NETWORK CENTRALITY AND INFORMAL INSTITUTIONS: EVIDENCE FROM A LAB EXPERIMENT IN THE FIELD

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

Abstract. While social closeness mitigates contractual incompleteness, we examine how communities can enlist third parties to improve cooperation between socially distant pairs. Network-central members may be particularly effective at this role through two channels: information and enforcement. We conduct modified trust games (with and without third parties) in 40 Indian villages to measure the effectiveness of central third parties. Assigning a punisher at the 75th percentile of the centrality distribution (versus the 25th) increases efficiency by 21%. 2/5 of the effect is attributed to information and 3/5 to enforcement. Central punishers are most valuable when senders and receivers are socially distant.

JEL Classification Codes: L14, Z13, O17, D02

Keywords: Social networks, Institutional structure, Informal contracting

1. Introduction

The literature on social networks has paid particular attention to "central" agents. They have been shown to effectively spread information, influence opinions and propagate shocks (e.g., Katz and Lazarsfeld, 1970; DeMarzo et al., 2003; Ballester et al., 2006; Golub and Jackson, 2010; Acemoglu et al., 2012; Banerjee et al., 2013; Cai et al., 2014). In each of these settings, the importance of a central agent is derived from the fact that her actions (or beliefs) influence her immediate neighbors. Those neighbors in turn influence their neighbors, echoing the effect throughout the network.

However, network centrality has important implications well beyond its role in information diffusion that has been studied in prior work. In particular, we focus on the fact that if two agents differ in how well their opinions spread in the network, they may have asymmetrical influence on each other's reputations. Rather than treat information transmission as the end goal, we ask how the centrality of agents can provide incentives for cooperative behavior and thus influence contract enforcement. Specifically, we focus on the role that third parties can play to improve the scope of transactions that require trust (e.g., loan repayment, quality of

---

Date: December 2015.

We thank Daron Acemoglu, Abhijit Banerjee, Sam Bowles, Luke Coffman, Esther Duflo, Pascaline Dupas, Matt Elliott, Itay Fainmesser, Garance Genicot, Ben Golub, Matthew O. Jackson, Jacob Leshno, Markus Möbius, Muriel Niederle, Rajiv Sethi, Ryan Sheely, Adam Szeidl and Juan Pablo Xandri. Important feedback was provided by numerous seminar participants. Manaswini Rao, Gowri Nagraj and Dhruv Parthasarathy provided excellent research assistance. We thank the Russell Sage Foundation, the George and Obie Shultz Fund, and the MIT CIS for financial support. Breza and Chandrasekhar thank the NSF GRFP and Larreguy thanks the Bank of Spain and Caja Madrid Foundation.

[†]Columbia Business School, Division of Finance and Economics. Email: `ebreza@columbia.edu`.

[‡]Stanford University, Department of Economics and NBER. Email: `arungc@stanford.edu`.

[⋆]Harvard University, Department of Government. Email: `hlarreguy@fas.harvard.edu`.

products supplied to a buyer, payment for services rendered).[1] This is an important question because 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.[2] 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.[3] In our setting, only 16% of possible contracting partners are of social distance two or closer. Thus, there is a natural role for third party institutions to contribute to contract enforcement. To study this, we conduct a lab-in-the-field experiment where two parties engage in an investment game, sometimes in the presence of a third party who may or may not be able to levy fines.

There are two main mechanisms through which central members of a community could be best suited to serve in the role of arbiter. First, because of their ability to influence others, they may be able to credibly threaten to punish non-cooperation well into the future, after the transaction has occurred (i.e., in the setting of this paper, after the experiment has finished). A more central agent observing an unfair transaction can more adversely affect the offender's reputation amongst other villagers, for example. Further, because, this central agent is also more likely to interact directly with the offender in the future, any sanctions levied by the central agent in the future may be more costly. We call this the *information effect*, referring to the fact that in some of our lab games the third party takes no action in the experiment and only gains information from monitoring the interaction. The third party's information effect comes through (the threat of) any reputational or punitive measures taken after the game has ended.[4]

Second, some institutional arrangements may allow for community members to levy fines or other types of direct punishments against offending parties (i.e., in the setting of this paper, during the course of the experiment), and central individuals may be well-suited to mete out these types of direct (in-game) punishments.[5] Again, because central agents have an outsized ability to affect the reputations of others and to distribute or withhold favors, they may also face less retaliation against direct punishments than do peripheral agents. We call this the *enforcement effect*.

[1]Despite the lack of empirical work in this space, it is related to some recent work in theory. For instance, Fainmesser (2014) shows that in environments where bilateral traders cannot maintain efficient exchange, central enough intermediaries can exploit their position to facilitate trade. Relatedly, there has been theoretical work on how network position affects reputational channels in bargaining and trading models (Abreu and Manea, 2012).

[2]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

[3]Examples include co-investment opportunities that arise with fellow community members who are not one's friend, job referrals (Beaman and Magruder 2012, Pallais and Sands 2014, Adelman et al. 2015) where the referrer must vouch for the referee, contributions to a public good involving more than one's local set of social links.

[4]See Appendix A for a discussion of evidence of this mechanism from field interviews.

[5]Examples may include monetary fines or exclusion from public goods such as grazing lands, labor market opportunities, or self-help savings or borrowing groups. Again, see Appendix A for field evidence.

We investigate these issues through lab-in-the field experiments in 40 rural villages in India, where formal contract enforcement is weak. Because we are precisely interested in understanding how positions of individuals in the actual community are brought into the games, all experimental sessions are non-anonymized. We use the network data of Banerjee et al. (2013), which provides information on the network characteristics of each player. In each experimental session, subjects play two- and three-party investment games (also called trust games, modeled after Berg et al. 1995, Charness et al. 2008, and Fehr and Fischbacher 2004) with high stakes.[6] Our goal is to determine if outcomes can improve under the presence of central third-party intervention and whether these improvements are due to the information or enforcement channels.

To measure the information effect, in a subset of games we assign a third party to simply observe the sender and receiver transfers. The logic is that such a third party can observe misbehavior and can either pass this information to others after the experiment or interact in the future with the sender or the receiver – possibly sanctioning – outside the game. Of course, those to whom she passes information may also apply sanctions in the future. As mentioned above we call this the information effect, noting that it potentially has an extra-game punitive component.[7] To isolate the enforcement mechanism, in a random set of games we further give the third party the ability to punish the receiver within the game. In this treatment, the third party watches, as before, but additionally can levy a fine – an observable punishment – on the receiver during the experimental session. This treatment gives the third party an immediate, verifiable punishment technology in addition to the information-based, extra-game punishment tools available in the information-only treatment. In both third-party treatments, we are particularly interested in how the effectiveness of the third party may increase with network centrality and decrease with the social proximity of the two players. By differencing these treatments, we can identify the institutional, within-game enforcement channel separately from the extra-game information channel alone.

We find that, 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 within the game. Decomposing this into enforcement and information, 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 information channel (an effect of 9% of the mean).

[6]Our work complements an experimental literature that investigates whether agents being members of the same group influences punishment (Goette et al., 2006; Bernhard et al., 2006; Chen and Li, 2009; Goette et al., 2012), discussed in Section 6.

[7]We name this the *information effect* because the third party simply observes (information about) the transactions within the game and cannot take any actions during the experiments. Any effects of the information treatment, however, derive from the third party's (actual or perceived) actions outside of the game, which again may include spreading information to others or meting out extra-game punishments. We are interested in capturing precisely this in the information effect.

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

Third, we show other natural measures of social hierarchy cannot match the effectiveness of a network-central punisher. Third parties (monitors and punishers) of elite status or high caste do not improve outcomes. Additionally, note that all of our network results condition on these demographic variables (as well as wealth, education and age) interacted with treatment, meaning that the network effects are not simply proxying for classical notions of social hierarchies or demographic variables.[8]

Finally, adding up over all triples of sender, receiver and third party, on average adding a third party to a two-party investment game neither increases nor decreases sender transfers. However, this masks a striking and predictable heterogeneity. Socially close senders and receivers achieve better outcomes on their own, and consistent with previous work by Cardenas et al. (2000) and Fehr and Gächter (2002), adding a third party generates some crowd-out. However, for otherwise inefficient pairs of socially distant senders and receivers, the third-party institution may improve outcomes. 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. This finding also resonates with the work by Cardenas et al. (2000) and Fehr and Gächter (2002) because peripheral third-parties introduce incentives that crowd out other-regarding behavior but are not powerful enough to induce cooperation. Similarly, Ostrom (1992) and Bowles and Gintis (2002) discuss how external intervention, when poorly designed, may undermine the extent to which members of the community can use social capital to solve collective action problems. Our results show that only a small set of people in the social fabric are able to overcome this negative effect.

When working with networks data, it is useful to clarify what we mean by a network interaction and how we define a central agent. A link between households in our data captures whether respondents indicate in a survey a strong social or financial relationship. Surely in village communities, any two arbitrary households interact on occasion, even in absence of a direct link in our data (e.g, one may gossip with someone who is merely an acquaintance at the local tea shop, one may learn of a job opportunity indirectly through a friend's relative). We interpret the empirical network data as a way to parametrize interactions; an individual is more likely to pass information to or meet with direct contacts, is less likely to pass information to or meet friends of friends, and is even less likely to interact with friends of friends of friends, and so on. Under such a framework for interactions on a network, observe that

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

certain households will be more central than others (reaching directly or indirectly more individuals).[9] This perspective is useful to think about the heterogeneous value of third parties by centrality in our experiment. The information effect follows from the fact that a more central third party passes information more widely ceteris paribus. The enforcement effect is a consequence of this as well – a more central agent can affect reputations of others to a greater degree, is better positioned to give out (or withhold) favors, and could be subject to less retaliation outside of the game. Our experiment has been designed to explore carefully these ideas.

We chose the investment game with a third party for its simplicity.[10] To study the arbitration of exchange without contracts, at minimum three parties are required: two involved in exchange and an arbiter. Additionally, this is a parsimonious way to represent a number of real-world institutions. For example, consider an individual borrowing from a peer to invest in a profitable venture who defaults. Or, an individual in a rotating savings and credit association (RoSCA) who misses a payment after having received a lump sum from the other RoSCA members in an earlier period. In the case of default, aggrieved parties often turn to community leaders to help resolve disputes. Punishment may involve exclusion from future loans, from future participation in self-help group activities or from future labor market opportunities, etc.[11] The types of interactions we study in our experimental games are contained in many complex informal institutions commonly observed in developing countries (Harper, 2011; Sandefur and Siddiqi, 2015). In particular, in India these include older informal institutions such as *nyaya panchayats* or informal village courts (Galanter and Meschievitz, 1982) but also more modern ones such as informal arbitrators in RoSCAs, self help groups (SHGs), village savings and loan associations (VSLAs), and microfinance borrowing groups (MFIs). In each of these cases, actions by members are observable to others and have direct payoff consequences to a subset of the group. Further, the (often outside) informal arbitrators usually do take actions to sanction member misconduct.

To summarize, without formal contracts, the fact that socially proximate individuals have repeated, systematic interactions allows them to sustain transactions that require trust. But broader interactions – involving members from the community at a distance – require mediation (e.g. a local adjudicator that arbitrates between lenders and debtors in default or a community member who monitors the use of common resources). 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 an information

[9]This simple, physical description of interactions among villagers suggests the use of eigenvector centrality as our measure of importance. This certainly is not the only sensible way to model interactions, and different models would generate predictions for a slightly different notions of centrality. However, the core idea would be the same.

[10]In our setting, we focus on direct interactions between parties lacking formal contracts and ask about the effects of third-party observers/punishers. In our setting, the sender and receiver directly interact. Related work by Coffman (2011) investigates whether third parties are more likely to perceive the same selfish behavior as a norm violation if the sender and receiver directly interact, as opposed to interacting through an intermediary.

[11]See Appendix A for more real-world examples elicited during one-on-one conversations with 60 villagers.

effect, their effectiveness comes largely from the fact that they are better equipped to utilize punishment technologies. Further, the centrality effect is conditional on demographic-by-treatment controls: thus it is not proxying for leadership status, caste or gender. This tells us that network position is crucial in institutional efficiency, and network importance may be able to substitute for a formal contracting technology.

2. Experiment

2.1. **Setting.** 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 hours' drive from Bangalore. The villages are independent with a median pairwise distance of 46 kilometers. Twenty-four individuals aged 18 to 50 were recruited from each village. We chose these subjects as we had access to village census demographics as well as unique social network data. The data set is described more below in Section 3 and in Banerjee et al. (2013). The network represents social connections between individuals in a village with 12 dimensions of possible links: e.g., relatives, friends, and creditors. As such, we have extremely detailed data on social linkages, not only between our subjects but also about the embedding of the individuals in the network 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 she belongs to any 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.[12] To construct a wealth index, we use survey information on house size, electrification, building materials, and toilet amenities.

Our experiment has three treatments (T1-3) described in detail below. T1 is a sender-receiver investment game while T2 and T3 add the presence of a third party. For every treatment, we randomized individuals into the roles of sender, receiver and third party and most subjects played each role at least once. Additionally, we ensured and told participants that no individual would participate in any game with any other individual more than once.[13]

2.2. **Recruitment and Implementation.** We randomly chose a subset of approximately 18 households to invite to participate from each village. We visited all invited households two days before the experimental sessions and told them the location and the starting time. On the morning of the experiment, invited households were given first priority to participate.[14] The remaining slots were filled by walking around the entire village and making announcements about the experiment. We did not allow multiple members of the same household to participate.

---

[12]In India*, gram panchayats* are local government institutions at the village or small town level and *anganwadi* centers are local educational and health centers.

[13]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.

[14]Approximately 20% of these households did ultimately participate.

During each experimental session, participants played five or six total rounds of the three experimental treatments: two rounds of each of T2 and T3 and one or two rounds of T1. In each round, players were randomly assigned to the roles of sender (S) with endowment Rs. 60, receiver (R) with endowment Rs. 60, and third party (T) with endowment Rs. 100. The games were played in random order.[15] 14 surveyors moderated the experiments, each overseeing only one group of participants at a time. Before the experiment began, we trained all of the subjects on the rules of each of the variants of the game and informed them about the number of experimental rounds that they would play.

We divided the public space into 14 "stations", one for each surveyor who would administer the experiment to a group of two or three subjects. Each round, the team of surveyors (re)explained the rules of the game played that round. Then, the subjects were all informed which station they were randomly assigned to. After arriving at the station, the subjects were randomly assigned roles (S or R for the two-party case and S, R or T for the three-party case) by drawing chits from a bag. Once the roles had been assigned, the surveyors asked, sequentially, the sender, receiver, and third party (if applicable) their decisions which were observed by all participants at the station. We intentionally did not prohibit communication. Because the surveyor was present and conducted the round in the manner described, in practice there was not much, if any, discussion among the participants.[16]

Surveyors were randomly assigned to different groups, and thus their presence is orthogonal to the treatments and participant characteristics. While there might be a concern that surveyors also played 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. 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.

After playing all rounds, participants were given their ending wealth values for one randomly-chosen game plus a fixed participation fee of Rs. 20. The average payoff was

[15]As a result, for all regression specifications, we can include both game and round fixed effects.

[16]We did not conduct the strategy method. Recall that our lab-in-the-field was conducted by having a makeshift lab in each of 40 independent villages. This involved using public spaces, schools, dairy cooperatives, and marriage halls, among other locations. Further, to ensure independence in rounds, 14 surveyors were in charge of running the game and each was matched to a pair or triple. However, managing a paper-based strategy method implementation that was kept private from other subjects proved to be impossible. Further, we allowed smoother amounts of transfers, which would have made the strategy method even more difficult, because – for instance – the third party with punishment would have required filling in a table with all possible sender amounts (7 amounts, even in deciles), all possible conditional receiver amounts (for instance, even by deciles it would be 24 possible cells for the maximal receiver amount), and then all combinations. So this would be up to 7×24 cells per decision, which would have been impossible to implement. Similarly, our initial lab protocol set out to measure the interim beliefs about the actions of other players. However, it was not possible to find a laboratory methodology given our lab-in-the-field setting that would keep these beliefs private.

~Rs. 110, or approximately three-fourths of a daily agricultural wage. We stress that these are large stakes – nearly a day's wage for an experiment lasting no more than an hour.[17]

2.3. **Two-Party Experimental Benchmark.** We begin with the two-party investment game (T1) as a benchmark and consider an environment with two parties who do not have access to formal contracting. Later, we add a third-party institution.

In this benchmark game, two participants are selected at random and are assigned the roles of S and R.[18] 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, if subjects are selfish materialists and only care about the static game payoff, R will return $\tau_R = 0$ in any Nash Equilibrium, and by backward induction, S never makes a positive transfer. If instead subjects have other-regarding preferences with common knowledge about them, the Nash Equilibrium would encompass positive transfers τ_S and τ_R. Importantly, regardless of individuals' social preferences, any efficient outcome is one where $\tau_S = 60$. Thus, as long as social preferences do not alone induce full efficiency, the 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 can reach in the absence of formal contracting.

Previous research in both the lab and the field suggests that social proximity may substitute for formal contracting institutions.[19] 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 $d(S, R)$ in our specific setting and also measure how far contracting outcomes between socially distant (S,R) pairs are from the efficient level.

2.4. **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 pairs are socially distant. There are two natural roles that a third party can play: one of information diffusion and one of enforcement. A third party can observe (mis)-behavior on the part of one of the two vested parties and can pass information to others in

[17] 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.

[18] Randomization into roles is essential to avoid confounding the treatments with individual characteristics. E.g., Karlan (2005) and Schechter (2007) argue that sender behavior might be confounded by individual risk-preferences.

[19] 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.

the community.[20] These others, or even the third party herself, may in the future interact with R, and thus, information about (mis)-behavior in the experiment passed on by the third party may affect R's future payoffs. In this way R's behavior in the game might shape the outcome of subsequent interactions with others outside the game through reputation. Ultimately, this information channel mimics the sort of interaction that happens in these communities day-to-day.[21]

Beyond an information effect, there might also be an enforcement effect. In some cases, the third party may be endowed with a direct, observable, and verifiable punitive mechanism (e.g., levying a fine) that can be exacted on a wrong-doing party. Moreover, 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. Note that employing a third-party punisher induces both information and enforcement.

Our design is structured to disentangle mechanisms of information diffusion from mechanisms of enforcement. We ask both if information and enforcement can (separately) improve outcomes, and if it matters who plays the role of observer 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 information variant (T2), we assign an individual T to watch the play of S and R but do not allow her to take any actions within the game. This does not preclude T, as well as any other individual outside the game, from sanctioning or rewarding R outside the game. Such a potential effect of R's payoff is what we understand as the incentives induced by the third-party observation.

In the enforcement variant (T3), we assign and individual T to watch the play of S and R, and we also allow her to take a costly punitive action in a third stage of the game. Specifically, T can spend her own resources to levy a monetary punishment on R. For every Rs. 1 spent by T, R loses Rs. 4.[22] While the information treatment does not preclude punishment of R's actions *outside* of the game, this punishment that might arise from monitoring is potentially much less observable to R. Further, when punishment occurs *within* the game, the third party must decide *herself* whether to use our provided punishment tool to take an immediate, verifiable action.[23]

Comparing T2 to T1 allows us to capture the information effect, while comparing T3 to T2 isolates the enforcement effect.

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

[21]These future payoffs may come from informal loans, advice, job opportunities, help in navigating government bureaucracy, etc. See Appendix A for evidence from conversations with 60 villagers.

[22]We did not vary 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.

[23]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.

2.5. **Network Environment and Framework.** Before analyzing our experimental results, it is necessary to present our framework for modeling interactions on a network and to describe how this framework maps generally to the analysis of empirical social network data and specifically to our experimental design.

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. The empirical social network data used in this analysis describes whether two arbitrary households, i and j, are linked. In our data, we define a link between i and j by whether household i or j names the other as a partner in either an exchange of financial or material goods or in a social activity (see Section 3.1, for a more detailed description of the network survey). The set of links in each village is summarized by an adjacency matrix $\mathbf{A}$, where $A_{ij} \in \{0, 1\}$.

Our binary definition of a link is in some sense arbitrary, as it is surely the case in village life that members of households who are not directly linked in the network – that is $A_{ij} = 0$ – may still have an opportunity to interact. What is essential for our representation $\mathbf{A}$ is that interactions are considerably more likely with those nodes we define to be direct links, less so for neighbors of neighbors, and even less so for individuals farther away in the network.

Our framework for interpreting network-based interactions is simple. We start by noticing, broadly speaking, that there are two main types of interactions in our networks. First, an agent can pass information to another agent. We suppose that this happens stochastically within each period, with information traveling from node i to j (or from j to i) with some fixed probability θ. Second, agents may meet others. Clearly individuals should be more likely to meet their friends than their friends of friends. A simple and plausible model for this type of interaction is to suppose that every node i travels to a neighboring node with probability θ, to a neighbor's neighbor with probability θ^2 (if there is only one such path there), and so on.

What is the expected number of times that a packet of information originating at i reaches all other agents? What is the expected number of meetings that i has with all other agents? The answer to both of these questions is given by $DC_i(\theta, T)$, defined as

$$DC_i(\theta, T) = \left[\sum_{t=1}^{T} (\theta \mathbf{A})^t \cdot 1\right]_i .$$

It is useful to highlight that as $T \to \infty$, $DC_i(\theta, T)$ converges to the eigenvector centrality of agent i (see Banerjee et al. (2014)). Given our framework for how information flows and how meetings occur in the network, we thus use eigenvector centrality as the relevant notion of importance for our empirical analysis.[24] Again, in the above framework for network

[24] We note that the goal of our paper, of course, is not to claim that this is the exact nature of social interactions. Surely slight alterations to the physical model of interactions correspond to slightly different notions of what it means to be important in the network. The key point is that once equipped with a simple framework describing how agents in the society interact, it sheds light on why we may be prone to see differences across treatments based on the network position of the parties.

interactions, eigenvector centrality captures the importance of an agent both in terms of information dissemination and physical meetings. Further, this has nothing to do with the agents participating in the experiment itself and has only to do with the assumed physical interactions on the network.

We can use this simple framework to think about basic questions such as the value of receiving a favor from a given node. Consider a favor, such as informing someone of a desirable and suitable job opportunity. Notice that holding other factors constant, a more central node generates better favors for an arbitrary agent: the central node is more likely to encounter the agent and is more likely to hear about suitable opportunities for that agent. Examples of such favors from our setting include giving a loan, helping out someone in a time of need, passing on information, providing information about job opportunities, providing job opportunities themselves, etc. For more detail see Appendix A.

Now let us return to the analysis of our experiment, where agents S, R and T, interact according to our various treatments. We are interested in how the location of T in the network, relative to R and S affects the efficiency of the transfers in the game. First, let us consider the information effect, identified by T2. The basic intuition is straightforward. A more central T corresponds, as noted above, to being able to pass information more widely throughout the network. Therefore, there may be incentives for receivers to return more and for the senders, in turn, to send more. Similarly, though distinctly, both S and R are simply more likely to interact more with a more central T, ceteris paribus, which is also directly implied by the above framework. Thus, the players may potentially have incentives to build reputation with more central T. It is not the aim of this paper to parse between these features but instead note simply that if a more central agent is part of an institution, without adding any formal punishment mechanism, the observation of misbehavior together with the threat that word of this may get out in the community could provide sufficiently steep incentives for agents to behave well. It should go without saying that it is not necessary for third parties to actually talk about the specific choices within the experiment with others. It could just be that the sender understands that her actions may change the third party's view of the sender, and it is more valuable to maintain a better standing in the eyes of a more central third party.

In fact, conversations with villagers, as described in Appendix A, demonstrate that individuals believe that information will spread more widely when the third party is more central in the network. To operationalize this, we asked each respondent to consider a hypothetical situation where one party wrongs another in the presence of a third party. We asked them to consider two cases of this hypothetical scenario. First, the third party is a randomly chosen name from the village (case 1) and second, we provide them with the name of a randomly chosen central individual in the network (case 2). We find that 94.7% of the time respondents feel that more people in the village come to hear of the infraction in case 2 as compared to case 1.

In a follow-up field experiment, Breza and Chandrasekhar (2015) show the importance of the information effect generated by central individuals in another context. The authors facilitate the opening of savings accounts for 1,300 desired savers and also help participants to set a realistic six-month savings goal for themselves. Importantly, some savers are then randomly assigned to receive what Breza and Chandrasekhar (2015) call a peer monitor, who receives bi-weekly updates of the saver's progress, The findings indicate that an average monitor increases total savings by 35%, and increasing the monitor's network centrality by one standard deviation increases savings by 14%. Breza and Chandrasekhar (2015) also present supplemental evidence that the information effect is likely driving their result. First, 63% of monitors report telling others about the saver's progress, and 25% of savers can identify specific instances of the monitor passing this information to others. Second, over one year later, arbitrarily surveyed villagers are more likely to know if the saver exceeded her goal if the monitor was central. Further, when asked to assess the general ability of the savers to achieve goals, these survey respondents are more likely to have a negative opinion of individuals who missed their goals or positive opinions of individuals who reached their goals when the assigned monitor was more central.

Second, consider the enforcement effect, identified by T3-T2. For similar reasons we expect central individuals to be better able to levy direct fines within the game. In game T3, when T decides whether to punish R within the game, she must trade off any benefits from punishing (e.g., upholding a social norm, signaling to others) with any costs (e.g., retaliation by R). Because central agents are better able to affect the reputations of others and because they are better positioned to give out (or withhold) favors, they may also be less subject to retaliation outside of the game by the individuals that they punish within the game. One intuitive example is that a picked-on student in a middle school may make for a poor hall monitor; she would likely be incapable of assessing fines on other kids for fear of retaliation on the playground (e.g., being ostracized). Another example comes from Mexico, where the local police often wear masks, hiding their identities, in order to prevent retaliation by the criminals whom they arrest (Malkin, 2009).

An interesting distinction between T2 and T3 is that in T2 neither S nor R can verify the actions that T takes to punish any party through, say, communicating a lowered reputation or opinion of her. On the other hand, in T3 the parties directly observe T's decision to punish R. This means that T knows that R (or even S) could in principle react and retaliate to this verifiable action. Consider then the following scenario holding the identity of R fixed as well as some arbitrary degree of misbehavior committed by R. Now let's compare the scenario with a central T to that of a peripheral T when in both cases T is given the verifiable punishment technology. From our discussion above, it is clear that R stands to lose more by retaliating against a more central T than a peripheral T. This potentially provides a central T with a freedom to exercise a fine that a peripheral T may not have.

While we find retaliation to be the most plausible explanation for our observed enforcement effect, surely there are other channels that may also generate similar patterns. Examples include a distaste for or pleasure from punishing that varies with T 's centrality, signaling to others through punishing (central T may derive different benefits), or heterogeneity in the norms for who in a community has the authority to punish others. It is not our goal to parse between all possible explanations, so when we describe the enforcement effect, our interpretation pools this set of channels. What is important is that our design allows us to ask if randomly giving a punishment tool to a third-party monitor improves outcomes, and whether this improvement is differential across central and peripheral third parties. The answer may have important implications for institutional design.

We also 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.[25]

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.[26]

Elites include *gram panchayat* members, *anganwadi* teachers, etc. While ex ante powerful individuals may be more likely to reach these positions, the roles themselves may also generate influence. 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 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). We may expect that women might be less respected as third parties, especially when they have the ability to punish.

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. We have a census of every individual in every village and network data collected across the following 12 dimensions: "(1) those who visit the respondents' home, (2) those whose homes the respondent visits, (3) kin, (4) non-relatives with whom the

[25] It is also possible that each of these characteristics is correlated with network position. We explore these correlations in Section 3.3.

[26] 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).

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). 100% of households were included in each village census, and on average 48% of households were administered the network survey module. Households could name as many responses per question as they wanted. As all households were included in the census, even links with households that did not participate in the network survey were often mentioned.[27]

Careful analysis of empirical network data requires both many networks (because of the interdependencies within each network) as well as high-quality network data without missing much information (see Chandrasekhar and Lewis (2011)). Having access to many networks allows for network-level fixed effects and standard error clustering at a high level, while high sampling rates allow for the study of more complicated network statistics such as minimum path length and eigenvector centrality. Our research design with 40 networks and information about 73% of all possible links between households allows us both of these and is at the frontier of the existing studies. Other networks studies are conducted either with considerably fewer networks with higher sampling rates (e.g., Karlan et al. (2009) (2) and Leider et al. (2009) (1)), or with more networks with significantly lower sampling rates (e.g., Attanasio et al. (2012) (70)). Moreover, Chandrasekhar and Lewis (2011) conduct Monte Carlo simulations that indicate that the eigenvector centrality effects (on which we focus our analysis) are likely to be attenuated in the presence of partial network sampling.

To construct our network $\mathbf{A}$, we build an undirected, unweighted graph taking the union over the twelve dimensions 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 of this decision. Any two households are linked if any member has any relationship with anyone else. This is reasonable as the multiple dimensions are highly correlated. The union network ensures that we take into account any possible meaningful relationship, without constructing an ad hoc weighting procedure.

3.2. **Outcomes.** 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.[28] The modal transfer is 20, with

[27]We are only missing any links between two households that both did not participate in the networks module. This means that the probability of missing a link is 0.52^2.

[28]Participants could make transfers in increments of Rs. 1.

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). 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.

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 received by R. Note that most of the receivers transfer weakly less than the amount received from the sender, leaving receivers with quantities at least as high as their initial endowments.[29] Only 5% of games ended with the receiver sending more back to the sender than was initially received. 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 under selfish preferences.

While we do show the distribution of receiver transfers in Figure 2, we do not use receiver transfers as outcomes in our regression analysis. It is tempting to analyze both receiver behavior (amount returned as a function of transfer, conditional on treatment and network position of all) and third party behavior (whether there was a fine and the amount, conditional on receiver behavior, sender behavior, treatment and network position of all). The core problem is that receiver and third party behavior are subject to endogeneity concerns.[30] Had we conducted the strategy method, we could have used that information. Unfortunately, as noted in Footnote 16, it was difficult to conduct the strategy method in the makeshift labs which we had to temporarily create in each of the 40 independent villages.[31]

3.3. **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.[32] About 67% of the participants are high caste, which includes

[29] At least before the punishment decision is made.

[30] Both the sender and the receiver are forward looking, and as such, they anticipate the behavior of the receiver and the third party as function of the network position of all individuals. As a consequence, the behavior of the sender and receiver already incorporates equilibrium play and cannot be considered as exogenous to understand the behavior of the receiver and the third party.

[31] In the explanation provided in the paper, holding S and R play fixed, the punishment strategy is increasing in third party centrality. The main problem in testing this, however, is that we only observe equilibrium responses by the sender and receiver to the punishment threats posed by the third parties. When the receiver decides how much to return, she considers the response of the punisher. Therefore, if punishment is indeed increasing in centrality (conditional on R transfers), there may be a centrality level above which the receiver becomes unwilling to incur punishment in equilibrium and raises her transfer. Thus, depending on where this threshold is set, we may observe equilibrium punishment either increasing or decreasing in centrality. It is easy to write a simple model generating this phenomenon. A similar issue is true when we consider receiver behavior. However, sender transfers are by construction immune to this, as they do not depend on previous in-game play but only on treatment assignment, network position, covariates and unobservables.

[32] This means that on average, an individual had attended 8th standard, which is the last year of primary education.

general or "otherwise backwards" (OBC) castes.[33] Finally, 20% of participants are a leader or local elite.

Turning to network characteristics, 96.5% of pairs are reachable (there exists a path through the network connecting the two).[34] Between the reachable pairs, the maximum social distance is 8, while the average social distance is approximately 3.6. The average eigenvector centrality is 0.02 with a standard deviation of 0.04, indicating that there is substantial heterogeneity in social importance across individuals. Compared to the average network characteristics from Banerjee et al. (2013), our sample of participants comes from slightly more central households – the average centrality quantile in the experimental sample is 0.59 compared to a population-wide mean of 0.5. However, the average path length between participants is slightly longer than in the overall population – the sample average is 3.6, while the population average is approximately 2.8.

Finally, we explore the relationship between the network and demographic characteristics. To do this, in Table 4 we present a correlation matrix of as well as a principal component matrix. While the various functions of eigenvector centrality are highly correlated,[35] 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 provides suggestive evidence as to why, even when conditioning on treatment-physical covariates interactions in regressions, our network-based results remain extremely robust.

Similarly, in Table 5, 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 reinforces 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 pair interaction, 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.

[33] There are three standard caste categories in India: general merit (GM); scheduled caste and scheduled tribe (SC/ST); and other backward caste (OBC). The SC/ST group is traditionally the most disadvantaged. Our indicator for "high caste" groups the GM and OBC designations.

[34] We condition our sample on this set of reachable pairs for all of the analysis.

[35] 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.

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 C), wealth of all three 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 information channel. $\gamma_{T3} - \gamma_{T2}$ measures the relative return to centrality of T through the enforcement channel (net of the information 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 an information channel, an enforcement channel or both.

Observe that we are chiefly interested in heterogeneous treatment effects based on third-party 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 (as we do) that among all observable characteristics only network-central punishers improve efficiency, then we can be confident that informal enforcement mechanisms will only work well when the right individual – a central individual – is chosen.

5. Results

5.1. **Preliminaries.** Table 2 presents how transfers vary by treatment. The outcome variable, as described above, is the transfer made by the sender to the receiver. The omitted category is the two-party game. In Column 1 we include no controls for session fixed effects, surveyor fixed effects, round fixed effects, and sequence of game fixed effects. In Column 2 we include all these so-called experimental controls. We find no significant effects of introducing a third party nor of giving the third party the punishment technology, and if anything there is a noisy negative point estimate. As alluded to above, in this research agenda we are interested in the (predictable) heterogeneity underlying this: are central third parties particularly good? is there less of a need for third parties when the transacting parties are socially proximate?

5.2. **Network centrality.** We now address the main question 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 within-game enforcement as compared to an information effect?

Table 3 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. We consider centrality percentile instead of absolute centrality to have a comparable measure of importance across villages, while we focus on the median for the sake of power. Finally, in columns 1 and 3 we do not use demographic-by-treatment controls, which we introduce in columns 2 and 4.

We present only the coefficients of third party centrality interacted with treatment or demographics of the third party interacted with treatment because of the numerous number of coefficients. Notice that even when considering columns 1 and 3, where we only include network variables, there are at least 15 other variables (centrality of R and S and their interactions with treatment, all pairwise proximities between (S, R, T) and all interaction of these with treatment, and treatment dummies) in the regression. Columns 2 and 4 now adds numerous more interactions with demographics. In Appendix Table B.2, we present the entire table for columns 1 and 3 up to experimental controls (we omit presenting variables consisting of round, sequence, and surveyor fixed effects).[36]

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 from the third party information effect at the top of the inter-quartile range of the centrality distribution.

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 our 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

[36]The much larger table for columns 2 and 4 with all terms are available upon request.

with treatment-varying effects.[37] Columns 1 and 3 show the same regression results without demographic controls as well. The results for the enforcement channel are remarkably robust. However, removing the controls generates considerable noise in the estimates of the information 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 + information treatment (T3) than information alone (T2). The difference in these effects is statistically significant in most specifications.[38]

Finally, we can ask what happens if a sender-receiver pair is randomly assigned the least central third party? Note that one cannot read this off the main effects on treatment from Appendix Table B.2, because the treatment dummies alone show effects where we set all other network characteristics to zero as well. However, we can add up appropriately (which is equivalent to estimating heterogeneous effects by third party centrality alone) and find that assigning the least central third party monitor generates a Rs. 1.08 decline in transfers from S to R, though not statistically significant (p-value 0.56). Further, adding the least central third party who can punish generates an additional Rs. 4.87 decline in transfers which is statistically significant (p-value 0.018). Note that from Table 2 we know that the net effect of a third party is a statistically insignificant negative, so given that highly central third parties have sizable and significant effects relative to the mean, it goes without saying that the worst possible third parties actually reduce efficiency. Though exploring the precise mechanism of this crowd-out is beyond the scope of this paper, note that it is consistent with ideas from Cardenas et al. (2000) and Fehr and Gächter (2002) that intrinsic motivation may be crowded out by introducing external incentives.[39]

Taken together we find 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 information 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 separate the information 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 predicts efficiency and to decompose this efficiency effect along information and enforcement channels.

5.3. **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 with network position. Our

---

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

[38]We can also provide evidence that the global notion of network centrality (eigenvector centrality) is driving the effect, and not a more local measure (degree). Available upon request are tables that include degree interacted with treatment in the main specification.

[39]Crowd-out may also arise from more complicated mechanisms. For example, the sender may balk at the idea of an otherwise peripheral and irrelevant member of their community exercising control in this setting. Alternatively, peripheral punishers may have less information about relationships between the sender and receiver and may have a harder time interpreting low receiver transfers.

main specification includes demographic-by-treatment regressors, which allow us to control for them when looking at centrality effects.

As noted in Table 4, demographic variables tend to be much more correlated with other demographic variables and less with network variables. Similarly, in the principal component analysis in Table 5, we observe that demographic and network characteristics pick up different dimensions of variance. This suggest that it is unlikely that network variables are going to merely proxy for other standard demographic features.

We further 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 information channel or through the enforcement channel, unlike the case of network centrality.

Finally, in Appendix Table B.1, we present results demographic variable by demographic variable (neither conditioning on each other nor on network centrality). Even in these cases, there is no association between elite, caste or gender status of the third party on the sender's transfer. This also implies that the effectiveness of an arbiter is likely not coming from any particular elite, caste or gender status.

5.4. **When can third parties help?** Surely, the presence of third parties should be more useful when there is more scope for improvement, as in the case of socially distant sender-receiver pairs. 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 ask in what contexts will a third-party institution be most successful. Are third parties 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 Table 2, that on average introducing a third-party monitor or punisher has no net effects, and that peripheral third parties can even make things worse. Despite third parties having no detectable mean effects, the central 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 for 10 different game configurations. Both panels include the two-party games (left-most in each grouping) alongside results from the games with third parties with high (center in each grouping) and low (right-most in each grouping) centrality.[40] Panel A considers the results from the games where third parties are punishers (information + enforcement effect), while Panel B shows the results from the games where third parties only monitor (information effect). We further consider cases where S and R are of close social proximity (left groupings) versus far social proximity (right groupings).[41] The bar charts then illustrate many of our key networks results but also allow for comparisons between the three- and two-party games.

[40]We say that T is of high centrality if she has above-median eigenvector centrality among the individuals playing from her own village.

[41]We say that S and R are close if their distance is at most two.

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 distant. 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 Cardenas et al. (2000) and Fehr and Gächter (2002), we observe that the extrinsic incentives introduced by the third-party punisher (but not the monitor) crowd out transfers when S and R are socially close. Cardenas et al. (2000) and Fehr and Gächter (2002) show that individuals confronted with regulation and incentives contracts exhibit less other-regarding behavior, which manifests to a larger extent among socially close individuals. Related, Bowles and Gintis (2002) and Ostrom (1992) describe situations under which external intervention can undermine existing community-based solutions.

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:

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

$Closeness(S,R)_{jgv}$ 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.

Table 6 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.

For the furthest pairs, going from the 25th to the 75th percentile in centrality of the third party punisher 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).

Note 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 information channels becomes difficult. While we are unable to reject that there is no information effect due to noise, we can only separate the effect of enforcement from the (noisily estimated) effect of information in the specifications with the full demographic controls (columns 2 and 4).

The exercise here provides evidence 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 distant contracting pairs occur when the individual with the punishment technology is socially important. Let us return to our example of a farmer using another villager's tractor and paying for the service after harvest. If the individuals are socially close, this arrangement should function well. If the individuals are socially far, the transaction is unlikely to occur on credit without some institutional arrangement. If the arbiter had been (in a counterfactual) a peripheral member of the social network, this would be particularly ineffective. However, if the arbiter is someone who is central in the network, then the efficiency-enhancing transaction should be more likely to occur.

6. Discussion

We now take stock of the nascent but vibrant experimental literature studying real-world social networks.[42] In one of the earliest strands of this literature, researchers have studied the interaction of other-regarding preferences with social networks. Several studies have used variations of the dictator game to identify why people give to others (Leider et al. 2009, Ligon and Schechter 2012, Goeree et al. 2010, and Branas-Garza et al. 2010). These studies generally find that directed altruism is an important component of dictator transfers and that transfers generally increase with social closeness and with the degree of the dictator. There is also evidence that friends arrive at better trust game outcomes than strangers (Glaeser et al., 2000; Barr, 2003; Binzel and Fehr, 2013).

A second strand of the literature investigates the extent to which social networks can substitute for commitment contracts. This is particularly important in informal insurance relationships. Chandrasekhar et al. (2013) randomly assign both the identities of pairs of individuals in a network and the available contract structures available to them in multi-period laboratory games. They find that only when parties are socially close are they able to overcome contractual incompleteness.

A third strand focuses instead on how individuals endogenously sort into groups (Barr and Genicot 2008, Barr et al. 2012, and Attanasio et al. 2012). The research finds evidence that individuals tend to select socially proximate group members, which is consistent

[42] Our discussion focuses on complete information environments. Another strand of the literature explores hidden income. While classical models suggest that proximity may relax information constraints (Cole and Kocherlakota 2001, Chandrasekhar et al. 2011) there may be countervailing forces including outsized demands from friends and family (Banerjee and Mullainathan 2010; Jakiela and Ozier 2012).

with the second strand of the literature. If the most efficient outcomes under contractual incompleteness are precisely when individuals are proximate, then when given a choice of interaction, we should expect such a pattern.

Taken together, this literature finds benefits of networks in incomplete contracting environments, but also deep limitations. Agents are cognizant that efficiency is only possible amongst those who are proximate. So, they select precisely those who are socially close as their partners. This means that in many environments, particularly those where diversification may be important, people may choose not to transact with those from other groups. Note that in our sample 84% of pairs of individuals are socially distant, and therefore precluded from efficient transaction without commitment contracts. This is where we are able to make a new contribution, by focusing on when arrangements between parties that otherwise would not be happening at a pair level would be facilitated by a third party institution in the network.

This work complements an experimental literature that investigates whether agents being members of the same group influences punishment (Goette et al., 2006; Bernhard et al., 2006; Chen and Li, 2009; Goette et al., 2012). For instance, Goette et al. (2006) assigns individuals randomly to different platoons in the Swiss army for four weeks during training and then has subjects play a simultaneous prisoner's dilemma game possibly with third-party punishment. Consistent with our findings, they find that agents randomly assigned to the same platoon cooperate more and third parties are more willing to enforce the cooperation norm towards fellow-platoon members. We explore how the location of the third party within the social fabric, the agent's centrality, influences her efficiency or lack thereof and attempt to begin to unpackage the mechanism.

Using random matching of pairs and triples, we measure the causal effect of having a third party in an incomplete contracting setting and show that the network centrality of this agent influences efficiency. We further decompose the efficiency increase into an information 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). The largest efficiency increase occurs when senders and receivers are socially distant, unable to maintain efficient levels autonomously. We show that not every member is equally well-equipped to be part of a local institution. While the most central individuals do improve transfers, the wrong (i.e., peripheral) third party can actually decrease transfers. We also find that demographics 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.

Of course our findings raise several questions for future research: When individuals endogenously form groups to engage in a transaction, would the presence of an institution run by a central third party enable them to transact efficiently with those who are less socially proximate? Further, when a group needs to endogenously form a third party institution, do they select socially central individuals?

It is important to note that 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, 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. Importantly, as Appendix Table B.3 indicates, the conclusion of our analysis regarding the effect of a central third-party monitor or punisher remains unchanged if we account for the efficiency cost of the inflicted punishment.

Our results are consistent with the motivating idea that external intervention is only necessary when repeated game dynamics or social preferences are not 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.

References

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

Abreu, D. and M. Manea (2012): "Bargaining and efficiency in networks," *Journal of Economic Theory*, 147, 43–70.

Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi (2012): "The network origins of aggregate fluctuations," *Econometrica*, 80, 1977–2016.

Adelman, S., V. Hoffman, M. Mobius, and T. Rosenblat (2015): "Do Friends Help Friends Get Jobs?" *Working Paper*.

Attanasio, O., A. Barr, J. C. Cardenas, G. Genicot, and C. Meghir (2012): "Risk Pooling, Risk Preferences, and Social Networ," *American Economic Journal: Applied Economics*, 4, 134–167.

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

Ballester, C., A. Calvó-Armengol, and Y. Zenou (2006): "Who's who in networks. wanted: the key player," *Econometrica*, 74, 1403–1417.

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. and S. Mullainathan (2010): "The shape of temptation: Implications for the economic lives of the poor," *NBER Working Paper*.

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

Barr, A. (2003): "Trust and expected trustworthiness: experimental evidence from zimbabwean villages," *The Economic Journal*, 113, 614–630.

Barr, A., M. Dekker, and M. Fafchamps (2012): "Who Shares Risk with Whom under Different Enforcement Mechanisms?" *Economic Development and Cultural Change*, 60(4), 677–706.

Barr, A. and G. Genicot (2008): "Risk Sharing, Commitment, and Information: An Experimental Analysis," *Journal of the European Economic Association*, 6, 1151–1185.

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.

Beaman, L. and J. Magruder (2012): "Who gets the job referral? Evidence from a social networks experiment," *The American Economic Review*, 102, 3574–3593.

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

Bernhard, H., E. Fehr, and U. Fischbacher (2006): "Group affiliation and altruistic norm enforcement," *The American Economic Review*, 217–221.

Binzel, C. and D. Fehr (2013): "Social distance and trust: Experimental evidence from a slum in Cairo," *Journal of Development Economics*, 103, 99–106.

Bowles, S. and H. Gintis (2002): "Social capital and community governance*," *The Economic Journal*, 112, F419–F436.

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

Branas-Garza, P., R. Cobo-Reyes, M. P. Espinosa, N. Jiménez, J. Kovářík, and G. Ponti (2010): "Altruism and social integration," *Games and Economic Behavior*, 69, 249–257.

Breza, E. and A. Chandrasekhar (2015): "Savings Monitors," *Working Paper*.

Cai, J., A. de Janvry, and E. Sadoulet (2014): "Social Networks and the Decision to Insure," *American Economic Journal: Applied Economics*.

Cardenas, J. C., J. Stranlund, and C. Willis (2000): "Local Environmental Control and Institutional Crowding-Out," *World Development*, 28(10), 1719–1733.

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

——— (2013): "Can Networks Substitute for Contracts? Evidence From a Lab Experiment in the Field," *Working Paper*.

Chandrasekhar, A. G. and R. Lewis (2011): "Econometrics of Sampled Networks," 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.

Chen, Y. and S. X. Li (2009): "Group identity and social preferences," *The American Economic Review*, 431–457.

Coffman, L. C. (2011): "Intermediation reduces punishment (and reward)," *American Economic Journal: Microeconomics*, 77–106.

Cole, H. and N. Kocherlakota (2001): "Effcient Allocations with Hidden Income and Hidden Storage," *Restud*, 68, 523–42.

DeMarzo, P., D. Vayanos, and J. Zwiebel (2003): "Persuasion Bias, Social Influence, and Unidimensional Opinions*," *Quarterly journal of economics*, 118, 909–968.

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*.

Galanter, M. and C. S. Meschievitz (1982): *The Politics of Informal Justice*, New York: Academic Press, chap. In Search of Nyaya Panchayats: The Politics of a Moribund Institution.

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.

GOETTE, L., D. HUFFMAN, AND S. MEIER (2006): "The Impact of Group Membership on Cooperation and Norm Enforcement: Evidence Using Random Assignment to Real Social Groups," *The American Economic Review*, 96, 212.

GOETTE, L., D. HUFFMAN, S. MEIER, AND M. SUTTER (2012): "Competition between organizational groups: Its impact on altruistic and antisocial motivations," *Management science*, 58, 948–960.

GOLUB, B. AND M. JACKSON (2010): "Naive Learning in Social Networks and the Wisdom of Crowds," *American Economic Journal: Microeconomics*, 2, 112–149.

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

HARPER, E. (2011): *Customary Justice: From Program Design to Impact Evaluation*, Rome: International Development Law Organization.

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.

JAKIELA, P. AND O. OZIER (2012): "Does Africa Need a Rotten Kin Theorem? Experimental Evidence from Village Economies," .

KARLAN, D. (2005): "Using experimental economics to measure social capital and predict financial decisions," *The American Economic Review*, 95, 1688–1699.

KARLAN, D., M. MOBIUS, T. ROSENBLAT, AND A. SZEIDL (2009): "Trust and Social Collateral*," *Quarterly Journal of Economics*, 124, 1307–1361.

KATZ, E. AND P. F. LAZARSFELD (1970): *Personal Influence, The part played by people in the flow of mass communications*, Transaction Publishers.

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," .

MALKIN, E. (2009): "Revenge in Drug War Chills Mexico," *The New York Times*, December 22.

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.

Ostrom, E. (1992): "Community and the endogenous solution of commons problems," *Journal of Theoretical Politics*, 4, 343–351.

Pallais, A. and E. G. Sands (2014): "Why the Referential Treatment? Evidence from Field Experiments on Referrals," *Working Paper*.

Sandefur, J. and B. Siddiqi (2015): "Delivering Justice to the Poor: Theory and Experimental Evidence from Liberia," Working paper.

Schechter, L. (2007): "Risk Aversion and Expected-Utility Theory: A Calibration Exercise," *Journal of Risk and Uncertainty*, 35(1), 67–76.

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

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

Appendix A. Field Interviews

We spoke with 60 individuals in 4 different villages to gain insight into the mechanisms that villagers perceive are operating when there are disputes between two individuals, possibly involving a third party arbiter. We held our conversations in villages that were not included in our experimental sample to avoid priming them about our game.

The villagers offered us many real-world examples of situations where two community members, say A and B, might engage in a transaction where B has the opportunity to wrong A. The most common example, B defaulting on an informal loan given by A, most closely matches the investment games that were played in our sessions. However, many other types of scenarios also emerged that could loosely map to our experiment. For example, A and B might be engaged in a dispute about land or jewelry (both of which often serve as collateral for informal loans given by A), B may be the representative of a local dairy co-op who is withholding payments from farmers, B might be a trustee of a local temple misappropriating funds, or B could be a service provider practicing caste-based discrimination.

Further, we wanted to understand any perceived reputational consequences of B "wronging" A, and whether villagers feel others talk about these issues. To push this idea even more, we randomly supplied each respondent with two names: the name of a randomly chosen individual from the community, and the name of a central individual in the village network.[43] By asking respondents about what would happen if the random or central individual was the third party who observed an infraction, we can get some feel for what the villagers are thinking. We are interested in whether villagers feel that central individuals will spread information further. Our focus group analysis is consistent with this interpretation.

A first question concerns how the third party would update if she learned that B wronged A. Free-form responses demonstrated that villagers felt that the third party may not respect the perpetrator, may not trust the perpetrator, may think that the perpetrator is a cheat or – if serious enough – a thief, or more generally that the perpetrator is a "bad person".

We then asked what the social consequences might be to B. One consequence is that the third party might spread the information to others. When the identity of the third party supplied to the respondent was the central individual's name, the respondents on average said that 52.9 (standard error 5.8) other households would come to hear about the infraction. However, when the name supplied was a random draw, then on average respondents said that 28.2 (standard error 6.0) other households would come to hear about the infraction. Further, since each respondent was asked to respond to this question for both a central and a random individual, we can see what share of individuals said that more households would come to hear when the central individual was the third party as compared to the random individual. In fact, 94.7% of respondents said that more people would hear if the central individual was the third party.

Furthermore, this diffusion effect is likely a composition of either the third party directly informing others or others further continuing the gossip. When considering a question about

[43] The central name was supplied by using the gossip nomination procedure of Banerjee et al. (2014).

what happens when a RoSCA member fails to repay, 53.3% of respondents said that the fellow group members would tell others in the village and 23.3% of individuals stated that if others came to learn of this behavior they would also gossip about it with their friends and acquaintances, both of which are very high rates. While it is not surprising that the indirect gossip rate is less than the direct gossip rate, that it is nearly 1/4 indicates that even indirect information diffusion is important.

Next, we explored the social and economic consequences when others come to hear about the fact that B wronged A. Among other things, respondents indicated that they "would not associate myself with his company," "would never help him," "will not talk to the person," "will not have financial transactions with the person" and "will not help the person with money in the future". Further, when asked about the case where B is a member of a RoSCA and does not make the stipulated payments, 43% said that fellow group members would withhold future favors and 60% of other villagers who come to hear about this would withhold future favors. Meanwhile 80% of fellow group members and 83.3% of other villagers who come to hear about the infraction will withhold future loans. These responses are similar to another question as to what happens if B defaults on a large bilateral loan from A. Interestingly 23.3% said that they would not tell the wrongdoer about future job opportunities. This number jumps to 56.7% when we ask about B not doing a job properly for A as opposed to being a defaulting RoSCA member. In short, it seems that the reputational consequences are severe and in the future there will likely be some form of social sanctioning: favors, loans and information about jobs may be withheld. When we asked in informal conversations if central or peripheral members had better access to favors or other opportunities in the village (using the same method described above), all villagers we talked to indicated that the central person clearly had better access to resources. Naturally, it would be harder to retaliate against these well-positioned third parties.

Figures

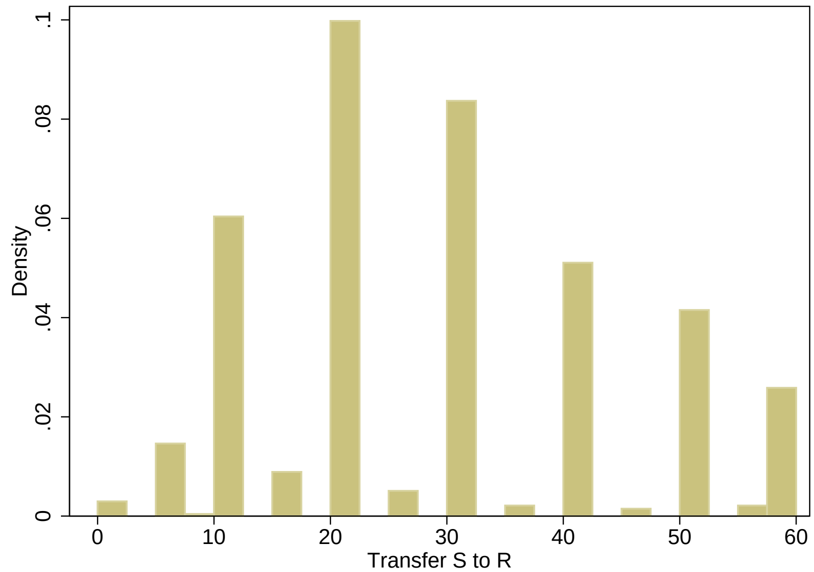

Figure 1. Distribution of transfers from sender to receiver

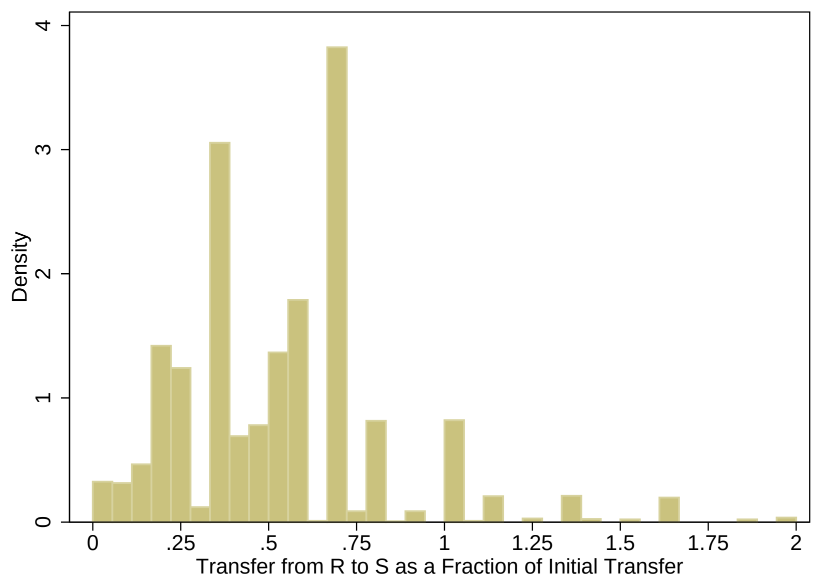

Figure 2. Distribution of transfers from receiver to sender as a fraction of the initial transfer received by the receiver

Panel A: Third-Party Punishers (Information + Enforcement Channel)

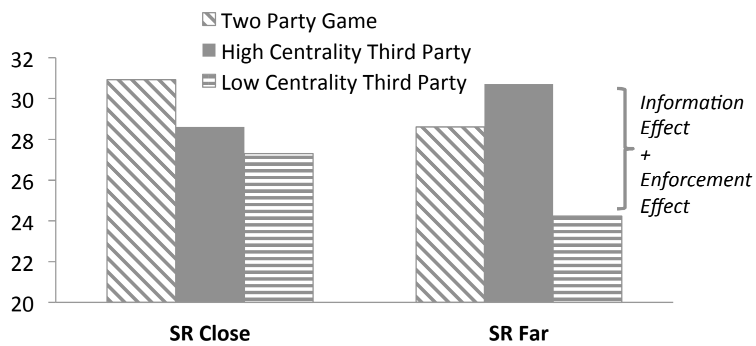

Panel B: Third-Party Monitors (Information Channel)

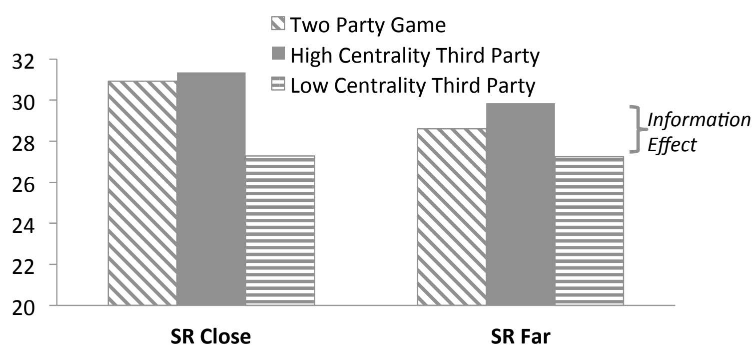

Figure 3. Sender Transfers by Game and Punisher Characteristics.

In all bar charts, the y-axis represents the transfer. In each grouping, the left-most bar shows the transfers in the two-party game, the middle bar shows the transfers in the three-party game with a third party of high centrality, and the right-most bar shows the transfers in the three-party game with a third party of low centrality. SR close denotes S-R pairs of social distance one or two. SR far includes all other S-R pairs. Panel A compares the game with two players to the game with a third-party punisher. Panel B compares the game with two players to the game with a third-party monitor.

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,T)	1136	3.5722	1.1043
Social Distance (R,T)	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,T) level (depending on the game). The sample is restricted to the set of reachable pairs.

Table 2. Sender Behavior

	(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 Omitted 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.

Table 3. 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 4. 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 5. 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 6. When do third parties help?

	(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 B. Supplementary Tables

Table B.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 B.2. Sender's Transfers and Importance Characteristics

	(1)	(3)
	Eigenvector Centrality Measure	
Outcome Variable: Transfer S to R	Quantile	High vs. Low
Third Party Centrality: Monitoring	1.338	2.728**
	(2.688)	(1.370)
Third Party Centrality: Punishment	8.063***	4.943***
	(2.537)	(1.297)
Sender Centrality	-7.194**	-3.318**
	(3.00)	(1.475)
Sender Centrality: Monitoring	7.159	3.442
	(4.472)	(2.376)
Sender Centrality: Punishment	1.787	0.352
	(3.948)	(2.057)
Receiver Centrality	-1.136	-1.800
	(3.107)	(1.38)
Receiver Centrality: Monitoring	4.394	3.520*
	(3.901)	(1.839)
Receiver Centrality: Punishment	-2.350	0.759
	(4.029)	(1.797)
Social Closeness (S,R)	1.601*	1.544*
	(0.848)	(0.785)
Social Closeness (S,R): Monitoring	-1.977**	-1.723**
	(0.900)	(0.850)
Social Closeness (S,R): Punishment	-0.792	-0.944
	(1.116)	(1.061)
Social Closeness (S,J): Monitoring	0.576	0.366
	(0.640)	(0.593)
Social Closeness (S,J): Punishment	-0.574	-0.450
	(0.602)	(0.581)
Social Closeness (R,J): Monitoring	-0.477	-0.740
	(0.649)	(0.622)
Social Closeness (R,J): Punishment	-0.400	-0.527
	(0.573)	(0.581)
T2: Game with Monitoring	1.018	4.543
	(4.675)	(4.442)
T3: Game with Monitoring and Punishment	2.435	4.398
	(5.361)	(5.347)
Controls: Experimental	Yes	Yes
Controls: Demographic	No	No
Observations	1,752	1,752
R-squared	0.243	0.248

Note: Regressions include observations from T1, T2, and T3. In all columns, the outcome variable is the amount transferred from S to R. In column (1) the centrality measure used is the village quantile ranking of eigenvector centrality. In column (2), 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 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 B.3. Total Surplus and Importance Characteristics

	(1)	(2)	(3)	(4)
	Eigenvector Centrality Measure			
Outcome Variable: Surplus	Quantile	Quantile	High vs. Low	High vs. Low
Third Party Centrality: Monitoring	2.851	10.41*	5.321*	6.105**
	(5.378)	(6.132)	(2.738)	(2.993)
Third Party Centrality: Punishment	13.83***	20.68***	8.754***	11.43***
	(5.024)	(6.182)	(2.690)	(3.334)
Third Party Elite: Monitoring		1.924		1.897
		(3.919)		(3.823)
Third Party Elite: Punishment		-1.489		-1.673
		(3.990)		(4.064)
Third Party High Caste: Monitoring		4.363		4.491
		(3.385)		(3.368)
Third Party High Caste: Punishment		-2.721		-2.956
		(3.514)		(3.422)
Third Party Female: Monitoring		0.383		0.381
		(2.923)		(2.798)
Third Party Female: Punishment		-2.904		-3.333
		(3.510)		(3.599)
Tests for Monitoring=Punishment: p-values				
Third Party Centrality	0.0716	0.184	0.2909	0.1632
Third Party Elite		0.5228		0.5043
Third Party High Caste		0.1007		0.086
Third Party Female		0.492		0.4359
Controls: Experimental	Yes	Yes	Yes	Yes
Controls: Demographic	No	Yes	No	Yes
Observations	1,751	1,515	1,751	1,515
R-squared	0.263	0.313	0.268	0.318

Note: Regressions include observations from T1, T2, and T3. In all columns, the outcome variable is the surplus (2 times the amount transferred from S to R minus four times the cost of the inflicted punishment). 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.

Appendix C. 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 set of reachable pairs. 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 C.1, $\gamma(i,j) = 2$ and $\gamma(i,k) = \infty$.

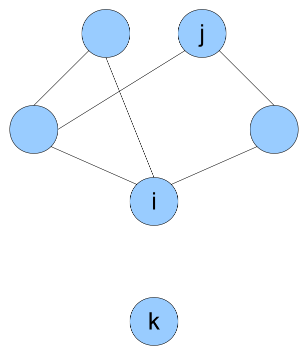

Figure C.1. Path lengths i, j and i, k

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 C.2(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 C.2(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.[44] Intuitively, this measure defines the importance of a node as proportional to the sum over each of its network neighbors'

[44]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.

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 C.2(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.

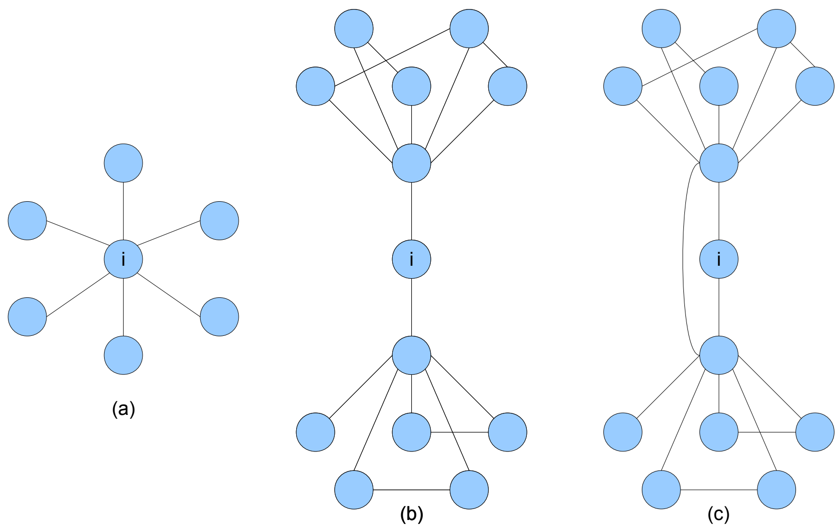

FIGURE C.2. Centrality of node i