NBER WORKING PAPER SERIES

SOCIAL STRUCTURE AND INSTITUTIONAL DESIGN: EVIDENCE FROM A LAB EXPERIMENT IN THE FIELD

Emily Breza
Arun G. Chandrasekhar
Horacio Larreguy

Working Paper 20309
http://www.nber.org/papers/w20309

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2014

We thank Daron Acemoglu, Abhijit Banerjee, Sam Bowles, Esther Duflo, Matt Elliott, Itay Fainmesser, Garance Genicot, Ben Golub, Matthew O. Jackson, Jacob Leshno, Markus Mobius, Rajiv Sethi, and Adam Szeidl for detailed discussion. Important feedback was provided by seminar participants at Berkeley (SEEDEC), Brown (Workshop on Networks and Applied Micro), Columbia Business School, CIRANO Workshop on Trade and Finance, and the Santa Fe Institute. Manaswini Rao, Gowri Nagraj 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 views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Social Structure and Institutional Design: Evidence from a Lab Experiment in the Field
Emily Breza, Arun G. Chandrasekhar, and Horacio Larreguy
NBER Working Paper No. 20309
July 2014
JEL No. D02,L14,O17,Z13

ABSTRACT

In settings with poor formal contract enforcement, profitable investments are likely unrealized. While social closeness can mitigate contractual incompleteness, we examine how to improve the preponderance of cases where contracting parties cannot rely upon social ties. We ask if a community can enlist members to monitor transactions or punish offending parties.

We conduct a laboratory experiment in 40 Indian villages, with 960 non-anonymized subjects, where we have social network data. Participants play modified sender-receiver investment games, with and without third-party monitors and punishers. We examine whether network centrality of the third party increases efficiency of interaction. Furthermore, we decompose the efficiency increase into a monitoring channel (central third parties are valuable since they may influence reputations) and an enforcement channel (central third parties may be more able to punish without fear of retaliation).

Assigning a third party at the 75th percentile of the centrality distribution (as compared to the 25th) increases efficiency by 21% relative to the mean: we attribute 2/5 of the effect to monitoring and 3/5 to enforcement. The largest efficiency increase occurs when senders and receivers are socially distant, unable to maintain efficient levels autonomously. Results cannot be explained by demographics such as elite status, caste, wealth or gender.

Our findings show not every member is equally well-equipped to be part of a local institution. Knowing that a central third party observes their interaction increases sender-receiver efficiency. More importantly, to be able to punish someone, the third party must be important in the community.

Emily Breza
Columbia Business School
Division of Finance and Economics
3022 Broadway, Uris Hall 821
New York, NY 10027
ebreza@columbia.edu

Arun G. Chandrasekhar
Department of Economics
Stanford University
579 Serra Mall
Stanford, CA 94305
and NBER
arungc@stanford.edu

Horacio Larreguy
Harvard University
Department of Government
1737 Cambridge Street
CGIS Knafel Building 408
Cambridge, MA 02138
hlarreguy@fas.harvard.edu

1. Introduction

Developing countries are often plagued by weak legal enforcement of contracts. For example, in India it takes 1,420 days on average to resolve a commercial dispute between two businesses.[1] Strategies for mitigating these weak institutional environments may include entering into relational contracts or turning to individuals with extra-legal authority to help resolve disputes (Dixit 2011; Greif 1993; Greif et al. 1994). This paper focuses on the role that third-parties can play in helping to improve the scope of trade. Namely, we ask which members of a local community are best suited to serve in the role of arbiter and which enforcement tools these individuals should be given.

A network represents a pattern of economic and social relationships between agents. Each agent tends to only interact with a subset of others, and the pattern of local interactions woven together paints a picture of a global social structure. While social proximity – that is, being close in a social network – can help to sustain interactions, when parties are socially distant, mutually beneficial transactions may remain infeasible (Bowles and Gintis 2004; Chandrasekhar et al. 2013; Goeree et al. 2010; Glaeser et al. 2000; Grimm and Mengel 2009; Leider et al. 2009; and Ligon and Schechter 2012). Relying only on social closeness to substitute for formal enforcement has its limitations: it precludes any number of useful transactions involving the vast set of people who aren't closely linked.[2] In our setting, only 16% of possible contracting partners are of social distance two or closer. Thus, there is a natural role for institutions to contribute to contract enforcement.

In this paper, we address whether, an authority figure from the community can help to sustain cooperative behavior. In particular, we focus on whether the network importance (centrality[3]) of a local third-party monitor or enforcer influences the extent to which the individual can serve as an effective outside institution.[4] Network centrality captures the ease with which an individual passes and acquires information

[1]In contrast, it takes 370 days in the United States, where only half the cost is incurred. Source: doingbusiness.org

[2]Examples include co-investment opportunities that arise with fellow community members who are not one's friend, job referrals where the referrer must vouch for the referee, contributions to a public good involving more than one's local set of social links.

[3]We will use eigenvector centrality due to its natural interpretation in a simple information passing or stochastic meeting story, which we discuss below. We do not mean to suggest that this is the only (or exact) notion of network centrality relevant in the interactions we study.

[4]To our knowledge, this issue has not been studied in the literature. A related, though distinct analysis comes from the theoretical literature; Fainmesser (2014) shows that in environments where bilateral traders cannot maintain efficient exchange, central enough intermediaries can exploit their position to facilitate trade.

to and from others (see e.g., Jackson 2008 and Golub and Lever 2010). Central individuals may make effective arbiters because of their ability to spread information about the transactions to the rest of the community (reputation). It may also be the case that they are better able to wield tools for punishing violations (enforcement): less encumbered by the fear of retaliation by others, they may be able to threaten effectively with a punishment. We conduct a high-stakes laboratory experiment in the field designed to separate between these two channels. We also compare network-based authority to other naturally-arising hierarchy measures in the community such as caste, elite status, and gender.

We play two- and three-party investment games (modeled after Berg et al. 1995; Charness et al. 2008; and Fehr and Fischbacher 2004) in 40 rural Indian villages. Our setting is unique because of the availability of detailed demographic and network information available for most community members collected in previous work by Banerjee et al. (2013). We are precisely interested in understanding how positions of individuals in the actual community are brought into the games, and therefore, all experimental sessions are non-anonymized. This way, we bring individuals' deep and persistent relationships into the lab, which cannot be artificially induced in an anonymous lab experiment.

The two-party game is a natural benchmark for measuring the extent to which social proximity can alone improve outcomes. In this game, a sender (S) chooses how much of her endowment to send to the receiver (R), and any transfers triple in value. The receiver then determines how much to return to the sender. The initial amount sent by the receiver measures the efficiency of the transaction.

Our goal is to determine if outcomes can improve under the presence of specific types of third-party intervention and whether these improvements are due to reputational or enforcement channels. For this, we add two additional treatments. To investigate reputation, we assign a third party (T) to simply observe the sender and receiver transfers. The logic is that a monitor can observe misbehavior and she can either pass this information to others or, more subtly, may interact in the future with the sender or the receiver in other walks of life in the future. More central third parties are better positioned to pass information to a greater share of the population and are also more likely to interact with the sender or receiver in the future. To capture enforcement, we further give the third party the ability to punish the receiver. In this treatment, the third party both has the monitoring/reputation channel but additionally can levy a fine – an observable punishment – on the receiver in the game itself. In both cases, we are particularly interested in how the effectiveness of the third party may increase with her centrality in the network. By differencing these treatments, we can identify enforcement channels (as a whole) separated from the monitoring effect alone.

Our core results are as follows. First, third parties who are more central increase the level of efficiency relative to their less central counterparts. Specifically, sender transfers increase substantially – by 21% of the mean when the third party is at the 75th percentile as compared to the 25th percentile of the centrality distribution – when the third party is given the ability to punish. Decomposing this into enforcement and monitoring channels, we find that about 3/5 of the effect comes from the enforcement channel (an effect of about 12% of the mean) whereas 2/5 of the effect comes from the monitoring channel (an effect of about 9% of the mean).

Second, third parties are precisely more effective when the sender and receiver are socially distant. Thus, in exactly the environment identified by the preceding literature where lack of commitment contracts severely hurts efficiency, we see that introducing a central third party who can levy a fine generates great benefits to efficiency.

Third, we also show that other natural measures of social hierarchy cannot match the effectiveness of a network-central punisher. Third parties (monitors and punishers) of elite status, of high caste, or male do not improve outcomes. Additionally, note that all of our our networks results condition on these demographic variables[5] interacted with treatment, meaning that the network effects are not simply proxying for classical notions of social hierarchies.[6]

Finally, adding up over all triples of sender, receiver and third party, we find that on average adding a third party to a two-party investment game neither increases nor decreases sender transfers. But, as mentioned above, this masks a striking and predictable heterogeneity in line with the literature. Socially close senders and receivers achieve better outcomes on their own, and consistent with previous work by Fehr and Gächter (2002), adding a third party generates some crowdout. However, for otherwise inefficient pairs of socially far senders and receivers, the third-party institution may improve outcomes. We find that high-centrality punishers are able to substitute for the lack of social proximity between the sender-receiver pairs; far sender-receiver pairs under a high centrality third-party punisher behave in a manner similar to socially-close sender-receiver pairs. However, relative to the two-party game, adding a punisher who is peripheral in the social network is detrimental to efficiency and crowds out transfers.

[5]As well as wealth, education and age.

[6]Our enforcement channel is particularly robust to the inclusion or omission of any control, though our monitoring channel effects depend on the controls and become very noisy without them.

Without formal contracts, the fact that individuals have repeated, systematic interactions with a local set of (socially proximate) individuals is more than enough to sustain cooperative behavior. But broader interactions – involving members from the community at a distance– require mediation. Not all individuals in a community are equally well-equipped to handle this role. Specifically, those who are central in the social network are the best arbiters. While they do exhibit a monitoring effect, their effectiveness appears to come largely from the fact that they are better equipped to utilize the punishment technology given to them. Further, the centrality effect is conditional on demographic-by-treatment controls: thus it is not proxying for leadership status, caste status or gender. This tells us that network position is crucial in institutional design. Further, from a policy perspective, our results suggest that selecting informal arbiters on the basis of caste or occupational status alone may not be so effective. Network importance instead may be able to substitute for a formal contracting technology.

Structure of the Paper. The remainder of the paper is organized as follows. In section 2, we describe the experimental design. Section 3 contains an overview of our subject pool as well as a summary of the available network and demographic data.. We introduce our empirical strategy in Section 4. In section 5 we present the results, and section 6 concludes.

2. Experiment

Our experiments were conducted in the summer and fall of 2010 in 40 villages in Karnataka, India which range from a 1.5 to 3 hour's drive from Bangalore. The villages are independent – the median distance is 46 kilometers between two villages. In each village, 24 individuals aged 18 to 50 were recruited to take part in the experiments. As an incentive to attend, participants were paid a show-up fee of INR 20 and were told they would have the opportunity to win additional money. We chose these individuals as we had access to village census demographics as well as unique social network data, previously collected in part by the authors. The data set is described more below in Section 3 and in detail in Banerjee et al. (2013). 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. (2013) and Chandrasekhar et al. (2013). 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.

The survey data also includes information about caste, elite status, wealth proxies, gender, age, and educational attainment. An individual is of low caste for our analysis if they belong to any of the historically disadvantaged scheduled castes or scheduled tribes (SC/ST). A 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.[7] To construct a wealth index, we use survey information on house size, electrification, building materials, and toilet amenities.

Our experimental design has three treatments (T1-3). As described below, for every treatment, we use individuals who are randomized into roles: sender, receiver and third party. Additionally, we took care to ensure that no individual participated in any game with any other individual more than once.

2.1. **Two-Party Experimental Benchmark.** We consider an environment where there are two parties who do not have access to formal contracting. We begin with an important benchmark, the two-party investment game (T1). Later, we will add to this a third-party institution.

In this benchmark game, two participants are selected at random and are assigned the roles of S and R. The players are given endowments of Rs. 60 each. S moves first and must decide how much of her endowment to transfer to R ($\tau_S \in [0, 60]$). The receiver then receives $3 \times \tau_S$ (i.e., the size of the transfer triples).

In the second stage of the game, R decides how much of his or her wealth from the game to return to S ($\tau_R \in [0, 60 + 3\tau_S]$). Here, the final payoff for S is $60 - \tau_S + \tau_R$, and the final payoff for R is $60 + 3\tau_S - \tau_R$. Note that in this two-period game, R will return $\tau_R = 0$ in any Nash Equilibrium, and by backward induction, S never makes a strictly positive transfer τ_S, However, any efficient outcome is one where $\tau_S = 60$. This setting mimics a situation where there are efficiency gains from cooperating or co-investing, but there are no formal contracting tools to enforce positive transfers ex post. This game allows us to document the degree of efficiency a pair of individuals alone can reach.

Previous research in both the lab and the field suggests that social proximity may substitute for formal contracting institutions.[8] Much of the focus of previous work

[7] *Gram panchayats* are local government institutions at the village or small town level in India. *Anganwadi* centers are local educational and health centers in India.

[8] See Chandrasekhar et al. (2013), Bowles and Gintis (2004), Glaeser et al. (2000), Goeree et al. (2010), Leider et al. (2009), and Ligon and Schecter (2009) for lab evidence and McMillan and Woodruff (1999), among others. for field evidence.

has been on whether social proximity generates good behavior because of altruism or reciprocity, though it is not our aim here to separate such things. Let $d(i,j)$ denote the minimum path length between individuals i and j, which is the social distance between the two parties. A unit increase in social distance reduces social proximity by 1. We estimate the value of social distance $d(S,R)$ between S and R in our specific setting and also measure how far contracting outcomes between socially distant (S,R) pairs are from the efficient level.

2.2. **Third-Party Intervention.** Our goal is to understand if contracting outcomes can be improved by the involvement of a third party from the community, especially if the contracting parties are socially distant. There are two natural roles that a third party can play: one of monitoring or reputation and one of punishment or enforcement. A third party can observe (mis)-behavior on the part of one of the two vested parties and can therefore pass information to others in the community.[9] These others, or even the third party herself, may in the future interact with R. In this way R's behavior in the S-R interaction might shape the outcome of subsequent interactions with others through a reputational channel. Note that this sort of reputational punishment of R by the third party is usually not verifiable outside of the game as the third party cannot commit to passing specific messages to others. Ultimately, we are precisely interested in this monitoring channel as it mimics the sort of interaction that happens in these communities day-to-day.

Beyond a monitoring effect, there might also be an enforcement effect. It some cases, the third party may be endowed with a direct and verifiable punitive mechanism (e.g., levying a fine) that can be exacted on a wrong-doing party. Moreover, we note that it may be that not everyone is equally well-equipped to wield the punishment. For example, some individuals may be better suited to withstand retaliation by the punished party. Similarly, some may be viewed as being more fair or possessing the authority to punish others. We should note that employing a third-party punisher induces both monitoring and enforcement of the two-party exchange.

Our experimental design is structured to disentangle mechanisms of reputation from mechanisms of enforcement. We ask both if reputation and enforcement can (separately) improve outcomes, and if it matters who is given the role of monitor or punisher. Participants play two different games with third-party involvement. Again, for these games, the 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 the reputation variant (T2), we assign an individual T to watch the play of S and R but do not

[9]In our setting, such information transmission would occur outside of the experimental sessions.

allow her to take any actions within the game. In the enforcement variant (T3), we assign and individual T to watch the play of S and R and also allow her to take a costly punitive action in a third stage of the game. Specifically, T has the option to spend his or her own resources to levy a monetary punishment on R. For every Rs. 1 spent by T, R loses Rs. 4.[10] In both T2 and T3, the third party is given an endowment of Rs. 100.[11]

Comparing T2 to T1 allows us to capture the monitoring or reputational effect, while comparing T3 to T2 isolates the enforcement effect. We can also examine heterogeneous effects of adding monitors or punishers of different types to the baseline game.

2.3. **Social Structures.**

Network Centrality. A village social network is a description of interactions that tend to occur between households. These interactions can arise for myriad purposes – leisure, financial transactions, informational exchange, etc. – and can depend on invariant features such as caste and geography. Importantly, the patterns of interaction are not transient and represent a deep and persistent structure of exchange. In the experiment that follows, we want to know whether individuals who are important in a network sense –who tend to be highly interactive or from whom information tends to flow – will serve as efficiency-enhancing third parties.

We focus on eigenvector centrality, which measures how well an individual is able to spread and collect information through the social network.[12] Eigenvector centrality is a recursive measure of importance wherein an individual's centrality is proportional to the sum of each of her neighbor's centralities.[13]

Eigenvector centrality is relevant to our setting because it tells us which individuals are better suited to pass information about what happened in the game and which individuals are more likely to be relevant in future encounters. It also captures an individual's ability to exert social influence both within and outside of the game. A very simple framework from Banerjee et al. (2013) explains the relevance of this measure in our setting. Nonetheless, we do not mean to suggest that other notions of centrality are irrelevant.

[10] We did not vary experimentally the Rs. 4 punishment cost, and we do not claim that this is the efficiency-maximizing punishment technology. We leave the determination of the optimal punishment function for future work.

[11] 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.

[12] See Appendix B for more discussion about this and related measures.

[13] Empirical network papers employing eigenvector centrality include Hochberg et al. (2007), Banerjee et al. (2013), and Schechter et al. (2011).

Consider a simple process where every individual interacts with her neighbor in a network (given by adjacency matrix $\mathbf{A}$, where $A_{ij} \in \{0,1\}$ indicates whether i and j are linked or not) with probability p. This means that a piece of information from i to neighbor j travels with probability p, or that i may go to j's house with probability p. This process repeats itself for T periods. Then the expected number of times that information goes from i to all other nodes (or similarly the expected number of times i visits all other nodes) is given by

$$DC\left(p,T\right) := \sum_{t=1}^{T} \left(p\mathbf{A}\right)^{t} \cdot \mathbf{1}.$$

This measure is called diffusion centrality, and an entry DC_i represents the expected number of times that i has informed (or visited) all others. Banerjee et al. (2014) show that as $T \to \infty$ if $p \geq \frac{1}{\lambda_1(\mathbf{A})}$, diffusion centrality converges to eigenvector centrality.

Because a central monitor is well-suited to pass information, any monitoring effect should be more powerful when the third party is central, rather than peripheral. This may provide incentives for the receivers to return more and for the senders, in turn, to send more. Similarly, the likelihood of any sender or receiver interacting with the third party in the future is higher if the third party is central, potentially providing incentives for building reputation with the third party. Both of these mechanisms are present in our experimental treatment with a third-party monitor.

In the cases with third-party punishment, centrality may permit for greater independence for the punisher. For example, the costs outside the game of punishing the receiver may be especially high for punishers of low centrality; receivers could more easily retaliate against these types of punishers. Similarly, members of society may only grant the authority to punish to important enough members of the community.

Our design allows us to ask if randomly giving a punishment tool to a third-party monitor improves outcomes, and if this improvement is differential across central and peripheral third parties. If the monitor channel improves transfers, we predict that a more central monitor will generate more efficiency, and similarly, if the enforcement channel improves transfers, we predict that a more central punisher will be associated with more efficiency.

Other Types of Hierarchical Authority. We also consider the interactions of alternative types of social importance with the play of the game. In our data, we observe measures of caste status, elite status, and gender, each of which has been examined previously in the literature in other contexts. These characteristics might also correlate with third-party efficacy.[14]

Low caste status indicates that an individual belongs to a scheduled caste or tribe. Historically, these groups were disadvantaged in their access to education and employment opportunities (Munshi and Rosenzweig, 2006) and were limited in their economic and social mobility. Accordingly, one might expect high caste individuals to be viewed as more important decision-makers in the village, and they may be more immune to retaliation than members of the lower castes. Some laboratory evidence also suggests that low caste compared to high caste individuals punish norm violations less often and less severely (Hoff et al. 2008, 2011).

Those we term as elites in our setting include *gram panchayat* members, *anganwadi* teachers, etc. *Gram panchayat* members are individuals with great social and political power. While ex ante powerful individuals may be more likely to reach these positions, the roles themselves may also generate influence. School teachers and headmasters are also widely known and respected for the role they play in the education of the village and are often informed of new government programs for children and families before any other individuals. Local elites, many of whom are already endowed with formal authority might be well-suited to serve as arbiters. However, leaders may also be prone to elite capture and may be more sensitive to retaliation by their constituents (Ball et al. 2001; Abrams et al. 2012; von Essen and Ranehillii 2012).

Finally, we can also explore if the gender of the third party matters for punisher efficacy. India is one of the countries with the lowest sex ratios in the world (Sen 1992) and has recently implemented policies of reservation for women to correct for historical discrimination (Beaman et al. 2009). As such, we may expect that women might be less respected as third parties, especially when they have the ability to punish. We explore whether all these alternative types of social importance also make effectual monitors and punishers.

3. Data

3.1. **Network Data.** We chose our sample frame from villages where we have access to village census demographics as well as unique social network data, previously collected in part by the authors. In our data, we have a census of every individual in

[14]It is also possible that each of these characteristics is correlated with network position. We explore these correlations in Section 3.4.

every village and we have network data collected across the following 12 dimensions: "(1) those who visit the respondents' home, (2) those whose homes the respondent visits, (3) kin in the village, (4) non-relatives with whom the respondent socializes, (5) those from whom the respondent receives medical advice, (6) those from whom the respondent would borrow money, (7) those to whom the respondent would lend money, (8) those from whom the respondent would borrow material goods (kerosene, rice, etc.), (9) those to whom the respondent would lend material goods, (10) those from whom the respondent gets advice, (11) those to whom the respondent gives advice, and (12) those whom the respondent goes to pray with (at a temple, church, or mosque)" (Banerjee et al., 2013). 46% of households were administered the network survey module, and households could name as many responses per question as they wanted and name as many individuals from their village in response to the questions. 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, across many networks. This level of detail across so many independent networks is atypical.

To construct our network $\mathbf{A}$, we build an undirected, unweighted graph taking the union over the twelve dimensions, and we construct the network at the household level. This is consistent with the Banerjee et al. (2013) and Chandrasekhar et al. (2013) treatment of the data, and that work has a lengthier discussion on this decision. That is, any two households are linked if any member has any relationship with anyone else. This is reasonable as, in our villages, the multiple dimensions are highly correlated so the union network ensures that we take into account any possible meaningful relationship, without constructing an ad hoc weighting procedure. Going forward, we use this $\mathbf{A}$ to represent the network of the village and we construct distances between nodes and centralities of nodes accordingly.

3.2. **Recruitment and Implementation.** To recruit participants in each village, we randomly chose a subset of approximately 15 households to invite to participate. This randomization ensured that the villagers perceived our recruitment to be fair. We visited all invited households two days before the experimental sessions and informed them about the opportunity to participate. We then told them the location and the starting time of the games and guaranteed that they would have priority to participate. When we arrived in the villages at the specified time, in most cases many individuals both invitees and non-invitees were already waiting for us. Before registering the participants, we also walked through the village to remind households to participate.

Once registered for participation, each participant played five to six total rounds of three experimental treatments. During each round, players were randomly assigned to one of the three roles: 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 participants at a time. Surveyors were randomly assigned to different groups and thus their presence should be orthogonal to the treatments and the characteristics of the participants. Additionally, while there might be the concern that surveyors also play the role of monitors, they had no ties in the villages where we conducted our experiments and were prohibited from commenting on the behavior they observed during the experiments to other villagers. Further, they were present for all three treatments. Any cross-treatment or heterogeneous effects we measure are therefore net of any effects from the surveyor's presence.

Individuals played the different games in random order.[15] We also ensured that no two individuals played any game with one another more than once. This restriction limited the number of rounds we could play with the participants from each village.

Participant played two rounds of each of T2 and T3. Half of participants played one round of T1, while the other half played two rounds. After playing all of the rounds, participants were each given their ending wealth values for one randomly-chosen game plus a fixed participation fee of Rs. 20. The average payoff was approximately Rs. 110, or approximately three-fourths of a daily agricultural wage. We stress that these are extremely large stakes – nearly a day's wage for an experiment lasting no more than an hour.[16]

3.3. **Outcomes.** In our analysis, we focus on the initial transfer made by the sender to the receiver as our key outcome. This transfer level encodes efficiency because it is one-to-one with the entire size of the pie that the receiver then chooses how to allocate. Before analyzing the treatment effects and network effects, it is helpful to first observe the overall outcomes from the experimental sessions. The data include 1,888 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. 28.4. A zero transfer is only observed in 13 cases. The efficient transfer of Rs. 60 is observed 122 times (~6% of games).

[15]As a result, for all regression specifications, we can include both game and round fixed effects.
[16]For reference, assuming \$50,000 GDP per capita and an individual working 5 days a week for 52 weeks, this would scale to stakes of \$144 for participating in the experiment.
[17]Participants could make transfers in increments of Rs. 1.

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. 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.

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. 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.

3.4. **Sample Statistics.** Table 1 presents the descriptive statistics. Our analysis sample includes 930 participants from 1888 two- and three-party game observations. 59% of the participants are female, and the average education level is 8.15 years of schooling with a standard deviation of 4.32.[19] About 67% of the participants are general or "otherwise backwards" (OBC) caste.[20] Finally, 20% of households have a leader or local elite.

Turning to network characteristics, 96.5% of paris are reachable (there exists a path through the network connecting the two).[21] Between the reachable pairs, the maximum social distance is 8, while the average social distance is approximately 3.5.[22] The average eigenvector centrality is 0.02 with a standard deviation of 0.04, indicating that there is substantial heterogeneity in an individual's social importance.

Finally, we explore the relationship between the network characteristics and demographic covariates. To do this, in Table 2 we present a correlation matrix of several network and demographic covariates, as well as a principal component matrix. While the various functions of eigenvector centrality are highly correlated

[18]At least before the punishment decision is made.
[19]This means that on average, an individual had attended 8th standard.
[20]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. Our indicator for "high caste" groups the GM and OBC designations.
[21]We condition our sample on this set of reachable pairs for all of the analysis.
[22]Appendix B contains a glossary formally describing the network statistics used.

(mechanically, so),[23] and while there is non-zero correlation between network centrality and other physical covariates (e.g., wealth, elite status, caste), the correlation is not very high. This is interesting in that it provides some suggestive evidence as to why, even when conditioning on treatment-physical covariates interactions in regressions, our network-based results remain extremely robust. This bolsters the idea that the topological features of the network are instrumental in the results that we are describing.

Similarly, in Table 3, we present the first three vectors of a principal component decomposition of the importance characteristics. The decomposition contains six different measures of importance: eigenvector centrality, elite status, high caste, wealth, educational attainment, and gender. The five variables separate along three distinct dimensions. Eigenvector centrality is the main constituent of the first principal component, caste, wealth and education are all key contributors to the second principal component, and elite status and gender appear in the third principal component. This suggests that network centrality does have content distinct from the other demographic characteristics.

4. Empirical Strategy

We are interested in examining how the efficiency of a session, as measured by the sender's transfer τ_S, responds to whether there is a third party, whether the third party has access to a punishment technology and how these answers depend on the centrality of the third party in the network,

Specifically, our analysis uses regressions of the following form:

$$\begin{aligned} \tau_{S,rgjv} \quad = \quad & \alpha + \beta_{T2} \cdot \mathbf{1}_{\{g=T2\}} + \gamma_{T2} \cdot \mathbf{1}_{\{g=T2\}} \cdot e_{T,jgv} \\ & + \beta_{T3} \cdot \mathbf{1}_{\{g=T3\}} + \gamma_{T3} \cdot \mathbf{1}_{\{g=T3\}} \cdot e_{T,jgv} \\ & + \delta'_{T2} W_{jgv} \cdot \mathbf{1}_{\{g=T2\}} + \delta'_{T3} W_{jgv} \cdot \mathbf{1}_{\{g=T3\}} \\ & + \eta' X_{jgv} + \mu_r + \mu_{vg} + \epsilon_{rgv}. \end{aligned} \tag{4.1}$$

Here r is round, j is the triple of players (SRT), g is game, and v is village. $e_{T,jgv}$ is the eigenvector centrality of the third party T_j, and W_j is a triple of leadership status, caste, and gender of T_j. Finally, X_j is a vector of other demographics and network controls for all parties (e.g., centrality, leadership status, caste and gender of S and R, social closeness between all pairs [see Appendix B], wealth of all three

[23]Note that the table displays the raw correlations across individuals and villages. The quantile rankings of wealth and centrality are constructed using within-village rankings.

parties, and education of all three parties, as well as interactions of each of the aforementioned variables with game dummies), μ_r is a round fixed effect, and μ_{vg} is a village-experimental session fixed effect.

We are particularly interested in γ_{T2} and γ_{T3}. Observe that γ_{T2} measures the effect of an increase in centrality on efficiency through the monitoring channel. $\gamma_{T3} - \gamma_{T2}$ measures the relative return to centrality of T through the enforcement channel (net of the monitoring channel). Central questions are whether $\gamma_{T2} \geq 0$ and $\gamma_{T3} > \gamma_{T2}$.

Similarly, turning to the question of whether leaders, high caste member, or females make relatively better third-party institutions, we are interested in parameter vectors δ_{T2} and δ_{T3}. These capture whether leaders, high caste members or females generate efficiency as third-party institutions and whether this effect comes from a monitoring channel, an enforcement channel or both.

Observe that we are chiefly interested in heterogeneous treatment effects based on network position and in cross-treatment heterogeneous network effects. While these parameters are identified given our randomized experimental design, we do acknowledge that the networks themselves are not randomly assigned. People who are central might differ from people who are peripheral on numerous dimensions. In the analysis, we are both able to ask if other demographic measures of social importance can replicate the centrality results and to control for all available demographic characteristics in our main regression specification. If we find (which we do) that among all observable characteristics only network-central punishers improve efficiency, then we can be confident that for institutional design purposes, we should target high centrality individuals when building informal enforcement mechanisms.

5. Results

5.1. **Network centrality.** We now address the central theme of our paper: how does introducing a third-party institution influence the efficiency of outcomes? Do more central third parties generate more efficiency and, if so, how much can be attributed to an enforcement channel as compared to a monitoring channel?

Table 4 presents our main results – the specification described in (4.1). Columns 1-4 use two different versions of our centrality measure. Columns 1-2 present the results using centrality percentile, while columns 3-4 use an indicator for whether the third party is above the 50th percentile of centrality in the sample distribution. Finally, in columns 1 and 3 we do not use demographic-by-treatment controls, which we introduce in columns 2 and 4.

First, in column 2 we see that being randomly assigned a third party in the monitoring treatment who is at the 75th percentile of the centrality distribution as compared to the 25th percentile corresponds to a Rs. 2.58 increase (9% relative to the mean). This means that there is a modest efficiency gain simply from the monitoring channel.

Second, we turn to the enforcement channel. Being randomly assigned a third party at the 75th percentile instead of the 25th percentile of the centrality distribution corresponds to an *additional* Rs. 3.4 (12% relative to the mean) increase in transfers (column 4). In sum, having a third party who can punish who is at the top of the inter-quartile range of the centrality distribution increases transfers by about 21% relative to the mean, and nearly 3/5 of the effect is coming from the enforcement channel.

It is important to note that these are effects that are conditional on demographic and other social network controls interacted with treatment. This is our preferred specification since it directly controls for confounds such as leadership, caste, gender, wealth, age, education of all parties, sender and receiver centrality, and social closeness between all parties with treatment-varying effects.[24] Columns 1 and 3 show the same regression results without demographic controls as well. We find that the results for the enforcement channel are remarkably robust. However, removing the controls generates considerable noise in the estimates of the monitoring channel in column 1, and we are unable to reject zero effect in that specification. In each of the four specifications, the effect of being paired with a highly central third party is greater in magnitude in the enforcement + monitoring treatment (T3) than in the game with monitoring alone (T2). The difference in these effects is statistically significant at standard levels in three of the four specifications.

Taken together we find extremely stable and robust evidence that there is a large enforcement channel through which network centrality is associated with efficient outcomes. While we do find a modest effect size for the monitoring channel as well, this result is more sensitive to the inclusion of controls. Ultimately, by randomly grouping individuals and randomly varying the institutional setup (designed to parse the monitoring channel from the enforcement channel) while controlling for demographic-by-treatment covariates, we are able to take a reasonable measurement of how network centrality of a third party influences efficiency and to decompose it along monitoring and enforcement channels.

5.2. **Leadership, caste and gender.** 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

[24]The entire table with numerous coefficients is available upon request.

with network position. Additionally, understanding whether elite status, caste or gender status influences the effectiveness of a third party is interesting in its own right. Thus, our main specification includes demographic-by-treatment regressors, which allow us to control for these effects when looking at centrality effects but whose point estimates are themselves important.

As noted in Table 4, demographic variables tend to be much more correlated with other demographic variables and less with network variables. Similarly, when looking at the principal component analysis in Table 3, we observe that demographic and network characteristics pick up different dimensions of variance in a principal component analysis. This provides a priori evidence that it is likely that network variables are likely not going to merely proxy for other standard demographic features.

Turning to the effect of these demographic variables in their own right, we see that none of elite, caste nor gender replicates the patterns observed with the network characteristics. Specifically, there is no detectable effect of any of these demographic features either through the monitoring channel or through the enforcement channel, unlike the case of network centrality.

Finally, in Appendix Table A.1, we present results demographic variable by demographic variable (neither conditioning on each other nor on network centrality). We find that even in these cases, there is no association between elite, caste or gender status of the third party on the sender's transfer. This again confirms that none of the demographic importance measures is able to replicate the effects of being assigned to a network-central third party. This also implies that when trying to design informal enforcement institutions, it is not enough to try to select individuals to play the role of arbiter based on their demographic characteristics alone.

5.3. **Institutional design.** A prediction of our framework is that the third party should have the most scope to help precisely when senders and receivers cannot autonomously engage in efficient behavior. This is what we examine below. The following exercise serves two purposes. First, it functions as an over-identification test: it tells us where to look for another effect consistent with our story. Second, it allows us to address institutional design questions. Are third-party institutions only relevant and/or influential when contracting parties are socially distant and when the third party is more central? The latter point is particularly important as we show in Appendix Table A.2, that on average introducing a third-party monitor or punisher has no net effects. Despite having no mean effects, the third party may be particularly beneficial, as suspected, exactly when S and R are unable to maintain efficient levels on their own without commitment contracts.

Figure 3 illustrates this graphically, plotting sender transfers[25] 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 punisher is of high centrality (middle bar in each grouping) versus low centrality (right-most bar in each grouping).[26] The bar charts illustrate many of our key networks results but also allow for comparisons between the three- and two-party 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. Another striking pattern is that 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 considerably lower (11.9% relative to the mean) than the two-party outcome. However, when the punisher is central in the network, transfers are 9.1% higher (relative to the mean) than the two-party outcome. Moreover, this level of transfer is comparable to (and not statistically distinguishable from) the two-party outcome when the sender and receiver are socially close. We observe a similar, yet weaker pattern when S and R are far, and when there is a central vs. peripheral third-party monitor. Finally, because socially close S and R experience larger transfers in the absence of third-party enforcement, there is both less scope for a third party to improve but also more scope for it to crowd out efficiency. Consistent with Fehr and Gächter (2002), we observe that the extrinsic incentives introduced only by the third-party punisher but not the monitor crowd out transfers when S and R are socially close.

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

[26]We say that S and R are close if their distance is at most two, and we say T is of high centrality if she has above-median eigenvector centrality among the individuals playing from her own village.

We now do this exercise formally, conditioning on various fixed effects and demographic-by-treatment controls as in (4.1). We use regressions of the form:

$$
\begin{aligned}
\tau_{S,rgjv} \quad = \quad & \alpha + \beta_{T2} \cdot \mathbf{1}_{\{g=T2\}} + \phi_{T2} \cdot \mathbf{1}_{\{g=T2\}} \cdot e_{T,jgv} \\
& + \theta_{T2} \cdot \mathbf{1}_{\{g=T2\}} \cdot e_{T,jgv} \cdot Closeness(S,R)_{jgv} \\
& + \beta_{T3} \cdot \mathbf{1}_{\{g=T3\}} + \phi_{T3} \cdot \mathbf{1}_{\{g=T3\}} \cdot e_{T,jgv} \\
& + \theta_{T3} \cdot \mathbf{1}_{\{g=T3\}} \cdot e_{T,jgv} \cdot Closeness(S,R)_{jgv} \\
& + \delta'_{T2} W_{jgv} \cdot \mathbf{1}_{\{g=T2\}} + \delta'_{T3} W_{jgv} \cdot \mathbf{1}_{\{g=T3\}} \\
& + \eta' X_{jgv} + \mu_r + \mu_{vg} + \epsilon_{rgv}.
\end{aligned} \tag{5.1}
$$

$Closeness(S,R)_{jgv}$ is a number describing how socially proximate S and R are. Specifically, it is defined as $\max_{i',j'} d(i',j') - d(S,R)$. Thus, a value of zero indicates minimal closeness (maximal distance) and maximal closeness is simply $\max_{i',j'} d(i',j') - 1$. This is useful because ϕ_{T2} and ϕ_{T3} therefore encode how the centrality of the third party differentially influences efficiency when S and R are furthest away. X_j is as in (4.1), so recall that it includes interactions of $Closeness(S,R)$ with game dummies (and in fact all pairs of social closeness interacted with game dummies).

Table 5 presents the results. In Columns 1-2 we show results for centrality quantile, and columns 3-4 display a dummy for whether the third party is above the median centrality quantile. Columns 1 and 3 do not condition on demographic-by-treatment controls whereas columns 2 and 4 do.

We find that for the furthest individuals, going from the 25th to the 75th percentile in centrality of the third party with punishment corresponded to an increase of sender transfers of Rs. 17.6 (a 62% increase relative to the mean). However, S and R being closer by two steps (e.g., distance 2 instead of distance 4) corresponds to the third-party punisher's centrality being less valuable: the inter-quartile effect is now only an increase of 12.33 (a 43.4% increase relative to the mean). It is worth noting that we have very little power in this specification since we are looking at a triple interaction, and so decomposing the effects into enforcement versus monitoring channels becomes difficult. While we are unable to reject that there is no monitoring effect due to considerable noise, we can only separate the effect of enforcement from the (noisily estimated) effect of monitoring in the specifications with the full demographic controls. (columns 2 and 4)

The exercise here provides evidence supporting the idea 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.

6. Conclusion

A large literature highlights the fact that agents can maintain efficient transactions with those with whom they are socially close despite a lack of formal institutions. Thus, a crucial question is how to develop informal institutions that can improve efficiency even for distant contracting pairs. Namely, who in the social network can intervene to help facilitate trade, and which tools should that individual be given?

Using a lab-in-the-field design across 40 villages where we have detailed network data, we identify that the centrality of the third party greatly affects the efficiency of a sender-receiver interaction in a simple investment game. Additionally, by comparing a treatment where the third party simply observes the interaction with one where the third party can additionally punish the receiver, we distinguish between a pure monitoring channel and an enforcement channel.

We find that while both mechanisms are at play, the enforcement channel is robustly larger than the monitoring channel. This suggests that not all individuals are equally well-equipped to take on the role of a punisher. For instance, a less central individual may face retaliation outside of the dispute arbitration in a way that a more central individual need not worry about.

Further, our results are consistent with the motivating idea that external institutions are only necessary when repeated game dynamics or social preferences aren't enough to generate cooperation. The introduction of a central third party does little relative to the two-party game outcome when the sender and receiver are socially close. However, when the pair is distant, they have considerably less scope for cooperation and it is precisely in this situation where having a central third party truly pays off.

Our results present a lower bound on the potential for a carefully chosen individual to help, as we did not optimize the choice of the cost per unit punishment and did not vary experimentally the type of enforcement technology given to the third party. While on average introducing a third-party monitor or punisher has no net effects on efficiency, as Appendix Table A.2 indicates, this could simply be an artifact of the lack of optimality of our punishment technology. Clearly, one could think about constructing an optimal punishment technology that would yield weakly larger transfers than the ones presented here. However, this is beyond the scope of this paper and is left to future research.

Finally, for policy what is important is identifying the "right" individual in the community. We have shown that observable characteristics such as elite status, caste, and gender are not particularly good predictors of who makes for an effective monitor or judge. It is only the network centrality that is a good predictor of monitor or punisher success. This means that simply relying on demographic and occupational traits to build an effective informal institution may fall short.

References

Abrams, D., M. Bertrand, and S. Mullainathan (2012): "Do Judges Vary in their Treatment of Race?" *Journal of Legal Studies*, 41, 347–383.

Ball, S., C. Eckel, P. Grossman, and W. Zame (2001): "Status In Markets," *The Quarterly Journal of Economics*, 116, 161–188.

Banerjee, A., A. Chandrasekhar, E. Duflo, and M. Jackson (2013): "Diffusion of Microfinance," *Science*, 341, DOI: 10.1126/science.1236498, July 26 2013.

Banerjee, A. V., A. Chandrasekhar, E. Duflo, and M. O. Jackson (2014): "Gossip: Identifying Central Individuals in a Social Network," *Available at SSRN 2425379*.

Beaman, L., R. Chattopadhyay, E. Duflo, R. Pande, and P. Topalova (2009): "Powerful women: does exposure reduce bias?" *The Quarterly Journal of Economics*, 124, 1497–1540.

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

Bowles, S. and H. Gintis (2004): "Persistent parochialism: trust and exclusion in ethnic networks," *Journal of Economic Behavior & Organization*, 55, 1–23.

Chandrasekhar, A. G., C. Kinnan, and H. Larreguy (2013): "Can Networks Substitute for Contracts? Evidence From a Lab Experiment in the Field," *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.

Dixit, A. K. (2011): *Lawlessness and Economics: Alternative Modes of Governance*, Princeton: Princeton University Press.

Fainmesser, I. (2014): "Exclusive Intermediation," *Working Paper*.

Fehr, E. and U. Fischbacher (2004): "Third-party punishment and social norms," *Evolution and human behavior*, 25, 63–87.

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

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.

Golub, B. and C. Lever (2010): "The leverage of weak ties: how linking groups affects inequality," *Working Paper*.

Greif, A. (1993): "Contract Enforceability and Economic Institutions in Early Trade: the Maghribi Traders' Coalition," *American Economic Review*, 83(3), 525–48.

Greif, A., P. Milgrom, and B. R. Weingast (1994): "Coordination, Commitment, and Enforcement: The Case of the Merchant Guild," *Journal of Political Economy*, 102(4), 745–776.

Grimm, V. and F. Mengel (2009): "Cooperation in viscous populations - experimental evidence," *Games and Economic Behavior*, 66, 202–220.

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 (2008): "Spite and Development," *American Economic Review*, 98, 494–499.

——— (2011): "Caste and Punishment: the Legacy of Caste Culture in Norm Enforcement," *The Economic Journal*, 121, 449–475.

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

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

Ligon, E. and L. Schechter (2012): "Motives for sharing in social networks," *Journal of Development Economics*, 99, 13–26.

Ligon, E. and L. Schecter (2009): "The Value of Social Networks in Rural Paraguay," .

McMillan, J. and C. Woodruff (1999): "Interfirm relationships and informal credit in Vietnam," *The Quarterly Journal of Economics*, 114, 1285–1320.

Munshi, K. and M. Rosenzweig (2006): "Traditional institutions meet the modern world: Caste, gender, and schooling choice in a globalizing economy," *The American Economic Review*, 1225–1252.

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

Sen, A. (1992): "Missing women." *BMJ: British Medical Journal*, 304, 587.

von Essen, E. and E. Ranehillii (2012): "Punishment and Status," .

Figures

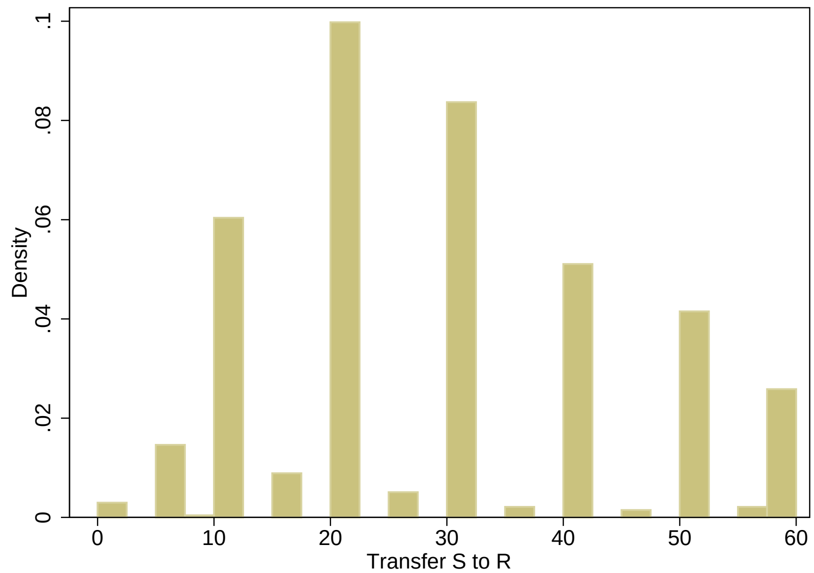

Figure 1. Distribution of transfers from sender to receiver

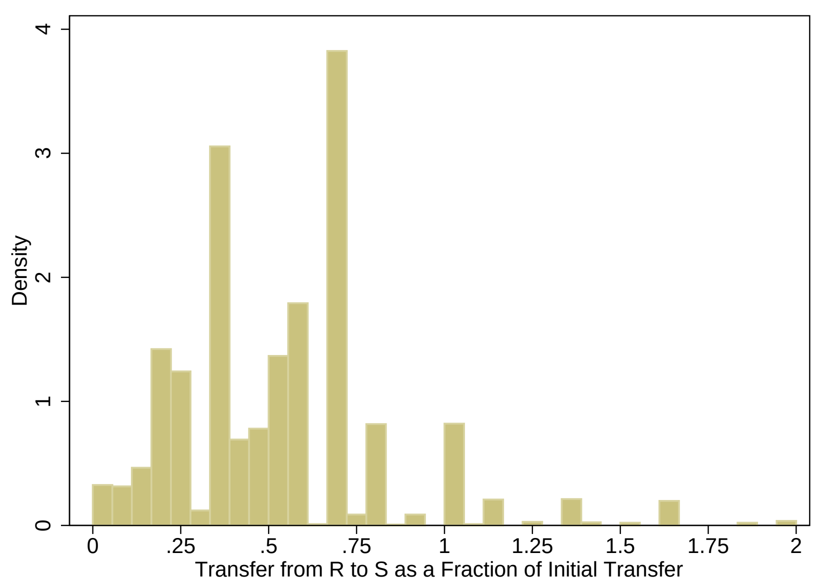

Figure 2. Distribution of transfers from receiver to sender as a fraction of the initial transfer

Panel A: Third-Party Monitors

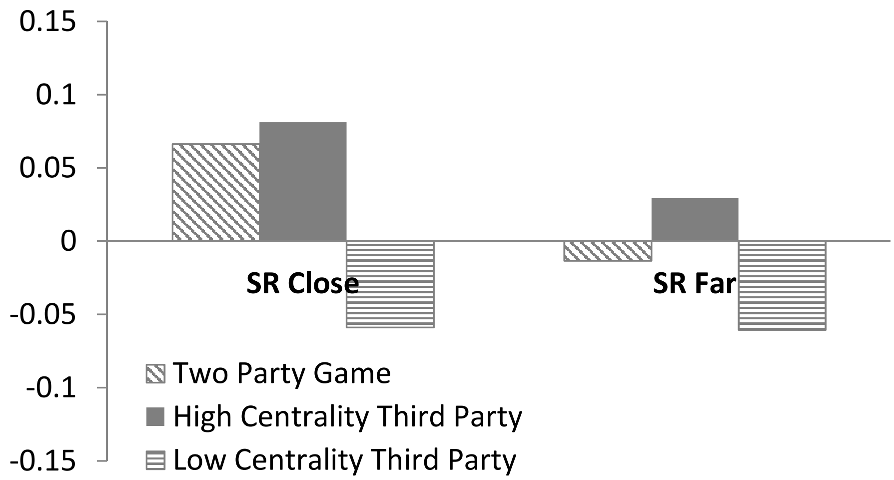

Panel B: Third-Party Punishers

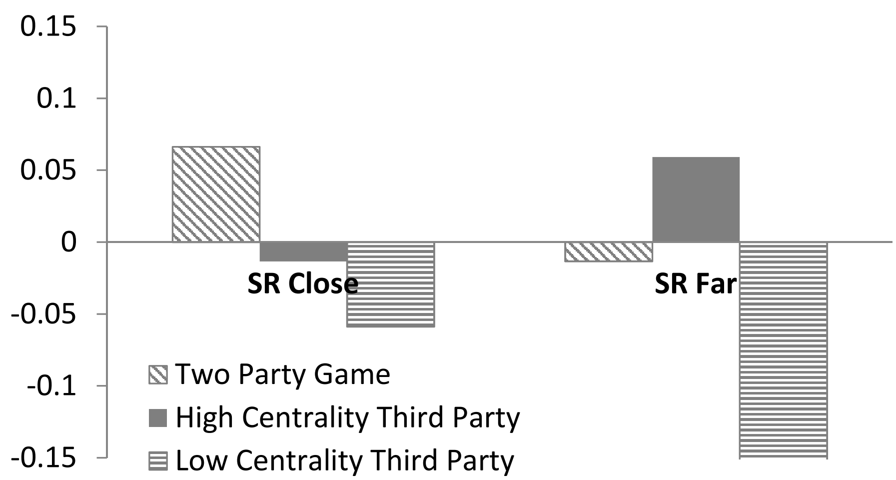

FIGURE 3. Normalized Sender Transfers by Game and Punisher Characteristics.

In all bar charts, the y-axis represents the percentage increase/decrease of the transfer in a cell relative to the average transfer size across all games (Rs. 28.43). In each grouping, the left-most bar shows percentage increase/decrease in transfers in the two-party game, the middle bar shows the percentage change in transfers in the three-party game with a third party of high centrality, and the right-most bar shows the percentage change in transfers in the three-party game with a third party of low centrality. Panel A compares the game with two players to the game with a third-party monitor. Panel B compares the game with two players to the game with a third-party punisher.

Tables

Table 1. Summary Statistics

	Number	Mean	Std. Dev.
Participant Characteristics			
Elite	930	0.1978	0.3986
High Caste	891	0.6667	0.4717
Female	930	0.5935	0.4914
Wealth Quantile (Village)	924	0.5296	0.2745
Education	899	8.1479	4.3237
Eigenvector Centrality Level	917	0.0225	0.0362
Eigenvector Centrality Quantile (Village)	917	0.5950	0.2652
High vs. Low Eigenvector Centrality	917	0.5300	0.4994
Group Characteristics and Outcomes			
Social Distance (S,R)	1790	3.5564	1.1387
Social Distance (S,J)	1136	3.5722	1.1043
Social Distance (R,J)	1134	3.5829	1.1218
Transfer S to R	1888	28.4370	15.3265
Fraction of S Transfer Returned by R	1874	0.5233	0.3524

Note: This table provides sample statistics for the experimental subjects. The participant characteristics are based on the sample of individuals who played our experimental games. The group characteristics and outcomes capture traits and transfers at the (S,R) or (S,R,J) level (depending on the game). The sample is restricted to the giant component of the social network.

Table 2. Correlation Matrix of Importance Measures

Correlations	Elite	Caste	Female	Wealth	Educ.	Eig. Cent.	Eig. Quantile	Eig. HvL
Participant Characteristics								
Elite	1.000							
High Caste	0.071	1.000						
Female	0.094	-0.040	1.000					
Wealth Quantile (Village)	0.165	0.298	-0.039	1.000				
Education	0.052	0.123	-0.220	0.198	1.000			
Eigenvector Centrality Level	0.117	0.095	0.070	0.173	-0.072	1.000		
Eigenvector Centrality Quantile (Village)	0.100	0.040	0.077	0.113	-0.174	0.635	1.000	
High vs. Low Eigenvector Centrality	0.079	0.034	0.110	0.059	-0.171	0.487	0.856	1.000

Note: This table presents the raw correlations (across individuals and villages) of the participant characteristics. The wealth and eigenvector centrality quantiles are all calculated within-village as is the High vs. Low Eigenvector Centrality measure.

Table 3. Principal Component Decomposition of Importance Measures

	Principal Components		
	1st PC	2nd PC	3rd PC
Participant Characteristics			
Elite	0.1343	0.2470	0.5854
High Caste	0.0802	0.5057	0.1148
Female	0.1241	-0.2658	0.7115
Wealth Quantile (Village)	0.1364	0.5799	0.1540
Education	-0.1512	0.5081	-0.2324
Eigenvector Centrality Level	0.4968	0.0816	-0.0942
Eigenvector Centrality Quantile (Village)	0.5969	-0.0516	-0.1683
High vs. Low Eigenvector Centrality	0.5618	-0.0940	-0.1515
Eigenvalue	2.4476	1.5264	1.0835

Note: The columns display the first three principal components in the principal component decomposition.

TABLE 4. Sender's Transfers and Importance Characteristics

	(1)	(2)	(3)	(4)
	Eigenvector Centrality Measure			
Outcome Variable: Transfer S to R	Quantile	Quantile	High vs. Low	High vs. Low
Third Party Centrality: Monitoring	1.338	5.154*	2.728**	3.182**
	(2.688)	(3.093)	(1.370)	(1.525)
Third Party Centrality: Punishment	8.063***	11.98***	4.943***	6.214***
	(2.537)	(3.121)	(1.297)	(1.575)
Third Party Elite: Monitoring		0.907		0.885
		(1.903)		(1.854)
Third Party Elite: Punishment		-0.521		-0.679
		(1.915)		(1.956)
Third Party High Caste: Monitoring		2.121		2.200
		(1.700)		(1.692)
Third Party High Caste: Punishment		-0.292		-0.332
		(1.620)		(1.585)
Third Party Female: Monitoring		0.134		0.125
		(1.435)		(1.372)
Third Party Female: Punishment		-1.923		-2.115
		(1.499)		(1.520)
Tests for Monitoring=Punishment: p-values				
Third Party Centrality	0.0298	0.0832	0.1683	0.1070
Third Party Elite		0.5877		0.5537
Third Party High Caste		0.2542		0.2398
Third Party Female		0.3395		0.2916
Controls: Experimental	Yes	Yes	Yes	Yes
Controls: Demographic	No	Yes	No	Yes
Observations	1,752	1,515	1,752	1,515
R-squared	0.243	0.295	0.248	0.298

Note: Regressions include observations from T1, T2, and T3. In all columns, the outcome variable is the amount transferred from S to R. In columns (1) and (2), the centrality measure used is the village quantile ranking of eigenvector centrality. In columns (3) and (4), the centrality measure is an indicator for being above the median of the eigenvector centrality of all game participants. Standard errors are clustered at the experimental session level, and all specifications include experimental session fixed effects and treatment fixed effects. All specifications include the following network controls: Centrality of S, Centrality of S: Monitoring, Centrality of S: Punishment, Centrality of R, Centrality of R: Monitoring, Centrality of R: Punishment, Social Closeness (S,R), Social Closeness (S,R): Monitoring, Social Closeness (S,R): Punishment, Social Closeness (S,T): Monitoring, Social Closeness (S,T): Punishment, Social Closeness (R,T): Monitoring, Social Closeness (R,T): Punishment. All columns additionally include experimental controls for sequence of games in session, round, and surveyor fixed effects. Columns (2) and (4) also include controls and their full interactions with treatment for wealth, age, education, and indicator for whether each pair of participants are members of the same household. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 5. Institutional Design

	(1)	(2)	(3)	(4)
	Eigenvector Centrality Measure			
Outcome Variable: Transfer S to R	Quantile	Quantile	High vs. Low	High vs. Low
Third Party Centrality: Monitoring	1.888	1.199	2.305	2.305
	(7.739)	(9.453)	(6.092)	(6.092)
Third Party Centrality: Punishment	25.43**	35.14***	19.44***	19.44***
	(9.732)	(9.937)	(5.278)	(5.278)
Social Closeness (S,R)	1.605*	1.312	1.554*	1.290*
	(0.847)	(0.791)	(0.783)	(0.733)
Social Closeness (S,R): Monitoring	-1.904	-2.287	-1.618*	-1.501
	(1.386)	(1.399)	(0.968)	(1.000)
Social Closeness (S,R): Punishment	1.547	1.894	0.389	0.345
	(1.631)	(1.702)	(1.142)	(1.231)
Third Party Centrality x Social Closeness (S,R): Monitoring	-0.128	0.915	-0.195	0.159
	(1.835)	(1.974)	(1.051)	(1.300)
Third Party Centrality x Social Closeness (S,R): Punishment	-3.901*	-5.245**	-2.520**	-2.924***
	(2.072)	(2.142)	(1.081)	(1.086)
Tests for Monitoring=Punishment: p-values				
Third Party Centrality	0.0508	0.0135	0.0685	0.0300
Third Party Centrality x Social Closeness (S,R)	0.1850	0.0344	0.1334	0.0641
Controls: Experimental	Yes	Yes	Yes	Yes
Controls: Demographic	No	Yes	No	Yes
Observations	1,752	1,515	1,752	1,515
R-squared	0.245	0.297	0.250	0.302

Note: Regressions include observations from T1, T2, and T3. In all columns, the outcome variable is the amount transferred from S to R. In columns (1) and (2), the centrality measure used is the village quantile ranking of eigenvector centrality. In columns (3) and (4), the centrality measure is an indicator for being above the median of the eigenvector centrality of all game participants. Standard errors are clustered at the experimental session level, and all specifications include experimental session fixed effects and treatment fixed effects. All specifications include the following network controls: Centrality of S, Centrality of S: Monitoring, Centrality of S: Punishment, Centrality of R, Centrality of R: Monitoring, Centrality of R: Punishment, Social Closeness (S,T): Monitoring, Social Closeness (S,T): Punishment, Social Closeness (R,T): Monitoring, Social Closeness (R,T): Punishment. All columns additionally include experimental controls for sequence of games in session, round, and surveyor fixed effects. Columns (2) and (4) include controls and their full interactions with treatment for each player for the following demographic characteristics: caste, elite status, gender, wealth, age, education, and indicator for whether each pair of participants are members of the same household. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Online Appendix

Appendix A. Tables

Table A.1. Sender Transfers and Third-Party Demographic Characteristics

	(1)	(2)	(3)
		Characteristic	
Outcome Variable: Transfer S to R	Elite	High Caste	Female
Third Party Characteristic: Monitoring	0.396	1.734	0.454
	(1.637)	(1.317)	(1.093)
Third Party Characteristic: Punishment	-0.430	0.386	-0.947
	(1.774)	(1.439)	(1.225)
Tests for Monitoring=Punishment: p-values			
Third Party Centrality	0.6453	0.661	0.3652
Controls: Experimental	Yes	Yes	Yes
Observations	1,844	1,665	1,884
R-squared	0.233	0.238	0.233

Note: Regressions include observations from T1, T2, and T3. In all columns, the outcome variable is the amount transferred from S to R. Each column plots regressions only considering one demographic characteristic including: elite status, high caste, and female gender. Standard errors are clustered at the experimental session level, and all specifications include experimental session fixed effects and treatment fixed effects. Each specifications include the full set of demographic controls for the given characteristic: Characteristic of S, Characteristic of S: Monitoring, Characteristic of S: Punishment, Characteristic of R, Characteristic of R: Monitoring, Characteristic of R: Punishment. All columns additionally include experimental controls for sequence of games in session, round, and surveyor fixed effects. *** p<0.01, ** p<0.05, * p<0.1.

Table A.2. Sender Behavior and Total Payoffs

	(1)	(2)
	Transfer S to R	Transfer S to R
T2: Game with Monitoring	-0.206	0.548
	(1.392)	(1.499)
T3: Game with Monitoring and Punishment	-1.567	-1.153
	(1.271)	(1.319)
Mean of Ommitted Category: Two-Party Game	29.081	29.081
Controls: Experimental	No	Yes
Observations	1,888	1,884
R-squared	0.219	0.230

Note: In all columns, the outcome variable is the amount transferred from S to R. Standard errors are clustered at the experimental session level, and all specifications include experimental session fixed effects and treatment fixed effects. Column (2) additionally includes experimental controls for sequence of games in session, round and surveyor fixed effects. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Appendix B. 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.

Path Length and Social Closeness. 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 $\gamma(i,j)=\infty$, though in our analysis we focus on the giant component of the graph. The *social closeness* between i and j is defined as $\max_{i',j'} d\left(i',j'\right) - d\left(i,j\right)$. This defines a measure of how close the two nodes are with 0 meaning that the path is of maximal length and $\max_{i',j'} d\left(i',j'\right) - 1$ meaning that they share an edge. In figure B.1, $\gamma(i,j)=2$ and $\gamma(i,k)=\infty$.

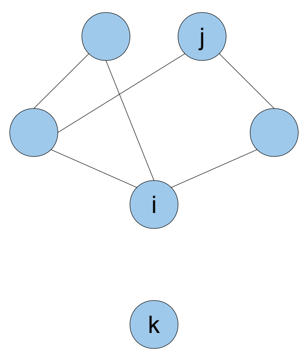

Figure B.1. Path lengths i,j and i,k

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

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.[27] Intuitively, this measure defines the importance of a node as proportional to the sum over each of its network neighbors' importances. 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. This measure captures how well information flows through a particular node in a transmission process. Relative to betweenness, 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 B(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.

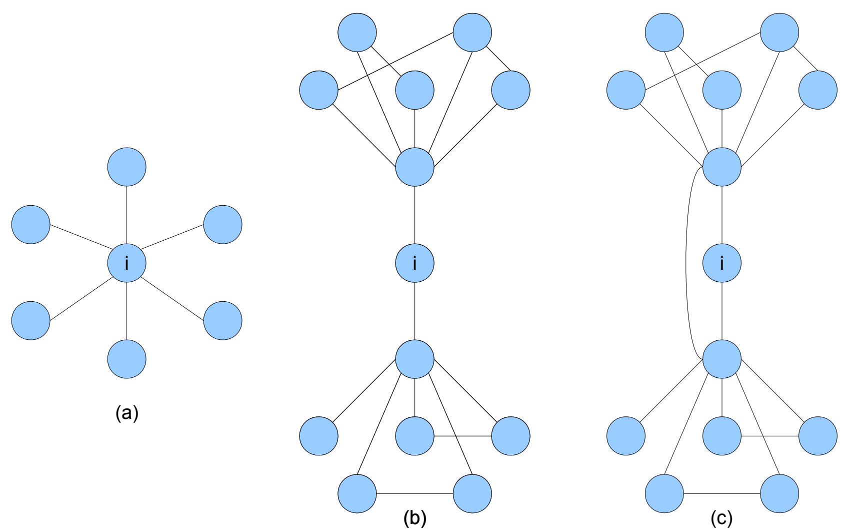

FIGURE B.2. Centrality of node i

[27]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.