ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

*An International Society for the Advancement of Economic Theory in its Relation to Statistics and Mathematics*

http://www.econometricsociety.org/

*Econometrica*, Vol. 87, No. 1 (January, 2019), 175–216

# SOCIAL NETWORKS, REPUTATION, AND COMMITMENT: EVIDENCE FROM A SAVINGS MONITORS EXPERIMENT

EMILY BREZA
*Department of Economics, Harvard University, NBER, and J-PAL*

ARUN G. CHANDRASEKHAR
*Department of Economics, Stanford University, NBER, and J-PAL*

*Econometrica*, Vol. 87, No. 1 (January, 2019), 175–216

# SOCIAL NETWORKS, REPUTATION, AND COMMITMENT: EVIDENCE FROM A SAVINGS MONITORS EXPERIMENT

EMILY BREZA
Department of Economics, Harvard University, NBER, and J-PAL

ARUN G. CHANDRASEKHAR
Department of Economics, Stanford University, NBER, and J-PAL

We conduct an experiment to study whether individuals save more when information about the progress toward their self-set savings goal is shared with another village member (a "monitor"). We develop a reputational framework to explore how a monitor's effectiveness depends on her network position. Savers who care about whether others perceive them as responsible should save more with central monitors, who more widely disseminate information, and proximate monitors, who pass information to individuals with whom the saver interacts frequently. We randomly assign monitors to savers and find that monitors on average increase savings by 36%. Consistent with the framework, more central and proximate monitors lead to larger increases in savings. Moreover, information flows through the network, with 63% of monitors telling others about the saver's progress. Fifteen months after the conclusion of the experiment, other villagers have updated their beliefs about the saver's responsibility in response to the intervention.

KEYWORDS: Commitment, reputation, savings, social networks.

## 1. INTRODUCTION

PEER EFFECTS have been found in a range of settings from schooling to exercise to savings. The literature has traditionally focused on cleanly identifying the reduced form effect, asking how an individual's savings or academic performance depends on the savings or academic performance of her peers. Individuals may be affected by the actions of their peers through a variety of channels. Using the example of group savings, the literature has shown that peer effects operate through channels such as (a) learning how to use financial products; (b) reminders; (c) posting a bond; or (d) reference-dependent preferences ("keeping up with the Joneses") (see Jack and Suri (2014), Cai, de Janvry, and Sadoulet (2015), Bryan, Karlan, and Zinman (2015), Kast and Pomeranz (2018), Beaman, Karlan, and Thuysbaert (2014), Beshears, Choi, Laibson, Madrian, and Milkman (2015), Munshi (2014), Karlan, McConnell, Mullainathan, and Zinman (2016), Bursztyn, Ederer, Ferman, and Yuchtman (2014), and Banerjee, Chandrasekhar, Duflo, and Jackson (2013)).

Emily Breza: ebreza@fas.harvard.edu
Arun G. Chandrasekhar: arungc@stanford.edu
We thank Monisha Mason, Lisa Nestor, Shobha Dundi, and Gowri Nagraj for outstanding assistance in the field. We thank Robert Townsend, Esther Duflo, Abhijit Banerjee, and Matt Jackson for their tremendous guidance. We also thank Ran Abramitzky, Sumit Agarwal, Attila Ambrus, John Beshears, Gabriel Carroll, Pascaline Dupas, Itay Fainmesser, Erica Field, Andy Foster, Ben Golub, Rema Hanna, Andrew Hertzberg, Stephan Meier, Markus Mobius, Kaivan Munshi, Rohini Pande, Luigi Pistaferri, Dina Pomeranz, Jonathan Robinson, Duncan Thomas, Chris Udry, and Xiao Yu Wang for helpful discussions, as well as numerous seminar and conference participants. We are grateful to CFSP, NSF SES-1156182, the Chazen Institute at Columbia Business School, CIBER and the Earth Clinic at Columbia University, the Russell Sage Behavioral Economics Small Grant, and SEED for generous support. Finally, we both thank the NSF GRFP. A previous version of this paper was titled "Savings Monitors."

 https://doi.org/10.3982/ECTA13683

However, less has been written on whether peer effects may arise from individuals wanting to impress others through their actions.

This paper focuses on this last channel that is likely present in many applications—that when actions are observable to others, they may come with reputational benefits. Further, those benefits may depend on the network position of the observer given that building reputation may be more valuable with some members of the community than others.[1]

The potential for reputation-based peer effects is particularly widespread in development economics. For example, in theorizing about repayment incentives in joint liability microcredit, Besley and Coate (1995) described a social punishment that depends on reporting poor behavior and admonishment by others, writing:

> "the contributing member may admonish his partner for causing him or her discomfort and material loss. He might also report this behavior to others in the village, thus augmenting the admonishment felt. Such behavior is typical of the close-knit communities in some LDCs."

Peer-driven financial institutions, such as rotating savings and credit associations (RoSCAs), self-help groups (SHGs), and village savings and loan associations (VSLAs), are ubiquitous in the developing world and are thought to, in part, work on this principle. However, it is hard to get traction on how these institutions work, let alone isolate the reputational channel: they are complicated objects of anywhere from five to 30 individuals, with endogenous group formation and forces beyond simple reputation effects contributing to good behavior.[2]

We design and implement a savings field experiment to focus on the reputational force highlighted by Besley and Coate (1995) and begin to unpack the black box of peer effects in informal financial institutions. To make the problem tractable, we construct a simplified "institution" of one saver who desires to save matched to one observer and induce random group formation. Importantly, our setting is naturalistic, mimicking the business correspondent (BC) model, which is commonly used by banks in India to service rural customers.

Specifically, we conduct an experiment across 60 villages in rural Karnataka, India, where we have complete network data for almost all households in every village. We assist 1,300 individuals to review their finances, set a six-month savings goal, and open a formal account at a bank or post office. A random group of savers is selected from each village, and each saver receives a (different) partner for the duration of the experiment, whom we call a monitor. In 30 randomly-selected villages, we randomly assign individuals from a pre-specified pool to serve as monitors, and in the remaining 30 villages, using random serial dictatorship, we allow savers to select their respective monitors from the pre-specified monitor pool. In all cases, the monitor receives bi-weekly information about the saver's target account savings. As monitors are drawn from a random pool of villagers,

[1] Our experiment was inspired by the earlier lab-in-the-field study Breza, Chandrasekhar, and Larreguy (2015), where we explored the efficiency of transfers in non-anonymous sender–receiver investment games with a third-party observer. Villagers were assigned to one of three treatments: (1) sender–receiver game, (2) sender–receiver game with a third party who observes the interaction, but takes no action of her own, or (3) sender–receiver game with a third party who observes the interaction and can levy a fine against the receiver. The interaction is fully non-anonymous. We were interested in how the network position of the third party influences the efficiency of the transaction. When a more central third-party observes the transaction, efficiency increases significantly (as seen from comparing (2) to (1) for more versus less central third parties). Further, the beneficial effect of centrality is greater when the third party is also given an observable punishment technology.

[2] Reputation effects may help to explain, for example, why researchers have documented peer effects in microfinance groups even in the absence of contractual joint liability (Breza (2014), Feigenberg, Field, and Pande (2013)).

they vary in their position in the village network: some are more central (i.e., more connected directly or indirectly) than others, and some have closer relationships (i.e., proximity through the network) with the saver. Using the 30 villages in which we randomly assign saver–monitor pairs, we study how the network position of randomly-assigned monitors influences savings behavior. Further, we use the 30 villages in which savers choose their monitors to benchmark how much agents save under endogenous group formation.

Why might the monitor's position in the network be important? Each monitor learns about the saver's progress. The monitor may, in turn, pass that information or any opinion she has made on to others. Thus, the monitor's position within the village network may determine how far and to whom her opinion may spread. For example, more central agents—that is, better connected (directly or indirectly) to a larger set of people—are well-suited to broadcast information. In turn, they may make more effective monitors, ceteris paribus, as the saver has more to gain by impressing them. Similarly, a socially proximate monitor may be more likely to speak to others with whom the saver is likely to interact in the future. Therefore, by telling individuals who are more relevant to the saver's future interactions, proximate monitors may also be more valuable.

To help clarify these issues and identify those aspects of the otherwise-complex network on which to focus our empirical analysis, we develop a simple signaling model. In this model, we assume that savers gain utility from interacting in the future with individuals who have heard about their successes.[3] Here, the network plays two roles: information is disseminated from the monitor through the network, and future interactions between the saver and other villagers (including the monitor) occur through the network. We show that a saver is incentivized to save more when randomly assigned to a more central monitor or to one that is more proximate to her.

Equipped with this framework, we pair our experiment with detailed network data collected in part by the authors in previous work (Banerjee, Chandrasekhar, Duflo, and Jackson (forthcoming)). These household-level network data comprise 12 dimensions of interactions across *all* potential pairs of households in each of the 60 study villages.[4] Two moments of the network data emerge from our model and we focus on both: monitor (eigenvector) centrality, which captures how much information emanating from a monitor should spread in the network, and the social proximity between the saver–monitor pair, which is the inverse of the shortest path length through the network.[5] The framework also generates a model-specific network statistic that drives savings incentives, which we take to the data. This model-specific statistic is increasing in both monitor centrality and saver–monitor proximity. We do not claim that our experiment shows that this mechanism is the sole force driving why monitor centrality and proximity should affect savings, but we do provide evidence consistent with such a story playing a role.

Savings is an ideal application for our experiment for several reasons. First, we require a setting where reputation is important. Anecdotal and survey evidence from the study villages suggests that a large fraction of villagers indeed want to save more, and that showing one can save more is a sign of responsibility. Furthermore, in economic models, caring

[3]There are many microfoundations for such an effect. Successful savers may gain an improved reputation for being responsible, for example. Alternately, agents may feel embarrassment or shame when interacting with people who have learned of their shortcomings.

[4]The network data we use here are very detailed. With data surveyed from 89% of households in each village, the probability of not surveying either member of a pair is $0.09^2$. So there are data on $1-0.09^2 \approx 0.988$ share of possible links in the network. We use what is called the OR network, drawing a link between two households if either named the other.

[5]See Katz and Lazarsfeld (1970), Ballester, Calvó-Armengol, and Zenou (2006), Banerjee et al. (2013), and Golub and Jackson (2010).

about the future is often what makes someone trustworthy—savings is a strong signal of precisely this. Second, we want to be able to accurately measure the outcome variable. Certainly savings in a bank account is easy to observe and we can verify the data through passbooks. Third, we desire a context that is naturalistic. Savings is an obvious application in which to study public commitments, as many of the informal financial products commonly observed in developing countries (and in the study villages) and discussed above incorporate groups of individuals from the same social network and rely on mechanisms that are likely to include mutual monitoring/observation (Besley and Coate (1995), Beaman, Karlan, and Thuysbaert (2014), Besley, Coate, and Loury (1993), Karlan (2007), Giné and Karlan (2014), Bryan, Karlan, and Zinman (2015), and Breza (2014)).[6,7]

Finally, chronic under-saving is an important issue in developing and developed countries alike. The desire to save is widespread, but many are unable due to lack of access, lack of commitment, or lack of attention (Ashraf, Karlan, and Yin (2006), Brune, Giné, Goldberg, and Yang (2016), Karlan et al. (2016), Thaler and Benartzi (2004), and Beshears, Choi, Laibson, Madrian, and Sakong (2011), for example). Our intervention can be interpreted as a special kind of commitment savings device where the characteristics of the monitor determine how well it performs. Research has also shown that increased savings has numerous benefits including increased investment, working capital, income, and even labor supply, and can improve the ability for households to overcome shocks (Dupas and Robinson (2013b, 2013a), Prina (2015), Schaner (2018), and Kast and Pomeranz (2018)). We can explore these issues in the short run (6 months) and medium run (21 months) through the lens of our study.

Our empirical analysis has four components. First, using the data from the 30 villages in which we randomly assign monitors, we establish that receiving an arbitrary monitor increases total savings across formal and informal vehicles by 36%. As predicted by our model, the largest increases are generated by more central monitors as well as more proximate monitors. Increases of one standard deviation in monitor centrality and proximity, respectively, correspond to increases in savings of 14% and 16%. Similarly, a one standard deviation increase in the model-specific network statistic corresponds to an increase in savings of 33.5%.

Second, we make use of novel supplemental data to support the reputational story. We show that monitors indeed speak to others about the saver, and 40% of savers even hear gossip about themselves through back-channels. Moreover, 15 months after the conclusion of the intervention, the opinions of randomly-selected households about a saver's performance and ability to follow through on self-set goals are related to the centrality of

[6]Our paper is related to Kast and Pomeranz (2018), who conducted an experiment layering a peer savings scheme on top of an existing microfinance borrowing group. Members were motivated to save by making public commitments in front of one another along with public contributions. The authors found large positive effects on savings. In a second experiment, SMS-based reminders to save led to similar savings impacts. A distinction with our setting is that in their's, the monitors were both co-borrowers with the savers and were savers themselves.

[7]The paper is also related to a larger dialogue about the role of social capital and observability in informal financial institutions (Platteau (2000)). Jakiela and Ozier (2016) provided lab-in-the-field evidence that peer observability can have a dark side, causing some women to inefficiently distort their behavior to avoid paying a so-called "kin tax." Brune, Giné, Goldberg, and Yang (2016) and Dupas and Robinson (2013a) argued that social pressure may be one mechanism underlying the results of their respective savings field experiments as well. One natural feature of opting into any group-driven financial product, including our experimental sample, is that, by definition, some component of savings is observable. Therefore, it is likely that individuals who worry the most about unwanted demands from others would be the least likely to participate.

that saver's randomly-assigned monitor. To our knowledge, this is the first time such an exercise has been conducted in the literature.

Third, we provide evidence that our intervention caused lasting and positive average impacts on participant households. We show that the increases in savings caused by our intervention come from increases in labor supply and decreases in unnecessary expenditures. Fifteen months after the end of the intervention, we show that subjects randomly assigned to monitors report declines in the incidence of unmitigated shocks.[8] Moreover, the increases in savings persist 15 months after the intervention. Taken together, these results suggest that monitors, especially central and proximate ones, help savers to direct financial slack toward savings, rather than wasteful expenditures or leisure, result in improved risk-coping, and yield persistent increases in savings, likely held as buffer stocks.[9]

Fourth, in the 30 villages in which savers could choose their monitors, we find that monitored savers perform approximately as well as their random assignment village counterparts. Further, we find that non-monitored savers in the endogenous assignment villages save substantially more than the non-monitored savers in the random assignment villages. Non-monitored savers in endogenous selection villages completely "catch up" to the saving levels of the monitored community members. This suggests that monitoring can affect the saver's propensity to save, but may also spill over to friends of the saver. Indeed, we find suggestive evidence that more unplanned conversations about savings take place in endogenous selection villages, many of which were likely overheard by non-monitored savers.

In short, our study points to the idea that reputations matter, and they matter heterogeneously within the broader village network. The experiment provides a context in which agents could respond to our monitor treatment using an important economic vehicle—formal savings—that itself stood to generate real benefits to our subjects. That the increased savings allowed our subjects to better respond to health and household shocks indicates that the monitor treatment effect was strong enough not only to change savings behavior directly but to also yield measurable and meaningful economic consequences over the next year. Again, we caution that reputation is of course not the only force that could matter for our findings, but our analysis points to it being a relevant channel.

The remainder of the paper is organized as follows. Section 2 contains a description of the experimental design, setting, and data. In Section 3, we provide a parsimonious model that shows why it is natural to focus on centrality and proximity. Section 4 presents the results for the villages where monitors were randomly assigned to savers, and Section 5 presents a discussion of threats to validity. We discuss savings balances in the villages with endogenous monitor assignment in Section 6, and Section 7 is a conclusion. Robustness exercises and extensions can be found in the Supplemental Material (Breza and Chandrasekhar (2019)) and the Auxiliary Online Appendix (available on the authors' webpages at https://stanford.edu/∼arungc/BC_aux.pdf).

## 2. DATA AND EXPERIMENTAL DESIGN

The requirements for our study are threefold: (1) detailed social network data; (2) the presence of financial institutions where study participants can open accounts and save;

[8]We denote a shock to be an event such as a personal health shock, bovine health shock, or other unexpected household expenditure where the household did not have enough cash on hand to cover the cost. We show that the incidence of reporting an above-median number of shocks drops for individuals assigned to a monitor.

[9]The most common savings goal purposes listed at baseline were unforeseen expenditures and emergencies.

(3) experimental variation in the nature of the saver–monitor relationship. Our final sample includes approximately 3,000 participants from 60 villages in rural Karnataka, India.

### 2.1. *Network and Demographic Data*

We chose to set our experiment in villages that coincide with the social network and demographic data set previously collected, in part by the authors (and also described in Banerjee et al. (forthcoming)). In our field experiment, we match participants to this unique data set.

Banerjee et al. (forthcoming) collected network data from 89.14% of the 16,476 households living in those villages. The data concern "12 types of interactions for a given survey respondent: (1) whose houses he or she visits, (2) who visit his or her house, (3) his or her relatives in the village, (4) non-relatives who socialize with him or her, (5) who gives him or her medical advice, (6) from whom he or she borrows money, (7) to whom he or she lends money, (8) from whom he or she borrows material goods (e.g., kerosene, rice), (9) to whom he or she lends material goods, (10) from whom he or she gets advice, (11) to whom he or she gives advice, (12) with whom he or she goes to pray (e.g., at a temple, church, or mosque)." This provides a rich description of the interactions across households. We construct one network for each village at the household level. A link exists between households if any member of a household is linked to any other member of another household in at least one of the 12 ways. Network-level summary statistics are displayed in Auxiliary Appendix Table K.1.

As such, we have extremely detailed data on social linkages, not only between our experimental participants but also about the embedding of the individuals in the social fabric at large. Every village is associated with a social network. Following the extensive work on this data, we assume that this is an undirected, unweighted network (see, e.g., Banerjee et al. (2013), Jackson, Barraquer, and Tan (2012), Chandrasekhar, Kinnan, and Larreguy (2018) for discussion).

Moreover, the survey data include information about caste, elite status, house amenities, and the GPS coordinates of respondents' homes. In the local context, 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. All our analyses include measures of network effects conditional on these numerous observables.

Given the richness and uniqueness of the data collected by Banerjee et al. (2013) (used to analyze the diffusion of microfinance), other projects have been conducted in subsets of the 75 villages, namely, three half-day lab-in-the-field experiments (Chandrasekhar, Kinnan, and Larreguy (2018), Breza, Chandrasekhar, and Larreguy (2015), and Chandrasekhar, Larreguy, and Xandri (2012)). Auxiliary Appendix M describes these other studies, discusses treatment balance with respect to them, and also shows robustness to saver-level controls for prior participation. The results do not change with the addition of such controls, which is to be expected, as these were several-hour lab-in-the-field experiments that should not interact with a financial inclusion program years later.

### 2.2. *Bank and Post Office Accounts*

In addition to the social network data, a key requirement of our study is convenient access to bank and post office branches for all participants. In each village, we identify one bank branch and the local post office branch to offer as choices to the savers. We select bank branches that satisfy several criteria: are located within 5 km of the village, offer "no-frills" savings accounts,[10] and agree to expedite our savings applications and process them in bulk.[11] Out of the 75 villages surveyed by Banerjee et al. (forthcoming), 60 villages satisfy these criteria and constitute our final sample.

Each village in India is well-served by the postal branch network. Branches are generally within a 3km walk of each village in our sample. Thirty-five percent of our study villages have a post office branch within the village boundary. We offer the post office choice because women often feel uncomfortable traveling to bank branches but feel more comfortable transacting with the local post master. Some individuals prefer bank accounts because those accounts make it easier to obtain bank credit in the future. Both the bank and postal savings accounts have very similar product characteristics. In each type of account, the minimum balance is typically Rs. 100 ($\sim$\$2), and there are no account maintenance or withdrawal fees. The interest rates on the bank accounts range from 3% to 4.5%, and the interest rate on the postal accounts is 4%, making real interest rates negative, consistent with a story of precautionary savings being a primary motive here.[12] Users of both types of accounts are given a passbook, which is an official document containing the account information, the name, address, and photo of the account holder, and the record of all account activity (deposits and withdrawals). Entries in the passbooks are stamped with an official seal by branch personnel and cannot be forged. "No-frills" bank and post office accounts have no formal commitment features. Participants in the study are allowed to withdraw freely from the accounts at any time. The only source of commitment comes from the presence of a monitor, described below. Our monitor treatments introduce an "informal commitment" in the parlance of Bryan, Karlan, and Nelson (2010).

### 2.3. *Experimental Design*

Figures 1 and 2 represent our experimental design and Figure 3 presents a time line. Study participants are randomly selected from an existing village census database, collected in conjunction with the network survey, and then randomly assigned to be part of our saver group, monitor group, or pure control (Figure 1(B)).

All potential treatment savers and monitors who are interested in participating (Figure 1(C)) are administered a short baseline survey, which includes basic questions on account ownership, income sources, and desire to save. Only a quarter of households had a bank or post office account at baseline. We aim to test whether monitors can increase savings balances and also increase the use of already-accessible interest-bearing bank savings accounts.

Next, potential savers establish a six-month savings plan. Importantly, this plan is established before the saver knows whether she is assigned to the non-monitored treatment or one of the monitored treatments. Moreover, the saver does not know whether the village is assigned to endogenous monitor selection or random monitor selection. The saver does know about the arms of the experiment: that the village may have random or endogenous choice of monitor, and it may be the case that irrespective of that, they may randomly be assigned to not receive a monitor. They make their plans and participation decisions knowing this.

The process of setting a savings goal includes listing all expected income sources and expenses month by month for six months. Savers are prompted to make their savings

[10]"No-frills" accounts generally have no minimum balance, charge no user fees, and require a minimal initial deposit, which is generally around Rs. 100 (\$2).

[11]The 5 km distance restriction meant that we were not able to work with only one bank, and instead opened accounts at branches of six different banks.

[12]Inflation rates are as follows: 2010: 9.47%, 2011: 6.49%, 2012: 11.17%, 2013: 9.13% (Inflation.eu, 2017).

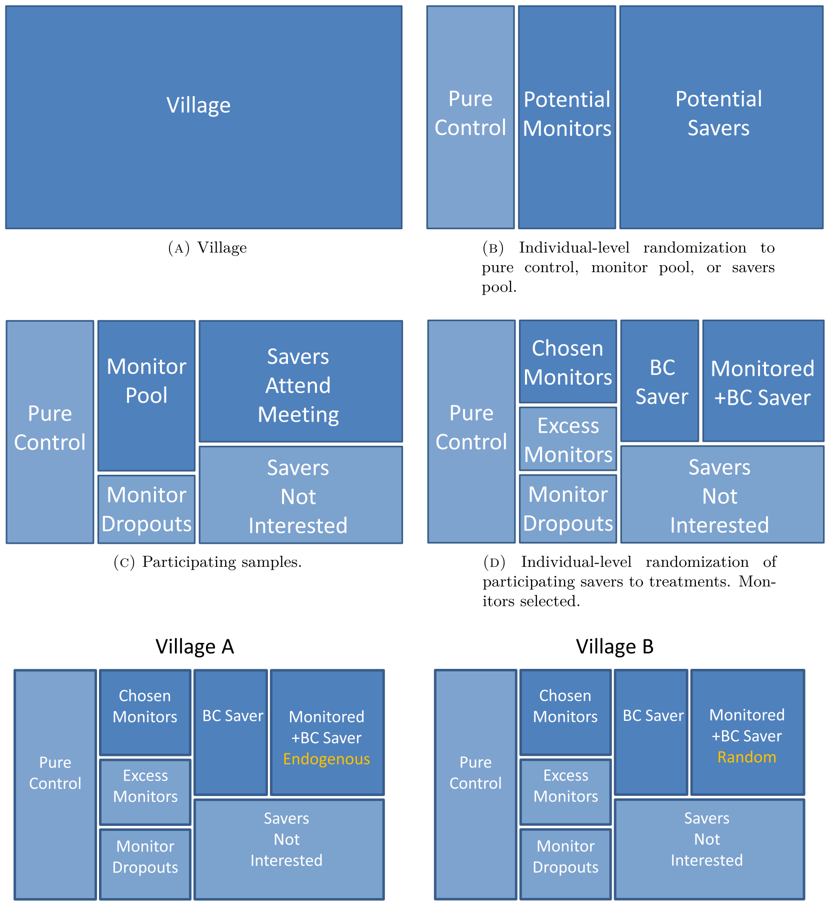


(A) Village

(B) Individual-level randomization to pure control, monitor pool, or savers pool.

(C) Participating samples.

(D) Individual-level randomization of participating savers to treatments. Monitors selected.

(E) Village-level randomization. Village A is randomly assigned to endogenous monitoring treatment. Village B is randomly assigned to exogenous monitoring treatment.

FIGURE 1.—Experimental design and randomization. Notes: "BC Saver" refers to our non-monitored treatment (T1) described in Section 2.

goals concrete, and we record the desired uses of the savings at the end of the six-month period. Individuals are then invited to a village-level meeting in which study participation is finalized and treatment assignments are made. Potential monitors are also invited to attend the village meeting and are told that if selected, they can earn a small participation fee and incentive payment for participating.

Our sample frame for randomization is the 57% of savers who self-select into attending the village meeting (see Figure 2). We use two different data sources in Supplemental Material Tables C.1 and C.2 to explore correlates with participation. In Supplemental Material Table C.1, we use our responses to the short baseline survey to compare the participants with non-participants.[13] In Supplemental Material Table C.2, we compare the participating households with the full set of non-participants using the village census data

[13]If, during the initial visit, the potential savers tell the enumerators that they are not interested, then the baseline survey is not completed and they are not invited to the meeting. However, they are included in Figure 2.

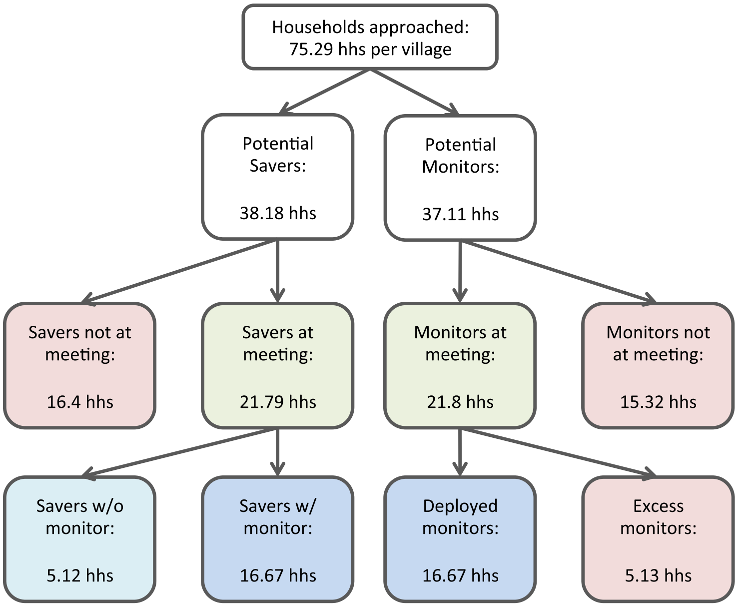


FIGURE 2.—Sample description. Notes: This figure shows the average number of households per village in our sample in each cell. There are 60 villages in the sample, with an average population of 222.12 households. Every village has an average of 146.83 pure control households that were not approached at all before or during the savings intervention.

collected by Banerjee et al. (forthcoming). Both tables show that participants disproportionately come from poorer households with a desire to save. Landless laborers are more likely, while salaried government workers are less likely to select into the sample. Moreover, the stated saving goals of participants are 8% smaller in size. This is consistent with poorer individuals having a harder time meeting their savings goals on their own. We also observe that households that actively save at a regular frequency and in which at least one member has a formal account are more likely to participate. Finally, individuals with exposure to RoSCAs and SHGs are almost 10 percentage points more likely to participate. This is a nice feature because these are the types of people who are prone to participate in social finance. If, after hearing about the potential monitoring treatments, there are villagers concerned about rotten-kin type of forces (Jakiela and Ozier (2016)), they would be likely to self-select out of our study as well as SHGs, which also render one's savings visible to a group.

From the pool of consenting participants and attendees of the village meeting, we randomly assign savers to one of three treatments (see Figure 1(E)):

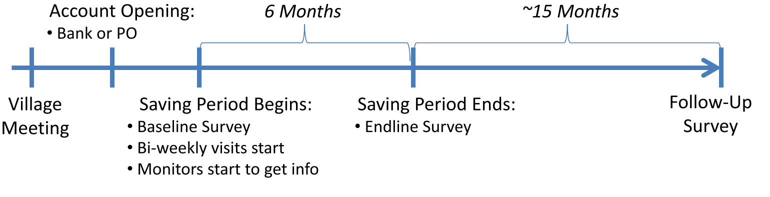


FIGURE 3.—Time line of experiment.

T1. Non-monitored treatment: Randomization at the individual level. The individual-level treatment resembles a financial institution already in use in India called *business correspondents*. Agents of the bank travel to villages to provide direct in-home customer service. This includes account opening procedures and deposit-taking. However, we were not legally able to collect deposits ourselves as researchers.

T2. Monitoring with random matching: Villages randomly assigned to have random assignment of monitors, individuals within a village randomly assigned to whether or not they receive monitors, and savers randomly matched to monitors.

T3. Monitoring with endogenous matching: Villages randomly assigned to have choice of monitors, individuals within a village randomly assigned to whether or not they receive monitors, and savers choose monitors by random serial dictatorship.

All individuals who attend the village meeting are assisted in account opening by our survey team. Savers are allowed to choose to open a bank or a post office account or to use an existing account, if applicable. We help savers to assemble all of the necessary paperwork and "know your client" (KYC) identification documents for account opening and submit the applications in bulk. The savings period begins when all of the savers have received their new bank or post office account passbooks.

All savers in the individual treatment (T1, T2, and T3) are visited on a fortnightly basis. Our surveyors check the post office or bank passbooks and record balances and any transactions made in the previous 14 days and also remind savers of their goals.[14] This process gives us a reliable measure of savings in the target account on a regular basis. These home visits also serve as strong reminders to save. We should note that in no treatment do our surveyors collect deposits on behalf of the savers.[15]

In our peer treatment with random matching (T2), we randomize the assignment of monitors to savers. In each village, a surplus of monitors turned up to the village meeting, so there were more than enough monitors for each T2 (or T3) saver. Every two weeks, after surveyors visit the T2 savers, they then visit the homes of the monitors. During these visits, the monitors are shown the savings balances and transaction records of their savers and are also reminded of each saver's goal. Thus, our intervention intermediates information between the saver and monitor. At the end of the savings period, monitors receive incentives based on the success of their savers. Monitors are paid Rs. 50 if the saver reaches at least half the goal, and an additional Rs. 150 if the monitor reaches the full goal.[16]

The peer treatment with endogenous matching (T3) is identical to T2, except for the method of assigning monitors. In this treatment, individuals are allowed to choose their monitor from the pool of all potential monitors attending the meeting. We only allowed one saver per monitor, so we randomized the order in which savers could choose. There was excess supply of monitors, so even the last saver in line had many choices. It is important to note here that the pool of potential monitors is recruited in an identical fashion in both sub-treatment groups (2) and (3). Table I presents summary statistics for the sample that attended the village meeting and also shows baseline differences between T1, T2, and T3.

[14]We were not able to obtain administrative data from the banks and post offices due to the large number of different institutions (Post office + branches of six different banks).

[15]This is one important difference between our product and the typical business correspondent model.

[16]Monitors also receive a participation fee at the time of the village meeting. Monitors that are ultimately selected receive Rs. 50 and those who are not ultimately selected receive Rs. 20. We had initially wanted to vary experimentally the size of the monitor incentives, but the required sample size was not feasible given our budget and the number of villages with both network data and a nearby bank branch willing to expedite our account opening. We investigate whether these incentives could be driving our results in Section 5.

TABLE I

SUMMARY STATISTICS, TREATMENT ASSIGNMENT, AND ATTRITION[a]

| | Treatment (Village Meeting Sample) | | | | Treatment (Endline Sample) | | | |
|---|---|---|---|---|---|---|---|---|
| | (1) Mean of Non-Monitored Savers | (2) Diff. Random vs. No Monitor | (3) Diff. Endogenous vs. No Monitor | (4) Diff. Endogenous vs. Random | (5) Mean of Non-Monitored Savers | (6) Diff. Random vs. No Monitor | (7) Diff. Endogenous vs. No Monitor | (8) Diff. Endogenous vs. Random |
| Age | 33.094 | −0.147 | 0.158 | 0.306 | 33.454 | 0.041 | 0.025 | −0.016 |
| | (0.385) | (0.458) | (0.528) | (0.458) | (0.387) | (0.454) | (0.551) | (0.468) |
| Female | 0.756 | −0.041 | −0.025 | 0.016 | 0.780 | −0.041 | −0.025 | 0.016 |
| | (0.024) | (0.032) | (0.034) | (0.031) | (0.026) | (0.030) | (0.034) | (0.029) |
| Married | 0.857 | −0.029 | −0.024 | 0.004 | 0.875 | −0.033 | −0.041 | −0.008 |
| | (0.019) | (0.021) | (0.027) | (0.025) | (0.020) | (0.022) | (0.027) | (0.025) |
| Widowed | 0.036 | 0.010 | 0.015 | 0.006 | 0.033 | 0.017 | 0.025 | 0.007 |
| | (0.010) | (0.013) | (0.016) | (0.016) | (0.010) | (0.014) | (0.017) | (0.017) |
| Positive Savings 6 mos Prior to Baseline | 0.717 | 0.024 | 0.012 | −0.013 | 0.725 | 0.018 | 0.016 | −0.002 |
| | (0.032) | (0.035) | (0.036) | (0.042) | (0.033) | (0.037) | (0.037) | (0.043) |
| Has Post Office or Bank Acct. at Baseline | 0.378 | −0.011 | 0.040 | 0.052 | 0.396 | −0.010 | 0.024 | 0.034 |
| | (0.032) | (0.036) | (0.034) | (0.040) | (0.033) | (0.037) | (0.035) | (0.041) |
| Has BPL Card | 0.840 | 0.020 | −0.002 | −0.021 | 0.842 | 0.015 | 0.011 | −0.004 |
| | (0.021) | (0.025) | (0.027) | (0.029) | (0.024) | (0.030) | (0.028) | (0.029) |
| Savings Goal | 1,837.622 | −239.078 | 131.062 | 370.139 | 1751.282 | −207.686 | 185.768 | 393.454 |
| | (117.135) | (117.395) | (165.180) | (178.183) | (126.558) | (117.814) | (166.190) | (180.173) |
| Savings Goal (1% outliers trimmed) | 1,649.503 | −106.480 | −55.071 | 51.410 | 1,537.687 | −35.748 | 34.536 | 70.284 |
| | (76.036) | (78.989) | (100.964) | (94.230) | (75.687) | (81.396) | (103.380) | (96.577) |
| Log Savings Goal | 7.253 | −0.063 | 0.009 | 0.072 | 7.209 | −0.035 | 0.047 | 0.082 |
| | (0.040) | (0.041) | (0.046) | (0.053) | (0.042) | (0.041) | (0.048) | (0.054) |
| Saver Centrality | 1.700 | 0.151 | 0.255 | 0.104 | 1.709 | 0.168 | 0.237 | 0.069 |
| | (0.058) | (0.076) | (0.072) | (0.094) | (0.055) | (0.081) | (0.069) | (0.099) |
| Saver Centrality (1% outliers trimmed) | 1.690 | 0.113 | 0.229 | 0.116 | 1.709 | 0.112 | 0.214 | 0.102 |
| | (0.056) | (0.071) | (0.071) | (0.088) | (0.055) | (0.077) | (0.066) | (0.093) |

(*Continues*)

TABLE I—*Continued*

| | Treatment (Endline Sample) | | | |
|---|---|---|---|---|
| | (1) Mean of Non-Monitored Savers | (2) Diff. Random vs. No Monitor | (3) Diff. Endogenous vs. No Monitor | (4) Diff. Endogenous vs. Random |
| Endline Survey Administered (Non-Attriters): No Fixed Effects | 0.889 (0.018) | −0.027 (0.025) | −0.004 (0.022) | 0.023 (0.028) |
| Endline Survey Administered (Non-Attriters): Village Fixed Effects | 0.887 (0.015) | −0.004 (0.029) | −0.025 (0.025) | −0.021 (0.039) |
| 15-Month Follow-Up Survey Administered (Non-Attriters): No Fixed Effects | 0.893 (0.018) | −0.000 (0.024) | 0.035 (0.021) | 0.036 (0.025) |
| 15-Month Follow-Up Survey Administered (Non-Attriters): Village Fixed Effects | 0.896 (0.014) | 0.014 (0.029) | 0.010 (0.023) | −0.004 (0.037) |

[a]Table displays summary statistics and treatment balance of baseline characteristics. Columns (1)–(4) consider the full sample of savers who opted into the village meeting, while columns (5)–(8) consider the sample for which we have endline survey responses. Columns (1) and (5) present means (and standard deviations) for the non-monitored savers. Columns (2) and (6) present the differences between random monitor treatment savers with the non-monitored savers. Columns (3) and (7) compare the endogenous monitor treatment with non-monitored savers. Columns (4) and (8) compare monitored savers under random versus endogenous assignment. The village meeting sample (columns (1)–(4)) includes 1,307 observations, except when 1% outliers are trimmed. With 1% outliers trimmed, the sample is 1,286 observations for Savings Goal and 1,294 for Saver Centrality. The endline sample (columns (5)–(8)) includes 1,146 observations, except when 1% outliers are trimmed. With 1% outliers trimmed, the endline sample is 1,127 observations for Savings Goal and 1,136 for Saver Centrality. Standard errors from regressions of baseline characteristic on treatment are reported in parentheses. Standard errors are clustered at the village level.

For the most part, our sample is balanced across receiving a monitor or not and being in a random choice or exogenous choice village. Unfortunately, by chance, there is a small imbalance in saver centrality, which is driven by outliers in the random villages but persistent in the endogenous villages.[17] Note that the vast majority of the paper is concerned with the random choice villages, so this is less of a concern, and further, everywhere where controls are included, we control for saver centrality. We also unfortunately have a small imbalance in the size of the savings goals across treatments. However, Table I shows that this imbalance is caused by a few large outliers. Auxiliary Appendix N shows that all of our conclusions are unchanged if we drop the outliers and focus on the balanced sample of savings goals.

For our endogenous matching treatment, we chose to implement random serial dictatorship (RSD). Savers were ordered at random and were able to then select their monitors. This was a natural choice for several reasons. First, this mechanism is easy to implement in practice and therefore policy relevant. It is easy to explain to villagers and, owing to its randomness, it seems to be equitable. There was no resistance whatsoever to implementing such a scheme. Second, this design is easier to analyze given the randomization of the choice order. Additionally, it allows us to systematically explore which network aspects are valued when an individual selects a monitor. Third, there is an equivalence between RSD and various other matching schemes with trading which reach the core.[18]

At the end of the 6-month savings period, we administer an endline to all savers and monitors. We record the ending balance in the target accounts from the saver's passbook. We also collect complete savings information across all savings vehicles (including other formal accounts, other informal institutions, "under the mattress," etc.) to make sure that any results are not just coming from the composition of savings. Importantly, we also administer this endline survey to all available attriters or dropouts.[19] Approximately 16% of savers dropped out of our experiment at some point after the village meeting, many of which never opened a target account for the savings period. We were able to survey approximately 70% of the dropouts in our endline follow-up survey and obtain information about their ending savings balances and other key outcomes. Table I also shows differences in the final sampled population decomposed between T1, T2, and T3. We find no differential attrition across the sample of savers captured in our endline data.[20]

Finally, we administer a follow-up survey 15 months after the end of the savings period to the set of savers attending the village meeting. In addition to questions on savings balances, the survey contains retrospective questions about how the savers saved and how frequently they spoke with their monitors. It also contains questions in the style of Dupas

[17]As is well-known, the centrality distribution has a long right tail.

[18]Consider two allocation mechanisms. Say each agent has strict preferences over the monitors. The first mechanism is RSD. The second is when the monitors are randomly allocated to the various agents and then trading is allowed. In this exchange economy, there is a unique allocation in the core and it can be attained by a top trading cycle (TTC) algorithm. Results in Abdulkadiroğlu and Sönmez (2003), Carroll (2012), and Pathak and Sethuraman (2011) show that various versions of RSD and TTC are equivalent: the mechanisms give rise to the same *probability distribution over allocations* irrespective of the preferences.

[19]We also surveyed a random subset of those who were chosen to be savers but who were not interested in savings, and a random subset of the pure control group.

[20]Another reason for why there is no differential attrition, in addition to the high rate of endline survey participation, is the nature of the attrition itself. One common reason for dropping out is a lack of the "know your client" (KYC) legal documentation required for opening a bank account (20% of dropouts). The most frequent reason for dropping out is dis-interest in saving. Further, the composition of why savers drop out is virtually the same (and statistically indistinguishable) for monitored and unmonitored savers.

and Robinson (2013b) about unmitigated shocks sustained in the 15 months after the savings period ended. We ask respondents about transfers made with their monitors after the end of the savings period and also friendships made through the monitor. The follow-up survey also contains questions about each respondent's beliefs about the savings behavior and level of responsibility of 12 other arbitrarily-selected savers. Auxiliary Appendix Table K.2 includes control group means for variables from both of our endline surveys used in the analysis.

## 3. FRAMEWORK

In this section, we present the framework, which guides our empirical analysis. The details of the formal model are presented in Appendix A.

### 3.1. *Motivation*

Clearly, one could tell several stories about why our experiment affects information flow and savings. Underlying the framework is our hypothesis that people expect that they are more likely to be perceived as responsible if they save more in the experiment. Future benefits in the village, such as access to jobs or informal loans or other leadership opportunities, accompany such positive perceptions. In 2016, to provide support for our framework, we conducted a survey of 128 randomly-selected subjects across eight villages in our study area that had never participated in this nor any prior study conducted by either of the authors. We shared with them a vignette of our experiment in order to gather perceptions of villagers not exposed to our experiment. We described our savings monitors study and asked subjects about whether information about savings progress would spread, how that might depend on which monitor was assigned, whether savers would in turn save different amounts depending on who the monitor was, and whether this might lead to returns in terms of favors or hearing about job opportunities in the future. Figure 4 presents the results.

First, subjects believe that people will talk. Seventy-three percent of the respondents say that the monitor would spread information about poor savings of the saver and 93% of respondents say that the monitor would spread information about successful savings by the saver.

Second, subjects perceive that savers would perform better with a more central monitor. To operationalize this, we used the technique from Banerjee et al. (forthcoming) to elicit names of central individuals in the village from a separate sample of villagers without collecting detailed network data. We provided each respondent with the name of one randomly-chosen villager and the name of one of the central villagers. Importantly, we did not explain to the villagers how we obtained those names or comment on the centrality of the named individuals. We then asked respondents if either named individual was the monitor, which would generate more savings. Sixty-three percent believe the central monitor would generate more savings, 21% believe the average monitor would generate more, and 16% say they could not decide.

Third, subjects were asked about how much information about the saver's savings would spread under each of the two monitors: the majority perceive that there would be more information flow under the central monitors. Sixty-six percent of the village would come to find out if there was a central monitor but only 41% if there was an average monitor.

Fourth, survey responses suggest that subjects are cognizant that even successful savers will often fall short of stated goals. Given a goal of Rs. 1,500, under an average monitor they predict Rs. 819 in savings but Rs. 1,132 under a central monitor.

**Panel A: Saver vs. Monitor**

1). How likely is:
  A) the Saver
  B) the Monitor
 to spread information to others if the Saver who had a goal of Rs. 1,500 saves a high amount (Rs. 1,500) or a low amount (Rs. 100)?

**Panel B: Average vs. Central Monitor**

2). Will the Saver save more with an Average monitor or a Central monitor?

3). Suppose that a Saver has a goal of Rs. 1,500. How much will the Saver save with:
  A) an Average Monitor
  B) a Central Monitor

4). What fraction of the village will come to learn of the Saver's savings if she is assigned:
  A) an Average Monitor
  B) a Central Monitor

**Panel C: Successful vs. Unsucessful Saver**

5). If given the choice between a saver with:
  A) High Savings (Rs. 1,000)
  B) Low Savings (Rs. 100)
 who would you select for each of the following opportunities:
  i) Supervisor Job
  ii) Organizer of Village Event
  iii) Collector of Funds for Village
  iv) Job that requires manual labor

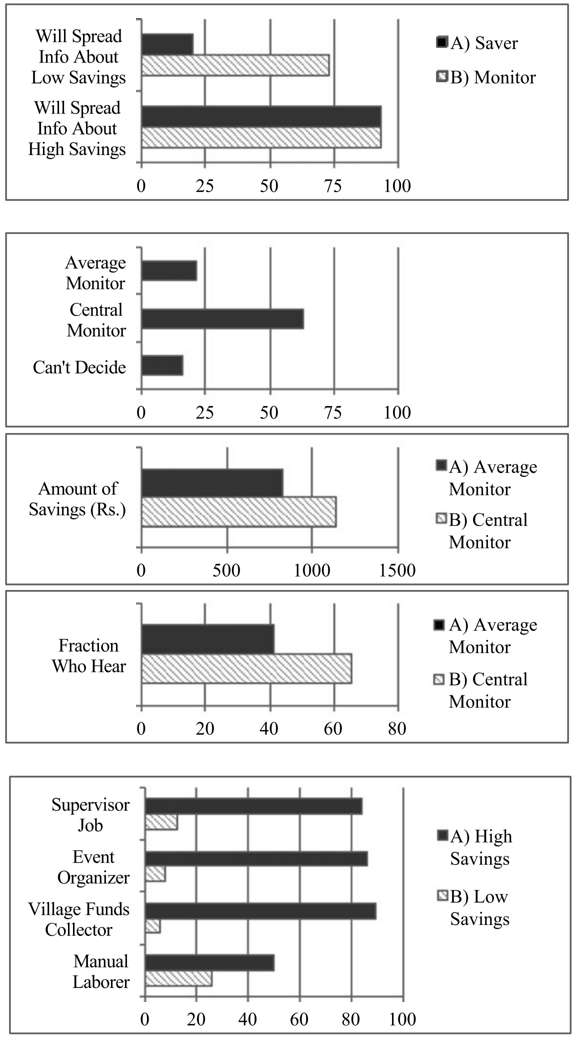


FIGURE 4.—Supplemental survey evidence. Notes: Surveys conducted with 128 individuals across eight villages. The villages were all in the study districts and were selected to be comparable to the study villages. Before the surveys were asked, four randomly selected households were selected to conduct the gossip questionnaire from Banerjee et al. (forthcoming). In the questions presented in Panel B, actual names of villagers were given for the Average Monitor and the Central Monitor. The Average Monitor name was selected by visiting houses according to the right-hand rule. The name of the Central monitor was obtained from the gossip questionnaires.

Fifth, subjects recognize that better savers would experience better rewards for their savings behavior in the future. If given the choice between a saver who saved a large versus a small amount (Rs. 1,000 versus Rs. 100, with a goal of Rs. 1,500), they would predominantly be more likely to take the more successful saver for a supervisor job, an event organizer, or to be a village funds collector. On the other hand, for a manual laborer, the choice seems more even, as one may expect. The evidence suggests that respondents are more willing to allocate jobs that require responsibility to those who saved more, consistent with the interpretation that respondents interpret saving more in the experiment as a signal of responsibility.[21]

Thus, this is a setting where savers understand that monitors would talk about their progress, that there are returns to reputation, that central monitors spread information more widely, and that recognizing this, savers would work to save more if given such a monitor.

### 3.2. *Sketch of the Model*

As the survey responses suggest, the reputational consequences of monitors likely constitute an important source of motivation to save in our experiment.[22] Moreover, the impacts of the monitor may be heterogeneous by network position. To explore this, we embed the standard signaling model of Spence (1973) in a network.

In our model, agents decide whether to save with a signaling motive in mind.[23] We assume that responsible individuals have a lower cost of savings, and that an agent receives utility benefits from interacting with others in the future who consider her responsible. In our setting, once monitors learn about the savings decision of the saver, they can (and do) pass that information on to others in the network. The set of people who will ultimately observe the signal sent by the saver is a function of the monitor's position in the network. Thus, when the agent decides to save, she will consider both how many people will learn about her actions and the likelihood of interacting with such individuals in the future.

The novelty of the model comes from taking an otherwise complicated object—agents interacting on a network—and modeling it naturalistically. The model demonstrates that we can focus on two aspects of the agents' network position: monitor centrality and saver–monitor proximity. We also obtain a new network measure that predicts how a monitor's network position relative to the saver impacts savings that we can also take to the data.

#### *Types*

There is a community of $n$ individuals. Each individual $i$ has a type $\theta_i \in \{H, L\}$, where $\theta_i$ is independently distributed and $\mathrm{P}(\theta_i = H) = \frac{1}{2}$. Type $H$ denotes a responsible agent. Each individual needs to pay a type-specific cost, $c_{\theta_i}$ to save, where $c_H < c_L$. This cost maps to the effort required for an agent to overcoming (with effort) her time inconsistency, temptations, or inattention issues, and these costs are lower for responsible types. Thus, being responsible is helpful for saving but also for other types of day-to-day behaviors and interactions in the community.[24]

[21] An alternative story is that savings does not contain useful information about responsibility but serves purely as a coordination device. There may be a norm "we do not hire people who have low savings as we do not have enough jobs to go around and we need a rule for deciding who gets the job, even though savings is unrelated to job performance." That the respondents tend to pick higher savers systematically for jobs that require responsibility, but not for manual labor, at least suggests that such a coordination story is unlikely.

[22] In fact, it was actually a member of a village in a different study's focus group who originally suggested the core experimental design to us, citing the idea that reputation about individuals accumulating savings when they commit to do so could be leveraged to help encourage savings behavior.

[23] As in Spence (1973), we abstract from the direct benefits savers might receive from the act of saving alone.

[24] In fact, survey data show that a randomly-chosen individual is 6pp more likely to believe that an individual who reached her goal is responsible (mean 0.46) relative to an individual who did not reach her goal. Anecdotal evidence presented in Supplemental Material Appendix B suggests that this influences how people will think of the saver in a labor market situation in the future, consistent with the survey evidence described above.

*Timing, Actions, and Payoffs*

Our model has two periods. The first summarizes the entire 6-month savings period of our experiment, during which the agent makes her savings decision, $s_i \in \{1, 0\}$. For simplicity, we take this decision to be binary; she either saves a large or a small amount over the entire period. Next, the monitor is informed of the saver's decision, $s_i$. The monitor can then diffuse this information through the network. The monitor's conversations are not strategic and are independent of $s_i$. Through this diffusion process, which is governed by the network structure, a subset of the community will be informed of $s_i$ at the end of the first period.

The second period of the model captures future interactions with villagers following the intervention's end. Savers interact with members of the community, again in a process governed by network structure, and receive payoffs from each interaction. Because these future interactions may take many forms under a wide set of circumstances, we model the payoffs in a reduced form way. The saver's payoff is equal to the third party's posterior expectation of the saver's productivity. Let $y(s_i)$ denote this posterior expectation.

The saver saves $s_i = 1$ if and only if the expected increase in payoff to saving the high amount exceeds the cost of doing so:

$$q_{ij}\big[y(1) - y(0)\big] > c_{\theta_i}, \tag{3.1}$$

where $q_{ij}$ is the expected number of individuals that saver $i$ with monitor $j$ will encounter in the future who also would have heard directly or indirectly about the saver's choice of savings. Thus, equation (3.1) captures the fact that the saver only benefits from choosing $s_i = 1$ when she interacts with informed third parties.

*Network Interactions*

We model network interactions in a parsimonious and natural way. In the first stage, the monitor meets her friend with some probability, meets her friend's friend with a lesser probability, and so on. Similarly, in the second stage, the saver meets any of her friends with some probability, meets a friend's friend with a lower probability (i.e., needs to meet a friend and also be referred to the friend's friend), and so on.

So, the probability that the agent meets some third party in the future, who will in turn have heard about the saver's choice of savings $s_i$ through the network, will depend on the network structure and the position of the saver $i$ and the monitor $j$ in the network. By modeling information flow from monitors to others in the network and the possibility of the saver running into the third party through the network in this way, we compute:

$$\begin{aligned} q_{ij} &= n \cdot \text{Social Proximity of Monitor and Saver} \\ &\quad + \frac{1}{n} \cdot \text{Monitor Centrality} \times \text{Saver Centrality} \end{aligned}$$

in the manner as described precisely in Appendix A.

*Equilibrium*

Under simple parameter assumptions,[25] we show there is a Perfect Bayesian equilibrium and a cutoff $\widehat{q}$ such that

$$s_i = \begin{cases} 1 & \text{if } \theta_i = H \text{ and } q_{ij} \geq \widehat{q}, \\ 0 & \text{if } \theta_i = H \text{ and } q_{ij} < \widehat{q}, \text{ and} \\ 0 & \text{if } \theta_i = L. \end{cases}$$

With this stylized structure on interactions, the signaling model predicts that in the data, we should see more savings where we have randomly assigned a saver to a monitor with higher $q_{ij}$. Ceteris paribus, $q_{ij}$ is higher when (1) saver–monitor proximity is higher or (2) monitor centrality is higher. Therefore, random assignment to higher centrality monitors or monitors closer to the saver in the network should yield higher savings.[26]

## 4. THE VALUE OF CENTRAL AND PROXIMATE MONITORS

### 4.1. *Random Monitors*

Our main results analyze how the centrality, proximity, and combined model-based regressor value of randomly-assigned monitors influence savings. Before turning to this, we briefly discuss the average impact of monitors relative to the baseline treatment bundle (non-monitored treatment). The main outcome of interest is the log of total savings across all accounts. Conceptually, this is the key outcome, as subjects could simply move funds from other places into the target formal account.[27]

Table II presents the effects on the log of total savings across all formal and informal savings vehicles (columns (1)–(3)) and then repeats this for whether the goal was reached (column (4)) and the savings in the target account (column (5)). In column (2), we also include village fixed effects as well as saver controls for saving goal, age, marital status, number of children, preference for bank or post office account, baseline bank or post office account ownership, caste, elite status, number of rooms in home, and type of electrical connection.

In column (3), we take a strict approach and use machine learning to select what out of the long list of controls we should include, which could potentially account for why we are seeing a treatment effect. This is the new technique called double post-LASSO of Belloni, Chernozhukov, and Hansen (2014a) (see also Belloni, Chernozhukov, and Hansen (2014b)). The idea is straightforward—because networks are not randomly assigned, and because we have many characteristics for which we could control, we allow machine learning (specifically LASSO) to pick out which covariates to include in the final regression specification. Here, our goal is to regress an outcome $y$ on a treatment $T$, observing a large vector of $X$'s. We use LASSO twice: first $y$ on $X$ to select $X^{\mathrm{RF}}$ (the reduced form regression) and second $T$ on $X$ to select $X^{\mathrm{FS}}$ (the first-stage regression). Taking the union of these selected regressors as $X^\star = X^{\mathrm{RF}} \cup X^{\mathrm{FS}}$, in a final step, we regress $y$ on $T$ and $X^\star$.

[25] If parameters are such that the cost of accumulating savings exceeds any possible posterior update about one's productivity, then of course everyone has $s_i = 0$.

[26] Note, if there was no signaling effect altogether because either productivity type was unrelated to cost of savings or because costs were systematically too high or too low for both parties and therefore only pooling, the regression would have no slope since all agents pool on $s_i = 0$.

[27] We developed our survey instrument after conducting numerous conversations about savings in similar communities. We attempted to be as comphrensive as possible in enumerating both formal and informal savings, including banks, post offices, SHGs, RoSCAs, MFIs, insurance schemes, cash at home, and so on.

TABLE II

EFFECT OF RANDOM MONITORS ON SAVINGS[a]

| Variables | (1) Log Total Savings | (2) Log Total Savings | (3) Log Total Savings | (4) Reached Goal | (5) Savings in Target A/c |
|---|---|---|---|---|---|
| Monitor Treatment: Random Assignment | 0.358 | 0.279 | 0.309 | 0.067 | 298.116 |
| | (0.150) | (0.161) | (0.151) | (0.031) | (130.053) |
| | [0.103, 0.614] | [0.005, 0.552] | [0.053, 0.565] | [0.015, 0.120] | [77.140, 519.092] |
| Observations | 549 | 549 | 549 | 679 | 679 |
| Fixed Effects | None | Village | | | |
| Controls | None | Saver | Double-Post LASSO | Double-Post LASSO | Double-Post LASSO |
| Non-Monitored Mean | 7.670 | 7.670 | 7.670 | 0.0719 | 336 |

[a]Table shows the effects of receiving a randomly allocated monitor on log total savings in the 30 random assignment villages. Total savings is the amount saved across all savings vehicles—the target account and any other account, both formal and informal including money held "under the mattress." Reached Goal is a dummy for if the saver reached their saving goal. Savings in Target A/c is the amount saved in the target account. Sample constrained to individuals who answered our questionnaire. Saver controls include the following saver characteristics: savings goal, saver centrality, age, marital status, number of children, preference for bank or post office account, whether the individual has a bank or post office account at baseline, caste, elite status, number of rooms in the home, and type of electrical connection. The double-post LASSO specification in columns (3)–(5) consider all saver controls and individual village fixed effects in the possible control set. Standard errors (clustered at the village level) are reported in parentheses. 90% confidence intervals are reported in brackets.

The coefficient on $T$ resulting from this procedure is the estimated causal treatment effect. The idea is that if some component of observables explained either treatment or the outcome variable, and therefore could explain the relationship of $T$ to $y$, then we allow that component to be selected. Of course, because the monitoring treatment is random, the double post-LASSO procedure for estimating the treatment effect of receiving any random monitor deals with covariate imbalance. However, when estimating the effect of monitor network position ($q_{ij}$ or proximity to monitor and centrality of monitor), below, double post-LASSO allows us to look at how the relationship of monitor network position and savings is affected or explained away by the other characteristics that the double post-LASSO selects.

Columns (1)–(3) present qualitatively similar results. We describe the results from column (3) and find that being randomly assigned to a monitor leads to a 0.309 log point increase in the total savings across all accounts. This corresponds to a 36% increase in savings across all savings vehicles of the households.[28,29]

Finally, in columns (4) and (5), we repeat the exercise of column (3) with two other outcome variables. We show that the results hold for both whether the saver reached his goal (a 93% increase) and the savings in the target account (88.7%).

Given these large impacts on overall savings, we next explore whether this increase is driven by a few individuals dramatically increasing their savings or by individuals across the group of savers more broadly. In Panel A of Figure 5, we plot the cumulative distribution functions of the log of total savings normalized by the savings goal for monitored versus non-monitored savers. The figure suggests that the average treatment effects are

[28]Given that we find an increase in savings across all accounts, we need not fear that the treatment effects are simply the cause of moving funds from one location into the target account.

[29]We have also checked whether participating in our experiment and receiving a bank account and bi-weekly visits increases total savings. We find very small statistically insignificant effects in Supplemental Material Appendix J.

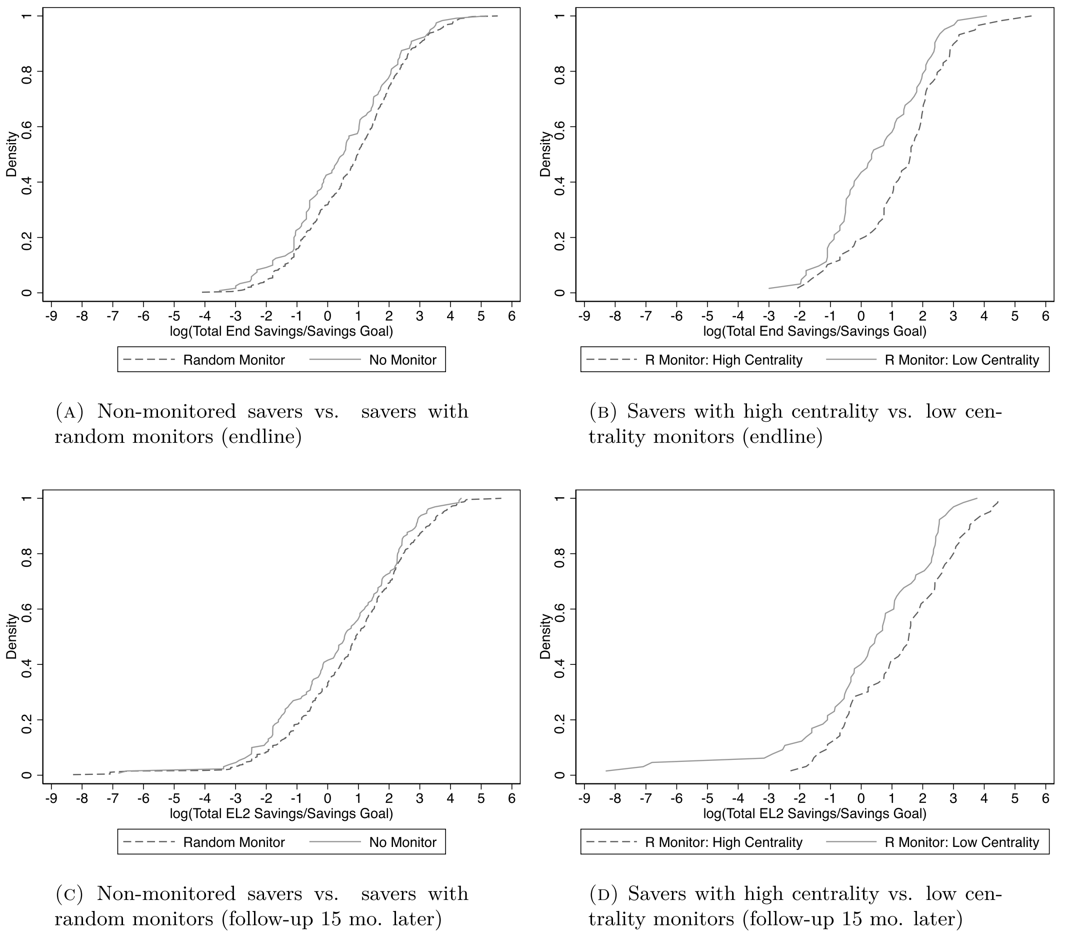


(A) Non-monitored savers vs. savers with random monitors (endline)

(B) Savers with high centrality vs. low centrality monitors (endline)

(C) Non-monitored savers vs. savers with random monitors (follow-up 15 mo. later)

(D) Savers with high centrality vs. low centrality monitors (follow-up 15 mo. later)

FIGURE 5.—Distributions (CDF) of log(Total Savings/Savings goal) by treatment. Notes: The panels plot the CDFs of $\log(\frac{TotalSavings}{SavingsGoal})$ for different experimental subsamples. Panels A and B present results for the endline. Panels C and D present results from the follow-up survey 15 months after the end of the experiment. In Panel A, we plot the CDFs for the non-monitored savers and the monitored savers, both in random assignment villages. $p = 0.081$ from a Kolmogorov–Smirnov test for the difference in distributions. In Panel B, we plot the CDFs for the monitored savers in random assignment villages with high versus low centrality. Here high centrality is defined as top 15% of monitor centrality and low as bottom 15%. $p = 0.01$ for a Kolmogorov–Smirnov test for the difference in distributions. In Panel A, we plot the CDFs for the non-monitored savers and the monitored savers, both in random assignment villages. $p = 0.168$ from a Kolmogorov–Smirnov test for the difference in distributions. In Panel B, we plot the CDFs for the monitored savers in random assignment villages with high versus low centrality. Here high centrality is defined as top 15% of centrality and low as bottom 15%. $p = 0.061$ from a Kolmogorov–Smirnov test for the difference in distributions.

not simply capturing large increases experienced by a small number of savers in the tail of the distribution. Indeed, the intervention shifts savers to save more across the entire distribution (a Kolmogorov–Smirnov test rejects that the CDFs are the same with $p = 0.08$).

## 4.2. *Monitor Centrality and Proximity*

We now turn to our main results: how $q_{ij}$ (the model-based regressor), monitor centrality, and saver–monitor proximity influence saving behavior. Table III presents the results

TABLE III

TOTAL SAVINGS BY NETWORK POSITION OF RANDOM MONITOR[a]

| Variables | (1) Log Total Savings | (2) Log Total Savings | (3) Log Total Savings | (4) Log Total Savings | (5) Log Total Savings | (6) Log Total Savings | (7) Reached Goal | (8) Savings in Target A/c |
|---|---|---|---|---|---|---|---|---|
| Monitor Centrality | 0.180 | | 0.136 | | 0.155 | | | |
| | (0.073) | | (0.072) | | (0.067) | | | |
| | [0.056, 0.305] | | [0.013, 0.259] | | [0.040, 0.269] | | | |
| Saver-Monitor Proximity | | 1.049 | 0.879 | | 1.110 | | | |
| | | (0.349) | (0.330) | | (0.290) | | | |
| | | [0.457, 1.642] | [0.318, 1.441] | | [0.618, 1.603] | | | |
| Model-Based Regressor ($q_{ij}$) | | | | 0.219 | | 0.290 | 0.042 | 154.935 |
| | | | | (0.118) | | (0.106) | (0.021) | (93.203) |
| | | | | [0.018, 0.419] | | [0.109, 0.470] | [0.007, 0.077] | [−3.428, 313.299] |
| Observations | 426 | 426 | 426 | 426 | 426 | 426 | 526 | 526 |
| Fixed Effects | Village | Village | Village | Village | | | | |
| Controls | Saver and Monitor | Saver and Monitor | Saver and Monitor | Saver and Monitor | Double-Post LASSO | Double-Post LASSO | Double-Post LASSO | Double-Post LASSO |
| Depvar mean | 8.029 | 8.029 | 8.029 | 8.029 | 8.029 | 8.029 | 0.141 | 629 |

[a]Table shows impacts on log total savings by monitor network position. Total savings is the amount saved across all savings vehicles—the target account and any other account, both formal and informal including money held "under the mattress." Reached Goal is a dummy for if the saver reached their saving goal. Savings in Target A/c is the amount saved in the target account. Sample constrained to savers who received a monitor in the 30 random assignment villages, who answered our questionnaire. The variable "Model-Based Regressor" is defined as $q_{ij}$ in the framework. Saver and Monitor controls include savings goal and saver centrality, along with the following variables for each monitor and saver: age, marital status, number of children, preference for bank or post office account (saver only), whether the individual has a bank or post office account at baseline, caste, elite status, number of rooms in the home, and type of electrical connection. Saver and Monitor controls additionally include the geographical distance between their homes. The double-post LASSO specifications in columns (5)–(8) consider all saver and monitor controls and individual village fixed effects in the possible control set. Standard errors (clustered at the village level) are reported in parentheses. 90% confidence intervals are reported in brackets.

of regressions of log total savings across all accounts on monitor centrality, saver–monitor proximity, and $q_{ij}$, as well as a battery of controls.

Columns (1)–(4) look at monitor centrality, saver–monitor proximity, both, and $q$ (the model-based regressor), and all include village fixed effects, controls for the savings goal, saver centrality, and controls for saver and monitor characteristics including age, marital status, number of children, preference for bank or post office account (saver), baseline bank or post office account ownership, caste, elite status, number of rooms in home, and type of electrical connection, along with geographic distance between the homes of the saver and monitor.

In columns (5) and (6), we repeat the same exercises of columns (3) and (4), but use double post-LASSO of Belloni, Chernozhukov, and Hansen (2014a). These provide our preferred estimates, though the results are comparable across the board. They are preferred because double post-LASSO employs a selection of regressors such that if some combination of covariates was effectively driving the effect on savings and we attributed it to networks, then the selector would include these and would actually kill the network effect we estimate. On the other hand, if regressors that predict neither the treatment (network position of monitor relative to saver) nor the outcome are being included, then it simply adds noise.

Consistent with our model, we find that being assigned to a central monitor or a proximate monitor generates large increases in savings. Namely, a one-standard-deviation increase in the centrality of the monitor corresponds to a 0.155 log point increase in the log total savings, or a 16.8% increase—a large effect (column (5)). Further, in Panel B of Figure 5, we explore the distributional effects of receiving a high centrality monitor versus a low centrality monitor. Receiving a high centrality monitor does shift most of the distribution to the right, again suggesting that increases are not only driven by a small number of highly-impacted savers (a Kolmogorov–Smirnov test rejects equality of the CDFs, $p = 0.01$).

Turning to proximity, moving from a monitor of distance 3 to 2 leads to a 20.3% increase in the total savings across all accounts—again a large effect. Finally, in column (6), we look at the effect of our model-based regressor, $q_{ij}$. A one-standard-deviation increase in the model-based regressor corresponds to a 33.6% increase in savings.

Columns (7) and (8) reproduce the results for the model-based regressor, $q_{ij}$, with two other outcome variables. In column (7), we show that a one-standard-deviation increase in $q_{ij}$ corresponds to a 4.2pp increase in the likelihood of a goal being met, which is two-thirds the effect size of being assigned an average monitor. Similarly, we see in column (8) that an increase in $q_{ij}$ corresponds to an increase in the amount saved in the target account.

Thus, we show that randomly assigning a better monitor in terms of the model ($q_{ij}$), or randomly assigning a more central and more proximate monitors encourages savings across all accounts, including both formal and informal. That these results hold controlling for numerous demographic characteristics of both savers and monitors suggests that observables that may be correlated with network position cannot explain our proximity and centrality results. The covariate controls described above include caste group fixed effects and even the geographic distance between homes of savers and monitors. Traits such as these could have been thought to be driving the network effect through omitted variables, but our results are estimated conditional on this variation and a machine learning technique actually jettisons a number of controls and improves our estimates. Magnitudes and significance are essentially the same even when entirely removing this bevy of characteristics (see Supplemental Material Appendix F, Table F.3), which bolsters the idea that the effects are not driven by these characteristics.

#### 4.2.1. *Multigraphs: Investigating Multiple Link-Types*

We next investigate whether the observed patterns are driven by a specific slice of the multigraph. It is possible that the financial component or the advice component of the network could be driving the effects. This would be true if, for example, financial information were only passed between individuals conducting financial transactions with one another. Table IV presents a version of our main specification, but allows centrality, proximity, and the model-based regressor to vary by relationship type. While the results get

TABLE IV

TOTAL SAVINGS BY NETWORK POSITION OF RANDOM MONITOR: MULTIGRAPH ANALYSIS[a]

| Variables | (1) Log Total Savings | (2) Log Total Savings | (3) Log Total Savings | (4) Log Total Savings |
|---|---|---|---|---|
| Monitor Centrality | 0.183 | | 0.141 | |
| | (0.101) | | (0.096) | |
| | [0.012, 0.354] | | [−0.021, 0.304] | |
| Monitor Centrality: Financial Network | −0.000 | | −0.015 | |
| | (0.154) | | (0.147) | |
| | [−0.262, 0.261] | | [−0.264, 0.234] | |
| Monitor Centrality: Advice Network | −0.004 | | −0.000 | |
| | (0.121) | | (0.112) | |
| | [−0.209, 0.202] | | [−0.191, 0.191] | |
| Saver-Monitor Proximity | | 0.731 | 0.634 | |
| | | (0.546) | (0.515) | |
| | | [−0.196, 1.658] | [−0.241, 1.509] | |
| Saver-Monitor Proximity: Financial Network | | 0.311 | 0.295 | |
| | | (0.930) | (0.936) | |
| | | [−1.270, 1.891] | [−1.296, 1.887] | |
| Saver-Monitor Proximity: Advice Network | | 0.237 | 0.147 | |
| | | (0.747) | (0.773) | |
| | | [−1.032, 1.507] | [−1.166, 1.459] | |
| Model-Based Regressor: Full Network ($q_{ij}$) | | | | 0.234 |
| | | | | (0.176) |
| | | | | [−0.065, 0.532] |
| Model-Based Regressor: Financial Network | | | | −0.005 |
| | | | | (0.189) |
| | | | | [−0.326, 0.316] |
| Model-Based Regressor: Advice Network | | | | −0.013 |
| | | | | (0.163) |
| | | | | [−0.290, 0.264] |
| Observations | 426 | 426 | 426 | 426 |
| Fixed Effects | Village | Village | Village | Village |
| Controls | Saver and Monitor | Saver and Monitor | Saver and Monitor | Saver and Monitor |
| Depvar Mean | 8.029 | 8.029 | 8.029 | 8.029 |

[a]Table shows impacts on log total savings by monitor network position, using different definitions of link-types. Total savings is the amount saved across all savings vehicles—the target account and any other account, both formal and informal including money held "under the mattress." Reached Goal is a dummy for if the saver reached their saving goal. Savings in Target A/c is the amount saved in the target account. Sample constrained to savers who received a monitor in the 30 random assignment villages, who answered our questionnaire. The variable "Model-Based Regressor" is defined as $q_{ij}$ in the framework. Saver and Monitor controls include savings goal and saver centrality, along with the following variables for each monitor and saver: age, marital status, number of children, preference for bank or post office account (saver only), whether the individual has a bank or post office account at baseline, caste, elite status, number of rooms in the home, and type of electrical connection. Saver and Monitor controls additionally include the geographical distance between their homes. Standard errors (clustered at the village level) are reported in parentheses. 90% confidence intervals are reported in brackets.

noisy, we find that only the centrality, proximity, and model-based regressor in the union of all relationships appear to matter. This is natural as individuals could pass information across link types: for instance, to a coworker, who then tells a friend, who then tells his neighbor about the information when borrowing rice.

### 4.3. *Effect of Central Monitors on Beliefs About Savers*

We next make use of novel supplemental data to provide evidence in support of the reputational mechanism of Section 3. One necessary condition for reputation to be at play is for the monitors and other community members to actually discuss the savings of participants. In fact, more than 60% report doing so in the last two weeks of the savings period. Further, 40% of monitored savers also report that the monitor passed information about their progress to others.[30]

Moreover, we attempt to track this information flow from monitors to other community members. Our follow-up survey, administered 15 months after the end of the intervention, asks respondents their views about a randomly chosen set of 12 savers who participated in our experiment. Namely, we capture a measure of responsibility—whether the saver is viewed generally (in avenues beyond savings alone) as being good at meeting self-set goals. We test whether community members update their beliefs about the saver's ability to meet goals more in response to their behavior in our experiment when the monitor is central.

Table V presents the results of this exercise. We examine a regression of whether the interviewee updated her beliefs about the saver's responsibility in the direction of the saver's savings goal attainment on the centrality of the randomly assigned monitor as well as the proximity between the interviewee and the saver's monitor (columns (1)–(3)). We repeat this exercise, changing the outcome variable to whether the interviewee knows correctly whether the saver attained her goal (columns (4)–(6)). Our preferred outcome is the update in the responsibility metric: if the monitor is more central, a random interviewee in the village is more likely to have a better view of the saver's general responsibility if she succeeded or a worse view of it if the saver failed. Our regression specifications include no fixed effects, village fixed effects, or interviewee fixed effects, the latter of which therefore captures variation within an interviewee but across randomly assigned saver–monitor pairs. We find that if a saver is randomly assigned a more central monitor, the respondent is more likely to believe that the saver is responsible and also is more likely to know if the saver reached her goal.

While interesting, this dynamic is not necessary for our story. Specifically, it need not be the case that the information has already or immediately spread. What is important in our framework is that when the saver impresses the monitor, there may be benefits at some point when a new opportunity arises (much like sending a letter of recommendation).

This is an admittedly imperfect exercise. We use self-reported data on whether people chat about others, whether people hear gossip about themselves through back channels, and several questions about respondents' perspectives on other savers' responsibility profiles and savings habits in the experiment. The usual caveats about self-reported data apply here and, further, we are not making a causal claim that this shift in beliefs exactly corresponds to the shift in savings. Nonetheless, we want to emphasize that the evidence

[30]We asked "did your monitor tell others about your savings goal or about your progress toward trying to meet it?" This is striking because it requires enough communication such that savers hear gossip about themselves. Our own reflection suggests it is rare for people to gossip about a subject in front of him/her.

TABLE V

BELIEFS ABOUT SAVERS AND MONITOR CENTRALITY[a]

| | Respondent's Beliefs About the Saver | | | | | |
|---|---|---|---|---|---|---|
| | Responsibility | | | Goal Attainment | | |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Monitor Centrality | 0.043 | 0.040 | 0.037 | 0.027 | 0.021 | 0.021 |
| | (0.014) | (0.014) | (0.015) | (0.010) | (0.008) | (0.008) |
| | [0.019, 0.066] | [0.016, 0.063] | [0.013, 0.062] | [0.010, 0.043] | [0.007, 0.034] | [0.007, 0.035] |
| Respondent-Monitor Proximity | 0.046 | 0.014 | 0.031 | 0.005 | −0.003 | −0.003 |
| | (0.043) | (0.038) | (0.035) | (0.019) | (0.020) | (0.024) |
| | [−0.028, 0.119] | [−0.050, 0.078] | [−0.029, 0.091] | [−0.027, 0.037] | [−0.037, 0.030] | [−0.043, 0.038] |
| Observations | 4,743 | 4,743 | 4,743 | 4,743 | 4,743 | 4,743 |
| Fixed Effects | None | Village | Respondent | None | Village | Respondent |
| Controls | Saver and Monitor | Saver and Monitor | Saver and Monitor | Saver and Monitor | Saver and Monitor | Saver and Monitor |
| Depvar Mean | 0.241 | 0.241 | 0.241 | 0.0614 | 0.0614 | 0.0614 |

[a]Table explores beliefs of 615 respondents across the 30 random villages, each of whom was asked in the 15-month follow-up survey to rate approximately 8 randomly selected savers who had a monitor from their village. "Responsibility" is constructed as 1(Saver reached goal) $*$ 1(Respondent indicates saver is good or very good at meeting goals) + (1 − 1(Saver reached goal)) $*$ 1(Respondent indicates saver is mediocre, bad, or very bad at meeting goals). "Goal Attainment" measures whether the saver reached her goal and the respondent correctly believed this to be true. Saver and Monitor controls include savings goal and saver centrality, along with the following variables for each monitor and saver: age, marital status, number of children, preference for bank or post office account (saver only), whether the individual has a bank or post office account at baseline, caste, elite status, number of rooms in the home, and type of electrical connection. Saver and Monitor controls additionally include the geographical distance between their homes. All specifications include controls for saver and respondent centrality and for the proximity between the respondent and monitor. Columns (2) and (5) include village fixed effects. Columns (3) and (6) include respondent fixed effects. Standard errors (clustered at the village level) are reported in parentheses. 90% confidence intervals are reported in brackets.

provided here is (a) largely consistent with our framework, (b) mostly self-consistent, and (c) agrees with the anecdotal evidence provided in Supplemental Material Appendix B. Given the difficulties in digging into such a mechanism in a networks setting, we argue that this simple idea—simply asking whether conversations happened, asking whether people changed their views of others, etc.—which has not been used much in this literature, has tremendous value for this research program.

Consistent with the perception effects, conversations with study participants and other villagers support the idea that reputational mechanisms are at play in our experiment. In fact, our experimental design was based, in part, on a conversation with a gentleman in a rural village. In Appendix B of the Supplemental Material, we present short excerpts of conversations with participants that we recorded. Many villagers described wanting to impress their monitor in general and paying special attention when that monitor was important. Some respondents gave us specific examples of why impressing the monitor would be helpful in the future. Finally, we remind the readers that in our follow-up survey across 128 subjects in eight new villages, the responses were consistent with what has been documented here (Figure 4).

### 4.4. *Longer-Run Impacts*

Given that our treatment increased total savings across all accounts, we next ask whether we can detect any lasting benefits of the additional savings. This is a difficult question, so to address this, in our 15-month follow-up survey, we asked subjects about their ability to cope with various shocks. Given that our intervention helped savers to increase their stock of savings, we can ask if, in the subsequent 15 months, they were less likely to be in a situation where they did not have money to cope with a shock.[31]

Specifically, we posed a series of questions to the savers as to whether they faced a specific hardship for which they did not have enough savings to purchase a remedy (e.g., falling ill and being unable to purchase medicine). Table VI presents the results. We measure effects on the total number of unmitigated shocks (columns (1)–(2)), whether the household experienced fewer unmitigated shocks than the median (columns (3)–(4)), incidence of unmitigated health shocks (columns (5)–(6)), and incidence of unmitigated household consumption shocks (columns (7)–(8)). Specifications are shown with and without village fixed effects, and all regressions use the standard saver controls. We find that being randomly assigned a monitor leads to a decline in the rate at which individuals face a shock and are unable to purchase a remedy. For instance, there is a 0.202 decline in the total number of shocks (on a base of 1.774, column (1), $p = 0.12$). Further, there is a 7.6pp decline in the probability that a household has greater than median number of instances where they were unable to cope with the shock ($p = 0.076$, column (3)). We find suggestive, though not statistically significant, effects when we look at health and household expenditures as separate categories. We acknowledge that the types of shocks that the intervention helped savers to mitigate are likely of modest scale.[32] The key point is that there are, nonetheless, tangible benefits of savings for situations like these. Note that it could also be another channel: the tangible benefit of improved reputation, which may cause others to be willing to help the saver in times of need.

[31]Note that this could arise for two reasons. First, and perhaps the ex ante more likely reason, agents would have more money to deal with the same distribution of shocks. Second, agents could have invested in shock mitigation. Our analysis does not need to take a stand.

[32]We are not claiming that the gains in savings had large persistent health benefits, for example.

TABLE VI

SHOCK MITIGATION FOR MONITORED SAVERS IN RANDOM VILLAGES[a]

| Variables | (1) Shocks Total Number | (2) Shocks Total Number | (3) Shocks Above Median | (4) Shocks Above Median | (5) Shocks Health | (6) Shocks Health | (7) Shocks HH Expenditure | (8) Shocks HH Expenditure | (9) Final Endline Savings log(Tot. Sav.) 15 mos. | (10) Final Endline Savings log(Tot. Sav.) 15 mos. |
|---|---|---|---|---|---|---|---|---|---|---|
| Monitor Treatment: Random Assignment | −0.202 | −0.254 | −0.076 | −0.095 | −0.072 | −0.099 | −0.044 | −0.062 | 0.315 | 0.289 |
| | (0.129) | (0.131) | (0.042) | (0.044) | (0.062) | (0.067) | (0.040) | (0.043) | (0.200) | (0.195) |
| | [−0.417, 0.013] | [−0.473, −0.036] | [−0.147, −0.006] | [−0.170, −0.021] | [−0.175, 0.031] | [−0.211, 0.012] | [−0.111, 0.023] | [−0.134, 0.009] | [−0.020, 0.650] | [−0.037, 0.615] |
| Observations | 1,173 | 1,173 | 1,173 | 1,173 | 1,173 | 1,173 | 1,173 | 1,173 | 1,172 | 1,172 |
| Fixed Effects | Village | None | Village | None | Village | None | Village | None | Village | Village |
| Controls | Saver | Saver | Saver | Saver | Saver | Saver | Saver | Saver | Saver | Saver |
| Non-monitored Mean | 1.774 | 1.774 | 0.579 | 0.579 | 0.857 | 0.857 | 0.496 | 0.496 | 7.652 | 7.652 |

[a]Table presents effects of a monitor on shock mitigation in the 30 random allocation villages. All outcome variables are measured in the 15-month follow-up survey. The outcome variables in columns (1)–(8) are all measures of unmitigated shocks experienced by the savers between the end of the six-month savings period and the survey. Columns (9)–(10) report log total savings across all accounts at the time of the survey. The total number of shocks measures the number of unmitigated shocks experienced, including deaths, family illnesses, health shocks, livestock shocks, and unexpected HH expenditures. Sample constrained to all savers in the sample who answered our questionnaire. If a response was missing for a category, the observation is missing in the regression. Controls include the following saver characteristics: log savings goal, saver centrality, age, marital status, number of children, preference for bank or post office account, whether the individual has a bank or post office account at baseline, caste, elite status, number of rooms in the home, type of electrical connection, a dummy for endogenous village (when no village fixed effects), and a dummy for endogenous monitor. Columns (1), (3), (5), (7), and (9) include village fixed effects. Standard errors (clustered at the village level) are reported in parentheses. 90% confidence intervals are reported in brackets.

Finally, in the last two columns of Table VI, we present the effects of the random monitor treatment on log savings balances 15 months after the intervention. Remarkably, while a bit noisier, the size of the increase in savings is as large as that reported in Table II ($p = 0.12$, column (9)). This suggests that individuals are able to maintain their savings even after the monitors are no longer actively receiving information. Panels C and D of Figure 5 show that the increases in savings across the distribution are still apparent 15 months later.

We have documented increases in savings both during and considerably after our experiment. One concern the reader may have is whether there is widespread over-savings. We emphasize our results are for total savings across all household accounts; recall, only the target account is revealed to the monitor. So the fact that we find effects here suggests that over-savings is unlikely. There are also other reasons why we believe over-savings was not a widespread issue: (1) only information in the target account is provided to the monitors, (2) the main effects on total savings are persistent; (3) recall people who were worried about undue negative social pressure from potential future monitoring could privately opt out of ever participating in the program; (4) individuals who were worried about excess pressure from their assigned monitors could have opted out at any time, ending the revelation of information, and chose not to;[33] (5) as discussed below, we demonstrate increased labor supply likely drives the savings increases rather than taking costly loans or reducing essential consumption; (6) we document positive effects on substantive outcomes like shock mitigation.

## 5. THREATS TO VALIDITY

### 5.1. *Negligibility of Monitor Incentives*

There are two natural questions one may ask when it comes to monitors in this study. First, is it the case that the presence of the monitor causes individuals to unwind their savings from other accounts? Second, does the fact that the monitors received a small incentive drive the results?

We show evidence against both of these hypotheses. Conditional on reaching her goal, a saver exceeds 200% of her goal in 65% of the cases. Further, over 75% of individuals who reach their goal in the target account save in excess of 200% of their target amount across all accounts. This suggests that individuals are not likely subject to undue pressure. They save immensely, mostly do not bunch at their goal, and do not unwind across other accounts.

Turning to the monitor incentives, we do the following. Recall that the monitor incentive function has two discontinuities. In addition to the payment made at the full goal, we added a second discontinuity at the half goal to generate a test. In terms of personal value to the saver, the incentives above and below the half goal should be smooth. So testing for bunching above this threshold should identify how the monitor incentive may have differentially led to behavior nudging people across the threshold. Turning to the full goal amount, this is a mix of potential monitor incentives but also natural incentives to simply reach one's stated goal: they may be saving up for something specific, and furthermore, after all, it is a goal. Both are natural motivations to bunch at the goal.

Table VII presents the results. The outcome variable is a dummy for whether the saver who is in the window of the specified value (1/2 or full goal) has saved weakly more

[33]This is of course subject to the caveat that opting out itself could harm one's reputation.

TABLE VII

NO EVIDENCE OF BUNCHING OR GAMING[a]

| Variables | (1) Exceeded Payment Threshold | (2) Exceeded Payment Threshold | (3) Exceeded Payment Threshold | (4) Exceeded Payment Threshold |
|---|---|---|---|---|
| Monitor R | −0.257 | −0.257 | −0.330 | −0.306 |
| × In Window of Half Goal | (0.136) | (0.135) | (0.143) | (0.140) |
| | [−0.486, −0.028] | [−0.484, −0.030] | [−0.570, −0.090] | [−0.541, −0.072] |
| Monitor E | −0.216 | −0.216 | −0.265 | −0.253 |
| × In Window of Half Goal | (0.134) | (0.133) | (0.141) | (0.137) |
| | [−0.441, 0.009] | [−0.439, 0.007] | [−0.502, −0.028] | [−0.483, −0.022] |
| In Window of Half Goal | 0.900 | 0.900 | 0.765 | 0.765 |
| | (0.090) | (0.089) | (0.105) | (0.105) |
| | [0.749, 1.051] | [0.751, 1.049] | [0.589, 0.941] | [0.589, 0.941] |
| Monitor R | 0.257 | 0.514 | 0.470 | 0.492 |
| × In Window of Full Goal | (0.237) | (0.136) | (0.117) | (0.112) |
| | [−0.142, 0.656] | [0.286, 0.742] | [0.273, 0.667] | [0.304, 0.681] |
| Monitor E | 0.100 | 0.286 | 0.304 | 0.288 |
| × In Window of Full Goal | (0.269) | (0.158) | (0.146) | (0.138) |
| | [−0.353, 0.553] | [0.020, 0.552] | [0.060, 0.549] | [0.057, 0.520] |
| In Window of Full Goal | 0.600 | 0.286 | 0.222 | 0.200 |
| | (0.228) | (0.119) | (0.098) | (0.090) |
| | [0.216, 0.984] | [0.085, 0.486] | [0.058, 0.387] | [0.048, 0.352] |
| Observations | 89 | 115 | 175 | 184 |
| Size of Window Around Half/Full | ± 50/100 | ± 66/200 | ± 100/300 | ± 116.66/350 |
| Depvar Mean | 0.742 | 0.652 | 0.537 | 0.527 |

[a]Table explores different windows of the half goal and the full goal thresholds by treatment. Column (1) considers windows equal to the size of the monitor's incentive: Rs. 50 around the half goal and Rs. 150 of the full goal. In column (2), these windows are Rs. 66 and Rs. 200, respectively. In column (3), the windows are Rs. 100 and Rs. 300 around the half and full goal thresholds. Finally, in column (4), these windows are Rs. 116.66 and Rs. 350. Standard errors (clustered at the village level) are reported in parentheses. 90% confidence intervals are reported in brackets.

than the value. In column (1), we look at the 1/2 goal and full goal savings amounts for each saver and look within a window of the bonus (Rs. 50 or Rs. 150) of each. The first three rows constitute our test of interest as they focus on the 1/2 goal mark. We see that unmonitored savers, conditional on being in the window around the 1/2 goal, are 90% likely to be weakly greater than 1/2 their goal. This drops by 25.7pp ($p = 0.066$) or 21.6pp ($p = 0.11$) when one has a random or endogenous monitor, respectively (column (1)). This suggests that if anything, the fact that the monitor may have an incentive makes it less likely for the saver to just clear the threshold. Of course, we interpret this as the monitor incentive having no meaningful effect, not that it disincentivizes clearing the threshold.

In columns (2)–(4), we repeat the exercise scaling the window by 3/2, 2, and 7/3 (so Rs. 66/Rs. 200, Rs. 100/Rs. 300, and Rs. 116/Rs. 350, respectively). Notice that the set of observations in the window do not change across columns (1) and (2) and similarly (3) and (4) for the 1/2 goal mark. Our results remain essentially the same and we gain precision for the endogenous monitoring case. Notice that the endogenous and random monitoring case cannot be distinguished from half the savers on either side of the window.

Overall, this rejects the bunching hypothesis since, first, in the monitored groups, it is as good as random that people are on either side of the window but, further, if anything, the unmonitored group is significantly more likely to bunch on the right of the 1/2 goal mark despite not facing any monitor incentives by definition. We believe that this is a good test of the impact of monitor incentives because 1/2 goal is not a particularly salient milestone for the saver aside from the monitor incentive.

Rows 4–6 present the same estimates but for the full goal. Note that, by construction, there is likely to be more bunching here (ex ante) simply because individuals set goals for themselves and they may be saving for specific goods. With the most conservative window, we find that 60% of unmonitored savers fall at or within Rs. 150 above the goal, whereas that fraction is 86% and 70% for the monitored savers. As the window widens, the share at or above the goal stays roughly the same under monitoring and declines to 20%–29% for unmonitored savers. This is not surprising because bunching should happen irrespective of the incentive. Also, as the window widens, one is adding in the treatment effect.

Because we find so little evidence of gaming, we believe that our results would still hold even in absence of financial incentives, but an experimental test is required to confirm this.

## 5.2. *Robustness of Our Results*

We now describe the results of two robustness exercises. First, we might be worried about measurement error: it is important to see that, in fact, savings were achieved and also that we can at least partially understand the source of the increased savings. Second, because we do have survey attrition in the sample, we show robustness of our results to corrections for that attrition.

In Supplemental Material Appendix G, we deal with the first concern and examine how the savers saved. We tackle this in two ways: first, using a detailed expenditure survey in the sixth month of our savings period and second, using a retrospective survey in our follow-up fifteen months after the savings period ended. Table G.1 Panel A presents the results of the first exercise. We find that being assigned a random monitor leads to noisy 9.6% decline in total expenditures ($p = 0.12$). In levels, with a 5% winsorizing to deal with outliers, there is a Rs. 534 decline in total expenditures ($p = 0.068$). We see, consistent with anecdotes and our retrospective survey, evidence of decline in festival expenditure (by Rs. 242, $p = 0.065$), decline in transportation expenditure (by Rs. 147, $p = 0.054$), and an increase in tea consumption (by Rs. 30, $p = 0.084$, which is a common drink to take on the job).

Panel B provides a more-powered view, albeit through a retrospective survey. Assignment to a random monitor corresponds to a claimed 7.2pp increase in labor supply on a base of 27.1% ($p = 0.037$), a 2.5pp increase in business profits on a base of 4.75% ($p = 0.12$), and 8.1pp ($p = 0.083$) reduction in unnecessary expenditures on a base of 21.5% for the non-monitored. Reassuringly, while it seems more work and better budgeting led to savings, there is no statistically significant increase in borrowing money from one's network, no reduction in transfers to others, and no borrowing to save.

Throughout the paper, we drop observations for which we do not have total savings information from our main total savings regressions. Recall from Table I that attrition is balanced across monitored and non-monitored savers. Nonetheless, our main regression estimates might be conservative if monitors disproportionately caused individuals with large savings balances to attrit from the study and those with small savings balances to remain, or they might be overstated if monitors caused the better (worse) savers to remain (attrit). Thus, we conduct an exercise using Lee bounds in Supplemental Material Appendix F.1.

Table F.2 presents the results. Looking at both the effect of having a random monitor and, conditional on the random monitoring sample, the effect of having a monitor with a high value of the model-based regressor generates lower bounds that are comparable to our main regression estimates (e.g., a lower bound of 0.31 on the value of having a high model-based regressor monitor as compared to a point estimate of 0.451). Our bounds are noisy: we have used female as a binary predictor of attrition (see Table F.1) to tighten the bound.

## 6. ENDOGENOUS MONITORS

### 6.1. *Endogenous Monitors Benchmark*

The previous results suggest that a social planner interested in maximizing savings could "optimize" the allocation of monitors to savers using the network as an input. However, such an allocation mechanism is likely infeasible for most real-world institutions. In many of the informal peer-based financial arrangements that are commonly found in developing countries, individuals endogenously sort into groups.[34] Thus, measuring savings under endogenous monitor choice is a useful policy-relevant benchmark. For obvious reasons, causally determining what drives choice is beyond the scope of the paper. That is neither our aim nor what we claim to measure. Our goal, rather, is to assess how well the community does when left to its own devices.

To measure this benchmark, we analyze the savings outcomes in the 30 villages with endogenous monitor choice. It is important to note that a priori, savings could be higher or lower in endogenous relative to random monitor assignment. On the one hand, savers might completely unwind all savings benefits of monitors. It may also be the case that some individuals feel constrained socially in their ability to choose their preferred monitor.[35] On the other hand, individuals might arrive at the "optimal" savings-maximizing allocation of savers to monitors. Thus, any outcome between "optimality" and full unwinding is feasible.

Table VIII presents the log total savings of participants in endogenous and random choice villages with and without monitors. In column (1), we include village fixed effects, so the estimated coefficients measure the effects of receiving a monitor relative to non-monitored savers in the same village. As before, being assigned a monitor in the random assignment villages leads to increased total savings. However, the savers who picked their own monitors in endogenous assignment villages save no more than the savers who were not assigned to receive a monitor in the same endogenous assignment villages (insignificant $-0.139$).

In column (2), we remove the village fixed effects, so that we can also compare non-monitored villagers in endogenous assignment villages to non-monitored villagers in random assignment villages. Even when we remove the fixed effects, being assigned a random monitor in the random assignment villages increases total savings. Further, being able to choose one's own monitor in the endogenous assignment villages also leads to a 29.8% $(\exp(0.376 - 0.115))$ increase in total savings relative to a non-monitored villager in a random assignment village. However, note that we also find that the non-monitored villagers in endogenous assignment villages save considerably more than their non-monitored counterparts in the random assignment villages. Of note, the savings increases of those

[34]Examples include Stickk.com, which asks individuals to choose a "referee" to monitor their progress toward a goal. Also, MFIs, ROSCAs, and SHGs often involve endogenous group formation. We also note that a financial institution in India approached us to implement a similar program in one of their urban branches.

[35]For example, low caste individuals may feel uncomfortable choosing high caste monitors. Similarly, low income day laborers may feel that they are not entitled to pick important people in the village.

TABLE VIII

RANDOM VERSUS ENDOGENOUS MONITORS[a]

| Variables | (1) Log Total Savings | (2) Log Total Savings | (3) Reached Goal | (4) Reached Goal | (5) Savings in Target A/c | (6) Savings in Target A/c |
|---|---|---|---|---|---|---|
| Random Assignment Village × Monitor | 0.281 | 0.286 | 0.062 | 0.063 | 322.135 | 302.048 |
| | (0.149) | (0.146) | (0.030) | (0.030) | (137.604) | (133.419) |
| | [0.031, 0.531] | [0.042, 0.531] | [0.012, 0.111] | [0.012, 0.114] | [92.187, 552.083] | [79.092, 525.003] |
| Endogenous Assignment Village × Monitor | −0.139 | −0.115 | 0.058 | 0.062 | 195.237 | 202.514 |
| | (0.157) | (0.140) | (0.023) | (0.022) | (208.812) | (212.886) |
| | [−0.401, 0.122] | [−0.350, 0.120] | [0.021, 0.096] | [0.026, 0.098] | [−153.708, 544.182] | [−153.239, 558.267] |
| Endogenous Assignment Village | | 0.376 | | −0.005 | | 252.516 |
| | | (0.204) | | (0.033) | | (207.196) |
| | | [0.036, 0.717] | | [−0.060, 0.050] | | [−93.728, 598.760] |
| Observations | 1,061 | 1,061 | 1,298 | 1,298 | 1,298 | 1,298 |
| Fixed Effects | Village | None | Village | None | Village | None |
| Controls | Saver | Saver | Saver | Saver | Saver | Saver |
| Depvar Mean | 8.026 | 8.026 | 0.125 | 0.125 | 678.483 | 678.483 |
| Endog. Assign. × Monitor + Endog. Assign. = 0 | | 0.164 | | 0.076 | | 0.006 |
| Endog. Assign. × Monitor + Endog. Assign. = Random Assign. × Monitor | | 0.866 | | 0.840 | | 0.387 |

[a]Table reports effects of receiving a monitor in random versus endogenous allocation villages. Total savings is the amount saved across all savings vehicles—the target account and any other account—by the saver. Reached Goal is a dummy for if the saver reached their saving goal. Savings in Target A/c is the amount saved in the target account. Sample includes savers who responded to our questionnaire. Controls include the following saver characteristics: savings goal, saver centrality, age, marital status, number of children, preference for bank or post office account, whether the individual has a bank or post office account at baseline, caste, elite status, number of rooms in the home, and type of electrical connection. Columns (1), (3), and (5) include village fixed effects, while columns (2), (4), and (6) do not. Standard errors (clustered at the village level) are reported in parentheses. 90% confidence intervals are reported in brackets.

randomly assigned to monitors, able to choose their monitors in endogenous assignment villages, and non-monitored in endogenous assignment villages are not distinguishable from each other. When we turn to the effects of random and endogenous monitors on goal attainment (columns (3) and (4)), there we see that again, endogenous and random monitors generate similar levels of goal attainment. However, we do not observe a goal attainment spillover onto the non-monitored savers in the endogenous villages. Columns (5) and (6) repeat the exercise for the target savings account.

Why the non-monitored savers save more in endogenous choice villages is an interesting question. Given that we did not expect such an outcome, we can only speculate as to the exact mechanism. We think that it is most likely that endogenous choice led to an increase in the number of conversations in the village about savings. For example, we observe that savers ran into their endogenously chosen monitors more than their randomly assigned monitors (5.1 versus 4.0 times per fortnight—difference significant at the 1% level).[36] These conversations may have motivated some of the non-monitored savers to save more. In Supplemental Material Appendix H, we conduct an exercise to test for spillovers from monitored to non-monitored savers.[37] We do find evidence that the monitors, and especially the high centrality monitors, affect the savings of the friends of their savers.[38] Better understanding these spillovers is an interesting direction for future research.

Thus, the results show that the community does well implementing our informal peer-based financial product on its own. So even if it is not feasible to optimize the matching of savers to monitors, the community can still benefit from an endogenous implementation.

### 6.2. *Exploring Choice*

While our experiment was not designed to fully unpack monitor choice, we end by exploring one specific aspect of choice. To do this, we extend our signaling model in Supplemental Material Appendix I.1 to develop intuitions for which individuals might pick more central and proximate monitors and where choice order may matter. The model extension also provides a framework for thinking about who might self-select into the experiment.

We consider agents of both heterogeneous quality and centrality, who first decide whether or not to opt into the experiment, knowing that if they do, they will be assigned to BC, random monitoring, or endogenous monitoring. In the endogenous treatment, agents also choose their monitors from the available pool, and agents know this. Our extended model shows the complexities in modeling choice in the endogenous treatment, even abstracting away from the likely forces that may also affect choice (whether people are amicable, forgiving, encouraging, etc.). We focus on one specific subtlety—that $H$ types have an incentive to enter our experiment and choose highly central monitors in the endogenous treatment, if they are available, whereas not only do $L$ types want to choose low centrality monitors to avoid being revealed, but highly central $L$ types may not even opt into the experiment.[39] Therefore, when we look at choice, the theory suggests that

[36]In contrast, planned meetings between savers and their monitors changed by a much smaller, insignificant amount (2.5 vs. 2.3, $p$-value 0.4).

[37]We also show that allowing for such spillovers does not change our main results in the random allocation villages. The logic is that having a friend who is randomly assigned a monitor, conditional on participating, is orthogonal to receiving a monitor oneself or that monitor's location in the network.

[38]Finally, it is also possible that the ability for savers to choose their own monitors increases the desirability of the program and the buy-in of the village.

[39]Consistent with this observation, saver centrality is correlated with total savings in our data, conditional on savings goal, though this may be a spurious correlation for a variety of other obvious reasons.

high centrality savers should be more likely to choose better monitors. Further, if high centrality monitors are scarce, there should be a relationship between choosing early and choosing more central monitors, but only among the highly central. This is indeed what we show in Supplemental Material Appendix I.1.

## 7. CONCLUSION

Reputations matter. Our subjects enunciate this both in direct surveys and through their economic decisions. When it is known that information about their savings is transmitted to others in the community, participants increase their savings in meaningful enough amounts that they are better able to mitigate shocks.

But reputation in *whose* eyes also matters, and the social network provides an apt lens to examine this. Individuals benefit from impressing their monitors because, in the future, they might need to rely either on the monitor directly or on parties who have come to learn about them from the monitor. This motive is undoubtedly asymmetric. Certain sets of people interact more or less frequently with others. A network perspective puts discipline on how reputational stakes may vary with the position of one's monitor in the community.

Our field experiment is carefully designed to quantify impacts on a measurable and important behavior—savings. Further, we collect evidence pertaining to how the households saved, whether the savings had follow-on benefits, and whether the savings accumulation persisted. We make a methodological contribution toward measuring reputation by tracking the information flow itself from the monitors to other members of the community.

The findings of this experiment speak to a general discussion in development economics about the nature and role of social sanctions that may support informal financial institutions. In our simplified setup, a benefit or sanction is simply getting a good or bad name as demonstrated by one's effort to accumulate savings. We show that monitors do pass on information, savers desire to be perceived as responsible, and savers make payments into the monitored accounts. These findings document empirically the forces alluded to in the literature (e.g., Besley and Coate (1995), Munshi (2014)). Further, because the degree of information that is passed on is correlated in a convincing manner with the network position of the monitor, the identity of who in a community can leverage this reputational motive is an important factor when considering whether networks can sustain good behavior.

The forces described here are likely to operate in settings where a primary barrier to savings accumulation is a failure of responsibility (including things like time inconsistency or inattention). Further, it is essential that social reputation carries weight, or more generally, network relationships with community members carry great weight. These conditions are likely to hold in exactly the types of communities that are able to sustain RoSCAS, SHGs, VSLAs, and other types of informal financial structures. However, contexts where the primary barriers include lack of access to financial institutions and urban settings where individuals often have access to markets or a number of alternatives beyond their networks, are unlikely to be ones where our mechanism would be strongly at play.

## APPENDIX A: FORMAL MODEL

### A.1. *Description of the Signaling Environment*

There are $n$ agents and each has privately known type $\theta_i \in \{H, L\}$. We assume each agent's type is drawn, i.i.d., to be $\theta_i = H$ with probability $1/2$ and $\theta_i = L$ with probability

1/2. This data generating process is commonly known. A subset $m \ll n$ of agents participate in a signaling game and are called "savers." Another subset of $m$ agents are matched bijectively to the set of savers, and we call them "monitors." The remaining $n-2m$ agents have no designated role. We assume that whether a given individual participates in the signaling game is private at the outset, so the fact that $i$ is a saver is known only to her monitor $j$, and vice versa.

The game proceeds in two phases. Each saver $i$ decides whether to take a potentially costly action ($s_i = 1$) at cost $c_{\theta_i}$ or not ($s_i = 0$) at no cost. We assume high types find the action less costly, $c_H < c_L$. Each agent also has a productivity $A_{\theta_i}$ with $A_H > A_L$. That is, high types are productive whereas low types are not. We assume $c_L > \bar{q}(A_H - A_L) > c_H$, for some $\bar{q} > 0$.

Let $\mathbf{P} = \{p_{kl} : k, l \in \{1, \dots, n\}\}$ be a matrix of probabilities and is commonly known by all agents. $p_{kl}$ will denote the probability that $k$ meets $l$ in a given phase of the game *and* that an interaction between them materializes. In phase 1, this means that they meet and $k$ passes information to $l$. In phase 2, this means that they meet and $l$ offers $k$ a payoff opportunity. The foundations of $\mathbf{P}$ will be described below.

In phase 1, when selecting $s_i$, every saver $i$ knows that her choice will be observed by a unique monitor $j$. The information about the saver's action is communicated to others in the community as follows. Independently with probability $p_{kl}$, $k$ meets $l$ and passes a piece of information to $l$. If $j$ is the monitor and $k$ is some party, this means that $k$ becomes informed of $s_i$ with probability $p_{jk}$ and this is independently drawn across all $k \neq j, i$. Therefore, in phase 1, after $i$ chooses $s_i$, every agent $k \neq j, i$ in the network is informed, independently, about $i$'s decision with probability $p_{jk}$ and receives no information with probability $1 - p_{jk}$. Let $r_k$ denote an indicator for whether $k$ received this information.

In the second phase of the game, with probability $p_{kl}$, $k$ meets $l$ and $l$ offers $k$ a payoff which depends on $l$'s belief about $k$'s productivity. So if $i$ is the saver and $k$ is the third party, this means that, with probability $p_{ik}$, they meet and $k$ offers $i$ a payoff that depends on $k$'s inference of $i$'s private type. Thus, every saver $i$ then meets each agent $k$, independently, with probability $p_{ik}$. If the meeting happens, agent $k$ then offers agent $i$ a payoff which corresponds to her belief about $i$'s productivity: $\mathrm{E}_k[A_{\theta_i} | s_i, r_k]$. Note that this will depend on whether $k$ has received information about whether $i$ participates in the game as well as the choice if she does participate.

The expected payoff of saver $i$ with monitor $j$ for choice $s_i$ is given by

$$U(s_i) := \sum_k p_{ik} p_{jk} \mathrm{E}_k[A_{\theta_i} | s_i, r_k = 1] + \sum_k (1 - p_{jk}) p_{ik} \mathrm{E}_k[A_{\theta_i} | s_i, r_k = 0] - s_i c_{\theta_i}.$$

It will be useful to define $q_{ij} := \sum_k p_{ik} p_{jk}$. We assume the probabilities are such that $q_{ij} < \bar{q}$ for any $i, j$.[40]

### A.2. *Relationship to Our Experiment*

The mapping to our experiment is as follows. In the first phase, a potential saver decides whether to save a low ($s_i = 0$) or a high ($s_i = 1$) amount. This decision sends a signal to the monitor as to whether the saver is responsible or not. The type $\theta_i$ represents responsibility.

[40]Note that this utility formulation allows for $k = i$, and thus, for the saver to meet herself and give herself a payoff. We show in Auxiliary Appendix Q that the contribution of these self-meeting terms to the overall utility is negligible, vanishing as the size of the graph gets large, for economically relevant networks. Thus, the conclusions look identical to first order when omitting the terms from the utility function.

The idea is that it is relatively costlier for irresponsible individuals to overcome their time inconsistency, temptations, or inattention and accrue high savings. Single crossing means more responsible people are better able to overcome their time inconsistency, given by $c_H < c_L$.

In the second phase, the saver has a future interaction with a fellow community member from the village network. The saver again meets a community member through the graph. The returns to this interaction can depend on whether this community member knows about the saver's "type" via the signaling process in period 1. This is because responsible people are more productive: so $A_H > A_L$. If the member of the community believes the individual is irresponsible, she offers the saver a lower return and the saver has less to gain in the second period since she receives the low wage. Otherwise, if the member believes that she is responsible, she offers the saver the high wage. However, it is possible that the community member simply has not heard any rumor about the individual's type whatsoever, in which case the saver receives a pooled wage, which we normalize to $\frac{A_H+A_L}{2}$. In a typical signaling model, the costly signal is always transmitted to the market. Here the signal is more likely to be transmitted to someone who will give the saver a payoff if the saver is more likely to meet this person and this person is more likely to have heard directly/indirectly about the information via the monitor, captured by the $p_{ik}p_{jk}$ term. Below, we describe a model based on the social network that provides a concrete specification for $p_{ik}$ and $p_{jk}$.

## A.3. *The Network Environment and Interpretation of $q_{ij}$*

### A.3.1. *The Network Environment*

In order to model the network environment, it is necessary to first define what we mean by a network interaction. Our perspective is informed by our data. 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. For instance, 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, etc. Therefore, we interpret the network as a medium through which we can parameterize interactions; an individual is more likely to meet with direct contacts, is less likely to meet friends of friends, and is even less likely to interact with friends of friends of friends, and so on.

In our interaction model, agents interact in an undirected, unweighted graph with associated adjacency matrix $\mathbf{A}$. Each element $A_{kl}$ is a binary indicator for a strong social or financial relationship between nodes $k$ and $l$.

We suppose meetings happen stochastically, with a node $k$ meeting node $l$ when $A_{kl} = 1$ with some fixed probability $\theta$, node $k$ meeting some node $m$ (not $k$'s neighbor) who is $l$'s neighbor with probability $\theta^2$, and so on. The model is parsimonious, depending on the single parameter $\theta$.

The total expected number of times that $kl$ meet can therefore be given by

$$\mathbf{M}(\mathbf{A}, \theta) := \left[\sum_{t=1}^{T}(\theta\mathbf{A})^t\right],$$

which is a matrix with entries $M_{kl}$. Observe that the right-hand side counts the expected number of times node $k$ encounters node $l$ and takes into account the potentially numerous paths between $k$ and $l$ of lengths $t = 1, \dots, T$.

Assume every time a meeting happens between some nodes $j$ and $k$, a successful event happens with sufficiently small probability $\beta$, so that $O(\beta^2)$ terms below can be ignored. In phase 1, this event is simply that $j$ passes $k$ a piece of information. For example, if $j$ is a monitor and $k$ is a third party, this information is the savings $s_i$ of saver $i$. Let $x_{jk}$ be the random variable that counts the number of meetings between $j$ and $k$ mediated by the above process. Then $\mathrm{E}[x_{jk}] = M_{ij}$. Then the ex ante probability that $k$ learns the information $s_i$ from $j$ is given by, using the binomial approximation,

$$\mathrm{E}\big[1-(1-\beta)^{x_{jk}}\big] = \beta\mathrm{E}[x_{jk}] + O\big(\beta^2\big) = \beta M_{jk} + O\big(\beta^2\big).$$

So we have that

$$\mathrm{P}\big(k \text{ learns } j\text{'s information}\big) = \beta M_{jk} + O\big(\beta^2\big).$$

Turning to phase 2, here the event after the meeting between two agents $i$ and $k$ is that with some small enough probability, again $\beta$ for parsimony, $k$ offers $i$ a payoff which depends on $k$'s inference about $i$'s type from whatever he hears in phase 1. Then again we can define $x_{ik}$ to be the random variable that counts the number of meetings between $j$ and $k$ mediated by the above process. Of course, again the ex ante probability that $k$ encounters $i$ and has a payoff to offer is given by

$$\mathrm{E}\big[1-(1-\beta)^{x_{ik}}\big] = \beta\mathrm{E}[x_{ik}] + O\big(\beta^2\big) = \beta M_{ik} + O\big(\beta^2\big).$$

So we have that

$$\mathrm{P}(k \text{ offers a payoff to } i) = \beta M_{jk} + O\big(\beta^2\big).$$

Because, in this simplified setup, the expressions for meeting and then interacting in each phase are identical, let us define

$$p_{kl} := \mathrm{P}\big(l \text{ learns } k\text{'s information}\big) = \mathrm{P}(k \text{ offers a payoff to } l).$$

This is the probability that $k$ and $l$ meet and interact in a particular stage of the game. Therefore, we have

$$p_{kl} \propto \left[\sum_{t=1}^{T}(\theta\mathbf{A})^t\right]_{kl},$$

where the constant of proportionality is $\beta$.[41]

Let $\mathbf{P}$ denote the full matrix with entries $p_{ij}$. This formulation equips us with expressions for the key probabilities in the signaling model: $p_{jk}$, the probability that monitor $j$ transmits his observation of $i$'s savings to third party $k$, and $p_{ik}$, the probability that saver $i$ encounters third party $k$ for a payoff.

Given this framework for interactions on a network, observe that certain households will be more central than others (reaching directly or indirectly more individuals). As will become clear, this has nothing to do with the strategic interactions themselves but rather only with the assumed physical interactions on the network.

[41]While somewhat awkward, note that this makes $p_{kk}$ be proportional to the number of $\theta$-weighted paths from one to oneself. As described in Auxiliary Appendix Q, this term is negligible (of lesser order) in all our calculations, so we keep it this way to allow us to work with standard network statistics.

It is useful to formally define

$$DC(\mathbf{A}, \theta) := \sum_{t=1}^{T} (\theta\mathbf{A})^t \cdot 1$$

as the *diffusion centrality*. This is the notion of centrality that emerges from our simple model of interaction on a network. Note that $DC(\mathbf{A}, \theta)$ is related to the commonly studied eigenvector centrality in the following way. Let $\lambda_1$ be the first (maximal) eigenvalue corresponding to the matrix $\mathbf{A}$ and let $e(\mathbf{A})$ be the corresponding eigenvector. Taking the limit as $T \to \infty$ with $\theta \geq \frac{1}{\lambda_1}$ leads to a vector $\lim_{T\to\infty} \frac{\sum_{t=1}^{T}(\theta\mathbf{A})^t \cdot 1}{\sum_{t=1}^{T}(\theta\lambda_1)^t} = e(\mathbf{A})$, the eigenvector centrality.[42]

We also consider measures of social proximity between nodes $i$ and $j$ in the graph. Note that if two agents are closer in the graph, the rows of $\mathbf{P}$ corresponding to those agents must be more correlated. This is because if $i$ and $j$ are neighbors, any path to a given $k$ of length $\ell$ from $i$ to $k$ must be either of length $\ell + 1$, $\ell$, or $\ell - 1$ from $j$ to $k$. This notion of proximity, $\operatorname{cov}(p_{i\cdot}, p_{j\cdot})$, is key for our proofs, below. Note, additionally, that this measure is correlated with the more standard inverse-distance measure used in the network literature, $1/d(i, j)$, where $d(i, j)$ is the shortest path between $i$ and $j$.

This certainly is not the only sensible way to model interactions, and different models would generate predictions for slightly different notions of centrality. However, the core idea would be the same. 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.

### A.3.2. *Computing $q_{ij}$*

We now show that we can write $q_{ij}$, where $i \neq j$, in terms of the centralities and proximity, as defined above:

$$\begin{aligned} q_{ij} &= \sum_k p_{ik} \cdot p_{jk} = n \cdot \operatorname{cov}(p_{i\cdot}, p_{j\cdot}) + \frac{1}{n} \sum_k p_{jk} \sum_k p_{ik} \\ &= n \cdot \operatorname{cov}(p_{i\cdot}, p_{j\cdot}) + \frac{1}{n} DC_j \cdot DC_i. \end{aligned}$$

## A.4. *Analysis*

Let us define $\widehat{q} := \frac{c_H}{A_H - A_L}$. By assumption, note that $\widehat{q} < 1$. Let $r_k$ denote a dummy variable for whether a third party $k$ in the network hears a report about $i$ (that $i$ was a saver, the amount $s_i$, and the identity of monitor $j$).

LEMMA A.1: *Under the maintained assumptions, there is a Perfect Bayesian equilibrium such that*

[42]This is the same modeling structure used in Banerjee et al. (2013). For a more general discussion about eigenvector centrality in network economic models, see Jackson (2008). See also DeMarzo, Vayanos, and Zwiebel (2003), Golub and Jackson (2010, 2012), and Hagen and Kahng (1992).

(1) *for H types*,

$$s_i = \begin{cases} 1 & \text{if } q_{ij} \geq \widehat{q}, \\ 0 & \text{otherwise}. \end{cases}$$

(2) *for L types*, $s_i = 0$ *irrespective of* $q_{ij}$.
*This PBE is supported by beliefs by each third party* $k$,

- $\mathrm{P}(\theta_i = H | r_k = 0) = \frac{1}{2}$,
- $\mathrm{P}(\theta_i = H | s_i = 1, r_k = 1, q_{ij} \geq \widehat{q}) = 1$ *and* $\mathrm{P}(\theta_i = H | s_i = 0, r_k = 1, q_{ij} \geq \widehat{q}) = 0$,
- $\mathrm{P}(\theta_i = H | s_i = 1, r_k = 1, q_{ij} < \widehat{q}) = x$ *and* $\mathrm{P}(\theta_i = H | s_i = 0, r_k = 1, q_{ij} < \widehat{q}) = \frac{1}{2}$, *for any* $x \in [0, 1]$.

PROOF: Consider a saver of type $\theta_i$. She chooses $s_i = 1$ if and only if

$$\begin{aligned} &\sum_k p_{ik} p_{jk} \mathrm{E}_k[A_{\theta_i} | s_i = 1, r_k = 1] + \sum_k (1 - p_{jk}) p_{ik} \mathrm{E}_k[A_{\theta_i} | s_i = 1, r_k = 0] \\ &\quad > \sum_k p_{ik} p_{jk} \mathrm{E}_k[A_{\theta_i} | s_i = 0, r_k = 1] + \sum_k (1 - p_{jk}) p_{ik} \mathrm{E}_k[A_{\theta_i} | s_i = 0, r_k = 0]. \end{aligned}$$

Since $\mathrm{E}_k[A_{\theta_i} | s_i = 1, r_k = 0] = \mathrm{E}_k[A_{\theta_i} | s_i = 0, r_k = 0]$, she chooses $s_i = 1$ if and only if

$$\sum_k p_{ik} p_{jk} \mathrm{E}_k[A_{\theta_i} | s_i = 1, r_k = 1] - c_{\theta_i} > \sum_k p_{ik} p_{jk} \mathrm{E}_k[A_{\theta_i} | s_i = 0, r_k = 1].$$

Let $\mu_k(s_i) := \mathrm{P}(\theta_i = H | s_i, r_k = 1)$ be defined as the posterior that $k$ has about $i$'s type given she has received a report and observes $s_i$. Notice that because network structure is common knowledge and all third parties $k$ hold the same prior about $\theta_i$, we can write $\mathrm{E}_k[A_{\theta_i} | s_i, r_k] = \mathrm{E}[A_{\theta_i} | s_i, r_k]$ and $\mu_k(s) = \mu(s)$ for any $k$ with $r_k = 1$.

We can then write

$$\mathrm{E}[A_{\theta_i} | s_i, r_k = 1] = \mu(s_i)(A_H - A_L) + A_L \quad \text{and} \quad \mathrm{E}[A_{\theta_i} | s_i, r_k = 0] = \frac{A_H + A_L}{2}.$$

From the above, $s_i = 1$ if and only if

$$\big(\mu(1) - \mu(0)\big)(A_H - A_L) \sum_k p_{ik} p_{jk} \geq c_{\theta_i},$$

which can be written, since $q_{ij} = \sum_k p_{ik} p_{jk}$, as $q_{ij} \geq \frac{c_{\theta_i}}{\Delta\mu \Delta A_\theta}$.

In this equilibrium, if $q_{ij} > \widehat{q}$, then $\mu(s_i = 1) = 1$ and $\mu(s_i = 0) = 0$. (Recall that network structure is common knowledge, so $q_{ij}$ is known to $k$ when making this calculation.) Observe that

$$\frac{c_L}{(A_H - A_L)(1 - 0)} > \bar{q} > q_{ij} \geq \widehat{q} = \frac{c_H}{(A_H - A_L)(1 - 0)}.$$

Therefore, no $L$-type finds it worthwhile to try to secure a reputation gain by investing in $s_i = 1$, because $c_L$ is too high.

Meanwhile, if $q_{ij} < \widehat{q}$, then

$$q_{ij} < \widehat{q} = \frac{c_H}{(A_H - A_L)(1 - 0)} < \frac{c_L}{A_H - A_L}$$

and therefore neither type finds it worthwhile, even for the maximal reputation gain, $(1-0)$, to have $s_i = 1$. In this case, $\mu(s_i = 0) = \frac{1}{2}$ and $\mu(s_i = 1)$ can take any value since even the maximal increase in reputation (1-0) would not make it worthwhile. *Q.E.D.*

This result immediately implies the following.

PROPOSITION A.2: *Under the maintained assumptions and the Perfect Bayesian equilibrium described in Lemma A.1*, $\mathrm{P}(s_i = 1|q_{ij})$ *is a* (*weakly*) *monotonically increasing function in* $q_{ij}$. *Consequently*, $\mathrm{P}(s_i = 1|q_{ij})$ *must be* (*weakly*) *monotonically increasing in both social proximity*, $\mathrm{cov}(p_{i\cdot}, p_{j\cdot})$, *and monitor centrality*, $DC_j$.

PROOF: By random assignment of monitors $j$ to savers $i$, and by orthogonality of private type $\theta_i$ to network position, it follows that half of those with a sufficiently high $q_{ij}$ must be of low type, so $\mathrm{P}(s_i = 1|q_{ij}) = \frac{1}{2}$ if $q_{ij} \geq \widehat{q}$. Meanwhile, if $q_{ij} < \widehat{q}$, $\mathrm{P}(s_i|q_{ij}) = 0$. As shown above, since

$$q_{ij} = n \cdot \mathrm{cov}(p_{i\cdot}, p_{j\cdot}) + \frac{1}{n} DC_j \cdot DC_i,$$

ceteris paribus an increase in either $\mathrm{cov}(p_{i\cdot}, p_{j\cdot})$ or $DC_j$ increases $q_{ij}$. *Q.E.D.*

## A.5. *Interpretation of Results*

Our framework suggests that we should focus our empirical analysis on two network features: centrality, in particular eigenvector centrality which follows directly from the model, and proximity. We have the following predictions: (1) as $q_{ij}$ increases, a greater proportion of savers should be saving high amounts; (2) as monitor centrality increases, a greater proportion of savers should be saving high amounts; (3) as saver–monitor proximity increases, a greater proportion of savers should be saving high amounts. These directly motivate regressions of savings on network position, as conducted in the paper. Figure A.1 presents an example where $\mathrm{cov}(p_{i\cdot}, p_{j\cdot})$ is varied between saver $i$ and monitor $j$ but $DC_j$ is held fixed. This is to give an idea of how to envision holding distance fixed as we vary centrality, or vice versa.

A reasonable question to raise is whether individuals already know each others' types, especially those who are socially close. We think that there is significant scope for learning about even a close individual's type for several reasons. The first piece of evidence comes from our data. Fifteen months after our intervention, individuals were asked to rate 12 random subjects about whether the subjects reached their goals, as well as answer several questions concerning their level of responsibility. The respondents were no more likely to rate their unmonitored friends (who reached their goal throughout the experiment) as responsible as more distant individuals despite there being a positive correlation on average

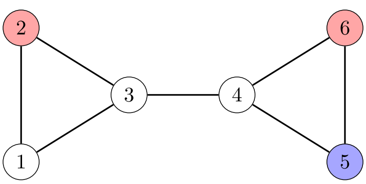


FIGURE A.1.—Let node 5 be the saver and let nodes 2 and 6 be potential monitors. This presents a situation where $DC_2 = DC_6$, by symmetry, but clearly $\mathrm{cov}(p_{5\cdot}, p_{2\cdot}) \neq \mathrm{cov}(p_{5\cdot}, p_{6\cdot})$.

between responsibility and goal reaching. If anything, they were slightly worse at rating their friends. Second, the work of Alatas, Banerjee, Chandrasekhar, Hanna, and Olken (2016) examines how well individuals are able to rank others' wealth in their communities. While individuals are slightly better at ranking those to whom they are socially closer, the error rates are still very high, indicating highly imperfect local information. Third, we have anecdotal evidence from our subjects that indicates that there is scope to build reputation among even their friends, neighbors, or important individuals in their communities. Thus, while it is entirely possible ex ante for the scope for reputation building to be lower among the socially proximate (due to heterogeneous priors), our own prior is that this is unlikely to be the case.

## REFERENCES


ABDULKADIROĞLU, A., AND T. SÖNMEZ (2003): "School Choice: A Mechanism Design Approach," *American Economic Review*, 93 (3), 729–747. [187]

ALATAS, V., A. BANERJEE, A. G. CHANDRASEKHAR, R. HANNA, AND B. A. OLKEN (2016): "Network Structure and the Aggregation of Information: Theory and Evidence from Indonesia," *American Economic Review*, 106 (7), 1663–1704. [215]

ASHRAF, N., D. KARLAN, AND W. YIN (2006): "Tying Odysseus to the Mast: Evidence From a Commitment Savings Product in the Philippines," *Quarterly Journal of Economics*, 121 (2), 635–672. [178]

BALLESTER, C., A. CALVÓ-ARMENGOL, AND Y. ZENOU (2006): "Who's Who in Networks. Wanted: The Key Player," *Econometrica*, 74 (5), 1403–1417. [177]

BANERJEE, A., A. G. CHANDRASEKHAR, E. DUFLO, AND M. JACKSON (2013): "The Diffusion of Microfinance," *Science*, 341 (6144). [175,177,180,212]

BANERJEE, A., A. G. CHANDRASEKHAR, E. DUFLO, AND M. O. JACKSON (Forthcoming): "Using Gossip to Spread Information: Theory and Evidence from Two Randomized Controlled Trials" *Review of Economic Studies*. [177,180,181,183,188,189]

BEAMAN, L., D. KARLAN, AND B. THUYSBAERT (2014): "Saving for a (not so) Rainy Day: A Randomized Evaluation of Savings Groups in Mali," NBER Working Paper. [175,178]

BELLONI, A., V. CHERNOZHUKOV, AND C. HANSEN (2014a): "High-Dimensional Methods and Inference on Structural and Treatment Effects," *The Journal of Economic Perspectives*, 28 (2), 29–50. [192,196]

——— (2014b): "Inference on Treatment Effects After Selection Among High-Dimensional Controls," *The Review of Economic Studies*, 81 (2), 608–650. [192]

BESHEARS, J., J. J. CHOI, D. LAIBSON, B. MADRIAN, AND J. SAKONG (2011): "Self Control and Liquidity: How to Design a Commitment Contract," RAND Corporation Working Paper. [178]

BESHEARS, J. L., J. J. CHOI, D. LAIBSON, B. C. MADRIAN, AND K. L. MILKMAN (2015): "The Effect of Providing Peer Information on Retirement Savings Decisions," *Journal of Finance*, 70 (3), 1161–1201. [175]

BESLEY, T., AND S. COATE (1995): "Group Lending, Repayment Incentives and Social Collateral," *Journal of Development Economics*, 46 (1), 1–18. [176,178,208]

BESLEY, T., S. COATE, AND G. LOURY (1993): "The Economics of Rotating Savings and Credit Associations," *The American Economic Review*, 83 (4), 792–810. [178]

BREZA, E. (2014): "Peer Effects and Loan Repayment: Evidence From the Krishna Default Criris," Working Paper. [176,178]

BREZA, E., A. G. CHANDRASEKHAR (2019): "Supplement to 'Social Networks, Reputation, and Commitment: Evidence From a Savings Monitors Experiment'," *Econometrica Supplemental Material*, 87, https://doi.org/10.3982/ECTA13683. [179]

BREZA, E., A. G. CHANDRASEKHAR, AND H. LARREGUY (2015): "Network Centrality and Institutional Design: Evidence From a Lab Experiment in the Field," NBER Working Paper 20309. [176,180]

BRUNE, L., X. GINÉ, J. GOLDBERG, AND D. YANG (2016): "Facilitating Savings for Agriculture: Field Experimental Evidence From Malawi," *Economic Development and Cultural Change*, 64 (2), 187–220. [178]

BRYAN, G., D. KARLAN, AND S. NELSON (2010): "Commitment Devices," *Annu. Rev. Econ.*, 2, 671–698. [181]

BRYAN, G., D. KARLAN, AND J. ZINMAN (2015): "Referrals: Peer Screening and Enforcement in a Consumer Credit Field Experiment," *American Economic Journal: Microeconomics*, 7 (3), 174–204. [175,178]

BURSZTYN, L., F. EDERER, B. FERMAN, AND N. YUCHTMAN (2014): "Understanding Mechanisms Underlying Peer Effects: Evidence From a Field Experiment on Financial Decisions," *Econometrica*, 82 (4), 1273–1301. [175]

CAI, J., A. DE JANVRY, AND E. SADOULET (2015): "Social Networks and the Decision to Insure," *American Economic Journal: Applied Economics*, 7 (2), 81–108. [175]

CARROLL, G. (2012): "When Are Local Incentive Constraints Sufficient?" *Econometrica*, 80 (2), 661–686. [187]

CHANDRASEKHAR, A. G., C. KINNAN, AND H. LARREGUY (2018): "Social Networks as Contract Enforcement: Evidence from a Lab Experiment in the Field," *American Economic Journal: Applied Economics*, 10 (4), 43–78. [180]

CHANDRASEKHAR, A. G., H. LARREGUY, AND J. XANDRI (2012): "Testing Models of Social Learning on Networks: Evidence From a Framed Field Experiment," Working Paper. [180]

DEMARZO, P. M., D. VAYANOS, AND J. ZWIEBEL (2003): "Persuasion Bias, Social Influence, and Unidimensional Opinions," *The Quarterly Journal of Economics*, 118 (3), 909–968. [212]

DUPAS, P., AND J. ROBINSON (2013a): "Savings Constraints and Microenterprise Development: Evidence From a Field Experiment in Kenya," *American Economic Journal: Applied Economics*, 5 (1), 163–192. [178]

——— (2013b): "Why Don't the Poor Save More? Evidence From Health Savings Experiments," *The American Economic Review*, 103 (4), 1138–1171. [178,187,188]

FEIGENBERG, B., E. FIELD, AND R. PANDE (2013): "The Economic Returns to Social Interaction: Experimental Evidence From Microfinance," *Review of Economic Studies*, 80 (4), 1459–1483. [176]

GINÉ, X., AND D. KARLAN (2006): "Group versus Individual Liability: Short and Long Term Evidence From Philippine Microcredit Lending Groups," *Journal of Development Economics*, 107, 65–83. [178]

GOLUB, B., AND M. O. JACKSON (2010): "Naive Learning in Social Networks and the Wisdom of Crowds," *American Economic Journal: Microeconomics*, 2 (1), 112–149. [177,212]

——— (2012): "How Homophily Affects the Speed of Learning and Best-Response Dynamics," *The Quarterly Journal of Economics*, 127 (3), 1287–1338. [212]

HAGEN, L., AND A. KAHNG (1992): "New Spectral Methods for Ratio Cut Partitioning and Clustering," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 11 (9), 1074–1085. [212]

INFLATION.EU (2017): "Historic Inflation India—CPI Inflation," http://www.inflation.eu/inflation-rates/india/historic-inflation/cpi-inflation-india.aspx. [181]

JACK, W., AND T. SURI (2014): "Risk Sharing and Transactions Costs: Evidence From Kenya's Mobile Money Revolution," *The American Economic Review*, 104 (1), 183–223. [175]

JACKSON, M., T. BARRAQUER, AND X. TAN (2012): "Social Capital and Social Quilts: Network Patterns of Favor Exchange", *The American Economic Review*, 102 (5), 1857–1897. [180]

JACKSON, M. O. (2008): *Social and Economic Networks*. Princeton University Press. [212]

JAKIELA, P., AND O. OZIER (2016): "Does Africa Need a Rotten Kin Theorem? Experimental Evidence From Village Economies," *The Review of Economic Studies*, 83 (1), 231–268. [178,183]

KARLAN, D. (2007): "Social Connections and Group Banking," *The Economic Journal*, 117 (517), F52–F84. [178]

KARLAN, D., M. MCCONNELL, S. MULLAINATHAN, AND J. ZINMAN (2016): "Getting to the Top of Mind: How Reminders Increase Saving," *Management Science*, 62 (12), 3393–3672. [175,178]

KAST, F., AND D. POMERANZ (2018): "Saving more in groups: Field experimental evidence from Chile," *Journal of Development Economics*, 133, 275–294. [175,178]

KATZ, E., AND P. F. LAZARSFELD (1970): "Personal Influence, the Part Played by People in the Flow of Mass Communications," Transaction Publishers. [177]

MUNSHI, K. (2014): "Community Networks and the Process of Development," *The Journal of Economic Perspectives*, 28 (4), 49–76. [175,208]

PATHAK, P. A., AND J. SETHURAMAN (2011): "Lotteries in Student Assignment: An Equivalence Result," *Theoretical Economics*, 6 (1), 1–17. [187]

PLATTEAU, J.-P. (2000): "Egalitarian Norms and Economic Development," in *Institutions, Social Norms, and Economic Development*. Harwood Academic, 189–240. [178]

PRINA, S. (2015): "Banking the Poor via Savings Accounts: Evidence From a Field Experiment," *Journal of Development Economics*, 115, 16–31. [178]

SCHANER, S. (2018): "The Persistent Power of Behavioral Change: Long-Run Impacts of Temporary Savings Subsidies for the Poor," *American Economic Journal: Applied Economics*, 10 (3), 67–100. [178]

SPENCE, M. (1973): "Job Market Signaling," *The Quarterly Journal of Economics*, 87 (3), 355–374. [190]

THALER, R. H., AND S. BENARTZI (2004): "Save More Tomorrow: Using Behavioral Economics to Increase Employee Saving," *Journal of political Economy*, 112 (51), S164–S187. [178]

*Co-editor Joel Sobel handled this manuscript.*

*Manuscript received 29 July, 2015; final version accepted 9 July, 2018; available online 26 July, 2018.*