SAVINGS MONITORS

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

Abstract. We conduct a field experiment 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 "best" 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. In markets with credit constraints, household savings are important for the accumulation of assets. Further, cash savings can provide a buffer so that investments need not be liquidated during adverse shocks. In developing countries, however, rural households do not appear to save adequately. The opportunity to save in local bank branches is not the main obstacle, as the formal financial sector has extensive geographical reach in many settings such as rural India. For instance, in our study region most households live within 3 miles of a bank branch. Moreover, these branches offer no-frills savings accounts with zero or extremely low minimum balance requirements as well as no restrictions on or fees for making withdrawals.

Date: August 2013 – Preliminary and incomplete, please do not cite without an author's consent.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 are grateful to CFSP, NSF, the Chazen Institute at Columbia Business School, and CIBER at Columbia University for generous support. We also thank Robert Townsend, Esther Duflo, and Abhijit Banerjee for their tremendous guidance. 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.

One financial product offered by some banks and favored by policymakers is the use of *Business Correspondents* – that is, an individual hired by banks to visit villagers at their homes and collect savings deposits. Our study builds on this particular model.

Behavioral economics research has highlighted three obstacles to good savings behavior: (a) inattention, (b) time inconsistency, and (c) a particular special case, temptation. Experimental evidence has shown that product innovations such as reminders and commitment devices can help to mitigate these channels (Karlan et al. (2010); Ashraf et al. (2006); Banerjee et al. (2009)).

We explore two interventions to help to individuals potentially overcome these behavioral biases and increase their savings balances. First, we provide a treatment bundle that is reminiscent of the Business Correspondent model. In this treatment, we assist households in opening an account, help individuals construct an attainable 6-month savings goal, and visit savers fortnightly to check in on their progress. This baseline treatment offers built-in reminders, which by themselves should help combat inattention. It does not, however, tackle time inconsistency and temptation.

As such, even with reminders and goal setting, individuals still may not save according to their targets. In our second set of treatments we ask to what extent the social network can help to increase goal attainment beyond the reminders treatment. An individual is assigned a monitor, and the monitor is notified on a regular basis of the saver's progress in attaining her savings goal.[1] The core idea is simple: individuals maintain reputation within their communities. Their ability to present themselves to their fellow community members as responsible individuals who can attain their savings goals may influence their reputation and therefore other walks of their life. This peer intervention is not unlike making use of social capital to help generate incentives to maintain repayment behavior in a microfinance setting.

We offer two different versions of the monitoring treatment. Savers selected to receive a monitor are further randomized to either receive a randomly-selected monitor or to choose their own monitor from a pre-specified pool of villagers.[2] By comparing across treatments, we can ask whether individuals do better at attaining their goals

[1]In developed (and some developing) countries, technology has been used to help to close this gap (e.g., automatic withdrawals and commitment mechanisms such as stickk.com). Where technology-based solutions are not readily available or banks may be unwilling to be financially innovative and introduce new commitment-savings devices, it may be cost effective to call upon existing informal institutions such as the network.

[2]The randomization is at the village level.

when they are free to choose their own monitors. The answer is unclear ex ante. On the one hand it may be the case that individuals are better able to choose monitors who enforce better behavior, while on the other hand individuals may purposefully choose monitors who may "go easy" on them, or who may stand to learn little about their reputation from their savings behavior. In fact, individuals may feel unable to choose optimal monitors – important community members who ex ante do not have a strong knowledge about how responsible the saver is – as it may be construed as a burden. Meanwhile, random assignment alleviates this burden as the banking institution is making this request.

We are interested in whether certain community members are likely to be better monitors than others. Given that we anticipate a reputational mechanism, we hypothesize that savers who are assigned monitors that are central in the network are likely to see the largest gains. To study this we make use of use extensive social network and covariate data (from Banerjee et al. (2011)) to understand which types of saver-monitor pairs are most successful. Further, we analyze how the success of the monitor changes with the social distance between savers and monitors. The direction of the social distance effect is ambiguous under the theory. Distant pairs may have incentives to use the monitoring relationship to build new, possibly valuable, network connections. However, close pairs may have lower monitoring costs. We also investigate whether individuals choose the "optimal" monitor from their set of choices.

We find that the business correspondent treatment bundle (henceforth BC) has at most a modest effect on savings accumulation (10%). Individuals in the BC treatment have a low rate of meeting their established long-run savings goals (7%). However, adding a monitor has a large effect on goal attainment (6pp increase) and total savings balances (34% increase). We also provide evidence that individuals save more when they are assigned to a random monitor versus a self-chosen monitor. Finally, we find that the network position of the monitor matters considerably for savings goal attainment and total balances. Namely, having a monitor who is important or central in the social network corresponds to larger treatmtent effects. Going from the 5th to 95th percentile of the monitor centrality distribution in the exogenous treatment is associated with an 8pp increase in goal attainment and a 43% increase in total savings balances. 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. The centrality and proximity results hold conditional on each other, and the analysis includes controls for caste, wealth, occupation and village fixed effects.

Our findings suggest that individuals in the endogenous monitor treatments are picking the "right" monitors; those who are chosen to select their monitors earlier in the matching phase tend to perform better. However, the allocation of monitors to savers through the endogenous selection process generates a socially suboptimal outcome. Essentially, while ex ante worse savers may require differentially better monitors, the endogenous treatment tends to allocate these more central monitors to central savers. We interpret this as peripheral individuals not having the social capital to request, in the endogenous treatment, high centrality individuals to be their monitors. Meanwhile, an institution with random assignment of monitors matches high central monitors with peripheral savers who otherwise may not have had the capital to make such a request.

Ultimately, this peer-based product is an easy-to-implement and cost-effective business model that can be used in conjunction with business correspondents. Through our randomized research design we are able to examine whether this peer-based commitment device is a viable policy solution, both when individuals choose their monitors and when their monitor is randomly assigned. Our product utilizes established peer networks, present in every rural community, to help individuals improve savings capacity and potentially accumulate business assets.

Previous Literature. The benefits of savings are well-documented in the economics literature. Dupas and Robinson (2009), 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.[3] Rational hypotheses include distrust of banks, limited access to safe storage of funds, high transactions costs and high discount rates.[4]

[3]See Laibson (1997)and Banerjee and Mullainathan (2010) among many others.

[4]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.

Much 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.

Recent research has also underscored the importance of social networks in many spheres of life. Social networks have been shown to diffuse information (Banerjee et al. (2011) for microfinance, Conley and Udry (2010) for agriculture, Kremer and Miguel (2007) for health, and Beaman and Magruder (2009) for job referrals), help sustain cooperation (Chandrasekhar et al. (2011) and Breza et al. (2011)), and aid in coordination (Choi et al. (2011)).[5] Peer dynamics have been incorporated into many formal and informal financial products such as ROSCAs, Self Help Groups (SHGs), and Grameen Bank-style microlenders. In a related lab experiment conducted in the field, Breza et al. (2011) provide evidence that reputation may matter significantly and that its effect may vary significantly with the centrality of parties within a social network. Breza et al. (2011) 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. The present paper builds on this sort of micro-finding, leveraging similar forces for a real world savings product.

In our setting, we ask if social monitoring can help individuals from "defaulting" on their own goals. In other words, can peers substitute for a formal commitment device? The closest paper to ours is Kast et al. (2012), who analyze the effects of peer monitoring using both self-help group peers and self-selected "savings buddies". They find that monitoring (where peers choose their own monitors_ substantially increases savings balances as well as the frequency of deposits. However, they find that the peer pressure mechanism itself may not be essential for increasing savings, and that reminders may equally successful.

We aim to isolate and focus on the role of peer monitoring and identify those peers who are most effective in helping their savers to save. In each or our treatments,

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

individuals receive a bundle of many sevices including account opening, goal setting, and bi-weekly visits from one of our enumerators (implicitly serving as a reminder to save). For some savers we add either an endogenous or randomly-selcted monitor from the saver's own village. Further, we use detailed social network data to analyze the heterogeneous benefits of different types of monitors. This allows us to identify whether there are peer monitoring based effects that influence goal attainment beyond a reminders effect. Our findings also shed light on how local community structures might be able to complement formal institutions most effectively.

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 our hypotheses, while section 4 displays preliminary results. Section 5 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 8 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. However, our baseline shows that the use of these branches is quite low. Only a quarter of households had an account at baseline. Figure 1 shows the baseline intensity of use of available savings vehicles separately for male and female savers. On average, potential savers keep a large fraction of their 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.

We are keen 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. Thus, 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. (2011) 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. (2011) and Chandrasekhar et al. (2011). 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 $G_r = (V_r, E_r)$ consisting of a set V_r of vertices (households) and E_r of edges (denoting whether two vertices are linked or not in the village). To every network we associate an adjaceny matrix $\mathbf{A}_r$ which represents whether or not two vertices are linked. That is, $\mathbf{A}_{r,ij} = 1$ if and only if $ij \in E$.

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.

2.2. **Experimental Design.** Figure 2 pictorially represents our experimental design. Study participants are randomly selected from an existing village census database and then randomly assigned to be part of our saver group or monitor group. A random subset of those individuals who are not selected into our study comprise the pure control group (T0).

All potential treatment savers and monitors who are interested in participating are administered a baseline survey, which includes questions on historic savings behaviour, 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. For potential saver households, a six-month savings plan is established. 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. 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. 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.

The three treatments assigned to savers are:

(1) Individual or business correspondent bundle treatment (Randomized at the individual level)
(2) Peer treatment with random matching (Randomized at the village level)
(3) Peer treatment with endogenous matching (Randomized at the village level)

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 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.[6]

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.

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. We collect complete savings information 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.

[6]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.

3. Framework

We can compare the treatment means in groups 1, 2, and 3 to each other and as well as to the pure control group. Comparing treatment 1 to the pure control gives the causal effect of our bundle of opening an account, eliciting a savings goal, and generating reminders. Comparing pooled treatments 2 and 3 to treatment 1 gives the incremental benefit of having a monitor. Comparing treatments 3 versus 2 gives the value of allowing agents to choose their monitors relative to receiving an average monitor from the available pool.

We can also measure the effect of the business correspondent bundle but comparing the end savings balances in the pure control group to a weighted sum of the outcomes of the individual treatment group and the never-takers. Because the business correspondent treatment includes fortnightly reminders, we hypothesize a positive treatment effect on savings.

We hypothesize that individuals in a monitored treatment will be more likely to reach their savings goal and more likely to accumulate business assets. However it is unclear whether the endogenously-chosen monitors should outperform or underperform the random monitors. On one hand, individuals may be well-positioned to know who might be the ideal savings monitor. Therefore, allowing for choice should increase effectiveness. However, individuals may feel social pressures to choose friends, who may not be in a position to "enforce" savings behavior. It is also possible that individuals do not know who would be optimal. On the other hand, random individuals may be so unimportant in an individual's financial life that they have to ability to help individuals reach their financial goals. Saver-monitor pairs might live far apart, and may never have the opportunity to interact. If these types of matches are common, then the random monitor might not be so effective. Finally, it may be the case that individuals understand who the best monitors are in a community, but that they may be mis-allocated.

We also hypothesize that the identity of the monitor should affect the success of the monitor-saver relationship. Our random monitor allocation allows us to test the heterogeneous causal effects of monitoring by monitor-saver characteristics. We utilize detailed covariate and social network data (from Banerjee et al. (2011)) to study the effects of receiving a monitor with different network characteristics.

We focus on two notions of the social network: eigenvector centrality and social proximity. We use the eigenvector centrality of an individual's household in the social network as a measure of her importance. Formally,

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

where λ_1 is the maximal eigenvalue of the adjacency matrix and ζ 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. In our context this suggests that a saver has more reputation at stake when not meeting her goal in front of a more central monitor. Furthermore, one might imagine that individuals in a community meet via a stochastic matching process directed by the network. In such a context, an individual is differentially more likely to meet those who are more eigenvector-central. Again this suggests that gaining a positive reputation in the eyes of a more central monitor will have greater value. For a more general discussion about eigenvector centrality in network economic models, see Jackson (2008).[7]

The social distance between two individuals is the shortest path (if one exists and is finite) between the two through the graph.[8] Formally,

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

It may be the case that more socially proximate individuals have better prior knowledge about a saver's responsibility and therefore has less to learn based on the saver's performance. On the other hand, a more socially proximate monitor is more likely to have repeated interactions in other walks of life with the saver and therefore the reputation formed is likely to help leverage numerous other interactions. As such, a priori there is no obvious sign to the relationship between savings goal attainment and the distance between saver and monitor in the network.

Ultimately, we hypothesize the central individuals will be especially good monitors, even conditional on distance and other demographic covariates. We are interested in signing the effect of social proximity.

Finally, we can measure how well the endogenously chosen monitors perform relative to the "optimal" monitors identified from the random matching treatment.

---

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

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

4. Results

4.1. **Treatment Level Effects.** We present data from all 60 villages for many outcomes. However, for some outcome variables including total savings balances across all vehicles, we are missing outcomes from 5 villages. We plan to incorporate this information into the analysis soon. Figure 3 presents the histogram of savings goals, censoring the top 5%.[9] There are a few large outliers (maximum goal Rs. 26,000), so the mean of Rs. 1,788, is larger than the median, which is Rs. 1,200. 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.

First, we explore the effects of the business correspondent treatment bundle on savings attainment in Table 2. We should first note that these regressions are underpowered. The data that we will obtain shortly will increase the sample size in all of these regression specifications. The table suggests that if anything, the business correspondent treatment has a positive, but modest impact on savings outcomes. The effect of the treatment on overal savings balances is positive and is significant at the 10% level. The specification in logs suggests that the intervention increases savings by 10%, although with the limited data set, the coefficient is not statistically significant at standard levels.

Table 3 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

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

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 4 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 3, 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 5 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. The sample size for the regressions using total savings is substantially smaller than for the target account regressions. This is due to the fact that we are still missing data from 4 villages for those outcomes. We envision an increase in power once all of the data is included.

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 0.86 log points. These magnitudes are quite large. There is weak evidence (insignificant coefficients) that being paired with a relative may increase goal attainment, but does not increase total savings by as much as an unrelated direct link.

Moving from the 5th to 95th percentile of the eigenvector centrality of the monitor (0.12 units) increases savings goal attainment by 8 pp. This effect is of a similar magnitude to the effect of assigning an average monitor. There are also significant effects on total savings outcomes across all vehicles. Moving from a low to high centrality monitor increases total savings by 0.5 log points. 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 potential individual is not randomly assigned (although the matches between saver and monitor are), we believe that the network is actually impacting the success of the monitoring relationships. The regressions in Table 5 all contain controls for saver and monitor demographic characteristics, proxies for saver and monitor wealth, and controls for the saver's savings goals. It is suggestive that the network effects matter even after controlling for these indicators.

4.3. **Endogenous Monitor Selection.** The analysis of the randomly matched savers and monitors suggests that individuals in the endogenous group should be better off picking monitors of close social proximity and high network centrality. We investigate whether individuals actually make these choices in Figures 4 and 5.

Figure 4 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. 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. Recall that not all available monitors are assigned a saver in both groups. Thus, the patterns in social proximity could be rationalized by either a different overall set of monitors being chosen, rematching of the same individuals to different savers, or both. 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. It almost appears that if a relative is in the monitor pool, savers in the endogenous group are almost

Figure 5 displays a histogram of the quantiles of the selected monitors in the endogenous versus random treatments. In the population, the quantile of eigenvector centrality should be distributed uniformly by construction. However, note that choosing to be a monitor selects for individuals who are on average more central than the overall population, even in the random matching treatment. Further, the mass in the endogenous treatment shifts even further to the right, with individuals actively choosing more central monitors than random. A regression of centrality quantile on treatment type indicates that endogenously chosen monitors are on average 5 percentiles higher in the centrality distribution than randomly chosen monitors. Because centrality is only monitor-specific and doesn't depend on the identity of the saver, these histograms suggest that the actual pool of chosen monitors is changing between the two treatments.

Furthermore, we can ask how the savers perceive the effort and value of their monitors both in the random and endogenous groups. Figure 6 shows the beliefs of the savers regarding how hard their monitors worked during the six month savings period, and Figure 7 shows the preferences of savers if they were able to hypothetically choose a different monitor. The dark bars indicate the random treatment, while the lighter colored bars indicate the endogenous treatment. On average, savers in the endogenous treatment believe that their monitors worked harder. The difference is half of a category, and is statistically significant at the 1% level. There is also a large spike at "Very Low" effort for the random group. Conversations with savers in the random group suggest that some socially distant monitors never communicated with their savers. It is less likely that endogenously-chosen monitors chose to never engage with their savers.

Ex post, we also find that savers in the endogenous treatment group are more likely (and statistically significantly so) to report that they would want to choose the same monitor if they had the opportunity to choose in the future. In the random treatment group, instead of choosing the same monitor, many savers instead choose family, friends or neighbors.

In sum these results suggest that individuals do in fact choose those monitors who should help them succeed at saving more. Individuals also seem happy with their choices ex post. However, one puzzle remains. Why do we observe individuals in the endogenous group saving less than in the random group in Table 4?

To investigate this question, we analyze whether perceived high quality monitors may be scarce in the endogenous monitor group (relative to the random monitor group). 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. Table 6 displays the effects of choosing late on savings goal attainment. In each column, we regress goal attainment on the badge order quantile. The badge order quantile ranges from 0 to 1, with 0 denoting the first individual to receive a monitor in each session. 1 denotes the last individual to receive a monitor in each session. The table shows that there is no statistically significant effect of badge order on goal attainment in the random monitor group. However, in the endogenous treatment group, chosing late decreases the likelihood of goal attainment. The difference in these effects between treatment and control is also significant.

This suggests that at the end of the village meeting, in endogenous treatments, good monitors are scarce. However, the random monitor pool is no different from the endogenous monitor pool. This suggests that the allocation of monitors is important and, moreover, individuals recognize who the "star" monitors are likely to be. While the random treatment distributes the desirable monitors more equitably throughout the saver distribution, endogenous treatment – due to assortativity in the network structure combined with a reluctance to ask socially distant individuals to be one's monitor – tends to have high centrality monitors paired with high centrality savers, who may simply have less to gain from the relationship as compared to a low centrality saver.

5. Conclusion

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 peer monitoring has sizeable effects in increasing savings goal attainment. Further, randomly chosen monitors are more effective than individually chosen monitors.[10] 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. However, connecting peripheral individuals with important monitors may be even more beneficial as the peripheral individuals may stand to gain the most. As such, self-sorting may lead to inefficient, if not unequal, outcomes owed to the increasing social cost faced by peripheral individuals to approach central 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.

[10]We caution that individuals may understand that, for instance, due to unforseen circumstances they may be unable to make their goals and such may be unwilling to risk losing reputation in the eyes of a more central monitor. As such, in our endline we ask whether individuals would have wanted a different monitor and, if so, whom.

References

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 (2011): "The Diffusion of Microfinance," MIT working paper.

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," .

Breza, E., A. G. Chandrasekhar, and H. Larreguy (2011): "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*.

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

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

Conley, T. and C. Udry (2010): "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 (2009): "Savings Constraints and Microenterprise Development: Evidence from a Field Experiment in Kenya," NBER Working Paper 14693.

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.

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*.

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.

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.

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

Figures

Figure 1. Intensity of Use of Available Savings Vehicles

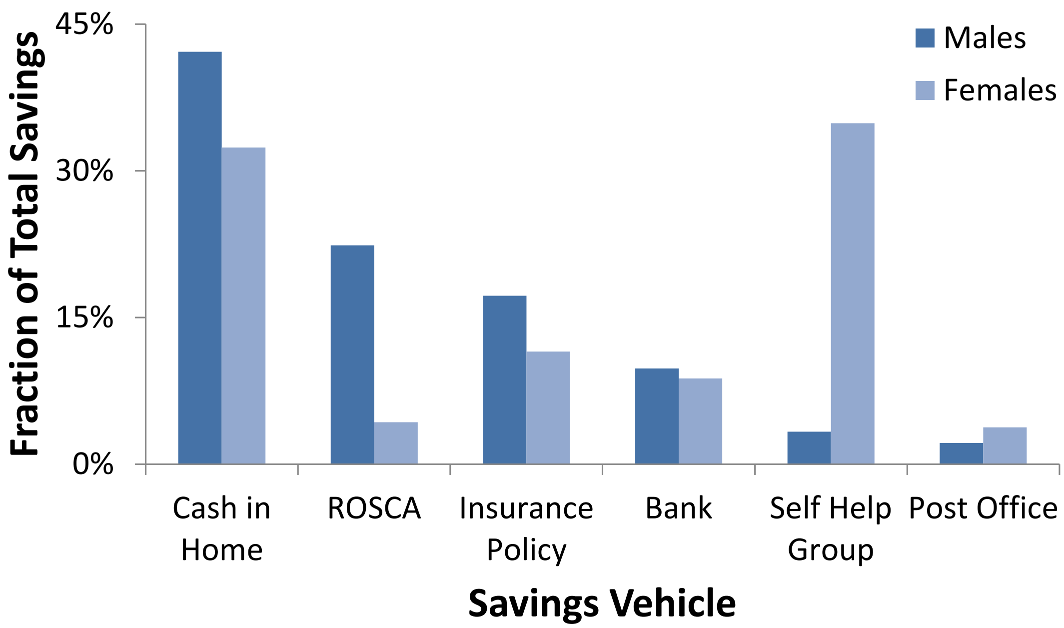

Figure 2. Experimental Design and Randomization

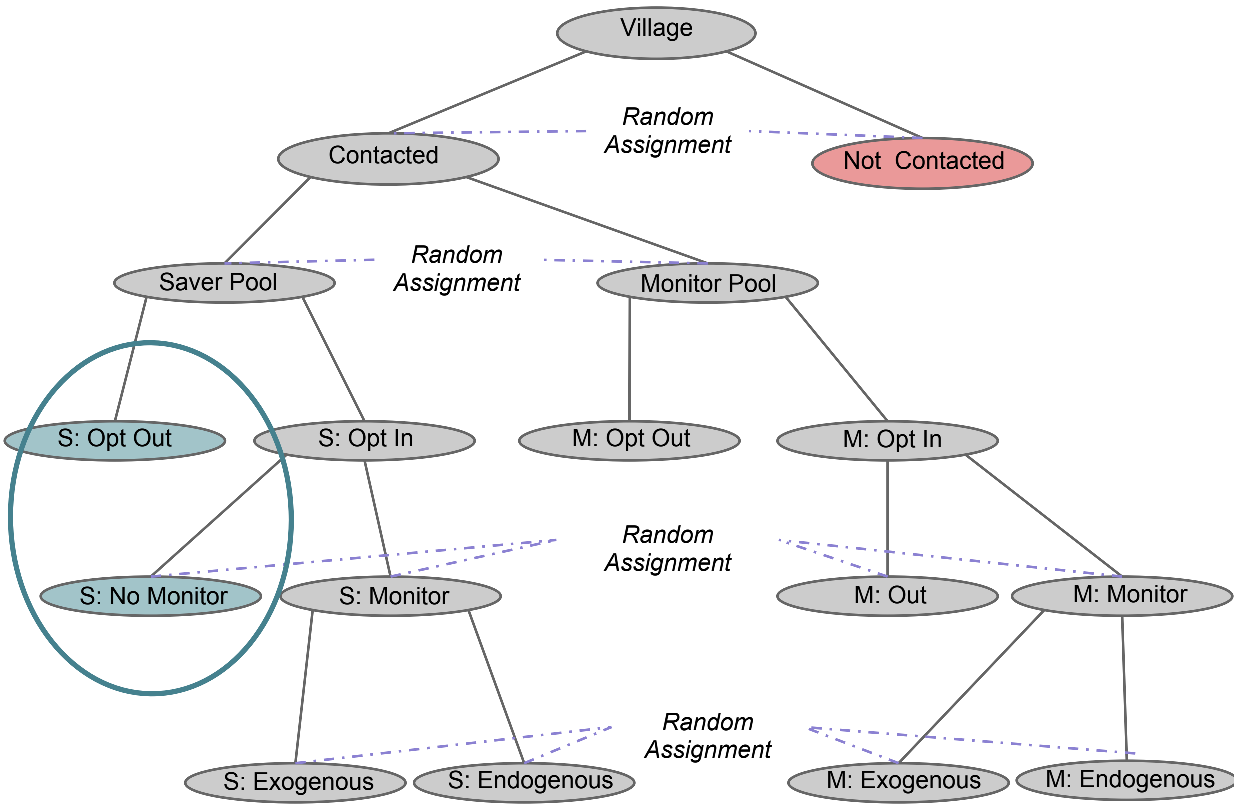

FIGURE 3. Histogram of Baseline Savings Goals

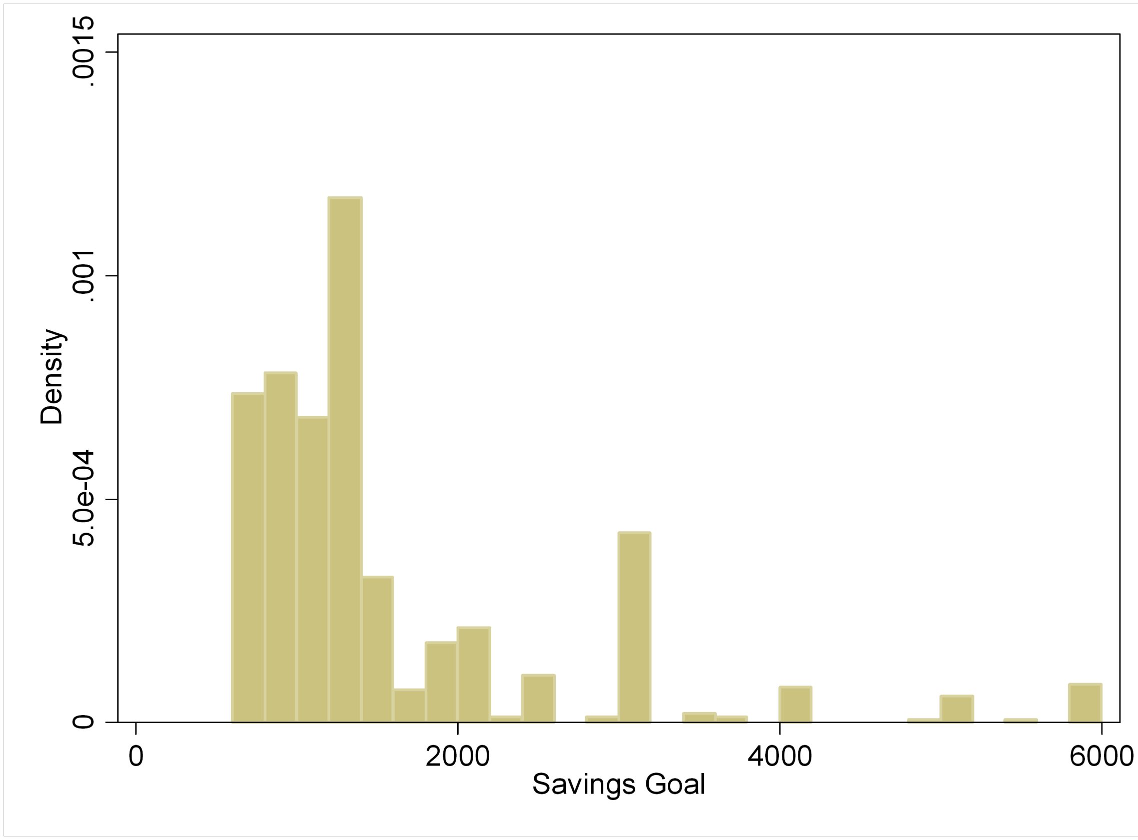

The figure shows the distribution of the baseline savings goals. We clip the top 5% tail of the distribution to make the figure more readable.

Panel A: Endogenous Treatment

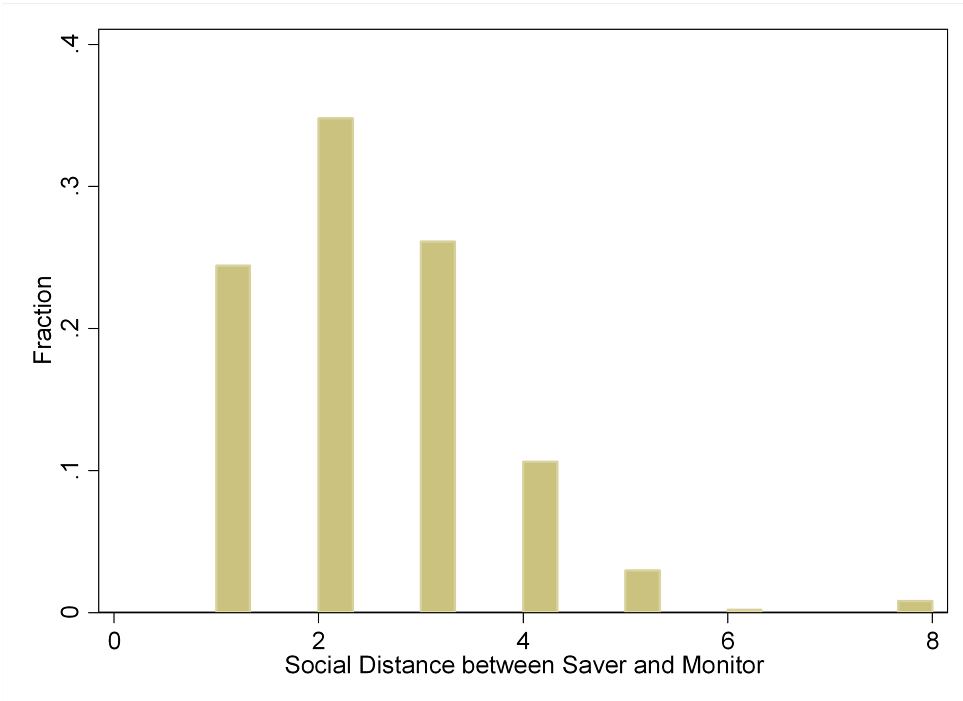

Panel B: Random Treatment

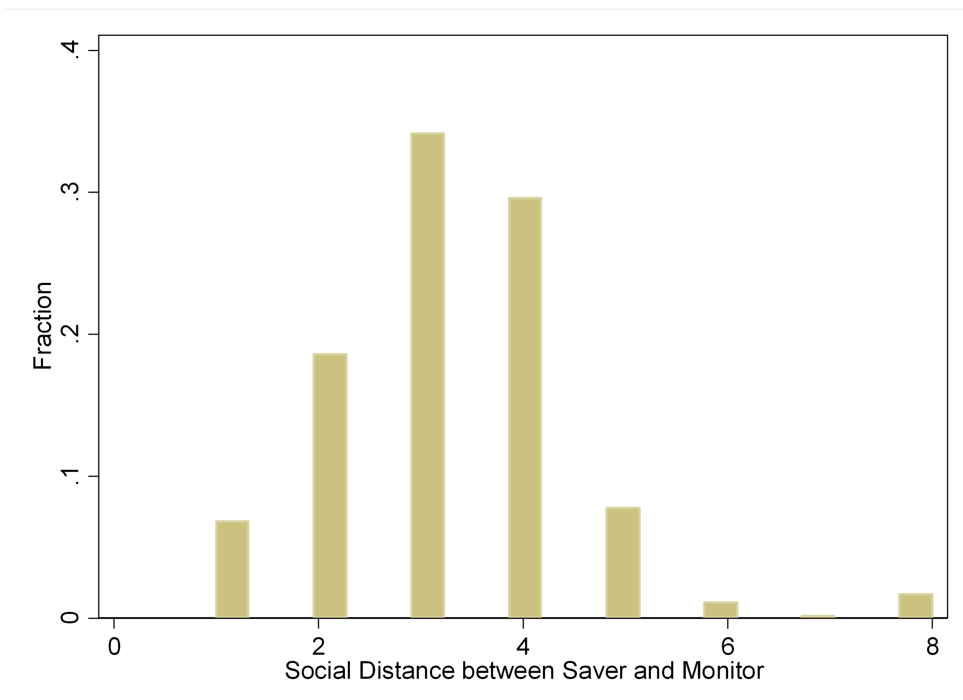

FIGURE 4. Social Distance Between Saver and Monitor by Treatment

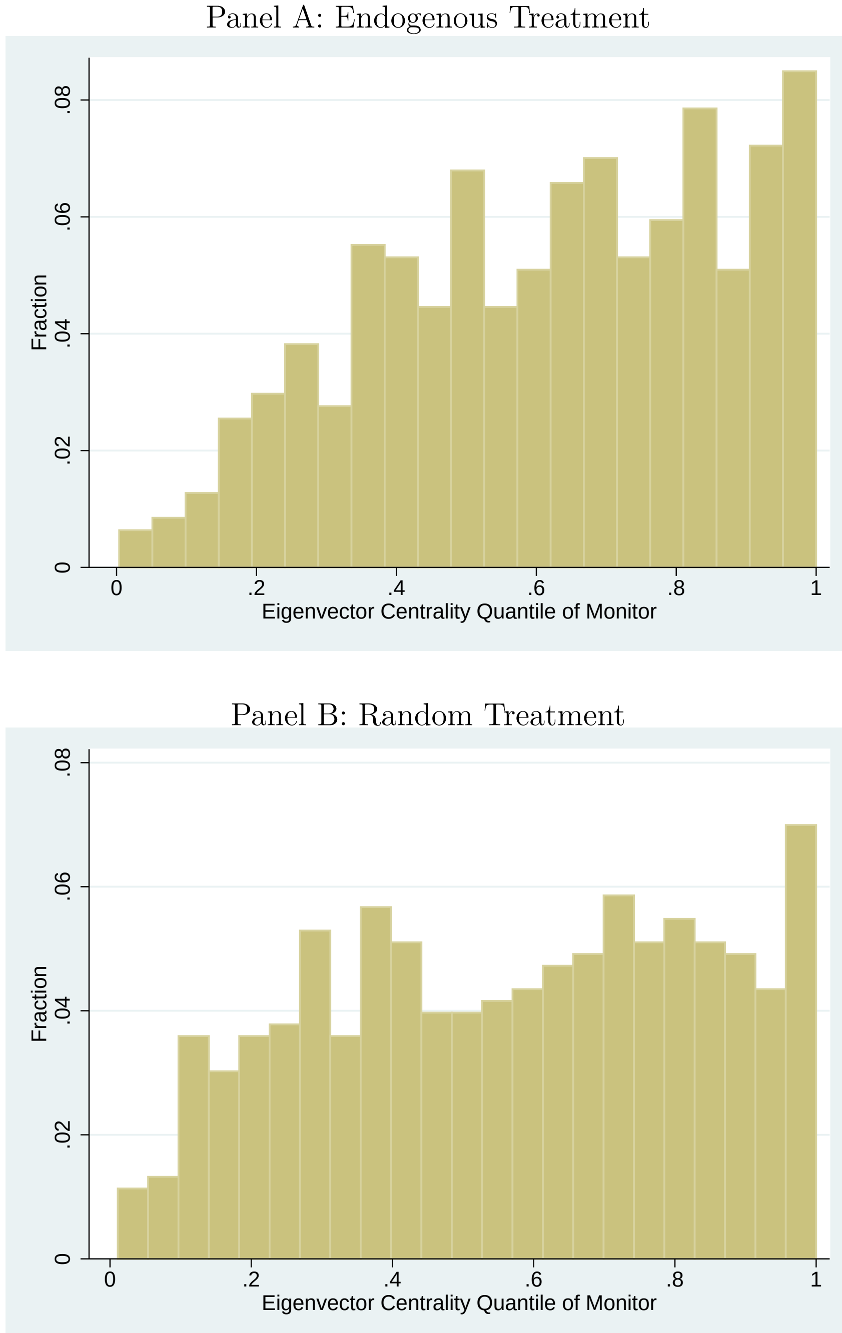

Figure 5. Centrality Quantile of Monitor by Treatment

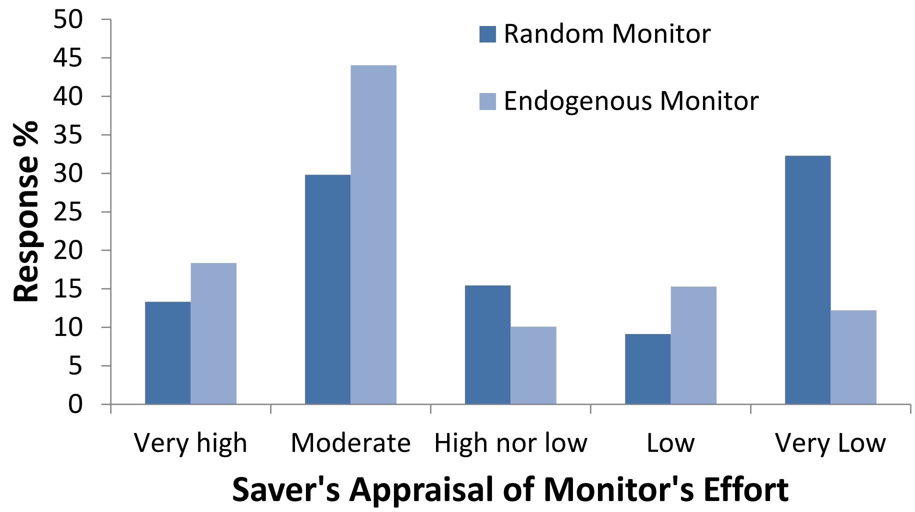

Figure 6. Monitor Effort Appraisal by Savers

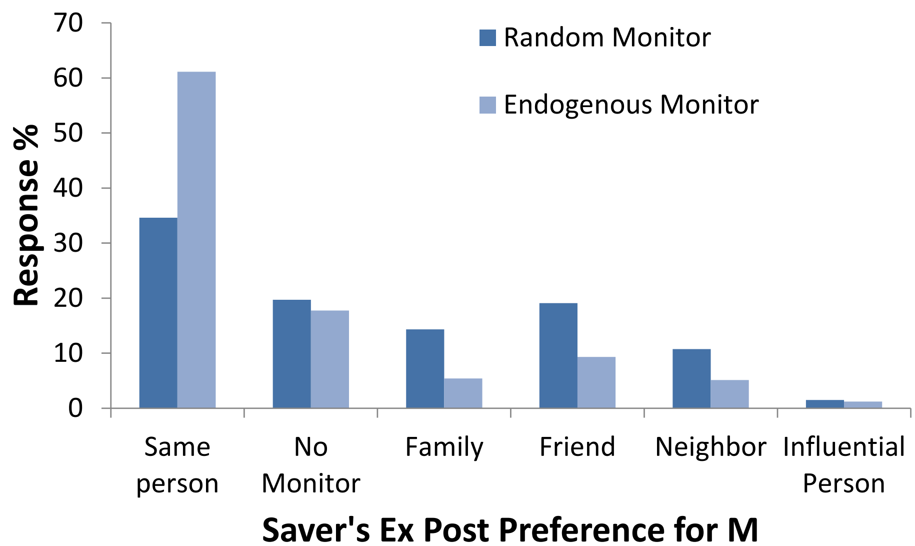

Figure 7. Saver's Preference for Monitors in Future

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)	
Saving Goal (5% outliers trimmed)	1,443***	9.437	-10.13	1,246	3,032***	-1,197	3.844	971
	(54.09)	(62.90)	(60.46)		(433.9)	(1,243)	(568.7)	

Table 2. Business Correspondent Treatment Bundle Effects

	(1) Total Savings	(2) Total Savings	(3) Log Total Savings	(4) Log Total Savings
Business Correspondent Treatment	1,336	1,377*	0.137	0.128
	(1,070)	(745.0)	(0.0987)	(0.0995)
Winsorizing (Total Savings)	No	1%	No	1%
Observations	1,809	1,793	1,772	1,788
R-squared	0.041	0.040	0.064	0.063

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. Goal Attainment Treatment Effects

Panel A	(1) Reached Goal	(2) Reached Half Goal	(3) Savings Balance	(4) Savings Balance	(5) Excess Savings	(6) Excess Savings
Monitor	0.0572***	0.0627**	315.4	321.1	343.2	364.8
	(0.0206)	(0.0258)	(249.8)	(254.2)	(274.5)	(262.0)
Saving Goal	-8.94e-06***	-1.78e-05***	0.0901	0.200		
	(3.33e-06)	(4.53e-06)	(0.0649)	(0.151)		
Constant	0.0993***	0.185***	446.0**	269.0	-1,238***	-1,049***
	(0.0146)	(0.0176)	(205.4)	(278.7)	(209.9)	(200.3)
Winsorized (Saving Goal)	No	No	No	1%	No	1%
Observations	1,184	1,184	1,184	1,163	1,184	1,163
R-squared	0.128	0.145	0.050	0.052	0.055	0.050

Panel B	(1) Reached Goal	(2) Reached Half Goal	(3) Savings Balance	(4) Savings Balance	(5) Excess Savings	(6) Excess Savings
Monitor: Endogenous	0.0489**	0.0562	180.9	233.9	-120.0	167.4
	(0.0235)	(0.0374)	(483.3)	(488.4)	(478.9)	(482.8)
Montor: Random	0.0652*	0.0689*	445.2**	405.5**	789.8***	555.7**
	(0.0330)	(0.0356)	(186.1)	(201.4)	(251.1)	(212.4)
Saving Goal	-8.82e-06***	-1.77e-05***	0.0921	0.202		
	(3.26e-06)	(4.49e-06)	(0.0671)	(0.154)		
Constant	0.0987***	0.184***	437.0**	262.0	-1,257***	-1,059***
	(0.0150)	(0.0175)	(197.2)	(280.1)	(198.6)	(190.5)
Winsorized (Saving Goal)	No	No	No	1%	No	1%
Observations	1,184	1,184	1,184	1,163	1,184	1,163
R-squared	0.128	0.145	0.050	0.052	0.057	0.050

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 4. Ending Total Savings Treatment Effects

Panel A	(1) Log End Savings	(2) Log End Savings	(3) Log End Savings	(4) Log End Savings
Monitor	0.345**	0.344**	0.346***	0.333***
	(0.135)	(0.135)	(0.116)	(0.116)
Log Saving Goal	0.442***	0.378***	0.183*	0.148
	(0.103)	(0.105)	(0.0940)	(0.0983)
Constant	4.610***	5.042***	2.269***	2.592***
	(0.759)	(0.776)	(0.606)	(0.635)
Winsorized (Saving Goal)	1%	1%	1%	1%
Winsorized (End Balance)	No	1%	No	1%
Baseline Savings Controls	No	No	Yes	Yes
Observations	845	840	806	802
R-squared	0.135	0.132	0.402	0.392

Panel B	(1) Log End Savings	(2) Log End Savings	(3) Log End Savings	(4) Log End Savings
Monitor: Endogenous	0.104	0.114	0.289	0.269
	(0.194)	(0.195)	(0.178)	(0.179)
Montor: Random	0.611***	0.596***	0.409***	0.403***
	(0.170)	(0.170)	(0.145)	(0.145)
Log Saving Goal	0.445***	0.382***	0.185*	0.150
	(0.102)	(0.103)	(0.0933)	(0.0975)
Constant	4.565***	5.000***	2.258***	2.580***
	(0.752)	(0.765)	(0.598)	(0.625)
Winsorized (Saving Goal)	1%	1%	1%	1%
Winsorized (End Balance)	No	1%	No	1%
Baseline Savings Controls	No	No	Yes	Yes
Observations	845	840	806	802
R-squared	0.139	0.136	0.403	0.393

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

	(1) Reached Goal	(2) Reached Half Goal	(3) XS Savings Target Acct.	(4) Total End Savings	(5) Log End Savings
A: Centrality Only					
Eigenvector Centrality of Monitor	0.810*	0.502	760.6	53,918*	4.341**
	(0.452)	(0.432)	(1,614)	(29,745)	(2.008)
B: Proximity Only					
Social Proximity of Monitor and Saver	0.173*	0.0725	449.5**	13,150	0.868**
	(0.0870)	(0.128)	(207.1)	(9,006)	(0.368)
C: Proximity and Relatives					
Social Proximity of Monitor and Saver	0.156**	0.0526	362.9	16,181	0.991*
	(0.0697)	(0.118)	(250.8)	(11,878)	(0.496)
Monitor and Saver are Relatives	0.0404	0.0463	201.3	-6,624	-0.270
	(0.119)	(0.125)	(357.1)	(7,385)	(0.495)
D: Centrality and Proximity					
Eigenvector Centrality of Monitor	0.652	0.430	275.8	42,433	3.659*
	(0.462)	(0.429)	(1,712)	(25,118)	(2.114)
Social Proximity of Monitor and Saver	0.152	0.0681	465.1*	12,200	0.753*
	(0.0944)	(0.135)	(258.1)	(8,882)	(0.402)
Winsorized (Saving Goal)	1%	1%	1%	1%	1%
Controls	Yes	Yes	Yes	Yes	Yes
Observations	426	426	422	315	313

A, B, C, and D represent 4 different sets of regression specifications. Controls for monitor and saver demographics and assets along with savings goals are included in each regression. 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 6. Monitor Selection Order and Goal Attainment: Random and Endogenous

Dependent Variable: Reached Goal	(1) Monitor Random	(2) Monitor Endogenous	(3) Monitor Random	(4) Monitor Endogenous
Badge Order Quantile	0.0770	-0.0753	0.0733	-0.0838*
	(0.0723)	(0.0463)	(0.0728)	(0.0481)
Controls	Sav. Goal, Bank/PO	Sav. Goal, Bank/PO	Full	Full
Observations	429	382	429	382
R-squared	0.166	0.137	0.170	0.152

All regressions include village fixed effects. Standard errors are clustered at the village level. *** p<0.01, ** p<0.05, * p<0.1