SAVINGS MONITORS

EMILY BREZA† AND ARUN G. CHANDRASEKHAR‡

Abstract. We conduct a field experiment in India to explore two interventions to help individuals to increase their savings balances. First, we design a financial product based on the popular business correspondent model, which includes frequent reminders, assistance in account opening, and the setting of a six-month savings goal. Second, we measure the effectiveness of adding a peer monitoring component to this basic bundle and test whether the local social network can help to increase the penetration of the formal banking system. We ask whether having a monitor substitutes for a formal commitment device, whether individuals choose the most effective monitors, and moreover, whether some community members are better than others at encouraging financial capability.

JEL Classification Codes: D03, D14, O16, Z13

Keywords: Savings, Social Networks, Peer Monitoring

1. Introduction

Increasing a household's capacity to save can have large effects on a range of economic outcomes.[1] Despite the importance of savings, a large literature documents that rural households do not appear to save adequately.[2] One possible solution explored by the literature is a technology-driven solution consisting of commitment devices or reminders platforms.

A distinct, but important technology at the disposal of communities – especially in communities such as rural villages – is the social and economic network of interactions. Numerous informal financial instruments such as informal insurance, rotating savings credit associations (RoSCAs), village savings associations, self-help groups, and merchant guilds

Date: November 2013 – Preliminary and incomplete, please do not cite without an author's consent.

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, Gabriel Carroll, Pascaline Dupas, Dina Pomeranz, Chris Udry, Rohini Pande, Jonathan Robinson, for helpful discussions, as well as seminar and conference participants at the Santa Fe Institute, SITE, and the European Conference on Household Finance (Einaudi) for their insightful comments. We are grateful to CFSP, NSF, the Chazen Institute at Columbia Business School, CIBER at Columbia University, 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.

†Columbia Business School. Email: eb2840@columbia.edu.

‡Stanford University. Email: arungc@stanford.edu.

[1] See Dupas and Robinson (2013b) and Kast and Pomeranz (2013).

[2] Mechanisms include time inconsistency, inattention, and lack of opportunities. For example Ashraf et al. (2006), Karlan et al. (2010), and Dupas and Robinson (2013a).

have leveraged the social network in order to sustain constrained-efficient behavior in the absence of formal institutions.[3] Despite the ubiquitous role of social networks, there has been little emphasis both in (a) thinking about institutional design that exploits social network structure and (b) measuring experimentally the role that social networks play on determining outcomes in such relationships. A central challenge is that such an analyses require both detailed network data across many communities as well as experimental variation that randomly assigns members within a network. We address these challenges by conducting a peer monitoring field experiment to study how powerful a village social network can be in encouraging formal savings behavior. Specifically, we conduct our experiment across 60 villages where we have detailed social network data between all households across numerous dimensions. We both randomly and endogenously assign monitors – who observe an individual's savings progress – to savers within a village and study how variation in a saver-monitor network influences savings behavior.

Individuals within a village engage in repeated interactions with each other on a wide variety of dimensions: e.g., talking about new technologies or job opportunities, making transfers to each other, sharing and maintaining public goods.[4] Network-based models of these interactions do not treat agents as exchangeable; importantly, some agents are more central than others – they directly or indirectly are involved in more information exchanges or financial transactions. There has been growing empirical evidence that the variation in village network centrality is important in predicting economic outcomes in a variety of contexts.[5] We aim to understand if, when given the role of monitor, some members of the network, namely the more central members, are also better suited to help others to save. We are able to ask this question in a high stakes field setting where we are also able to vary experimentally the centralities of our monitors.

To understand whether monitors can help individuals reach their savings goals and, importantly, whether central monitors can generate further expansion in savings, we designed the following experiment. Individuals were randomly assigned to one of three savings treatments (or assigned to a pure control group). Before knowing which treatment an individual was assigned to, she would write down a six-month savings goal. The first group was offered a savings device following the common Business Correspondents model (henceforth BC bundle). In this treatment, we assisted households in opening an account, helped individuals construct an attainable 6-month savings goal, and visited savers fortnightly to check in on their progress. The second group was offered the BC bundle but additionally was assigned a random monitor from their village. After every

[3]See, for instance Greif (1993).
[4]See Townsend (1994); Beaman and Magruder (2009); Conley and Udry (2010a); Fafchamps and Lund (2003).
[5]For examples, see Banerjee et al. (2013); Breza et al. (2013); Chandrasekhar et al. (2013); Kinnan and Townsend (2010).

fortnightly visit to the saver's house, the monitor would be informed about the saver's progress towards her savings goal. The third group was offered the BC bundle, but in this case individuals were able to choose their monitors. The fourth group consisted of pure control – those who were offered no BC bundle nor a monitor. Individuals in the three treatment groups set their six-month savings goals before being assigned to any treatment; the control group never set a goal nor received any feature of the BC bundle.

Our main results are as follows. First, we find that randomly assigned monitors substantially increase the probability of reaching one's savings goal (6 pp increase relative to a 7% rate of reaching one's savings goal in the individual level BC treatment); moreover, the overall savings balances across all accounts increases by more than 35% relative to the BC treatment. Randomly assigned monitors, on average, generate better results than both the individual treatment and the endogenous treatment.

Second, there are substantial network-effects. Specifically, randomly assigned central monitors generate considerably larger increases in savings attainment. A one standard deviation increase in monitor centrality corresponds to a Rs. 1400 or 18% increase in savings relative to the BC treatment. Additionally, being assigned a monitor of social distance one is associated with an 8.5pp increase in goal attainment and 43% increase in total savings balances relative to a monitor of social distance two. Importantly, the centrality and proximity results hold conditional on each other, and the analysis includes numerous demographic controls including caste, wealth, occupation and village fixed effects.

We stress that these network findings are robust to the inclusion of controls for numerous observable characteristics including wealth proxies, caste, and village leadership status to show that the network effects are present conditional on these observables.

Third, endogenous monitors do not generate gains relative to the individual-level treatment. While individuals do select more socially close monitors and central monitors, the endogenous treatment does no better than the individual level BC treatment.

Our fourth main result is that the BC bundle itself generates large effects. Individuals offered the business correspondent treatment save 25% more with 38% of invited savers taking-up the treatment. However, individuals in the BC treatment exhibit a low rate of meeting their savings goals (7%) in their formal accounts. We should note that all treatment groups include this bundle, so the treatment effects of monitoring, for example, are measured above and beyond the effects of the BC package.

We interpret our network results using a simple signaling model. Because numerous economic transactions require collaboration with some fellow community member, opinions that other community members hold may be extremely valuable. A saver's ability to present herself to her to others as a responsible individual who can attain her goals may influence how members of the community treat her in other contexts.[6] From a networks

[6]This can be interpreted quite loosely: one who faces less time inconsistency, one who is able to follow-through on goals set, one who is less plagued by inattention, etc.

perspective, we know that information propagated by central individuals is more likely to percolate throughout the village network.[7] This implies that it may be more valuable to generate a favorable opinion from a network-central monitor than from a socially peripheral monitor.

The model's key predictions are that (i) central monitors should generate better savings behavior in the random treatment and (ii) socially proximate monitors should generate better savings behavior in the random treatment. The core findings are consistent with this model.

We explore alternative mechanisms as well. For instance, we provide additional evidence that the effect is likely not due to misattributing demographic effects to network effects (we control for numerous observables at our disposal) nor is it likely due to mimicry (individuals' savings behavior are uncorrelated with their monitors' savings behavior).

Finally, we are able to ask whether more savings are generated when the savers themselves are given the ability to choose (publicly) their own monitors. Our findings suggest that the random treatment causes more savings, and that the excess savings do not appear to come from overmonitoring by central monitors. This suggests that there may be frictions in the community based on social relationsihps that limit the choice sets of savers when they themselves must choose actively.

Our results contribute to two literatures: (i) analysis of savings in developing countries and (ii) analysis of the importance of social networks in developing economies.

Both the desire of the poor to save as well as the benefits of savings are well-documented in the economics literature.It has been established empirically that the poor have a desire to save but are unable due to lack of access, lack of commitment, or lack of attention.[8]. Dupas and Robinson (2013a), Brune et al. (2013) and others provide experimental evidence that increased savings can increase investment, working capital, income, and even labor supply. There are many hypotheses that attempt to explain why individuals fail to save "enough." Many behavioral explanations such as time inconsistency, temptations, or inattention can produce undersaving.[9] Rational hypotheses include distrust of banks, limited access to safe storage of funds, high transactions costs and high discount rates.[10]

---

[7]See Jackson (2008) for an extensive review of the literature on such models. For empirical evidence that information about the characteristics of fellow villagers appears to have diffused through the community, see Alatas et al. (2012).

[8]See Dupas and Robinson (2013b) for a more detailed discussion. Examples include Shipton (1990) who document villagers creating their own lockboxes, Rutherford and Arora (2009) who provides evidence that the poor desire to save more, Collins et al. (2009) who documents numerous savings vehicles used by the poor in south Asia and South Africa, and Dupas and Robinson (2013b).

[9]See Laibson (1997), Bernheim et al. (1999), Bernheim et al. (2013) and Banerjee and Mullainathan (2010) among many others.

[10]Innovations such as mobile money have begun to reduce the costs of using savings accounts. Schaner (2011), for example, finds that lowering transactions costs dramatically increases savings account usage, but differentially by gender and measures of household bargaining power. Dupas and Robinson (2013a) show that opening savings accounts increased savings accumulation for women, not men, in their context.

Much of the recent empirical evidence has focused on the behavioral channels. Ashraf et al. (2006) demonstrate that commitment accounts offered by banks can dramatically increase savings balances. Brune et al. (2013) show that opening savings accounts (with some of the features of commitment devices) leads to increased savings and also a range of business improvements. Offering reminders to save to combat inattention has proved particularly effective in the field. Karlan et al. (2010) and Kast et al. (2012) find large effects of reminders on savings. Kast et al. (2012) also test a peer monitoring technology using self help group peers and also self-selected "savings buddies". They find that the peer treatment works, but no better than the reminders. Our individual level BC treatment bundle and endogenous monitor treatment confirm these findings. The endogenous monitors are, loosely, analogous to self-selected "savings buddies". While they do increase savings relative to a pure control group, they are no more effective than an individual BC bundle (which has reminders built in by construction). However, our exogenous monitor treatment and results on the massive impacts of monitor centrality are wholly new to the literature. They suggest that, indeed, there are large gains to be had by leveraging peer monitors if done in the right way.

Meanwhile, recent research has also underscored the importance of social networks in understanding developing economies. Village networks broadly provide two services: (i) they serve as the surface on which information diffusion and social learning happens and (ii) they allow agents within the network to sustain cooperative behavior even with limited access to formal contracts. Empirical analyses have shown the importance of network structure in spreading information in a variety of contexts in development: health (Kremer and Miguel, 2007), job referrals (Beaman and Magruder, 2009), agriculture (Conley and Udry, 2010b), consumer choice (Miller and Mobarak, 2013), microfinance (Banerjee et al., 2013), and insurance Cai et al. (2013)). Banerjee et al. (2013) show that in an information diffusion process, the network centrality[11] of the source of the information (here about a new microfinance product) affects the extent to which information about microfinance spreads through the village network.

Aside from diffusion, lab experiments and lab experiments in the field have shown that networks can also help to sustain cooperation (Chandrasekhar et al. (2013) and Breza et al. (2013)), aid in coordination (Choi et al. (2011), Kearns et al. (2006)), and facilitate trade (Choi et al, (2013), Judd and Kearns (2008)).[12] In a related lab experiment conducted in the field, Breza et al. (2013) provide evidence that signaling may matter significantly and that its effect may vary significantly with the centrality of parties within a social network.

They interpret this as evidence that women are savings constrained. Chin et al. (2012) finds large effects of account opening assistance in a population of Mexican migrants in Texas.

[11]They show that eigenvector-type centrality notions are best associated with the depth of information spreading.

[12]These experimental studies all reference a large body of theoretical work. See Jackson (2008) for a general discussion.

Breza et al. (2013) study the addition of a third party judge to a sender-receiver investment game and show that when the judge, who has access to a punishment technology, is more central in the network, this generates significantly more efficient behavior in the sender-receiver interaction. Our savings intervention attempts to extend the results of the lab experiment in the field to a long-run, high-stakes application of saving in rural communities.

The bulk of the empirical literature has focused on information transmission in various forms (e.g., information about health products, financial products, job opportunities) or on lab-type analyses of coordination or overcoming lack-of-commitment. Our analysis focuses on a different, but important aspect in our view: a social network can be used as a signaling device. The threat of reputation loss is greater with certain individuals in the network and therefore individuals are more inclined to behave correctly in their presence. Our study makes use of a rigorous randomized controlled trial to study network effects in such an environment, which has a natural signaling interpretation.

More broadly, our analysis contributes to an understanding of informal institutions such as RoSCAs and self-help groups. We interpret our findings as providing evidence about the extent to which signaling may matter in sustaining cooperative behavior in these institutions. Additionally, our findings suggest that more attention ought to be paid to thinking about how networks may be powerful tools in policy design. Given the size of the network effects, it raises the question as to how one may obtain network data in a cost-effective manner (a subject that the authors and others are pursuing in future research).

The remainder of the paper is organized as follows. In section 2, we describe the experimental subjects, network and survey data sources, along with the experimental design. In section 3, we present a model of signaling on networks, while section 4 displays our main results. Section 5 examines robustness and extensions, and section 6 concludes.

2. Data and Experimental Design

2.1. **Setting.** Our overall sample frame consists of 3,000 individuals from 60 villages located in Karnataka, India which range from a 1.5 to 3 hour's drive from Bangalore. Each study village is within 5 km of a physical bank branch, and each bank branch is legally mandated to offer an interest-bearing "no frills" account with no minimum balances or transactions fees. Additionally, 35% of the study villages had a post office within the village boundary. The post offices offer interest-bearing "no frills" accounts as well. However, our baseline shows that the use of these branches is quite low. Only a quarter of households had an account at baseline. Figure 3 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 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.

As our primary goal is to understand which members of society serve the role of savings monitor most effectively and which individuals are chosen by their peers to play the role of monitor, we chose the study villages to coincide with the demographic and social network data set previously collected in part by the authors. The data is described in detail in Banerjee et al. (2013) and Jackson et al. (2010). In our field experiment, we match participants to this unique data set.

The graph represents social connections between individuals in a village with twelve dimensions of possible links, including relatives, friends, creditors, debtors, advisors, and religious company. We work with an undirected and unweighted network, taking the union across these dimensions, following Banerjee et al. (2013) and Chandrasekhar et al. (2013). As such, we have extremely detailed data on social linkages, not only between our experimental participants but also about the embedding of the individuals in the social fabric at large. We can use the different dimensions of relationships to differentiate an individual's risk sharing network as well. 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. $G_r = (V_r, E_r)$ is a graph consisting of vertices $V_r = \{1, ..., n\}$ 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 and the GPS coordinates of respondent 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.** Figure 1 pictorially represents our experimental design and Figure 2 presents a timeline. Study participants are randomly selected from an existing village census database 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 baseline survey, which includes questions on historic savings behavior, income sources and desire to save. We keep track of non-takers and survey them at the end of the six-month savings period, when we also survey the pure control group.

Potential savers establish a six-month savings plan. Importantly, this plan is established before the saver knows whether they are assigned to the individual treatment (BC bundle) or a monitor treatment. 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 up to Rs. 300 (~$6) for participating.

From the pool of consenting participants and attendees of the village meeting, we randomly assign savers to one of three treatments. The three treatments assigned to savers are (see Figure 1.E)[13]:

T1 : Individual or business correspondent bundle treatment (Randomized at the individual level)

T2 : Peer monitoring with random matching (Randomized at the village level)

T3 : Peer monitoring with endogenous matching (Randomized at the village level)

Note that we do observe attrition between the baseline and the village meeting. However conditional on attendance, attrition at this stage of the experiment is uncorrelated with final treatment status. In order to compare the basic individual treatment with the pure control group, we survey attriters to record their endline savings balances.

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 one post office to offer as choices to the savers. Savers are allowed to choose one or another. Savers who already have bank or post office accounts are offered the chance to open another account. The post office accounts are opened at the nearest post office branch location, generally within a 3km walk of each village. In fact 35% of 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, and agree to expedite our savings applications and process them in bulk. 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 identification documents for account opening and submit the applications in bulk.

Savers in the individual treatment (T1) are visited on a fortnightly basis. Our surveyors check the post office or bank passbooks and record balances and any transactions made

[13] Let T0 denote the pure control treatment.

in the previous 14 days and also remind savers of their goals. These home visits serve as strong reminders to save. Some participants report that these visits are very motivating. We should note that in no treatment do our surveyors collect deposits on behalf of the savers. This is the one large departure from the business correspondent model. As a result, our estimates should serve as a lower bound of the effects of that model on savings.

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. Savers in T2 are also visited fortnightly by our surveyors. However, our surveyors then pay visits to the homes of the monitors. During these visits, the monitors are shown the savings balances and transactions of their savers. At the end of the savings period, monitors receive incentives based on the success of their savers. Monitors are paid Rs. 100 if the saver reaches at least half the goal, and an additional Rs. 200 if the monitor reaches the full goal.[14]

The peer treatment with endogenous matching (T3) is identical to T2, except for the means of assigning monitors to savers. In this treatment, individuals are allowed to choose their monitor from the pool of potential monitors. 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 identical 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.

Figure 4 presents the histogram of savings goals, censoring the top 5%.[15] 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% outliners. In most 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.

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, 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. Randomization order is clearly exogenous and therefore

[14]We had initially wanted to offer treatments without incentives for monitoring, 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 process.

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

establishing causal effects becomes simpler. Additionally, it allows us to systematically explore what 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.

Naturally, the degree to which the environment studied in the allocation mechanism literature decribes our environment to first order determines the degree to whch this intuition is relevant for our case. Of course, the result also suggest that 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.

At the end of the 6-month savings period, we administer an endline to all savers and monitors to record total household savings and also information about interactions between savers and monitors (Figure 2). Again we 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. Finally, we also conduct an endline survey of the pure control group and of a random subset of savings group attriters. The key variable in this group is again total individual and household savings. 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. Table 1 also shows differences in the final sampled population decomposed between T1, T2, and T3.

3. Framework

We provide a simple, stylized framework in order to highlight the role of the embedding of the saver-monitor pair in the network. Section 3.1 introduces key attributes of the network that we use in our analysis: monitor centrality and saver-monitor social proximity. By construction, our exogenous treatment randomly assigns relative network positions of savers and monitors and enables us to identify the effect of centrality and proximity, conditional on all observables. These effects by themselves are of both policy interest and also contribute to our basic understanding of the social netwrok underpinnings of savings peer effects.the role of social networks influencing the extent of a peer effect in savings behavior.

To further clarify our thinking, in section 3.2 we model the effect of a saver-monitor's embedding i'n the network through a two-period game. The first period captures the saver-monitor interaction and establishes whether the saver saves a high amount or a low amount. The second period captures the world outside the savings exercise – it represents future interactions within the village. The most natural model for this is a simple Spence (1973) style signaling model. By attaining a high level of savings (or meeting one's goal), the saver demonstrates to the monitor that she is responsible. (In fact, it was actually a member of a village who originally suggested this experimental design to us, citing the idea that reputation about individuals hitting their savings goals may be leveraged to help encourage saving accumulation.)

The core idea of the model is that individuals come in two kinds: responsible or irresponsible. A responsible individual can be interpreted as an individual who is able to overcome (with effort) her time inconsistency, temptations, or inattention issues. Of course, an individual's responsibility matters for interactions across all walks of life. It is important to other individuals when considering futuer interactions with the saver.

An individual's ability to accumulate high savings (as is her goal) will signal her type. In the first period, the saver decides whether to save a high or low amount. This decision is observed by the monitor, who then may (stochastically) pass this information to other community members. In the second period, which represents future interactions of the saver with the remainder of her community, the saver has the opportunity to engage in a task with another member of society (whom she meets through the network). The return to the task depends on the saver's type. If she is a responsible type then the return is high whereas it is lower if the is irresponsible.

Our model, therefore, is an extremely simple network-based signaling model. Individuals pass information or meet along the network according to a simple stochastic process. Individuals interact (or pass information) randomly through the network along edges, independently, with some probability. With this simple structure on interactions, the signaling model predicts that, ceteris paribus, (a) higher central monitors generate a greater share of savers saving high amounts and (b) greater social proximity between savers and monitors generates a greater share of the saver saving a high amount.

While the model highlights the natural story – signaling – as a primary reason as to why we might expect monitor centrality and saver-monitor proximity to both be positively correlated with a saver's savings outcome, there are clearly other stories that can be naturally incorporated into such a framework. For instance, an individual may value demonstrating responsibility to a more central monitor simply because she is more likely to meet the more central monitor in the future, conditional on a distance. To take another example, individuals may get warm glow from being responsible in front of another, and simply gain more warm glow from behaving well in front of a more network-central individual since they may be more likely to pass on this information. Ultimately, our claim is not that the mechanism is signaling and only signaling, in the specific way we have described, but that the framework provided here suggests a natural and parsemonious framing of reputation-based channels.

3.1. **Networks: A Primer.** We are interested in how the village network structure impacts the saver-monitor interaction and subsequent savings. We focus on two notions of the social network: eigenvector centrality and social proximity. These will come up in the next subsection where we introduce the formal model. We define them first here and describe some intuition. We use the eigenvector centrality of an individual's household in the social network as a measure of her importance. Formally,

$$\lambda_1 (\mathbf{A}) e = \mathbf{A} e$$

where λ_1 is the maximal eigenvalue of the adjacency matrix and $e = e(\mathbf{A})$ is the associated eigenvector. Importantly, this notion of centrality represents a node's importance in a random-walk process through the graph. As such, it measures a node's importance in information transmission: more central nodes may be better able to pass information. For a more general discussion about eigenvector centrality in network economic models, see Jackson (2008).[16]

The social distance between two individuals is the shortest path (if one exists and is finite) between the two through the graph.[17] Formally,

$$\gamma_{ij} (\mathbf{A}) = \operatorname{argmin}_{\ell \in \mathbb{N}} \mathbf{1} \left\{ \left[A^{\ell} \right]_{ij} > 0 \right\}.$$

The core idea that we seek to explore is whether central monitors generate better savings behavior and whether far monitors generate worse savings behavior. We describe our model in the next subsection.

The model we develop will depend only on the network structure and a single parameter θ which represents the probability of two nodes interacting. 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 (and meetings) along the network occur in the following manner. Given $\mathbf{A}$, there is some probability θ,

[16] See also DeMarzo et al. (2003),Golub and Jackson (2009), Golub and Jackson (2010), Hagen and Kahng (1992).

[17] Two nodes are reachable if they are on the same connected component and clearly distance conditional on reachability is the relevant concept.

that a given piece of information crosses link ij.[18] Let

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

and observe that it counts the expected number of times a piece of information starting from node j hits node k. The advantage of this form is that it allows us to make predictions on how both centrality of a given node and the distance between the nodes will impact our signaling game. This is the same modeling structure used in Banerjee et al. (2013). As noted there, taking the limit as $T \to \infty$ with $\theta \geq \frac{1}{\lambda_1}$ leads to a vector $\lim_{T\to\infty} \sum_{t=1}^{T} (\theta\mathbf{A})^t \cdot 1 \propto e(\mathbf{A})$.

3.2. **A Signaling Model of Peer Monitoring.** Consider a two-period signaling game, wherein parties are embedded in a social network. A saver i is either a "low" or "high" type. Low types are irresponsible and find it differentially more costly to accumulate the high levels of savings in the face of time inconsistency, temptations or inattention. 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 stochastically observes the saver's outcome. Let o_{ji} denote this observation and assume $o_{ji} = 1$ with probability $p_{ij}(\mathbf{A},\theta)$ and $o_{ji} = 0$ otherwise.
- $t = 1$: 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.
- $t = 2$: 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 an irresponsible individual to overcome their time inconsistency, temptations or inattention and accrue high savings. However, monitors may imperfectly observe this. Even though our surveyors do inform them bi-weekly, one can imagine that there is inattention on the part of the monitor, or that the surveyor visit serves as a reminder to the monitor to act. Hence

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

whether j observes i's savings depends on whether j meets i through the graph in our model.

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 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 leads to transitions from only pooling equilibria to separating equilibria (applying the intuitive criterion).

Finally, we describe how we treat the probabilities. In what follows, we take $T \to \infty$ and treat the probability of a meeting (or information transmission) between u and v as being proportional to this limit object. That is, it will be proportional to the expected number of times (say) that a piece of information starting from u and traversing links with probability θ will hit u. This is a common strategy in network analysis and makes the problem particularly parsimonious.

Lemma 3.1. *A separating equilibrium exists if*

$$y_H - \frac{c_H}{p_{ij}\left(\mathbf{A},\theta\right)\cdot e_j\left(\mathbf{A}\right)\cdot e_i\left(\mathbf{A}\right)} > y_L > y_H - \frac{c_L}{p_{ij}\left(\mathbf{A},\theta\right)\cdot e_j\left(\mathbf{A}\right)\cdot e_i\left(\mathbf{A}\right)}.$$

Otherwise, only a pooling equilibrium exists in which all types pool on $s_i = 0$.

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

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

corresponds to a separating equilibrium.[19]

Next we decompose P into its constitutent parts. The expected number of times that a given node k receives a signal sourced from j is given by $\left[\sum_t \left(\theta A\right)^t\right]_{jk}$. Integrating over

[19]This follows from the fact that

$$P y_H + (1-P)\,0 - c_H > P y_L + (1-P)\,0 > P y_H + (1-P)\,0 - c_L.$$

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.

all the k, and taking $T \to \infty$, we have

$$\lim_{T\to\infty}\left[\sum_t (\theta\mathbf{A})^t \cdot 1\right]_j \propto e_j(\mathbf{A},\theta).$$

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

$$\lim_{T\to\infty}\left[\sum_t (\theta\mathbf{A})^t \cdot 1\right]_i \propto e_i(\mathbf{A},\theta).$$

Therefore, the payoffs given a monitor are

$$\mathrm{E}\left[o_{ji}r_{jk}m_{ik}|j\right]y_H - c_H > \mathrm{E}\left[o_{ji}r_{jk}m_{ik}|j\right]y_L > \mathrm{E}\left[o_{ji}r_{jk}m_{ik}|j\right]y_H - c_L.$$

By independence we have

$$\begin{aligned}
\mathrm{E}\left[o_{ij}r_{jk}m_{ik}|j\right] &= \mathrm{E}\left[o_{ji}|j\right]\mathrm{E}\left[r_{jk}|j\right]\mathrm{E}\left[m_{ik}|j\right] \\
&= p_{ij}(\mathbf{A},\theta)\cdot\sum_k\left[\sum(\theta\mathbf{A})^t\right]_{jk}\cdot\sum_k\left[\sum(\theta\mathbf{A})^t\right]_{ik} \\
&= p_{ij}(\mathbf{A},\theta)\cdot e_j(\mathbf{A},\theta)\cdot e_i(\mathbf{A},\theta).
\end{aligned}$$

□

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 IC.

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

1. *Centrality of monitor:*
   - (a) *Holding p_{ij} and e_i fixed, for e_j sufficiently high, there exists a separating equilibrium with high types attaining s_H.*
   - (b) *With the same parameters fixed, for e_j sufficiently low, there remains only a pooling equilibrium on s_L.*
2. *Saver-monitor proximity:*
   - (a) *Holding e_i and e_j fixed, for p_{ij} sufficiently high, there exists a separating equilibrium with high types attaining s_H.*
   - (b) *With the same parameters fixed, for p_{ij} sufficiently low, there remains only a pooling equilibrium on s_L.*

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.

4. Results

4.1. **Average Treatment Effects.** Table 2 presents treatment effects of the monitoring treatments versus the baseline bundle of setting a savings goal and receiving biweekly visits from our surveyors. Panel A provides OLS estimates of the pooled monitor treatment effect, while Panel B decomposes the results between random and endogenous monitors. In the baseline treatment group, only 7.9% of savers reach their goals. We will be able to measure the impact of this baseline reminders treatment once data entry concludes, but given this low level, the results are unlikely to be very economically significant. However, we find that savers with monitors increase goal attainment by 72% relative to the baseline reminders group. Column 1 gives the pooled treatment effect of adding a monitor, which is approximately 6 percentage points in the LPM regression. Analyzed separately, both monitor treatments significantly increase goal atainment in a comparable magnitude, (6.52pp in the random treatment vs 4.89pp in the endogenous treatment relative to reminders group). Column 2 displays the likelihood of reaching half or more of the savings goal, and the treatment effects look similar. While in the treatment group, the likelihood of this outcome is more than double the full goal attainment likelihood, the treatment effects are quite similar.

Columns 3 through 6 show that the monitoring treatments have consequences for levels of savings. Column 4 and 5 show OLS regressions of total savings (in levels) in the target account on treatment status, while columns 5 and 6 show results on savings in excess of the saving goal (in levels). We run these specifications in levels due to the high number of individuals saving Rs 0 in their target accounts. In both columns 4 and 6, we drop observations in the top 1 percentile of saving goals. Qualitatively, the results are similar in all four columns. While the pooled monitor treatments increase savings by more than Rs 300, the coefficients are not significant. Separating the effects of the endogenous versus the random monitors, we observe that the random monitor treatment has larger and statistically significant effects on each measure of savings.

We find quite meaningful impacts of receiving a monitor on reaching the savings goal in the target account, but do these results really represent increases in overall savings, or are they simply the product of moving funds from other vehicles into the target accounts? Table 3 details the treatment effects of having a monitor on total end savings of the saver across all savings vehicles. Because savings balances are generally non-zero and because they vary dramatically across the sample, we run the regressions in the log of the ending savings balance. Our results confirm the pattern of table 2, that on average, monitors have a large effect on savings (33% increase), and that this effect tends to be even larger when the monitor is randomly assigned (40-60% increase). This finding suggests that monitors are actualy playing a role in household savings accumulation, rather than savings reallocation.

We also provide suggestive evidence that the average random monitor is more effective than the average monitor who is selected by the saver. The rates of goal attainment are comparable in the random and endogenous treatments (extensive margin) , and in some specifications, the savings amounts (intensive margin) are detectably larger. Adding a random monitor to the baseline treatment increases savings by Rs 790, a 0.32 standard deviation effect size.

4.2. **Effective Monitor Characteristics.** We find that the position of the (randomly chosen) monitor within the social network has large effects on savings behavior. We focus on two notions of network position: first, the social distance between the households of the saver and monitor and second, the eignevector centrality of the monitor's household. Table 4 presents the results of regressions of savings outcomes on network statistics. Each panel in the table contains a different regression specification. The outcomes of interest are goal attainment, half goal attainment, excess savings (savings - goal) in the target account, total savings across all vehicles, and log savings across all vehicles.

Broadly, we find that both monitors with close social proximity to the savers and monitors with high network centrality improve the savings outcomes of the savers. Moving from an unreachable monitor to a monitor of social proximity 1 (direct links) increases goal attainment by 17 percentage points and increases total savings by 1.3 log points (Panel B). These magnitudes are quite large and robust to controling for centrality as well. There is weak evidence (insignificant coefficients) that being paired with a relative may even increase goal attainment above and beyond simply being paired with a direct link.

A one standard deviation increase in eigenvector centrality of the monitor (which is the normalized regressor presented in the table) increases savings goal attainment by 3.3 pp. This effect is of a similar magnitude to half the effect of assigning an average monitor. There are also significant effects on total savings outcomes across all vehicles. A one standard deviation increase in eigenvector centrality of the monitor increases total savings by 0.22 log points. In levels, this corresponds to an increase in total end savings of Rs. 1689. The effects of being assigned to a socially close or a high centrality monitor on total savings also remain significant in regressions that include both proximity and centrality characteristics simultaneously.

While the network position of any individual is not randomly assigned (although the matches between saver and monitor are), we believe that the network is impacting the success of the monitoring relationships. The regressions in Table 4 all contain controls for numerous saver and monitor demographic characteristics, proxies for saver and monitor wealth, and controls for the saver's savings goals.

In sum, the evidence suggests that both conditional on each other and a number of observable characteristics, monitor centrality and saver-monitor proximity matter immensely and in the directions anticipated by the model. The magnitude of these effects – especially the centrality effects – suggests that there may be large returns to policy-makers designing products that leverage the network in this manner.

4.3. **Endogenous Monitor Selection.** Turning to the endogenous treatment, we begin by asking whether savers choose monitors that are both more central and socially close when they have higher choice priority. Figures 5 and 6 provide graphical evidence. Figure 5 presents histograms of the social distance between saver-monitor pairs in both the endogenous and random 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.[20] Similarly, Figure 6 shows that endogenously-chosen monitors are also more central in the network than randomly matched monitors. Recall that not all available monitors are assigned a saver in both treatments. Thus, this gap is in part mechanical assuming (as in the model) that savers try to pick more central monitors when available.

Having demonstrated that in the endogenous treatments, individuals both choose more central monitors and more proximate monitors, we study whether individuals with lower choice priority have worse outcomes in the endogenous treatment. Observe that in the random treatment, choice priority (which now has no meaning since the pairings were random) should be uncorrelated with outcomes.[21] Table 6 displays whether individuals reached their savings goals as a function of their choice quantile. Columns (4) and (6) specifically focus on whether an individual had a low choice quantile, defined as being in the bottom third of the distribution. Having a low choice quantile in the endogenous treatment reduces the probability of reaching one's goal by 10pp, relative to a mean of 12.5%. Column (6) shows that there is a robust statistical difference in having a low choice quantile in the endogenous treatment versus random treatment (as expected, since choice quantile should have no appreciable effect in the random treatment).

Ultimately, this suggests that towards the end of the village meeting, good monitors are scarce in the endogenous treatments. Taken together with the results from Table 2, which show that the endogenous treatment does worse than the random treatment, this suggests that the allocation of monitors may be important.

There are two plausible and related mechanisms to explain such a finding. We think that it is likely that the act of outside agents choosing the pairs facilitated success in the random treatment relative to the endogenous. First, individuals may not feel that they

[20] A regression of social proximity on treatment type indicates that endogenous saver-monitor pairs are 0.14 proximity units closer than random pairs. The coefficient is significant at the 1% level. It is also the case that endogenous group members are much more likely to choose relatives. While only 3% of monitors are related to their savers in the random group, the fraction jumps to 14% in the endogenous group.

[21] Recall that participants were assigned to their monitors in random order in both treatment groups. While the order of assignment should not affect the quality of the monitor in the random group, individuals who select late in the endogenous group have fewer potential monitors from which to choose.

are able to ask their preferred monitors to help them to save. Peripheral individuals may be shy (or pay a psychological cost) to ask a central individual to monitor. Alternately, central monitors with many social obligations may similarly not be able to choose their preferred monitors because of binding obligations with others. Second, the very nature of our randomization may have changed the meaning of the savings monitor. Our randmization device may have given a more formal sort of authority to the monitors, for example.[22] We should note that the language to discuss the monitor and the scripts of the visits to the homes of the savers and monitors were not at all different between the random and the endogenous groups.. Both mechanisms suggest, however, that in our setting, centralizing and anonymizing the choice process fostered better outcomes. This may support interventionist policy even when building institutions from existing social structures.

Finally, we note that our endogenous treatment results echo the findings in Kast et al. (2013). In their savings buddies treatments, where individuals choose buddies with whom they save, while the treatment is effective, it is no more effective than an individual-level reminder treatment. Similarly, we find that endogenous monitors do help increase savings (relative to pure control), though not relative to the individual-level business correspondents bundle (which has an in-built reminder mechanism).

4.4. **Business Correspondents.** We find that the effects of the business correspondent treatment are quite large. Table 7 shows the intent to treat estimates of the business correspondent bundle on overall endline savings and also on log savings. In three of the four specifications, the treatment effect is positive and both economically and statistically significant. This is true in both specifications trimming 1% outliers and also in the log savings specifications. The 0.23 coefficient in the log savings specification suggests that the treatment increased savings balances by approximately 25% across the potential savers group. We note that 37.7% of potential savers actually took up one of the savings interventions. This would suggest a treatment on the treated estimate of approximately 0.6 log points (~81% increase). However, this ToT estimate should be taken as an upper bound. Many of the potential savers who opted to not attend the village meeting did receive our help in projecting future income and expenses and also in devising a savings goal. To the extent that this process is driving the results, then the ITT estimate of 0.23 is a better ToT approximation.

These results fit well with the existing evidence that account opening assistance (Schaner (2011), Dupas and Robinson (2013a), Chin et al. (2012), and Brune et al. (2013)) and reminders (Karlan and Zinman (2010) and Kast et al. (2012)) can be quite effective at spurring savings. This might suggest that there are positive spillovers from the BC model even if there is limited saving in the bank itself.

[22] See Fehr and Wilkening (2012) and the literature on process-based preferences.

5. Welfare, Robustness and Extensions

5.1. **Welfare.** Making inferences about saver welfare is quite difficult in settings such as ours as the key barriers to savings are likely to be behavioral.[23]

Our biggest hurdles in translating our experimental results into welfare judgments are twofold. First, we have demonstrated that having a more central monitor improves savings. However, a natural worry is that central monitors may present too strong of an incentive device to save. Having too rigid of a monitor may cause households to save during a negative income shock rather than smooth. Another, related worry is that by asymmetrically revealing savings within a target account, the saver may face undue pressure to make transfers to their monitor (or others) as a function of the stated amounts in the target account. Second, interpreting endogenous monitor choice as optimal is also difficult without knowing more about the underlying choice framework. Further, it is unlikely that a saver is able to fully understand the counterfactual savings (or utility) path under a different treatment or a different monitor.[24]

We argue that there are likely large welfare gains due to introducing a savings monitor – especially a high centrality savings monitor. First, we note that over 90% of savers who reach their goals actually surpass them. Second, over 65% of savers actually save in excess of 200% of their stated goals in their target accounts. This provides evidence that individuals are not facing undue pressure to save at any cost nor are they likely to be facing undue pressure to make undesired transfers to others as a function of the observable savings in the target account. If these were deep concerns, we should expect to see many more individuals reach the goal exactly and not surpass it so immensely. Recall that the framing of our experiment, by its very nature, was about a saver attempting to hit an ex-ante formulated goal, not to double it.

Third, the increased savings are not solely from the target account. In fact, 80% of individuals save in excess of 175% of their savings goal across all accounts (and 75% save in excess of 200% of their savings goal). Savers who hit their goals also tended to save considerably more in other accounts (other formal and informal vehicles). Recall that during the monitoring intervention, our research team did not provide information to the monitor about the saver's savings in any other account. The monitors were only told the verified amounts in the saver's savings *in the target account*. This implies that savers always had the option of moving savings away from other vehicles and into the target account. That is, if the monitor became too rigid, the saver could always dissave out of another savings vehicle such as cash under the mattress. In fact, empirical

[23] Kőszegi and Rabin (2007), Kőszegi and Rabin (2008), and Beshears et al. (2008) reflect on situations when revealed preference arguments fail and where observing choices of agents is not sufficient to compute the welfare of those agents.

[24] In a follow-up survey, we plan to investigate whether savers did have accurate beliefs about counterfactual monitoring parirings.

studies such as Giné et al. (2010) have shown how individuals can sometimes unwind their commitment devices when given the opportunity. The fact that when given a more central monitor, savers considerably increase their overall savings across accounts *outside* of the target accounts and moreover also surpass their stated goals considerably in their target accounts, suggests massive and willful gains in savings.

Taken together, the evidence suggests that savings monitors introduced welfare gains for those individuals who were induced to reach their goals by their monitors.[25] To not believe there are welfare gains for this group, one has to believe that individuals feel compelled to increase savings across other (unobserved) accounts while also being willing to demonstrate considerably higher savings in the target account as well, even though the individual would be better off doing neither. Given evidence presented, this does not seem likely.

5.2. **Incentives.** Due to power considerations, we were not able to vary experimentally the incentives offered to monitors for the goal attainment of savers. It is true that the incentives are relatively large when compared to the distribution of savings goals. One natural concern may be that savers and monitors collude because of the stakes of the incentives. This collusion may also be easier to facilitate as a function of the social network. After the completion of the savings period, some monitors were asked about the incentives. Several monitors shared that they viewed savings as being a personal decision on the part of the saver and that the incentives did not motivate them. Other monitors failed to claim their incentive prizes at the end of the experiment. A few other savers withdrew from their target accounts the week before the saving period ended, thus reducing the incentive paid to the monitors. Further, we do not see very much evidence of collusion in the data. The monitor only has incentives to collude, and similarly the saver only has the ability to hold up the monitor when the savings of the saver are marginal to the incentive. This occurs within Rs. 200 of the full saivngs goal and within Rs. 100 of the half saivngs goal.[26] Again very few savers hit their goals exactly (the minimum requirement for the monitor to receive the incentive), and savers saved not only in their target accounts, but across many vehicles. Gamers receiving no actual savings impetus from their monitors could have instead transferred funds temporarily into the target account from other accounts to ensure the payment of the monitor incentive. Finally, Table 9 shows that there savers do not rush to make deposits in the last two weeks of the savings period, as might be the case if monitors were lending funds to savers. Also, there is no evidence of bunching in the goal attainment data around Rs. 200 of the savings goal.

[25]Under our model, it is true that individuals who do not reach their goals send signals that they are of a "bad" type when they do not reach their goals. We cannot measure the costs of this signal in our setting. We should also note that our model abstracts from savers receiving any direct benefit from savings.

[26]Recall that monitors earned Rs. 300 if the saver reached her full goal, and Rs 100 if the saver reached half of her goal.

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.

5.3. **Alternative explanations.** We now discuss and rule-out several alterantive explanations for our results.

5.3.1. *Demographic characteristics caste.* To confirm that the network effects are not proxying for various demographic characteristics, we include a number of controls in our key tables. We control for various wealth indicators as well as caste. The network effects are robust even conditional on these controls suggesting that they are not simply proxies for this. Additionally, this fact – that network effects were not simply proxying for wealth, caste, and other demographics – has been demonstrated in other work by the authors as well (Breza et al. (2013)).

5.3.2. *Warm glow.* One alternate hypothesis is that savers save, not to generate information for their monitors, but simply to experience the warm glow of reaching a goal in a public way. For warm glow to explain our results, savers would need to feel more warm glow in the eyes of exactly the most network central individuals in the village, controling for demographic characteristics. While we cannot rule out such mechanisms in our setting, we do feel that the literature suggests that there is economic content in showing the monitor that the saver can reach his or her goal.

5.3.3. *Mistaken accounting.* A common issue in savings-related research is that mistakes in accounting – that is, changes in the composition of savings across accounts – are attributed to increases (or decreases) in savings. To investigate this, we make use of the fact that we obtained detailed expenditure data in the last month across several ex-ante defined categories (which we considered potential temptation goods in the baseline). Our expenditure categories are not complete, and we further did not ask savers about labor supply decisions. While this is an imperfect solution, since we do not have access to all expenditure data across all accounts, it is a partial solution and we hope to find reductions in some categories. If individuals work harder to save more,[27] then this will generate an underestimate of sources of savings.

Table 11 in the Online Appendix demonstrates that during the last month of the savings period (the only time when we happened to survey expenditures in detail), there were statistically significant declines in expenditures on festivals (Rs. 215), transportation (Rs. 141) and a statistically significant increase in tea (Rs. 27). While overall the expenditure data is somewhat noisy across the various categories, these numbers square with a back-of-the-envelope accounting exercise. Across all measures the point estimate

[27] Banerjee et al. (2009) find that households work harder when they receive microfinance loans. One explanation is that they increase labor supply to cover the weekly loan installment payments.

on log expenditures is 0.06 (though insignificant). Recall that the mean level of savings was Rs. 1700. This corresponds to an average savings of Rs. 283 per month, which is not dissimilar to the magnitude we find in this calculation. This, together with the fact that we find treatment effects not only in our target accounts but also across other (and all) savings vehicles, suggests that the identified effects are not just products of a mistaken accounting exercise which changes composition of savings across accounts, but a real, substantive change in savings behavior.

5.3.4. *Observable savings.* One may be concerned that is that part of the effect can be coming from the fact that we have introduced observable savings into a game between individuals where they may face hidden income plus hidden savings frictions. We sketch out the relevant environments and then turn to the data. For simplicity, consider two environments. In the first environment, say the saver typically borrows from the individual assigned to be her monitor. Access to an observable savings technology allows her to demonstrate repayment capacity as the stock of savings can serve as a sort of collateral or co-investment. Therefore, this would have a positive effect. (The only case where we would find the opposite effect is if the saver was in strategic default against the monitor.) In the next environment, say that the saver typically lends to the monitor. One might be worried that this allows the monitor to observe that the saver has the capacity to make a specific transfer in a given state when the saver otherwise may not have desired this information revelation. Finally, from the perspective of an informal insurance with hidden income and hidden savings perspective, insurance transfers are not possible and all transfers would be in the form of loans (Cole and Kocherlakota (2001)). However, introducing observable savings allows the saver to credibly place a lower bound on income realizations and therefore may help insurance-motivated transfers.

We note that we have detailed network data as to whether individuals have a financial relationship with their monitors. We look separately for effects when the saver borrows from the monitor or lends to the monitor and find no differential effects (Table 10 in the Online Appendix). We are not claiming, of course, that there is absolutely no partial effect of relaxing the information frictions by introducing a financial product that increases (partial) observability of income. However, we are arguing that in our context we are unable to detect any differential net effects.

5.3.5. *Mimicry.* A natural worry is that the saver-monitor relationships lends itself to mimicry. Perhaps more central monitors are pre-disposed to having better savings behavior and the monitor centrality effect is simply proxying for mimicry. Table 9 in the Online Appendix demonstrates that monitor baseline savings and projected savings over the subsequent six-month horizon have no statistically significant relationship with saver outcomes. This demonstrates that mimicry is not the likely channel and that having a monitor with strong saving habits does not generate better saver outcomes.

6. Conclusion and Discussion

We conduct a field experiment in rural Indian villages with the goal of increasing the use of savings vehicles. We test two product innovations that employ combinations of reminders, goal setting and peer monitoring. We find that both our business correspondent bundle and peer monitoring have sizeable effects in increasing savings goal attainment. Further, randomly chosen monitors appear to be more effective than individually chosen monitors. Additionally, among the randomly chosen monitors, there is substantial variation in their efficacy. Monitors that are in the upper tercile of eigenvector centrality generate a 20% rate of reaching one's six-month savings goal, relative to a benchmark of only 7% under the reminders only treatment. The results suggest that a simple modification of the business correspondents model of savings collection, combined with a carefully selected monitor to partner with the saver, can immensely encourage better savings behavior.

More generally, our findings suggest that in settings involving peer monitoring, intervention may be required to obtain socially optimal results. Socially important individuals may find it easier to ask favors of other high status individuals. The results here suggest that programs which engage in formal assignment of parties to roles may help alleviate some of this cost and therefore generate more efficient outcomes overall.

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*.

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.

Banerjee, A., A. G. Chandrasekhar, E. Duflo, and M. Jackson (2013): "The Diffusion of Microfinance," *Science*, 341 (6144).

Banerjee, A., E. Duflo, R. Glennerster, and C. Kinnan (2009): "The miracle of microfinance? Evidence from a randomized evaluation," *Department of Economics, Massachusetts Institute of Technology (MIT) Working Paper, May*.

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

Beaman, L. and J. Magruder (2009): "Who gets the job referral? Evidence from a social networks experiment," .

Bernheim, B. D., D. Ray, and S. Yeltekin (2013): "Poverty and Self-control," Tech. rep., National Bureau of Economic Research Working Paper.

Bernheim, D., D. Ray, and S. Yeltekin (1999): "Self-control, saving, and the low asset trap," .

Beshears, J., J. J. Choi, D. Laibson, and B. C. Madrian (2008): "How are preferences revealed?" *Journal of Public Economics*, 92, 1787–1794.

Breza, E., A. G. Chandrasekhar, and H. Larreguy (2013): "Punishment on Networks: Evidence from a Framed Field Experiment," *Working Paper*.

Brune, L., X. Giné, J. Goldberg, and D. Yang (2013): "Commitments to Save: A Field Experiment in Rural Malawi," *Working Paper*.

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*.

Chin, A., L. Karkoviata, and N. Wilcox (2012): "Impact of Bank Accounts on Migrant Savings and Remittances: Evidence from a Field Experiment," *Mimeo*.

Choi, S., D. Gale, S. Kariv, and T. Palfrey (2011): "Network architecture, salience and coordination," *Games and Economic Behavior*, 73, 76–90.

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

Collins, D., J. Morduch, S. Rutherford, and O. Ruthven (2009): *Portfolios of the poor: how the world's poor live on two dollars a day*, Princeton University Press.

Conley, T. and C. Udry (2010a): "Learning about a new technology: Pineapple in Ghana," *The American Economic Review*, 100, 35–69.

——— (2010b): "Learning about a new technology: Pineapple in Ghana," *The American Economic Review*, 100, 35–69.

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

Fafchamps, M. and S. Lund (2003): "Risk-sharing networks in rural Philippines," *Journal of Development Economics*, 71, 261–287.

Fehr, E. and T. Wilkening (2012): "The lure of authority: Motivation and incentive effects of power," .

Giné, X., J. Goldberg, D. Silverman, and D. Yang (2010): "Revising commitments: Time preference and time-inconsistency in the field," *Unpublished manuscript, Department of Economics, University of Michigan, Ann Arbor, MI*.

Golub, B. and M. Jackson (2009): "How homophily affects learning and diffusion in networks," *Working Papers*.

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

Greif, A. (1993): "Contract enforceability and economic institutions in early trade: The Maghribi traders' coalition," *The American economic review*, 525–548.

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.

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., M. McConnell, S. Mullainathan, and J. Zinman (2010): "Getting to the Top of Mind: How Reminders Increase Saving," *NBER Working Paper*.

Karlan, D. and J. Zinman (2010): "Expanding Microenterprise Credit Access: Using Randomized Supply Decisions to Estimate the Impacts in Manila," Innovations for Poverty Action 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," Tech. rep., National Bureau of Economic Research.

Kast, F. and D. Pomeranz (2013): "Do Savings Constraints Lead to Indebtedness? Experimental Evidence from Access to Formal Savings Accounts in Chile," *Harvard Business School Working Paper*, 14-001.

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

Kőszegi, B. and M. Rabin (2007): "Mistakes in choice-based welfare analysis," *The American economic review*, 97, 477–481.

——— (2008): "Choices, situations, and happiness," *Journal of Public Economics*, 92, 1821–1832.

Kremer, M. and E. Miguel (2007): "The Illusion of Sustainability*," *The Quarterly Journal of Economics*, 122, 1007–1065.

Laibson, D. (1997): "Golden eggs and hyperbolic discounting," *The Quarterly Journal of Economics*, 112, 443–478.

Miller, G. and A. M. Mobarak (2013): "Intra-Household Externalities and Low Demand for a New Technology: Experimental Evidence on Improved Cookstoves," *Working Paper*.

Pathak, P. A. and J. Sethuraman (2011): "Lotteries in student assignment: An equivalence result," *Theoretical Economics*, 6, 1–17.

Rutherford, S. and S. Arora (2009): *The poor and their money*, Practical Action Publ.

Schaner, S. G. (2011): "The Cost of Convenience? Transaction Costs, Bargaining Power, and Savings Account Use in Kenya," Tech. rep., Mimeo.

Shipton, P. (1990): "How Gambians save - and what their strategies imply for international aid." Tech. rep., WPS 395, Agricultural and Rural Development Department, The World Bank.

Townsend, R. M. (1994): "Risk and Insurance in Village India," *Econometrica*, 62, 539–591.

Figures

Figure 1. Experimental Design and Randomization

(a) Village

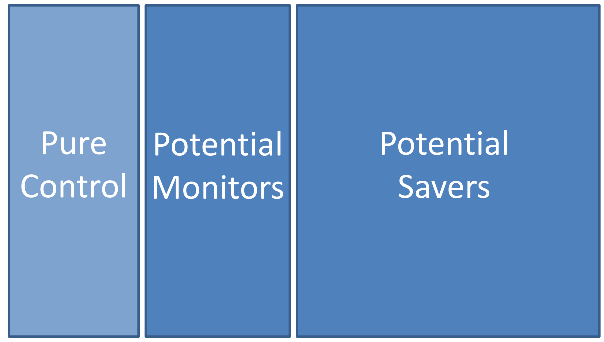

(b) Randomization of individuals to pure control, monitor pool or savers pool.

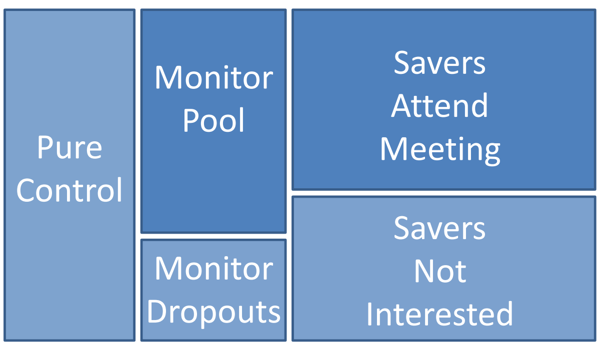

(c) Participating samples.

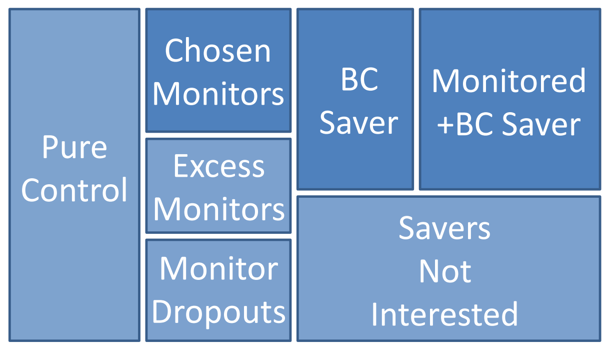

(d) Allocation of savers and monitors to treatments.

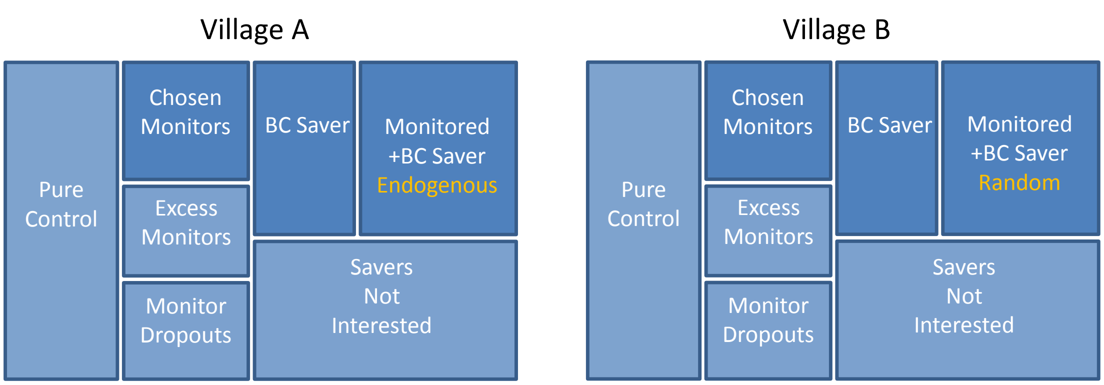

(e) Village-level randomization. Village A is randomly assigned to endogenous monitoring treatment. Village B is randomly assigned to exogenous monitoring treatment.

FIGURE 2. Timeline of Experiment

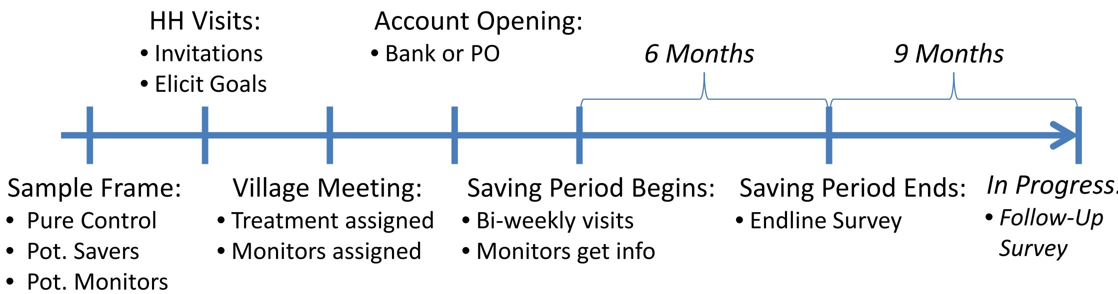

FIGURE 3. Intensity of Use of Available Savings Vehicles

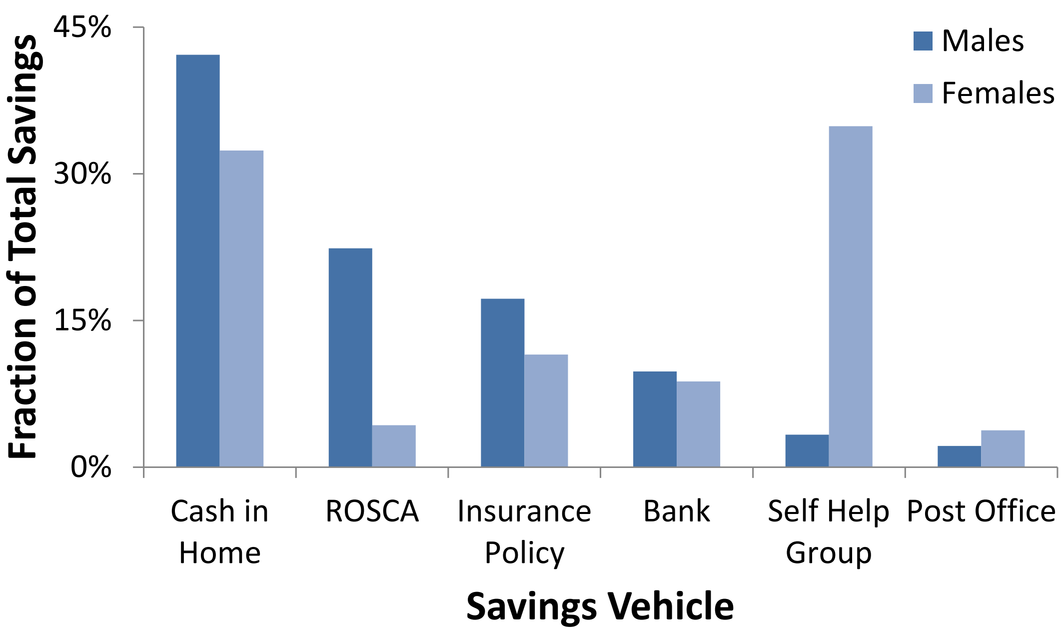

Figure 4. Histogram of Baseline Savings Goals

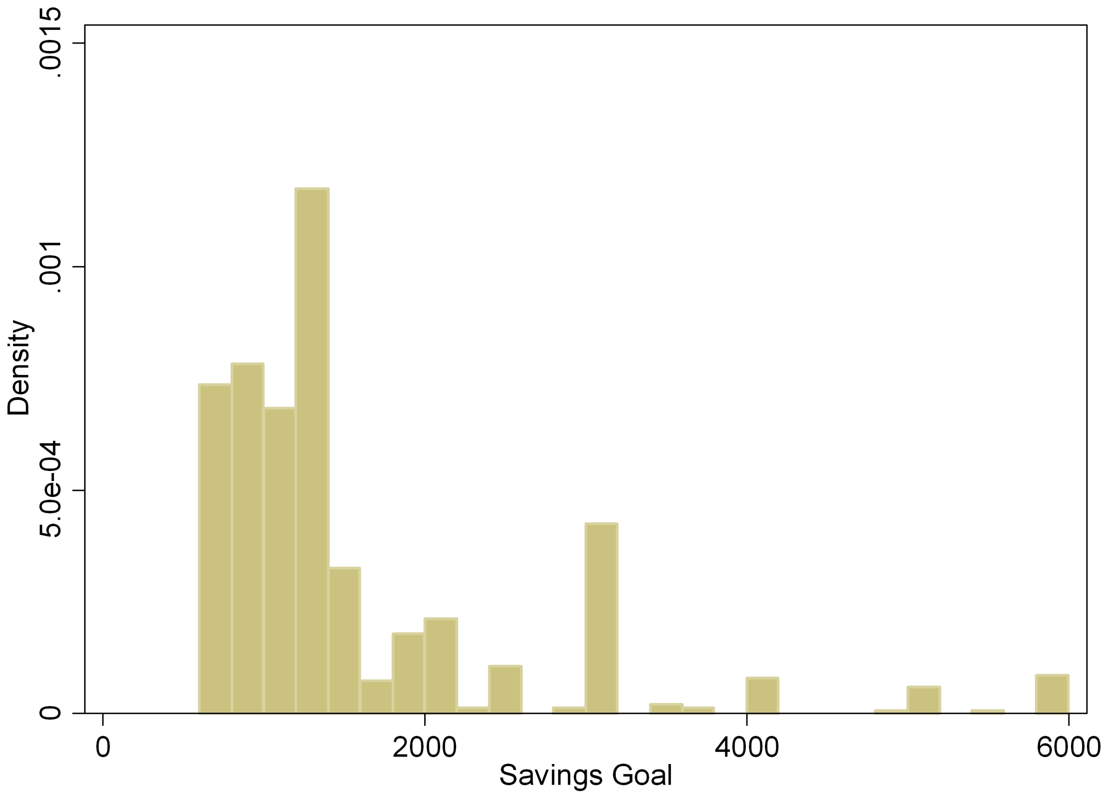

The figure shows the distribution of the baseline savings goals. We clip the top 5% tail of the distribution to make the figure more readable.

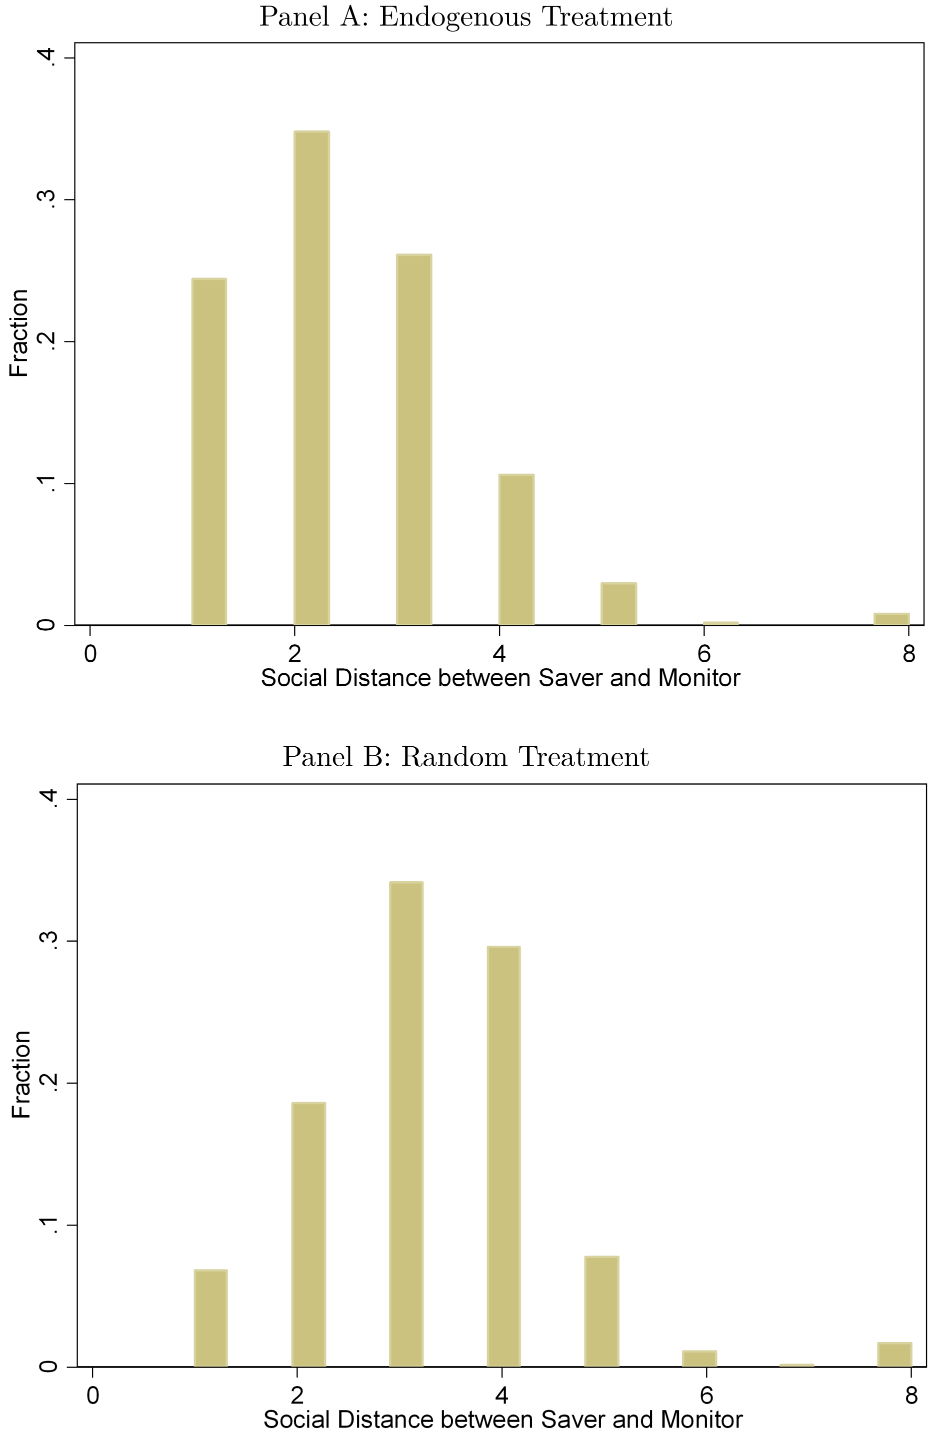

FIGURE 5. Social Distance Between Saver and Monitor by Treatment

Panel A: Endogenous Treatment

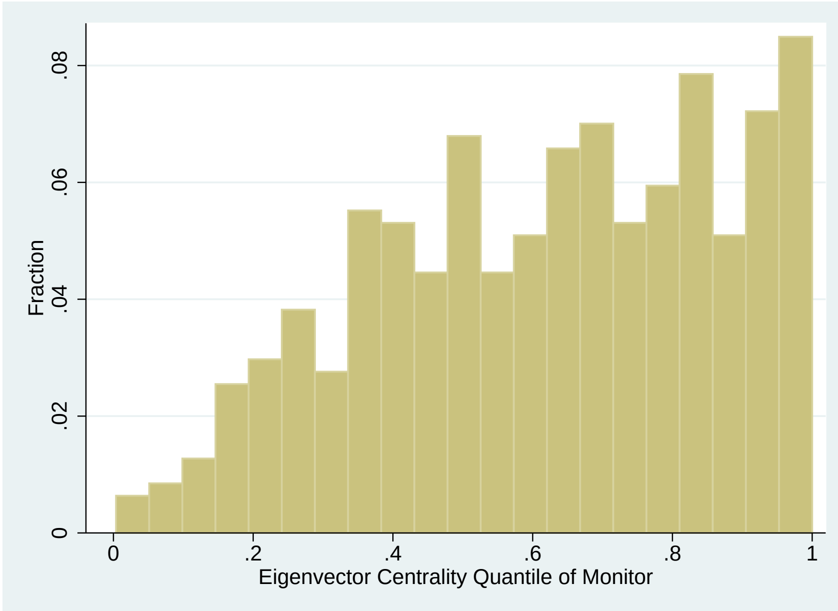

Panel B: Random Treatment

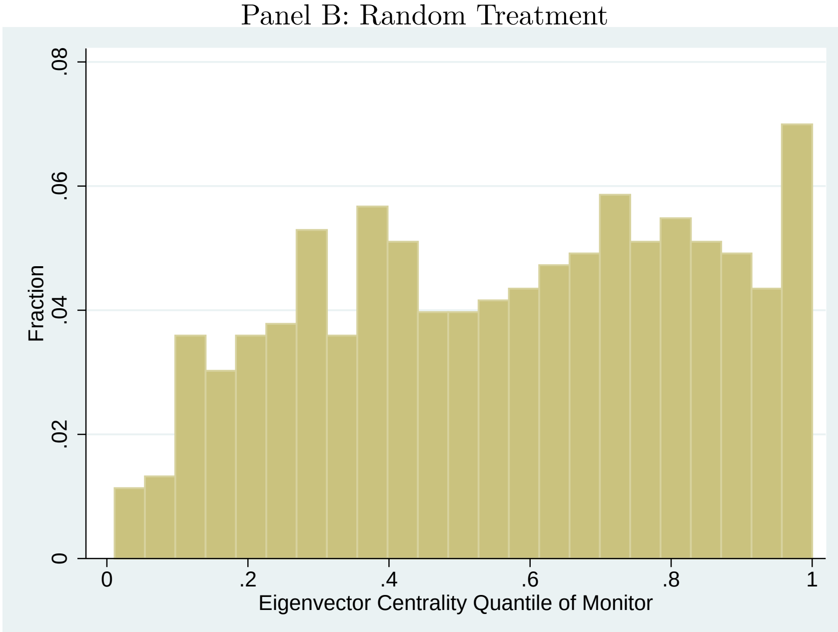

FIGURE 6. Centrality Quantile of Monitor by Treatment

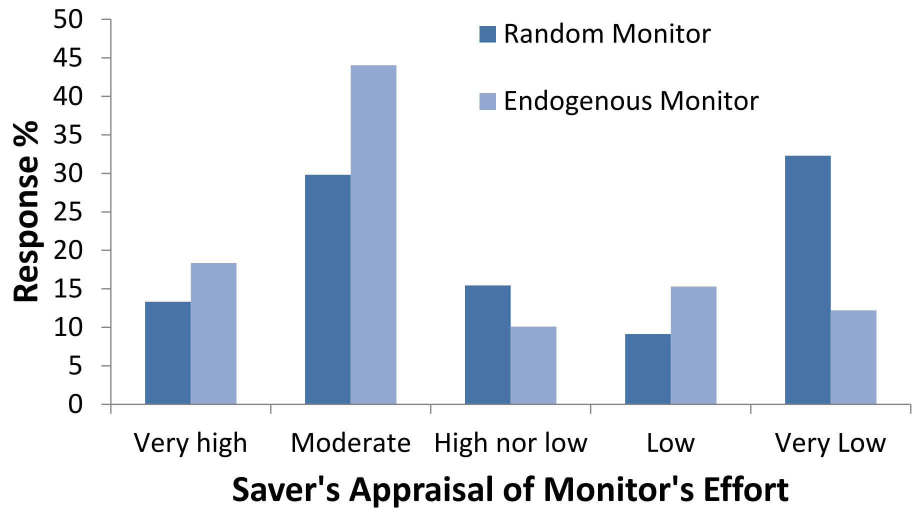

FIGURE 7. Monitor Effort Appraisal by Savers

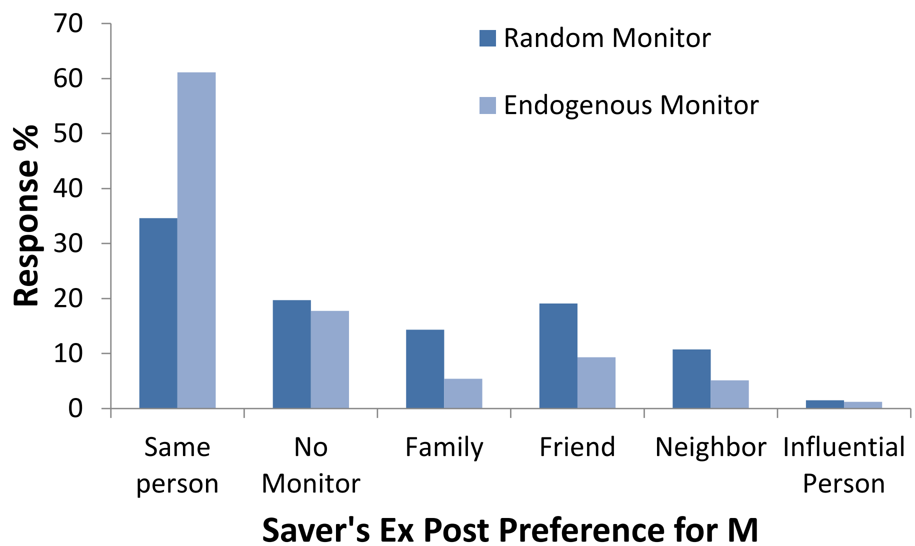

FIGURE 8. Saver's Preference for Monitors in Future

Tables

Table 1. Summary Statistics, Treatment Assignment, and Attrition

	Treatment (Village Meeting Sample)				Treatment (Endline Sample)			
Dependent Variable	Baseline	Endogenous	Random	Obs.	Baseline	Endogenous	Random	Obs.
Age	33.09	0.207	-0.147	1,304	33.11	0.348	0.360	971
	(0.385)	(0.527)	(0.458)		(0.441)	(0.621)	(0.496)	
Female	0.756	-0.0249	-0.0411	1,304	0.773	-0.0150	-0.0392	971
	(0.0243)	(0.0340)	(0.0316)		(0.0289)	(0.0350)	(0.0370)	
Married	0.857	-0.0255	-0.0287	1,304	0.863	-0.0390	-0.0120	971
	(0.0192)	(0.0272)	(0.0208)		(0.0236)	(0.0314)	(0.0247)	
Widowed	0.0358	0.0155	0.00954	1,304	0.0386	0.0247	0.00671	971
	(0.00984)	(0.0162)	(0.0126)		(0.0117)	(0.0202)	(0.0151)	
Positive Savings in Prior 6 Months	0.717	0.0163	0.0244	1,304	0.730	0.0252	0.00105	971
	(0.0319)	(0.0366)	(0.0346)		(0.0376)	(0.0458)	(0.0387)	
Has Bank Account at Baseline	0.293	0.0573	0.0150	1,304	0.305	0.0589	0.0286	971
	(0.0304)	(0.0349)	(0.0347)		(0.0348)	(0.0405)	(0.0403)	
Has Post Office Account at Baseline	0.134	-0.0466*	-0.00748	1,304	0.133	-0.00908	-0.0317	971
	(0.0223)	(0.0243)	(0.0245)		(0.0256)	(0.0282)	(0.0305)	
Has BPL Card	0.840	0.0197	0.00363	1,304	0.820	0.0150	0.0203	971
	(0.0211)	(0.0251)	(0.0266)		(0.0266)	(0.0327)	(0.0302)	
Predicted Income - Predicted Expenses	3,175	-204.6	-961.4	1,304	1,828	205.8	-269.4*	971
	(349.8)	(607.4)	(947.5)		(148.4)	(191.0)	(139.5)	
Saving Goal	1,838	-239.1**	132.8	1,304	1,578	27.25	-70.65	953
	(117.1)	(117.4)	(167.0)		(88.45)	(116.1)	(95.57)	
Saving Goal (1% outliers trimmed)	1,650	-106.5	-55.77	1,283	1,404	33.34	31.87	928
	(76.04)	(78.99)	(102.0)		(64.33)	(68.11)	(71.63)	

Table 2. Goal Attainment Treatment Effects

Panel A	(1) Reached Goal	(2) Reached Goal	(3) Reached Half Goal	(4) Savings Balance	(5) Savings Balance	(6) Excess Savings	(7) Excess Savings
Monitor	0.0613***	0.0610***	0.0485*	347.3	352.4	381.0	397.1
	(0.0195)	(0.0196)	(0.0245)	(232.7)	(236.5)	(256.1)	(244.5)
Saving Goal		-9.03e-06***	-1.80e-05***	0.0923	0.198		
		(3.29e-06)	(4.50e-06)	(0.0654)	(0.149)		
Constant	0.0783***	0.0947***	0.193***	407.9**	239.6	-1,240***	-1,065***
	(0.0149)	(0.0139)	(0.0170)	(196.7)	(273.5)	(195.9)	(187.0)
Winsorized (Saving Goal)	No	No	No	No	1%	No	1%
Observations	1,302	1,302	1,302	1,302	1,281	1,302	1,281
R-squared	0.124	0.127	0.144	0.048	0.050	0.054	0.049

Panel B	(1) Reached Goal	(2) Reached Goal	(3) Reached Half Goal	(4) Savings Balance	(5) Savings Balance	(6) Excess Savings	(7) Excess Savings
Monitor: Endogenous	0.0567**	0.0594**	0.0415	284.1	331.2	13.59	266.1
	(0.0242)	(0.0244)	(0.0347)	(447.4)	(452.2)	(450.5)	(451.0)
Montor: Random	0.0658**	0.0625**	0.0554	409.2**	373.3*	740.2***	525.4***
	(0.0304)	(0.0303)	(0.0344)	(174.5)	(188.3)	(230.0)	(194.0)
Saving Goal		-9.01e-06***	-1.79e-05***	0.0932	0.199		
		(3.23e-06)	(4.46e-06)	(0.0673)	(0.152)		
Constant	0.0782***	0.0946***	0.193***	404.0**	238.1	-1,252***	-1,070***
	(0.0151)	(0.0141)	(0.0170)	(190.7)	(274.7)	(187.3)	(179.4)
Winsorized (Saving Goal)	No	No	No	No	1%	No	1%
Observations	1,302	1,302	1,302	1,302	1,281	1,302	1,281
R-squared	0.124	0.127	0.145	0.048	0.050	0.055	0.049

Regressions in both panels are based on savings balances accumulated in the treatment bank or post office account. Dropouts are assumed to not have used these accounts (as many never opened an account). All regressions include village fixed effects. Standard errors are clustered at the village level. *** p<0.01, ** p<0.05, * p<0.1

Table 3. Ending Total Savings Treatment Effects

Panel A	(1) Log End Savings	(2) Log End Savings	(3) Log End Savings	(4) Log End Savings
Monitor	0.119	0.132	0.125	0.143
	(0.121)	(0.121)	(0.101)	(0.104)
Log Saving Goal	0.473***	0.413***	0.276***	0.232**
	(0.101)	(0.102)	(0.0873)	(0.0910)
Constant	4.516***	4.914***	2.174***	2.527***
	(0.728)	(0.734)	(0.620)	(0.638)
Winsorized (Saving Goal)	1%	1%	1%	1%
Winsorized (End Balance)	No	1%	No	1%
Baseline Savings Controls	No	No	Yes	Yes
Observations	1,045	1,038	1,045	1,038
R-squared	0.138	0.136	0.368	0.365

Panel B	(1) Log End Savings	(2) Log End Savings	(3) Log End Savings	(4) Log End Savings
Monitor: Endogenous	-0.122	-0.115	-0.0457	-0.0289
	(0.177)	(0.176)	(0.151)	(0.159)
Montor: Random	0.370**	0.389**	0.304**	0.323**
	(0.151)	(0.151)	(0.131)	(0.130)
Log Saving Goal	0.478***	0.418***	0.280***	0.236**
	(0.101)	(0.102)	(0.0873)	(0.0910)
Constant	4.468***	4.860***	2.143***	2.494***
	(0.724)	(0.728)	(0.618)	(0.635)
Winsorized (Saving Goal)	1%	1%	1%	1%
Winsorized (End Balance)	No	1%	No	1%
Baseline Savings Controls	No	No	Yes	Yes
Observations	1,045	1,038	1,045	1,038
R-squared	0.142	0.140	0.370	0.367

Regressions in both panels are based on total endline savings balances accumulated in any savings vehicle. Data includes responses for ~70% of treatment dropouts. All regressions include village fixed effects. Standard errors are clustered at the village level. *** p<0.01, ** p<0.05, * p<0.1

Table 4. Goal Attainment and Total Savings by Network Position of Random Monitor

	(1) Reached Goal	(2) Reached Half Goal	(3) Excess Savings	(4) Total End Savings	(5) Log End Savings
A: Centrality Only					
Eigenvector Centrality of Monitor	0.0327*	0.0193	29.81	1,689**	0.218***
	(0.0163)	(0.0154)	(57.00)	(671.2)	(0.0741)
B: Proximity Only					
Social Proximity of Monitor and Saver	0.179**	0.0856	436.1**	8,333**	1.283***
	(0.0798)	(0.117)	(193.6)	(3,050)	(0.387)
C: Proximity and Relatives					
Social Proximity of Monitor and Saver	0.173**	0.0824	382.7	8,496*	1.318**
	(0.0706)	(0.114)	(239.0)	(4,376)	(0.480)
Monitor and Saver are Relatives	0.0154	0.00771	128.2	-368.1	-0.0778
	(0.103)	(0.112)	(323.6)	(5,268)	(0.488)
D: Centrality and Proximity					
Eigenvector Centrality of Monitor	0.0271	0.0165	13.47	1,407**	0.171**
	(0.0169)	(0.0155)	(60.20)	(649.3)	(0.0735)
Social Proximity of Monitor and Saver	0.151*	0.0751	437.2*	6,535**	1.026**
	(0.0858)	(0.122)	(234.6)	(3,016)	(0.378)
Winsorized (Saving Goal)	1%	1%	1%	1%	1%
Controls	Yes	Yes	Yes	Yes	Yes
Observations	459	459	455	395	395

A, B, C, and D represent 4 different sets of regression specifications. Controls for monitor and saver demographics (including wealth proxies) along with savings goals are included in each regression. In column 3, Excess Savings is measured as the total savings in the target account net of the saving goal. Dropouts are assumed to not have used the target accounts (as many never opened an account). All regressions include village fixed effects. Standard errors are clustered at the village level. *** p<0.01, ** p<0.05, * p<0.1

Table 5. Monitor Characteristics in Endogenous vs. Random Choice

	(1) Monitor Centrality	(2) Monitor Centrality	(3) SM Inverse Social Dist.	(4) SM Relatives	(5) SM Same Caste	(6) SM Same Caste
Monitor: Endogenous	0.221**	-0.0540	0.149***	0.108***	0.229***	0.204***
	(0.0841)	(0.151)	(0.0181)	(0.0187)	(0.0488)	(0.0588)
Monitor: Endogenous * Saver Centrality		0.137*				
		(0.0767)				
Eigenvector Centrality of Saver (# of Std. Dev.)		0.0654				
		(0.0633)				
Monitor: Endogenous * Saver High Caste						0.0696
						(0.103)
Saver High Caste Indicator						-0.177**
						(0.0798)
Constant	1.759***	1.638***	0.516***	0.0342***	0.484***	0.552***
	(0.0608)	(0.115)	(0.0100)	(0.00898)	(0.0334)	(0.0391)
Observations	1,000	1,000	998	998	1,000	1,000
R-squared	0.011	0.031	0.089	0.037	0.054	0.076

Regressions in all columns restrict the sample to only treatments with either a random or endogenous monitor. Because monitor type is assigned at the village level, we cannot include village fixed effects. Standard errors are clustered at the village level. *** p<0.01, ** p<0.05, * p<0.1

Table 6. Goal Attainment and Monitor Choice Set

Dependent Variable: Reached Goal	(1) Endogenous Treatment	(2) Random Treatment	(3) E or R Treatment	(4) Endogenous Treatment	(5) Random Treatment	(6) E or R Treatment
Choice Quantile	0.0559	-0.0539	-0.0759			
	(0.0607)	(0.0914)	(0.0768)			
Low Choice Quantile Indicator				-0.0993**	0.0197	0.0557
				(0.0384)	(0.0506)	(0.0395)
Monitor: Endogenous * Choice Quantile			0.130			
			(0.0933)			
Monitor: Endogenous * Low Choice Quantile						-0.147***
						(0.0554)
Monitor: Endogenous			-0.0721			0.0468
			(0.0652)			(0.0286)
Saving Goal	-1.84e-05**	-2.56e-05	-2.21e-05***	-1.80e-05**	-2.62e-05	-1.77e-05***
	(8.21e-06)	(1.67e-05)	(7.71e-06)	(8.12e-06)	(1.62e-05)	(6.32e-06)
Constant	0.139***	0.211***	0.218***	0.186***	0.176***	0.138***
	(0.0401)	(0.0464)	(0.0470)	(0.0139)	(0.0277)	(0.0249)
Winsorized (Saving Goal)	1%	1%	1%	1%	1%	1%
Village Fixed Effects	Yes	Yes	No	Yes	Yes	No
Observations	457	522	979	457	523	1,281
R-squared	0.119	0.172	0.007	0.126	0.171	0.010

Columns 1 and 4 show results from regressions restricting the sample to the endogenous monitor treatment. Columns 2 and 5 restrict the sample to the random monitor treatment. Columns 3 and 6 are differences in differences regressions and contain observations from both monitored treatments. Because random and endogenous monitor types are assigned at the village level, columns 3 and 6 do not include village fixed effects. Standard errors are clustered at the village level. *** p<0.01, ** p<0.05, * p<0.1

Table 7. Business Correspondent Treatment Bundle Effects

	Total Savings	Total Savings	Log Total Savings	Log Total Savings
"BC" Treatment ITT	1,789	2,129**	0.231**	0.242**
	(1,100)	(851.5)	(0.102)	(0.101)
Constant	8,200***	6,444***	7.440***	7.392***
	(833.9)	(645.2)	(0.0751)	(0.0741)
Winsorizing (Total Savings)	No	1%	No	1%
Observations	1,856	1,839	1,835	1,818
R-squared	0.044	0.061	0.068	0.074

Appendix A. For Online Publication

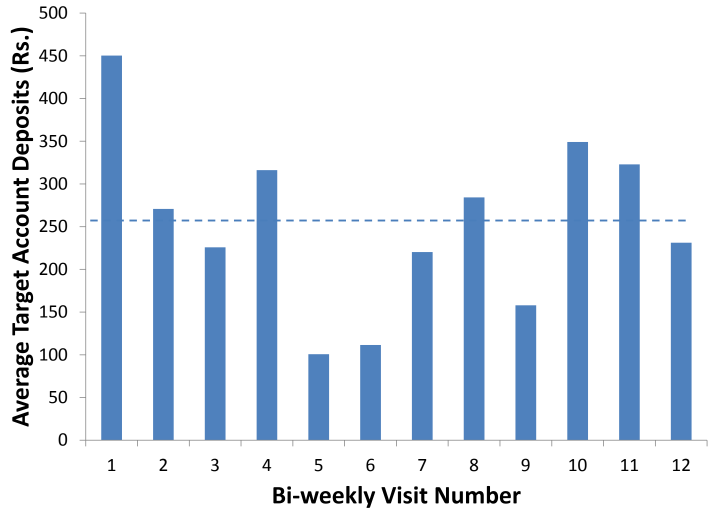

Figure 9. Aggregate Timing of Savings Deposits

Table 8. Correlates of Goal Attainment and Endline Savings in the Non-Monitored Treatment Group

Saver Characteristics	(1) Reached Goal	(2) Reached Half Goal	(3) Excess Savings	(4) Total End Savings	(5) Total End Savings	(6) Log End Savings
Eigenvector Centrality (# of Std. Dev.)	-0.000705	0.0510	-479.5	-2,330	-1,950	-0.0931
	(0.0286)	(0.0422)	(692.2)	(2,010)	(1,746)	(0.189)
Age	1.75e-05	0.000801	-32.97	-240.8	-105.4	0.00280
	(0.00261)	(0.00448)	(40.06)	(321.4)	(153.5)	(0.0218)
Female	0.00296	0.0392	-125.0	-635.6	-4,204	0.0419
	(0.0507)	(0.0822)	(372.9)	(6,188)	(4,849)	(0.567)
Married	-0.00326	0.0582	681.4	4,193	3,606	-0.111
	(0.0590)	(0.0789)	(1,116)	(4,866)	(4,364)	(0.516)
Number of Children	0.00614	-0.0128	65.85	-412.8	-1,120	0.147
	(0.0224)	(0.0307)	(177.1)	(1,705)	(1,545)	(0.180)
Number of Rooms	0.0307	0.0604*	-87.48	34.18	27.80	0.0464
	(0.0225)	(0.0313)	(209.1)	(1,122)	(955.5)	(0.153)
Electrical Connection	-0.00939	-0.0261	-805.7	-7,262*	-3,678	-0.476
	(0.0474)	(0.0542)	(772.3)	(3,763)	(2,721)	(0.321)
Post Office Account Indicator	0.138*	0.0506	389.5	5,518	4,813	0.549
	(0.0689)	(0.111)	(734.3)	(3,938)	(3,944)	(0.502)
Has Bank or PO Account at Baseline	0.0254	0.0239	886.9	7,390**	6,868**	0.812**
	(0.0587)	(0.0761)	(1,094)	(3,166)	(3,101)	(0.350)
Predicted Income During Saving Period	-1.03e-06**	-1.38e-06**	0.00698	0.0766	0.0541*	4.93e-06
	(4.05e-07)	(6.86e-07)	(0.0124)	(0.0615)	(0.0314)	(3.74e-06)
Saving Goal	-8.70e-06	-3.61e-05	-0.783***	2.464	2.455*	0.000304***
	(1.83e-05)	(2.35e-05)	(0.282)	(1.497)	(1.278)	(0.000114)
Constant	-0.184	-0.104	1,526	15,973	7,172	6.946***
	(0.193)	(0.276)	(2,666)	(17,757)	(10,796)	(1.354)
Winsorized (Saving Goal)	1%	1%	1%	1%	1%	1%
Winsorized (End Balance)	No	No	No	No	1%	No
Observations	260	260	260	226	224	225
R-squared	0.332	0.317	0.355	0.330	0.406	0.421

The regression sample includes the Non-Monitored treatment group. All regressions include village fixed effects. Standard errors are clustered at the village level. *** p<0.01, ** p<0.05, * p<0.1

Table 9. Saver Outcomes and Baseline Monitor Savings for R treatment

Dependent Variables: Saver Performance	(1) Reached Goal	(2) Reached Half Goal	(3) Excess Savings	(4) Total End Savings	(5) Log End Savings
Monitor Saved in Formal Account Previous 6 Mos	0.0213	0.0328	-57.13	-2,669	-0.0151
	(0.0491)	(0.0389)	(174.5)	(2,119)	(0.192)
Monitor Savings Goal: Fraction of Projected Income	-0.0172	-0.0108	180.5	2,953	0.0817
	(0.0578)	(0.0583)	(201.1)	(2,547)	(0.261)
Monitor Baseline Log Savings	-0.00488	-0.00576	-4.813	246.1	0.0589
	(0.00918)	(0.00844)	(43.91)	(391.7)	(0.0453)
Constant	0.209	0.182	312.8	-3,620	6.859***
	(0.224)	(0.211)	(834.2)	(10,163)	(1.069)
Winsorized (Saving Goal)	1%	1%	1%	1%	1%
Controls	Yes	Yes	Yes	Yes	Yes
Observations	421	421	417	366	366
R-squared	0.212	0.235	0.462	0.251	0.294

Regressions are run for only the group of savers randomly assigned to monitors. Monitor and Saver demographic controls are included in each specification. All regressions include village fixed effects. Standard errors are clustered at the village level. *** p<0.01, ** p<0.05, * p<0.1

Table 10. Borrowing and Lending Relationships in R treatment

	(1) Reached Goal	(2) Reached Half Goal	(3) Excess Savings	(4) Total End Savings	(5) Log End Savings
Monitor and Saver: Any Direct Relationship	0.0948	0.0431	156.5	2,785	0.467*
	(0.0615)	(0.0751)	(153.9)	(2,691)	(0.253)
Monitor and Saver: Lending Relationship	0.0461	0.0293	323.3	6,608	0.297
	(0.127)	(0.124)	(338.7)	(6,256)	(0.600)
Monitor and Saver: Borrowing Relationship	-0.0883	-0.0398	-390.1	-1,900	0.298
	(0.111)	(0.125)	(318.2)	(7,899)	(0.700)
Constant	0.0856	0.110	239.5	-2,823	7.069***
	(0.208)	(0.219)	(857.2)	(9,783)	(1.096)
Winsorized (Saving Goal)	1%	1%	1%	1%	1%
Controls	Yes	Yes	Yes	Yes	Yes
Observations	459	459	455	395	395
R-squared	0.216	0.234	0.456	0.241	0.295

Regressions are run for only the group of savers randomly assigned to monitors. Monitor and Saver demographic controls are included in each specification. All regressions include village fixed effects. Standard errors are clustered at the village level. *** p<0.01, ** p<0.05, * p<0.1

Table 11. Expenditures in Last Month of Savings Period by Treatment Status

Dependent Variable: Expenditure in Last Month of Saving Period	(1) Log Measured Expenditures	(2) Pan	(3) Festivatls	(4) Tea	(5) Meals	(6) Eggs and Dairy	(7) Other Food	(8) Transport	(9) Phone
Monitor: Endogenous	0.0476	-23.14	0.705	18.19	32.68	60.97	-96.45	11.93	-4.464
	(0.0881)	(26.62)	(94.29)	(39.92)	(33.13)	(38.42)	(88.67)	(87.44)	(18.37)
Monitor: Random	-0.0669	15.56	-215.1*	27.23*	20.79	-75.02	-158.4	-140.5*	0.00955
	(0.0708)	(26.68)	(128.4)	(15.76)	(39.54)	(62.54)	(126.3)	(70.99)	(29.81)
Constant	8.548***	145.1**	415.8**	261.0***	459.5***	661.4***	1,293***	530.0***	282.3***
	(0.138)	(55.15)	(182.3)	(73.15)	(82.18)	(104.9)	(236.4)	(165.3)	(42.07)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,000	1,134	1,133	1,135	1,121	1,128	1,134	1,135	1,126
R-squared	0.178	0.093	0.260	0.096	0.124	0.170	0.363	0.160	0.176

Dependent variables are measures of expenditures of households in last month of the 6 month savings period. All regressions include village fixed effects. Standard errors are clustered at the village level. *** p<0.01, ** p<0.05, * p<0.1