

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

EMILY BREZA[†] AND ARUN G. CHANDRASEKHAR[‡]

ABSTRACT. We study whether individuals save more when information about their savings is shared with another village member (a “monitor”). We focus on whether the monitor’s effectiveness depends on her network position. Central monitors may be better able to disseminate information, and more proximate monitors may pass information to individuals who interact with the saver frequently. In 30 villages, we randomly assign monitors. Average monitors increase savings by 35%. A one-standard deviation more central monitor increases savings by 14%; increasing proximity from social distance three to two increases savings by 16%. The increased savings persist over a year after the intervention’s end, and monitored savers better respond to shocks. In 30 other villages, savers choose their monitors. Proximate and central monitors are chosen.

Information flows. 63% of monitors tell others about the saver’s progress. 15 months later, others know more about the saver’s progress and believe she is responsible if the saver was assigned a more central monitor.

JEL CLASSIFICATION CODES: D14, D83, L14, O16, Z13

KEYWORDS: Commitment, Reputation, Savings, Social Networks

1. INTRODUCTION

Peer effects have been found in a wide 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

Date: July 2015.

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 Matthew O. 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 seminar and conference participants at the Santa Fe Institute, SITE, the European Conference on Household Finance (Einaudi), the NBER Corporate Finance meeting, Duke, IFPRI, the Stanford RAIN Seminar, Cambridge, USC, the AEA Annual Meeting for their insightful comments, and the Harvard Conference on Information Transmission in Networks. 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 the Stanford Institute for Innovation in Developing Economies for generous support. Finally, we both thank the National Science Foundation GRFP for support while in graduate school.

This version: July 2015. Previous title “Savings Monitors.”

[†]Columbia Business School. Email: ebreza@columbia.edu.

[‡]Stanford University and NBER. Email: arungc@stanford.edu.

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 et al. \(2013\)](#), [Bryan et al. \(2012\)](#), [Kast et al. \(2012\)](#), [Beaman et al. \(2014\)](#), [Beshears et al. \(2015\)](#), [Munshi \(2014\)](#), [Karlan et al. \(2010\)](#), [Bursztyn et al. \(2013\)](#), [Banerjee et al. \(2013\)](#), for example). 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 peer effect 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

To explore this, we conduct an experiment across 60 villages in rural Karnataka, India, where we have complete network data for almost all households in each and every village. We focus on savings as our application and 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, they naturally 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, using the 30 villages in which savers choose their monitors, we investigate whether agents do in fact choose more central and proximate monitors and whether they save more.

¹The present field experiment was inspired by the earlier lab-in-the-field study [Breza et al. \(2015\)](#). There, we explore the efficiency of transfers in non-anonymous sender-receiver investment games with a third-party observer. Villagers are 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, and we are interested in how the network position of the third party influences the efficiency of the transaction. The very fact that a more central third-party observes the transaction increases efficiency significantly (as seen from comparing (2) to (1) for more versus less central third parties). This suggests that what we call the information effect, just the third party observing the interaction in the investment game, is stronger for more central agents. Further, the beneficial effect of centrality is even greater when the third party is also given an observable punishment technology.

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 – i.e., 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.²

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 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 extremely detailed network data collected in part by the authors in previous work (Banerjee et al. (2014)). This household-level network data comprises 12 dimensions of interactions across *all* potential pairs of households in each of the 60 study villages.⁴ Again, because the social network is an extremely complicated object, we focus our analysis on the two moments of the network data that emerge from our model: 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

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

²At the same time, it could be the case that individuals have heterogeneous priors about the saver based on position in the network.

³There 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 people who have learned of their shortcomings.

⁴The network data we use here is essentially complete, with data surveyed from 90% of households in each village. Thus, the probability of not surveying either member of a linked pair is 0.1^2 . This means that there is data on essentially $1 - 0.1^2 \approx 0.99$ share of possible links in the network. In the networks literature, we use what is called the OR network, drawing a link between two households if either named the other as a partner.

⁵See Katz and Lazarsfeld (1970), DeMarzo et al. (2003), Ballester et al. (2006), Banerjee et al. (2013), and Golub and Jackson (2010).

more is a sign of responsibility. 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) such as RoSCAs, SHGs and MFIs 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 et al. (2014), Besley et al. (1993), Karlan (2007), Giné and Karlan (2006), Bryan et al. (2012), and Breza (2014)).

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 et al. (2006), Brune et al. (2013), Karlan et al. (2010), Thaler and Benartzi (2004), and Beshears et al. (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), Dupas and Robinson (2013a), Brune et al. (2013), Prina (2013), Schaner (2014), and Kast and Pomeranz (2014)). 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 five main components. First, using the data from the 30 villages in which we randomly assign monitors to savers, we establish that receiving an arbitrary monitor increases total savings across formal and informal vehicles by 35%. 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%.

Second, we make use of unique 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 that saver's randomly-assigned monitor.

Third, we provide evidence that our intervention caused lasting and positive 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 shocks where they did not have enough cash to

respond which we measure following Dupas and Robinson (2013a).⁶ 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 towards savings, rather than wasteful expenditures or leisure, result in improved risk-coping, and yield persistent increases in savings, likely held as buffer stocks.⁷

Fourth, in the 30 villages in which savers could choose their monitors, we find that individuals pick more central monitors as well as more proximate monitors relative to random assignment and to a benchmark where savers simply pick their closest available friend. This suggests that savers are more sophisticated when making their monitor choices. We also find that these savers perform at least as well as the savers who are randomly assigned to monitors.

Fifth and finally, we find that non-monitored savers in the endogenous assignment villages save substantially more than the non-monitored savers in the random assignment villages. In fact, the non-monitored savers in the 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 the behavior of the friends of the saver. We show that indeed savers save more when their friends are randomly assigned to a more central monitor (in the random assignment villages). Given that the monitors are more central on average in the endogenous villages, the commensurate increase in spillovers can explain the observed “catch up” among the non-monitored savers.

In short, our study points to the idea that reputations matter, and that 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.

More generally, our findings are relevant for understanding many financial arrangements and products that enlist social capital. In particular, our study speaks to the large literature on how rotating savings and credit associations (RoSCAs) function. RoSCAs are quite complicated objects to study, as they involve many agents making simultaneous decisions that become observable to the rest of the group. Notice that the structure of our experiment is (by design) like a simplified slice of a RoSCA, an individual makes payments into a pot over time, and those payments are made observable to another individual. Here, we address

⁶That is, the incidence of reporting a situation where they had an above-median number of shocks (including personal health shocks, bovine health shocks, or other unexpected household expenditure and did not have enough cash on hand to cover the cost) drops for individuals assigned to a monitor.

⁷The most common savings goal purposes listed by savers at baseline were unforeseen expenditures and emergencies.

one piece of the RoSCA repayment incentive – how the opinion formed about whether an individual is able to be financially responsible motivates the agent.

The remainder of the paper is organized as follows. Section 2 describes 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 supplemental evidence to help to interpret these results. Section 6 discusses the villages with endogenous monitor assignment, Section 7 touches on policy implications, and Section 8 concludes.

2. DATA AND EXPERIMENTAL DESIGN

2.1. Setting and Data. The requirements for our study are threefold: 1) the presence of financial institutions where study participants can save; 2) detailed social networks data; and 3) experimental variation in the nature of the saver-monitor relationship. Our final sample includes 3,000 participants from 60 villages in rural Karnataka, India that meet our criteria.

We chose to set our experiment in 60 villages from rural Karnataka, India that coincide with the social network and demographic data set previously collected in other work, in part by the authors (and also described in Banerjee et al. (2014)). In our field experiment, we match participants to this unique data set.

Banerjee et al. (2014) collected network data from 89.14% of the 16,476 households living in those villages. The data concerns “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 pattern of interactions across households. Using this data we construct one network for each village at the household level and indicate that 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.⁸ Network-level summary statistics are displayed in Appendix Table B.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. We use the following notation: we have a collection of R villages, indexed by r and N_r individuals per village. Every village is associated with a social network G_r .

⁸The main idea here is that individuals can communicate if they interact in any of the 12 ways. This is the network of potential communications and a good description of which individuals are likely to interact (in one of several ways in a day to day sense) with others.

$G_r = (V_r, E_r)$ is a graph consisting of vertices $V_r = \{1, \dots, n\}$ and edges E_r where $ij \in E_r$ means that households i and j are linked. Following the extensive work on this data, we assume that this is an undirected, unweighted network: households are linked or are not linked and $ij \in E_r \iff ji \in E_r$ (see, e.g., Banerjee et al. (2013), Jackson et al. (2010), Chandrasekhar et al. (2013) for discussion). We use $\mathbf{A}_r := \mathbf{A}(G_r)$ to denote the adjacency matrix. This is a matrix with $A_{ij} = 1 \{ij \in E_r\}$.

Moreover, the survey data includes information about caste, elite status, measures of house amenities and the GPS coordinates of respondent’s homes. In the local cultural 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 study network effects conditional on these numerous observables.

2.2. Experimental Design. Figures 1 and 2 pictorially 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 questions on historic savings behavior, income sources and desire to save.

Our baseline data shows that the use of these branches is quite low. Only a quarter of households had a bank or post office account at baseline. Figure B.1 shows the baseline intensity of use of available savings vehicles separately for male and female savers. On average, potential savers keep a large fraction of their liquid savings in cash stored inside the house. For women, one third of savings is kept in self help groups (SHGs), while ROSCAs and insurance policies (generally through Life Insurance Corporation of India) are popular among men. Only 10% of savings are kept in formal bank accounts. 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 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 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.⁹

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)¹⁰:

T1 : Non-monitored treatment (Randomization at the individual level)¹¹

T2 : Peer monitoring with random matching (Randomization to receive a monitor at the individual level, randomization of monitor assignment at the village level)

T3 : Peer monitoring with endogenous matching (Randomization to receive a monitor at the individual level, randomization of monitor assignment at the village level)

All individuals who attend the village meeting are assisted in account opening by our survey team. In each village, we identify one bank branch and the local post office branch to offer as choices to the savers. Savers are allowed to choose either option. Savers who already have bank or post office accounts are offered the chance to open another account if they wish. The local post office branches are generally within a 3km walk of each village, and 35% of our study villages have a post office branch within the village boundary. We select bank branches that satisfy several criteria: within 5km of the village, offer “no-frills” savings accounts,¹² and agree to expedite our savings applications and process them in bulk.¹³ We offer the post office choice because women often feel uncomfortable traveling to bank branches but feel much more comfortable transacting with the local post master. On the other hand, some individuals greatly prefer bank accounts because those accounts make it easier to obtain bank credit in the future. We help savers to assemble all of the necessary paper work and “know your client” (KYC) identification documents for account opening and submit the applications in bulk.

Both the bank and the postal savings accounts have very similar product characteristics. First, in each type of account, the minimum balance is very small, typically Rs. 100 (~\$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%. “No frills” bank and post office accounts have no formal commitment features, and participants in the study are allowed to withdraw freely from the accounts at any time. We stress that in

⁹If the potential savers are not interested in savings, then they do not participate in the goal elicitation and they are not invited to the village meeting. Our sample frame for randomization is the set of savers attending the village meeting.

¹⁰Let T0 denote the pure control treatment.

¹¹We sometimes also refer to this treatment as the Business Correspondent treatment (BC). This is because the individual level treatment resembles a financial institution already in use in India called *business correspondents*. In this institution 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. We discuss this further in Section 7.

¹²“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).

¹³The 5km distance restriction meant that we were not able to work with only one bank, and instead opened accounts at branches of six different banks.

our study, the only source of commitment comes from the presence of a monitor in T2 and T3. Thus, our monitor treatments introduce an “informal commitment” in the parlance of Bryan et al. (2010).

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, both deposits and withdrawals. Entries in the passbooks are stamped with an official seal by branch personnel and cannot be forged. 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.¹⁵ We should note that in no treatment do our surveyors collect deposits on behalf of the savers.

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 to savers. In this treatment, individuals are allowed to choose their monitor from the pool of all potential monitors attending the village meeting. We only allowed one saver per monitor, so we randomized the order in which savers could choose. Again, 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 1 presents summary statistics for the sample that attended the village meeting and also shows baseline differences between T1, T2, and T3.

¹⁴We 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).

¹⁵Some participants even report that these visits are very motivating.

¹⁶Monitors 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.

Figure 4 presents the histogram of savings goals, censoring the top 5%.¹⁷ There are a few large outliers (maximum goal Rs. 26,000), so the mean of Rs. 1838 shrinks to Rs. 1650 when we trim 1% outliers. In all specifications of our key results we drop the top 1% of savings goal observations. While Rs. 600 may seem small on face value, it is equivalent to 3-5% of household income for the poorer members of the sample. It is also equal to the amount that could be saved if each household member saved instead of drinking one cup of tea each day. Ananth et al. (2007) suggest that some individuals with high returns to savings may nonetheless have a hard time saving even small amounts.

For our endogenous matching treatment, we chose to implement random serial dictatorship (RSD). Here, 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, it is rather intuitive, 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. Does an individual select a more network-central monitor? Does an individual select a socially close monitor? Third, and perhaps most importantly, there is a deep matching literature establishing equivalence between RSD and various other matching schemes with trading which reach the core. Specifically, consider two allocation mechanisms in an environment of n savers and n monitors, and say each agent has strict preferences over the monitors. The first mechanism is RSD. The second is when the monitors are (for instance) randomly allocated to the various agents and then trading is allowed. In this (now) 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, where equivalence means the mechanisms give rise to the same *probability distribution over allocations* irrespective of the preferences of the agents. These results both characterize optimality of RSD as well as provide a justification for real-world use.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

¹⁷Note that the minimum goal is Rs. 600, the lower bound of allowed goals for participants.

¹⁸Naturally, the degree to which the environment studied in the allocation mechanism literature describes our environment to first order determines the degree to which this intuition is relevant for our setting. Of course, the result also suggests that at best, deviations from the espoused theoretical framework lead to welfare differences between RSD and TTC that are at best likely to be ambiguous as opposed to being ordered.

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 1 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 from the savings period about how the savers saved, how frequently they spoke with their monitors, and whether they made any financial transfers with their monitors. It also contains questions about shocks sustained by the respondent in the 15 months after the savings period in the style of Dupas and Robinson (2013b). 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. Appendix Table B.2 includes control group means for all of the variables from both of our endline surveys used in the analysis.

3. FRAMEWORK

Empirically studying how the position of the monitor in the social network affects the saver’s behavior can be challenging. To identify those aspects of the otherwise-complex network on which to focus our empirical analysis, we develop a simple two-period model. The model focuses on the signaling/reputation aspect and how that interacts with networks. For parsimony and in the tradition of classical signaling models, we abstract from the direct value of savings itself.²¹ In this model, we assume that savers can gain utility from interacting in the future, after the experiment, with individuals who have heard about their successes during the experiment. The novelty of the model comes from how we take an otherwise complicated object – agents playing a game on a network – and in a naturalistic way discover that we can focus on two aspects of the agents’ network position for our analysis.

¹⁹We 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.

²⁰Another 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 the virtually the same (and statistically indistinguishable) for monitored and un-monitored savers.

²¹Similarly, in the canonical example of Spence (1973) education only has a signaling value even though surely education also has a direct value (human capital investment) in the real world.

The model’s first period summarizes the entire 6-month savings period of our experiment. The saver first decides how much to save (taken as the total savings over the 6 months of the intervention).²² Recall that members of the research team communicate the saver’s progress to each monitor. Thus, in the model, the saver’s monitor is immediately informed of the amount saved. The monitor can then pass this information or any opinion she has formulated about the saver on to others.²³ The extent to which information flows through society is governed by the network structure (i.e., the monitor informs her friend with some probability, the friend informs another friend with some probability, etc.) In sum, at the end of the first period, a subset of the community is informed about the saver.

The second period of the model captures future interactions of the agent with the village following the end of the intervention.²⁴ Savers interact with members of the community, again in a process governed by the network structure and receive a payoff from each interaction. We model this naturalistically, so meeting a given friend may happen with some probability, meeting a friend’s friend happens with a lower probability (i.e., need to meet a friend and also need to be referred to the friend’s friend), as a stylistic way to parametrize future interactions through the network. Because these future interactions may take many forms and may occur under a wide set of circumstances, we model the payoffs in a reduced form way. The basic idea, of course, is that when encountering someone in the future, an agent’s payoff is weakly (if not strictly) higher if she has saved more. This may be because the successful saver demonstrates her capacity to keep her commitments, signals responsibility, feels less embarrassment or shame about her own shortcomings, etc.

Our shorthand to operationalize this idea is to use a standard [Spence \(1973\)](#) signaling model. By saving a higher amount the saver demonstrates to the monitor that she is responsible. After all, enrolling in the experiment in the first place indicates a demand for commitment to help accumulate savings. (In fact, it was actually a member of a village in a different study’s focus group who originally suggested the 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.)

For simplicity, individuals come in two kinds: responsible and irresponsible. A responsible individual can be interpreted as one who is able to overcome (with effort) her time inconsistency, temptations, or inattention issues. Of course, an individual’s responsibility

²²For simplicity we assume this is either a low or high amount, though certainly extending this to the continuum is straightforward and yields the same predictions.

²³As we show in Section 5.1, survey evidence documents that monitors do indeed pass such information to others and that, further, many savers have even heard, through back-channels, about others talking about the saver’s progress.

²⁴Again, surely agents also experience a direct benefit from savings. Our results in Section 5 indeed document such a benefit. We focus on the signaling aspect for parsimony.

matters for interactions across all walks of life.²⁵ For example, individuals in our study villages rely regularly on others for loans, jobs, insurance and information.

Thus, our model is an extremely simple signaling game situated in a network where individuals pass information or meet along the network according to a simple stochastic process. The signaling game is entirely standard; our contribution is in modeling the interpersonal interaction and information exchange to observe what pieces of agents' network positions influence the incentives in the signaling game. Individuals interact (or pass information) randomly through the network along edges, independently, with some probability.²⁶ With this stylized structure on interactions, the signaling model predicts that, *ceteris paribus*, being randomly assigned (a) a more central monitor leads to a greater share of savers saving a high amount and (b) a more socially proximate monitor leads to a greater share of savers saving a high amount.

3.1. Network Interactions.

The Physical Environment. It is useful to clarify exactly what we mean by a network interaction and how we define a central agent. 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 parametrize interactions; an individual is more likely to pass information to or meet with direct contacts, is less likely to pass information to or meet friends of friends, and is even less likely to interact with friends of friends of friends, and so on.

Notice that, broadly speaking, there are two main types of interactions relevant to our setting. First, an agent can pass information to another agent. We suppose that this happens stochastically within each period, with information traveling from node i to j (or from j to i) with some fixed probability (θ). Second, agents may meet others. Clearly individuals should be more likely to meet their friends than their friends of friends. A simple and plausible model for this type of interaction is to suppose that every node i travels to a neighboring node with probability θ , to a neighbor's neighbor with probability

²⁵In fact, survey data shows 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 Appendix Section A suggests that this influences how people will think of the saver in a labor market situation in the future.

²⁶Further, individuals have the same prior beliefs over all individuals that they may encounter in the future. We assume homogenous priors here to isolate the information diffusion component of the model. Surely it is possible that *ex ante* more proximate individuals have less to learn about one another. This could lead to non-monotone returns to signaling to close individuals.

θ^2 (if there is only one such path there), and so on. This parsimonious story motivates our model's physical environment.

In our model, agents in an undirected, unweighted graph with associated adjacency matrix \mathbf{A} interact. The model is simple, essentially depending on the single parameter θ which represents the probability of any two nodes in the network interacting either through information passing or a physical meeting. We use $p_{uv}(\mathbf{A}, \theta)$ to denote the probability that nodes u and v interact in a particular stage of the game. We micro-found this through a simple model of interaction on a network. All information passing (and meetings) along the network occur in the following manner. Given \mathbf{A} , there is some probability θ , that a given piece of information crosses any given link ij .²⁷ Let us define

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

where the constant of proportionality is not relevant for the model but ensures that the term is a probability. Observe that the right-hand side counts the expected number of times a piece of information starting from node u hits node v and takes into account the potentially numerous paths information may take between u and v . Let \mathbf{P} denote the full matrix with entries p_{uv} .

Key Quantities: Centrality and Proximity. Given a 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.

It is useful to formally define two quantities:

- (1) $DC(\mathbf{A}, \theta) := \sum_{t=1}^T (\theta \mathbf{A})^t \cdot \mathbf{1}$ is called the *diffusion centrality* with T rounds of communication. Taking the limit as $T \rightarrow \infty$ with $\theta \geq \frac{1}{\lambda_1}$ leads to a vector $\lim_{T \rightarrow \infty} \sum_{t=1}^T (\theta \mathbf{A})^t \cdot \mathbf{1} \propto e(\mathbf{A})$, the eigenvector centrality.²⁸ This object is a vector where $DC_i(\mathbf{A}, \theta)$ gives the expected number of times information starting from a given node i hits all other nodes in the graph, with stochasticity parametrized by θ . Note that this also – equivalently – gives the expected number of times that i interacts in total with all other nodes over T periods. This is the notion of centrality that emerges from our simple model of interaction on a network.
- (2) Let the *distance* between i and j in the graph be the length of the shortest path between them. Let *proximity* be $1/d(i, j)$. Then it is worth observing that if two agents are closer in the graph, the rows of \mathbf{P} corresponding to those agents must be

²⁷Assume that $\theta \geq \frac{1}{\lambda_1(\mathbf{A})}$.

²⁸This 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 et al. (2003), Golub and Jackson (2012), Golub and Jackson (2010), and Hagen and Kahng (1992).

more correlated. This is because if i and j are neighbors, any path to a given k of length ℓ from i to k must be either of length $\ell + 1$, ℓ or $\ell - 1$ from j to k .

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.

3.2. A Signaling Model of Peer Monitoring. Consider a two-period game played among network members. A saver i is either a “low” or “high” type. Low types (“irresponsible”) find it differentially more costly to accumulate the high levels of savings.²⁹ Let the cost of savings be denoted $c_H < c_L$ for the high or low type, respectively. In the future, when a responsible type engages in a productive activity, her productivity is y_H whereas it is y_L for an irresponsible type.

The timing is as follows:

- $t = 0$: saver i picks a level of savings $s_i \in \{s_H, s_L\}$.
- $t = 0.5$: monitor j of i observes the saver’s outcome.
- $t = 1$: monitor j of i diffuses information about s_i throughout the village. A given individual k in the village has heard this with probability $p_{jk}(\mathbf{A}, \theta)$. Let $r_{jk} \in \{0, 1\}$ denote j having successfully reported to k .³⁰
- $t = 2$: saver meets a random individual in the future with probability $p_{ik}(\mathbf{A}, \theta)$. Let $m_{ik} \in \{0, 1\}$ denote this meeting. This individual offers i a wage contract in a competitive labor market, where wage offers (given signals of type) are equal to productivity. A responsible individual has productivity y_H whereas an irresponsible individual has productivity y_L .

The interpretation is as follows. In the first period, a potential saver decides whether to save a high or low amount. This decision sends a signal to the monitor as to whether the saver is responsible or not. The idea is that it is relatively costlier for irresponsible individuals to overcome their time inconsistency, temptations or inattention and accrue high savings.

In the second period, 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. If the member of the community knows the individual is irresponsible, the saver has less to gain in the second period since

²⁹This cost may include present bias, temptations or inattention. In a different, but equivalent framing, low type might also encode individuals who naturally feel less shame when they don’t save very much and interact with informed agents in the future.

³⁰Technically, this should be a monotone function of the expected number of times information starting from j has hit k , since this term alone is not a probability. Our notation is a simple shorthand which preserves the qualitative aspects of the result. For related work see King (2015).

she receives the low wage. Otherwise, if the member knows that she is responsible, she receives 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 0.

The remainder of our analysis is focused on the separating equilibrium of the Spence signaling game (if parameters are in a range where a separating equilibrium exists). The justification for this comes from the intuitive criterion (Cho and Kreps (1987)). We are interested in how changes in individuals' relative network positions lead to transitions from only pooling equilibria to separating equilibria (applying the intuitive criterion).

The first result characterizes whether a separating equilibrium exists depends on these quantities.

Lemma 3.1. *A separating equilibrium exists if*

$$y_H - \frac{c_H}{\text{cov}(p_i, p_j) + DC_j(\mathbf{A}, \theta) \cdot DC_i(\mathbf{A}, \theta)} > y_L > y_H - \frac{c_L}{\text{cov}(p_i, p_j) + DC_j(\mathbf{A}, \theta) \cdot DC_i(\mathbf{A}, \theta)}.$$

Otherwise, only a pooling equilibrium exists in which all types pool on s_L .

Proof. Let q denote the probability that a randomly chosen member of the village has observed the signaling outcome. Here q is a reduced form for the three probabilities discussed above. Then it is straightforward to see that

$$y_H - \frac{c_H}{q} > y_L > y_H - \frac{c_L}{q}$$

corresponds to a separating equilibrium.³¹

Next we decompose q into its constituent parts in our model (derived below):

$$q = \text{cov}(p_i, p_j) + DC_j(\mathbf{A}) \cdot DC_i(\mathbf{A}).$$

The expected number of times that a given node k receives a signal sourced from j is given by p_{jk} . Integrating over all the k , we have³²

$$\left[\sum_t (\theta \mathbf{A})^t \cdot \mathbf{1} \right]_j = DC_j(\mathbf{A}, \theta).$$

Meanwhile, the probability that i will meet a given k is given by the analogous expression and therefore again we have

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

³¹This follows from the fact that

$$qy_H + (1-q)0 - c_H > qy_L + (1-q)0 > qy_H + (1-q)0 - c_L.$$

The first inequality ensures that the high type saves high and the second inequality ensures that the low type does not find it worthwhile to do so. This also clarifies why the normalization at 0 is innocuous.

³²In this derivation we ignore the constant of proportionality (or assume that it is 1) for parsimony. This has no consequence for the result.

Therefore, the payoffs given a monitor are

$$\mathbb{E}[r_{jk}m_{ik}|j]y_H - c_H > \mathbb{E}[r_{jk}m_{ik}|j]y_L > \mathbb{E}[r_{jk}m_{ik}|j]y_H - c_L.$$

It is clear that we can write

$$\begin{aligned} \mathbb{E}[r_{jk}m_{ik}|j] &= \sum_k p_{jk}(\mathbf{A}, \theta) p_{ik}(\mathbf{A}, \theta) \\ &= \text{cov}(p_{i\cdot}, p_{j\cdot}) + \sum_k p_{jk} \times \sum_k p_{jk} \\ &= \text{cov}(p_{i\cdot}, p_{j\cdot}) + DC_j(\mathbf{A}, \theta) \cdot DC_i(\mathbf{A}, \theta). \end{aligned}$$

The condition above ensures that a responsible type will invest in high savings and an irresponsible type will not find it worthwhile to represent herself as a responsible type, exploiting single crossing. Of course, by inspection one can see for low or high enough probabilities, it should be impossible to satisfy at least one if not both ICs. \square

Proposition 3.2. *The following describe how monitor effectiveness varies with network position.*

- (1) *Centrality of monitor:*
 - (a) *Holding $\text{cov}(p_{i\cdot}, p_{j\cdot})$ and DC_i fixed, for DC_j sufficiently high, there exists a separating equilibrium with high types attaining s_H .*
 - (b) *With the same parameters fixed, for DC_j sufficiently low, there remains only a pooling equilibrium on s_L .*
- (2) *Saver-monitor proximity:*
 - (a) *Holding DC_i and DC_j fixed, for $\text{cov}(p_{i\cdot}, p_{j\cdot})$ sufficiently high, there exists a separating equilibrium with high types attaining s_H .*
 - (b) *With the same parameters fixed, for $\text{cov}(p_{i\cdot}, p_{j\cdot})$ sufficiently low, there remains only a pooling equilibrium on s_L .*

It is worth noting that one can simply interpret $\text{cov}(p_{i\cdot}, p_{j\cdot})$ as the proximity between nodes i and j in the graph. The rows $p_{i\cdot}$ count the expected number of paths between i and some k , and as the distance between j and i decreases, the rows become more similar.

Thus, we have the following predictions for our monitors: (1) as monitor centrality increases, a greater proportion of savers should be saving high amounts; (2) as saver-monitor proximity increases, a greater proportion of savers should be saving high amounts.

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 own data. 15 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 between responsibility and goal reaching. If anything, they were slightly worse at rating their friends. Second, the work of [Alatas et al. \(2012\)](#) 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 indicate that there is scope, in their view, to build reputation among even their friends, neighbors or important individuals in their communities.³³ 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.

In short, we have described a simple physical process by which both meetings and information transmission occur. Our framework suggests that we should focus our empirical analysis on two features of the network, centrality, in particular eigenvector centrality which follows directly from the model, and proximity.

4. THE VALUE OF CENTRAL AND PROXIMATE MONITORS

4.1. Random Monitors. Our main results analyze how the centrality and proximity of randomly-assigned monitors influence savings behavior. Before turning to this, we briefly discuss the average impact of monitors relative to the baseline treatment bundle (non-monitored treatment). Table 2 presents the results, showing effects on the log of total savings across all accounts in columns 1-2 and the amount saved in the target account (which was made visible to the monitors) in columns 3-4. 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, whether saver has a bank or post office account at baseline, caste, elite status, number of rooms in the home and type of electrical connection.

We find that being randomly assigned to a monitor leads to a 0.3 log point increase in the total savings across all accounts. This corresponds to a 35% increase in savings across all savings vehicles of the households. Target account savings also increase by approximately Rs. 360. This again is a large effect, more than doubling the savings in the non-monitored group (mean 342.7).³⁴ In Appendix Table B.3, we show similar impacts on goal attainment (6.3pp increase, corresponding to an 80% increase in the likelihood relative to non-monitored savers).

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

³³See quotes from participant savers in Appendix Section A.

³⁴We have also checked whether being included in our experiment and receiving a bank account and bi-weekly reminder visits increases total savings. We find very small statistically insignificant effects. Analysis available upon request.

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 vs. non-monitored savers. The figure suggests that the average treatment effects are 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.

In sum, having a randomly-assigned monitor helps increase savings significantly.

4.2. Monitor Centrality and Proximity. We now turn to our main results: how monitor centrality and proximity influence saving behavior. Table 3 presents the results of regressions of log total savings across all accounts on monitor centrality, saver-monitor proximity, and a battery of controls.³⁵ In column 1 we condition on both centrality and proximity, though in columns 2-3 we present results separately by network characteristic. All regressions 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 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. We also control for the geographic distance between the homes of the saver and monitor.³⁶

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 0.13 log point increase in the log total savings, or a 14% increase – a large effect. Further, in Panel B of Figure 5, we explore the distributional effects of receiving a high centrality monitor vs. 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.

Turning to proximity, moving from a monitor of distance two to three leads to a 16% increase in the total savings across all accounts – again a large effect.³⁷

In Appendix Table B.4 we present the results of monitor centrality on the incidence of the saver reaching her goal. We find that a one standard deviation more central monitor corresponds to a 2.9pp increase in the likelihood of a goal being met, which is just under

³⁵Again, we stress that all accounts includes all formal and informal accounts, including savings “under the mattress.”

³⁶Results are also robust to controls for measures of the baseline savings behavior of the monitor, available on request.

³⁷In Appendix Table B.9 we measure the effect of being paired with a direct social connection and investigate whether this effect is stronger if the saver and monitor also have financial ties. We find large positive effects of receiving a monitor of social distance one, but no differential effects from receiving a monitor who also has financial ties. We define a link as having a financial component if the nodes report borrowing or lending small amounts of money or material goods to one another. In our sample, 27% of direct links have a financial component. Further, 86% of financial relationships are reciprocated – i.e., savers and monitor both borrow and lend to one another. The prevalence of reciprocated financial relationships is unsurprising given the strong risk sharing motives present in village India (Townsend (1994)). Thus we focus on symmetric financial connections rather than directed lending-only or borrowing-only relationships.

half the effect size of being assigned an average centrality monitor. Similarly going from a monitor of distance two to three results in a 2pp increase in the likelihood of a goal being met.

Thus we show that randomly assigning 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. Furthermore, magnitudes and significance are essentially the same even when entirely removing this bevy of characteristics (available upon request), which bolsters the idea that the effects are not driven by these characteristics.

In Appendix C we look at the same results as those in Table 3, but where we also allow centrality and proximity to vary by relationship type. We find that only the centrality and proximity in the union of all relationship types matters. This is natural as individuals could pass information across link types: for instance, to her coworker, who then tells a friend, who then tells his neighbor when borrowing rice about the information.

Finally, it is worth noting that given our focus on total savings, our results are likely not driven by mistaken accounting. Namely, individuals could have unwound the savings encouragement in the experiment by moving savings from some other account into the target account. (Recall that the monitor only was informed about the saver’s progress in the target account.) By looking at results across all savings vehicles, we indeed document increases in total savings inside and outside of the target accounts.

5. SUPPLEMENTAL EVIDENCE

Given our main result, that the monitoring treatment is more effective under a highly central and socially proximate monitor, we now provide supplemental evidence in support of the reputational mechanism of Section 3. We also describe how the savers managed to save during the study period and look for lasting impacts of our intervention after the end of the six month savings period. Finally we present tests that suggest that the monitor incentives are not likely to be driving our key results.

5.1. Perception Effects. We make use of novel supplemental data to support the idea that savings and reputation are linked and that some monitors are better able to affect reputations than others. 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 knowing through back-channels that the monitor passed information about their progress to others.³⁸

Moreover, we attempt to track this information flow from monitors to other members of the community. 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. First, we ask whether the subject knows if each of these 12 individuals reached their savings goal or not. Second, 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 here whether community members are better informed about the savings behaviors of others when the saver has a monitor and when that monitor is central. Further, we ask 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 4 presents the results of this exercise. We examine a regression of whether the interviewee knows correctly whether the saver reached her goal conditional 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 updated her beliefs about the general ability of the saver to reach her goals in the direction of the saver’s savings goal attainment (Columns 4-6).

The basic idea is that if the monitor is more central, it may be the case that a random interviewee in the village is more likely to have both heard of the saver’s success and also have a better view of the saver’s general responsibility.

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 indeed if a saver is randomly assigned a more central monitor, the respondent is more likely to know if the saver reached her goal and also is more likely to believe that the saver is good at meeting her goals.

We note that 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 in the future when a new opportunity arises (much like sending out a letter of recommendation).

It should go without saying that 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

³⁸This is particularly striking because it requires enough communication and leakage to occur such that savers hear gossip about themselves. We encourage the reader to reflect on how often this happens in their own lives.

other savers' responsibility profiles and savings habits in the experiment. The usual caveats about self-reported data certainly 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 provided here is (a) largely consistent with our framework/story, (b) mostly self-consistent, and (c) agrees with the anecdotal evidence provided in Appendix A. Further, 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.

5.2. Anecdotal Evidence. 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 Section A, 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.

5.3. How Did They Save? We can use our detailed survey data to explore how savers who were randomly selected to receive monitors managed to increase their savings balances over the duration of the experiment.

Panel A of Table 5 looks at detailed data on expenditures taken from the last month of the 6-month savings period. It presents regressions of expenditures on various categories on random assignment to monitors. As before, regressions include village fixed effects and the standard set of controls for savers. We find that being assigned a monitor leads to significant declines in expenditures on festivals (by Rs. 223) and transportation (by Rs. 154), and a minor increase in expenditure on tea (by Rs. 35.13).³⁹ A back of the envelope calculation shows that this can roughly explain the increased savings of those assigned a monitor: (Rs. 154 + Rs. 223 - Rs. 35) = Rs. 342) while the increase in total savings across all accounts for the last month is approximately Rs. 200.

Panel B uses retrospective data from the 15-month follow-up survey which asked what measures agents took to accumulate savings. We find that being assigned a random monitor led to a 7.12pp increase in the probability that a saver increased labor supply, a 2.02pp increase in savings from business profits, and a 7.9pp increase in the probability of cutting unnecessary expenditures. There is no evidence of changes to interpersonal transfers. Further, very few individuals report taking any sort of loan in order to save.

The results in Table 5 corroborate our overall monitor treatment effects and suggest that increases in savings mainly came at from natural places of slack in the household's budget,

³⁹The increase in tea consumption may be a direct result of the increase in labor supply.

namely leisure, unnecessary expenditures and festivals. That individuals are not borrowing to save in response to the monitor treatment adds to the evidence provided below that savers were likely not facing undue pressure from their monitors and in fact on average gained from genuine increases in savings.

5.4. Benefits from Saving. Given that our treatment increased total savings across all household accounts, a natural question to ask is whether we can detect any lasting benefits of the accrued savings caused by the monitoring treatments. This is a difficult question, so to address this, in our 15-month follow-up survey, we adopt the methods proposed by Dupas and Robinson (2013b). 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 be able to cope with a shock.⁴⁰

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 6 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 (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 0.199 decline in the total number of shocks (on a base of 1.769, column 1). 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. 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.⁴¹ The key point is that there are, nonetheless, tangible benefits of savings for situations like these.

Finally, in the last two columns of Panel B, we present the effects of the random monitor treatment on log savings balances 15 months after the intervention. Remarkably, the size of the increase in savings is as large as that reported in Table 2. This suggests that individuals are able to maintain their savings even after the monitors are no longer actively receiving information. Appendix Figure B.2 shows that the increases in savings across the distribution are still apparent 15 months later.

⁴⁰Note 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 conceivably have invested in shock mitigation. Like Dupas and Robinson (2013b), our analysis does not need to take a stand on this.

⁴¹We are not claiming that the gains in savings had large persistent health benefits, for example.

Taken together, these findings serve to show that there was truly an increase in savings (since they were better able to make purchases to cope with shocks) that persisted after the intervention and moreover show that there were important, real consequences of the increased savings.

5.5. 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 don't unwind across other accounts.

Turning to the monitor incentives, we do the following exercise. 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. Note that in terms of personal value to the saver, the incentives above and below the half goal should be smooth. So looking at 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, notice this is a mix of potential monitor incentives but also natural incentives to simply reach one's stated goal.

Table 7 presents the results. 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. Column 1 presents the group without a monitor (so by construction there is no monitor incentive), column 2 presents data from subjects with a random monitor and column 3 presents data from subjects with an endogenously chosen monitor. When we look at the half goal marker we find the share above 1/2 goal in the unmonitored group is 90%, whereas the numbers drop to 64% and 68% in the cases with monitors. 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 more likely to bunch on the right of the 1/2 goal mark despite not facing any monitor incentives by definition.⁴² 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.

Turning to the full goal, first note that by construction there is likely to be more bunching here (ex ante) simply because individuals set goals for themselves. We do observe that 60% of unmonitored savers fall at or within Rs. 150 above the goal, whereas this fraction is 86% and 70% for the monitored savers.

⁴²We find similar patterns when we zoom out to a window size of Rs. 100. Available upon request.

Finally, as seen in Figure B.3, we observe that savers do not rush to make deposits in the last two weeks of the savings period, as is might be the case if monitors were lending funds to savers. This supplements the evidence that there is no evidence of bunching or gaming in the neighborhood of 1/2 goal attainment.

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

6. ENDOGENOUS MONITORS

Given the effects on randomly assigned monitors, we now turn to the case of endogenous monitor selection.

6.1. Central and Proximate Monitors are Chosen. Recall that in every village there is a surplus of potential monitors. Therefore, we can ask if the network characteristics of the selected monitors in the endogenous-selection villages differ from those in the random-selection villages. We begin by asking whether savers in the endogenous-choice villages choose monitors that are both more central and socially close. Panels A and B of Figure 6 present histograms of the social distance between saver-monitor pairs in both the random and endogenous treatments. The figure clearly shows that individuals are more likely to choose friends and individuals of closer social proximity in the endogenous treatment relative to the proportions achieved by random matching. A regression of social proximity on treatment type indicates that endogenous saver-monitor pairs are 0.14 proximity units closer than random pairs. Similarly, Panels A and B of Figure 7 show that endogenously-chosen monitors are also more central in the network than randomly-matched monitors.

This evidence is consistent with the idea that agents may be sophisticated about the commitment value of having more central monitors. However, a natural alternative could be that individuals are simply choosing the closest available network connection. Because more central individuals are, by construction, closer to others, it could very well be the case that there is entirely no sophistication on the part of the choosers. To understand this, we run an algorithm mimicking this “closest neighbor” choice sequence.⁴³ Recall that savers were allowed to choose in random order. Given this order, the algorithm assumes that the saver picks the monitor of the closest social proximity from the available pool. If there are multiple potential monitors available of the same social proximity, then the algorithm draws the monitor uniformly at random. Once a monitor is chosen, that individual is no longer available in the choice set for subsequent savers. We present the resulting distributions of social distance and centrality in Panel C of Figures 6 and 7. Certainly the resulting distribution is more heavily weighted towards more central monitors than in the random monitor

⁴³We thank Attila Ambrus for this suggestion.

distribution. In fact, this does generate a distribution of centrality comparable to the endogenous choice villages. However, the marginal distribution of distances to partners does not match the observed distribution in the endogenous villages very well. Saver-monitor pairs are too close under the algorithm – nearly 75% of savers would have picked monitors at distance 1 with virtually all other savers picking monitors at distance 2. However, we find that in the observed endogenous distribution, 37% of savers pick monitors at distance 1, 50% at distance 2, and 12% at distance 3. The social distance distribution generated by the “closest neighbor” algorithm is easily rejected using a Kolmogorov-Smirnov test. Taken together, this suggests that agents are picking more central monitors than random, and in a way that can’t be explained simply by picking their closest available choice.

6.2. Effects of Endogenous Monitors. Given that when randomly assigned, central and proximate monitors are more effective, and when allowed to choose, savers pick more central and proximate monitors, it is natural to ask if savers perform better under endogenous choice. We explore the average treatment effects across all three treatments in Table 8. We measure the treatments effects separately for savers with randomly-chosen and endogenously-chosen monitors. In column 2, we include village fixed effects. Thus, the estimated coefficients measured the effects of receiving a monitor relative to savers in the same village. Here, we see that the savers who were able to pick their own monitors save no more across all accounts than the savers who were not assigned to receive a monitor (insignificant coefficient -0.0830).

However, when removing the village fixed effects, column 1 suggests that the negative coefficient can be explained by a large spillover effect on the control group. Relative to the non-monitored savers in the villages with random monitor assignment (omitted category), non-monitored savers in the endogenous choice villages increase savings by 0.35 log points. Moreover, the total savings effect is not statistically different from that of monitored savers in either monitor treatment.⁴⁴

Thus, we find that endogenous monitors are about as good as having a randomly assigned monitor and, more interestingly, that even the unmonitored individuals in endogenous villages do about as well.

6.3. Measuring and Calibrating Monitor Spillovers. Next, we use our experimental variation in monitor assignment in the random villages to look for spillovers from monitored to non-monitored savers. Non-monitored savers in both random- and endogenous-selection villages may be affected if their friends receive monitors and may experience larger spillovers if those monitors are especially effective.⁴⁵ The random variation in both the assignment of

⁴⁴Appendix Table B.7 investigates the effects of random and endogenous monitors on goal attainment. 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.

⁴⁵This could happen for a variety of reasons. For instance, “keeping up with the Joneses”, increased motivation to save, receiving reminders from the friend’s monitor, an overhearing more conversations about savings,

savers to treatment groups and of monitors to savers in the random selection villages allows us to measure such causal spillover effects.

Our experiment also contains random variation in the community-level composition of monitors. We have already shown that the monitors in the endogenous-selection villages are of higher centrality than those in the random-selection villages. Thus, we can ask if a change in monitor composition might lead to greater spillovers onto the non-monitored savers. Moreover, to what extent might this compositional difference also help to explain the large observed differences in the savings of non-monitored individuals between the endogenous- and random- selection villages displayed in Table 8? This calibration is the aim of this section.

We begin by looking for spillovers onto non-monitored savers in the villages with random monitor selection. We use the following regression specification:

$$(6.1) \quad y_{ir} = \alpha_r + \beta_1 \sum_j A_{ij,r} SM_j + \beta_2 \sum_j A_{ij,r} SM_j MC_j + \beta_3 \sum_j A_{ij,r} AttSaver_j + \delta' X_{ir} + \epsilon_{ir}.$$

This estimating equation allows the savings of non-monitored individuals to depend on having more friends randomly assigned to receive a monitor (SM) and having more friends randomly assigned to receive a central monitor (SM*MC).⁴⁶ All of this is conditional on the number of friends participating as savers in the experiment and fixed effects for the number of friends. The standard set of controls is included in X .⁴⁷ When we move from the random-selection to the endogenous-selection villages, the distribution of most of the explanatory variables is held constant by virtue of randomization. The only variable that differs is MC. Thus, we are particularly interested in β_2 .

Table 9 presents the results estimating Equation 6.1. We find that a one standard deviation increase in the sum of monitor centralities received by one's friends corresponds a 1.3 log point increase in total savings across all accounts.

We can now measure the expected change in log total savings for the non-monitored savers when moving from random assignment to endogenous assignment villages. The incremental change in the sum of the monitor centralities of one's friends is 0.036. Multiplying by the regression coefficient (10.24) we find that the predicted change is 0.368 log points. Note

etc. Our model in Section 3.2 abstracts from this to focus on the signaling story, just like it abstracts from the direct value of savings itself.

⁴⁶More generally, all agents – un-monitored and monitored – may face this type of spillover effect. Also note that the variation identifying the peer effect is orthogonal to the treatment status of the saver. In Appendix Table B.8, we show that, unsurprisingly, we can replicate our main results while accounting for these spillovers in our main regressions.

⁴⁷Recall that the standard controls include savings goal, gender, age, marital status, widow status, caste, elite status, material measures of wealth, whether the saver had a pre-existing bank or PO account, preference for bank or PO account during the savings period, and village fixed effects.

that this explains essentially the entire gap of 0.35 log points from Table 8. Furthermore, the standard errors on the computation are such that we cannot reject equality (the null that we can explain the entire gap). Spillovers alone can explain the pattern of “catch-up” that we observe in the endogenous selection villages.

To summarize, we perform an accounting exercise in the endogenous villages to see whether the better composition of more central monitors in the endogenous villages should naturally have more spillovers onto the unmonitored savers and therefore account for their enhanced behavior in the endogenous villages. In order to estimate this spillover, we use the data from the random assignment villages. Because monitors are randomly assigned here, we can directly estimate the reduced form of how one’s friends’ monitor centralities influences one’s own savings behavior. And then we can ask if we compute the implied spillover under the distribution of centralities under endogenous monitor choice, whether this can explain why unmonitored savers do better in the endogenous villages.

In measuring these spillovers, we are conducting a very different type of analysis from the core exercise of the paper. The spillovers measured here likely bundle many different channels of influence including (but not limited to) “keeping up with the Joneses”, increased motivation to save, receiving reminders from the friend’s monitor, and overhearing more conversations about savings.⁴⁸ Our goal here is simply to use the reduced form evidence of spillovers to reconcile the large savings differences of non-monitored savers between random and endogenous villages. We make no attempt to unpack or identify any precise mechanism underlying this reduced form. Anecdotal evidence suggests that conversations between savers and monitors tend to take place in public and are likely to be overheard by the saver’s friends.

7. DISCUSSION: TOWARDS POLICY

While the aim of our experiment was not to test an already scalable and cost-effective financial product, we do believe that our experiment provides lessons for policy-makers and for financial institutions.

We show that on average, our peer monitoring treatment increases savings substantially and also results in follow-on economic benefits over the subsequent year. While this positive effect is even stronger when the “right” members of the community are placed in the role of monitor (i.e., central and proximate individuals), we also demonstrate that our endogenous villages also experience positive savings impacts, and that those impacts spill over onto non-monitored savers. This suggests that even if it is not practical to fully optimize the matching of savers to monitors, the community can still benefit from more decentralized product designs that can be low cost and low touch for an implementing organization.

⁴⁸This multitude of channels is also the reason why we do not try to estimate and instrument a more structured model of spillovers in the spirit of [Bramouille et al. \(2009\)](#).

The findings from our study can help to provide a window into the operation of one of the oldest and most ubiquitous informal financial products around the world, RoSCAs (rotating savings and credit associations).⁴⁹ Similar to our setting, RoSCAs can provide social commitment for individuals who are unable to save on their own (Gugerty (2007)). Though simple to operate in practice, RoSCAs are actually quite difficult to study empirically. They are typically composed of upwards of ten endogenously-selected individuals, greatly complicating any theoretical or econometric analysis. As a result, empirical work on whether and why RoSCAs work tends to be observational. Our paper sheds light on the group dynamic by isolating and randomizing participants into pairs. Further, our results are useful for understanding how group composition affects the strength of social enforcement.

Recent research has documented the beneficial impacts of RoSCA-like structures on both savings and other outcomes. In one of the first RCTs measuring the impacts of village savings and loan associations (VSLAs), Beaman et al. (2014) show that the technical assistance provided by the implementing NGO helps endogenously formed groups accumulate savings and lend those savings to one another. Similar to our effects on medium-run savings and shock mitigation, they show that the main economic impacts of their VSLA intervention are on food security, consumption smoothing and buffer stock savings.

Our paper is also related to Kast et al. (2012) which looks at two experiments. In the first, SHG members (established and maintained by a microlender) were motivated to encourage each other to save by making public commitments in front of the other SHG members. They found large effects but also found that the results of the second experiment, where SHG members received SMS-based reminders to save, could rationalize the effects of the first experiment. An interesting distinction with our setting is that in their's, the monitors were both co-borrowers with the savers and were savers themselves by construction.⁵⁰

One important policy consideration is the cost of implementing and scaling a peer monitoring product. Our specific treatments were implemented with research goals in mind, and were never meant to be profitable or scalable. However, we do think that there are many opportunities for financial institutions to reduce the costs of product delivery. One of our largest costs was personnel. In order for the research team to have more control over the implementation, we chose to send individuals to each village on a bi-weekly basis to meet the savers, physically verify the passbooks, and pass the relevant information on to the monitors. Many financial institutions in India already use the Business Correspondent (BC) model, in which agents of the bank travel to villages to provide direct in-home customer service. This includes account opening procedures and deposit-taking. One could easily

⁴⁹Other financial products popular in developing countries such as SHGs, VSLAs, and MFIs also borrow many of the social aspects of RoSCAs.

⁵⁰We should also note that the subjects were recruited from a different population in our study. The Kast et al. (2012) subjects were all borrowing in pre-existing groups while they were saving towards their goals.

imagine a small tweak to this model, where the BC could intermediate information to others in the village after his pre-specified appointments. Further, banks could use technologies such as SMS to implement a peer monitoring scheme.

The other main cost associated with our intervention was the incentive given to monitors. First, as discussed previously, we think that the incentives had negligible effects on savings outcomes. Second, we certainly did not attempt to “optimize” the size of these incentives.⁵¹ Nevertheless in the endogenous monitor case, the aggregate monitor incentives paid to participants correspond to a 6% semi-annual interest rate on all additional savings that were caused by our interventions, which – while not cheap – is not outlandish.⁵² Experimenting with the size of the incentives would likely yield significant cost reductions.

In sum, the research done here and in related work suggests that there are likely to be minimally invasive, low-cost policies that financial institutions could adopt to encourage savings. This suggests a cheap way forward exploiting the large benefits of peer effects. This also raises interesting policy questions regarding whether policymakers should encourage savings through SHGs or goal-setting through business correspondents?

8. CONCLUSION

Reputations matter. Our subjects enunciate this both verbally and through their economic decisions. When 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 from the monitor. This motive to impress is undoubtedly asymmetric in communities. Certain sets of people interact more or less frequently with others, and a network perspective puts discipline on thinking about 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 economically important behavior – savings. Further, we collect evidence pertaining to how the households managed to save, whether the savings had follow-on benefits, and whether the savings accumulation persisted into the future. We acknowledge the difficulty in isolating reputation and make a contribution by attempting to track the information flow itself from the monitors to other members of the community. The study’s results point to minimally

⁵¹In fact, we believe that the optimal incentive would be close to, if not equal to, zero.

⁵²To reach this 6% value, we first calculate the aggregate payments that we made to monitors in the Endogenous villages. We then calculate the excess savings across all savings vehicles that were caused by our treatments. We include both the direct effects of receiving a monitor on savings and also the spillovers onto non-monitored savers.

invasive, low-cost policies that financial institutions could easily adopt in order to encourage savings.

REFERENCES


- ABDULKADIROĞLU, A. AND T. SÖNMEZ (2003): "School choice: A mechanism design approach," *American economic review*, 729–747.
- ALATAS, V., A. BANERJEE, A. G. CHANDRASEKHAR, R. HANNA, AND B. A. OLKEN (2012): "Network structure and the aggregation of information: Theory and evidence from Indonesia," *NBER Working Paper 18351*.
- ANANTH, B., D. KARLAN, AND S. MULLAINATHAN (2007): "Microentrepreneurs and Their Money: Three Anomalies," *Working Paper*.
- 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, 635 – 72.
- BALLESTER, C., A. CALVÓ-ARMENGOL, AND Y. ZENOU (2006): "Who's who in networks. wanted: the key player," *Econometrica*, 74, 1403–1417.
- BANERJEE, A., A. G. CHANDRASEKHAR, E. DUFLO, AND M. JACKSON (2013): "The Diffusion of Microfinance," *Science*, 341 (6144).
- BANERJEE, A., A. G. CHANDRASEKHAR, E. DUFLO, AND M. O. JACKSON (2014): "Gossip: Identifying central individuals in a social network," *NBER Working Paper 20422*.
- 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*.
- 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*.
- 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, Forthcoming*.
- BESLEY, T. AND S. COATE (1995): "Group lending, repayment incentives and social collateral," *Journal of development economics*, 46, 1–18.
- BESLEY, T., S. COATE, AND G. LOURY (1993): "The economics of rotating savings and credit associations," *The American Economic Review*, 792–810.
- BRAMOULLE, Y., H. DJEBBARI, AND B. FORTIN (2009): "Identification of peer effects through social networks," *Journal of Econometrics*, 150, 41–55.
- BREZA, E. (2014): "Peer Effects and Loan Repayment: Evidence from the Krishna Default Crisis," *Working Paper*.
- 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*.
- BRUNE, L., X. GINÉ, J. GOLDBERG, AND D. YANG (2013): "Commitments to Save: A Field Experiment in Rural Malawi," *Working Paper*.

- BRYAN, G., D. KARLAN, AND S. NELSON (2010): “Commitment devices,” *Annu. Rev. Econ.*, 2, 671–698.
- BRYAN, G. T., D. KARLAN, AND J. ZINMAN (2012): “Referrals: Peer Screening and Enforcement in a Consumer Credit Field Experiment,” *NBER Working Paper*.
- BURSZTYN, L., F. EDERER, B. FERMAN, AND N. YUCHTMAN (2013): “Understanding mechanisms underlying peer effects: Evidence from a field experiment on financial decisions,” *Econometrica*.
- CAI, J., A. DE JANVRY, AND E. SADOULET (2013): “Social Networks and the Decision to Insure,” *Working Paper*.
- CARROLL, G. (2012): “When Are Local Incentive Constraints Sufficient?” *Econometrica*, 80, 661–686.
- CHANDRASEKHAR, A. G., C. KINNAN, AND H. LARREGUY (2013): “Informal Insurance, Social Networks, and Savings Access: Evidence from a lab experiment in the field,” *Working Paper*.
- CHO, I.-K. AND D. M. KREPS (1987): “Signaling games and stable equilibria,” *The Quarterly Journal of Economics*, 179–221.
- DEMARZO, P. M., D. VAYANOS, AND J. ZWIEBEL (2003): “Persuasion bias, social influence, and unidimensional opinions,” *The Quarterly Journal of Economics*, 909–968.
- DUPAS, P. AND J. ROBINSON (2013a): “Savings Constraints and Microenterprise Development: Evidence from a Field Experiment in Kenya,” *American Economic Journal: Applied Economics*, 5, 163–92.
- (2013b): “Why Don’t the Poor Save More? Evidence from Health Savings Experiments,” *The American Economic Review*, 103, 1138–71.
- GINÉ, X. AND D. KARLAN (2006): “Group versus Individual Liability: Evidence from a Field Experiment in the Philippines,” Yale University Economic Growth Center working paper 940.
- GOLUB, B. AND M. O. JACKSON (2010): “Naive learning in social networks and the wisdom of crowds,” *American Economic Journal: Microeconomics*, 112–149.
- (2012): “How Homophily Affects the Speed of Learning and Best-Response Dynamics,” *The Quarterly Journal of Economics*, 127, 1287–1338.
- GUGERTY, M. K. (2007): “You Can’t Save Alone: Commitment in Rotating Savings and Credit Associations in Kenya,” *Economic Development and Cultural Change*, 55, 251–282.
- HAGEN, L. AND A. KAHNG (1992): “New spectral methods for ratio cut partitioning and clustering,” *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 11, 1074–1085.
- JACK, W. AND T. SURI (2014): “Risk Sharing and Transactions Costs: Evidence from Kenya’s Mobile Money Revolution,” *The American Economic Review*, 104, 183–223.
- JACKSON, M., T. BARRAQUER, AND X. TAN (2010): “Social Capital and Social Quilts: Network Patterns of Favor Exchange,” .

- JACKSON, M. O. (2008): *Social and Economic Networks*, Princeton University Press.
- KARLAN, D. (2007): "Social connections and group banking*," *The Economic Journal*, 117, F52–F84.
- KARLAN, D., M. MCCONNELL, S. MULLAINATHAN, AND J. ZINMAN (2010): "Getting to the Top of Mind: How Reminders Increase Saving," *NBER Working Paper*.
- KAST, F., S. MEIER, AND D. POMERANZ (2012): "Under-savers anonymous: Evidence on self-help groups and peer pressure as a savings commitment device," *NBER Working Paper*.
- KAST, F. AND D. POMERANZ (2014): "Saving More to Borrow Less: Experimental Evidence from Access to Formal Savings Accounts in Chile," *NBER Working Paper*.
- KATZ, E. AND P. F. LAZARSELD (1970): *Personal Influence, The part played by people in the flow of mass communications*, Transaction Publishers.
- KING, M. (2015): "Who can you trust? Reputation and cooperation in networks," *Mimeo*.
- MUNSHI, K. (2014): "Community networks and the process of development," *The Journal of Economic Perspectives*, 49–76.
- PATHAK, P. A. AND J. SETHURAMAN (2011): "Lotteries in student assignment: An equivalence result," *Theoretical Economics*, 6, 1–17.
- PRINA, S. (2013): "Banking the poor via savings accounts: Evidence from a field experiment," *Working Paper*.
- SCHANER, S. (2014): "The Persistent Power of Behavioral Change: Long-Run Impacts of Temporary Savings Subsidies for the Poor," *Working Paper*.
- SPENCE, M. (1973): "Job market signaling," *The quarterly journal of Economics*, 355–374.
- THALER, R. H. AND S. BENARTZI (2004): "Save More Tomorrow: Using behavioral economics to increase employee saving," *Journal of political Economy*, 112, S164–S187.
- TOWNSEND, R. M. (1994): "Risk and Insurance in Village India," *Econometrica*, 62, 539–591.

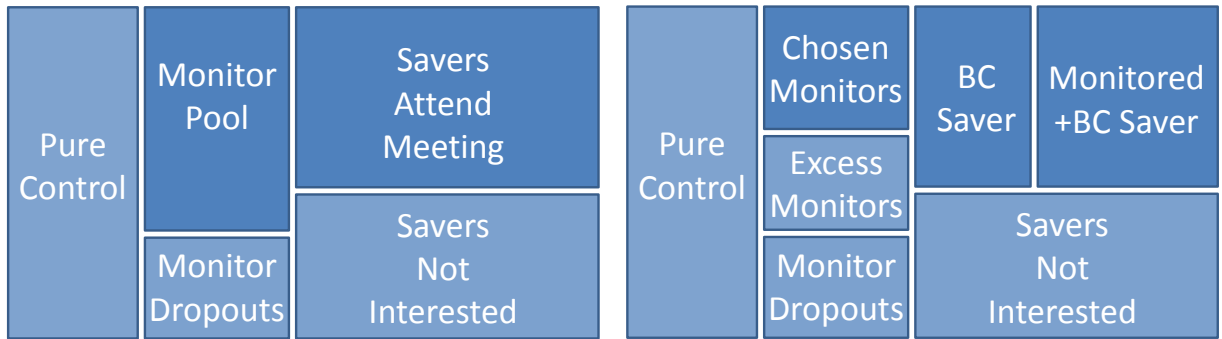
FIGURES

FIGURE 1. Experimental Design and Randomization




(A) Village

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



(C) Participating samples.

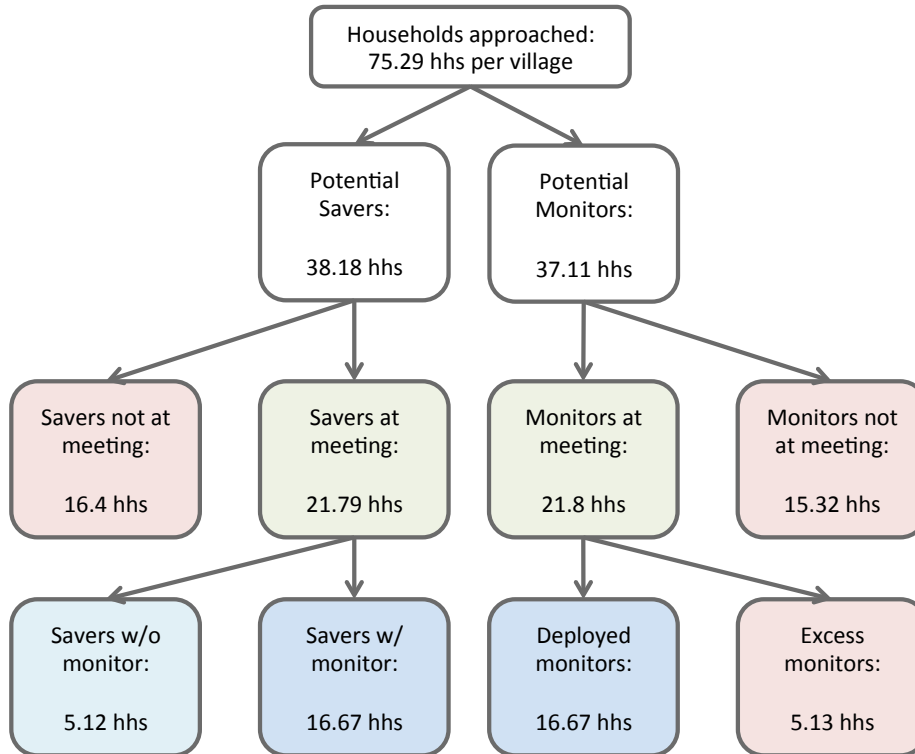
(D) Individual-level randomization of participating savers to treatments. Monitors selected.



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

“BC Saver” refers to our non-monitored treatment (T1) described in section 2.

FIGURE 2. Sample Description



This figure shows the average number of households per village in our sample in each cell. We have 60 villages in our sample.

FIGURE 3. Time Line of Experiment






FIGURE 4. Histogram of Baseline Savings Goals



The figure shows the distribution of the baseline savings goals. We clip the top 5% tail of the distribution to make the figure more readable.

FIGURE 5. Distributions (CDF) of $\log(\text{Total Savings}/\text{Savings goal})$ by Treatment

Non-Monitored Savers vs. Savers with Random Monitors



Savers with High Centrality vs. Low Centrality Monitors




FIGURE 6. Social Distance Between Saver and Monitor by Treatment





FIGURE 7. Centrality Quantile of Monitor by Treatment



TABLES

TABLE 1. Summary Statistics, Treatment Assignment, and Attrition

<i>Dependent Variable</i>	Treatment (Village Meeting Sample)			Treatment (Endline Sample)			Obs.
	Mean of Non-Monitored Savers	Diff. Random vs. No Monitor	Diff. Endogenous vs. No Monitor	Mean of Non-Monitored Savers	Diff. Random vs. No Monitor	Diff. Endogenous vs. No Monitor	
Age	33.09 (0.385)	-0.147 (0.458)	0.158 (0.528)	33.45 (0.387)	0.0414 (0.454)	0.0254 (0.551)	1,146
Female	0.756 (0.0243)	-0.0411 (0.0316)	-0.0253 (0.0343)	0.78 (0.0261)	-0.0412 (0.0304)	-0.0248 (0.0339)	1,146
Married	0.857 (0.0192)	-0.0287 (0.0208)	-0.0244 (0.0272)	0.875 (0.0204)	-0.0334 (0.0218)	-0.0409 (0.0268)	1,146
Widowed	0.0358 (0.00984)	0.00954 (0.0126)	0.0151 (0.0161)	0.033 (0.0101)	0.0175 (0.0137)	0.0246 (0.0174)	1,146
Positive Savings 6 Mos Prior to Baseline	0.717 (0.0319)	0.0244 (0.0346)	0.0116 (0.0358)	0.725 (0.0333)	0.0181 (0.0371)	0.0157 (0.0370)	1,146
Has Post Office or Bank Acct. at Baseline	0.378 (0.0316)	-0.0111 (0.0362)	0.0404 (0.0340)	0.396 (0.0329)	-0.00964 (0.0367)	0.0241 (0.0354)	1,146
Has BPL Card	0.84 (0.0211)	0.0197 (0.0251)	-0.00175 (0.0269)	0.842 (0.0235)	0.0150 (0.0296)	0.0112 (0.0280)	1,146
Savings Goal	1838 (117.1)	-239.1 (117.4)	131.1 (165.2)	1751 (126.6)	-207.7 (117.8)	185.8 (166.2)	1,146
Savings Goal (1% outliers trimmed)	1650 (76.04)	-106.5 (78.99)	-55.07 (101.0)	1538 (75.69)	-35.75 (81.40)	34.54 (103.4)	1,127
Log Savings Goal	7.253 (0.0398)	-0.0631 (0.0415)	0.00868 (0.0464)	7.209 (0.0421)	-0.0350 (0.0408)	0.0467 (0.0476)	1,146
Projected Income - Projected Expenses	3175 (349.8)	-204.6 (607.4)	-975.5 (943.8)	2878 (376.1)	-295.0 (596.0)	-1,103 (1,035)	1,146
Endline Survey Administered (Non-Attriters) <i>No Fixed Effects</i>				0.889 (0.0179)	-0.0272 (0.0252)	-0.00390 (0.0219)	1,307
Endline Survey Administered (Non-Attriters) <i>Village Fixed Effects</i>				0.887 (0.0150)	-0.00413 (0.0295)	-0.0248 (0.0254)	1,307
15-Month Follow-Up Survey Administered (Non-Attriters) <i>No Fixed Effects</i>				0.893 (0.0175)	-0.000259 (0.0243)	0.0353 (0.0212)	1,307
15-Month Follow-Up Survey Administered (Non-Attriters) <i>Village Fixed Effects</i>				0.896 (0.0143)	0.0139 (0.0287)	0.0101 (0.0232)	1,307

TABLE 2. Effect of Random Monitors on Savings

	(1)	(2)	(3)	(4)
<i>Dependent Variable</i>	Log Total Savings	Log Total Savings	Target Account Savings	Target Account Savings
Monitor Treatment: Random Assignment	0.301 (0.155)	0.370 (0.148)	360.6 (174.5)	368.2 (188.4)
Observations	544	544	673	673
R-squared	0.123	0.025	0.031	0.013
Dependent Variable Mean (Omitted Group)	7.647	7.647	342.7	342.7
Fixed Effects	Village	Village	Village	Village
Controls	Saver	No	Saver	No

Notes: 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” – by the saver. Target account savings is the amount that is saved in the target account associated with the experiment. Sample constrained to 30 villages where monitors are randomly assigned. Columns (3) and (4) use administrative data for all 673 savers who were in our sample at baseline. Columns (1) and (2) use survey data where 544 savers answered our questionnaire. Controls include the following saver characteristics: savings goal, 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. All regressions include village fixed effects. Standard errors clustered at the village level.

TABLE 3. Total Savings by Network Position of Random Monitor

	(1)	(2)	(3)
<i>Dependent Variable</i>	Log Total Savings	Log Total Savings	Log Total Savings
Monitor Centrality	0.134 (0.0729)	0.178 (0.0736)	
Saver-Monitor Proximity	0.865 (0.334)		1.032 (0.352)
Observations	424	424	424
R-squared	0.161	0.150	0.155
Fixed Effects	Village	Village	Village
Controls	Saver, Monitor	Saver, Monitor	Saver, Monitor

Notes: Total savings is the amount saved across all savings vehicle – the target account and any other account, both formal and informal including money held “under the mattress” – by the saver. Sample constrained to savers who received a monitor in the 30 villages where monitors are randomly assigned and where the savers answered our questionnaire. Controls include savings goal, and 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. We also control for the geographical distance between the homes of the saver and monitor. All regressions include village fixed effects. Standard errors clustered at the village level.

TABLE 4. Beliefs About Savers and Monitor Centrality

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Dependent Variable: Beliefs about Saver</i>	Reached Goal	Reached Goal	Reached Goal	Good at Meeting Goals	Good at Meeting Goals	Good at Meeting Goals
Monitor Centrality	0.0206 (0.00937)	0.0157 (0.00804)	0.0157 (0.00854)	0.0389 (0.0144)	0.0374 (0.0140)	0.0353 (0.0148)
Respondent-Monitor Proximity	0.00357 (0.0194)	-0.00252 (0.0196)	-0.00160 (0.0239)	0.0476 (0.0422)	0.0181 (0.0366)	0.0360 (0.0342)
Observations	4,743	4,743	4,743	4,743	4,743	4,743
R-squared	0.026	0.020	0.342	0.030	0.023	0.314
Fixed Effects	No	Village	Respondent	No	Village	Respondent
Controls	Saver	Saver	Saver	Saver	Saver	Saver

Notes: The dependent variables are measured in the 15 month follow-up survey and capture the beliefs of respondents about savers in the random monitor villages. “Reached Goal” measures whether the saver reached her goal and the respondent correctly believed this to be true. “Good at Meeting Goals” is constructed as $1(\text{Saver reached goal}) * 1(\text{Respondent indicates saver is good or very good at meeting goals}) + (1 - 1(\text{Saver reached goal})) * 1(\text{Respondent indicates saver is mediocre, bad or very bad at meeting goals})$. Sample uses 615 respondents across the 30 random villages, each of whom was asked to rate approximately 8 randomly selected savers who had a monitor from their village in the manner described above. Controls include the following saver characteristics: savings goal, 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 2 and 5 include village fixed effects. Columns 3 and 6 include respondent fixed effects. Standard errors clustered at the village level.

TABLE 5. How Did the Savers Save?

Panel A: Expenditures During Month 6 of Savings Period

<i>Dependent Variable</i>	Log Expenditures	Festivals	Pan	Tea	Meals Away	Eggs and Meat	Other Food	Transport.	Entertainment and Phone
Random Monitor	-0.0662 (0.0708)	-223.3 (126.9)	17.88 (25.88)	35.13 (17.75)	19.16 (38.54)	-53.22 (62.91)	-150.0 (129.6)	-153.5 (80.79)	-2.719 (29.01)
Observations	981	1,114	1,114	1,115	1,101	1,108	1,114	1,115	1,106
R-squared	0.028	0.028	0.018	0.006	0.042	0.055	0.075	0.046	0.057
Fixed Effects Controls	Village Saver	Village Saver	Village Saver	Village Saver	Village Saver	Village Saver	Village Saver	Village Saver	Village Saver

Panel B: Retrospective Assessment from Follow-Up Survey

<i>Dependent Variable</i>	(1) Increased Labor Supply	(2) Business Profits	(3) Cut Unnecessary Expenditures	(4) Money from Family and Friends	(5) Reduced Transfers to Others	(6) Took a Loan
Random Monitor	0.0712 (0.0332)	0.0202 (0.0156)	0.0787 (0.0422)	-0.0227 (0.0346)	0.0148 (0.0120)	-0.0222 (0.0190)
Dep. Var. Mean	0.15	0.03	0.15	0.19	0.01	0.04
Observations	1,026	1,026	1,026	1,026	1,026	1,026
R-squared	0.055	0.026	0.020	0.056	0.016	0.014
Fixed Effects Controls	Village Saver	Village Saver	Village Saver	Village Saver	Village Saver	Village Saver

Notes: Panel A measures the effect of receiving a randomly assigned monitor on selected measures of expenditures in the sixth month of the savings period measured at the end of the monitoring intervention. Panel B reports survey responses from the 15 month follow-survey. 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. The total expenditure outcome variable takes the intersection of all responses. Controls include the following saver characteristics: savings goal, 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, forecasted expenditure at baseline, number of rooms in the home, type of electrical connection and an indicator for endogenous monitor. All regressions include village fixed effects. Standard errors clustered at the village level.

TABLE 6. Shock Mitigation for Monitored Savers in Random Villages

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
<i>Dependent Variable: Shocks</i>	Total	Total	Greater than	Greater than	Health	Health	HH	HH	log(Tot. Sav.)	log(Tot. Sav.)
	Number	Number	Median	Median	Health	Health	Expenditure	Expenditure	15 mos.	15 mos.
Monitor Treatment: Random Assignment	-0.199 (0.128)	-0.249 (0.131)	-0.0757 (0.0416)	-0.0944 (0.0441)	-0.0752 (0.0615)	-0.103 (0.0670)	-0.0521 (0.0384)	-0.0721 (0.0419)	0.324 (0.196)	0.290 (0.190)
Observations	1,153	1,153	1,153	1,153	1,153	1,153	1,153	1,153	1,152	1,152
R-squared	0.021	0.021	0.019	0.016	0.020	0.020	0.010	0.011	0.074	0.083
Mean of Dep. Var (Control)	1.769	1.769	0.577	0.577	0.862	0.862	0.500	0.500	3.779	4.264
Fixed Effects	Village	No	Village	No	Village	No	Village	No	Village	No
Controls	Saver	Saver	Saver	Saver	Saver	Saver	Saver	Saver	Saver	Saver

Notes: The outcome variables are all measures of shocks experienced by the savers between the end of the six month savings period and the 15 month follow-up survey. The total number of shocks measures the number of types of shocks experienced, including deaths, family illnesses, health shocks causing missed work, livestock shocks, 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. The total expenditure outcome variable takes the intersection of all responses. Controls include the following saver characteristics: savings goal, 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, and 7 include village fixed effects. Standard errors clustered at the village level.

TABLE 7. No Evidence of Bunching or Gaming

	(1)	(2)	(3)
	Number of Savers by Treatment		
	No Monitor	Random Monitor	Endogenous Monitor
<i>Within Rs. 50 of Half Goal</i>			
Target acct savings € [half goal-50, half goal)	1	5	6
Target acct savings € [half goal, half goal+50]	9	9	13
Target acct savings € [half goal-50, half goal+50]	10	14	19
Fraction above half goal given near half goal	90%	64%	68%
	(1)	(2)	(3)
	Number of Savers by Treatment		
	No Monitor	Random Monitor	Endogenous Monitor
<i>Within Rs. 150 of Full Goal</i>			
Target acct savings € [goal-150, goal)	2	3	6
Target acct savings € [goal, goal+150]	3	18	14
Target acct savings € [goal-150, goal+150]	5	21	20
Fraction above goal given near goal	60%	86%	70%

Notes: The top panel shows the number of savers reaching within Rs. 50 of 1/2 of the savings goal by treatment. Rs. 50 is the size of the incentive received by the monitor if the saver reaches or surpasses 1/2 of the goal. The bottom panel shows the number of savers reaching within Rs. 150 of the full savings goal by monitor status. Rs. 150 is the size of the additional incentive received by the monitor if the saver reaches or surpasses the full goal.

TABLE 8. Random vs. Endogenous Monitors

	(1)	(2)
<i>Dependent Variable</i>	Log Total Savings	Log Total Savings
Monitor Treatment: Random Assignment Village	0.298 (0.145)	0.298 (0.148)
Monitor Treatment: Endogenous Assignment Village	-0.0722 (0.146)	-0.0830 (0.161)
Non-Monitored Treatment: Endogenous Assignment Village	0.354 (0.202)	
Observations	1,042	1,042
R-squared	0.145	0.125
Fixed Effects	No	Village
Controls	Saver	Saver

Notes: Total savings is the amount saved across all savings vehicle – the target account and any other account – by the saver. Sample includes all 1042 savers in our sample who responded to our questionnaire. Controls include the following saver characteristics: savings goal, 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. All regressions include village fixed effects. Standard errors clustered at the village level.

TABLE 9. Spillovers from Monitored Savers: Non-Monitored Sample

	(1)	(2)
<i>Dependent Variable</i>	Log Total Savings	Log Total Savings
Number of Friends Assigned a Monitor	0.833 (0.315)	0.372 (0.287)
Sum of Centralities of Friends' Monitors		10.24 (3.749)
Number of Friends Attending Meeting	0.795 (1.240)	0.650 (1.257)
Observations	123	123
R-squared	0.541	0.584
Fixed Effects	Village	Village
Controls	Saver	Saver

Notes: Total savings is the amount saved across all savings vehicle – the target account and any other account – by the saver. The mean of the variable 'Sum of Centralities of Friends' Monitors' is 0.14 with a standard deviation of .13. Sample is restricted to the 123 non-monitored savers in the 30 villages who responded to our questionnaire. Controls include the following saver characteristics: savings goal, 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. All regressions include village fixed effects. Standard errors clustered at the village level.

APPENDIX A. SUPPLEMENTAL APPENDIX: QUOTES

“For those who want to save in a bank or post office account but do not have the habit of doing so, having a monitor may help... Having a more important person as a monitor may help in comparison to a person who is not well known by people in the village. A person may save more if it is an important person knowing they might get more benefits from this person later on.” – Subject 1

“If the monitor was a very important person in the village, and the saver did not meet a goal that she set, the monitor would lose trust in the saver. The monitor will feel that if in the future he or his friends gives her some job or tasks or responsibilities, the saver may not fulfill them.” – Subject 2

“When paired with an important person, they will save more to build the monitor’s confidence in them. That way the person builds trust with me [sic]... If the person does not fulfill savings, the monitor will be disappointed and think ‘I used to place trust in that person but now I can’t’. They would speak less to the saver and feel ‘cheated to trust’ [sic]. They may tell others... But if someone is too irresponsible then monitor or no monitor, the saver will not save.” – Subject 3

“People will only reach their goals if their monitors are family, friends, neighbors, or important people.” - Subject 4

“I would like to choose the important person except if there are close friends. Then I may hesitate if I do not know him well.” – Subject 5

APPENDIX B. SUPPLEMENTAL APPENDIX: TABLES AND FIGURES





FIGURE B.2. Distributions (CDF) of $\log(\text{Total Savings}/\text{Savings goal})$ by Treatment

Non-Monitored Savers vs. Savers with Random Monitors



Savers with High Centrality vs. Low Centrality Monitors




FIGURE B.3. Aggregate Timing of Savings Deposits




TABLE B.1. Endline Survey Summary Statistics: Non-Monitored Savers

<i>Summary Statistics: Sample Villages</i>	Obs.	Mean	Std. Dev.
<i>Network Characteristics</i>			
Number of Households	60	222.12	65.85
Average Degree	60	17.57	3.96
Average Clustering	60	0.30	0.05
Average Path Length	60	2.34	0.19

TABLE B.2. Endline Survey Summary Statistics: Non-Monitored Savers

<i>Summary Statistics: Non-Monitored Savers, R Villages</i>	Obs.	Mean	Std. Dev.
<i>Endline Survey: Conclusion of Intervention</i>			
Total Savings	123	8890.44	17616.18
Log Total Savings	123	7.67	1.83
Log Total Expenditures (past month)	120	8.62	0.79
<i>Expenditure Categories (past month):</i>			
Festivals	133	824.81	1335.38
Pan	132	197.73	219.89
Tea	132	277.05	227.49
Meals Away	131	259.39	478.09
Eggs and Meat	131	606.72	783.06
Other Food	133	1526.92	1347.66
Transport	132	641.14	1061.75
Entertainment and Phone	133	244.64	213.72
<i>Final Endline Survey: 15 Months Following Conclusion</i>			
Total Savings	133	9263.29	16124.83
Log Total Savings	133	7.65	2.08
<i>How the Savers Saved:</i>			
Increased Labor Supply	117	0.15	0.36
Business Profits	117	0.03	0.18
Cut Unnecessary Expenditures	117	0.15	0.35
Money from Family and Friends	117	0.19	0.39
Reduced Transfers to Others	117	0.01	0.09
Took a Loan	117	0.04	0.20
<i>Shocks</i>			
Total Number of Shocks	133	1.77	1.43
Greater than Median Number of Shocks	133	0.58	0.50
Health Shock Indicator	133	0.86	0.66
HH Expenditure Shock Indicator	133	0.50	0.50
<i>Beliefs about Non-Monitored Savers in R Villages</i>			
Reached Goal	2141	0.03	0.18
Good at Meeting Goals	2141	0.21	0.41
Responsibility Index Raw Data (5 point scale)	1467	2.00	0.97

TABLE B.3. Effect of Random Monitors on Goal Attainment

<i>Dependent Variable</i>	(1)	(2)
	Reached Goal	Reached Goal
Monitor Treatment: Random Assignment	0.0630 (0.0316)	0.0606 (0.0313)
Observations	673	673
R-squared	0.021	0.012
Dependent Variable Mean (Omitted Group)	0.073	0.073
Fixed Effects	Village	Village
Controls	Saver	No

Notes: Reached Goal is a dummy for whether the saver (weakly) exceeded her savings goal. Sample constrained to 30 villages where monitors are randomly assigned and all 673 savers who were in our sample at baseline. Controls include the following saver characteristics: savings goal, 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. All regressions include village fixed effects. Standard errors clustered at the village level.

TABLE B.4. Goal Attainment Network Position of Random Monitor

<i>Dependent Variable</i>	(1)	(2)	(3)
	Reached Goal	Reached Goal	Reached Goal
Monitor Centrality	0.0288 (0.0162)	0.0339 (0.0156)	
Saver-Monitor Proximity	0.118 (0.0718)		0.147 (0.0698)
Observations	523	523	523
R-squared	0.053	0.048	0.046
Fixed Effects	Village	Village	Village
Controls	Saver, Monitor	Saver, Monitor	Saver, Monitor

Notes: Reached Goal is a dummy for whether the saver (weakly) exceeded her savings goal. Sample constrained to all 523 monitored savers in the random villages. Controls include savings goal, and 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. We also control for the geographical distance between the homes of the saver and monitor. All regressions include village fixed effects. Standard errors clustered at the village level.

TABLE B.5. How Did the Savers Save?

Panel A: Expenditures During Month 6 of Savings Period

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
<i>Dependent Variable</i>	Log Expenditures	Festivals	Pan	Tea	Meals Away	Eggs and Meat	Other Food	Transport.	Entertainment and Phone
Monitor Treatment: Random Assignment	-0.0655 (0.0716)	-224.1 (130.2)	17.04 (25.57)	37.58 (18.83)	19.75 (40.97)	-47.80 (62.68)	-160.7 (130.9)	-159.4 (79.93)	-2.623 (29.35)
Observations	522	578	581	580	573	574	578	578	577
R-squared	0.068	0.035	0.021	0.027	0.060	0.097	0.112	0.084	0.057
Fixed Effects	Village	Village	Village	Village	Village	Village	Village	Village	Village
Controls	Saver	Saver	Saver	Saver	Saver	Saver	Saver	Saver	Saver

Panel B: Retrospective Assessment from Follow-Up Survey

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Dependent Variable</i>	Increased Labor Supply	Business Profits	Cut Unnecessary Expenditures	Money from Family and Friends	Reduced Transfers to Others	Took a Loan
Monitor Treatment: Random Assignment	0.0680 (0.0332)	0.0224 (0.0163)	0.0761 (0.0447)	-0.0244 (0.0342)	0.0164 (0.0119)	-0.0224 (0.0192)
Observations	528	528	528	528	528	528
R-squared	0.074	0.056	0.025	0.057	0.055	0.031
Fixed Effects	Village	Village	Village	Village	Village	Village
Controls	Saver	Saver	Saver	Saver	Saver	Saver

Notes: Panel A measures the effect of receiving a randomly assigned monitor on selected measures of expenditures in the sixth month of the savings period measured at the end of the monitoring intervention. Panel B reports survey responses from the 15 month follow-up survey. Sample constrained to all savers in the 30 villages where monitors are randomly assigned and where the savers answered our questionnaire. If a response was missing for a category, the observation is missing in the regression. The total expenditure outcome variable takes the intersection of all responses. Controls include the following saver characteristics: savings goal, 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, forecasted expenditure at baseline, number of rooms in the home and type of electrical connection. All regressions include village fixed effects. Standard errors clustered at the village level.

TABLE B.6. Shock Mitigation for Monitored Savers in Random Villages

<i>Dependent Variable: Shocks</i>	(1) Total Number	(2) Total Number	(3) Greater than Median	(4) Greater than Median	(5) Health	(6) Health	(7) HH Expenditure	(8) HH Expenditure	(9) log(Tot. Sav.) 15 mos.	(10) log(Tot. Sav.) 15 mos.
Monitor Treatment: Random Assignment	-0.189 (0.128)	-0.228 (0.128)	-0.0751 (0.0422)	-0.0903 (0.0434)	-0.0769 (0.0628)	-0.0992 (0.0686)	-0.0484 (0.0391)	-0.0676 (0.0416)	0.294 (0.207)	0.255 (0.195)
Observations	599	599	599	599	599	599	599	599	596	596
R-squared	0.030	0.031	0.031	0.026	0.030	0.030	0.016	0.022	0.087	0.103
Mean of Dep. Var (Control)	1.769	1.769	0.577	0.577	0.862	0.862	0.500	0.500	3.779	4.264
Fixed Effects	Village	No	Village	No	Village	No	Village	No	Village	No
Controls	Saver	Saver	Saver	Saver	Saver	Saver	Saver	Saver	Saver	Saver

Notes: The outcome variables are all measures of shocks experienced by the savers between the end of the six month savings period and the 15 month follow-up survey. The total number of shocks measures the number of types of shocks experienced, including deaths, family illnesses, health shocks causing missed work, livestock shocks, unexpected HH expenditures. Sample constrained to all savers in the 30 villages where monitors are randomly assigned and where the savers answered our questionnaire. If a response was missing for a category, the observation is missing in the regression. The total expenditure outcome variable takes the intersection of all responses. Controls include the following saver characteristics: savings goal, 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, 5, and 7 include village fixed effects. Standard errors clustered at the village level.

TABLE B.7. Random vs. Endogenous Monitors

<i>Dependent Variable</i>	(1) Reached Goal	(2) Reached Goal
Monitor Treatment: Random Assignment	0.0619 (0.0322)	0.0601 (0.0311)
Monitor Treatment: Endogenous Assignment	0.0646 (0.0219)	0.0598 (0.0229)
Village w/ Endogenous Monitor Selection	-0.00497 (0.0334)	
Observations	1,286	1,286
R-squared	0.025	0.023
Fixed Effects	No	Village
Controls	Saver	Saver

Notes: Reached Goal is a dummy for whether the saver (weakly) exceeded her savings goal. Controls include the following saver characteristics: savings goal, 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. All regressions include village fixed effects. Standard errors clustered at the village level.

TABLE B.8. Random Monitor Treatment Effects: Robustness to Inclusion of Spillovers

	(1)	(2)	(3)
<i>Dependent Variable</i>	Log Total Savings	Log Total Savings	Log Total Savings
Monitor Centrality	0.145 (0.0727)	0.179 (0.0718)	
Saver-Monitor Proximity	0.928 (0.367)		1.077 (0.378)
Observations	424	424	424
R-squared	0.299	0.289	0.292
Fixed Effects	Village	Village	Village
Controls	Saver, Monitor	Saver, Monitor	Saver, Monitor

Notes: Total savings is the amount saved across all savings vehicle – the target account and any other account, both formal and informal including money held “under the mattress” – by the saver. Sample constrained to savers who received a monitor in the 30 villages where monitors are randomly assigned and where the savers answered our questionnaire. Controls include savings goal, and 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. They also include the number of friends, the number of monitored friends, and the sum of the centralities of the monitors of the friends. We also control for the geographical distance between the homes of the saver and monitor. All regressions include village fixed effects. Standard errors clustered at the village level.

TABLE B.9. Random Monitor Analysis: Financial Component of Network Only

	(1)	(2)
<i>Dependent Variable</i>	Log Total Savings	Log Total Savings
Saver and Monitor Direct Friends: Any Relationship	0.564 (0.210)	0.522 (0.255)
Saver and Monitor Direct Friends: Borrowing or Lending Relationship		0.145 (0.435)
Observations	424	424
R-squared	0.144	0.145
Fixed Effects	Village	Village
Controls	Saver, Monitor	Saver, Monitor

Notes: Total savings is the amount saved across all savings vehicle – the target account and any other account, both formal and informal including money held “under the mattress” – by the saver. We define a link as having a financial component if the nodes report borrowing or lending small amounts of money or material goods to one another. In our sample, 27% of direct links have a financial component. Controls include savings goal, and 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. We also control for the geographical distance between the homes of the saver and monitor. All regressions include village fixed effects. Standard errors clustered at the village level.

APPENDIX C. MULTIGRAPH ANALYSIS

To this point, we have used the union network in all of our analysis, that is the network encoding all types of interactions. The basic logic is that since an individual could pass information to her coworker, who then tells her friend, who then tells his neighbor when borrowing rice about the information, the union network may give the best picture of what type of centrality and proximity matter. However, a natural question to ask is whether what really matters is centrality or proximity in the network for a specific type of relationship.

To study this we examine on two dimensions of the network: the financial graph constructed by only considering borrowing or lending links among households and the advice graph constructed by only considering ties where one household reported a willingness to seek advice from other households.

We then repeat our main regressions, where instead of merely looking at the centrality of the monitor or the proximity between the saver and monitor from the union network, we also include the centrality and proximity constructed from the financial and advice networks. Table C.1 presents the results. In column 2 we include monitor centralities from the union network, the financial network and the advice network. We find that the centrality from the union network is strongly related to saver behavior, whereas conditional on this there is no significant differential effect due to centrality in the financial or advice network. In column 3

we repeat this exercise where we include saver-monitor proximities from the union, financial and advice networks. Again we find that the only significant relationship comes from the union network and conditional on this there is no differential effect from the financial or advice networks. Finally in column 1 we include all six variables together. Here due to the relative correlation of all the variables involved, the results are very noisy. But the point estimates for the coefficients on the centrality and proximity variables constructed from the union networks remain similar, bolstering our conclusion that to first order we can simply consider the “all” network.

TABLE C.1. Total Savings by Network Position of Random Monitor

	(1)	(2)	(3)
<i>Dependent Variable</i>	Log Total Savings	Log Total Savings	Log Total Savings
Monitor Centrality: Full Network	0.139 (0.0960)	0.180 (0.100)	
Monitor Centrality: Financial Network	-0.0111 (0.146)	0.00246 (0.154)	
Monitor Centrality: Advice Network	-0.00228 (0.112)	-0.00589 (0.121)	
Saver-Monitor Proximity: Full Network	0.642 (0.515)		0.814 (0.521)
Saver-Monitor Proximity: Financial Network	0.258 (0.928)		0.320 (0.938)
Saver-Monitor Proximity: Advice Network	0.147 (0.770)		0.202 (0.761)
Observations	424	424	424
R-squared	0.164	0.152	0.152
Fixed Effects	Village	Village	Village
Controls	Saver, Monitor	Saver, Monitor	Saver, Monitor

Notes: Total savings is the amount saved across all savings vehicle – the target account and any other account, both formal and informal including money held “under the mattress” – by the saver. Sample constrained to savers who received a monitor in the 30 villages where monitors are randomly assigned and where the savers answered our questionnaire. Controls include savings goal, and 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. We also control for the geographical distance between the homes of the saver and monitor. All regressions include village fixed effects. Standard errors clustered at the village level.