwise backwards" (OBC) caste.[10] Finally, 20% of households have a leader.

Turning to network characteristics, the average social proximity between pairs (the inverse of the social distance) in our experiment is 0.31.[11] The maximum social distance, when it is finite, is 7, and the minimum is 1. 96% of pairs are reachable (there exists a path through the network connecting the two). The average degree (number of friends) is 9.84 with a standard deviation of 6.62, indicating that there is substantial heterogeneity in an individual's number of connections.

[6]Solving $60 - \tau_S + \tau_R = 60 + \alpha\tau_S - \tau_R$ yields the result.
[7]Notice that if $\alpha = 2$, then we would not be able to separate between (2) and (3).
[8]See Banerjee et al. (2012) for an analysis of the diffusion of information about the game.
[9]This means that on average, an individual had attended 8th standard.
[10]There are three standard caste categories in India: general merit (GM); scheduled caste and scheduled tribe (SCST); and other backward caste (OBC). The SCST group is traditionally the most disadvantaged.
[11]Appendix A contains a glossary formally describing the network statistics used.

3. Framework

In this section, we first motivate the choice of the network characteristics – centrality and social proximity– on which we consider heterogeneous treatment effects on. Second, we develop a simple model that captures the effect of the third party punisher's centrality, which is the main focus of our paper.

3.1. **Network Characteristics.** We are interested in how the network relationships between agents impact economic outcomes. Moreover, it may be the case that agents choose members of society to serve as enforcers of contracting norms. As these parties themselves are embedded in the social network, it raises the question of which network characteristics effective judges possess. Given the innumerable ways in which networks may affect economic interactions, for parsimony we focus on two natural network measures are the centrality of each individual and the distance between any two individuals.

The main measure we are interested in is the centrality of the individuals. There are several reasons why centrality should matter in our experiments. We focus our analysis on the centrality of the third party. Note that the third party has to actively take a decision in T3 (either punish or not punish R). Thus, in T3, the third party may gain or lose reputation in the eyes of the other participants based on her punishment decision.[12] The bulk of our analysis uses the eigenvector centrality as the notion of network centrality in our analysis. Eigenvector centrality is a recursive notion of importance wherein an individual's centrality is proportional to the sum of each of her neighbor's centralities.

We are also interested in is that of the inverse social distance (or social proximity). Let γ_{ij} denote the minimum path length between individuals i and j, we define social proximity as γ_{ij}^{-1}. Social proximity is commonly used in the experimental networks literature (e.g., Goeree et al. (2010); Leider et al. (2009)). We should expect that unregulated interactions between pairs of individuals should have more cooperative outcomes for those with high social proximity. The addition of a third party may cause efficiency to either increase or decrease depending on the social proximity of the parties involved. While a punisher might be better able to induce efficiency when the other two individuals are social far, social proximity between a punisher and sender might be detrimental for efficiency.

[12]This is not true in T2 because the third party merely is an observer. In T2 the third party's role is to potentially propagate information outside the experiment.

Finally, we also consider how demographic characteristics (e.g., caste and elite), which might confound with the network measures, interact with contracting between individuals. Two individuals belonging to the same caste group may operate much like social proximity. However, caste also has a power dimension. High caste individuals may be able to exercise power over low caste individuals. Moreover, being a member of the elite in a village could also affect the power dynamic between parties. The experimental predictions are similar to those for centrality. However, elites may be better at resource capture than network leaders.

Throughout the paper, we consider heterogeneous treatment effects based on network position. While these parameters are identified given our experimental design, it is important to use caution when interpreting the results. We caveat that networks are not randomly assigned. People who are central might differ from people who are peripheral on numerous dimensions. We supplement the networks data with covariates that might be correlated with centrality or social proximity such as wealth, leadership status in the community, and caste. Furthermore, in section 4.4, we build a case that the network position is what matters by ruling out these other covariates as the key drivers of our findings.

3.2. **Simple Model of Centrality and Punishment.** We present a simple two-period model of the three-player game to focus on the effect of the third party punisher's centrality. In the first period, two players play the sender-receiver game and the third player acts as a punisher. In the second period, either the sender or receiver plays a coordination game with the punisher. The first-period game represents our laboratory experiment, and the second-period coordination game represents the super-game individuals continue to play after they leave our experiment. Our framework captures the fact that individuals' behavior – in particular the punisher's behavior – may send a signal through actions in the lab that have consequences outside the lab..

We start by describing the strategies and payoffs. We then describe the types of punishers. Finally, we characterize the subgame perfect equilibria of the two-period game where the second-period coordination game provides incentives for punishers to build up reputation in the first-period game. We show that high centrality punishers generate more efficient outcomes – in the sense of higher sender transfers – than low centrality punishers.

Trust Game with Third Party Enforcement. In the first period, a sender (S), a receiver (R), and a judge (T) play the following game. At the beginning of the period, S receives an endowment E and decides how much to transfer to R, which we denote as τ_S. τ_S triples in value, and R receives $3 \cdot \tau_S$. Then, R decides how much to transfer back to S, denoted τ_R, and S receives τ_R. Without loss of generality, we restrict R to play one of three different strategies: $\tau_R \in \{0, \eta_M \cdot \tau_S, \eta_H \cdot \tau_S\}$. R can follow the high norm, η_H, the intermediate norm, η_M, or can return nothing. We assume $\eta_H > \eta_M$.[13]

We define d as the utility cost of a deviation from the high norm. We assume that this cost can take three exogenously determined[14] values, $d \in \{0, d_S, d_L\}$:

$$d = \begin{cases} 0 & \text{if } \tau_R = \eta_H \cdot \tau_{S,} \\ d_S & \text{if } \tau_R = \eta_M \cdot \tau_{S,} \\ d_L & \text{otherwise.} \end{cases}$$

We call d_S a *small norm deviation* and d_L a *large norm deviation*.

Next, T receives an endowment M and decides whether to punish R. The punishment is a binary decision, has a cost c to T, and incurs a penalty of fixed size κ to R. An unpunished deviation from the high norm causes a disutility $\theta_T \cdot d$ to T, where $\theta_T \in \{0, 1\}$. This assumption captures the idea that Ts have heterogeneous types and may experience utility from norm compliance.

The first period payoffs for all three players are as follows,

$$\begin{aligned} U_S &= E - \tau_S + \tau_R, \\ U_R &= 3 \cdot \tau_S - \tau_R - \kappa \cdot P, \\ U_T &= M - \theta_T \cdot d \cdot (1 - P) - c \cdot P, \end{aligned}$$

where it is worth noticing that the disutility that a T of type $\theta_T = 1$ experiences from a high norm deviation is increasing with the size of the deviation.

Additionally, we assume that $d_L > c > d_S$. Under such parameter constraints, in the absence of the second-period game, T would never have incentives to punish small deviations of the high norm in the first-period stage game.

[13]Here we abstract from the discussion in Section 2.3 and assume that receivers play one of three norms.

[14]Allowing for continuous τ_R and endogenous $d = d(\tau_S, \tau_R)$ does not change any of the intuition or conclusions of the model. We introduce this simplification for clarity.

Coordination Game. In the second stage, either S or R plays a coordination game with T:

		S or R	
T	Action	Low	High
	Low	$((1-\theta_T)\cdot L, 0)$	$((1-\theta_T)\cdot L, -\alpha)$
	High	$(-\alpha, 0)$	(β, β)

Payoffs satisfy $\alpha, \beta > 0$, $L > \beta$. Punishers with $\theta_T = 0$ never cooperate because playing *Low* is a dominant strategy. On the contrary, when $\theta_T = 1$, T cooperates when she expects the other player to also cooperate. In this case, (Low, Low) and $(High, High)$ are both Nash Equilibria in this static stage game. Consider an individual S or R who plays the above coordination game with T. Denote λ to be her belief that her counterpart T plays *High*. She will play *High* as long as $\lambda \geq \frac{\alpha}{\alpha+\beta}$. For future exposition, we denote $\frac{\alpha}{\alpha+\beta} = \gamma$.

Types of Players. For simplicity, we focus on the case where only T's have type heterogeneity. We denote π_T to be the prior that $\theta_T = 1$. This prior π_T is equal to the population share of high types. Additionally, we assume that T may be central or peripheral in the social network, and that individuals know the type θ_T with certainty for central Ts. In contrast, individuals have uncertainty about the type of peripheral Ts, and thus have a prior π_T over the probability that $\theta_T = 1$ for any given peripheral T. We assume that $\pi_T < \gamma$.

Timing. The timing of the two-period game is as follows,

- $t = 0$: Nature draws $T's$ type, θ_T, and whether she is peripheral or central
- $t = 1$: S, R and T play the trust game with third party enforcement,
  - S receives E and transfers τ_S,
  - R receives $3 \cdot \tau_S$ and transfers τ_R,
  - S receives τ_R,
  - T decides whether to punish R, and
  - period payoffs are realized.
- $t = 1.5$: S or R update π_T to form λ.
- $t = 2$: Either S or R play the coordination game with T,
  - S or R and T decide whether to play *Low* or *High*, and
  - period payoffs are realized.

Characterization of the Game. We look for subgame perfect equilibria, and characterize the game through backward induction.

Characterization of the Coordination Game in $t = 2$. As noted above, S or R's best response is to play $High$ as long as $\lambda \geq \gamma$. Thus, there are two possible cases, summarized in the following lemma.

Lemma 1. *If T is peripheral and $\lambda < \gamma$ or T is central and $\theta_T = 0$, then the unique stage-game Nash Equilibrium is S or R and T both play Low. However, if T is peripheral and $\lambda \geq \gamma$ or T is central and $\theta_T = 1$, then there exists a stage-game Nash Equilibrium such that S or R and T both play $High$.*

The proof is straightforward and omitted. Note that the stage game equilibrium where S or R and T coordinate on $High$ is the efficient static Nash Equilibrium of the stage game. We focus on the equilibrium of the two-period game where cooperation can be sustained, and consequently, where there are incentives for T to build up reputation in the first-period game.[15]

Characterization of T's Strategy in the Trust Game with Third-Party Enforcement. If nature draws $\theta_T = 0$, then $T's$ dominant strategy is to not punish regardless of τ_R. If nature draws $\theta_T = 1$, then we distinguish between the cases when T is central and when T is peripheral. When T is central, her type is already known and thus has no incentives to build reputation in the first period game. Given $d_L > c > d_S$, a central T's dominant strategy (when $\theta_T = 1$) is to only punish large deviations from the high norm.

The interesting case arises when T is peripheral and $\theta_T = 1$. In such a case, given that $d_L > c$, a peripheral T also punishes large deviations. However, T might also now have incentives to punish small deviations in the first stage in order to build reputation for the second period and separate from the $\theta_T = 0$ types.

When does such a separating equilibrium exist? The characterization of the third party's strategy in the period-one game is summarized in Lemma 2.

Lemma 2. *Assume that $d_L > c > d_S$, $d_S + \beta > c$,* and $\pi_S, \pi_R \geq \gamma > \pi_T$, then given the second-stage equilibrium strategies in Lemma 1, there exists an equilibrium punishment strategy where:[16]

[15] In all sub-cases, (Low, Low) is also a Nash Equilibrium of the second period stage game.

[16] There also exists a SPE where nobody ever cooperates in the second stage, and nobody ever punishes in the first stage.

(1) Any T with $\theta_T = 0$ never punishes, regardless of τ_R.
(2) Central T with $\theta_T = 1$ only punish large deviations from the high norm.
(3) Peripheral T with $\theta_T = 1$ punish large *and* small deviations from the high norm.

Proof. See Appendix B. □

Characterization of the Sender's and Receiver's Strategies in the Trust Game with Third-Party Enforcement. There are three cases we should consider: a) when S and R face a central T with $\theta_T = 1$, b) when S and R face a central T with $\theta_T = 0$, and c) when S and R face a peripheral T of unknown θ_T. Lemma 3 characterizes the strategies of S and R in the period-one game.

Lemma 3. *Assume that* $d_L > c > d_S$*,* $d_S + \beta > c$*,* $\pi_S, \pi_R \geq \gamma > \pi_T$, and $\eta_M > 1$. Then given the equilibrium punishment rule from Lemma 2:

(1) If S and R face a central T of type $\theta_T = 1$, S transfers $\dfrac{\kappa}{\eta_M}$ to R and R transfers κ back to S.
(2) If S and R face a central T of type $\theta_T = 0$, S transfers 0 to R and R transfers 0 back to S.
(3) If S and R face a peripheral T of unknown type, S transfers $\dfrac{\pi_T \cdot \kappa}{\eta_L}$ to R and R transfers $\pi_T \cdot \kappa$ back to S.

Proof. See Appendix B. □

The following proposition characterizes the strategies and outcomes of the separating equilibrium of the two-period game.

Proposition 4. *Assume that* $d_L > c > d_S$*,* $d_S + \beta > c$*,* $\pi_S, \pi_R \geq \gamma > \pi_T$, and $\eta_M > 1$.

(1) *If S and R face a central T of type $\theta_T = 1$, S transfers* $\dfrac{\kappa}{\eta_M}$ to R *in the period-one game*, *R transfers κ back to S, and T does not punish R. In the period-two game, S or R and T play $(High, High)$.*
(2) *If S and R face a central T of type $\theta_T = 0$, S transfers* 0 to R *in the period-one game*, *R transfers* 0 *back to S, and T does not punish R. S or R and T play (Low, Low).*
(3) *If S and R face a peripheral T, S transfers* $\dfrac{\pi_T \cdot \kappa}{\eta_H}$ to R *in the period-one game*, *R transfers $\pi_T \cdot \kappa$ back to S, and T does not punish R. S or R and T play (Low, Low).*

It is worth noting that central T's of type $\theta_T = 1$ are better able to provide incentives to cooperate than peripheral T's, and hence, they enhance the efficiency of the two-period game outcome. While a central T of type $\theta_T = 1$ is able to induce a transfer $\frac{\kappa}{\eta_M}$, a peripheral T is only able to induce a transfer $\frac{\pi_T \cdot \kappa}{\eta_H}$, where $\frac{\kappa}{\eta_M} > \frac{\pi_T \cdot \kappa}{\eta_H}$, since $\frac{\eta_H}{\pi_T} > \eta_H > \eta_M$.

Further, this implies that average sender transfers across all peripheral T's will be lower than the average sender transfers across all central T's. When T is central, the average transfer is $\bar{\tau}_S^{central} = \frac{\pi_T \cdot \kappa}{\eta_M}$ while $\bar{\tau}_S^{peripheral} = \frac{\pi_T \cdot \kappa}{\eta_H}$. Under our assumptions, $\bar{\tau}_S^{central} > \bar{\tau}_S^{peripheral}$.

4. Results

4.1. **Pooled Equilibrium Play.**

Before analyzing the treatment effects and network effects, it is helpful to first observe the overall outcomes from the experimental sessions. The data include 1,988 total games, and Figure 1 shows the distribution of initial transfers from S to R observed in all games pooled together. Almost all transfers are made in increments of Rs. 5 or Rs. 10.[17] The modal transfer is 20, with the mean occurring at Rs. 28.5. A zero transfer is only observed in 13 of the games. The efficient transfer of Rs. 60 is observed 122 times (~6% of games).

Moving to the receiver's response, Figure 2 shows the pooled distribution of transfers from R to S as a fraction of the initial transfer from S to R. Note that most of the receivers transfer weakly less than the amount sent by the sender, leaving receivers with quantities at least as high as their initial endowments[18]. Only 5% of games ended with the receiver sending more back to the sender than was initially transferred. Also note that there are two transfer levels with notably high frequencies occurring at $\frac{\tau_R}{\alpha\tau_S} = \frac{1}{3}$ and $\frac{\tau_R}{\alpha\tau_S} = \frac{2}{3}$. These values correspond to norms 2 and 4, "keep the surplus" and "split the pie." The receivers seem to adhere to some notion of fairness as described in the norms of section 2.3. The mean level of $\frac{\tau_R}{\alpha\tau_S}$ is approximately 0.5. Note that while, on average, both S and R gain relative to their initial endowments, approximately 25% of senders are worse off in monetary terms than if they had played the static Nash Equilibrium, $\tau_S = 0$.

[17]Participants could make transfers in increments of Rs. 1.
[18]At least before the punishment decision is made.

Figure 3 provides an alternate illustration of R's average response to S.[19] The graph plots a nonparametric approximation of $\frac{\tau_R}{\alpha\tau_S}$ as a function of τ_S. Surprisingly, very small initial transfers are rewarded with large return transfers (statistically indistinguishable from sending everything back). However, as $\tau_S > 20$ the overall relationship between initial transfer and amount returned is remarkably stable at approximately 0.5 in equilibrium.

The equilibrium punishments incurred by the judges in T3 can also teach us about the acceptable transfer norms in the participating villages. Figure 4 shows incurred punishments as a fraction of transfers returned from R to S. On the interval from 0 to 1, punishment is decreasing as a function of the fraction returned to the sender, as would be expected from a norm-enforcer. Returning nothing is associated with an average punishment amount of Rs. 10. This expected punishment declines dramatically as $\frac{\tau_R}{\alpha\tau_S}$ approaches 1. Above 1, punishment appears to be increasing, but is very noisy. In this range, punishment enforces an unfair outcome for receivers; their final payoffs are lower than their initial endowments.

These outcomes show that while players in the role of S tend to transfer amounts substantially greater than zero, most games are quite far from the efficient outcome. Further, sender transfers are quite heterogeneous. These outcomes also indicate the that receivers tend to focus on two of the norms from Section 2.3 and that when observed, punishment is decreasing in the size of the receiver's transfer, both of which are captured in our model. We next move to understand the extent to which the contracting structure and the social network can help S, R pairs to achieve more efficient outcomes and can help to explain the heterogeneity of game outcomes.

4.2. **Treatment Level Effects.** We begin by analyzing the game outcomes by treatment. Table 2 presents the payoffs and sender transfers by treatment. In each specification, the omitted treatment is the two-player game. Results in column 1 indicate that we cannot reject that the game in which T can only observe, but not punish, has different total payoffs than the baseline. However, the game in which T can both monitor and punish decreases total payoffs by Rs. 9.97. In game 3, the average punishment level is 8.28 and can mostly explain the decrease in payoffs. Columns 2 and 3 show the payoffs separated by S and R. The entire difference in total payoffs (column 1) across treatments is borne by R , which is again consistent

[19]We note that any relationship between player behavior and τ_S is endogenous. Therefore the plots in Figures 3 and 4 are descriptions of the equilibrium and are not causal effects. They ought to be interpreted with caution.

with the monetary punishments eroding payoffs. Column 4 looks at how the initial τ_S, which is a measure of efficiency, responds to treatments. None of the treatments has statistically distinguishable effects relative to the baseline. Columns 5 to 8 indicate that results in columns 1 to 4 are robust to the inclusion of several controls.

In light of our model, there are two opposing effects that come into play when we add a third party who can punish. On the one hand, the punisher should be better able to enforce the return of larger transfers, in turn encouraging higher initial transfers. On the other hand, low centrality judges may try to build reputation by punishing, even for very small norm violations. This type of behavior could therefore completely offset the positive effect of the punisher. Further, the presence of a third party may crowd out any pro-social behavior observed in the two-party games.

Note that the Charness et al. (2008) games are played with anonymous agents, so the social proximity of agents is 0, the relative network centralities of players can be thought of as 0, and the value of future interactions in the super-game can also be thought of as 0. In contrast, our experiments are played in a non-anonymized environment in which agents are entirely socially connected. Our networks exhibit small-world phenomena; the average proximity of senders and receivers is high (.32). Individuals have many opportunities to interact with one another outside of our laboratory games. Consequently, any network effects on game behavior are likely to be extremely salient and influence the main effects in our data. Relative to our results, we can think of the Charness et al. (2008) data as coming from socially distant pairs and triples of individuals who all have extremely high centrality in the network and who have no reputation-building motives with one-another. Therefore, the anonymous, socially distant punisher does not have any reputation-building incentives and only plays the role of norm enforcer. In our games, however, the punisher has two separate incentives for intervening in the game, and the sign of the effect on overall efficiency is ambiguous.

4.3. **Network Importance and Sender Transfers .** We now address the central theme of our paper: how social networks affect the ability for participants in an investment game to cooperate, and how giving punishment technologies to central versus peripheral individuals affects the efficiency of outcomes. We focus on the play of the senders, as they determine the efficiency of the outcomes. Table 3 displays our main network findings. We consider measures of network importance (top of table) and measures of network proximity (bottom of table).

4.3.1. *Centrality.* The top portion of Table 3 shows how the transfers of the sender change with the eigenvector centrality of the players.[20] We center our exposition on the columns that include the network characteristics of the judge, because those allow us to measure the causal effect of introducing a punishment technology.

As a preview of our main result, Figure 5 provides evidence that more central third party punishers are associated with higher transfers from S to R. Columns 2, 5, 8 and 11 of Table 3 confirm that, using two different measures of centrality, on average, the most central judge induces S to transfer approximately Rs. 3.0 more to R than the least central judge. Columns 3, 6, 9 and 12 indicate that the effect of judge centrality is only present in the treatment in which the judge has the ability to punish. In these specifications, the game with monitoring only is the omitted category. This result is in line with the theoretical prediction of the model that central judges are better able to provide incentives for cooperation. In the game with only monitoring, the judge does not take an action, and thus does not have any scope for building reputation. Our results confirm that this motive, which is captured in our model, is only present when the judge is given a punishment technology.

Moreover, we find that more central senders send less to receivers, a result which appears quite robust across all specifications with no demographic controls. While our simple model in 3.2 was not designed to capture reputation motives of the S and R, we can use the same logic to explain how the centrality of S affects transfers. If senders use their transfers to build reputation, we should observe that transfers are decreasing in the centrality of S because again, signals sent by peripheral individuals are more.

4.3.2. *Proximity.* While our biggest contribution is demonstrating the potential to employ high centrality third parties to possibly improve informal contracting outcomes, we also consider the role of social proximity in encouraging cooperation. Our results echo those of other related studies and suggest that social proximity does help to foster more efficient outcomes.

To preview of our results on social distance Figure 6 shows the total payoffs of S and R are increasing in sender-receiver social proximity. The bottom portion of Table 3 shows how the transfers of the sender change with the social proximity of the

[20]We include specifications with different sets of controls, and we also evaluate the regressions with different measures of eigenvector centrality: centrality quartile and an indicator for high versus low centrality.

players. Columns 1, 4, 7 and 10 provide evidence that, in the two-party game (T1), an increase in the social proximity between S and R corresponds to an increased transfer from S to R. S transfers approximately Rs. 8 more to R if they are at distance one as opposed to being socially unconnected, though the coefficient is only marginally significant.

The rest of the columns, which include the social proximity between the punisher and both the sender and the receiver, reflect the results for the games where we introduce a punisher (T2 and T3). We also find evidence that in T3 as opposed to T2, social proximity between S and T induces the sender to transfer less to the receiver. This appears to provide evidence for collusion between the sender and the judge. A sender-judge pair at social distance one corresponds to the sender transferring between Rs. 10 and Rs. 12 less to the receiver than a sender-judge pair who are not socially connected. The effect is only present in T3 where the judge is able to take an action but not the case in T2 where the judge can only monitor and no in-game action is required. Overall, while social proximity may improve outcomes between S and R, it appears that social proximity between the players and the judge may also undermine the punishment institution.

4.4. **Demographic Characteristics and Robustness.** Because the social networks in these 40 study villages are not randomly assigned, it is natural to ask whether our network effects are driven by other demographic characteristics that happen to be correlated with network position. We take three approaches. First, we show that the demographic and network characteristics pick up different dimensions of variance in a principal component analysis. Second, we show that other demographic characteristics that may be correlated with the network such as caste or elite status cannot replicate the patterns observed with the network characteristics. Third, we show that our main networks results do not change even if we control for all available demographic characteristics.

Additionally, four of the demographic characteristics in our data, elite status, high caste, wealth, and education, may represent power or a notion of hierarchy in the study villages. Therefore, we check if these three variables are driving our observed centrality effects.

Principal Component Analysis. In Table 4, we present a principal component decomposition of the importance characteristics. The decomposition contains five different measures of importance: eigenvector centrality, elite status, high caste, wealth, and educational attainment. The five variables separate along three distinct dimensions. Caste, wealth and education are all key contributors to the first principal component, eigenvector centrality is the main constituent of the second principal component, and elite status appears to be its own dimension in the third principal component. This suggests that network centrality does have content distinct from the other demographic characteristics.

Even though these measures may be correlated, the principal component decomposition suggests that the demographic measures are distinct from network centrality.

Demographic Characteristics. Panel A of Table 5 presents results for sender behavior as a function of the elite status of the participants. While there is no detectable effect of elite status on transfers in the baseline two-player game (columns 1 and 4), elite status does weakly affect how the game is played in the treatments with third parties (columns 2, 3, 5, and 6). Columns 2 and 4 indicate that, in the pooled games with a third-party, senders who are elites send approximately Rs 2 less to receivers. Moreover, columns 3 and 6 suggest that such an effect is driven by T2. This gives some evidence that resources are perhaps directed towards elites who exhibit their power in the presence of third-parties that cannot punish them.

Importantly for our central result, whether the judge is an elite does not affect sender transfers in either of the treatments with a third party, which suggests that it is unlikely that the results on punisher centrality are driven by the fact that more central people might belong to the elite.

The effects of caste composition on sender transfers are displayed in Panel B of Table 5. Because we only have caste information for a subset of villages, our analysis is quite underpowered. However, the results certainly do not suggest that the presence of a high caste punisher contributes to the sender's transfer.

Education and proxies for wealth also cannot replicate the result that adding an important punisher improves game outcomes. We do not include the regression results here, but they are available upon request. (See discussion below).

Robustness to Controls. While Table 5 indicate that caste and elite status cannot explain the effects of the centrality of the judge and the social proximity of the sender and receiver, we explore an expanded set of possible importance and proximity measures in Table 6. In the table, we use the same regression specification as in

Table 3, but we also include a full set of demographic controls, interacted with the treatment status where appropriate.

A natural covariate that might capture power relationships aside from elite status is wealth. To proxy for wealth, we construct a within-village ranking of households based on a principal component analysis of the size, construction materials, electrification, and type of toilet facilities in their homes. As one might expect, our wealth quantile ranking does correlate with the household's eigenvector centrality quantile.[21] Further, we include the educational attainment of the players in addition to our measures of elite status and caste.

Despite the positive relationship between centrality and the various other covariates that might capture power relationships, the results in Table 6 indicate that the effects of the centrality of the judge on sender transfers are robust to controlling for those covariates. The only effects that no longer survive when we add such controls is that of the sender's centrality.

Finally, note that we run these extended specifications using three different functional forms of eigenvector centrality as regressors. We continue to use the centrality quantile of each player in the village in columns 1 to 3, and an indicator for above-median eigenvector centrality in the experimental sample in columns 4 to 6. Additionally, we use the level of eigenvector centrality in columns 7 to 9. We find that the results are quite similar for all sets of specifications.

4.5. **Evaluating Institutional Design.** Finally, we can ask which combinations of contract enforcement mechanisms and network characteristics produce the most efficient sender transfers. Figure 7 plots sender transfers[22] for 10 different game configurations. Panel A includes two-party games (left-most in each grouping) and three-party games with monitors. Panel B includes the same two-party games (left-most in each grouping) alongside results from the games with punishers. We further consider cases where S and R are of close social proximity (left groupings) versus far social proximity (right groupings) and cases where the third-party judge is of high centrality (middle bar in each grouping) versus low centrality (right-most bar in each grouping). The bar charts illustrate many of our key networks results but also allow for comparisons between the three- and two-party games.

[21]The wealthiest household has a centrality ranking 13 percentage points higher than the poorest household. The relationship is significant at all standard levels.

[22]Normalized by the average sender transfer across all of the games.

The bar charts reinforce the result that in the two-party game, outcomes are better when the sender and the receiver are socially close (although in this specification, the difference is not significant). Another striking pattern is that in the games with a monitor, neither social closeness nor judge centrality appears to affect sender transfers. When S and R are socially close, even in the games with a punisher, don't produce results different from the average transfer. However, the identity of the punisher is extremely important when S and R are socially far. In these cases, when the punisher is peripheral in the network, sender transfers are significantly lower than the two-party outcome. However, when the punisher is central in the network, transfers are marginally significantly higher than the two-party outcome.

These results suggest that when the contracting parties are socially close, they can sustain reasonably good outcomes without outside intervention. However, when the contracting parties are socially distant, third parties who have the ability to take punitive actions may improve outcomes, so long as that individual is chosen carefully. In our setting, the best outcomes with socially far contracting pairs occur when the individual with the punishment technology is socially important.

5. Conclusion

We conduct laboratory experiments in the field with non-anonymized participants from real-life social networks to understand how different contracting environments affect the outcomes of joint investment games. We use detailed network data to further analyze how the social network characteristics of participants interact with the contracting environments to shape final payoffs. Our games are played among individuals from rural Indian villages, who can fully identify each other, thus making all past and future interactions between the participants relevant for how they play our games.

We focus on the role for third-party punishers to improve outcomes and explore how the network centrality of the punisher impacts sender transfers. We find that the punishment technology can help to improve outcomes, but only when the punisher is central in the social network. We show that the monitoring function of the third party cannot explain our results, nor can other demographic characteristics which may be stand-ins for importance in the community. Our findings are consistent with the model we develop where socially peripheral punishers may use their actions in the game to build reputation about their types in anticipation of future interactions outside of the game.

Our results provide a first analysis of how the local social network interacts with institutional design and informal contracting. However, we have only scratched the surface of this problem. Further work is neccesary to take these ideas from the laboratory and put them into practice.

References

Alatas, V., A. Banerjee, A. G. Chandrasekhar, B. Olken, and R. Hanna (2010): "Identifying the Poor: Social Networks and the Aggregation of Information," .

Banerjee, A., E. Breza, A. G. Chandrasekhar, E. Duflo, and M. Jackson (2012): "Come play with me: Experimental evidence of information diffusion about rival goods," *Work in Progress*.

Banerjee, A., A. G. Chandrasekhar, E. Duflo, and M. Jackson (2011): "The Diffusion of Microfinance," MIT working paper.

Berg, J., J. Dickhaut, and K. McCabe (1995): "Trust, reciprocity, and social history," *Games and Economic Behavior*, 10, 122–142.

Bertrand, M. and A. Schoar (2006): "The role of family in family firms," *The Journal of Economic Perspectives*, 20, 73–96.

Bohnet, I. and B. Frey (1999): "Social distance and other-regarding behavior in dictator games: Comment," *American Economic Review*, 335–339.

Burnham, T. (2003): "Engineering altruism: A theoretical and experimental investigation of anonymity and gift giving," *Journal of Economic Behavior & Organization*, 50, 133–144.

Chandrasekhar, A. G., C. Kinnan, and H. Larreguy (2011): "Informal Insurance, Social Networks, and Savings Access: Evidence from a lab experiment in the field," *MIT Working Paper*.

Charness, G., R. Cobo-Reyes, and N. Jiménez (2008): "An investment game with third-party intervention," *Journal of Economic Behavior & Organization*, 68, 18–28.

Charness, G. and U. Gneezy (2008): "What's in a name? Anonymity and social distance in dictator and ultimatum games," *Journal of Economic Behavior & Organization*, 68, 29–35.

Fehr, E. and S. Gächter (2002): "Do Incentive Contracts Undermine Voluntary Cooperation?" *Mimeo*.

Fehr, E., K. Hoff, and M. Kshetramade (2008): "Spite and development," *World*.

Fehr, E. and J. A. List (2004): "The hidden costs and returns of incentives - trust and trustworthiness among CEOs," *Journal of the European Economic Association*, 2, 743–771.

Feigenberg, B., E. Field, and R. Pande (2010): "Building Social Capital Through MicroFinance," *NBER Working Papers*.

Frey, B. S. and F. Oberholzer-Gee (1997): "The cost of price incentives: An empirical analysis of motivation crowding-out," *The American economic review*, 87, 746–755.

Glaeser, E., D. Laibson, J. Scheinkman, and C. Soutter (2000): "Measuring Trust," *Quarterly Journal of Economics*, 115, 811–846.

Goeree, J., M. McConnell, T. Mitchell, T. Tromp, and L. Yariv (2010): "The 1/d law of giving," *American Economic Journal: Microeconomics*, 2, 183–203.

Hochberg, Y., A. Ljungqvist, and Y. Lu (2007): "Whom you know matters: Venture capital networks and investment performance," *The Journal of Finance*, 62, 251–301.

Hoff, K., M. Kshetramade, and E. Fehr (2009): "Caste and Punishment: The Legacy of Caste Culture in Norm Enforcement," .

Hoffman, E., K. McCabe, and V. Smith (1996): "Social distance and other-regarding behavior in dictator games," *The American Economic Review*, 86, 653–660.

Jackson, M., T. Barraquer, and X. Tan (2010): "Social Capital and Social Quilts: Network Patterns of Favor Exchange," .

Jackson, M. O. (2008): *Social and Economic Networks*, Princeton University Press.

Kinnan, C. and R. Townsend (2010): "Kinship and Financial Networks, Formal Financial Access and Risk Reduction," *manuscript, Northwestern University and MIT*.

Leider, S., M. M. Möbius, T. Rosenblat, and Q.-A. Do (2009): "Directed Altruism and Enforced Reciprocity in Social Networks*," *Quarterly Journal of Economics*, 124, 1815–1851.

Rabin, M. (1998): "Psychology and economics," *Journal of economic literature*, 36, 11–46.

Schechter, L., A. Yuskavage, and U. Treasury (2011): "Reciprocated Versus Unreciprocated Sharing in Social Networks," *Working Paper*.

Figures

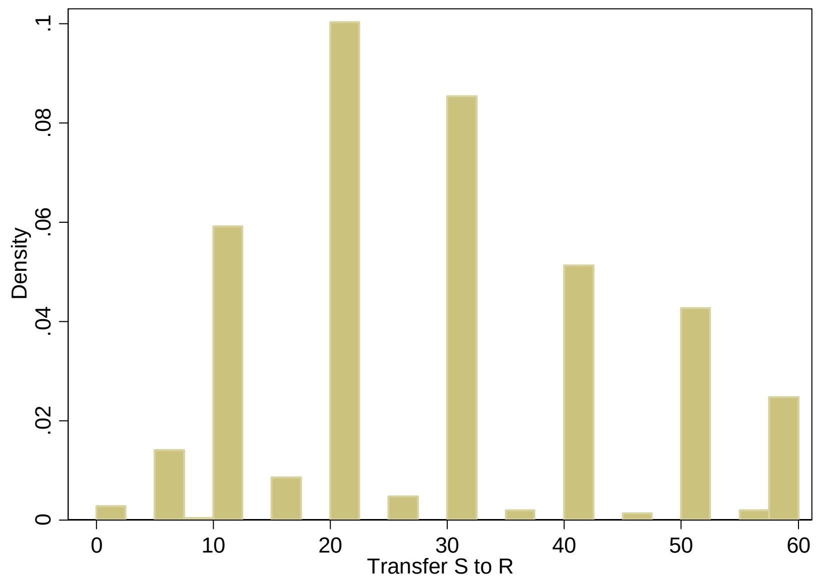

Figure 1. Distribution of transfers from sender

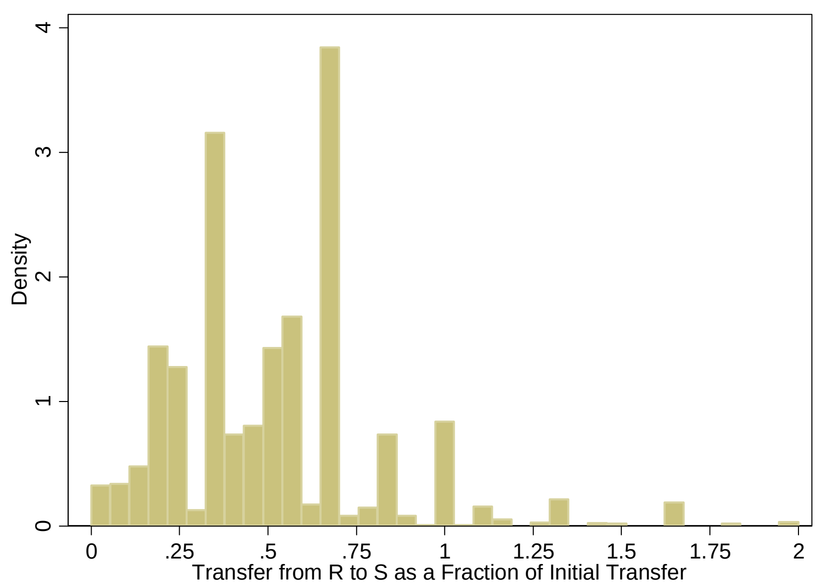

Figure 2. Distribution of transfers from receiver to sender as a fraction of the initial transfer

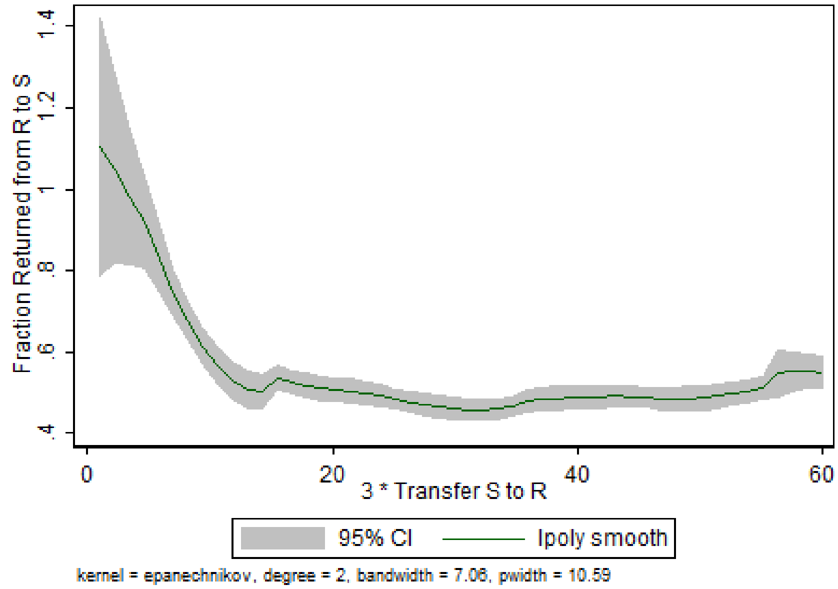

FIGURE 3. Fraction returned from R to S as a function of $\alpha \times$transfer S to R

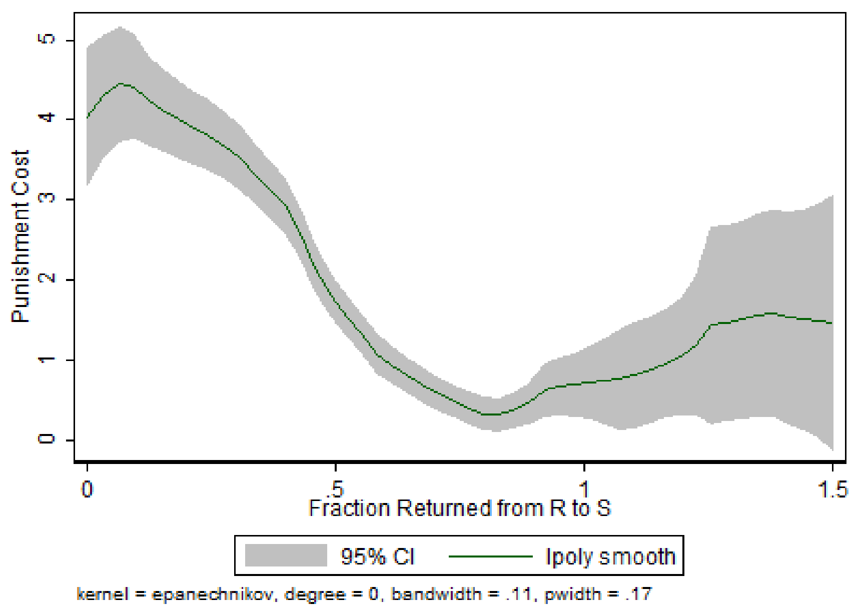

FIGURE 4. Punishment cost paid by T by fraction returned from R to S

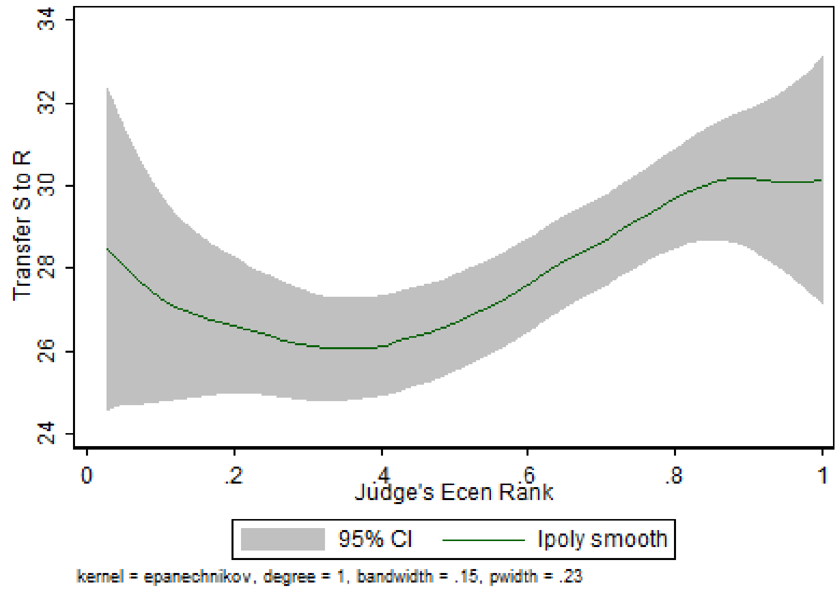

FIGURE 5. Transfer from S to R as a function of the percentile of the eigenvector centrality of T.

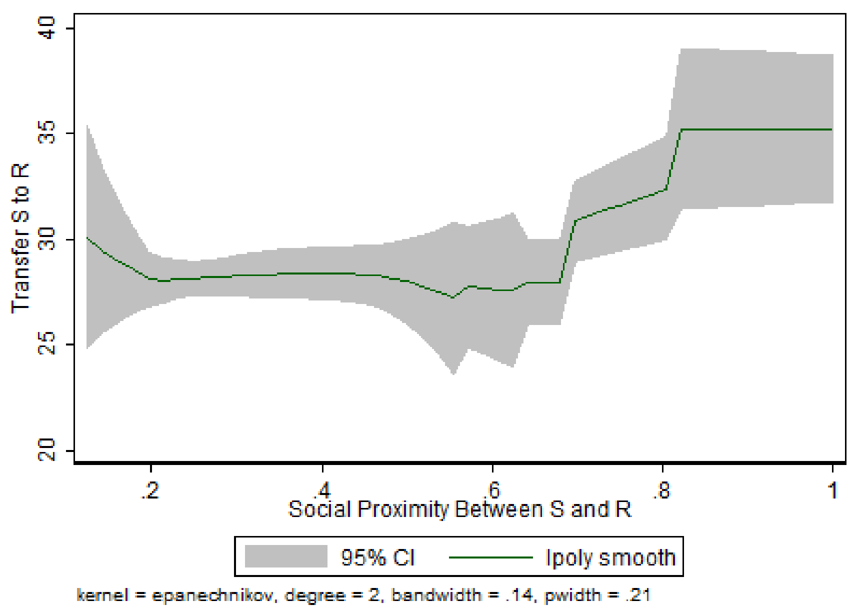

FIGURE 6. Transfer from S to R as a function of social proximity between S and R

Panel A: Third-Party Monitors

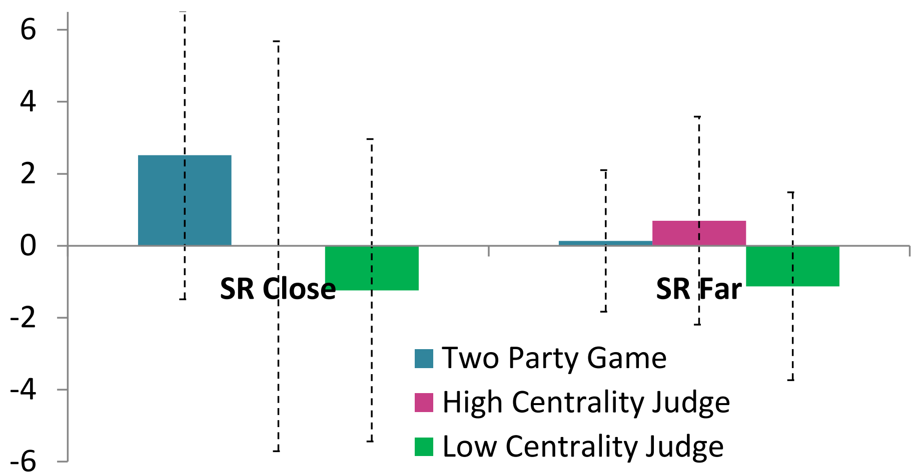

Panel B: Third-Party Punishers

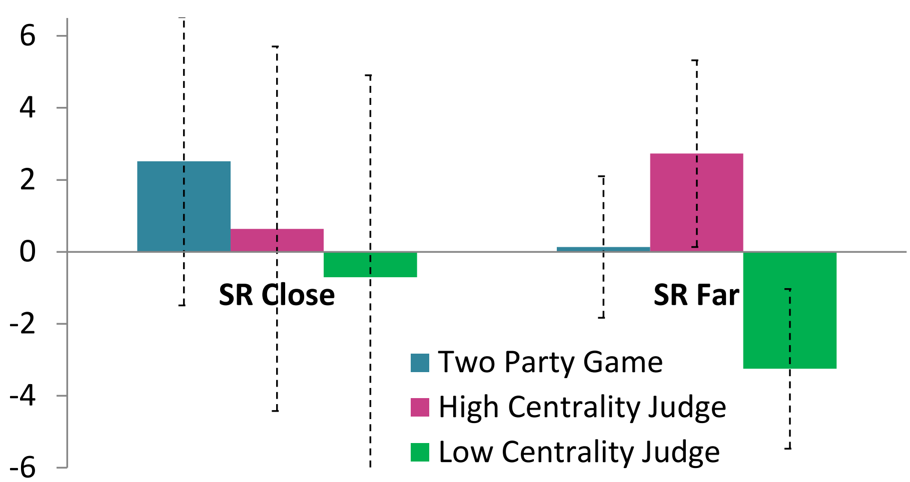

In all bar charts, the y-axis represents the average transfer from the sender to the receiver, normalized by the average transfer size. In each grouping, the left-most bar shows transfers in the two-party game, the middle bar shows transfers in the three-party game with a judge of high centrality, and the right-most bar shows transfers in the three-party game with a judge of low centrality.

Figure 7. Normalized Sender Transfers by Game and Punisher Characteristics

Tables

Table 1. Summary Statistics

	Mean	Std. Dev.
Age	29.95	8.14
Female	0.61	0.49
Education	8.14	4.31
High Caste	0.63	0.48
HH has a Leader	0.20	0.40
Average Proximity b/w Pairs	0.31	0.17
Average Reachability b/w Pairs	0.96	0.20
Average Degree	9.84	6.62
Average Eigenvector Centrality	0.02	0.04
Average Betweenness Centrality	0.00	0.01

Table 2. Sender Behavior and Total Payoffs

Outcome:	Total Payoffs	Payoff S	Payoff R	Transfer S to R	Total Payoffs	Payoff S	Payoff R	Transfer S to R
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Game w/ Monitoring	0.838	1.354	-1.180	-0.0270	2.054	0.843	0.624	0.655
	(2.735)	(1.386)	(2.247)	(1.320)	(2.769)	(1.371)	(2.516)	(1.345)
Game w/ Monitoring and Punishment	-9.969***	1.839	-12.47***	-1.117	-9.248***	1.545	-11.29***	-0.641
	(2.580)	(1.525)	(1.927)	(1.208)	(2.628)	(1.754)	(2.314)	(1.247)
Controls	No	No	No	No	Yes	Yes	Yes	Yes
Mean of Dep. Var. (2 Person Game)	177.31	73.693	104.11	28.944	177.305	73.693	104.11	28.944
Observations	1,988	1,986	1,986	1,987	1,984	1,982	1,982	1,983
R-squared	0.229	0.164	0.096	0.216	0.245	0.177	0.109	0.235

Standard errors are clustered at the Room Level. Columns (1) to (4) only include round fixed effects. Columns (5) to (8) include controls for block, order, round , and surveyor fixed effects. *** p<0.01, ** p<0.05, * p<0.1

Table 3. Sender's Transfers and Network Characteristics

		Transfer S to R	Transfer S to R	Transfer S to R	Transfer S to R	Transfer S to R	Transfer S to R	Transfer S to R	Transfer S to R	Transfer S to R	Transfer S to R	Transfer S to R	Transfer S to R
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Network Importance Measures	J Centrality		3.121*	-1.277		3.180*	-1.462		2.786***	1.107		2.879***	1.055
			(1.791)	(2.256)		(1.749)	(2.128)		(0.924)	(1.163)		(0.907)	(1.132)
	J Centrality * Punishment			8.643***			9.094***			3.186**			3.447**
				(2.948)			(2.917)			(1.524)			(1.521)
	S Centrality	-5.088**	-2.066	0.229	-5.576**	-1.616	0.692	-2.879**	-1.248	-0.0887	-3.046**	-0.999	0.126
		(2.472)	(2.083)	(2.733)	(2.336)	(2.059)	(2.747)	(1.303)	(1.081)	(1.446)	(1.280)	(1.073)	(1.468)
	S Centrality * Punishment			-5.357			-5.391*			-2.690			-2.601
				(3.289)			(3.236)			(1.645)			(1.644)
	R Centrality	-1.122	1.122	3.047	-1.448	0.831	2.902	-1.455	0.948	1.645	-1.510	0.714	1.504
		(3.383)	(2.004)	(2.740)	(3.171)	(2.004)	(2.782)	(1.495)	(0.902)	(1.377)	(1.450)	(0.883)	(1.402)
	R Centrality * Punishment			-3.791			-4.122			-1.470			-1.699
				(3.321)			(3.480)			(1.704)			(1.770)
Network Proximity Measures	Social Proximity S & R	7.834	2.893	-0.384	8.220*	3.059	-0.627	7.706*	2.884	0.942	7.787*	3.127	0.847
		(4.708)	(2.996)	(4.298)	(4.499)	(3.174)	(4.452)	(4.057)	(2.775)	(3.964)	(3.995)	(2.948)	(4.123)
	Social Proximity S & R * Punishment			7.505			8.202			5.104			5.681
				(5.751)			(5.856)			(5.624)			(5.716)
	Social Proximity S & J		-1.550	5.567		-1.424	5.255		-1.735	4.374		-1.628	4.025
			(2.679)	(4.514)		(2.635)	(4.427)		(2.568)	(4.129)		(2.536)	(4.118)
	Social Proximity S & J * Punishment			-12.82**			-12.25**			-11.09**			-10.50**
				(5.397)			(5.134)			(5.054)			(4.908)
	Social Proximity R & J		-3.756	-0.398		-3.832	-1.189		-4.374	-1.722		-4.464	-2.553
			(2.818)	(3.779)		(2.750)	(3.820)		(2.803)	(3.756)		(2.755)	(3.822)
	Social Proximity R & J * Punishment			-6.208			-4.680			-4.946			-3.338
				(4.839)			(5.032)			(4.942)			(5.134)
	Game with Punishment		-1.096	2.874		-0.939	2.127		-1.070	2.874		-0.900	2.132
			(1.079)	(3.059)		(1.088)	(2.969)		(1.071)	(2.379)		(1.082)	(2.255)
	Controls	Yes	Yes	Yes	No	No	No	Yes	Yes	Yes	No	No	No
	Eigenvector Centrality Measure	Quantile	Quantile	Quantile	Quantile	Quantile	Quantile	High/Low	High/Low	High/Low	High/Low	High/Low	High/Low
	Observations	672	1,173	1,173	675	1,174	1,174	672	1,173	1,173	675	1,174	1,174
	R-squared	0.331	0.346	0.357	0.312	0.332	0.343	0.333	0.351	0.359	0.314	0.337	0.345

In columns (1), (4), (7), and (10) only the game with no judge is included, and in columns (2), (3), (5), (6), (8), (9), (11), and (12) only games with any judge are included. Standard errors are clustered at the room level. All specifications contain room fixed effects. Columns (1) - (3) and (7) - (9) include controls for block, order round , and surveyor fixed effects. Columns (4) - (6) and (10) - (12) include round fixed effects. *** p<0.01, ** p<0.05, * p<0.1

Table 4. Principal Component Decomposition of Importance Measures

	Principal Components		
	1st PC	2nd PC	3rd PC
Centrality Quantile	-0.0207	0.7639	-0.1986
Elite	0.3516	0.2717	0.8931
High Caste	0.5711	0.1428	-0.326
Wealth Quantile	0.6113	0.1174	-0.2363
Education	0.4196	-0.5554	0.0299
Eigenvalue	1.6034	1.2217	0.8881

The columns represent the 1st, 2nd, and 3rd principal components in a principal component decomposition.

Table 5. Sender's Transfers, Elite Status, and Caste

		Transfer S to R	Transfer S to R	Transfer S to R	Transfer S to R	Transfer S to R	Transfer S to R
		(1)	(2)	(3)	(4)	(5)	(6)
	Panel A						
Demographic Importance Meas.	Judge's HH Elite		0.458	0.622		0.309	0.343
			(1.202)	(1.548)		(1.179)	(1.508)
	Judge's HH Elite * Punishment			-0.387			-0.137
				(2.171)			(2.164)
	Sender's HH Elite	0.167	-1.856*	-3.022**	0.250	-1.716*	-3.112**
		(1.481)	(0.973)	(1.298)	(1.499)	(0.978)	(1.347)
	Sender's HH Elite * Punishment			2.275			2.723
				(2.056)			(2.081)
	Receiver's HH Elite	1.472	1.428	1.407	1.331	1.389	1.433
		(1.763)	(0.900)	(1.396)	(1.723)	(0.900)	(1.375)
	Receiver's HH Elite * Punishment			0.0926			-0.0197
				(2.146)			(2.184)
	Game with Punishment		-1.332	-1.718		-1.128	-1.636
			(1.042)	(1.110)		(1.052)	(1.131)
	Controls	Yes	Yes	Yes	No	No	No
	Observations	699	1,240	1,240	702	1,241	1,241
	R-squared	0.326	0.325	0.326	0.309	0.312	0.313
	Panel B						
Demographic Importance Meas.	J High Caste		1.366	2.650		0.990	3.748
			(3.615)	(4.819)		(2.994)	(3.493)
	J High Caste * Punishment			-1.240			-4.108
				(5.717)			(5.296)
	S High Caste	1.935	-2.448	1.740	3.431	-3.075	1.918
		(3.967)	(3.864)	(5.319)	(3.938)	(3.611)	(5.617)
	S High Caste * Punishment			-7.679			-9.131
				(7.058)			(7.446)
	R High Caste	1.867	-4.183	-3.516	1.967	-4.283	-3.880
		(3.318)	(3.830)	(4.313)	(2.833)	(3.856)	(4.422)
	R High Caste * Punishment			-1.047			-0.874
				(5.872)			(6.257)
	Game with Punishment		-2.779	3.306		-2.986	5.625
			(3.480)	(6.009)		(2.908)	(6.234)
	Controls	Yes	Yes	Yes	No	No	No
	Observations	178	171	171	179	171	171
	R-squared	0.413	0.368	0.379	0.259	0.301	0.325

In columns (1) and (4) only the game with no judge is included, and in columns (2), (3), (5), and (6) only games with any judge are included. Standard errors are clustered at the room level. All specifications in Panel A contain room fixed effects, while specifications in Panel B contain village fixed effects. This is because of the small sample size in the caste regressions. Columns (1) - (3) include controls for block, order round , and surveyor fixed effects. Columns (4) - (6) include round fixed effects. *** p<0.01, ** p<0.05, * p<0.1

Table 6. Robustness to Demographic Controls and Eigenvector Centrality Measure

		Transfer S to R (1)	Transfer S to R (2)	Transfer S to R (3)	Transfer S to R (4)	Transfer S to R (5)	Transfer S to R (6)	Transfer S to R (7)	Transfer S to R (8)	Transfer S to R (9)
Network Importance Measures	J Centrality		3.908*	-0.230		3.090***	1.546		30.06*	-9.211
			(2.003)	(2.577)		(0.998)	(1.298)		(16.89)	(26.13)
	J Centrality * Punishment			9.476***			3.349*			71.90**
				(3.273)			(1.695)			(28.92)
	S Centrality	-4.553*	-1.327	1.387	-2.595*	-0.881	0.437	-25.17	-7.332	-9.746
		(2.598)	(2.240)	(2.875)	(1.413)	(1.163)	(1.524)	(18.05)	(13.84)	(14.88)
	S Centrality * Punishment			-5.240			-2.647			5.089
				(3.494)			(1.707)			(23.88)
	R Centrality	-1.349	1.226	3.412	-1.268	1.087	1.965	-19.28	11.16	26.28
		(3.482)	(2.150)	(3.005)	(1.525)	(0.981)	(1.520)	(22.86)	(15.41)	(16.52)
	R Centrality * Punishment			-4.203			-1.742			-31.41
				(3.451)			(1.800)			(25.36)
Network Proximity Measures	Social Proximity S & R	8.921*	2.655	0.698	8.675**	2.687	1.918	8.171*	2.182	1.416
		(4.900)	(3.036)	(4.397)	(4.336)	(2.879)	(4.120)	(4.371)	(2.760)	(3.858)
	Social Proximity S & R * Punishment			5.510			3.490			2.366
				(6.179)			(5.976)			(5.745)
	Social Proximity S & J		-1.357	5.652		-1.220	5.017		-1.395	6.324
			(2.922)	(4.813)		(2.849)	(4.405)		(2.711)	(4.568)
	Social Proximity S & J * Punishment			-11.94**			-10.28*			-12.42**
				(5.951)			(5.514)			(5.822)
	Social Proximity R & J		-4.473	-0.584		-4.978	-1.943		-3.617	-0.603
			(3.044)	(4.162)		(3.024)	(4.129)		(2.868)	(3.822)
	Social Proximity R & J * Punishment			-7.269			-5.504			-5.452
				(5.195)			(5.335)			(4.803)
	Game with Punishment		-1.142	2.836		-1.133	3.527		-1.178	3.571
			(1.116)	(4.839)		(1.105)	(4.379)		(1.119)	(4.156)
	Basic Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
	Demographic Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
	Eigenvector Centrality Measure	Quantile	Quantile	Quantile	High/Low	High/Low	High/Low	Level	Level	Level
	Observations	641	1,111	1,111	641	1,111	1,111	641	1,111	1,111
	R-squared	0.330	0.363	0.382	0.332	0.368	0.384	0.330	0.363	0.380

In columns (1), (4), and (7) only the game with no judge is included, and in columns (2), (3), (5), (6), (8), and (9) only games with any judge are included. Standard errors are clustered at the room level. All specifications contain village fixed effects due to small samples of caste and co-household observations. All specifications include controls for block, order round , and surveyor fixed effects. In columns (1) - (3), the centrality measure is the eigenvector centrality percentile within the village. In columns (4) - (6), the centrality measure is an indicator for above-median eigenvector centrality in the experimental sample. Columns (7) - (9) use the eigenvector centrality level. *** p<0.01, ** p<0.05, * p<0.1

Appendix A. glossary of Network Statistics

In this section we briefly discuss the network statistics used in the paper. Jackson (2008) contains an extensive discussion of these concepts which the reader may refer to for a more detailed reading.

Path Length and Social Proximity. The *path length* between nodes i and j is the length of the shortest walk between the two nodes. Denoted $\gamma(i,j)$, it is defined as $\gamma(i,j) := \min_{k \in \mathbb{N} \cup \infty} \left[A^k\right]_{ij} > 0$. If there is no such walk, notice that $\gamma(i,j) = \infty$. The *social proximity* between i and j is defined as $\gamma(i,j)^{-1}$ and defines a measure of how close the two nodes are with 0 meaning that there is no path between them and 1 meaning that they share an edge. In figure 8, $\gamma(i,j) = 2$ and $\gamma(i,k) = \infty$.

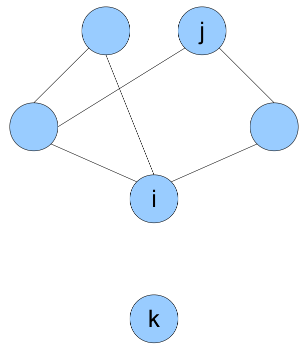

FIGURE 8. Path lengths i, j and i, k

Vertex characteristics. For completeness we discuss three basic notions of network importance from the graph theory literature: degree, betweenness centrality, and eigenvector centrality. The *degree* of node i is the number of links that the node has. In figure A(a), i has degree 6 while in (b) i has degree 2. While this is an intuitive notion of graphical importance, it misses a key feature that a node's ability to propagate information through a graph depends not only on the sheer number of connections it has, but also how important those connections are. Figure A(b) illustrates an example where it is clear that i is still a very important node, though a simple count of its friends does not carry that information. Both betweenness centrality and eigenvector centrality address this problem.

The *betweenness centrality* of i is defined as the share of all shortest paths between all other nodes $j, k \neq i$ which pass through i. This is a normalized measure which is useful when thinking about a propagation process traveling from node j to k as taking the shortest available path.

The *eigenvector centrality* of i is a recursive measure of network importance. Formally, it is defined as the ith component of the eigenvector corresponding to the maximal eigenvalue of the adjacency matrix representing the graph.[23] The intuition for its construction is that one may be interested in defining the importance of a node as proportional to the sum over each of its network neighbor's importance. By definition the vector of these importances must be an eigenvector of the adjacency matrix and restricting the importance measure to be positive means that the vector of importances must be the first eigenvector. Intuitively, this measure captures how

[23]The adjacency matrix A of an undirected, unweighted graph G is a symmetric matrix of 0s and 1s which represents whether nodes i and j have an edge.

well information flows through a particular node in a transmission process. Relative to betweenness centrality, a much lower premium is placed on a node being on the exact shortest path between two other nodes. We can see this by comparing figure A(b), where i has a high eigenvector centrality and high betweenness, to (c), where i still has a rather high eigenvector centrality but now has a 0 betweenness centrality since no shortest path passes through i.

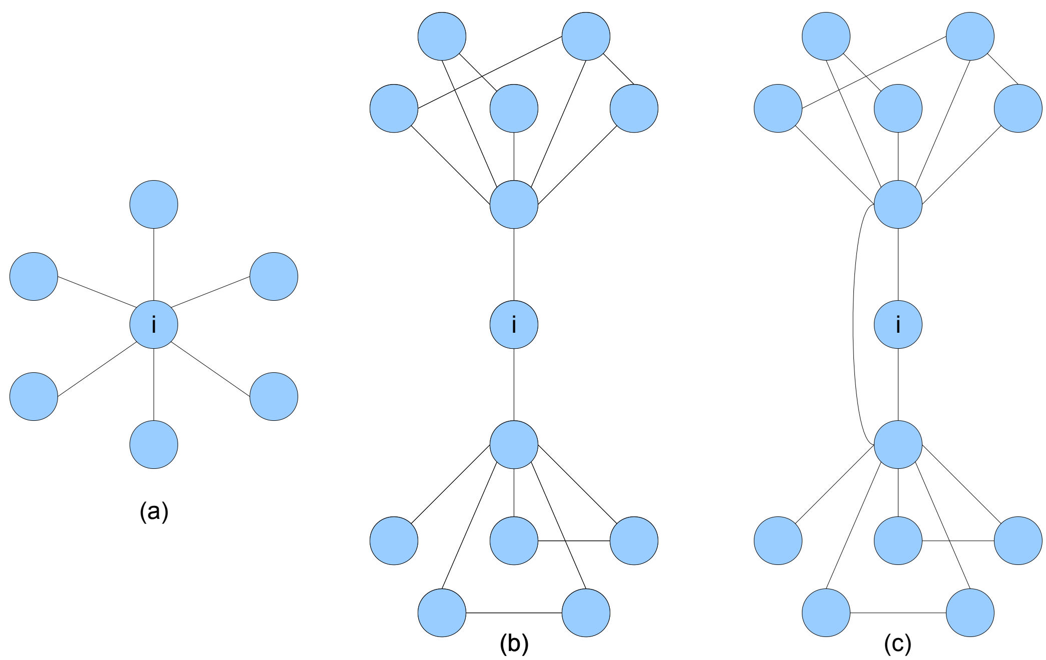

FIGURE 9. Centrality of node i

APPENDIX B. PROOFS

Proof of Lemma 2. Consider the case where peripheral Ts with $\theta_T = 1$ do not punish small deviations. In this case, there is a pooling equilibrium where both types of peripheral Ts do not punish small deviations. Consequently, there is no update about the type of peripheral Ts, λ=θ_T, and their payoff in the period two game is 0. Alternatively, consider the case where Ts with $\theta_T = 1$ do punish small deviations. In this case, there is a separating equilibrium between both types of peripheral. Thus, there is an update about the type of peripheralTs with $\theta_T = 1$, λ=1$>\gamma$, and their payoff in the period two game is β. Accordingly, peripheralTs with $\theta_T = 1$ would like to separate from peripheralTs with $\theta_T = 0$ when the benefits outweigh the costs, that is if $d_S + \beta > c$. □

Proof of Lemma 3. In case a), R anticipates that she is punished only if she transfers 0. The payoff of such a strategy is $3 \cdot \tau_S - \kappa$. On the contrary, R can avoid being punished if she transfers $\eta_M \cdot \tau_S$, consequently receiving a payoff of $(3 - \eta_M) \cdot \tau_S$. Accordingly, R transfers $\eta_M \cdot \tau_S$, as long as $\tau_S \leq \dfrac{\kappa}{\eta_M}$. Otherwise, R transfers 0 and anticipates a punishment of κ. In anticipation of R's strategy, WLOG S decides whether to transfer 0 or to transfer κ.[24] If S transfers $\dfrac{\kappa}{\eta_M}$, she receives a payoff of $E - \dfrac{\kappa}{\eta_M} + \kappa$. Otherwise, S transfers 0 and receives a payoff of E. Thus, S transfers $\dfrac{\kappa}{\eta_M}$ as long as $\eta_M > 1$.

In case b), when there is a central T of type $\theta_T = 0$, the solution is trivial. R expects no punishment regardless how much she sends back to S, and consequently, she transfers 0 to S. In anticipation, S transfers 0 to R.

In case c), there is a peripheral T, who is of type $\theta_T = 1$ with probabilityπ_T and of type $\theta_T = 0$ with probability $(1 - \pi_T)$. Consequently, R anticipates that with probability $(1 - \pi_T)$ she receives no punishment regardless how much she sends back to S. Additionally, R expects that with probabilityπ_T she will receive a punishment if she transfers either 0 or $\eta_M \cdot \tau_S$ back to S, but no punishment if she transfers $\eta_H \cdot \tau_S$. Using the fact that $\eta_M \cdot \tau_S$ is a dominated strategy relative to transferring 0 and a similar reasoning to the central T with$\theta_T = 1$ case, R transfers $\eta_H \cdot \tau_S$, as long as $\tau_S < \dfrac{\pi_T \cdot \kappa}{\eta_H}$. Otherwise, R transfers 0 and anticipates an expected punishment of $\pi_T \cdot \kappa$. Moreover, S transfers $\dfrac{\pi_T \cdot \kappa}{\eta_H}$ as long as $\eta_H > 1$. Otherwise, S transfers 0. □

[24] S's optimization problem has a corner solution.