

What Are the Headwaters of Formal Savings? Experimental Evidence from Sri Lanka

MICHAEL CALLEN

University of California San Diego

SURESH DE MEL

University of Peradeniya

CRAIG MCINTOSH

University of California San Diego

and

CHRISTOPHER WOODRUFF

University of Oxford

First version received March 2016; Editorial decision January 2019; Accepted March 2019 (Eds.)

The world's poor are seeing a rapid expansion in access to formal savings accounts. What is the source of savings when households are connected to a formal account? We combine a high-frequency panel survey spanning two and a half years with an experiment in which a Sri Lankan bank used mobile Point-of-Service (POS) terminals to collect deposits directly from households each week. We find that the headwaters of formal savings lie in sacrificed leisure time: households work more, and improved savings options generate an increase in labour effort in both self-employment and in the wage market. The results suggest that the labour allocation channel is an important mechanism linking savings opportunities to income.

Key words: Savings, Intertemporal labour supply, Household decision-making, High-frequency surveys.

JEL Codes: O16, D14, G21

1. INTRODUCTION

There is currently a strong push—in developed and developing countries alike—to provide financial services that bring the savings of the unbanked into formal institutions. The surge in interest stems from a recognition of the potential of formal savings to improve the efficiency of the overall economy by increasing the depth of financial markets, and from increased understanding of the potential of formal savings to improve the lives of people, most of them poor, who do not use banks. Half of adults worldwide remain unbanked, but 35% of the unbanked report obstacles to saving that can be overcome by improvements to products and regulation

The editor in charge of this paper was Jerome Adda.

(Demirgüç-Kunt and Klapper, 2012). Several recent studies suggest that overcoming the obstacles to savings can lead to very large positive effects on household expenditures (Kast *et al.*, 2012; Dupas and Robinson, 2013a; Prina, 2015; Brune *et al.*, 2016; Schaner, 2017).

The reasons behind these substantial impacts of formal savings on the well-being of the poor have been the subject of a great deal of theoretical and empirical work. Much attention has been focused on understanding why individuals fail to achieve their own savings goals, and how well-designed financial services help overcome these obstacles. Potential impediments to savings include problems with self-control arising from time-inconsistent preferences (Laibson, 1997) and other-control issues arising from pressure to make transfers to spouses or peers (Anderson and Baland, 2002). Commitment products (Thaler and Benartzi, 2004; Ashraf *et al.*, 2006) and savings accounts that provide women greater autonomy over savings (Ashraf, 2009; Schaner, 2017) may overcome these obstacles. Even those models that use a neoclassical framework to explain the large effects of recent savings experiments (such as Dupas and Robinson, 2013b) typically invoke production function non-convexities that prevent individuals from making high-return investments in productive inputs without savings instruments that allow them to accumulate enough capital. Under this framework, tools that help entrepreneurs save can have transformative effects by allowing people to make lumpy investments that provide higher steady-state returns and, thereby, enable them to escape low-investment poverty traps (Banerjee and Newman, 1993).

Studies of savings products targeting the poor have paid less attention to an older neoclassical explanation for the impact of new savings technologies: even with time-consistent preferences and a concave production function, savings interventions that generate an increase in the effective real interest rate will encourage intertemporal substitution, increasing the incentive to earn today so as to enjoy interest-driven consumption in the future. This logic is characteristic of workhorse macroeconomic models such as Lucas Jr and Rapping (1969) and Blanchard and Fischer (1989), and places emphasis on labour rather than investment as the pathway that connects improved savings options with increased incomes.

In this article, we present an experiment providing a powerful inducement to save. We combine the experiment with detailed, high-frequency surveys to capture the time path of changes in cash flows within households as money begins to flow into formal savings. Our experimental sample comprises 389 treatment and 406 control households who we survey periodically over two and a half years. We attempt to answer a simple question: what are the headwaters of formal savings? Different potential explanations imply distinct signatures in terms of the timing and nature of intra-household cash flows, allowing our high-frequency data to provide an informative account of the causal pathways behind savings impacts. If overcoming self-control problems related to contemporary consumption is key, we should find the headwaters of formal savings in decreased consumption. If other-control is central, then formal savings will flow from decreased net intra-household transfers and potentially from decreased use of collective informal savings products that increase savers' autonomy. If non-convexities in production are the driver, we will observe an initial transition period during which consumption, net transfers, or other investments decrease so as to generate savings balances, followed by lumpy investments in productive assets and then an increase in income resulting from improved marginal returns on capital. If the intertemporal substitution story is operative, we should instead see an immediate increase in labour as households respond to the higher interest rate. Of course, more than one of these mechanisms may be relevant.

We collaborated with a Sri Lankan bank to send staff door to door to offer households a weekly savings deposit collection service using a mobile Point-of-Service (POS) handset that prints out deposit receipts on the spot. The service was provided to randomly selected unbanked individuals

living at least 5 km from their nearest bank branch.¹ Starting 5 months before the collections began and continuing monthly for 13 months afterward, and again 19 and 25 months afterward, we collected detailed survey data, including income, expenditures, transfers, and labour market activity data. We use the detailed survey data to analyse changes in households' cash flows and labour allocations resulting from the savings treatment. The methodology we use is along the lines of the analysis in [Samphantharak and Townsend \(2009\)](#).

The weekly POS deposit collection service led to large increases in formal savings. The number of transactions with formal financial institutions per month quadrupled as a result of the treatment; the flow of savings into bank accounts almost doubled; and overall savings increased by more than 15%, or about US\$7 per month. The treatment generates an immediate increase in total household income, and in earned income, in particular. The increased income is used to fund the increased savings and to pay off loans. We see no change in transfers to or from other households, and no decrease in consumption.

The suggestion that increases in income are a source of savings rather than a result of savings leads us to examine labour market activity in more detail. The survey data for the main respondents show that labour hours, and, particularly, hours worked on the wage market jump within a month of the offering of the treatment.² Moreover, an immediate shift out of self-employment and into wage work occurs, with the occupational shift persisting over time. In aggregate, we find an increase in earnings from wage work, and no increase in self-employment income. The lack of any lag in the savings and income effects implies that the saving–invest–earn channel highlighted in the existing literature cannot be the only pathway through which savings accounts affect income. Instead, our data are consistent with the neoclassical model in which new savings options represent a change in intertemporal prices which affect the intertemporal allocation of consumption and, correspondingly, of labour ([Blanchard and Fischer, 1989](#)).³ Indeed, a calibration exercise suggests that the labour response we observe in the data is in line with what should be expected from the 5 percentage-point increase in real interest rates that is part of the account offered in our treatment.

This shift from self-employment to wage work, which ran counter to our initial expectations, is, again, consistent with the neoclassical model if households wishing to expand labour hours confront constant marginal returns on the wage market and diminishing marginal returns in self-employed microenterprises. To examine this possibility further, we split the sample into three subsamples based on activities reported during the baseline period: (i) individuals who are self-employed in the retail and service sectors, where we posit the ability to expand is limited given the very local markets typical in our rural setting; (ii) those active in the manufacturing sector, where sales are more easily made outside the local area; and (iii) those not involved in any self-employment. We find that those who were self-employed in retail and services are the source of all movement out of self-employment into wage work, and those who were self-employed in the manufacturing sector are the source of all increases in income and savings. The initial increase in

1. The intervention is designed to open new savings opportunity for unbanked individuals. We screened our experimental sample such that participants satisfied three criteria: (i) they received income at least weekly, *e.g.*, were self-employed or daily wage workers; (ii) they were accessible by a motorcycle and; (iii) they were unbanked, *i.e.*, they had no formal bank account or had not used a bank account in the previous month.

2. Nearly all wage jobs in this setting are day labour jobs, typically agricultural work that is paid daily in cash. Having a bank account does not directly increase access to these jobs. It is therefore unlikely that the increase in wage market hours is a mechanical consequence of having an account.

3. [Banerjee *et al.* \(2015\)](#) develop a model to structure an analysis of the effects of an expansion of a microcredit program in India. They also show that the increase in credit can lead to an increase in labour. Intuitively, their effect comes from the fact that credit makes the purchase of durable goods possible, and hence raises the marginal benefit of consumption. In the Blanchard and Fischer framework, savings accounts increase the labour effort immediately because the higher interest rate increases the benefit of future consumption.

income comes from an increase in hours worked, but also, apparently, from an increase in intensity of work per hour. The data provide evidence that the labour/leisure channel is an important piece of the answer to both the question of the headwaters of savings and the question of the channel through which the availability of savings services affects expenditures. Note that we are not the first to explore the labour channel. Dupas and Robinson (2013a) examine the impact of providing a non-interest-bearing account to self-employed workers in Kenya. The market vendors in their study make substantial use of the savings product, and Dupas and Robinson show that their working hours increase by just over 5% of the control group mean (p. 178, Table 3), an effect that is highly insignificant given their standard errors. Additionally, Schaner (2017) examines the effects of providing a high-powered savings subsidy on labour effort 3–4 years after the incentive expires. She finds a positive though insignificant increase of about 15% above the control group mean in hours worked in the business (Table 4). Bruhn and Love (2014) find that the opening of 800 Banco Azteca branches in rural Mexican communities—which substantially increased both formal borrowing and savings opportunities—increased employment by 1.4% over the pre-opening level (Table 4).

Labour hours may have increased for reasons other than the increase in interest rates. We examine the importance of several other channels. In a second experiment, we transitioned some of the control zones to collection through local lockboxes rather than door-to-door collection. We use this experiment to test whether the peer pressure and reminders that accompany door-to-door collection play a role in increasing income and savings. The data show that the lockbox treatment also led to increased labour market participation. These efforts, however, do not result in more savings for this group. This suggests some potential role for either peer pressure or limited attention.

The literature also points to savings accounts allowing participants to deal with other-control or self-control issues. Our data show no significant role for other-control, but we do find some evidence that the collection services disproportionately benefitted those most subject to self-control issues. To explore self-control, we use individual information on the time consistency of preferences at baseline, along with a list of experimental questions on the consumption of “sin goods”. Regarding other-control, the product allows the saver to interact privately with a bank staffer to make a deposit into the account solely in his or her own name, and thus may improve the ability of the account holder to deny the demands of others for transfers. Baseline survey questions on whether the participant has control over savings decisions in the household show that the treatment effect is not significantly larger for those initially lacking control over savings decisions.

Finally, we worried that the availability of formal savings mechanisms might crowd out informal savings. Rotating savings and credit associations, referred to locally as *seetus*, are a common means of informal savings in our sample. We exploit an experiment in which we randomized the intensity of treatment (Crépon *et al.*, 2013; Baird *et al.*, 2018) within the *seetus* used by study households at baseline to examine the interaction between formal and informal savings. We find that, if anything, our collection services and *seetus* prove to be complements in this setting. Overall, our results highlight the ability of savings to have a transformative effect on income even in an environment in which non-convexities in production do not appear to play a dominant role.

The remainder of the article is structured as follows: Section 2 presents the background and design of the experiment, and Section 3 develops a simple framework to organize the analysis. Section 4 presents the impact of the program on savings, the “headwaters” of personal liquidity, and on entrepreneurial activity. Section 5 investigates the extent to which increases in peer pressure or improvements in self- and other-control may also be responsible for generating our results, and Section 6 concludes.

2. SURVEY AND EXPERIMENTAL DESIGN

This section describes the procedure used to create an experimental sample and the protocol used to assign households to treatment. Our objective in sampling was to identify unbanked individuals who lived in remote households (accessible by motorcycle), and who received income at frequent intervals (and, as a result, were likely to benefit from weekly deposits). Figure 2 provides a complete overview of our study design.

2.1. *Sampling*

In 2008, Sri Lanka's National Savings Bank (NSB) began a program of remote deposit collection targeted at business owners within 1 km of NSB branches. The collections are made using a wireless POS terminal carried by a bank agent traveling by motorcycle. The POS terminal allows deposits to be recorded in the bank's accounts immediately; at the same time, the depositor receives a deposit confirmation showing the new account balance. Agents typically visit a given merchant once per week.

For this project, NSB agreed to provide branches in the towns of Bandarawela and Mahiyangana with an additional POS terminal and agent to make collections from more remotely located rural areas. While NSB's goal for the initial POS program was to reach highly liquid merchants, the goal of this project was reaching lower-income, under-banked households.

2.1.1. Sampling criteria. We agreed on selecting areas by first going 5–15 km from the branch along the main roads leading away from each of the two towns, and then taking households within 1 km of the main road. Figure 1 maps the study households around Bandarawela and Mahiyangana and shows their proximity to the road network and their distance to NSB branches; this confirms that nearly all households selected for the study lived between 5 km and 15 km from the nearest NSB branch.⁴ We began by conducting a listing exercise, visiting every household in the sampling area, and listing all economically active adults aged 18–65 in the households. The screening survey collected information on: the frequency of pay from employment; whether the individual had a bank account, and, in the case of individuals who did have a bank account, whether the individual had made any banking transaction within the past month; and whether the individual participated in one or more *setus*, the rotating savings and credit associations that are common in Sri Lanka. A total of 5,081 individuals in 3,657 households were listed during this exercise.

Before drawing the sample, we eliminated individuals whose profile made the weekly POS collection either infeasible or inappropriate. The first group included those living in houses that could not be reached by motorcycle. By this criterion, the house had to be within 500 feet of a point accessible to the agent via motorcycle. We also eliminated those who received income payments less frequently than once per week. The weekly collection is of most value where income streams are either daily or weekly. In practice, this means the sample is largely made up of individuals who are either self-employed or daily wage workers. This is important to keep in mind because the daily paid workers may have more ability to change hours worked than, for example, monthly salary workers. Finally, we eliminated those who had made a deposit in a bank account within the past month, since we expected the behaviour of those using the bank regularly would not be affected as much by the deposit collection services. Note that the last two criteria were applied

4. Supplementary Appendix Figure A1 depicts the distribution of distances between study participants and the nearest bank branch travelling along the nearest road.

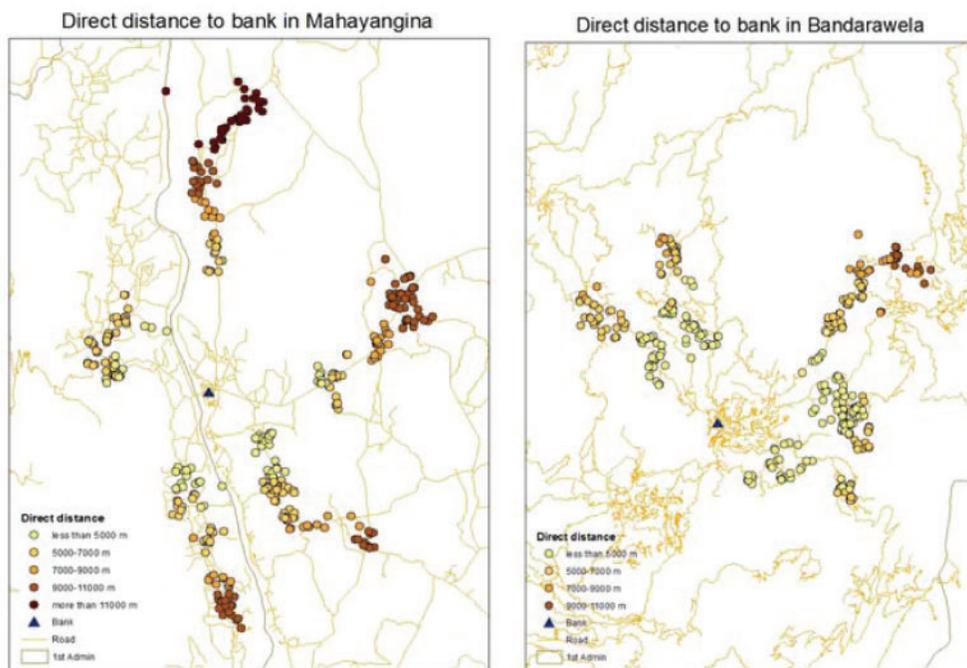


FIGURE 1

Maps of the experimental sample in Mahayangina and Bandarawela

at the individual level, so that if any listed member of the household received income payments at least weekly and had not used a bank account in the previous month, that individual remained eligible even if the criteria eliminated other household members. We also only sampled prime-age working individuals, defined as being between 18 and 65.

2.1.2. Creating zones. NSB quite reasonably felt that their agents could not refuse a deposit from anyone approaching them while making their rounds. In the region of our experiment, the households are scattered into small clusters rather than distinct villages. Therefore, we conducted the randomization at the level of clusters of houses. We identified 156 distinct clusters of households (hereafter referred to as “zones”), separated by at least a stretch of road long enough to make spillovers between treatment and control zones unlikely. In addition, in this area, agriculture is the primary activity in the sampled areas around both towns. Near Bandarawela, the main crops are tea and vegetables; near Mahiyangana, the main crops are rice and vegetables.

2.1.3. Sampling within zones. With these zones in mind, we selected the experimental sample in several stages. Because we were particularly interested in the interaction between formal and informal savings—specifically whether the introduction of deposit collection changes *seetu* participation—we began by selecting a sample of individuals who participate in a *seetu* in which all of the members reside within 1 km of the initial respondent. We refer to these *seetus* as “close *seetus*.” The first step was to select an individual in each cluster who is a member of a close *seetu*.⁵ There were 93 zones with at least one member of a close *seetu*. We randomly selected

5. These individuals also satisfied the sampling requirements of being unbanked and being accessible by motorcycle.

| | | | | | | | | |
|------------------------|---|------------|--|------------|--|------------|---|------------|
| LISTING | Originally Listed: Households within 5-15 km of bank branch, 1 km of main road 163 Zones 3,657 Households 5,080 Individuals | | | | | | | |
| | Determining Eligible Sample for Study: Imposing Study Eligibility Requirements: (1) Self-Employed (not full-time wage worker). (2) Home is accessible by Motorbike. (3) No formal bank transaction in past month. 156 Zones 1,469 Households 1,713 Individuals | | | | | | | |
| ELIGIBILITY & SAMPLING | Assigning Strata & Drawing Sample for Study: | | | | | | | |
| | S1: Members of 'Close Seetus' w/in 1 km: <i>N</i> =101 | | S2: Females who are not S1s, no duplicate households. <i>N</i> =228 | | S3: Males who are not S1s, no duplicate households. <i>N</i> =207 | | S4: Fellow <i>seetu</i> member of S1s who are not S2s or S3s, no duplicate households. <i>N</i> =259 | |
| TREATMENT ASSIGNMENT | Treatment Assignment at Zone Level | | | | | | | |
| | Allocated to Control, Surveyed at Baseline 78 Zones 406 Household & Individuals | | | | Allocated to Treatment, Surveyed at Baseline 78 Zones 389 Households & Individuals | | | |
| | S1: 53 | S2: 103 | S3: 109 | S4: 141 | S1: 48 | S2: 125 | S3: 98 | S4: 118 |
| ATTRITION | Lost due to panel attrition in the household survey by Round 20: 0 zones, 27 individuals & households | | | | | | | |
| ENDLINE SURVEY | Surveyed at Endline | | | | | | | |
| | Allocated to Control 78 Zones 390 Households & Individuals | | | | Allocated to Treatment 78 Zones 378 Households & Individuals | | | |
| | S1: 50 | S2: 99 | S3: 105 | S4: 136 | S1: 47 | S2: 120 | S3: 95 | S4: 116 |

FIGURE 2
Overview of study design

one member of a close *seetu* in each of these 93 zones, and then re-surveyed these 93 individuals to collect the list of names of other participants in their *seetu*. We selected only one individual per zone to ensure that seeds did not come from overlapping *seetus*. This list of names permitted us to “snowball” sample other members of their *seetu*. If individuals brought in via the snowball did not meet the other sampling requirements, we sampled them nonetheless.

The second step of the sampling was aimed at achieving gender balance. The listing yielded a much smaller number of women than men meeting the employment/earnings frequency criteria explained in Section 2.1.1 above. We also required that women eligible for sampling at this stage were not a member of a close *seetu*. We selected up to five women in each of the zones. There were more than five women in only two of the 156 zones, implying that we had selected essentially all the women in the comprehensive listing who qualified according to the earning frequency and bank transaction criteria. We then extended the sample by selecting an additional one or two men randomly from each zone.⁶ Finally, for each close *seetu* we then randomly selected 20% of its membership to participate in the survey. The resulting intake sample of 829 has 432 women and 397 men.

We sampled only one individual per household. Between August 2010, when the baseline was conducted, and November 2010, when the randomization was carried out, 34 respondents attrited from the survey, leading to an experimental sample of 795 who were included in the

6. Because some zones have no males, we select up to two additional males per zone and use those to fill the remaining sample.

randomization. Twelve households did not answer at least two post-treatment survey waves, our criterion for inclusion in the panel analysis, leaving us with a final sample of 783 for analysis. In all, 44% of the sample (344) are members of close *seetus*, with the remainder (439) coming from the broader random selection.

2.2. Treatment assignment

The randomization of collection services was carried out at the zone level, stratifying and balancing on data from the baseline survey. We stratified on quartiles of total savings balances, above/below median of household expenditures, and whether the zone included a close *seetu*. We then re-randomized to achieve balance on a set of variables related to savings behaviour and income generation.⁷ The study sample thus consists of 78 zones with 389 individuals for treatment and 78 zones with 406 individuals for control.⁸

In December 2010, we offered to open accounts at the local NSB branch for those selected into treatment.⁹ The account opening required that the individual come to the bank branch. During the first 2 weeks of the month, project research assistants were at the bank branches to help participants with the required paperwork. The minimum balance required to open an account is 500 LKR (~US\$4.50); we paid this for all treated individuals who opened an account.¹⁰ Of the 389 individuals selected for treatment, 347 (89%) opened an account. POS collection services began 13 December 2010, in both branches. One main road from town was serviced each day of the week.

2.2.1. Unbundling experiment. Six months into the main experiment, we began a series of unbundling experiments. In a randomly selected and well-balanced subset of the control and weekly home visit treatment arms, we worked with participating households to establish locked collection boxes in each selected cluster of houses. The lockboxes were visited by the NSB agent on a weekly or biweekly basis. Participants were given envelopes and stickers with which to seal the envelopes. On the stickers, they wrote their account name, account number, and the amount of the deposit. Initially the participants were given several stickers with their name and account number filled in by project staff; all they had to do was fill in the amount. Participants placed deposits in envelopes and dropped the envelope into the box. The NSB agent came, together with project staff, on a scheduled day and time with a key to open the box to collect the deposits made by participants. The agent carried his POS terminal with him, made the deposits on the spot, and then left the receipt for the deposit in a sealed envelope in a compartment built onto the side of the box. Participants then returned to collect the receipt. Thus, the deposits were made into the NSB

7. Specifically, to achieve balance we randomize using the ‘big stick’ approach, whereby randomizations are redrawn until the minimum p -value among a pre-specified set of balance variables is greater than some threshold (see Bruhn and McKenzie, 2009) for a discussion of the advantages and disadvantages of this approach. In our case, we select a threshold of 0.21. The variables on which we balanced were: agricultural employment, education, gender, present bias, monthly income, whether the individual had a formal bank account, balances of formal and informal savings, total and informal monthly savings amounts, and whether anyone in the household was either a grantor or recipient of an outstanding informal loan. We also examine whether our core results are robust to using randomization inference, where potential treatment assignments satisfy these conditions.

8. *Seetu* members who were brought into the study through the listing of all *seetu* members received treatment based on their zone of residence, not on the status of their respective *seetu* seed.

9. If individuals in treatment zones either had an existing account (and so were excluded during screening), or opened one during the study, they were eligible to use the deposit collection service.

10. The opening deposit is trivial relative to median income of the participants (22,000 LKR). The exchange rate during the course of the study was approximately 111 LKR per U.S. dollar.

exactly as in the door-to-door collection, but in a manner that was more efficient for the agent. In practice, since the agent arrived around the same time on the announced day, participants often came to the box at that time and made the deposit in person. [De Mel *et al.* \(2013\)](#) describe the impact on savings of the collection boxes by using the experiment to parse out the role of habit formation and salience in savings. Here, we use the additional experiment introducing weekly box collection to the control to try to adjudicate between mechanisms linking deposit collections to changes in labour supply, in a manner we discuss in Section 5 below.

2.2.2. Saturation experiment. Finally, we exploited the additional listing of *seetu* members described in Section 2.1.3. to introduce a randomized saturation experiment to investigate whether the introduction of a formal savings account affected how individuals interacted with their informal savings organizations. We began from the group of 39 close *seetus* that had been selected into the treatment; in each of these *seetus*, 20% of the membership participated in the household survey. We then offered treatment to (but did not survey) an additional 20% of the members in one-third of these groups, and to an additional 40% of members in another one-third of the groups. With respect to treatment, 45 of the close *seetus* are untreated, 13 are treated at 20%, 13 are treated at 40%, and 13 at 60% saturation. The comparison of the fortunes of the constant 20% of members studied across the randomized variation in the treatment saturation allows us to isolate externality effects of formal financial services on the welfare of the membership of informal savings networks.

2.3. *Surveys and administrative data*

We conducted a baseline and three additional pre-treatment surveys for the full sample at monthly frequency between August and November 2010. We randomly allocated 498 households to a sample in which we continued to conduct monthly surveys, and 297 households to a sample in which we conducted quarterly surveys. These both reduced survey costs and allowed us at least a partial test of whether survey frequency affects deposit or aggregate savings behaviour. For both monthly and quarterly survey groups, we conducted surveys at the defined frequencies through November 2011. We then conducted monthly surveys of everyone in December 2011 and January 2012, and longer-term follow-up surveys in July 2012 and January 2013. Thus, for the monthly survey sample, we have five pre-treatment surveys and 15 post-treatment (13 monthly and two semi-annual) surveys; for the quarterly survey sample, we have four pre-treatment and eight post-treatment (four at quarterly intervals, two monthly, and two semi-annual) surveys. To avoid confounding the primary results, in the initial analysis we drop the 192 treatment and 150 control individuals involved in the unbundling exercise described above as soon as their treatment status changes. That is, we retain them in the sample until the point when they change from their first treatment designation (either control or in-person deposit collection) to lock box deposit collection.

Supplementary Appendix Table A1 details the timing of the surveys and shows which surveys are included in the sample we use here. The result is a full 30 months of data for the core sample (92 zones; 18 months at high frequency) and, for a subsample of 64 zones, data from the 12 months prior to the beginning of the unbundling experiment. Results are very similar if we use only the sample of 92 zones, but the precision of the short-term estimates improves with the inclusion of the additional group that receives the core treatment for 6 months. Our analyses use individual-level fixed effects, and we cluster standard errors at both the zone and survey round using the method developed by [Cameron *et al.* \(2011\)](#) as a way of accounting for both the substantial autocorrelation present in high-frequency household data and the effect of local and

period-specific shocks. For balance tests using pre-treatment data, however, we cluster only at the zone level.

We undertook this project with the aim of answering the simple but vexing question: what is the root source of money that is newly brought into the formal financial sector? When people begin to use formal savings, what other behaviours in the household change to allow this liquidity to be deposited in a bank? Candidate explanations are: (i) that saved capital is substituted from cash in the mattress; (ii) that greater discipline from formal savings causes expenditures to decrease; (iii) that formal savings come at the cost of informal mutual insurance networks; or (iv) that individuals work more in order to save. The survey was designed with these sources in mind. The heart of the survey instrument is a cash flow analysis (see Supplementary Appendix C for the relevant survey questions) for both the sampled household and the sampled individual, who was always the respondent. To unpack the headwaters of formal finance, we need to be able to construct balances of financial flows at the household level. Thus, our survey was designed to capture monthly inflows (earned income, transfers from individuals, and other transfers) and outflows (expenditures, transfers to other individuals, savings and loan payments). While these modules are meant to comprehensively cover sources and uses, this module does not capture expenditures related to business and agricultural investment, which have their own detailed modules elsewhere in the survey.¹¹

Because we expect individual members of a household to have better information about their own earnings and transfers than those of other members of the household, we also ask about the respondent's individual cash flow. Questions asked of the individual respondent net out her/his transfers to and from other members of the household to calculate personal income, and then allocate this personal income to individual consumption items. Significantly, we ask the direct questions on the amount saved through various means (in the NSB account, in other formal accounts, in cash, in *seetus*, etc.) with regard to the individual respondent. Since savings decisions are likely made at the household level in many households, we make use of the aggregate household income and expenditure data. But we also use the more detailed individual data to focus on the outcomes of the individual respondent. The enumerators were trained to check that the sources of cash matched the uses of cash for the individual. Where the initial responses yielded differences, the enumerators pointed out the inconsistency and re-asked the income and expenditure questions.¹²

The decision to focus the detailed earnings questions on the activities of the participant him/herself represents a tradeoff. On the one hand, we focus on data the participant knows best. On the other hand, we will be somewhat limited in answering the "headwaters" question if the changes in income, expenditure, or savings come from changes in the behaviour of other members of the household. That is, if we identify that increases in savings in banks are associated with increases in transfers from the spouse, we know only indirectly whether the spouse increased his income (and if so, we do not know how) or decreased his formal or informal savings. But the aggregate household data allow us to identify the sources of changes in savings arising from income and expenditure patterns of other household members up to a point.

11. Both the monthly and quarterly surveys asked respondents about income and expenditure flows during the previous month. We asked about the previous month even in the quarterly survey to minimize recall issues. For investments, including business investments and financial investments in jewellery, loans, etc., the questions were asked with reference to the time since the previous survey (i.e. the previous month for the monthly sample and previous 3 months for the quarterly sample).

12. Samphantharak and Townsend (2009) assign mismatches in monthly income and expenditure to changes in cash holding. We measure cash holdings along with other financial assets. To the extent that individuals do not want to report increases in cash holdings, they might be tempted to increase reported expenditures instead. However, over years of conducting surveys in Sri Lanka, we have found respondents to be generally very willing to discuss financial matters openly, and we do not detect any patterns in the data which suggest significantly misreported cash balances.

In addition to the detailed survey data, we have administrative data from NSB for the accounts directly linked to the project. These detail each deposit and withdrawal, as well as other transactions (*e.g.* interest payments). We use the institutional data in this article to study the determinants of uptake and account usage and institutional balances; all remaining impact estimates rely exclusively on information from the household survey.¹³

2.4. *Balance and attrition*

Table 1 examines the balance of the experiment. Treatment was assigned at the zone level and stratified on whether a zone contained a close *seetu* (meaning that all the members of the *seetu* were resident in that zone), as well as on average household expenditures and baseline average savings in each zone. Stratification was done using baseline (round 1) data, and we test for balance using the remaining pre-treatment rounds (2–5). Table 1 shows balance on a range of variables measuring savings, income, and expenditure. Only one of the 35 difference-in-means tests, for monthly savings in cash, shows a statistically significant difference. However, while not significant at conventional levels, some baseline differences are fairly large. For example, earned income over the previous month is about 1,787 LKR lower in the treatment group compared to the control, which is similar in absolute value to the earned income effect (after winsorizing) that we estimate below. We use individual fixed effects throughout our analysis to help address potential imbalances.

Supplementary Appendix Table A2 examines correlates of survey attrition. Overall, the resurvey rates were high. Ignoring attrition from the sample that occurred prior to the assignment of treatment, only 12 of 795 individuals were not surveyed at least twice pre-treatment and twice post-treatment (our criteria for inclusion in the panel sample). We also examine the 20 individuals who did not complete the last high-frequency survey (wave 18 in January 2012). Most importantly, attrition does not differ in the treatment and control groups whether we examine simple differences between the two groups or whether we control for other covariates.

Despite the overall balance of the experiment, the number of zones to which treatment was assigned was relatively small (156), and, due to sampling rules intended to locate as many self-employed females as possible, there is substantial cross-zone variation within the sample in variables such as gender and *seetu* participation. For example, 28 zones contain all-male and 10 zones all-female core respondents. Potentially, due to this clustering of gender by zone, the experiment is imbalanced when we look among men only or among women only.¹⁴ Supplementary Appendix Figure A2 shows the densities of one of the most severely imbalanced variables, the log of personal income. While the imbalances are not visually severe, the figure illustrates how the problems within exclusively male and exclusively female subsamples counteract each other to lead to a full treatment/control comparison that is well balanced.¹⁵

13. Our analysis of a subsequent unbundling experiment takes place within the treatment group, and hence makes extensive use of the institutional data; see De Mel *et al.* (2013).

14. To perform randomization inference, we re-ran our treatment assignment protocol 1,000 times to obtain a corresponding number of artificial draws. This allows us to examine how the actual subgroup balance of our experiment compares to the distribution of subgroup treatment/control differentials that could have been obtained given our sample and randomization routine. This illustrates that we were simply unlucky in realizing a group of treatment females (males) that are substantially richer (poorer) than the control (p -value < 0.01 for highly imbalanced variables such as household income among males).

15. Interestingly, if we use only the data from the round 1 survey (on which the stratification was conducted), no significant signs of subgroup imbalance are observable. The imbalances result from a gap that opens between treatment and control within men and within women in rounds 2–5, and hence is only observable because we collected multiple high-frequency pre-treatment surveys as suggested by McKenzie (2012).

TABLE 1
Randomization verification

| Baseline outcome | Number of baseline surveys | | Pre-treatment mean | | Treatment-control differential | SE of difference |
|--|-------------------------------|-----------|-----------------------|-----------|-----------------------------------|---------------------|
| | Control | Treatment | Control | Treatment | | |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Number of transactions with a formal financial institution during previous month | 1,445 | 1,373 | 0.41 | 0.40 | -0.013 | (0.046) |
| Female | 1,445 | 1,373 | 0.52 | 0.54 | 0.016 | (0.038) |
| Discount factor (delta) | 1,445 | 1,373 | 0.92 | 0.93 | 0.0090 | (0.0095) |
| Time consistency factor (beta) | 1,445 | 1,373 | 0.97 | 0.99 | 0.020* | (0.011) |
| Total number of seetus in which respondent participated in previous month | 1,445 | 1,373 | 1.23 | 1.16 | -0.069 | (0.11) |
| Total payouts from seetus in previous month | 1,445 | 1,373 | 948 | 817 | -127.4 | (176.0) |
| Number of transactions with a financial institution other than NSB over previous month | 1,445 | 1,373 | 0.41 | 0.39 | -0.018 | (0.044) |
| Total amount saved through ROSCAs over previous month | 1,445 | 1,373 | 1,470 | 1,367 | -104.0 | (190.7) |
| Total amount withdrawn from ROSCAs over previous month | 1,445 | 1,373 | 1,126 | 1,011 | -115.1 | (181.1) |
| Total amount put into informal savings over previous month | 1,445 | 1,373 | 2,765 | 2,480 | -289.5 | (316.8) |
| Total amount withdrawn from informal savings over previous month | 1,445 | 1,373 | 1,940 | 1,738 | -206.7 | (267.9) |
| Total change in household cash balance over previous month | 1,445 | 1,373 | 585 | -125 | -712.7* | (397.1) |
| Total amount received in loans over previous month | 1,445 | 1,373 | 2,916 | 2,803 | -110.7 | (396.7) |
| Aggregate income over previous month | 1,445 | 1,373 | 22,792 | 22,144 | -657.7 | (969.2) |
| Aggregate consumption over previous month | 1,445 | 1,373 | 18,672 | 17,907 | -769.4 | (952.8) |
| Earned household income over previous month | 1,445 | 1,373 | 22,195 | 21,117 | -1081.7 | (1007.1) |
| Total household income over previous month | 1,445 | 1,373 | 29,754 | 28,111 | -1630.3 | (1413.6) |
| Personal income taken from informal savings over previous month | 1,445 | 1,373 | 13,229 | 11,767 | -1451.7 | (991.2) |
| Transfer from spouse over previous month | 1,445 | 1,373 | 3,713 | 4,043 | 323.9 | (401.8) |
| Transfer from outside the household over previous month | 1,445 | 1,373 | 3,589 | 3,437 | -150.0 | (445.7) |
| Total personal income over previous month | 1,445 | 1,373 | 21,098 | 19,819 | -1267.7 | (1163.2) |
| Transfers to spouse over previous month | 1,445 | 1,373 | 1,945 | 2,113 | 167.4 | (242.9) |
| Transfers to children over previous month | 1,445 | 1,373 | 260.66 | 253.92 | -8.08 | (36.5) |
| Transfers to anyone over previous month | 1,445 | 1,373 | 3,093 | 3,076 | -19.8 | (267.6) |
| Monthly savings into banks | 1,445 | 1,373 | 780 | 635 | -141.8 | (137.4) |
| Monthly savings in cash | 1,445 | 1,373 | 2,413 | 2,157 | -246.2 | (236.4) |
| Monthly savings through ROSCAs | 1,445 | 1,373 | 1,481 | 1,407 | -75.5 | (198.9) |
| Monthly savings through other means | 1,445 | 1,373 | 416.96 | 388.43 | -30.6 | (-60.1) |
| Total Monthly savings over previous month | 1,445 | 1,373 | 5,587 | 4,932 | -646.6 | (467.1) |
| Respondent reports being self-employed in non-ag enterprise during previous month | 1,445 | 1,373 | 0.51 | 0.53 | 0.019 | (0.043) |
| Wage hours worked | 1,445 | 1,373 | 5.95 | 5.65 | -0.30 | (0.96) |
| SE hours worked | 1,445 | 1,373 | 24.26 | 25.48 | 1.23 | (2.74) |
| Total asset purchases over previous month | 733 | 723 | 561 | 490 | -74.4 | (153.7) |
| Total value of business inventory | 733 | 723 | 35,491 | 41,796 | 6203.8 | (8115.8) |
| Total business expenses over previous month | 733 | 723 | 20,344 | 22,860 | 2484.9 | (3287.9) |
| Total business sales over previous month | 733 | 723 | 31,561 | 33,561 | 2013.0 | (4117.1) |
| Total business income over previous month | 733 | 723 | 10,500 | 10,734 | 233.7 | (1099.2) |

Notes: This table reports on the pre-treatment balance of the main outcomes of the study as well as a broad set of covariates. Variables are Winsorized at the 99% level. Data include 4 months of pre-treatment surveys. The treatment-control differential is estimated using a regression including fixed effects for survey wave, and clustering standard errors at the zone level (treatment unit). The unit of observation is the individual/month. Regressions pool samples surveyed monthly and samples surveyed quarterly. * significant at 10%; ** significant at 5%; *** significant at 1%.

Given this issue, we pursue several other approaches to test whether our experiment provides clean identification. First, we test for subgroup balance using the three self-employment classifications that are explored later in the analysis (self-employed in manufacturing, self-employed in trade and services, and not self-employed). Supplementary Appendix Table A3 demonstrates that the study is well-balanced within subgroups. Second, we exploit our multiple pre-treatment survey waves to subject the experiment to a test of parallel trends. Supplementary Appendix Table A4 estimates intercept and time-trend slope differentials between the treatment and the control prior to treatment. This table presents the results of 68 tests; none of the intercept terms are significant, one slope term is significant at 90% and one (total consumption) at 95%. This suggests that the pre-treatment rates of change in the treatment and control were broadly comparable. Nonetheless, because of the imbalance within gender, we avoid analysis by gender and discuss imbalance issues carefully when we move to the analysis of the *setu* experiment in Section 5.

3. THE INTERTEMPORAL SUPPLY OF LABOUR

A substantial literature in macroeconomics examines the role of interest rates in governing the intertemporal allocation of labour [Lucas Jr and Rapping \(1969\)](#). Central to these models is the idea that increasing the real interest rate raises the relative return to working in the present relative to working in the future. In this section, we provide a standard model to develop this point.

Consider an agent who consumes c_t and works h_t in period t with utility

$$u(c_t, l_t) = \frac{1}{(1+\rho)^t} [u(c_t) - h_t^\gamma]. \tag{3.1}$$

In period t the agent has wealth A_t , which is a function of the prior period's assets and wage income.

$$A_t = (1+r_t)A_{t-1} + w_t h_t - c_t. \tag{3.2}$$

Several simplifying assumptions are worth noting. First, we ignore the issue of risk in the income stream. Increasing risk would lead to precautionary savings, but, because our experiment changes the interest on savings in isolation, any effect of treatment on the level of precautionary savings is likely to be second order. Second, we assume that utility from consumption and disutility of effort are additively separable, implying that the marginal utility of consumption is independent of the disutility of effort. We make the standard assumption of positive and concave consumption in utility. Finally, we assume that wage rates are not affected by the experiment. Because the treatment group is a small part of the regional population, we take the increase in the interest rate as isolated from any effect on wages or prices. The Euler equation governing the intertemporal allocation of labour is:

$$\frac{h_t}{h_{t+1}} = \left[(1+r_{t+1}) \frac{w_t}{w_{t+1}} \frac{1}{(1+\rho)} \right]^{\frac{1}{\gamma-1}}. \tag{3.3}$$

If the POS savings treatment provides an improved savings option,¹⁷ then for the experimental treatment group, the increase in r , holding other parameters fixed, represents an increase in lifetime

16. The corresponding Euler equation for the intertemporal allocation of consumption is $\frac{u'(c_t)}{u'(c_{t+1})} = \frac{(1+r_{t+1})}{(1+\rho)}$.

17. The NSB accounts paid an interest rate of 5%/year. Plausibly, the effective interest rate on cash held at home (a common situation in our sample) is negative because money might be stolen or used by someone other than the respondent.

wealth for those in the treatment group with positive savings balances. The treatment thus implies an income effect in addition to the substitution effect. In this formulation, the elasticity of the labour ratio with respect to the gross interest rate $1+r_{t+1}$ is $\frac{1}{\gamma-1}$ and so is a function of the curvature of the disutility of effort function. The Euler equation communicates a standard feature of these models: if the cost of effort function is convex ($\gamma > 1$), then the substitution effect will dominate the income effect such that an increase in the future interest rate should lead to more work in the current period relative to the future period.¹⁸ That is, holding the period-specific wage rates constant, the increase in r will lead to an increase in hours in period t relative to period $t+1$.

4. RESULTS

The changes in labour supply and intertemporal consumption that come from the model are driven by an increase in the interest rate on savings. If the collection services did increase the effective interest rate realized by program participants, then we should observe both a robust take-up of the product, and an increase in total savings. These outcomes are the analogous to the first stage in the experiment. We begin, then, by looking at these preliminary outcomes.

4.1. Uptake and usage

Supplementary Appendix Table A5 uses pre-treatment averages (data from rounds 1–5) to predict subsequent uptake and usage of the accounts among the 389 individuals offered the treatment. Account usage appears progressive in that those most likely to make any deposit are low-income females who score poorly on the digitspan test. At the same time, a strong predictor of use of our accounts is a history of formal savings prior to the initiation of the experiment. Low-expenditure households use the accounts most frequently for both deposits and withdrawals, but none of our determinants are significant in explaining the total number of transactions in program accounts. Among those who make deposits, the two significant determinants of the log of deposits are having more education and having higher formal savings at baseline. Those with high baseline savings also make significantly more withdrawals from program accounts, however, leading to final savings balances that are not significantly higher than those who had no formal savings at baseline. Final savings balances are strongly progressive in expenditures (meaning that poor households save significantly more through the program than wealthier households even in absolute terms), and baseline *seetu* participation is also predictive of lower final savings balances. The progressivity of final savings balances is encouraging, indicating that weekly home visits with the POS terminals are particularly effective at generating deposits from low-income households.

4.2. Savings impacts

The primary objective of the deposit collections is to increase bank savings. We estimate the impact of deposit collections on banking activity running regressions of the form:

$$y_{izt} = \beta \cdot Treated_{zt} + \gamma_i + \eta_t + \varepsilon_{izt} \quad (4.1)$$

y_{izt} is the individual savings outcome, where i refers to individual, z to the zone, and t to the month of the study; $Treated_{zt}$ is a dummy variable equal to one for individuals in treatment zones

18. There are, similarly, income and substitution effects in consumption. The increase in r favours consuming tomorrow over today, but the increase in wealth increases aggregate lifetime consumption.

after month six of the study (January 2011), γ_i are individual fixed effects and η_t are month fixed effects.¹⁹

We note that all results presented in this article, other than the results in Table 1, are from regressions that include individual and month fixed effects. Table 2 provides estimates corresponding to this specification. Panel A columns (1)–(3) uses administrative savings data as outcomes and columns (4)–(7) uses survey data as outcomes. We note that, throughout the article, wherever we report control means, we do so for the untreated group pooling across all waves of data collection. Estimates in column (1) indicate that monthly bank transactions increase by 1.73 from a base of 0.01 and, correspondingly, in column (2) we see that monthly deposits increase from a base of 1.68 LKR by 1,143 LKR; column (3) shows that monthly deposits minus withdrawals increase by 425 LKR from a base of 0.85 LKR. The administrative data, therefore, indicate a substantial increase in program savings as a result of treatment. Turning to estimates based on survey outcomes, column (4) indicates that treatment quadruples reported transactions from a control average of 0.5 transactions per month to a treatment average of over two per month in the treatment group. Column (5) indicates a near doubling of bank savings. Column (6) suggests no corresponding reduction in informal and cash savings, and column (7) suggests indicates a 16.9% increase in total savings. The degree of the effect of deposit collections on monthly bank transactions is also depicted in Figure 3. Despite a large increase over the control group, the data show that treatment households only make deposits through the collectors during half of the weeks in which visits occur. Comparing columns (3) and column (5) suggests that treatment might also lead to an increase in deposits in other formal accounts. While we would be cautious about that conclusion—both because the difference of 265 LKR per month is not significantly different from zero and because the survey data may be noisier—the data suggest that the NSB savings are not coming directly from other formal savings accounts.

We re-randomized in this experiment to achieve balance. To address potential concerns for inference, Supplementary Appendix Table A6 repeats the estimates from Table 2, including p -values obtained using randomization inference.²⁰ This approach tests a null of no treatment effect for any participant in the study. Under this null, we can generate a sampling distribution for β simply by estimating specification (4.1) on any treatment assignment that satisfies our assignment protocol.²¹ This approach yields similar conclusions. In particular, we reject the null of no effect on total savings for any participant at the 5% level.

This preliminary look at savings therefore indicates that treatment leads to a statistically significant and economically meaningful increase in total savings. There is no evidence that the savings deposited in the NSB account are crowding out either other formal savings or informal savings. The increase in savings suggests that participants viewed the collection services as a substantial change in their portfolio of financial services, opening the possibility that the treatment led to more fundamental changes in their behaviour.

4.2.1. Savings impacts by time period. After the introduction of treatment, we followed the subjects for 2 years, allowing for an examination of whether participation fluctuated over time.

19. The specification therefore provides difference-in-difference estimates of the effect of treatment. Correspondingly, we omit a dummy variable equal to one for all subjects in the treatment group and a dummy variable equal to one for the post-treatment period, as those are fully accounted for by the individual and time fixed effects. Standard errors in all regressions are clustered at the zone level (to account for the design effect) and at the survey wave level (to account for contemporaneous shocks in outcomes) using the technique of Cameron *et al.* (2011).

20. We point the interested reader to Gerber and Green (2012) for a full description of this approach to inference.

21. In practice, we obtain 1,000 potential treatment assignments.

TABLE 2
Estimated impacts of deposit collections on savings outcomes

| | NSB institutional data | | | Survey data | | | |
|--------------------------|---------------------------------|----------------------|----------------------------------|-------------------------------|---------------------|-------------------------------|---------------------|
| | No. of program transactions (1) | Program deposits (2) | Program deposits-withdrawals (3) | Total no. of transactions (4) | Bank savings (5) | Informal and cash savings (6) | Total savings (7) |
| Panel A: Pooled | | | | | | | |
| Treated | 1.73*** (0.13) | 1143.5*** (158.9) | 424.9*** (95.1) | 1.57*** (0.15) | 689.6*** (198.6) | 42.2 (322.0) | 883.0** (376.7) |
| Observations | 9172 | 9172 | 9172 | 9172 | 9172 | 9172 | 9172 |
| R-squared | 0.68 | 0.28 | 0.090 | 0.58 | 0.25 | 0.32 | 0.33 |
| Number of HHs | 783 | 783 | 783 | 783 | 783 | 783 | 783 |
| Panel B: By period | | | | | | | |
| Treated: months 1-4 | 1.74*** (0.22) | 1112.3*** (187.1) | 660.2*** (178.6) | 1.52*** (0.17) | 660.6*** (220.1) | 217.1 (345.9) | 964.3** (428.1) |
| Treated: months 5-8 | 1.74*** (0.17) | 1060.8*** (190.7) | 269.6** (112.3) | 1.65*** (0.16) | 616.4*** (229.5) | -184.2 (471.6) | 634.8 (489.9) |
| Treated: months 9-13 | 1.69*** (0.23) | 1283.2*** (230.1) | 279.2* (159.8) | 1.55*** (0.29) | 818.5*** (271.0) | 35.9 (409.7) | 1036.8** (469.5) |
| 6 months post-treatment | 0.42*** (0.052) | 247.8** (125.3) | -527.8*** (164.8) | 0.32*** (0.11) | 436.0 (295.5) | 80.5 (495.9) | 926.2* (538.5) |
| 12 months post-treatment | 0.36*** (0.034) | 806.7*** (194.5) | -386.3* (207.7) | 0.19** (0.099) | 10.9 (222.6) | -204.0 (412.7) | -220.4 (480.3) |
| Control group mean | 0.01 | 1.68 | 0.85 | 0.51 | 756.81 | 4,307.43 | 5,213.32 |
| Observations | 10,055 | 10,055 | 10,055 | 10,055 | 10,055 | 10,055 | 10,055 |
| R-squared | 0.53 | 0.28 | 0.11 | 0.57 | 0.24 | 0.30 | 0.31 |
| Number of HHs | 783 | 783 | 783 | 783 | 783 | 783 | 783 |

Notes: This table reports on the effect of treatment on savings outcomes. Panel A reports results pooling across survey rounds 2-18. Panel B reports effects separately for the five indicated sets of survey rounds separately. Each column within a panel reports results of separate regressions using individual- and month-level fixed effects. Standard errors clustered at the zone (treatment unit) and survey wave level using the method of Cameron *et al.* (2011) are reported in parentheses. The unit of observation is the individual/month and the reported coefficients correspond to dummy variables equal to one for individuals offered weekly home visits from National Savings Bank deposit collectors in the months indicated. Columns (1)-(3) use institutional data from NSB to calculate cumulative monthly totals, including zeros for individuals without bank accounts. Columns (4)-(7) use survey data; (4) gives the cumulative number of financial transactions with any institution over the course of the month, while (5)-(7) give the total amount saved by the individual survey respondent over the previous month in each type of savings. Outcomes are Winsorized at the 99% level. Data include 4 months of pre-treatment data, 13 treated months, and two waves conducted 6 and 12 months after monthly home visits ended; the treatment changed at this point from door-to-door collection to collection from a lock box placed in the cluster of houses (see Section 2.2.1 for details). Panel A includes data through the end of weekly home visits (17 waves). Panel B includes the two post-treatment surveys (19 waves). Regressions pool samples surveyed monthly and samples surveyed quarterly. *significant at 10%; **significant at 5%; ***significant at 1%.

We estimate these effects using regressions of the form:

$$\begin{aligned}
 y_{izt} = & \phi_1 \cdot Treated_{z,1-4} + \phi_2 \cdot Treated_{z,5-8} + \phi_3 \cdot Treated_{z,9-13} \\
 & + \phi_4 \cdot Post_treated_{z,18} + \phi_5 \cdot Post_treated_{z,24} + \gamma_i + \eta_t + \varepsilon_{izt}
 \end{aligned} \quad (4.2)$$

Where variables and subscripts are defined similar to specification (4.1), except $Treated_{z,1-4}$ is a dummy variable equal to one for zones in the treatment group for months 1-4 of the post-treatment period (January-April 2011), $Treated_{z,5-8}$ is a dummy variable equal to one for zones in the treatment group for months 5-8 of the post-treatment period (May-August 2011), and so on. $Post_treated_{z,18}$ and $Post_treated_{z,24}$ are dummies for the treated group 6 and 12 months after the door-to-door treatment ended, and weekly home visits were replaced with weekly box pickup at the zone centre. These post-treatment coefficients measure the impact of weekly box pickup when individuals have first been habituated to save by a year of weekly home collections.

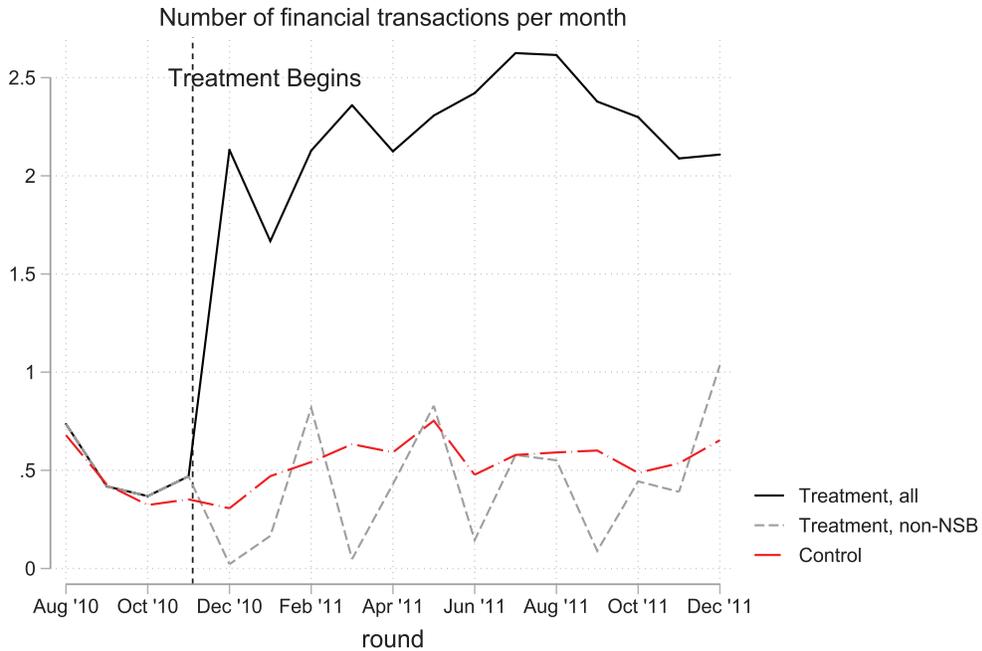


FIGURE 3
Treatment impacts on bank transactions

Examining the period-by-period impacts in Panel B of Table 2, we see that once the weekly home visits end and the zone switches to weekly box pickup we continue to observe significantly elevated financial transactions and NSB deposits relative to the control. The net flow of funds into NSB accounts reverses, however, and the treatment group eventually withdraws money from accounts at a higher rate than the control. Total savings flows drop back to the control level within 12 months, consistent either with households approaching a new steady state given a continued decrease in effective savings interest rates or with boxes having a more modest impact on total savings than deposit collection at the home.²²

4.3. *Headwaters*

Having verified that the experiment had the intended effect of increasing the aggregate savings of the participants, we can now ask the central question for which the experiment was designed: where does the increase in savings come from? The three immediate possibilities are: (i) a reduction in expenditures; (ii) an increase in earned income; (iii) and increase in net transfers from others inside or outside the household. Previous research gives some support for at least the first two possibilities. Banerjee *et al.* (2015) find that microfinance clients investing in microenterprises reduce expenditures on “temptation goods” and increase the purchase of durable goods. And both Dupas and Robinson (2013a) and Prina (2015) find that expenditures increase with the introduction of savings products that are similar to the one we examine. Both of these studies find incomes increased by an economically meaningful but statistically insignificant amount among those offered the savings accounts.²³ Our high-frequency data provide us with the

22. Section 5.1 provides additional evidence as to the impact of the collection boxes on household behaviour.

23. Dupas and Robinson show results on enterprise revenue rather than on profits.

statistical power to present a detailed answer to the “headwaters” question. We take a “sources and uses” accounting approach to analysing the effect of the POS treatment on outcomes, beginning by examining the sources and uses of funds at an aggregate household level. We estimate these impacts using a specification similar to equation (3.1) above, except with the cross-sectional unit, indexed by the subscript i , being the household rather than the individual. The first column of Table 3 shows the treatment effect on total household income from all sources, including earnings, remittances from domestic and foreign sources, and other flows (including loans, government program payments, and the like). Panel A shows the effect of treatment using the raw data, and Panel B shows the treatment effects using data winsorized at the 99th percentile. Treatment has a substantial positive effect, increasing total household income by roughly 10% of the mean in the control group; the measured effect is of similar magnitude using either the unwinsorized or winsorized data, though it is statistically significant only with the winsorized data.

Columns 2, 5, and 6 disaggregate the household income into three sources, which come directly from survey responses: earned income (Column 2), remittances from domestic and foreign sources (Column 5), and other household inflows, including loan receipts, *setu* payments, government program payments, and the like (Column 6). In the unwinsorized data, just over 90% of the increase in household income comes from earned sources. Winsorizing the data reduces the effect of earned income somewhat. Combined with a larger estimated treatment effect on total household income, this implies a larger role for other household inflows. However, even though the measured flow from other sources is higher in the winsorized data (Column 6, Panel B), the treatment effect on these flows remains insignificant.

Columns 3 and 4 split the increase in household earnings into changes in income reported by the respondent, and a residual, which we label changes in earnings from other household members.²⁴ Using either the unwinsorized or winsorized data, the change in the respondent’s income is about twice that of other household members. We return to a further analysis of the respondent’s income below. Understanding the source of these new (and fairly immediate) earnings is of key interest.

Columns 7–10 of Table 3 report the estimated treatment effect on uses of funds in the household. Reasonably, inflows could be consumed, saved, or transferred outside the household, and so we measure these dimensions of potential use. There are non-trivial measured effects on savings in both the unwinsorized and winsorized data, but the estimated savings effect is statistically significant only in the winsorized data. The unwinsorized data also show some evidence of an increase in loan repayments. In aggregate, the treatment effects on the sources and uses of funds balance reasonably well, with the changes in uses reflecting 70% of the increase in sources of funds in both the winsorized and unwinsorized data. One possible explanation for this “leakage” is that there are uses that are not categorized by respondents in any of these bins, as would be the case, for example, with investment in livestock or business inventories.

To present visual evidence of the distributional impact of the savings product, Figure 4 shows the un-winsorized cumulative density functions of the raw data for three of the main outcomes in Table 3, namely monthly bank savings, monthly total savings, and household income. The impacts on bank savings are clear throughout the distribution. For total savings, we see that the impacts are only visually positive above the median. For household income the impacts are less stark, and again the treatment effects are positive only above the median.

Panel C of Table 3 allows the treatment effect to vary across time by estimating the household-level analogue of specification (4.2) above. The save-invest-earn channel implies that income should increase little if at all initially, with larger increases occurring later after treatment. The labour reallocation effect should have the opposite pattern, with income rising immediately as

24. The baseline survey included a household roster for which the respondent reported the occupation, hours worked, and income for other household members.

TABLE 3
Examining the headwaters of formal savings

| Variables | Sources of income | | | | | Uses of income | | | | |
|---|------------------------|--------------------------------|-----------------------------|-----------------------------------|---|--|-------------------------|---------------------|--------------------------------|---------------------|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) |
| | Earned income | | | Transfers and loans | | Consumption and savings | | Transfers and loans | | |
| | Household income | Earned by household respondent | Earned by household members | Earned by other household members | Transfers from outside household (remittance) | Other household inflows (e.g. loan receipts) | consumption consumption | Total savings | Transfers to outside household | Loan payments |
| Panel A: Full sample, non-winsorized data (N = 783) | | | | | | | | | | |
| Treated months 1-13 | 2,634.9 (1842.6) | 2,372.0* (1270.0) | 1,646.9* (977.3) | 725.1 (894.2) | 103.0 (396.1) | 158.7 (1,315.7) | 279.2 (1,484.7) | 865.3 (650.1) | -92.5 (129.7) | 804.7* (425.5) |
| Observations | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 |
| | 0.241 | 0.271 | 0.302 | 0.295 | 0.180 | 0.157 | 0.167 | 0.223 | 0.072 | 0.142 |
| Panel B: Full sample, data winsorized at the 99th percentile (N = 783) | | | | | | | | | | |
| Treated months 1-13 | 3185.3** (1,240.6) | 1853.5** (768.3) | 989.0 (637.6) | 392.3 (578.9) | 68.6 (147.4) | 858.2 (914.4) | 899.6 (950.4) | 883.0** (376.7) | 3.31 (51.6) | 362.5 (232.7) |
| Observations | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 |
| R-squared | 0.360 | 0.425 | 0.378 | 0.518 | 0.243 | 0.202 | 0.290 | 0.325 | 0.255 | 0.198 |
| Panel C: Time-varying effects, data winsorized at the 99th percentile (N = 783) | | | | | | | | | | |
| Treated months 1-4 | 3,465.8* (1,810.9) | 2,018.8** (948.4) | 1,436.7* (753.1) | 49.3 (611.4) | 93.4 (159.3) | 1,025.9 (1,101.2) | 1,005.0 (1,177.7) | 964.3** (428.1) | -16.5 (49.8) | 558.1* (308.1) |
| Treated months 5-8 | 3270.6** (1,291.2) | 1823.4** (912.3) | 541.9 (868.7) | 800.6 (798.9) | 72.2 (212.1) | 930.2 (1,089.6) | 802.7 (1,049.5) | 634.8 (489.9) | 11.2 (75.4) | 277.8 (343.2) |
| Treated months 9-13 | 2,495.1 (1,700.3) | 1,493.8 (1,175.1) | 862.6 (821.4) | 279.3 (850.1) | 11.9 (234.7) | 553.6 (1,168.0) | 811.3 (1,182.7) | 1036.8** (469.5) | 22.8 (71.0) | 188.5 |
| 6 months post-treatment | 3,038.2* (1,657.2) | 1,854.9 (1,413.5) | 1,190.0 (1,273.1) | 743.4 (975.8) | 57.2 (253.7) | 606.4 (989.1) | 1,583.0 (1,151.1) | 926.2* (538.5) | -131.9* (72.0) | 735.8** (313.3) |
| Control Group mean | 29,393.04 (1,661.8) | 22,022.68 (1,139.0) | 9,879.11 (1,392.4) | 11,786.89 (1,151.7) | 611.10 (230.6) | 5,979.01 (1,102.6) | 16,316.90 (1,043.0) | 5,213.32 (480.3) | 353.37 (50.6) | 1,559.02 (282.1) |
| Observations | 10,055 | 10,055 | 10,055 | 10,055 | 10,096 | 10,096 | 10,055 | 10,055 | 10,055 | 10,055 |
| R-squared | 0.355 | 0.415 | 0.365 | 0.513 | 0.221 | 0.196 | 0.284 | 0.312 | 0.243 | 0.191 |

Notes: This table reports on the effect of treatment on financial outcomes. Panel A reports results pooling across survey rounds 2-18, with no Winsorizing of the outcome variables. Panel B reports effects for survey rounds 2-18, with the outcomes Winsorized at the 99th percentile. Panel C uses data Winsorized at 99% and presents estimates separately for the five indicated sets of survey rounds. Each column within a panel reports results of separate regressions using individual- and month-level fixed effects. Standard errors clustered at the zone (treatment unit) and survey wave level using the method of Cameron *et al.* (2011) are reported in parentheses. The unit of observation is the individual/month and the reported coefficients correspond to dummy variables equal to one for individuals offered weekly home visits from National Savings Bank deposit collectors in the months indicated. All columns in the table use survey data, and outcomes are monthly flows for the variables indicated. Data include 4 months of pre-treatment data, 13 treated months, and two waves conducted 6 and 12 months after monthly home visits ended; the treatment changed at this point from door-to-door collection to collection from a lock box placed in the cluster of houses (see Section 2.2.1 for details). Panels A and B include data through the end of weekly home visits (17 waves). Panel C includes the two post-treatment surveys (19 waves). Regressions pool samples surveyed monthly and samples surveyed quarterly. * significant at 10%; ** significant at 5%; *** significant at 1%.

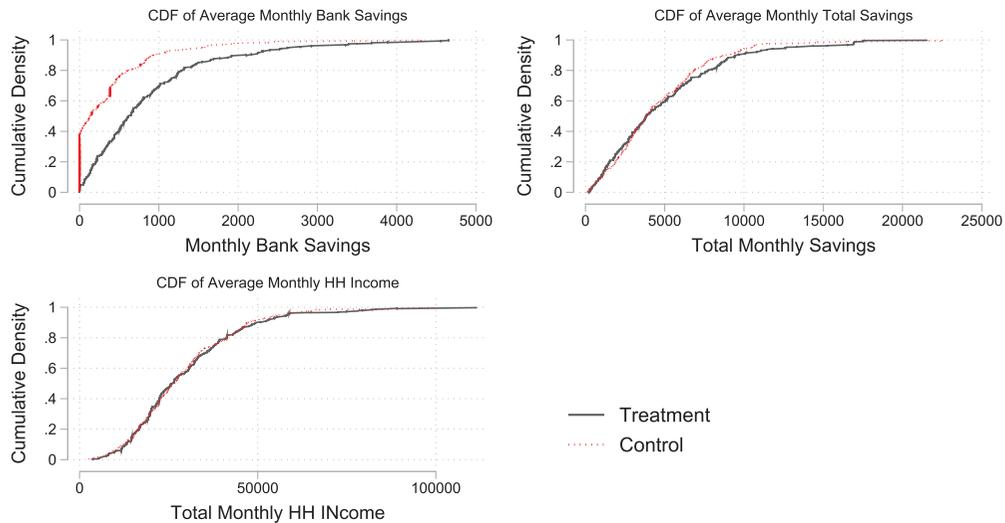


FIGURE 4
Cumulative density functions across treatment and control

households increase effort today. The data show a pattern more consistent with the labour-adjustment story, with very large immediate increases that do not grow over time, and may diminish by the last survey round. Figure 5 shows that (along with a large shock caused by widespread flooding), a treatment effect on business income is visible from the very first month of treatment. While impacts largely attenuate once weekly home visits stop, this pattern of treatment effects across time is the first piece of evidence in favour of labour reallocation as a mechanism linking new savings opportunities to increases in income.

Given that the savings increases we observe appear to be primarily driven by new income, we turn in Table 4 to examining the individual-level income data in more detail. We perform this analysis estimating regressions similar to specification (4.1) above, replacing the dependent variables with measures of the sources of individual-level income. We shift our focus to the respondent's behaviour because that is where we have the most detailed income data. In addition, the estimates in Table 3 indicate that the majority of the increase in earned income is due to the respondent.

Panel A of Table 4 shows the disaggregation of income for the full sample, with all income data winsorized at the 99th percentile. We begin by examining the effect of treatment on occupational choice. We find no evidence that the savings accounts operate by relaxing credit constraints, and, thereby, increasing entry into self-employment. On the contrary, we find that self-employment falls by 3.4 percentage points in the treatment group, relative to the control group. Engagement in wage work increases by two percentage points, though the effect is not statistically significant (p -value = 0.16). However, earnings from wage work (Column 4) and average hours in wage work (Column 7) both increase with the increase in wage hours being statistically significant. Note that all measures capture both intensive and extensive effects. Self-employment status, income, and hours are coded as zeros for those not self-employed, and hence the treatment effects are impacted by the entry into wage work and exit from self-employment reflected in Columns 1 and 3. Supplementary Appendix Table A7 uses Randomization Inference to calculate standard errors for Panel A of Table 4 and concurs with the two-way clustering in terms of the levels of significance of impacts on self-employment and wage hours.

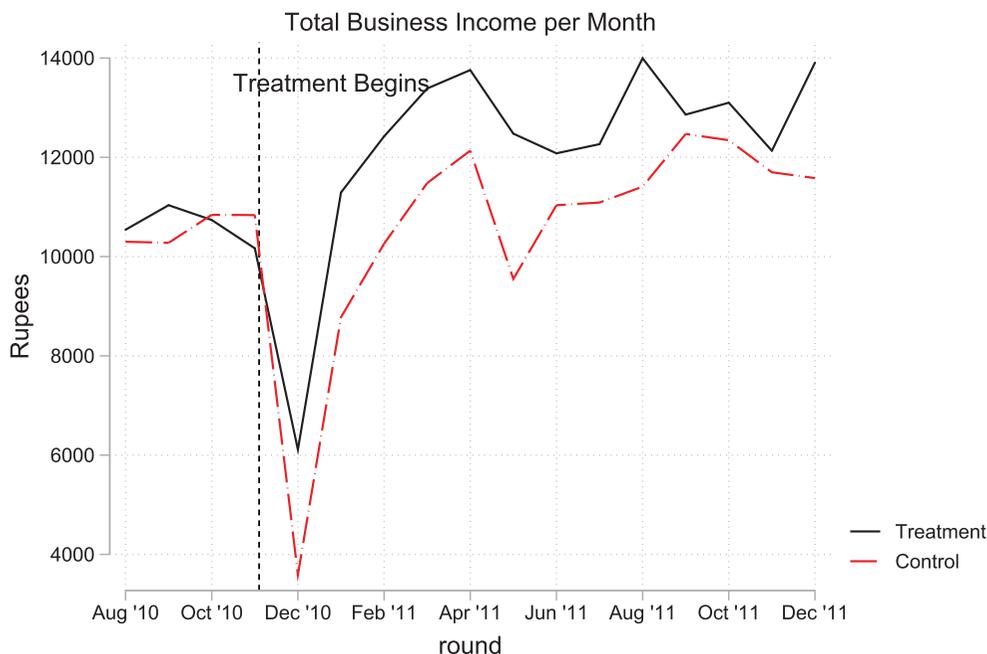


FIGURE 5
Impact on business income

We can also calculate an hourly treatment effect by comparing the increase in hours with the increase in earnings. Hours worked is reported on a weekly basis, and the income on a monthly basis. Thus, the increase of 1.5 h/week in wage work is thus approximately 6.4 h/month. This implies a wage at the margin of 49.4 LKR per hour (315 LKR / 6.4 h), an amount very similar to the average hourly wage of 50.5 LKR per hour in the control group. This is reassuring with regard to the consistency of the data.

In sum, we find that earned income increases immediately after treatment, and the magnitude of the treatment effect is either steady or decreasing across time. Treatment also leads to a shift out of self-employment and into wage work. These results are not consistent with the save-invest-earn channel of savings. They are consistent with an increase in labour effort when the deposit collections are introduced. To explore this channel a bit further, we step outside the experiment, and examine effects dividing the full sample into three subgroups.²⁵ We first split the sample into those who are self-employed and those who are not. We do this using the activities reported in the four baseline rounds, which took place before treatment status was known to the participants. This initial split is justified by the fact that the save-invest-earn channel can work through either the intensive or extensive (*e.g.* lower exit rates) channel among those self-employed at baseline, but only through the extensive margin among those not self-employed. There are 419 respondents active in self-employment in at least one baseline round, and 360 not active in self-employment.

25. The project started before the American Economic Association (AEA) Randomized Control Trial (RCT) registry was established. Hence, we have no public randomization plan. However, the stratification indicates that our initial analysis was built around the informal-formal savings channel. The data show little evidence for this channel in aggregate, though we return to this issue below. Though we balanced on expenditures, indicating that we saw this as a key outcome, we did not expect to find an immediate increase in income and a shift out of self-employment. The three subsamples we construct are, hence, non-experimental, and the analysis should be interpreted with this in mind.

TABLE 4
Estimated impacts of deposit collection on earnings outcomes

| Dependent variable: | (1) Y/N: wage employment? | (2) Y/N: Ag employment? | (3) Y/N: Self-employment? | (4) Wages earned | (5) Ag income earned | (6) Business income | (7) Wage hours | (8) SE hours |
|--|------------------------------|----------------------------|------------------------------|---------------------|-------------------------|------------------------|-------------------|-----------------|
| Panel A Full sample ($N=783$) | | | | | | | | |
| Treated months 1-13 | 0.020 (0.016) | -0.0000048 (0.014) | -0.034* (0.017) | 314.9 (212.7) | -15.4 (503.2) | 573.5 (469.0) | 1.52** (0.71) | -0.23 (1.00) |
| Control group mean | 0.16 | 0.59 | 0.53 | 1,263.89 | 2,566.01 | 5,834.98 | 6.00 | 24.84 |
| Observations | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 |
| R-squared | 0.86 | 0.92 | 0.92 | 0.76 | 0.21 | 0.68 | 0.83 | 0.85 |
| Panel B: Some self-employment and some manufacturing during pre-treatment period ($N=259$) | | | | | | | | |
| Treated months 1-13 | -0.019 (0.033) | -0.016 (0.026) | -0.0092 (0.027) | -204.1 (352.1) | -210.6 (759.4) | 2,808.5*** (1075.7) | -0.037 (1.38) | 3.02 (1.97) |
| Control group mean | 0.09 | 0.67 | 0.94 | 796.01 | 1,611.97 | 9,561.80 | 3.88 | 42.10 |
| Observations | 3,085 | 3,085 | 3,085 | 3,085 | 3,085 | 3,085 | 3,085 | 3,085 |
| R-squared | 0.80 | 0.89 | 0.65 | 0.62 | 0.13 | 0.56 | 0.78 | 0.68 |
| Panel C: Some self-employment but no manufacturing during pre-treatment period ($N=160$) | | | | | | | | |
| Treated months 1-13 | 0.10*** (0.032) | 0.030 (0.030) | -0.078** (0.037) | 1014.6** (420.8) | 293.9 (192.8) | -1141.3 (1164.9) | 4.40*** (1.54) | -3.07 (2.73) |
| Control group mean | 0.04 | 0.02 | 0.99 | 258.47 | -56.39 | 12,535.26 | 1.38 | 51.87 |
| Observations | 1,821 | 1,821 | 1,821 | 1,821 | 1,821 | 1,821 | 1,821 | 1,821 |
| R-squared | 0.74 | 0.63 | 0.51 | 0.66 | 0.11 | 0.61 | 0.71 | 0.63 |
| Panel D: No self-employment and no manufacturing during pre-treatment period ($N=364$) | | | | | | | | |
| Treated months 1-13 | 0.012 (0.020) | -0.0077 (0.021) | -0.026 (0.022) | 339.7 (241.4) | -257.0 (762.2) | -228.2 (342.6) | 1.30 (0.80) | -0.77 (0.79) |
| Control group mean | 0.25 | 0.72 | 0.05 | 1,988.71 | 4,251.10 | 535.13 | 9.30 | 1.74 |
| Observations | 4,266 | 4,266 | 4,266 | 4,266 | 4,266 | 4,266 | 4,266 | 4,266 |
| R-squared | 0.89 | 0.90 | 0.59 | 0.81 | 0.24 | 0.53 | 0.87 | 0.49 |

Notes: This table reports on the effect of treatment on employment outcomes; columns 4-9 are Winsorized at the 99th percentile. All regressions use data across survey rounds 2-18. Panel A uses the full sample of respondents in the study, Panel B uses all respondents who ever reported being self-employed in a manufacturing business in the five baseline rounds, Panel C uses all respondents who ever reported being self-employed in the five baseline rounds but not in a manufacturing business, and Panel D uses all respondents who never reported being self-employed in the five baseline rounds. Each column within a panel reports results of separate regressions using individual- and month-level fixed effects. Standard errors clustered at the zone (treatment unit) and survey wave level using the method of Cameron *et al.* (2011) are reported in parentheses. The unit of observation is the individual/month and the reported coefficients correspond to dummy variables equal to one for individuals offered weekly home visits from National Savings Bank deposit collectors in the months indicated. All regressions use survey data; columns (1)-(3) report whether an individual reported being employed in the described activities at any point during a month, and Columns (4)-(8) report monthly totals. Data include 4 months of pre-treatment data and 13 treated months, through the end of weekly home visits (17 waves). Regressions pool samples surveyed monthly and samples surveyed quarterly. * significant at 10%; ** significant at 5%; *** significant at 1%.

We then split the self-employment sample further into those active in manufacturing during baseline, and those active only in trade and services. Given the very rural nature of our sample, we expect those in trade and services may have less opportunity to increase demand, while those in manufacturing likely face fewer constraints because they can sell output in nearby towns.²⁶

Panels B-D of Table 4 report the treatment effects on activities of participants in the self-employed manufacturing, self-employed non-manufacturing, and not self-employed samples. Among those active in self-employed manufacturing in baseline (Panel B), we see no shift toward wage work, but instead a very large increase in reported business earnings. This is accompanied by

26. As it turns out, those active in manufacturing are also much more likely to report income from more than one activity during the baseline period. Two-thirds of those in manufacturing are also active either in wage labour or agriculture; this compares to less than 5% of those in trade and services. The differences in the subsamples may be due to these multiple activities or some other characteristic correlated with involvement in manufacturing.

a modest (and statistically insignificant) increase in hours worked. Together, these imply marginal earnings of 219 LKR per additional hour, about three times the sample mean of 72 LKR per hour among all self-employed. These data are thus consistent with an increase in both work intensity and work hours, as should be expected if the effective returns to labour increase.²⁷ Researchers have struggled to find direct measures of effort in self-employment activities, and we lack good measures, as well. So, we can only say that the reported increase in earnings is high relative to the reported increase in hours worked, and we speculate that this may reflect an increase in the intensity of effort.

In contrast, among those who are self-employed in trade and services (Panel C), we see a sharp exit effect (column 3), accompanied by an increase in wage work (column 1), in wage earnings (column 4), and in the hours worked in wage work (column 7). Again, the incremental earnings in wage work per incremental hour worked—54.9 LKR—is close to the sample average. Now, however, the same calculation on self-employment—measured as a decrease in earnings per decreased hour worked—is also much closer to the sample mean among the self-employed, at 88.2 LKR per hour. Note that this decrease is a LATE of those remaining in self-employment in the control group who would have exited with treatment. Thus, we should not expect to find any effect of effort, since the self-employed over whom this effect is measured are in the control group. While the trade and services self-employed are induced to make large movements to wage work, these movements are ultimately ineffective in increasing income: wage earnings rise, but business earnings fall by an almost identical amount. Finally, we see no effects on either activities or income earned among those not self-employed at baseline (Panel D). Supplementary Appendix Table A8 shows the main savings results disaggregated by the same subgroups as Table 4.

We now move to a period-by-period analysis of subgroup impacts. Just as Table 3 showed immediate impacts of the treatment on income, Table 5 demonstrates that an exit from self-employment and an increase in wage hours take place even in the first few months of treatment. Panels B and C show the time-varying effects for the two self-employment subsamples. Among the self-employed in manufacturing (Panel B), we see the immediate jump in business earnings, which fades very slightly over time until the last survey round, when it falls markedly and becomes statistically insignificant. In Panel C, we find that the exit from self-employment among those initially in trade and services increases throughout the post-treatment period. The treatment effect on reporting self-employment activity is -5.4% and significant in the initial post-treatment period. But this grows so that by the final two survey rounds, fully one in six of those in the treatment group who were initially self-employed has exited self-employment. Both reporting any wage earnings (Column 1) and the level of wage earnings (Column 4) increase steadily across time, but business earnings decrease by a comparable amount in the first 13 months, and by substantially more than the wage increase in the last two surveys.

It is reasonable to ask whether the differences in the earnings response across these three subsamples are mirrored in differences in savings behaviour. We examine this question in Table 6, where we return to the sources-and-uses analysis, this time by subsample. Panel A repeats the results for the full sample (with outcomes winsorized) in Table 3, again using data from survey round 2 (September 2010) to survey round 18 (January 2012). Panels B–D then show the self-employed manufacturing, self-employed trade and services, and not-self-employed subsamples. We see that there is a significant increase in savings (column 8) and loan repayments (column 10) only among the self-employed in manufacturing, where individuals were successful in increasing earnings. In this subsample, there is a small and insignificant change in consumption (column 7). Among the self-employed in trade and services, the only category showing a significant

27. See, for example, the estimated intensity effects found by Lazear (2000) when workers switched from hourly wages to piece rates.

TABLE 5
Estimated impacts of deposit collection on earnings outcomes by time period

| | (1) Y/N: Wage employment? | (2) Y/N: Ag employment? | (3) Y/N: Self- employment? | (4) Wages earned | (5) Ag income earned | (6) Business income | (7) Wage hours | (8) SE hours |
|--|---------------------------------|-------------------------------|----------------------------------|------------------------|----------------------------|---------------------------|----------------------|--------------------|
| Panel A: Full sample ($N=783$) | | | | | | | | |
| Treated months 1-4 | 0.013 (0.013) | -0.011 (0.012) | -0.034** (0.014) | 19.9 (194.0) | 449.9 (660.3) | 825.7* (457.9) | 1.14** (0.54) | -0.13 (0.91) |
| Treated months 5-8 | 0.024 (0.019) | 0.013 (0.018) | -0.028 (0.020) | 446.1* (246.8) | -528.6 (573.2) | 578.8 (524.7) | 1.44* (0.85) | -0.058 (1.08) |
| Treated months 9-13 | 0.027 (0.025) | -0.0016 (0.021) | -0.040 (0.026) | 583.8* (300.6) | -101.5 (582.5) | 210.3 (630.7) | 2.21** (1.10) | -0.63 (1.65) |
| 6 months post-treatment | 0.031 (0.027) | -0.012 (0.023) | -0.045 (0.032) | 884.0** (381.5) | 465.7 (925.8) | -473.4 (761.1) | 1.72 (1.11) | -4.05** (1.90) |
| 12 months post-treatment | 0.055* (0.028) | -0.0081 (0.024) | -0.028 (0.033) | 1165.9** (506.8) | -283.3 (1,272.7) | -1692.9*** (627.8) | 2.95** (1.38) | -1.17 (1.78) |
| Control group mean | 0.16 | 0.59 | 0.53 | 1,263.89 | 2,566.01 | 5,834.98 | 6.00 | 24.84 |
| Observations | 10,055 | 10,055 | 10,055 | 10,055 | 10,055 | 10,055 | 10,055 | 10,055 |
| R-squared | 0.852 | 0.909 | 0.914 | 0.743 | 0.203 | 0.667 | 0.817 | 0.839 |
| Panel B: Some self-employment and some manufacturing during baseline ($N=259$) | | | | | | | | |
| Treated months 1-4 | 0.0096 (0.030) | -0.042* (0.024) | -0.034 (0.022) | -298.6 (371.5) | -241.4 (1096.8) | 3106.1*** (902.0) | 0.69 (1.35) | 1.41 (1.96) |
| Treated months 5-8 | -0.035 (0.037) | 0.0053 (0.026) | 0.010 (0.028) | -185.7 (390.5) | -266.7 (732.7) | 2348.4** (1029.4) | -1.03 (1.51) | 4.37** (1.76) |
| Treated months 9-13 | -0.041 (0.038) | -0.0087 (0.037) | -0.00052 (0.039) | -79.9 (449.1) | -154.8 (1,339.0) | 2,867.6* (1679.3) | 0.15 (1.73) | 3.46 (2.95) |
| 6 months post-treatment | -0.047 (0.044) | 0.0066 (0.044) | -0.050 (0.048) | 309.8 (663.1) | -1368.8 (1,735.4) | 2,747.3** (1252.5) | -0.71 (2.07) | -1.01 (2.68) |
| 12 months post-treatment | -0.050 (0.043) | 0.019 (0.045) | -0.012 (0.049) | 56.3 (664.9) | 1,291.0 (3,207.2) | -1572.3 (1194.6) | -1.21 (1.99) | 0.49 (2.52) |
| Control group mean | 0.09 | 0.67 | 0.94 | 796.01 | 1,611.97 | 9,561.80 | 3.88 | 42.10 |
| Observations | 3,388 | 3,388 | 3,388 | 3,388 | 3,388 | 3,388 | 3,388 | 3,388 |
| R-squared | 0.790 | 0.881 | 0.634 | 0.587 | 0.130 | 0.542 | 0.755 | 0.664 |
| Panel C: Some self-employment but no manufacturing during baseline ($N=160$) | | | | | | | | |
| Treated months 1-4 | 0.063** (0.025) | 0.015 (0.020) | -0.054* (0.029) | 596.7* (327.7) | 495.7 (398.2) | -660.9 (1,427.5) | 2.66** (1.25) | -0.89 (3.13) |
| Treated months 5-8 | 0.12*** (0.032) | 0.045 (0.037) | -0.079** (0.038) | 1274.1*** (483.6) | 129.7 (198.5) | -899.8 (1289.2) | 5.29*** (1.55) | -3.52 (3.15) |
| Treated months 9-13 | 0.14*** (0.042) | 0.032 (0.042) | -0.11* (0.060) | 1382.3** (570.4) | 193.3 (180.0) | -2130.3 (1457.4) | 6.08*** (2.08) | -5.83 (4.12) |
| 6 months post-treatment | 0.15*** (0.044) | -0.016 (0.031) | -0.16** (0.065) | 1785.6*** (623.1) | -7.78 (135.0) | -6772.4*** (1,832.0) | 6.11*** (2.11) | -17.2*** (4.14) |
| 12 months post-treatment | 0.17*** (0.047) | -0.042 (0.034) | -0.16** (0.067) | 2,903.1*** (834.6) | 720.4 (449.9) | -6,098.3** (2,444.7) | 8.35*** (2.31) | -9.75** (4.75) |
| Control group mean | 0.04 | 0.02 | 0.99 | 258.47 | -56.39 | 12,535.26 | 1.38 | 51.87 |
| Observations | 1,979 | 1,979 | 1,979 | 1,979 | 1,979 | 1,979 | 1,979 | 1,979 |
| R-squared | 0.756 | 0.637 | 0.529 | 0.672 | 0.106 | 0.616 | 0.718 | 0.623 |

Notes: This table reports on the effect of treatment on employment outcomes; columns 4-9 are Winsorized at the 99th percentile. All regressions use data across survey rounds 2-20, which include four pre-treatment rounds, 13 treated months, and post-treatment rounds conducted 6 and 12 months after the weekly home visits ended. Panel A uses the full sample of respondents in the study, Panel B uses all respondents who ever reported being self-employed in a manufacturing business in the five baseline rounds, Panel C uses all respondents who ever reported being self-employed in the five baseline rounds but not in a manufacturing business, and Panel D uses all respondents who never reported being self-employed in the five baseline rounds. Each column within a panel reports results of separate regressions using individual- and month-level fixed effects. Standard errors clustered at the zone (treatment unit) and survey wave level using the method of Cameron *et al.* (2011) are reported in parentheses. The unit of observation is the individual/month and the reported coefficients correspond to dummy variables equal to one for individuals offered weekly home visits from National Savings Bank deposit collectors in the months indicated. All regressions use survey data; columns (1)-(3) report whether an individual reported being employed in the described activities at any point during a month, and Columns (4)-(8) report monthly totals. Regressions pool samples surveyed monthly and samples surveyed quarterly. * significant at 10%; ** significant at 5%; *** significant at 1%.

TABLE 5
(Continued)

| | (1) Y/N: Wage employment? | (2) Y/N: Ag employment? | (3) Y/N: Self- employment? | (4) Wages earned | (5) Ag income earned | (6) Business income | (7) Wage hours | (8) SE hours |
|--|---------------------------------|-------------------------------|----------------------------------|------------------------|----------------------------|---------------------------|----------------------|--------------------|
| Panel D: No self-employment and no manufacturing during baseline (<i>N</i> = 364) | | | | | | | | |
| Treated months 1-4 | -0.0089 (0.015) | -0.0015 (0.012) | -0.020 (0.019) | -37.7 (235.4) | 562.1 (939.8) | -156.4 (180.7) | 0.70 (0.50) | -0.53* (0.28) |
| Treated months 5-8 | 0.022 (0.024) | -0.0035 (0.030) | -0.027 (0.027) | 491.8* (297.0) | -1,152.6 (1,113.8) | -131.9 (515.5) | 1.42 (1.14) | -1.01 (0.97) |
| Treated months 9-13 | 0.028 (0.035) | -0.022 (0.030) | -0.033 (0.036) | 665.7* (390.6) | -367.9 (847.2) | -442.8 (428.5) | 1.99 (1.46) | -0.86 (1.62) |
| 6 months post-treatment | 0.038 (0.033) | -0.036 (0.035) | 0.0071 (0.039) | 958.1** (409.1) | 1,692.0 (1,216.7) | -392.5 (949.6) | 1.76* (1.05) | -0.96 (2.19) |
| 12 months post-treatment | 0.079** (0.036) | -0.027 (0.032) | 0.016 (0.037) | 1,236.1** (601.9) | -2523.4** (1,214.1) | -328.5 (558.9) | 3.53** (1.67) | 1.02 (1.60) |
| Control group mean | 0.25 | 0.72 | 0.05 | 1,988.71 | 4,251.10 | 535.13 | 9.30 | 1.74 |
| Observations | 4,688 | 4,688 | 4,688 | 4,688 | 4,688 | 4,688 | 4,688 | 4,688 |
| R-squared | 0.88 | 0.89 | 0.60 | 0.80 | 0.24 | 0.50 | 0.85 | 0.49 |

treatment effect is transfers outside the household; this category shows a decrease that, while significant at the 0.10 level, is relatively small in magnitude. Among those not self-employed, there is a marginally significant increase in other household inflows (column 6) and non-trivial, though highly insignificant, increases in savings and consumption. Thus, we find a significant increase in savings only where there is a significant increase in income, that is, among those who were self-employed in manufacturing at baseline. We see strong labour market effects among those self-employed in non-manufacturing activities, but these shifts generate neither increases in income nor increases in savings.

Survey-based data on financial behaviour are always reported with some noise.²⁸ To assess the robustness of our results, we repeat all of the core results without winsorizing, with winsorizing at the 95th percentile of the distribution, and using inverse hyperbolic sine transformations. These are shown in Supplementary Appendix Tables B1–B10.

5. WHY DO DEPOSIT COLLECTIONS INCREASE SAVINGS AND LABOUR EFFORT?

This section discusses and tests empirically, to the extent our data allow, several mechanisms that might link deposit collections to increased labour effort. Empirically, in response to deposit collection, weekly labour hours in wage work and self-employment increase by 1.29 h or 4.2% of the control group mean of 30.84 h (Table 4, Panel A, columns 7 and 8). In the non-manufacturing self-employed subsample, wage work and self-employment hours worked weekly increase by 1.33 h, a 2.5% increase over the control group mean of 53.25 h.

Are the increases we observe in labour hours reasonable given the interest paid on the account? One approach to answering this question is to recall that, in the simple model presented in Section 3 above, the elasticity of the labour ratio $\frac{h_t}{h_{t+1}}$ with respect gross interest rate $(1+r_{t+1})$ is just $\frac{1}{\gamma-1}$. Using a 5 percentage point increase in the real gross interest rate—that is, assuming that the

28. For example, we have one respondent who indicates receiving 1.65 m LKR as his monthly household income. At 16 times the 99th percentile, this very likely reflects measurement error, and is extreme enough that, even after taking the inverse hyperbolic sine transformation of this variable, it remains an outlier. Rather than selectively removing observations like this one, which very plausibly reflect recording error, we prefer winsorization.

TABLE 6
Estimated impacts of deposit collection on earnings outcomes by employment subgroup

| | Sources of income | | | | | Uses of income | | | | |
|---|-------------------------|-------------------------|--------------------------|--|---|--|-----------------------|-----------------------|--------------------------------|--------------------|
| | Household Income | Household earned income | Respondent earned income | Earned income: other household members | Transfers from outside household (remittance) | Other household inflows (e.g. loan receipts) | Household consumption | Total savings | Transfers to outside household | Loan payments |
| Panel A Full sample (N= 783) | | | | | | | | | | |
| Treated months 1-13 | 3,185.3** (1,240.6) | 1,853.5** (768.3) | 989.0 (637.6) | 392.3 (578.9) | 68.6 (147.4) | 858.2 (914.4) | 899.6 (950.4) | 883.0** (376.7) | 3.31 (51.6) | 362.5 (232.7) |
| Control mean | 29,307.7 | 21,566.3 | 9,789.4 | 11,460.3 | 633.3 | 6,290.6 | 16,229.3 | 5,280.4 | 368.0 | 1,597.0 |
| Observations | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 |
| R-squared | 0.36 | 0.43 | 0.38 | 0.52 | 0.24 | 0.20 | 0.29 | 0.33 | 0.26 | 0.20 |
| Panel B: Some self-employment and some manufacturing during baseline (N= 259) | | | | | | | | | | |
| Treated months 1-13 | 6,036.1*** (2,000.4) | 4,570.6*** (1,335.7) | 2,522.6** (1,138.3) | 1,722.4 (1,052.3) | 131.7 (303.7) | -0.25 (1,330.2) | 509.7 (1,447.1) | 1,901.3*** (559.3) | 46.1 (79.1) | 725.2** (283.6) |
| Control mean | 29,298.9 | 22,085.0 | 12,083.5 | 9,796.5 | 522.9 | 6,129.2 | 16,216.8 | 5,390.0 | 367.9 | 1,756.4 |
| Observations | 3,085 | 3,085 | 3,085 | 3,085 | 3,085 | 3,085 | 3,085 | 3,085 | 3,085 | 3,085 |
| R-squared | 0.40 | 0.44 | 0.35 | 0.48 | 0.24 | 0.20 | 0.29 | 0.35 | 0.26 | 0.16 |
| Panel C: Some self-employment but no manufacturing during baseline (N = 160) | | | | | | | | | | |
| Treated months 1-13 | 487.8 (3,087.0) | 912.8 (1,825.3) | 227.7 (1,184.8) | 258.1 (1,308.3) | -123.4 (314.6) | -179.2 (2,292.7) | 184.3 (1,831.8) | -148.6 (722.2) | -210.1 (127.8) | 387.8 (538.2) |
| Control mean | 28,631.3 | 21,820.7 | 12,294.0 | 9,196.3 | 470.4 | 5,277.1 | 16,933.1 | 5,008.4 | 296.3 | 1,095.1 |
| Observations | 1,821 | 1,821 | 1,821 | 1,821 | 1,821 | 1,821 | 1,821 | 1,821 | 1,821 | 1,821 |
| R-squared | 0.39 | 0.57 | 0.57 | 0.63 | 0.21 | 0.17 | 0.32 | 0.36 | 0.29 | 0.24 |
| Panel D: No self-employment during baseline (N= 364) | | | | | | | | | | |
| Treated months 1-13 | 2,239.5 (2,110.2) | 250.3 (1,255.8) | -4.65 (774.9) | -289.1 (1,032.6) | 107.3 (259.9) | 2,008.7* (1,216.3) | 1,550.8 (1,241.3) | 582.5 (838.3) | 69.2 (73.6) | 63.1 (427.3) |
| Control mean | 29,561.5 | 21,074.2 | 7,109.2 | 13,567.7 | 777.7 | 6,784.9 | 15,981.9 | 5,295.3 | 394.2 | 1,657.6 |
| Observations | 4,266 | 4,266 | 4,266 | 4,266 | 4,266 | 4,266 | 4,266 | 4,266 | 4,266 | 4,266 |
| R-squared | 0.33 | 0.36 | 0.31 | 0.48 | 0.26 | 0.22 | 0.28 | 0.30 | 0.24 | 0.21 |

Notes: This table reports on the effect of treatment on financial outcomes. All regressions use data across survey rounds 2 - 18, and outcomes are Winsorized at the 99% level. Panel A uses the full sample of respondents in the study, Panel B uses all respondents who ever reported being self-employed in a manufacturing business in the five baseline rounds, Panel C uses all respondents who ever reported being self-employed in the five baseline rounds but not in a manufacturing business, and Panel D uses all respondents who never reported being self-employed in the five baseline rounds. Each column within a panel reports results of separate regressions using individual- and month-level fixed effects. Standard errors clustered at the zone (treatment unit) and survey wave level using the method of Cameron *et al.* (2011) are reported in parentheses. The unit of observation is the individual/month and the reported coefficients correspond to dummy variables equal to one for individuals offered weekly home visits from National Savings Bank deposit collectors in the months indicated. All columns in the table use survey data, and outcomes are monthly flows for the variables indicated. Data include 4 months of pre-treatment data and 13 treated months. Regressions pool samples surveyed monthly and samples surveyed quarterly. * significant at 10%; ** significant at 5%; *** significant at 1%.

pre-intervention nominal interest rate was zero—and the sample-average increase in labour hours of 4.2% in the full sample, we find a value of γ equal to about 2.16. Using the 2.5% increase in the manufacturing subsample yields a γ of 2.95. Both of these are close to the quadratic disutility commonly assumed (often for convenience) in standard models and within the range of estimates for γ found in the literature (Beaudry and Van Wincoop, 1996, find an IES of 1, equivalent to $\gamma = 2$, and Augenblick *et al.*, 2015, estimate $\gamma = 1.6$ in an experiment over real effort tasks) suggesting that the increase in labour hours is reasonable given a 5 percentage point increase in real interest rates.

We can also address the importance of the income channel using a direct survey question we asked in the final monthly wave. In that survey, we asked participants in the treatment group several questions about how the savings treatment had affected their behaviour. Their responses are generally consistent with the outcomes and channels we identify in the data. Of those opening accounts, 78% reported that the intervention helped them to increase savings, and only 5% said they felt pressured to save by the door-to-door collection. Among those who said they had saved more, 38% said this was achieved primarily through increased income, 34% said this was achieved through decreased consumption, and 22% said this was achieved through changes in intra-household transfers. By far the most common source of increased income was said to be “more effort to increase sales” in self-employment or agriculture. This was reported by 50% of those saying that income had increased, with 13% reporting doing more work at a previously engaged wage-paying job, 5% reporting doing new wage work, and 18% reporting increased income from re-investing in self-employment or agriculture. Although we find some evidence of increased investment in enterprises in the final survey rounds, these self-reported narratives reinforce the idea that our intervention increased savings by raising incomes, principally through more labour effort.

While both the calibration and the survey evidence suggest that our results are broadly sensible relative to the neoclassical literature on intertemporal labour substitution, it is of course possible that labour effort and income increase for other reasons, or that other mechanisms amplify the effect of interest rates. We explore some of the other channels in this section.

In addition to the 5% interest rate, the deposit collection process has several features that could induce an increase in savings. First, in-person visits by bank staff might remind participants to save or exert peer pressure on participants. A growing literature suggests that people may fail to save because doing so is not salient (Taubinsky, 2013; Karlan *et al.*, 2016); and, correspondingly, they might be induced to do save by reminders. In-person visits may also generate peer pressure (Kast *et al.* (2012); Breza and Chandrasekhar (2015)). Second, the collection service has a commitment feature: we reduce the cost of making deposits, but not withdrawals. This feature constrains the ability of individuals to use their funds in the future since they must travel to make a withdrawal.²⁹ If participants are present-biased and sophisticated, they could use the account to constrain future behaviour. Finally, the product may reduce loss of savings either from theft or from others in the household demanding transfers. This is because once a deposit is made it is difficult to withdraw. Capture by others or loss through theft would imply that the effective pre-intervention interest rates (*e.g.* saving in cash) may have been negative, in turn implying that the increase in real interest rates was more than 5%.

29. A substantial literature in developing countries studies products that provide commitment against withdrawal (Ashraf *et al.*, 2006; Bryan *et al.*, 2010; Kast *et al.*, 2012; Dupas and Robinson, 2013b; Brune *et al.*, 2016; Blumenstock *et al.*, 2018). Indeed, one recent large experiment finds little effect of savings accounts without commitment features (Karlan *et al.*, 2016), and another with Malawian farmers suggests the potential benefits of commitment may outweigh the cost of reduced flexibility (Giné *et al.*, 2012).

Exploring these additional explanations can shed light on how expanding savings opportunities might increase labour effort and income. In addition, identifying the operative features of the treatment is useful for policy design.

Several elements of our design allow us to adjudicate between these potential explanations. First, after 6 months, we introduced lockboxes in a subset of our control zones. This intervention still reduces travel costs and protects against potential self-control and other-control impediments to saving. At the same time, in contrast, it has much more limited elements of peer pressure and provides much less of a reminder. Second, we elicit structural estimates of time preference using a hypothetical time-dated monetary payment protocol. This permits an examination of whether individuals who are more present-biased in a task designed to measure preferences exhibit larger savings, income, and labour responses to treatment.³⁰ We also have geolocations for the centre of each of the 156 zones in our sample, allowing us to measure the travel distance and travel time along roads from the zone to the bank (see Figure 1). Our sample was selected deliberately to lie between 5 km and 15 km from the bank. Conceivably, deposit collection reduces travel costs more substantially for participants who live farther from the bank. We therefore check for heterogeneous effects by travel distance. Finally, we examine the role of other-control using responses from baseline survey questions indicating whether the respondent had control over formal and informal savings decisions.

5.1. *Limited attention and peer pressure*

In July 2011, 6 months after deposit collections began in the treated sample, we introduced a lockbox treatment in 26 of our 78 *control* zones. As we discussed in Section 2.21 above, this treatment was integrated with the mobile POS technology. Deposit collectors visited the box periodically, added all deposits to user accounts, and left behind deposit receipts. Centralizing collection greatly decreased costs to the bank. We consulted households in the 26 zones to determine a secure location within their cluster of houses for the lockbox to be installed. Deposits in the boxes were collected by a bank agent weekly in 13 of the zones, and biweekly in the other 13 zones. Participants collected receipts for the deposits from envelopes in a compartment on the outside of the box. This introduction of a less personal collection service allows some scope to test for the importance of reminders and peer pressure. Importantly, the introduction of savings treatment to a new group of participants also provides a replication of the initial experiment in a different sample at a different time. We examine the effects of this intervention comparing the 13 zones given weekly box collection to the control, running regressions of the form:

$$y_{izt} = \beta \cdot \text{Lockbox}_{zt} + \gamma_i + \eta_t + \varepsilon_{izt}. \quad (5.1)$$

y_{izt} are the individual savings, income, or labour-related outcomes. The treatment effect is measured relative to respondents in the 52 control clusters that remained as undisturbed controls following the start of the lockbox experiment. Table 7 reports the corresponding estimates. Panel A reports effects on savings outcomes while Panel B reports effects on labour-related outcomes. In columns 1–3 of Panel A, we see that this treatment increased both the number and value of deposits, though by only about one-third as much as the deposit collection response reported in

30. The elicitation of preferences using time-dated monetary payments is more likely to accurately reveal present bias when conducted with credit-constrained individuals because the intertemporal fungibility of money may pose less of a confound (Cubitt and Read, 2007; Chabris *et al.*, 2008; Andreoni and Sprenger, 2012; Carvalho *et al.*, 2014; Augenblick *et al.*, 2015). Our sample, during the exercise, was selected specifically for lack of access to formal finance, and so it is plausible that the exercise in this environment could reflect present bias.

TABLE 7
Savings impacts comparing control to control-to-Box treatment (switches on in round 12)

| | NSB institutional data | | | Survey data | | | |
|--------------------------------------|-----------------------------|---------------------|------------------------------|---------------------------|-----------------|---------------------------|---------------|
| | No. of program transactions | Program deposits | Program deposits-withdrawals | Total no. of transactions | Bank savings | Informal and cash savings | Total savings |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| Panel A: Savings impacts | | | | | | | |
| Treated with box collection | 1.04*** | 446.1** | -33.3 | 1.06*** | 138.2 | -788.3* | -704.7 |
| | (0.23) | (179.0) | (208.2) | (0.25) | (244.0) | (418.7) | (520.9) |
| Control group mean | 0.098 | 39.4 | -4.53 | 0.62 | 740.8 | 4,253.9 | 5,192.2 |
| Observations | 4,890 | 4,890 | 4,890 | 4,890 | 4,890 | 4,890 | 4,890.00 |
| R-squared | 0.50 | 0.12 | 0.099 | 0.45 | 0.23 | 0.28 | 0.29 |
| Survey data | | | | | | | |
| | Y/N: Wage employment? | Y/N: Ag employment? | Y/N: Self-employment? | Wages earned | Business income | Wage hours | SE hours |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| Panel B: Earnings and labour impacts | | | | | | | |
| Treated with box collection | 0.063** | 0.00085 | -0.11*** | 782.7* | -929.9 | 2.94** | -2.27 |
| Control group mean | 0.15 | 0.59 | 0.55 | 1,158.4 | 5,831.8 | 5.72 | 25.7 |
| Observations | 4,890 | 4,890 | 4,890 | 4,890 | 4,890 | 4,890 | 4,890 |
| R-squared | 0.83 | 0.91 | 0.91 | 0.80 | 0.68 | 0.81 | 0.84 |

Notes: This table reports on the effect of the box collection treatment on savings outcomes (Panel A) and Earnings and Labour outcomes (Panel B). Results use data from round 2-18; the control group for this analysis remains untreated throughout, and the box treatment begins in round 11 when zones that had previously had no direct ability to save were given access to lock boxes in the zone centre that were collected weekly or biweekly by bank agents for NSB. Each column within a panel reports results of separate regressions using individual- and month-level fixed effects. Standard errors clustered at the zone (treatment unit) and survey wave level using the method of Cameron *et al.* (2011) are reported in parentheses. The unit of observation is the individual/month. Outcomes are Winsorized at the 99% level. Regressions pool samples surveyed monthly and samples surveyed quarterly. * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 2. In Columns 4–7, we see no clear impact on total savings, except perhaps that deposits crowded out informal savings. Therefore, the lockbox appears to have a less substantial effect on financial behaviour, and we find no clear evidence that the increase in NSB institutional deposits leads to an increase in overall savings. (Indeed, it appears to lead to some reduction in informal savings.)

However, the deposit collection boxes do appear to create a labour market response that is very similar to what we see in the initial treatment group. These results are reported in Panel B. Wage employment increases (column 1), self-employment decreases (column 3), wage income rises (column 4), and business income falls (column 5). Correspondingly, wage hours increase (column 6), and hours in self-employment decrease (column 7). These responses are similar to those seen in Panel C of Tables 4 and 6 (the non-manufacturing self-employed): a shift of labour toward wage work that is fully counterbalanced by a decline in business income. We interpret these results as showing that the lockbox savings products lead participants to make attempts to increase income, but these efforts are ultimately not successful, and, hence, lead to no additional income to channel into savings. Given sample sizes, we can only speculate why the lockbox treatment does not lead to an increase in income even among the self-employed in manufacturing. Perhaps these results suggest a link between salience provided by door-to-door collection and the production of overall income. Households may need the pressure of the weekly home visits to be successful in translating an increase in labour effort into an increase in total income. Overall, however, the

effect of the lockbox treatment broadly confirms the link between improved savings options and labour decisions.

5.2. *Self-control*

Much of the recent interest in the theory underlying the use of savings products is due to behavioural economics' focus on the difficulty that people have putting aside money from daily cash flows to meet savings goals. The literature provides several distinct reasons that individuals may fail to save, and many of these explanations suggest that a high-frequency deposit collection service can alter behaviour in important ways. Individuals with self-control problems will achieve a higher savings trajectory if the long-term self can lock the short-term self into a commitment [Laibson \(1997\)](#). Frequent deposit collecting allows people to get loose cash away from themselves and may permit them to control their own consumption of "temptation goods" ([Banerjee and Mullainathan, 2010](#)). Our savings product provides an interesting window on these behavioural dimensions: it features convenient deposit and inconvenient withdrawal, and therefore has a commitment dimension.

The fact that we find no evidence that treatment decreases expenditures suggests that standard arguments about controlling consumption cannot be playing a driving role in the overall effects of the program. The commitment feature of the product might allow present-biased individuals to select new consumption streams more aligned with long-run preferences. Through this channel, we would expect to see that income increases more, relative to controls while there would be ambiguous effects on aggregate consumption, with savings balances building up faster for the time inconsistent.

To test this, we measured present bias using a standard hypothetical multiple price list; our baseline survey asked the respondent "Suppose someone was going to pay you Rs. 1,500 1 month from now. He/she offers to pay you a lower amount today. What amount today would make you just as happy as receiving Rs. 1,500 in 1 month?" The answer to this question gives the total current discount factor, or the product of β and δ if intertemporal preferences take the quasi-hyperbolic form $U = E_0 \left[u(c_0) + \beta \sum_{t=1}^T \delta^t u(c_t) \right]$. We then ask the same question with the time scale moved out by 1 year (12 months from now versus 13 months), and the β parameter can be calculated by the ratio of the former to the latter answer.

In addition, because of the potential for deposit collection to control the consumption of "sin" or "temptation" goods, we conducted a survey experiment to induce truthful revelation of goods whose use may be stigmatized. The survey experiment asked respondents to pick (and not reveal to the enumerator), prior to each question about consumption of different sin goods, a card numbered from 1 to 10. If the number was between 4 and 10 the respondent was to answer "blue" if they consumed that sin good in the past month, and "red" if not, but if the number was between 1 and 3 they were to answer "blue" regardless of the correct response. In this way, enumerators can never infer the consumption status of a specific individual, but the sample prevalence of sin good consumption can be readily calculated. The same property of aggregation holds across sin goods for an individual; consequently the panel dependent variable used in the analysis is $\sum_j \frac{I_j(\text{blue}_j)}{0.7(J-0.3)}$, where $I_j(\text{blue}_j)$ is an indicator for a response of "blue" for each of the J sin goods in the survey experiment. We asked about purchase of alcohol, purchase of tobacco, and about gambling on horse races and on cards. These questions were asked in pre-treatment wave 3 and in post-treatment waves 10 and 18, and are included in Supplementary Appendix C.

Table 8 examines the heterogeneity of treatment effects by the extent of present bias. We estimate these effects using regressions of the form:

$$y_{izt} = \theta_1 \cdot \text{Treated}_{zt} + \theta_2 \cdot \text{Treated}_{zt} \cdot \text{Immediate_DF}_i + \gamma_i + \eta_t + \varepsilon_{izt} \quad (5.2)$$

TABLE 8
Does deposit collection increase savings more for individuals with a lower immediate discount factor?

| Dependent variable | Total household income | Wage hours | Self employed | Business income | Consumption | Consumption of "Sin" goods | Survey experiment "Sin" good consumption | Bank savings | Non-bank savings | Total savings |
|--|------------------------|------------|---------------|-----------------|-------------|----------------------------|--|--------------|------------------|---------------|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) |
| Immediate discount factor * Treated | -8,038.8 | -6.50 | 0.100 | -6,117.9 | 7,914.1 | -0.81 | 0.11 | -3,490.1** | -5,433.9** | -8,527.1** |
| | (14,686.6) | (7.10) | (0.19) | (5,042.2) | (9,917.6) | (314.6) | (0.49) | (1,774.1) | (2,661.4) | (3,867.2) |
| Treated at average value of Immediate_DF | 3,235.2** | 1.56** | -0.034** | 491.8 | 850.4 | 4.31 | -0.0076 | 711.3*** | 75.9 | 936.0** |
| | (1,281.2) | (0.71) | (0.017) | (443.0) | (962.8) | (48.6) | (0.045) | (202.1) | (321.1) | (382.4) |
| Control group mean | 29,307.7 | 5.91 | 0.53 | 5,559.0 | 16,229.3 | 492.6 | 0.15 | 765.1 | 4,360.5 | 5,280.4 |
| Observations | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 | 1,986 | 9,172 | 9,172 | 9,172 |
| R-squared | 0.36 | 0.83 | 0.92 | 0.70 | 0.29 | 0.53 | 0.44 | 0.25 | 0.32 | 0.33 |

Notes: This table reports on the heterogeneous effects of treatment on financial outcomes. The interaction 'Immediate Discount Factor * Treated' gives the differential effect of the treatment according to the structural estimate of the beta parameter. The elicitation protocol for this measure is described in Section 5.2. This beta parameter is demeaned before interaction, meaning that uninteracted treatment term gives the impact of the treatment at the average value of beta (0.98), and the estimate of the beta parameter is Winsorized at the 10% and 90% levels. Dependent variables are Winsorized at the 99% level. Each column reports results of separate regressions using individual- and month-level fixed effects. Standard errors clustered at the zone (treatment unit) and survey-wave level, using the method of Cameron *et al.* (2011), are reported in parentheses. The unit of observation is the individual/month and the reported coefficients correspond to dummy variables equal to one for individuals offered weekly home visits from National Savings Bank deposit collectors in the months indicated. All columns in the table use survey data, and outcomes are monthly flows for the variables indicated. Data include 4 months of pre-treatment data and 13 treated months. Regressions pool samples surveyed monthly and samples surveyed quarterly. * significant at 10%; ** significant at 5%; *** significant at 1%.

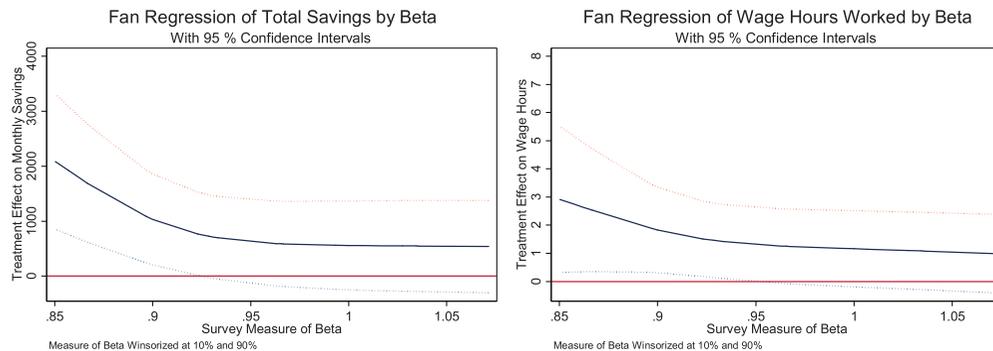


FIGURE 6
Fan regression by time inconsistency

the *Immediate Discount Factor* (DF_i), is our experimental measure of subjects' present bias corresponding to β in the quasi-hyperbolic utility function and the remaining variables are defined identically to specification (4.1) above. To improve the precision of the estimates, we winsorize *Immediate_DF* at the 10th and 90th percentiles. Given the values on the price list used to measure preferences, subjects at the 10th percentile have a *Immediate_DF* of 0.862 and subjects at the 90th percentile have a *Immediate_DF* of 1.071.³¹ We have demeaned *Immediate_DF*_{*i*} so that our estimate of θ_1 reflects the treatment effect at the mean *Immediate_DF*_{*i*} (which is *Immediate_DF*_{*i*} = 0.98 in both the winsorized and un-winsorized samples). From this average effect, θ_2 then provides the heterogeneity in impacts corresponding to a decrease in present bias (lower values of *Immediate_DF* imply more present bias).

The point estimates suggest that the positive effect of deposit collection on income (column 1), wage employment (column 2), business income (column 4), bank savings (column 8), non-bank savings (column 9), and total savings (column 10) are larger for present-biased individuals. For all three of the savings effects, the heterogeneity is significant at the 95% level implying that those with self-control problems find the home visits more useful in effectively increasing savings. To provide a sense of the difference, the estimated effect of deposit collection on total savings is 936 LKR at the mean of *Immediate_DF*_{*i*} while it is 1,422.04 LKR at the 10th percentile of the winsorized distribution (*Immediate_DF* = 0.923). Likewise, the estimated effect on total income is 3,235.2 LKR at mean *Immediate_DF*_{*i*} and 3,692.41 LKR at *Immediate_DF* = 0.923, though in this case the interaction is imprecisely measured. We present Fan regressions in Figure 6 that illustrate this same heterogeneity, showing that significant effects of the treatment on wage hours and overall savings are restricted to individuals who have measured immediate discount factor of less than about 0.925. The estimated effects on both our direct survey measure of "sin good" consumption, as well as the measure obtained using a randomized response technique, are small and imprecise.

These results suggest that the commitment features of deposit collection play some role in increasing savings by offsetting a reduction in savings due to present bias. Combining this with the intuition from standard models of intertemporal labour allocation—that increases in the value of one's labour provide an additional inducement to work in the current period—suggest that the commitment features of both the deposit collection and deposit boxes are a reason subjects are working harder. However, as Figure 6 shows, the data are consistent with an increase in

31. It is plausible that these respondents may have misunderstood the experimental task. Unwinsorized results are qualitatively similar, though less precisely estimated (results available on request).

savings for every level of beta. We additionally check for heterogeneity in treatment effects by distance to the bank along the most direct road (Supplementary Appendix Table A9), by numeracy (Supplementary Appendix Table A10), years of education (Supplementary Appendix Table A11), household size (Supplementary Appendix Table A12), the number of household members of working age (Supplementary Appendix Table A13), and the share of members of working age (Supplementary Appendix Table A14). There is limited apparent heterogeneity in deposit-collection impact along these dimensions (though the estimates are in cases imprecise) in linking deposit collection to increases in savings and in labour effort. Taken together with the results from the prior section, these tables suggest that behavioural explanations do play some role in moderating the effect of an interest rate shock on labour, income, and savings.

5.3. *Other-control*

When members of a household disagree on the optimal savings trajectory, private commercial savings may allow the pro-savings individual in the household to control the household savings rate (Ashraf *et al.*, 2006; Schaner, 2017). We examine the extent to which interpersonal conflict drives savings by testing for impact heterogeneity over the degree of financial control respondents had at baseline. In the second of our five baseline rounds, we asked the question, “Do you have control over the use of money needed . . . for savings (such as in financial institutions, *e.g.* banks) for yourself/your children?” We code a dummy for “savings no control” equal to one if the respondent reports no control over either formal or informal savings vehicles.³² This direct survey question provides a straightforward way of examining whether treatment had a stronger impact on individuals who lacked control over savings decisions at the time the intervention began.

The interaction analysis presented in Table 9 indicates that the impact of the treatment was not concentrated among those who lacked financial control. These regressions include individual and round fixed effects, the panel dummy for treatment, and the interaction between treatment and an indicator for “had no savings control at baseline”. Importantly, it does not appear that the effect of deposit collection on bank savings (column 8), non-bank savings (column 9), or total savings (column 10) is localized to this group. Correspondingly, there is no clear differential impact for this group for wage hours (column 2), self-employment (column 3), or business income (column 4). Effects on total income (column 1) and consumption (column 5) do appear to be larger for this group, though the interaction terms are not statistically significant. Curiously, there is some evidence that deposit collection increased consumption of sin goods in this group (columns 6 and 7). Collectively, this does not suggest a major role for other control. We emphasize that this context is atypical for a developing country in that financial control is not highly gendered. In our sample, 63% of men, and 54% of women, indicate having no control of either formal or informal savings decisions.

An alternate approach to this problem is presented in Supplementary Appendix Table A15, which draws on an idea from Anderson and Baland (2002) in which the share of household income earned by the respondent is taken as a measure of control over finances. Following their methodology, we analyse heterogeneity by terciles of the share of income earned by the respondent and find results opposite to those found in their article: total savings impacts are slightly larger for the top and bottom terciles than the middle tercile. While earnings impacts are larger for those with less income to begin with, there is no evidence of significant differential impacts on consumption,

32. This approach is intended to identify respondents with low degrees of spousal control. 456 of the 788 respondents in our sample (58%) report having no control. Results are similar if the analysis is conducted using only control over formal savings.

TABLE 9
Heterogeneous impacts by degree of control over savings at baseline

| Dependent variable | Total household income | Wage hours | Self employed | Business income | Consumption | Consumption of "Sin" goods | Survey experiment "Sin" good consumption | Bank savings | Non-bank savings | Total savings |
|---------------------------------|------------------------|-----------------|--------------------|------------------|-----------------------|----------------------------|--|--------------------|-------------------|------------------|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) |
| Treated | 1,739.90 (1,873.8) | 1.53 (0.94) | -0.036* (0.020) | 686.4 (727.2) | -69.0 (1,482.7) | 0.64 (65.2) | -0.099* (0.051) | 570.3** (267.2) | 138.6 (433.2) | 834.3 (572.4) |
| Treated * No savings control | 2,250.80 (1,819.8) | -0.21 (0.80) | -0.0071 (0.023) | 405.0 (799.2) | 1,456.40 (1,204.3) | 97.0** (43.3) | 0.16*** (0.051) | 201.0 (298.7) | -155.0 (408.5) | 92.6 (598.7) |
| Control group mean | 29,393.04 | 9.30 | 0.53 | 5,758.68 | 16,316.90 | 513.34 | 0.15 | 756.81 | 4,307.43 | 5,213.32 |
| Observations | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 | 9,172 | 1,986 | 9,172 | 9,172 | 9,172 |
| R-squared | 0.70 | 0.85 | 0.96 | 0.63 | 0.64 | 0.56 | 0.52 | 0.31 | 0.52 | 0.52 |

Notes: This table reports on the heterogeneous effects of treatment on financial outcomes. The interaction 'Immediate Discount Factor * Treated' gives the differential effect of the treatment according to the survey response estimate of the β parameter. The β parameter is demeaned before interaction, meaning that uninteracted treatment term gives the impact of the treatment at the average value of $\beta=0.98$. The estimate of the β parameter is Winsorized at the 10% and 90% levels. Dependent variables are Winsorized at the 99% level. Each column reports results of separate regressions using individual- and month-level fixed effects. Standard errors clustered at the zone (treatment unit) and survey-wave level, using the method of Cameron *et al.* (2011), are reported in parentheses. The unit of observation is the individual/month and the reported coefficients correspond to dummy variables equal to one for individuals offered weekly home visits from National Savings Bank deposit collectors in the months indicated. All columns in the table use survey data, and outcomes are monthly flows for the variables indicated. Data include 4 months of pre-treatment data and 13 treated months. Regressions pool samples surveyed monthly and samples surveyed quarterly. * significant at 10%; ** significant at 5%; *** significant at 1%.

transfers, or savings by income tercile. Thus, taking this approach, we again conclude that the impact of our mobile deposit collection does not arise as a result of improved ability to control the behaviour of other household members.

5.4. *The impact of formal accounts on informal savings Institutions*

At the time of designing the study we were concerned with the possibility that formal savings accounts, by creating a private and concealable way of storing wealth, might undermine the social insurance provided by informal savings institutions such as Rotating Savings and Credit Associations (ROSCAs) Besley *et al.* (1993). Such institutions are fragile in that they rely on social coercion to enforce contributions (Anderson *et al.*, 2009). By opening up new exit options, particularly for those with substantial savings liquidity, private accounts could spell trouble for the effectiveness of both informal saving and risk-pooling mechanisms. To address, this, we randomized the intensity of the weekly home treatment among *other* members (besides the survey respondent) for a set of *seetus* as described in Section 2.2.2. As an increasing share of these other members have access to private personal savings accounts, they may become more likely to renege on ROSCA contributions, and less likely to exert effort monitoring the behaviour of group members. Thus, if formal savings are a substitute for informal savings groups, we expect to see the attractiveness of ROSCA savings fall as an increasing share of the ROSCA members renege on their savings commitments and on their attempts to control the behaviour of fellow group members. At the extreme this could lead to group collapse, and on the intensive margin it could lead to a decline in the number of ROSCAs an individual chooses to participate in and the quantity of savings achieved through ROSCAs.

To understand the impact of *seetu* saturations on informal insurance, we use the monthly flows of transfers into the household and out of the household (we screened out cases in which a second household member was in the same *seetu* as the respondent, and so the interpersonal flows of informal insurance should be primarily inter-household). We also examine the effect of the *seetu* saturation experiment on overall savings. The sample for this analysis consists of the sampled 20% of the membership of each of the 84 close *seetus* drawn into the survey at baseline, and the experimental variation arises from whether an additional random 0, 20, or 40% of their ROSCA was offered the POS weekly home savings product.

We use two empirical specifications for each outcome. The first examines the overall intention-to-treat effect within the close *seetu* membership. The second includes a dummy for treatment and further includes the randomly chosen saturation within treatment *seetus*. In this specification the dummy for treatment gives the impact in the 20% treatment saturation *seetus*, and the slope term on the saturation gives the marginal effect on outcomes of treating (but not studying) an additional fraction of the membership beyond this. Because we have no individuals who are studied but not treated in treatment *seetus*, we are only able to estimate the Spillover on the Treated, and not the Spillover on the Not Treated (Baird *et al.*, 2018). The effects of the saturation should therefore be interpreted as the impact of increasing the fraction of the *seetu* given the deposit collection treatment, conditional upon the fact that the studied individual is treated.

Supplementary Appendix Table A16 shows the balance of the *seetu* saturation experiment. This randomization was conducted in a very small number of units; the saturation experiment is particularly subject to potential imbalance as each saturation cell contains only 13 *seetus*. We show balance over 15 tests, and find only one of them to be significant at the 10% level. Nonetheless there are some quantitatively large differences, particularly across the saturation amounts within the treatment. Overall savings is imbalanced across the distribution of the saturation. We therefore proceed to the analysis of the saturation experiment with some caution as the number of assignment units is small, and the balance not perfect.

TABLE 10
Seetu saturation experiment treatment effects

| | Number of Seetus currently participating in | | Monthly amount saved through Seetus | | Total monthly savings | |
|----------------------------|---|----------|-------------------------------------|----------|-----------------------|----------|
| Treated | 0.27* | 0.16 | 238.2 | 198.8 | 945.8* | 844.7 |
| | (0.14) | (0.13) | (198.3) | (190.1) | (546.3) | (585.5) |
| Seetu treatment saturation | | 1.42*** | | 509.4 | | 1304.7 |
| | | (0.53) | | (470.2) | | (2165.2) |
| Constant | 1.263*** | 1.273*** | 2,167*** | 2,170*** | 2,818*** | 2,827*** |
| | (0.08) | (0.08) | (146.40) | (147.60) | (441.50) | (443.10) |
| Control group mean | 1.81 | 1.81 | 2,032.60 | 2,032.60 | 5,311.90 | 5,311.90 |
| Observations | 3,255 | 3,255 | 3,255 | 3,255 | 3,255 | 3,255 |
| Number of HHs | 278 | 278 | 278 | 278 | 278 | 278 |

Notes: This table reports on the effect of treatment and the saturation of treatment in respondent's seetu (informal savings groups) on financial outcomes. All regressions use data across survey rounds 2-18, and outcomes are Winsorized at the 99% level. Standard errors clustered at the zone (treatment unit) and survey wave level using the method of Cameron *et al.* (2011) are reported in parentheses. All columns in the table use survey data, and outcomes are monthly flows for the variables indicated. Data include 4 months of pre-treatment data and 13 treated months. "Treatment" is a dummy indicating that the respondent household is receiving weekly home visits, and "Seetu Treatment Saturation" gives the fraction of the respondents seetu that also received the offer of treatment; this saturation was randomized to 0, 0.2, or 0.4 for seetus that had all members living in the same zone and did not contain more than 24 members. Each column reports results of separate regressions using individual- and month-level fixed effects. Treatment of non-study seetu members began in round 6. Regressions pool samples surveyed monthly and samples surveyed quarterly. *significant at 10%; **significant at 5%; ***significant at 1%.

The results of this analysis are given in Table 10. Contrary to the hypothesis that private individualized savings poses a threat to *seetu* participation, the treatment and the treatment saturation appear to lead to superior outcomes for the *seetus*. Treatment on average increases the number of *seetus* in which a respondent participates by 0.27 over a base of 1.8, an increase of more than 10%. Allowing for heterogeneity across the saturation rate indicates that where more individuals are treated, membership in *seetus* increases more rapidly. The coefficient of 1.42, significant at the 1% level, indicates that each 20% increase in the fraction of the *seetu* treated increases the number of *seetus* participated in by its members by a further 0.3. The monthly amount saved through *seetus* is higher in treatment than in control zones, and the saturation effect is again positive and large in absolute magnitude, although not significant. The total amount saved through the *seetu* is 238 LKR per month higher in treatment than control *seetus*, and the saturation of treatment in the informal savings groups is also positive (though insignificantly) correlated with total savings. Hence, the result of the randomized saturation experiment indicates that, far from posing a threat to *seetus*, the bank-driven deposit collecting appears to be fortifying informal savings groups.

6. CONCLUSION

We conducted an experiment with a powerful inducement to save: weekly visits by a deposit collector equipped with a wireless POS device capable of printing out account balance receipts on the spot. The experiment was implemented in Sri Lanka, a country with strong informal savings institutions and a tradition of using formal banks for "children's accounts" into which parents often begin to put savings at the birth of a child. Despite this, we find strong effects of the treatment not just on savings into program accounts, but into bank savings and overall savings as a whole. The data therefore suggest that, even though the sample had access to a wide range of formal and informal savings mechanisms, the simple low-transactions-cost account represented a significantly better option. We tracked the sample with monthly surveys including five pre-treatment waves, and a careful cash-flow accounting instrument designed to answer two

questions: First, what are the headwaters of formal savings? Second, what are the channels through which formal savings affects income and expenditures?

Our results confirm recent experimental findings that access to formal savings products can produce surprisingly large benefits for unbanked households, even when these accounts bear low real interest rates. Our detailed high-frequency surveys help to shed light on the mechanisms behind these savings increases. We find a large and immediate increase in income following treatment, allowing savings to increase without any accompanying decrease in consumption. The surveys reveal that this newfound income is not generated by increased investment in entrepreneurial activities; indeed, the savings product triggers an *exit* from self-employment. In aggregate, households work more hours in wage work. An increase in earned income explains a substantial majority of the increase in income.

The novelty of our results, then, is that we show substantial increases in savings and income without any foregone consumption, and without the save–invest–earn channel that is suggested by the literature on non-convexities in production. Earnings increase as soon as participants gain access to the deposit collection service, and they increase in the face of exit from self-employment. These overall sample impacts are made up of a number of distinct sub-group effects. When self-employed households with manufacturing businesses were offered weekly home deposits, they successfully expanded earnings by working harder in their own business; the higher income resulted in higher total savings. Among those with non-manufacturing businesses, households exited self-employment and increased labour in the wage market. The treatment effect on labour allocation is large in this group—one in six individuals have exited self-employment by the end of the study—but ultimately ineffective. The increase in wages is counterbalanced by a fall in business income; neither total income nor total savings improves. The impact of the weekly home box treatment largely mirrors this result. The effort to raise income by working harder in the wage market is unsuccessful.

Thus, only the subgroup successful in raising aggregate income was ultimately able to increase total savings: we find no headwaters for savings other than income. The implication is that new institutional savings options will be most effective in raising total savings when they are offered to households with potential to increase income, for example, by profitably expanding effort in self-employment. A high-salience inducement to save regularly may be important to induce participants to exert the continuous effort required to increase income through labour effort. We note that while we believe our subsample analysis illuminates the mechanisms, our main conclusions derive from the full sample analysis on which the experiment is based.

The centrality of income generation underlines the challenge in detecting impacts when they move through an outcome as difficult to detect as labour inputs. For example, Dupas and Robinson (2013b) find a positive treatment effect of their savings product of between 0.15 and 0.31 h worked per day in their samples. The effect we find is of similar magnitude—1.29 h per week, or ~ 0.25 h per day—but this effect is more precisely measured and is statistically significant. A substantial reason our study has the ability to detect this channel is the long, high-frequency panel we collected, providing us statistical power for highly stochastic outcomes McKenzie (2012).

In one sense it is surprising that a financial service offering nothing more than a low-cost means to save in a bank could increase household income. However, as in classical macroeconomic models of intertemporal labour allocation, the responses we observe in the data are consistent with those we would expect from individuals who experience a large increase in the interest they receive on savings. These effects appear to be larger for present-biased individuals, suggesting that the commitment dimension of the product may reinforce the interest rate effect. These features of the product drive up the benefit of working today to save for tomorrow, leading individuals to increase labour effort in the initial period to build up savings. We read the full-sample results, supported by

the subsample analysis, as suggesting that the labour channel, rather than the investment channel, deserves some attention when studying the benefits of saving.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

REFERENCES

- ANDERSON, S. and BALAND, J.-M. (2002), "The Economics of ROSCAs and Intrahousehold Resource Allocation", *The Quarterly Journal of Economics*, **117**, 963–995.
- ANDERSON, S., BALAND, J.-M. and MOENE, K. O. (2009), "Enforcement in Informal Savings Groups", *Journal of Development Economics*, **90**, 14–23.
- ANDREONI, J. and SPRENGER, C. (2012), "Estimating Time Preferences from Convex Budgets", *American Economic Review*, **102**, 3333–3356.
- ASHRAF, N. (2009), "Spousal Control and Intra-household Decision Making: An Experimental Study in the Philippines", *American Economic Review*, **99**, 1245–1277.
- ASHRAF, N., KARLAN, D. and YIN, W. (2006), "Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines", *The Quarterly Journal of Economics*, **121**, 635–672.
- AUGENBLICK, N., NIEDERLE, M. and SPRENGER, C. (2015), "Working Over Time: Dynamic Inconsistency in Real Effort Tasks", *The Quarterly Journal of Economics*, **130**, 1067–1115.
- BAIRD, S., BOHREN, J. A., MCINTOSH, C. and ÖZLER, B. (2018), "Optimal Design of Experiments in the Presence of Interference", *Review of Economics and Statistics*, **100**, 844–860.
- BANERJEE, A., DUFLO, E., GLENNERSTER, R. and KINNAN, C. (2015), "The Miracle of Microfinance? Evidence from a Randomized Evaluation", *American Economic Journal: Applied Economics*, **7**, 22–53.
- BANERJEE, A. and MULLAINATHAN, S. (2010), "The Shape of Temptation: Implications for the Economic Lives of the Poor" (Technical Report, National Bureau of Economic Research).
- BANERJEE, A. V. and NEWMAN, A. F. (1993), "Occupational Choice and the Process of Development", *Journal of Political Economy*, **101**, 274–298.
- BEAUDRY, P. and VAN WINCOOP, E. (1996), "The Intertemporal Elasticity of Substitution: An Exploration Using a US Panel of State Data", *Economica*, **1**, 495–512.
- BESLEY, T., COATE, S. and LOURY, G. (1993), "The Economics of Rotating Savings and Credit Associations", *The American Economic Review*, **1**, 792–810.
- BLANCHARD, O. J. and FISCHER, S. (1989), *Lectures on Macroeconomics*, (Cambridge, MA, USA: MIT Press).
- BLUMENSTOCK, J., CALLEN, M. and GHANI, T. (2018), "Why Do Defaults Affect Behavior? Experimental Evidence from Afghanistan", *The American Economic Review*, **108**, 2868–2901.
- BREZA, E. and CHANDRASEKHAR, A. G. (2015), "Social Networks, Reputation and Commitment: Evidence from a Savings Monitors Experiment", (Technical Report, National Bureau of Economic Research).
- BRUHN, M. and LOVE, I. (2014), "The Real Impact of Improved Access to Finance: Evidence from Mexico", *The Journal of Finance*, **69**, 1347–1376.
- BRUHN, M. and MCKENZIE, D. (2009), "In Pursuit of Balance: Randomization in Practice in Development Field Experiments", *American Economic Journal: Applied Economics*, **1**, 200–232.
- BRUNE, L., GINÉ, X., GOLDBERG, J. and YANG, D. (2016), "Facilitating Savings for Agriculture: Field Experimental Evidence from Malawi", *Economic Development and Cultural Change*, **64**, 187–220.
- BRYAN, G., KARLAN, D. and NELSON, S. (2010), "Commitment Devices", *Annual Review of Economics*, **2**, 671–698.
- CAMERON, A. C., GELBACH, J. B., and MILLER, D. L. (2011), "Robust inference with multiway clustering", *Journal of Business & Economic Statistics*, **29**, 238–249.
- CHABRIS, C. F., LAIBSON, D. and SCHULDIT, J. P. (2008), "Intertemporal Choice", in Durlauf, S. N., Blume, L. *et al.* (eds) *The New Palgrave Dictionary of Economics* (Basingstoke: Palgrave Macmillan).
- CRÉPON, B., DUFLO, E., GURGAND, M., RATHELOT, R. and ZAMORA, P. (2013), "Do Labor Market Policies Have Displacement Effects? Evidence from a Clustered Randomized Experiment", *The Quarterly Journal of Economics*, **128**, 531–580.
- CUBITT, R. P. and READ, D. (2007), "Can Intertemporal Choice Experiments Elicit Time Preferences for Consumption?", *Experimental Economics*, **10**, 369–389.
- DE MEL, S., MCINTOSH, C. and WOODRUFF, C. (2013), "Deposit Collecting: Unbundling the Role of Frequency, Salience, and Habit Formation in Generating Savings", *American Economic Review*, **103**, 387–92.
- DEMIRGÜÇ-KUNT, A. and KLAPPER, L. (2012), "Measuring Financial Inclusion", *Brookings Papers on Economic Activity*, **1**, 279–340.
- DUPAS, P. and ROBINSON, J. (2013a), "Savings Constraints and Microenterprise Development: Evidence from a Field Experiment in Kenya", *American Economic Journal: Applied Economics* **5**, 163–192.
- DUPAS, P. and ROBINSON, J. (2013b), "Why Don't the Poor Save More? Evidence from Health Savings Experiments", *American Economic Review*, **103**, 1138–1171.
- GERBER, A. S. and GREEN, D. P. (2012), *Field Experiments: Design, Analysis, and Interpretation*, (New York, NY, USA: WW Norton).

- GINÉ, X., GOLDBERG, J., SILVERMAN, D. and YANG, D. (2012), *Revising Commitments: Field Evidence on the Adjustment of Prior Choices* (Washington, DC, USA: The World Bank).
- KARLAN, D. S., DUPAS, P., ROBINSON, J. and UBFAL, D. (2016), “Banking the Unbanked? Evidence from Three Countries”, (Technical Report, SSRN Working Paper 2815092).
- KAST, F., MEIER, S. and POMERANZ, D. (2012), “Under-savers Anonymous: Evidence on Self-help Groups and Peer Pressure as a Savings Commitment Device”, (Technical Report, National Bureau of Economic Research).
- LAIBSON, D. (1997), “Golden Eggs and Hyperbolic Discounting”, *The Quarterly Journal of Economics*, **112**, 443–478.
- LAZEAR, E. P. (2000), “Performance Pay and Productivity”, *American Economic Review*, **90**, 1346–1361.
- LUCAS Jr, R. E. and RAPPING, L. A. (1969), “Real Wages, Employment, and Inflation”, *Journal of Political Economy*, **77**, 721–754.
- MCKENZIE, D. (2012), “Beyond Baseline and Follow-up: The Case for More t in Experiments”, *Journal of Development Economics*, **99**, 210–221.
- PRINA, S. (2015), “Banking the Poor Via Savings Accounts: Evidence from a Field Experiment”, *Journal of Development Economics*, **115**, 16–31.
- SAMPHANTHARAK, K. and TOWNSEND, R. (2009), *Households as Corporate Firms: An Analysis of Household Finance Using Integrated Constructing Financial Statements from Integrated Household Surveys and Corporate Financial Accounting*. (Cambridge, MA, USA: Cambridge University Press).
- SCHANER, S. (2017), “The Cost of Convenience? Transaction Costs, Bargaining Power, and Savings Account Use in Kenya”, *Journal of Human Resources*, **52**, 919–945.
- TAUBINSKY, D. (2013), “From Intentions to Actions: A Model and Experimental Evidence of Inattentive Choice”, (Working Paper).
- THALER, R. H. and BENARTZI, S. (2004), “Save More Tomorrow™: Using Behavioral Economics to Increase Employee Saving”, *Journal of Political Economy*, **112**, S164–S187.