

The Persistent Power of Behavioral Change: Long-Run Impacts of Temporary Savings Subsidies for the Poor*

Simone Schaner[†]

March 1, 2016

Abstract

I use a field experiment in rural Kenya to study how temporary incentives to save impact long-run economic outcomes. Study participants who were randomly selected to receive large temporary interest rates on an individual bank account had significantly more income and assets 2.5 years after the interest rates expired. These changes are much larger than the short-run impacts on experimental bank account use and almost entirely driven by increased rates of entrepreneurship. Temporary interest rates directed to joint bank accounts had no detectable long-run impacts on entrepreneurship or income, but increased investment in household public goods and led to greater spousal consensus over financial matters. The short-run effects of modest unconditional cash payments were similar to those of the interest rates, but the cash payments had no apparent long-run impact on economic outcomes.

*I thank Erica Field, Dean Karlan, Erzo F.P. Luttmer, Rohini Pande, Joshua Schwartzstein, Jonathan Zinman, and seminar participants at UC Berkeley, Boston University, Brown University, IFPRI, Notre Dame, Tufts University, UC San Diego, and the University of Missouri for numerous useful comments. This project would not have been possible without the tireless assistance, hard work, and commitment of many employees of Family Bank. I am particularly indebted to Victor Keriri Mwangi, Steve Mararo, and Michael Aswani Were. I also thank James Vancel, Noreen Makana, and Thomas Ginn for superb field management and the IPA field officers for their excellent assistance with the data collection. I gratefully acknowledge the financial support of the Russell Sage Foundation, the George and Obie Shultz Fund, MIT's Jameel Poverty Action Lab, the National Science Foundation's Graduate Research Fellowship, and the Yale Savings and Payments Research Fund at Innovations for Poverty Action, sponsored by a grant from the Bill & Melinda Gates Foundation. All errors are my own.

[†]Dartmouth College, JPAL, and BREAD. simone.schaner@dartmouth.edu

1 Introduction

In spite of recent progress, approximately 700 million people still live in extreme poverty, with 35 percent of individuals in Sub-Saharan Africa living in this state (Cruz et al. 2015). As such, understanding how to help poor families grow their incomes remains a top global policy priority. Recent empirical work has produced compelling evidence that one-time “big push” interventions, where individuals are given large grants (and sometimes additional support) either in cash or in kind, can raise income and consumption levels, even years after grants have been disbursed.¹ But are resource-intensive big pushes always necessary? At least some individuals in the developing world seem to have the potential to grow themselves out of poverty: recent experimental and non-experimental studies have found very large – on the order of 5-30 percent per month – marginal returns to capital among microenterprises in contexts as varied as Sri Lanka, Ghana, India, Mexico, and Uganda.² Other researchers have documented individuals regularly revolving debt at interest rates as high as 10 percent per day (Aleem 1990; Ananth et al. 2007; Banerjee and Duflo 2007). Yet business training and financial literacy programs that coach the poor on how to change their behavior without providing incentives or economic support generally deliver disappointing results (McKenzie and Woodruff 2014; Karlan et al. 2014).

A simple alternative to coaching and counseling is to directly pay for the desired behavior change, at least at first. In this paper, I ask whether temporary, but high-powered incentives to save can have persistent impacts on economic outcomes. To this end, I analyze the results of a field experiment that I conducted in rural Kenya in 2009, in which married couples were given the opportunity to open new formal bank accounts. All participating couples could open up to three accounts: an individual account in the name of each spouse and a joint account. Before participants decided which accounts to open, each account was randomly assigned a temporary interest rate, which lasted for six months and ranged from zero to 20 percent in annual terms. The randomization was conducted so that all interest rates were completely independent of one another – thus, the experiment created variation in overall incentives to save *and* variation in the way that couples were incentivized to save (i.e. jointly versus individually).

¹For evidence on cash and in-kind grants to entrepreneurs, see de Mel et al. (2008), de Mel et al. (2012), and Fafchamps et al. (2014). For evidence on small business grants plus training or technical assistance for small business see Macours et al. (2012) and Blattman et al. (2014). For evidence on ultra-poor graduation programs, which combine grants of productive assets (usually livestock) with intensive technical assistance, see Banerjee et al. (2015), Bandiera et al. (2015), and Blattman et al.. See Gertler et al. (2012) for the long-term impact of *Oportunidades*, Mexico’s well-known conditional cash transfer program.

²For Sri Lanka, see de Mel et al. (2008) and de Mel et al. (2012). For Ghana see Udry and Anagol (2006) and Fafchamps et al. (2014). For India see Banerjee and Duflo (2012) and Field et al. (2013). For Mexico see McKenzie and Woodruff (2008). For Uganda see Blattman et al. (2014).

In the short run, the interest rates had their intended effect: study participants were much more likely to open and use accounts with higher interest rates when the interest rates were active. Although these short-run responses are robust, they are modest in magnitude: moving from no interest to 20 percent interest increased average daily balances in the individual and joint accounts by \$1.31 and \$2.09 respectively. In the longer run (3 years after the experiment and 2.5 years after the interest rates expired), I find that the interest rates changed savings behavior in more notable ways, and that the nature of the impact depends how the household was incentivized to save.

When an individually-owned bank account received a high interest rate, assets were re-allocated towards the owner of that account and as a result both individual and overall household income increased. These gains are driven by growth in entrepreneurship – study participants who received the highest interest rate on their individual account were significantly (11 percentage points, or a 30 percent increase from the mean) more likely to be entrepreneurs and had substantially more business profit and capital at endline. One striking feature of the individual interest rate is that its long-run impacts are much larger than its short-run impacts on experimental bank account use. Moving from the lowest to the highest individual rate increased the short-run average daily balance in individual accounts by less than \$2, yet in the long run the interest rate increased *monthly* individual business profit by \$6.85 and business capital by \$33. While these long-run effects are large, they are quite consistent with treatment effects on short-term account use and the returns to capital suggested by the data, which are well within the range of returns found by other researchers.³

In contrast, I cannot reject the null hypothesis that the joint interest rate had no impact on overall income or assets. I do, however, find some evidence that the joint rate increased investment in “household” assets such as home renovations and livestock. Couples who received higher joint interest rates also reported greater levels of spousal agreement regarding consumption and savings decisions. This suggests that the joint interest rate compelled couples to work together towards a mutually-agreeable savings goal, while the individual interest rate spurred study participants to invest in more independent (and, based on the income results, higher return) ventures.

Why didn’t study participants simply revert back to their old savings and investment behaviors after the interest rates expired? Perhaps the most straightforward explanation is that the interest rates helped study participants accumulate more capital in the short run, which helped them grow output in the long run, e.g. due to non-convexities in the

³The individual rate’s long-run treatment effects on business profits and assets could be attained by making an additional \$0.75 investment in productive capital and reinvesting the proceeds at a 14 percent monthly return for the next 2.5 years, for example.

production function. To test this hypothesis, I contrast the impact of the interest rates to that of a randomly-assigned modest cash payment, which did not explicitly incentivize saving. Twenty percent of study respondents were selected for this “cash prize”, which was delivered as part of the baseline discount factor elicitation procedure. The cash payments ranged from \$0.13-\$3.75 with an average value of \$3.09, were delivered within three months of the baseline sessions, and could be deposited in respondents’ newly-opened bank accounts. Although the cash payment increased short-run bank account balances more than the interest rates, it had no detectable impact on long-run economic outcomes. I therefore argue that explanations for the effects of the interest rates that rely solely on the capital stock are unlikely. Moreover, I find no evidence that the interest rates helped study participants by easing external savings constraints or increasing access to bank credit.

The contrast between the interest rate and cash prize results suggests that explicitly incentivizing study participants to exert effort to save (in a specific type of bank account) in the short run may have been important. More specifically, making a conscious, concerted attempt to save could have helped individuals establish new savings and investment practices that persisted after the interest rates expired. A notable prediction of models of habit formation is that behavior change induced by temporary incentives can continue after those incentives are removed (Becker and Murphy 1988). The interest rates could have also induced individuals to make plans or adopt new financial heuristics to support long-run behavior change (Thaler 1999). Consistent with these hypotheses, I find that respondents treated with higher individual interest rates were 27 percent (7.6 percentage points) more likely to say that they “saved regularly” 3.5 years after the interest rates expired. Moreover, those treated with higher individual interest rates were 35 percent (5.4 percentage points) more likely to explicitly budget for business expenses, and this increase is entirely driven by growth in *downwardly-rigid* business budgets: that is, budgets in which the respondent stated he or she would not reduce the allocated amount to meet an unexpected expense. Although these results are not definitive, they reinforce the notion that behavioral mechanisms like habit formation and mental accounting were important for sustaining the impact of the individual interest rate.

The main contribution of this paper is to demonstrate that temporary financial incentives can generate persistent changes in savings and investment behavior, and that these changes can translate into impacts on income. To date, most research on temporary incentives for behavior change has focused on education or health-related behaviors like gym attendance or smoking in developed-country contexts. Here, studies that track outcomes after incentives expire tend to do so for a limited period (e.g. one year or less after the intervention ends) and usually find modest-to-no impacts in the longer run, with treatment effects decaying over

time.⁴ This paper provides leading evidence that temporary incentives *can* have amplified effects in the long run, at least in the financial domain. One reason for this difference could be the fact that some individuals in my sample had very high (and potentially very salient) returns to saving, which could provide dynamic feedback to help reinforce new behaviors.

My results also contribute more broadly to existing research on saving in the developing world. Researchers have evaluated the impacts of a range of savings products, including basic formal bank accounts (Dupas and Robinson 2013; Prina 2015), commitment savings accounts (Ashraf et al. 2006b; Brune et al. 2013; Dupas and Robinson 2014), savings accounts with reminders (Karlan et al. 2013), savings groups that leverage peer pressure and support (Kast et al. 2013; Dupas and Robinson 2014; Breza and Chandrasekhar 2015), and deposit collection services (Ashraf et al. 2006a; Callen et al. 2014), among others. Although this literature has identified a number of services that help individuals increase savings balances, there is little evidence as to whether, or under what conditions, behavioral changes persist after services are discontinued – my paper helps fill this gap.

It is also important to note that my results stand in sharp contrast to most literature on business training and financial literacy. These studies evaluate programs that utilize educational modules, rather than financial incentives, to alter financial behavior and generally find limited impacts on economic outcomes, especially in the long run (McKenzie and Woodruff 2014; Karlan et al. 2014). In this sense my results align with Cole et al. (2011), who find that financial subsidies are more effective than financial literacy training at boosting takeup and long-run use of savings accounts in Indonesia. I deepen this finding by shedding light on how financial subsidies impact outcomes beyond the use of bank accounts targeted by the subsidies.

Of course, my results should not be taken to imply that interest rates are a panacea. Two recent papers find that much lower interest rates – both in absolute and real terms – have no measured impact on economic outcomes (Kast et al. 2013; Karlan and Zinman 2014). Moreover, my study was conducted in just two districts in Kenya, and there are several aspects of the design (e.g. multiple interest rates drawn by lottery, a choice between bank accounts, spouses attending baseline experimental sessions together) that may have been

⁴In terms of financial behavior, De Mel et al. (2013) find evidence that deposit collection services helped Sri Lankan study participants form a savings habit, but they do not study impacts on outcomes beyond saving. Studies of incentives to exercise or lose weight generally find that initial short-run effects dwindle over time, especially when habits are interrupted by holiday breaks (Charness and Gneezy 2009; John et al. 2011; Acland and Levy 2015; Royer et al. 2012). Volpp et al. (2009) find that financial incentives to quit smoking have decaying impacts over the longer-run, while Giné et al. (2010) find that a commitment savings account tied to smoking behavior had small effects that lasted at least six months after savings were released. In education, Jackson (2010) finds that paying students for passing advanced placement tests improved standardized test scores and increased college matriculation.

important for my results. I therefore prefer to interpret my results as demonstrating that the *right incentives* targeted to the *right population* can have meaningful, lasting effects. Further work would be needed to generate additional evidence on optimal design and targeting of these incentives to maximize policy impact.

The remainder of the paper proceeds as follows: Section 2 begins by describing the experimental context and design, then Section 3 presents the main results. Section 4 provides additional discussion with a focus on mechanisms, and Section 5 concludes.

2 Experimental Design and Data

2.1 Experimental Context

The experiment was conducted between July and September 2009 in partnership with Family Bank of Kenya – Appendix Figure A1 illustrates a timeline of major experimental activities. Study participants were recruited from 19 communities in two districts near the bank’s Busia branch in Western Province. Even though six formal banks were operating in Busia at the time, most of these banks did not offer low-cost accounts suitable for a low-income clientele. Importantly, Family Bank had just begun to market a new, low-fee account at the onset of the experiment. In contrast to traditional Kenyan bank accounts, which required relatively large minimum balances (around Ksh 1,000, or US\$12.50 at a 2009 exchange rate of Ksh 80 per dollar) and charged monthly account maintenance fees, the new Family Bank account had a minimum operating balance of Ksh 100 (\$1.25), no maintenance fees, and no deposit fees. The only fees charged were for withdrawals, which cost Ksh 62 over-the-counter and Ksh 30 at the ATM. Like most bank accounts on the market, the new Family Bank account did not bear any interest.

Just one other bank apart from Family Bank offered a similar low-cost account when the experiment began. The bank accounts in this study were therefore a relatively new technology and many study participants were unfamiliar with them. However, when interpreting the results it is important to keep in mind that Kenya’s financial services landscape has evolved dramatically since 2009. By the time of the 2012 endline survey nearly all banks offered low-cost accounts. Banking services have also been integrated into mobile money products and banks have expanded their reach into rural areas via agency banking. Thus, while the experimental accounts dominated most other accounts on the market in 2009, this was no longer the case at the time of the endline survey.

2.2 Experimental Design

The experiment targeted married couples who did not have any accounts with Family Bank but expressed interest in opening one. Trained enumerators recruited couples in 19 communities surrounding the bank branch. Couples who expressed initial interest were issued invitations to attend a group meeting at a local primary school. All baseline interviews, account opening paperwork, and randomization activities were conducted at these meetings. Upon arrival, couples were informed that they could open up to three accounts with Family Bank at the meeting – a joint account, an individual account for the husband, and an individual account for the wife. All accounts opened at the meetings were funded with the Ksh 100 minimum balance to eliminate barriers to account use. This amount could not be withdrawn, so participants had little incentive to open accounts they knew they would never use.

Before deciding which accounts to open, participants were given the opportunity to draw a temporary 6-month interest rate for each account. This process was designed so that interest rates on the three accounts were independent of one another. All randomization was conducted in the field, with respondents drawing folded envelopes from plastic bins.⁵ Individual accounts were assigned an annual interest rate of either 0, 4, 12, or 20 percent with equal probability, while joint accounts were assigned an annual interest rate of either 4, 12, or 20 percent with equal probability. The interest rates in the experiment were purposely chosen to exceed market rates by a large margin, with the hope that such substantial subsidies would stimulate a short-run savings response. At the time, most formal financial institutions offered no interest on small-scale savings balances – at best participants could have earned 0.5-2 percent annually elsewhere.⁶

Since many study participants had little-to-no experience with banks, project staff carefully explained what an interest rate was, provided numerical examples for each interest rate, and explained that the promotion would only last for six months. While very few couples chose to open all three bank accounts, all couples opened at least one account – as a result 99 percent of study participants had access to either a newly-opened joint account or a newly-opened individual account in their own name (Appendix Table A1 shows the distribution of account opening choices). Thus, the experimental design allows me to study the impact of interest rates on different account types holding access to a new bank account constant. Participants were also given a pocket-sized card for each account that they opened, which

⁵Respondents took separate draws for each potential account. The field staff were carefully trained not to allow respondents more than one draw for each treatment. I find no evidence of protocol problems when comparing the empirical distribution of treatments to the theoretical distribution of treatments.

⁶For comparison, average inflation was 9.3 percent in 2009 and 2.0 percent in 2010 – therefore most bank accounts offered a negative real rate of return.

featured a reminder to save and, when applicable, the interest rate.⁷

Before leaving the meeting all individuals participated in a final drawing for a “cash prize”.⁸ This prize was the incentive for baseline questions on rates of time preference, which consisted of choices between a smaller monetary amount at time t and a larger amount at time $t + \tau$ (see the Data Appendix for additional detail). All individuals had a 20 percent chance of being selected for a cash prize – selected individuals then drew one of their time preference questions at random for payout. Payouts ranged from Ksh 10 to Ksh 300, with an average payout of Ksh 247.⁹ All payouts could be either picked up in person at the project field office or deposited automatically into the individual’s newly-opened bank account. In practice, 77 percent of individuals elected to have their payouts deposited into a bank account, even though accessing these funds would require payment of the Ksh 62 withdrawal fee (there was no fee to pick up funds at the field office, although respondents only had a one-month window to claim their cash). This suggests that most individuals saw some value to the bank accounts and intended to use them for saving.

2.3 Data and Randomization Verification

My analysis uses data from four sources. The first is a baseline survey conducted during the experimental sessions. The baseline collected basic demographic information, information on rates of time preference, and data on income and use of several popular savings devices. Second, I use three years of administrative data from the bank to get an accurate measure of short- and long-run use of the experimental accounts. This administrative data includes the date and amount of all transactions posted to experimental accounts. Finally, I use data from two waves of endline surveys. The first wave was conducted between August and November of 2012, approximately three years after the initial experiment. The wave 1 endline collected detailed information about respondents’ financial lives: in addition to basic demographic information it asked about income, savings, and debt by source as well as financial transfers and household decision-making. The wave 2 endline was conducted between July and August of 2013 and was much shorter. This wave was informed by results from the wave 1 endline and was explicitly designed to collect additional detail on study participants’ budgets and savings attitudes. The Data Appendix gives more detail on the data sources and how key variables used in the analysis were constructed.

⁷All cards, regardless of the interest rate, featured the following message in Swahili: “keep this card as a reminder to use your new Family Bank account to build your savings for the future.”

⁸Individuals were also selected for information sharing treatments and free ATM card treatments at this stage. I do not discuss these interventions further as they have no impact on the results in this paper.

⁹These amounts were designed to be substantial enough to ensure that individuals made choices carefully – for comparison the median weekly income at baseline was Ksh 500.

A total of 1,558 individuals (779 couples) opened 1,152 bank accounts during the initial experiment. The survey team was able to re-interview 1,417 (91 percent) of these individuals during the wave 1 endline and 1,331 (85 percent) of these individuals during the wave 2 endline. Table 1 presents baseline demographic characteristics and verifies that these characteristics are not systematically correlated with the interest rates. The first column of the table shows means and standard deviations of the variables of interest. The next three columns present the coefficients and standard errors from regressions of individual characteristics on the treatment of interest (either the individual interest rate, the joint interest rate, or the cash prize). For ease of interpretation I divide the interest rate variables by 20 before running regressions – thus, a change from 0 to 1 can be interpreted as the effect of moving from no interest rate to a 20 percent interest rate. This convention is maintained for the rest of the paper. The final row of the table presents p-values from chi-squared tests that the coefficient on the treatment is jointly equal to zero across all outcomes.

The first two rows of Table 1 verify that follow-up is uncorrelated with treatment status.¹⁰ The bulk of my analysis relies on the wave 1 endline – I therefore limit my attention to the 1,417 individuals with endline survey data for the remainder of this table. Overall, participants have relatively low levels of human capital, income, and financial access. While three quarters of individuals are literate, average educational attainment is low, at 6.79 years. The most common occupations are subsistence farming and small-scale entrepreneurship, each accounting for 42 percent of respondents.¹¹ Individuals reported an average income of Ksh 4,595 (\$57) per month, but the median income is much lower, at Ksh 2,167 (\$27).¹²

Just 22 percent of respondents reported owning a bank account at baseline. However, nearly all individuals reported saving in some way, with the most popular methods being storing cash at home (87 percent of respondents) and saving with a rotating savings and credit association, or ROSCA (58 percent of respondents). Individuals report saving roughly Ksh 1,500 in the bank and in savings and credit cooperatives (SACCOs), and keep an average of Ksh 849 at home.

Columns 2-4 of Table 1 show that the randomization functioned well, with none of the joint tests rejecting the null of no relationship at conventional levels. Some individual characteristics are correlated with the treatments, however – participants who received higher joint interest rates are slightly older, more likely to be polygamous, have more children,

¹⁰Moreover, there is no evidence of selective attrition – see Appendix Table A2.

¹¹I define entrepreneurship to include individuals operating an independent business. Common examples include market vendors, bicycle taxi drivers, shop owners, and commercial farmers.

¹²Income, savings, and debt measures in my sample are all highly skewed – I therefore topcode all variables denominated in Kenyan Shillings, in baseline, endline, and administrative bank account data at the 99th percentile.

and are more likely to participate in a ROSCA. Cash prize recipients are less likely to be subsistence farmers, have lower monthly incomes, and less SACCO savings. Since none of the joint tests reject and the randomization was not stratified, I do not control for any baseline characteristics in the main analysis. All the results are, however, essentially unchanged when controlling for all variables in Table 1.

3 Main Results

3.1 Impacts on Experimental Bank Account Use

Figure 1 illustrates the impact of the individual and joint interest rates on experimental account over time. Panel A graphs the share of individual (A.i) or joint (A.ii) accounts that received at least one deposit within a given quarter following account opening. Panel B performs the same exercise for withdrawals. Note that since the interest rate randomization was unconditional on account opening, unopened accounts are always kept in the sample and coded like unused accounts – this convention is held throughout the paper. Lighter lines correspond to higher temporary interest rates.¹³ There are three noteworthy patterns apparent in the graphs: First, rates of deposits fall off rapidly after the first quarter, which suggests that many individuals who experimented with the accounts failed to establish a regular savings practice. Second, both individual and joint accounts with higher interest rates were more likely to see deposits and withdrawals in the short run. Third, the individual interest rate had a persistent long-run effect on rates of account use. In contrast, the impact of the joint interest rate appears to dissipate after the first nine months.

Table 2 takes a more detailed look at treatment effects on experimental accounts. Panel A studies participants’ use of individual accounts. The underlying regression equation is as follows:

$$y_{ic} = \beta_0 + \beta_1 intI_{ic} + \beta_2 intI_{-ic} + \beta_3 intJ_c + \beta_4 cash_{ic} + \beta_5 cash_{-ic} + \varepsilon_{ac} \quad (1)$$

where y_{ic} is use of account i owned by couple c , $intI_{ic}$ is the interest rate on that account, $intI_{-ic}$ is the interest rate on the couple’s other individual account (owned by spouse $-i$), $intJ_c$ is the interest rate on the couple’s joint account, $cash_{ic}$ is a dummy variable equal to 1 if individual i was selected to receive a cash prize, and $cash_{-ic}$ indicates that i ’s spouse was selected to receive a cash prize. All standard errors are clustered at the couple level and

¹³In order to give a picture of account use absent the cash prize, I drop accounts randomly selected to be eligible for a cash prize when performing calculations for the first quarter. Results are very similar, but first quarter usage rates are higher, when including these accounts.

all variables denominated in Kenyan Shillings are top-coded at the 99th percentile. This specification lets me study direct own-account effects (β_1 , analogous to patterns in Figure 1), the extent to which spousal and joint interest rates crowd out own-individual-account use (β_2 and β_3), and the effect of own and spousal cash prizes on individual account use.

Panel A, column 1 shows that moving from no individual interest to 20 percent interest increased the share of open individual accounts by 17.4 percentage points (56 percent), increased the likelihood that the account would receive at least one deposit in the first 6 months by 9 percentage points (150 percent), and more than quadrupled total deposits over the course of the first 6 months (a treatment effect of Ksh 625, or \$7.81). The individual interest rate also significantly increased withdrawals, so effects on short-run ending and average daily balances are more modest, at just over Ksh 100 (\$1.25), a tripling of the mean for the zero-interest group.¹⁴ It is important to keep in mind, however, that just 16 percent of individual accounts in the 20 percent interest group received any deposits in the first 6 months. Hence, treatment effects for compliers are necessarily much larger. Finally, column 7 shows that accounts with the highest individual interest rate were 4.2 percentage points more likely to be used in the third year following account opening.¹⁵

Panel A also shows some (imprecise) evidence that other interest rates crowded out use of individual accounts: higher spousal and joint interest rates are consistently associated with lower rates of individual account use, though point estimates are only sporadically significant. Finally, Panel A illustrates that the cash prize had large impacts on rates of individual account use, with no crowd out from the spousal cash prize.

Panel B presents analogous treatment effects for joint accounts. Since there are two individual interest rates for each joint rate, I impose a restriction that both individual rates have the same impact on joint account use. I also impose a restriction that both cash prizes have the same impact on joint account use. These choices streamline the analysis and seem reasonable, as the restrictions are never rejected in practice (p-values range from 0.23-0.97). I implement this via the following regression equation:

$$y_{Jc} = \gamma_0 + \gamma_1 (intIH_c + intIW_c) + \gamma_2 intJ_c + \gamma_3 (cashH_c + cashW_c) + \varepsilon_{Jc} \quad (2)$$

where y_{Jc} is a measure of joint account use in couple c , $intIH_c$ and $intIW_c$ are the husband and wife's individual interest rates, and $cashH_c$ and $cashW_c$ are dummy variables indicating that the husband and wife were selected for a cash prize. Here, I find that the joint interest

¹⁴The treatment effect on deposits less the treatment effect on withdrawals does not add up to the treatment effect on ending balances due to fees and top-coding.

¹⁵I code an account as "used" if it received either a deposit or a withdrawal during the relevant time period.

rate significantly increased the probability that joint accounts were opened and used, though effects on actual balances and long-run use are only marginally significant. Panel B also shows that joint accounts were less likely to be opened when couples drew higher individual interest rates.

Although the crowd out estimates in Panels A and B are not robustly significant, they are important in practice. Panel C makes this clear by calculating treatment effects on couples' overall use of experimental accounts – regressions follow equation 2. Figure 2 collects point estimates from Panels A-C of Table 2 to visually depict crowd out. Figure 2 Panel A.i shows that going from no individual interest to 20 percent interest increased short-run total deposits in the treated individual account by Ksh 625, but reduced deposits in the spousal individual account by Ksh 193 and reduced joint account deposits by Ksh 263, leading to a small effect on the couples' total deposits across all experimental accounts. The net effect of the joint interest rate on total deposits is similarly small. Patterns for the average daily balance (Panel B), are similar. Overall, the figure underscores that the first order effect of the interest subsidies was to change *where* couples saved and, by extension, who within the couple did the saving. That said, I do find evidence that couples who received higher absolute interest rates saved more: Appendix Table A3 shows that the maximum interest rate available to the couple significantly increased all measures of experimental account use.

In contrast, I have very clear evidence that the cash prize increased overall use of experimental accounts. Panel C of Table 2 shows that the cash prize increased total deposits into experimental accounts by Ksh 904 (notably larger than the average cash payout of Ksh 247) and increased the average daily balance held across all experimental accounts by Ksh 218.

Finally, Table 2 and Figure 1 make it clear that the vast majority of study participants abandoned the experimental accounts in the long-run – just three percent of individual accounts (8 percent of individual accounts that were actually opened) and five percent of joint accounts (8 percent of opened accounts) were used in their third year. This, however, does not necessarily imply that most study participants abandoned formal banking altogether. Given the rapid change in Kenyan financial services over my study period, study participants may have transitioned to other formal banking products over time. The next sub-section turns to the two endlines to study impacts on overall bank account use and broader economic outcomes in the long run.

3.2 Impacts on Long-Run Economic Outcomes

The main objective of this paper is to study how the treatments impacted savings, income, debt, and intra-household dynamics after the interest rates expired. To do this, I make use of

data from the wave 1 endline, which was conducted three years after the initial experiment, or 2.5 years after the interest rates expired, and the wave 2 endline, conducted 3.5 years after the interest rates expired. All regressions are run at the individual level, following equation 1. This specification is useful because it allows me to study both own, spillover, and overall household-level effects of the individual interest rate and the cash prize. To that end, the tables present p-values from an F-test of the null hypothesis $\beta_1 + \beta_2 = 0$ – this lets me test whether the individual interest rate had a significant effect on average per-capita outcomes at the household level. In addition, I present p-values from F-tests of whether the net effect of the cash prize on household-level outcomes is equal to zero, i.e. $\beta_4 + \beta_5 = 0$, whether the individual and joint interest rates had equivalent impacts on household-level outcomes, i.e. $(\beta_1 + \beta_2) = \beta_3$, and whether the overall effect of the individual interest rate is equal to the overall effect of the cash prize, i.e. $(\beta_1 + \beta_2) = (\beta_4 + \beta_5)$.

Impacts on Savings and Income Table 3 presents treatment effects on overall bank account use, as well as broader economic outcomes including total assets, total debt, total income (all from the wave 1 endline), and self-reported savings practices (from the wave 2 endline).¹⁶ Column 1 of Table 3 shows that study participants who received the highest interest rate on their own individual account are 9 percentage points more likely to report having an individual or joint bank account with any bank. There is also marginally significant evidence that the joint interest rate increased overall rates of bank ownership. Hence the individual and joint interest rates did not simply keep people more attached to experimental accounts; they also kept people engaged in the banking sector more broadly. In contrast, there is no evidence that the cash prize increased rates of bank account ownership, but some evidence that it increased bank balances among account owners.

What is even more striking are results in columns 3 through 6, which illustrate impacts on total assets, debt, assets net debt, and income.¹⁷ Here we see that the individual interest rate reallocated assets within the household – study participants who received the highest individual interest rate reported Ksh 5,994 (\$75) more in total assets, while their spouses reported a Ksh 5,871 *reduction* in assets. I do interpret this finding with caution, however, since the effect of the spousal interest rate is not robust to taking a log transformation of total assets (while the direct impact of the individual interest rate is very robust to this trans-

¹⁶Throughout the paper I use the Central Bank of Kenya’s consumer price index to deflate all monetary amounts into 2009 values for comparability with the baseline.

¹⁷Total assets include cash saved in banks, SACCOs, mobile money accounts, at home, and in ROSCAs, as well as business assets, livestock, and “other assets”. When calculating livestock savings I record the value of all livestock owned by the household and assign half to each member of the couple, as livestock is an inherently joint method of saving. Total income includes returns from agricultural activities, business profits, wage earnings, and “other income”.

formation). Moreover, the asset reallocation appears to have been productivity enhancing: study participants who received higher individual interest rates report Ksh 1,137 (\$14) more income per month, a 27 percent increase relative to the no interest comparison group. This impact, however, does *not* reflect a reallocation within the household – the impact of the spousal interest rate on income is positive, small in magnitude, and not significantly different from zero. As a result, I am able to formally reject that the individual interest rate had no impact on household income.¹⁸ In order to give a sense of distributional impacts, Figure 3 graphs the CDF of bank savings, total assets, and monthly income by the individual interest rate (Panel A) and the joint interest rate (Panel B). In order to improve legibility of the graphs, I collapse the interest rate into “high” (12-20 percent) versus low (0-4 percent). Panel A makes it clear that the impacts of the individual interest rate are not driven by one part of the distribution. Consistent with this, I find that the main results in Table 4 are robust to a variety of top-coding and trimming schemes (see Appendix Table R2).

The last two columns of Table 3 use data from the wave 2 endline to study how the individual interest rate impacted savings attitudes and practices. There is no impact on the share of respondents who state that “saving is a priority for me”. But agreeing with this statement may not signal much: even though the vast majority (over 80 percent) of respondents state that saving is a priority, just 29 percent of respondents receiving no individual interest say they actually save regularly. Notably, the individual interest rate increases this proportion by 7.6 percentage points (27 percent). This suggests that the individual interest rate helped respondents sustain better savings practices in both the short and the long run, although this may have been offset by a decline in savings effort on the part of the spouse. Overall, the individual interest rate had large, meaningful impacts on study participants’ economic lives, well after the interest rates expired. In contrast, I find no robust, consistent patterns with respect to the joint interest rate or the cash prize.

Impacts on Business Outcomes What types of income and assets are driving the aggregate changes in Table 3? Table 4 shows that the individual interest rate had large and important effects on business outcomes, while Appendix Tables A4 and A5 show that the individual interest rate had no detectable impact on any other type of income or asset.¹⁹ Given this, I interpret the (somewhat smaller) impacts I find on business profits and capital as the main driver of the overall effect the interest rate had on income and assets.

Table 4 shows that receiving the highest individual interest rate increased study partic-

¹⁸Although the p-value on the joint test falls just short of the 5 percent significance level, I am easily able to reject at the 1 percent level if I impose the restriction that the spousal interest rate had no impact on individual income, which is clearly supported by the data.

¹⁹Appendix Table A6 illustrates impacts on debt.

participants' probability of having any business capital or profits by 8 percentage points (a 28 percent increase versus the no interest group) and increased the probability that a study participant reported entrepreneurship as his or her main occupation by 11 percentage points (30 percent). Moreover, the individual interest rate increased business capital by Ksh 2,651 (\$33) and business profits by Ksh 548 (\$7) per month. Appendix Figure A2 graphs the CDF of business profits and assets by the individual and joint interest rates. Here I find that the individual interest rate shifted the distribution of business profits everywhere to the left, while impacts on business assets are more concentrated in the lower part of the distribution.

The last two columns of Table 4 use data from the wave 2 endline to ask whether observed changes in business outcomes are accompanied by changes in financial management. Here I make use of a budgeting module that was the primary focus of the second endline: this module first asked individuals if they made spending decisions with a pre-planned budget in mind. Respondents were then asked to list each of their budget items and indicate whether the budgeted amounts were flexible (i.e. would the respondent reduce spending on the item to meet an unexpected Ksh 1,000 expense). I term budget items that would *not* be reduced to meet an unexpected expense as “downwardly rigid”. Here, I find that the individual interest rate increased the share of respondents explicitly budgeting for business by 5 percentage points (35 percent). This effect is entirely driven by growth in downwardly rigid budgeting. Even though just 6.6 percent of individuals in the lowest individual interest rate group had downwardly-rigid business budgets – which implies that most small-scale entrepreneurs in my sample do not budget for business in this way – the individual interest rate treatment nearly doubled this share. This, coupled with my earlier finding that the individual interest rate increased the share of people who save regularly, suggests that the individual interest rate had a sustained impact on financial behaviors that translated in to meaningful growth in economic resources.

I also find some (much weaker) evidence of positive spillover effects on spouses: spouses of study participants receiving higher individual interest rates were slightly more likely to report entrepreneurship as their main occupation and consequently reported more business income. Overall, I reject that the individual interest rate had no impact on household level business outcomes for all variables except business capital. In contrast, estimated impacts for the joint interest rate and the cash prize are generally close to zero and insignificant. I am able to formally reject that the impact of the individual interest rate is equal to the joint interest rate for half the variables on Table 4. I also reject that the individual interest rate and the cash prize had equivalent impacts for every outcome except business capital.

My data also reveal a great deal of business churning – just 50 percent of individuals who reported entrepreneurship as their main occupation at baseline reported the same at endline,

while 23 percent of non-entrepreneurs transitioned into entrepreneurship. Appendix Table A7 explores heterogeneity in main effects by baseline occupation. Here we see that the impact of the individual interest rate is concentrated among participants who were entrepreneurs at baseline. Hence the individual interest rate helped existing entrepreneurs stay in business, but did not facilitate much business creation. My results also suggest that the individual interest rate helped existing businesses grow – in order for the extensive margin to entirely account for the treatment effects on capital and profits, the average business “kept alive” by the higher interest rate would need to have been around the 90th percentile of the business profits and assets distributions. These results mirror recent findings from a number of capital drop experiments and microfinance evaluations, which suggest that only individuals with some pre-existing entrepreneurial skill are able to benefit from interventions to grow self-employment income (Crépon et al. 2011; Angelucci et al. 2013 Banerjee et al. 2013; Field et al. 2013; Fafchamps et al. 2014).

Impacts on “Joint” Household Outcomes Although I find no significant evidence that the joint interest rate impacted overall income and assets, I do find that it impacted outcomes beyond joint bank account use: Table 5 uses wave 1 endline data to explore how the interest rates affected measures that capture investment in household public goods and levels of spousal alignment. The first column of the table focuses on livestock holdings. Livestock are arguably the most important class of assets held by study households – nearly all households (95 percent) own some sort of livestock, and their value accounts for half of total assets. Livestock are also inherently joint investments, as they are easily observed and accessed by both members of a couple. Although the joint interest rate did not significantly impact average livestock holdings (see Appendix Table A4), Table 5 shows that a specification using the inverse hyperbolic sine transformation reveals that the joint rate did substantially increase livestock holdings at the lower end of the distribution.²⁰

It is common practice in the study area for households to make periodic, incremental investments in their homes. Home improvement expenditures are also inherently joint, as all members of the household benefit from them. Column 2 of Table 5 shows that nearly half of individuals reported making some investment in home renovation in the past year, and moving from the lowest to the highest joint interest rate increased this share by 6 percentage points (recall that the joint interest rate variable runs from 0.2-1, so I multiply the point estimate in Table 5 by 0.8 to get this effect). There is also evidence of increased home

²⁰The inverse hyperbolic sine of x is given by $\ln\left(x + (x^2 + 1)^{.5}\right) \approx \ln(2x)$. Results are very similar if I use $\ln(x + 1)$ instead. The impact of the joint interest rate on livestock holdings is confirmed by quantile regressions on level values, which reveal significant, positive impacts on the 28th to the 70th quantiles.

investment based on actual home quality. Column 3 shows that individuals who received higher joint interest rates were more likely to live in a home with a permanent (i.e. iron sheet as opposed to thatch) roof at endline.

The next two columns of Table 5 test whether the joint interest rate increased spousal alignment over decision making: at endline, all individuals were asked to rate, on a scale of 0 to 10, how much they and their spouse agreed about consumption and savings decisions.²¹ Columns 4 and 5 show that individuals who received higher joint interest rates reported greater levels of agreement about both topics. In spite of this, I do not observe any significant impacts on self-reported joint decision making regarding savings (column 6).

Taken together, the results in Table 5 suggest that the joint interest rate moderately increased investment in “joint” assets while slightly increasing spousal alignment. I caveat, however, that these results are much more modest in magnitude and less statistically significant than my earlier results for the individual interest rate. As a result, I cannot formally reject that the individual and joint interest rates had equivalent impacts on the outcomes in Table 5.

Robustness Since my analysis asks how several treatments impact a range of economic outcomes, it is important to ask whether inferences are robust to corrections for multiple hypothesis testing. Appendix A provides further robustness checks, showing that there is no evidence that the results are driven by reporting bias, and showing that the main results are robust to a variety of top-coding, trimming, and imputation strategies.

Figure 4 graphs how standard p-values in my main analysis compare to sharpened q-values that control the false discovery rate (FDR).²² I follow Anderson (2008) and use the two-step procedure described by Benjamini et al. (2006) to calculate q-values. The adjustment includes p-values from regression coefficients in all the main tables in this paper (except the randomization verification in Table 1), as well as coefficients in Appendix Tables A4-A6, which show income, assets, and debt by source. I group the results into families by treatment, pooling across all outcomes. For legibility, Figure 4 only plots results for traditional p-values below 0.20. I only display p- and q-values for the long-run results (Tables 3-5, Appendix Tables A4-A6) to graphically highlight their robustness.

Panel A of Figure 4 shows that all coefficients on the individual interest rate initially significant at the 1 percent level or better remain significant at the 5 percent level after the FDR adjustment, while results initially significant at the 5 percent level remain significant at the 5 or 10 percent level. In contrast, none of the initially-significant results for the spousal

²¹Individuals who were no longer married were asked refer back to when they were married.

²²The false discovery rate refers to the share of rejected null hypotheses that are type I errors; the q-value is the lowest FDR at which a hypothesis would be rejected.

interest rate, joint interest rate, or own cash prize remain significant after the adjustment, and just two coefficients on the spousal cash prize remain marginally significant.²³ Given this, I interpret my results for the joint interest rate with caution and focus on the effect of the individual interest rate in what follows.

3.3 Are the Long-Run Estimates Reasonable in Magnitude?

The estimated coefficients on the individual interest rate in Tables 3 and 4 are quite large in magnitude, especially when compared to the much more modest short-run effects found in Table 2. The average interest payout on open individual accounts with 20 percent interest was just Ksh 42 (\$0.53), yet 2.5 years later this group reported Ksh 548 more per month in business profit. Are these two observations compatible with one another? In this sub-section, I take the point estimates at face value and ask if they are both internally consistent and consistent with other studies on returns to capital in the developing world.

I address the latter question first. The ratio of the treatment effect on business profits to the treatment effect on business capital implies a monthly return to capital of 20.7 percent.²⁴ Although this is high, it is in line with existing estimates in the literature, which range from 4 percent per month (Blattman et al. 2014) to 33 percent per month (McKenzie and Woodruff 2008). It is also worth noting that my results are strikingly consistent with other evidence from Africa – Udry and Anagol (2006) estimate a 17-25 percent monthly return to pineapple cultivation in Ghana, while Fafchamps et al. (2014) find a 21-29 percent monthly return to a 150 cedi in-kind grant to Ghanaian microentrepreneurs.²⁵

While the *relative* changes I observe in income and assets are consistent with existing evidence, one must also ask whether the *absolute* changes make sense given the observed impacts on short-run bank account use. To help put my results in perspective, consider the following thought experiment: Assume that study participants who received large individual interest rates accumulated savings over the course of the six-month subsidy period, some of which they subsequently invested in productive activities. Call this “new” investment (or initial treatment effect on capital) K_0 . Further assume that the monthly return to capital is r , and that all subsequent capital accumulation is driven by reinvesting some of the returns

²³Moreover, the two outcomes are conceptually unrelated: the spousal cash prize is associated with a decline in the share of respondents saying they save regularly in Table 3 and a decline in debt owed to family members in Appendix Table A6.

²⁴The wave 1 endline survey did not ask about labor supply and individuals were not asked to report profits net of the value of their time. The “return” to capital I present is therefore simply the change in monthly income divided by the change in assets or business capital. To the extent that capital accumulation led participants to increase labor supply, these estimates give an upper bound on the true marginal return to capital.

²⁵I do, however, find larger returns than those estimated by Blattman et al. (2014) in Uganda.

on K_0 . If the marginal propensity to reinvest is i , then after t months the treatment effect on capital stocks would be $K_0(1 + ir)^t$. I observe study participants roughly 32 months after the interest rate expired. Given this, a 0.5 marginal propensity to reinvest, and a 20.7 percent monthly return to capital, a Ksh 114 (\$1.43) initial investment would be needed to generate my treatment effects on business profit and capital. If the marginal propensity to reinvest were 1, an initial investment of just Ksh 6 would be needed. These numbers compare favorably to the short-run treatment effects on bank account use in Table 2, where I find that the individual interest rate increased deposits into the individual account by Ksh 625 and the average daily balance in individual accounts by Ksh 110. The hypothetical initial investment amounts are also modest relative to the average amount of cash study participants reported keeping at home at baseline (Ksh 849). Appendix Table A8 summarizes these calculations and existing estimates of returns to capital in the literature.

Another important issue relevant for interpreting magnitudes is the number of potential compliers: when the individual interest rate was 20 percent, 49 percent of study participants opened an individual account. Just 16 percent of study participants opened *and* used the account. If one assumes that only people who opened an individual account benefitted from the interest rate, then the absolute maximum fraction of compliers in the population would be 49 percent, while a more reasonable maximum fraction of compliers would be closer to 16 percent.²⁶ If 16 percent of individuals complied, the short-run investment among compliers needed to generate my treatment effects would be either Ksh 713 (\$9 – when $i = 0.5$) or Ksh 38 (\$0.48 – when $i = 1$), which is still quite reasonable. Taking these observations together, I argue that the long run treatment effects are reasonable *assuming* the interest rates generated sustained behavior change that allowed returns to compound over time. Given that the short-run changes in experimental bank account use are relatively modest, it is much more difficult to justify the magnitude of the long-run results without a compounding mechanism.

That said, I also caveat that my long-run point estimates are drawn from sampling distributions with high variances. As such, the confidence intervals contain a range of much more modest (and more extreme) values. For example, a 95 percent confidence interval on the impact of the individual interest rate on business capital is [413, 4889], while the 95 percent confidence interval on monthly business profits is [156, 940]. Therefore a more conservative way to interpret the results is that I am able to reject that the individual

²⁶It is possible that some participants who opened an individual account but did not save still benefitted – if, for example, individuals accumulated a stock of money at home earmarked for the bank account but never got around to actually making a deposit. Given the experimental context, this is plausible: most individuals had to commute into town to use the bank, so making small frequent deposits was less practical than making a smaller number of large deposits.

interest rate had only moderate impacts on income and assets. Given that the average open individual account with a 20 percent interest rate earned just Ksh 42 in interest, even this interpretation is compelling.

4 Discussion and Mechanisms

4.1 Mechanisms

The results for the individual interest rate present an apparent puzzle: although its short-run impacts were modest, the treatment increased monthly household income by Ksh 1,137 in the long run (or Ksh 1,635 in nominal 2012 Shillings). Given a 2012 PPP exchange rate of Ksh 37 per US dollar, this is almost enough extra income to support an additional household member at a poverty line of PPP\$1.90/day. Why did study participants need the inducement of the interest rate to take advantage of such a lucrative income-generation opportunity? The rest of this section discusses potential mechanisms that could rationalize this result.

Capital Stock-Based Mechanisms One salient feature of my results is that the individual interest rate generated long-run *divergence* between the high- and low-interest groups – differences in savings behavior were amplified over time rather than reduced.²⁷ One of the simplest explanations for my results is that they are driven by the short-run changes in the capital stock generated by the temporary interest rate. This could create divergence if, for example, entrepreneurs have non-convex production functions, which in turn generate multiple steady-state levels of output (Azariadis and Stachurski 2005). If the interest rate helped people save past a nonconvexity then income and capital levels could have continued to diverge after the interest rate expired.

One issue with the capital-stock hypothesis is that the individual interest rate did not increase overall bank savings in the short run, while the cash prize did (Table 2, Panel C). Yet it is the interest rate, and not the cash prize, that had persistent impacts on economic outcomes. Given the fact that the cash prize had no impact on long run outcomes, yet a larger impact on overall short-run balances in experimental accounts, mechanisms working solely through the household-level capital stock seem unlikely.²⁸

²⁷Although I do not have short-run data on income and overall assets, this pattern is directly evident for experimental bank balances. The individual interest rate increased 6-month average daily balances in experimental individual accounts by Ksh 110, 3-year average daily balances in experimental individual accounts by Ksh 194, and overall endline bank balances by Ksh 1060.

²⁸A capital stock mechanism at the household-level could still be relevant if the mean effects in Table 2 mask important distributional differences in treatment effects. Quantile regression results, not shown here,

However, the individual interest rate *did* increase the amount of money stored in individual accounts, thereby changing responsibility for (and ownership of) saving within the household. Moreover, the individual interest rate's short-run implied local average treatment effects on *individual* account balances exceed those of the cash prize; hence a capital stock-based explanation could still be relevant if intra-household control of resources matters for how those resources are invested, e.g. due to innovations in bargaining power or other intra-household constraints (Mazzocco 2007; Ashraf 2009; Schaner 2015). To think through the plausibility of a capital-stock based mechanism more broadly, consider the utility of an individual in the experiment. I assume that this utility can be written as follows:

$$\begin{aligned}
V^i(A_t, R_t) &= \max_{c_t, b_t, k_t} u^i(c_t^i) + \delta V^i(A_{t+1}, R_{t+1}) \text{ subject to} \\
&\quad c_t' \mathbf{1} + b_t' \mathbf{1} + k_t' \mathbf{1} \leq A_t' \mathbf{1} \\
&\quad f(k_t) + R_t' b_t = A_{t+1}' \mathbf{1} \\
&\quad \text{intra-household constraints}
\end{aligned}$$

where $A_t = [A_t^i, A_t^{-i}, A_t^J]'$ represents individually and jointly owned assets available at the start of period t , $R_t = [R_t^i, R_t^{-i}, R_t^J]'$ are period t interest rates on bank accounts, $b_t = [b_t^i, b_t^{-i}, b_t^J]'$ is saving in each bank account, $k_t = [k_t^i, k_t^{-i}, k_t^J]'$ is individual and joint investment outside the bank (e.g. business capital or livestock), $f(k_t)$ is the household production function, $c_t = [c_t^i, c_t^{-i}]'$ is period t consumption, and $\mathbf{1}$ is a 3×1 vector of ones.

I do not place an explicit structure on the household bargaining process, in order to allow for the possibility that households are making decisions in a non-standard fashion. I only assume that individual i makes the best choice possible, subject to these constraints. To match the experiment, let each t interval be six months. In what follows I will contrast the impact of a shock to the individual interest rate (R_t^i) with the impact of a shock to individual assets (A_t^i) while keeping all other interest rates and assets levels fixed. To simplify notation, in what follows I will therefore write utility only in terms of these focal quantities – $V_t^i(A_t^i, R_t^i)$.

Let $V_t^i(A_t^i, \underline{R}_t^i)$ be utility when the period t individual interest rate is low and $V_t^i(A_t^i, \overline{R}_t^i)$ be utility when the period t interest rate is high. Next, note that if an individual were given

$$\gamma_t^i = \frac{1}{\underline{R}_t^i} \left(\overline{R}_t^i - \underline{R}_t^i \right) \overline{s}_t^i$$

in the first period (where \overline{s}_t^i is i 's optimal savings when the individual interest rate is \overline{R}_t^i),

suggest that this is not the case, however.

then the household could reproduce the high-interest allocation in the low-interest state of the world by simply increasing individual savings to \bar{s}_t^i . As long as the allocation under \bar{R}_t^i is available when the interest rate is low and spouse i is given γ_t^i (which seems reasonable in this case since spouse i has property rights over γ_t^i and savings in account i), then it must be that

$$V_t^i(A_t^i + \gamma_t^i, \underline{R}_t^i) \geq V_t^i(A_t^i, \bar{R}_t^i) > V_t^i(A_t^i, \underline{R}_t^i)$$

Consider the 16 percent of households who saved in the individual account when the individual interest rate was 20 percent – γ_t^i for this group is just Ksh 116 (\$1.45). Put another way, the above inequality implies that “complier” study participants would prefer an up-front payment of Ksh 116 to the higher interest rate. Recall that the cash prize, which had an average payout of Ksh 247, had no impact on long-run income. This in turn implies that study participants put little value (less than Ksh 247) on having substantially more income three years in the future, which would only be true if the discount factors and/or elasticity of intertemporal substitution are very low. Given that study participants *did* robustly respond to the temporary interest rates in the short run, this seems unlikely. I therefore argue that the combined facts of (a) a robust response to the short-run interest rates, (b) modest interest rate payouts, and (c) no long-run effects of the cash prize suggest that my results are difficult to explain with mechanisms that only work through the short-run capital stock.

Increased Use of Formal Bank Accounts Another possibility is that the individual interest rate worked by encouraging participants to adopt individual bank accounts, which in turn impacted business outcomes. Bank accounts may have been useful for savings-constrained individuals long after the interest rates expired.²⁹ For example, when production requires lumpy expenditures it could be helpful to store resources at the bank until it is time to purchase additional working capital. This could be particularly important if resources kept at home are “taxed” away by other members of the community or household (Platteau 2000; Anderson and Baland 2002; Jakiela and Ozier 2016).³⁰

Recall that all couples opened at least one bank account as part of the experiment and bank accounts were widely available on the market – the individual interest rate did not solve an access problem. However, the majority of study participants were unbanked at baseline – these individuals may not have totally understood bank accounts and their benefits. If

²⁹Dupas and Robinson (2013) find that giving female Kenyan entrepreneurs access to no-interest bank accounts had very large impacts on productive investment in the short-run (four to six months after account opening). This paper does not, however, study longer-run impacts.

³⁰Household constraints could also help explain why the individual and joint interest rates had different long-run impacts.

individuals learn about these benefits as they experience the accounts, temporary interest rates that increase experimentation could have long-run impacts. Appendix Table A9 tests this hypothesis by estimating the impact of the interest rates separately by baseline bank account ownership.³¹ The table shows that the individual interest rate benefitted banked respondents just as much as unbanked respondents, which suggests that the interest rate did not work by helping people learn about bank accounts.

Alternatively, increased bank account use in the short run may have facilitated future access to bank credit, which could have helped study participants build their businesses. However, Appendix Table A6 shows that recipients of higher individual interest rates actually report (insignificantly) *less* bank debt, and administrative data from the bank shows that just 1.5 percent of couples received a loan from Family Bank. It is therefore unlikely that the interest rates worked by connecting respondents with financial services that relieved savings or credit constraints.

Psychological Channels The final possibility I consider is that the individual interest rate operated through a (broadly defined) behavioral channel. The idea here is that actively responding to the interest rate in the short run may have had lasting impacts on psychological forces governing financial behavior. Consider, for example, a participant who wishes to take advantage of a 20 percent interest rate on his individual account. First, he must decide where to get deposits for the account: Will he reallocate his existing wealth? Or cut back on consumption? If the latter, he will have to decide what to cut back on and may need to devise new strategies to overcome temptation to spend money. Second, he must decide what to do with his new savings once the interest rate expires. Should he keep it in the bank for emergencies? Or invest the money in something productive, like his business or the family farm? The interest rate may have helped some individuals build better savings habits, or set up new savings and investment heuristics. If these changes persisted after the interest rates expired, then small short-run changes in behavior could compound over time, especially when returns to investment are high.

The idea behind Becker and Murphy (1988)'s model of habitual consumption is that individuals accumulate a "habit stock" as they engage in a particular behavior, and that the habit stock lowers the marginal cost of engaging in that behavior in future periods. This can give rise to multiple steady states, where price changes or other shocks can lead to dramatically different behavioral patterns in the long run. Recall that both the endline data on overall bank account use and administrative data on experimental account use shows that

³¹I focus on the inverse hyperbolic sine of income and assets because banked individuals reported substantially more resources at baseline and treatment effects scale roughly proportionally with baseline resources. Using level values does not change conclusions, however.

respondents who received higher individual interest rates were more likely to use individual bank accounts both during and after the promotional period. Moreover, 4 years after the initial experiment respondents treated with the highest individual interest rate were 7.6 percentage points (27 percent) more likely to say that they “saved regularly” (Table 3). This is consistent with the idea that the individual interest rate helped study participants build a stock of saving habits, which lowered the marginal cost of saving even after the interest rates expired. A related possibility is that the experience of saving (and investing) in the short-run helped individuals learn about their (high) marginal returns to capital, which in turn supported sustained savings and investment.

Mental accounting may have also played an important role in helping study participants sustain business investment after the interest subsidies expired. Thaler (1999) defines mental accounting very broadly to be “the set of cognitive operations used by individuals and households to organize, evaluate, and keep track of financial activities”. One aspect of mental accounting that is particularly relevant in my context is the idea that individuals mentally allocate income and wealth into a series of differentiated, non-fungible “accounts”, which help constrain financial decision-making. If the individual interest rate cued respondents to set up mental accounts for business, this could help explain why study participants were able to make continued investments in their businesses after the interest rates expired. My finding that the individual interest rate dramatically increased the share of respondents keeping downwardly rigid business budgets (Table 4) is consistent with a mental accounting hypothesis. Of course this evidence is only suggestive: the increased rates of business budgeting could simply be driven by higher rates of overall entrepreneurship. It is worth noting, however, that the majority of business owners do not keep downwardly rigid business accounts – hence, the fact that growth in business budgeting is driven entirely by this type of planning is striking.

Finally, the patterns in my data suggest that the nature of the interest rate may have been important for determining the type of savings and investments that households made. Small businesses are generally independently operated in this part of Kenya, while investments such as livestock and the home are inherently joint. Priming effects – where different contextual stimuli can lead individuals to make systematically different decisions – could have been particularly important in this context because all couples had the option to open both individual and joint accounts.³² Married couples attended the experiment together. They were therefore able to directly observe the interest rate for their own, their spouse’s, and the joint account. Thus, a particularly high joint rate could have easily cued the idea that “we need to start saving together”, while a particularly high individual interest rate could

³²See Kamenica (2012) for an economics-focused review of priming.

have made it more acceptable for the relevant account owner to redirect more resources towards his or her own enterprise. Indeed, anecdotal evidence from the experiment suggests that many couples saw the account opening choice as one of “going individual” versus “going joint” (recall that 55 percent of couples opened just a joint account and 30 percent of couples opened two individual accounts).

The wave 1 endline was specifically designed to shed light on the extent to which couples viewed income and assets as “individual” versus “joint”. For each type of income and asset, respondents were asked to report how much was individually earned/held versus jointly earned/shared with a spouse. In the previous tables, I have focused on overall measures of resources, which include individually-held resources plus half of jointly-held resources.³³ Table 6 breaks down my results for income and assets by individual and joint ownership. Here, we see that the impacts of the individual interest rate are entirely driven by growth in *individually-owned* income and assets. Although I observe no significant impacts for the joint interest rate, the coefficient on joint assets is meaningful in magnitude and just shy of marginal significance with a p-value of 0.125. Thus, one interpretation of my results is that the individual interest rate spurred participants to invest in high-return individual enterprises (hence leading to income growth), while the joint interest rate spurred participants to invest in lower-return jointly-held assets.

Summary Although I cannot pin down a single mechanism that can explain my findings, I am able to generate suggestive evidence that behavioral savings channels are important in my context. My preferred interpretation is that two forces interacted to generate my results: First, higher individual interest rates pushed couples to consider higher-return “individual” forms of investment beyond the bank accounts, while higher joint interest rates encouraged couples to save for lower-return “public” investments. Second, the experience of making a concerted effort to save in the short run helped individuals sustain greater levels of saving and investment in the long run. Here, a mix of factors could be at play, including habit formation, mental accounting, and learning about the return to saving.

5 Conclusion

I show that large short-term incentives to save have long-run implications for the economic lives of low-income, rural Kenyans. Moreover, it is not just the absolute value of the incentive that matters, but how it is delivered: incentivizing study participants to save in individually owned bank accounts increased rates of entrepreneurship and overall household income, while

³³This way, the average individual-level report of, e.g., business profit, reflects per-capita business profit.

incentivizing use of joint bank accounts increased investment in livestock and the home. In contrast, simply delivering an unconditional payment in the form of a modest cash prize had no impact on long-run outcomes. The magnitudes of some of the estimates for the individual interest rate are quite large – but given the high (though reasonable, given the literature) returns to capital in my sample, the effect sizes can be accounted for by simply assuming that people continuously reinvest a portion of the returns on a modest initial capital investment.

The main contribution of this paper is to demonstrate that (appropriately targeted) temporary financial incentives to save can lead to sustained behavior change that has long-run impacts on broad economic outcomes. Moreover, the contrast between my results for the interest rate and the cash prize suggests that incentivizing individuals to accumulate capital on their own can be more effective than simply giving out (at least very modest amounts of) capital without condition.

The cash prize results also suggest that the interest rates did not simply help people make a short-run push to save past a nonconvexity in the production function. Moreover, my results are inconsistent with the hypothesis that the interest rates helped respondents by exposing them to bank accounts or giving them access to credit. In contrast, I argue that my results are better reconciled by behavioral channels. The caveat here is that my findings are largely suggestive, so I cannot precisely pin down the mechanism by which the individual interest rate led to long-run growth in savings and entrepreneurship. Additional research would be needed to shed further light on the precise behavioral channels, such as priming, habit formation and mental accounting, that likely account for my results.

It is also important to ask *what*, exactly, about the interest rate treatments stimulated behavior change, especially since other studies that consider more market-oriented ranges of interest rates (up to 5 percent per year), find little-to-no impact on savings behavior (Karlan and Zinman 2013; Kast et al. 2013). There are several potential explanations for this difference: First, unlike those studies this paper evaluates the impact of interest rates that were very large (and well above market rates) both in nominal and real terms. Thus, high-powered incentives may be needed to bring about lasting behavior change. Second, the interest rates in my experiment were time-limited. The temporary nature of the treatments may have given individuals a deadline that helped them break a cycle of savings procrastination. Third, individuals in my study were given the opportunity to open and save in multiple bank accounts. As a result, a large part of the experiment’s impact on bank account use was driven by switching *between* accounts (i.e. individual versus joint). Simply subsidizing one type of account (e.g. as a promotion from a bank), could have different psychological effects, especially if “buy-in” from the spouse is needed to make real change, or if individuals make different choices in private as opposed to in full view of their spouse (Ashraf 2009).

Since my experiment bundled all the features listed above, I cannot speculate as to which features were most important for generating persistent behavior change. Further research focused on more scalable and standard short-term savings subsidy programs, would help fill this evidence gap. However, the simple fact that I reject that the interest rates had no impact offers important evidence that savings behaviors can be sticky, and that activating the “right” set of short-run behavioral changes can have important, persistent effects on overall income and asset holdings, especially when individuals have access to high-return investment opportunities. This result is striking, and stands in contrast to studies of temporary incentives in other areas, such as health. Thus, a key takeaway from my results is that some individuals have the capacity to improve their economic situation without sustained external assistance or a large, expensive “big push” intervention – rather, they simply need the “right push”.

References

- Acland, D. and M. Levy (2015). Naiveté, Projection Bias, and Habit-Formation in Gym Attendance. *Management Science* 61(1), 146–160.
- Aleem, I. (1990). Imperfect Information, Screening, and the Costs of Informal Lending: A Study of a Rural Credit Market in Pakistan. *The World Bank Economic Review* 4(3), 329–349.
- Ananth, B., D. Karlan, and S. Mullainathan (2007). Microentrepreneurs and Their Money: Three Anomalies.
- Anderson, M. L. (2008). Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects. *Journal of the American statistical Association* 103(484).
- Anderson, S. and J.-M. Baland (2002). The Economics of Roscas and Intrahousehold Resource Allocation. *Quarterly Journal of Economics* 117(3), 963–995.
- Angelucci, M., D. Karlan, and J. Zinman (2013). Win Some Lose Some? Evidence from a Randomized Microcredit Program Placement Experiment by Compartamos Banco. Dartmouth Working Paper.
- Ashraf, N. (2009). Spousal Control and Intra-household Decision Making: An Experimental Study in the Philippines. *American Economic Review* 99(4), 1245–1277.
- Ashraf, N., D. Karlan, and W. Yin (2006a). Deposit Collectors. *Advances in Economic Analysis and Policy* 6(2).
- Ashraf, N., D. Karlan, and W. Yin (2006b). Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines. *Quarterly Journal of Economics* 121(2), 635–672.
- Azariadis, C. and J. Stachurski (2005). *Handbook of Economic Growth*, Chapter 5, Poverty Traps. North-Holland.
- Bandiera, O., R. Burgess, N. Das, S. Gulesci, I. Rasul, and M. Sulaiman (2015). The Misallocation of Labor in Village Economies. Technical report, Suntory and Toyota International Centres for Economics and Related Disciplines, LSE.
- Banerjee, A. and E. Duflo (2012). Do Firms Want to Borrow More? Testing Credit Constraints Using a Directed Lending Program. MIT Working Paper.
- Banerjee, A., E. Duflo, R. Glennerster, and C. Kinnan (2013). The Miracle of Microfinance? Evidence From a Randomized Evaluation. MIT Working Paper.

- Banerjee, A., E. Duflo, N. Goldberg, D. Karlan, R. Osei, W. Parienté, J. Shapiro, B. Thuysbaert, and C. Udry (2015). A multifaceted program causes lasting progress for the very poor: Evidence from six countries. *Science* 348(6236), 1260799.
- Banerjee, A. V. and E. Duflo (Winter 2007). The Economic Lives of the Poor. *The Journal of Economic Perspectives* 21(27), 141–167.
- Becker, G. S. and K. M. Murphy (1988). A Theory of Rational Addiction. *The Journal of Political Economy*, 675–700.
- Benjamini, Y., A. M. Krieger, and D. Yekutieli (2006). Adaptive Linear Step-Up Procedures That Control the False Discovery Rate. *Biometrika* 93(3), 491–507.
- Blattman, C., N. Fiala, and S. Martinez (2014). Generating Skilled Self-Employment in Developing Countries: Experimental Evidence from Uganda. *The Quarterly Journal of Economics* 129(2), 697–752.
- Blattman, C., E. P. Green, J. C. Jamison, M. C. Lehmann, and J. Annan (forthcoming). The Returns to Microenterprise Support among the Ultra-Poor: A Field Experiment in Post-War Uganda. *American Economic Journal: Applied Economics*.
- Breza, E. and A. G. Chandrasekhar (2015). Social Networks, Reputation and Commitment: Evidence from a Savings Monitors Experiment. Technical report, National Bureau of Economic Research.
- Brune, L., X. Gine, J. Goldberg, and D. Yang (2013). Commitments to Save: A Field Experiment in Rural Malawi. Mimeo.
- Callen, M., S. De Mel, C. McIntosh, and C. Woodruff (2014). What are the Headwaters of Formal Savings? Experimental Evidence from Sri Lanka. Technical report, National Bureau of Economic Research.
- Charness, G. and U. Gneezy (2009). Incentives to Exercise. *Econometrica* 77(3), 909–931.
- Cole, S., T. Sampson, and B. Zia (2011). Prices or Knowledge? What Drives Demand for Financial Services in Emerging Markets? *The Journal of Finance* 66(6), 1933–1967.
- Cruz, M., J. Foster, B. Quillin, and P. Schellekens (2015). Ending Extreme Poverty and Sharing Prosperity: Progress and Policies. Policy research note, The World Bank.
- Crépon, B., E. Duflo, F. Devoto, and W. Pariente (2011). Impact of Microcredit in Rural Areas of Morocco: Evidence from a Randomized Evaluation. MIT Working Paper.
- De Mel, S., C. McIntosh, and C. Woodruff (2013). Deposit Collecting: Unbundling the Role of Frequency, Salience, and Habit Formation in Generating Savings. *The American Economic Review* 103(3), 387–392.

- de Mel, S., D. McKenzie, and C. Woodruff (2008). Returns to Capital in Microenterprises: Evidence from a Field Experiment. *The Quarterly Journal of Economics* 123(4), 1329–1372.
- de Mel, S., D. McKenzie, and C. Woodruff (2012). One-Time Transfers of Cash or Capital Have Long-Lasting Effects on Microenterprises in Sri Lanka. *Science* 335(6071), 962–966.
- Dupas, P. and J. Robinson (2013). Savings Constraints and Microenterprise Development: Evidence from a Field Experiment in Kenya. *American Economic Journal: Applied Economics* 5(1), 163–192.
- Dupas, P. and J. Robinson (2014). Why Don't the Poor Save More? Evidence from Health Savings Experiments. *American Economic Review* 103(4), 1138–1171.
- Fafchamps, M., D. McKenzie, S. R. Quinn, and C. Woodruff (2014). Microenterprise Growth and the Fly-paper Effect: Evidence from a Randomized Experiment in Ghana. *Journal of Development Economics* 106, 211–226.
- Field, E., R. Pande, J. Papp, and N. Rigol (2013). Does the Classic Microfinance Model Discourage Entrepreneurship Among the Poor? Experimental Evidence from India. *American Economic Review* 103(6), 2196–2226.
- Gertler, P. J., S. W. Martinez, and M. Rubio-Codina (2012). Investing Cash Transfers to Raise Long-Term Living Standards. *American Economic Journal: Applied Economics* 4(1), 164–92.
- Giné, X., D. Karlan, and J. Zinman (2010). Put Your Money Where Your Butt is: A Commitment Contract for Smoking Cessation. *American Economic Journal: Applied Economics*, 213–235.
- Jackson, C. K. (2010). A Little Now for a Lot Later: A Look at a Texas Advanced Placement Incentive Program. *Journal of Human Resources* 45(3), 591–639.
- Jakiela, P. and O. Ozier (2016). Does Africa Need a Rotten Kin Theorem? Experimental Evidence from Village Economics. *Review of Economic Studies* 83(1), 231–268. Mimeo.
- John, L. K., G. Loewenstein, A. B. Troxel, L. Norton, J. E. Fassbender, and K. G. Volpp (2011). Financial Incentives for Extended Weight Loss: A Randomized, Controlled trial. *Journal of General Internal Medicine* 26(6), 621–626.
- Kamenica, E. (2012). Behavioral Economics and Psychology of Incentives. *Annual Review of Economics* 4(1), 427–452.

- Karlan, D., M. McConnell, S. Mullainathan, and J. Zinman (2013). Getting to the Top of Mind: How Reminders Increase Saving.
- Karlan, D., A. L. Ratan, and J. Zinman (2014). Savings by and for the Poor: A Research Review and Agenda. *Review of Income and Wealth* 60(1), 36–78.
- Karlan, D. and J. Zinman (2013). Price and Control Elasticities of Demand for Savings. Yale Working Paper.
- Karlan, D. and J. Zinman (2014). Price and Control Elasticities of Demand for Savings. Technical report, Working Paper.
- Kast, F., S. Meier, and D. Pomeranz (2013). Under-Savers Anonymous: Evidence on Self-Help Groups and Peer Pressure as a Savings Commitment Device. NBER Working Paper No. 18417.
- Macours, K., P. Premand, and R. Vakis (2012). Transfers, diversification and household risk strategies: experimental evidence with lessons for climate change adaptation. *World Bank Policy Research Working Paper* (6053).
- Mazzocco, M. (2007). Household Intertemporal Behaviour: A Collective Characterization and a Test of Commitment. *Review of Economic Studies* 74(3), 857 – 895.
- McKenzie, D. and C. Woodruff (2008). Experimental Evidence on Returns to Capital and Access to Finance in Mexico. *World Bank Economic Review* 22(3), 457–482.
- McKenzie, D. and C. Woodruff (2014). What are We Learning From Business Training and Entrepreneurship Evaluations Around the Developing World? *The World Bank Research Observer* 29(1), 48–82.
- Platteau, J.-P. (2000). *Institutions, Social Norms, and Economic Development*. Amsterdam: Harwood Academic.
- Prina, S. (2015). Banking the Poor via Savings Accounts: Evidence From a Field Experiment. *Journal of Development Economics* 115, 16–31. Mimeo, Case Western Reserve University.
- Royer, H., M. F. Stehr, and J. R. Sydnor (2012, November). Incentives, Commitments and Habit Formation in Exercise: Evidence from a Field Experiment with Workers at a Fortune-500 Company. Working Paper 18580, National Bureau of Economic Research.
- Schaner, S. (2015). Do Opposites Detract? Intrahousehold Preference Heterogeneity and Inefficient Strategic Savings. *American Economic Journal: Applied Economics* 7(2), 135–174.

- Thaler, R. H. (1999). Mental Accounting Matters. *Journal of Behavioral Decision Making* 12, 183–206.
- Udry, C. and S. Anagol (2006). The Return to Capital in Ghana. Yale Working Paper.
- Volpp, K. G., A. B. Troxel, M. V. Pauly, H. A. Glick, A. Puig, D. A. Asch, R. Galvin, J. Zhu, F. Wan, J. DeGuzman, et al. (2009). A Randomized, Controlled Trial of Financial Incentives for Smoking Cessation. *New England Journal of Medicine* 360(7), 699–709.

Table 1. Attrition and Randomization Verification

Variable	(1)	(2)	(3)	(4)	(5)
	Mean	Interest Rates		Cash	N
		Individual	Joint	Prize	
No Follow Up: Wave 1	0.091 [0.287]	-0.021 (0.019)	-0.008 (0.024)	0.015 (0.019)	1558
No Follow Up: Wave 2	0.146 [0.353]	-0.033 (0.024)	-0.010 (0.030)	0.021 (0.023)	1558
Female	0.500 [0.500]	0.025 (0.035)	-0.003 (0.010)	0.000 (0.035)	1417
Age	40.6 [13.4]	-1.09 (0.926)	2.37* (1.37)	0.561 (0.920)	1417
Education (Years)	6.79 [3.94]	0.421 (0.276)	0.171 (0.370)	-0.070 (0.266)	1412
Literate	0.749 [0.434]	0.023 (0.030)	-0.010 (0.036)	-0.001 (0.030)	1417
Number Children	5.29 [3.46]	-0.020 (0.225)	0.710** (0.338)	-0.055 (0.220)	1414
Polygamous	0.235 [0.424]	-0.024 (0.028)	0.095* (0.049)	0.041 (0.031)	1408
Subsistence Farmer	0.424 [0.494]	0.003 (0.035)	-0.003 (0.046)	-0.066** (0.033)	1412
Entrepreneur	0.418 [0.493]	0.000 (0.035)	0.011 (0.044)	0.043 (0.034)	1412
Monthly Income	4595 [7524]	-198 (549)	-126 (669)	-823* (476)	1374
Has Bank Account	0.217 [0.413]	0.032 (0.029)	0.029 (0.036)	-0.006 (0.029)	1417
Has SACCO Account	0.035 [0.185]	0.019 (0.013)	0.007 (0.016)	-0.005 (0.012)	1413
ROSCA Participant	0.580 [0.494]	0.013 (0.034)	0.079* (0.042)	0.027 (0.034)	1417
Saves at Home	0.873 [0.333]	0.030 (0.023)	-0.044 (0.027)	0.029 (0.021)	1415
Bank Savings	1508 [5456]	550 (406)	-556 (461)	-18.1 (404)	1365
SACCO Savings	1437 [10831]	281 (753)	611 (961)	-1006* (550)	1411
Home Savings	849 [1730]	89.1 (121)	-165 (147)	-63.4 (117)	1385
I Mostly Save	0.427 [0.495]	-0.014 (0.035)	-0.031 (0.039)	0.036 (0.034)	1411
My Spouse Mostly Saves	0.353 [0.478]	-0.002 (0.034)	0.008 (0.038)	-0.037 (0.033)	1411
Impatient Now-Patient Later	0.211 [0.408]	0.010 (0.028)	-0.031 (0.033)	-0.051* (0.026)	1400
Patient Now-Impatient Later	0.294 [0.456]	-0.028 (0.031)	-0.003 (0.039)	0.018 (0.031)	1400
Weekly Discount Factor	0.756 [0.244]	-0.013 (0.017)	-0.028 (0.020)	-0.012 (0.018)	1417
Distance to Bank (Miles)	3.87 [2.15]	-0.167 (0.148)	0.156 (0.251)	-0.027 (0.140)	1417
P-value - Joint Test		0.576	0.468	0.200	

Notes: Standard deviations in brackets, robust standard errors clustered at the couple level in parentheses. All income and savings variables top-coded at the 99th percentile. In 2009 Ksh 80≈ \$1. The joint test is an F-test of the null hypotheses that the coefficients on the treatment variable across all equations/outcomes are jointly equal to zero. The individual interest rate is renormalized to run from 0-1, while the joint interest rate is renormalized to run from 0.2-1.

Table 2. Impacts on Experimental Bank Account Use

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	In First 6 Months:						
	Opened	Used	Total Deposits	Total Withdrawals	Ending Balance	Avg. Daily Balance	Used - Year 3
<i>Panel A. Individual Accounts</i>							
Individual Interest	0.174*** (0.033)	0.090*** (0.019)	625*** (194)	475*** (160)	105*** (33.8)	110*** (30.4)	0.042*** (0.013)
Spousal Interest	0.055* (0.033)	0.016 (0.018)	-193 (162)	-127 (132)	-57.4* (30.3)	-48.7* (25.4)	0.002 (0.012)
Joint Interest	-0.148*** (0.050)	-0.013 (0.026)	-254 (206)	-227 (172)	6.75 (34.8)	-12.1 (34.1)	-0.033** (0.015)
Cash Prize	0.029 (0.033)	0.267*** (0.030)	547*** (221)	408** (187)	98.7*** (35.8)	130*** (38.2)	0.025* (0.014)
Cash Prize - Spouse	0.016 (0.033)	0.049** (0.023)	205 (194)	126 (157)	48.5 (34.6)	50.0 (32.0)	0.008 (0.012)
DV Mean (0% Ind)	0.308	0.060	206	133	40.7	35.3	0.020
N	1558	1558	1558	1558	1558	1558	1558
<i>Panel B. Joint Accounts</i>							
Individual Interest	-0.085*** (0.031)	-0.024 (0.028)	-263 (237)	-212 (196)	-41.0 (50.5)	-58.9 (40.3)	0.002 (0.015)
Joint Interest	0.186*** (0.051)	0.131*** (0.048)	409 (457)	340 (363)	167* (91.1)	132* (74.7)	0.051** (0.025)
Cash Prize	-0.048 (0.032)	0.271*** (0.034)	152 (274)	-22.9 (220)	70.3 (46.2)	63.5 (39.6)	-0.011 (0.013)
DV Mean (4% Joint)	0.598	0.213	897	544	177	161	0.032
N	779	779	779	779	779	779	779
<i>Panel C. Overall Use of Experimental Accounts</i>							
Individual Interest		0.053* (0.028)	169 (306)	136 (254)	6.14 (59.4)	2.55 (50.1)	0.042** (0.021)
Joint Interest		0.071 (0.048)	-99.1 (600)	-113 (490)	180 (115)	108 (99.4)	-0.006 (0.035)
Cash Prize		0.479*** (0.028)	904** (421)	511 (351)	218*** (72.0)	243*** (68.4)	0.021 (0.023)
DV Mean (4% Joint)		0.402	1970	1333	319	331	0.112
N		779	779	779	779	779	779

Notes: Heteroskedasticity robust standard errors, clustered at the couple level when relevant, in parentheses. The individual interest rate is renormalized to run from 0-1, while the joint interest rate is renormalized to run from 0.2-1. Variables denominated in Kenyan Shillings are top-coded at the 99th percentile. In 2009 Ksh 80≈\$1. An account is coded as used in the first 6 months if it received any deposit or withdrawal over that time period. An account is coded as used in year 3 if it received any deposit or withdrawal in the third year following account opening. ***, **, and * indicate significance at the 1, 5, and 10 percent levels respectively.

Table 3. Long-Run (3-Year) Impacts on Overall Economic Outcomes

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Wave 1 Endline				Wave 2 Endline			
	Has Bank Account	Bank Savings	Total Assets	Total Debt	Net Assets	Income Last Month	Savings a Priority	Saves Regularly
Individual Interest	0.089*** (0.032)	796* (439)	5994** (2642)	-41.6 (1634)	7673*** (2782)	1137*** (477)	0.036 (0.028)	0.076*** (0.032)
Spouse's Interest	0.019 (0.032)	-480 (451)	-5871*** (2499)	2026 (1590)	-6578*** (2788)	139 (463)	0.011 (0.028)	-0.003 (0.032)
Joint Interest	0.072* (0.039)	-498 (524)	5095 (3283)	4268* (2392)	-1300 (3594)	558 (542)	0.057* (0.033)	-0.004 (0.042)
Cash Prize - Self	0.037 (0.031)	1252** (584)	-2005 (2468)	-1208 (1636)	-2424 (2681)	228 (436)	0.039 (0.025)	-0.025 (0.033)
Cash Prize - Spouse	-0.010 (0.031)	313 (425)	-982 (2599)	1859 (2086)	-554 (2790)	544 (451)	-0.001 (0.026)	-0.094*** (0.031)
P-value: Ind+Spouse=0	0.033**	0.591	0.976	0.383	0.792	0.065*	0.268	0.129
P-value: (Ind+Spouse)=Joint	0.560	0.322	0.359	0.503	0.676	0.387	0.855	0.224
P-value: Cash Self+Spouse=0	0.550	0.034**	0.455	0.806	0.473	0.243	0.342	0.011**
P-value: Net Ind.=Net Cash	0.259	0.179	0.551	0.681	0.451	0.574	0.887	0.004***
DV Mean (0% Ind.)	0.664	1530	21913	7724	13579	4265	0.817	0.286
DV Mean (4% Joint)	0.675	2017	24029	5864	19105	4657	0.813	0.343
N	1413	1237	1053	1394	1039	1279	1317	1322

Notes: Standard errors clustered at the couple level in parentheses. The individual interest rate is renormalized to run from 0-1, while the joint interest rate is renormalized to run from 0.2-1. All variables denominated in Kenyan Shillings are top-coded at the 99th percentile and deflated to 2009 values. In 2009 Ksh 80≈\$1. ***, **, and * indicate significance at the 1, 5, and 10 percent levels respectively.

Table 4. Long-Run Impact of Interest Subsidies on Entrepreneurial Activity

	(1)	(2)	(3)	(4)	(5)	(6)
	Wave 1 Endline			Wave 2 Endline		
	Main Occupation Entrepreneur	Any Business Profits or Capital	Business Capital	Business Profits	Has Business Budget	Downwardly Rigid Business Budget
Individual Interest	0.107*** (0.037)	0.079*** (0.032)	2651** (1142)	548*** (200)	0.054* (0.029)	0.059*** (0.021)
Spouse's Interest	0.064* (0.037)	0.036 (0.032)	-371 (1049)	343* (204)	0.049 (0.030)	0.022 (0.022)
Joint Interest	0.015 (0.045)	0.023 (0.042)	820 (1340)	254 (222)	0.028 (0.036)	-0.011 (0.027)
Cash Prize - Self	0.002 (0.034)	0.036 (0.033)	-746 (1026)	-187 (161)	-0.031 (0.028)	-0.017 (0.021)
Cash Prize - Spouse	-0.017 (0.034)	-0.058* (0.031)	15.8 (1106)	-21.5 (177)	-0.057** (0.027)	-0.019 (0.021)
P-value: Ind+Spouse=0	0.002***	0.017**	0.166	0.003***	0.020**	0.013**
P-value: (Ind+Spouse)=Joint	0.023**	0.149	0.524	0.081*	0.197	0.033**
P-value: Cash Self+Spouse=0	0.774	0.653	0.633	0.433	0.030**	0.259
P-value: Net Ind.=Net Cash	0.017**	0.057*	0.163	0.005***	0.002***	0.008***
DV Mean (0% Ind.)	0.361	0.284	3263	796	0.155	0.066
DV Mean (4% Joint)	0.458	0.355	4575	1036	0.205	0.117
N	1339	1417	1380	1368	1305	1305

Notes: Standard errors clustered at the couple level in parentheses. The individual interest rate is renormalized to run from 0-1, while the joint interest rate is renormalized to run from 0.2-1. All variables denominated in Kenyan Shillings are top-coded at the 99th percentile and deflated to 2009 values. In 2009 Ksh 80≈\$1 ***, **, and * indicate significance at the 1 5, and 10 percent levels respectively.

Table 5. Impact of Interest Subsidies on Public Goods and Spousal Agreement

	(1)	(2)	(3)	(4)	(5)	(6)
	Value Livestock (Hypersine)	Any Renovations Last Year	Home Has Permanent Roof	Agreement - How to Spend Money	Agreement - How Much to Save	Savings Decision Making -Decide Together
Individual Interest	0.096 (0.195)	0.018 (0.035)	0.003 (0.031)	-0.112 (0.177)	-0.066 (0.208)	-0.057* (0.034)
Spouse's Interest	-0.048 (0.193)	0.020 (0.035)	0.011 (0.031)	0.243 (0.172)	0.075 (0.210)	0.020 (0.034)
Joint Interest	0.594** (0.259)	0.087* (0.045)	0.092* (0.050)	0.386* (0.216)	0.407* (0.243)	0.005 (0.042)
Cash Prize - Self	-0.157 (0.185)	-0.065* (0.035)	-0.067** (0.032)	0.140 (0.167)	-0.019 (0.206)	0.003 (0.035)
Cash Prize - Spouse	-0.270 (0.189)	-0.010 (0.035)	-0.041 (0.031)	0.262 (0.168)	-0.067 (0.211)	-0.036 (0.034)
P-value: Ind+Spouse=0	0.891	0.487	0.802	0.617	0.976	0.468
P-value: (Ind+Spouse)=Joint	0.213	0.489	0.313	0.463	0.328	0.523
P-value: Cash Self+Spouse=0	0.196	0.165	0.073*	0.117	0.781	0.515
P-value: Net Ind.=Net Cash	0.307	0.145	0.160	0.459	0.833	0.965
DV Mean (0% Ind.)	8.78	0.468	0.715	7.55	7.34	0.432
DV Mean (4% Joint)	8.63	0.460	0.681	7.34	7.11	0.412
N	1366	1404	1411	1398	1397	1411

Notes: Standard errors clustered at the couple level in parentheses. The individual interest rate is renormalized to run from 0-1, while the joint interest rate is renormalized to run from 0.2-1. The agreement variables run from 0 (no agreement) to 10 (perfect agreement). The value of livestock is deflated to 2009 values. In 2009 Ksh 80≈\$1. ***, **, and * indicate significance at the 1, 5, and 10 percent levels respectively.

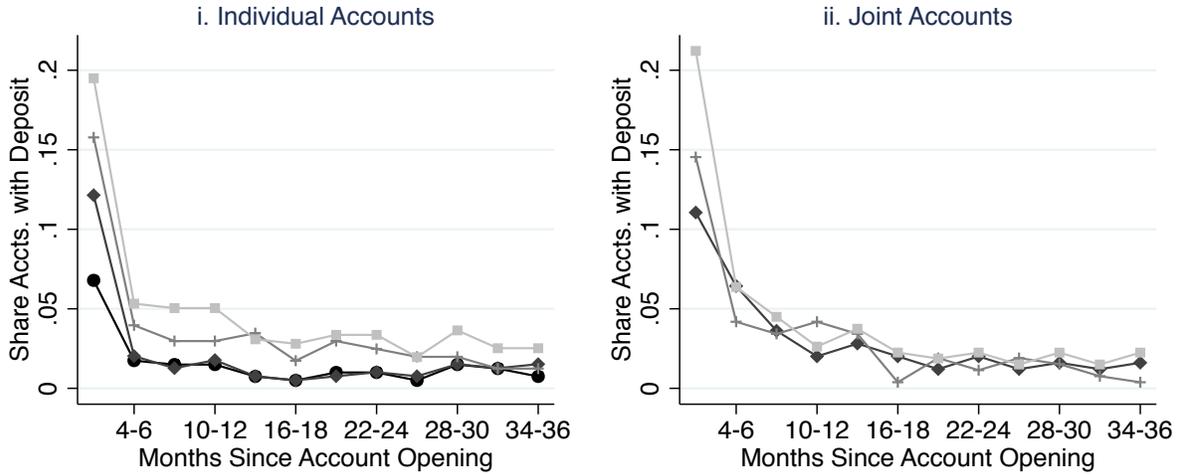
Table 6. Impact of Interest Rates on Individually-Held and Jointly-Held Resources

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Individually-Held Resources				Jointly Held Resources			
	Income - All Sources	Assets - All Sources	Business Income	Business Assets	Income - All Sources	Assets - All Sources	Business Income	Business Assets
Individual Interest	1142*** (411)	3536*** (1489)	569*** (191)	1712** (855)	-75.6 (98.4)	1642 (1592)	-1.97 (39.1)	603 (518)
Spouse's Interest	112 (406)	-2506* (1452)	265 (187)	28.0 (786)	33.8 (103)	-2051 (1529)	10.1 (45.1)	25.7 (493)
Joint Interest	391 (466)	252 (1822)	216 (208)	361 (1043)	-33.1 (139)	2896 (1904)	-3.01 (52.1)	84.9 (554)
Cash Prize - Self	253 (382)	755 (1468)	-181 (152)	729 (918)	-19.3 (93.4)	-1733 (1373)	10.3 (41.0)	-840** (376)
Cash Prize - Spouse	245 (390)	2031 (1629)	-128 (161)	175 (855)	220** (107)	-2159 (1416)	60.4 (45.2)	-258 (438)
P-value: Ind+Spouse=0	0.039**	0.609	0.003***	0.104	0.797	0.882	0.903	0.454
P-value: (Ind+Spouse)=Joint	0.241	0.791	0.077*	0.391	0.965	0.310	0.885	0.608
P-value: Cash Self+Spouse=0	0.381	0.212	0.217	0.461	0.219	0.091*	0.242	0.065*
P-value: Net Ind.=Net Cash	0.342	0.544	0.003***	0.586	0.285	0.341	0.432	0.070*
DV Mean (0% Ind.)	3110	7548	629	1998	903	13546	106	1056
DV Mean (4% Joint)	3550	10283	866	3081	943	13202	137	1335
N	1317	1207	1375	1397	1363	1223	1410	1399

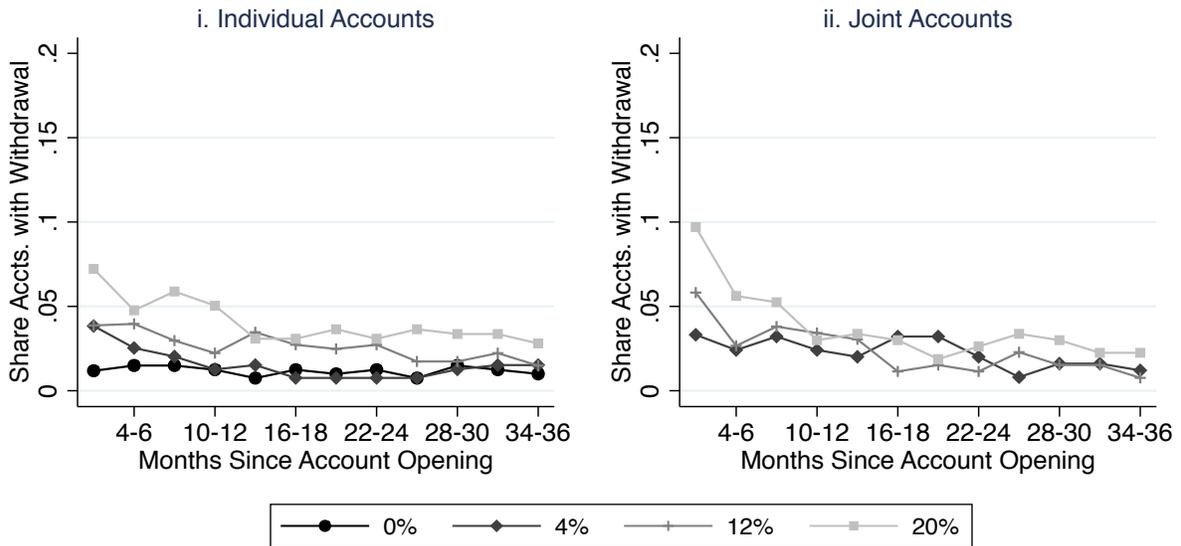
Notes: Standard errors clustered at the couple level in parentheses. During the endline respondents were asked to report whether resources were held individually or jointly with a spouse. Hence, total resources per person (outcomes in Table 3) are equal to individual resources plus half of joint resources. To be consistent with this coding and to keep outcomes in per capital terms, reported jointly-held resources are multiplied by one-half in this table. The individual interest rate is renormalized to run from 0-1, while the joint interest rate is renormalized to run from 0.2-1. All variables denominated in Kenyan Shillings are top-coded at the 99th percentile and deflated to 2009 values. In 2009 Ksh 80≈\$1. ***, **, and * indicate significance at the 1, 5, and 10 percent levels respectively.

Figure 1. Use of Experimental Accounts Over Time

A. Share Accounts with 1+ Deposits

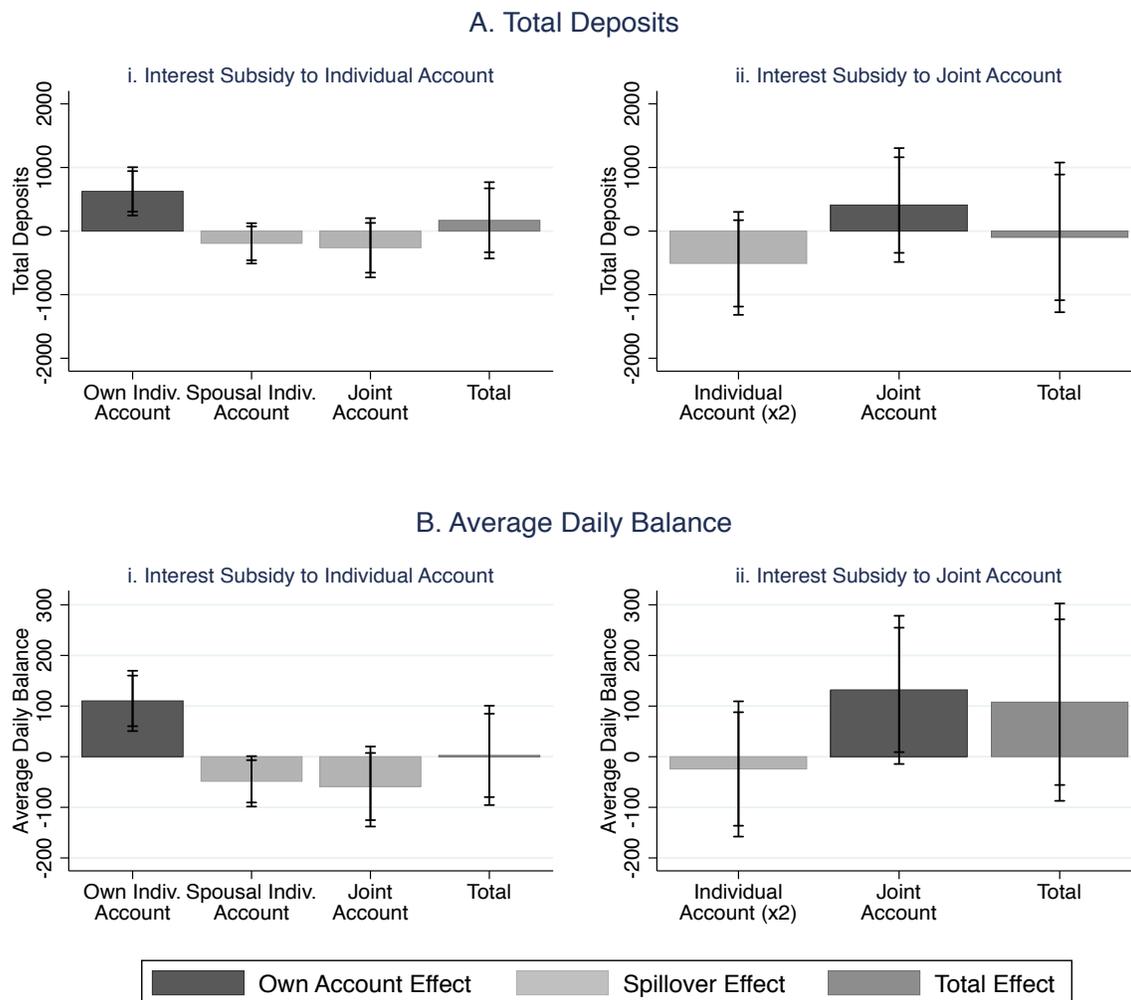


B. Share Accounts with 1+ Withdrawals



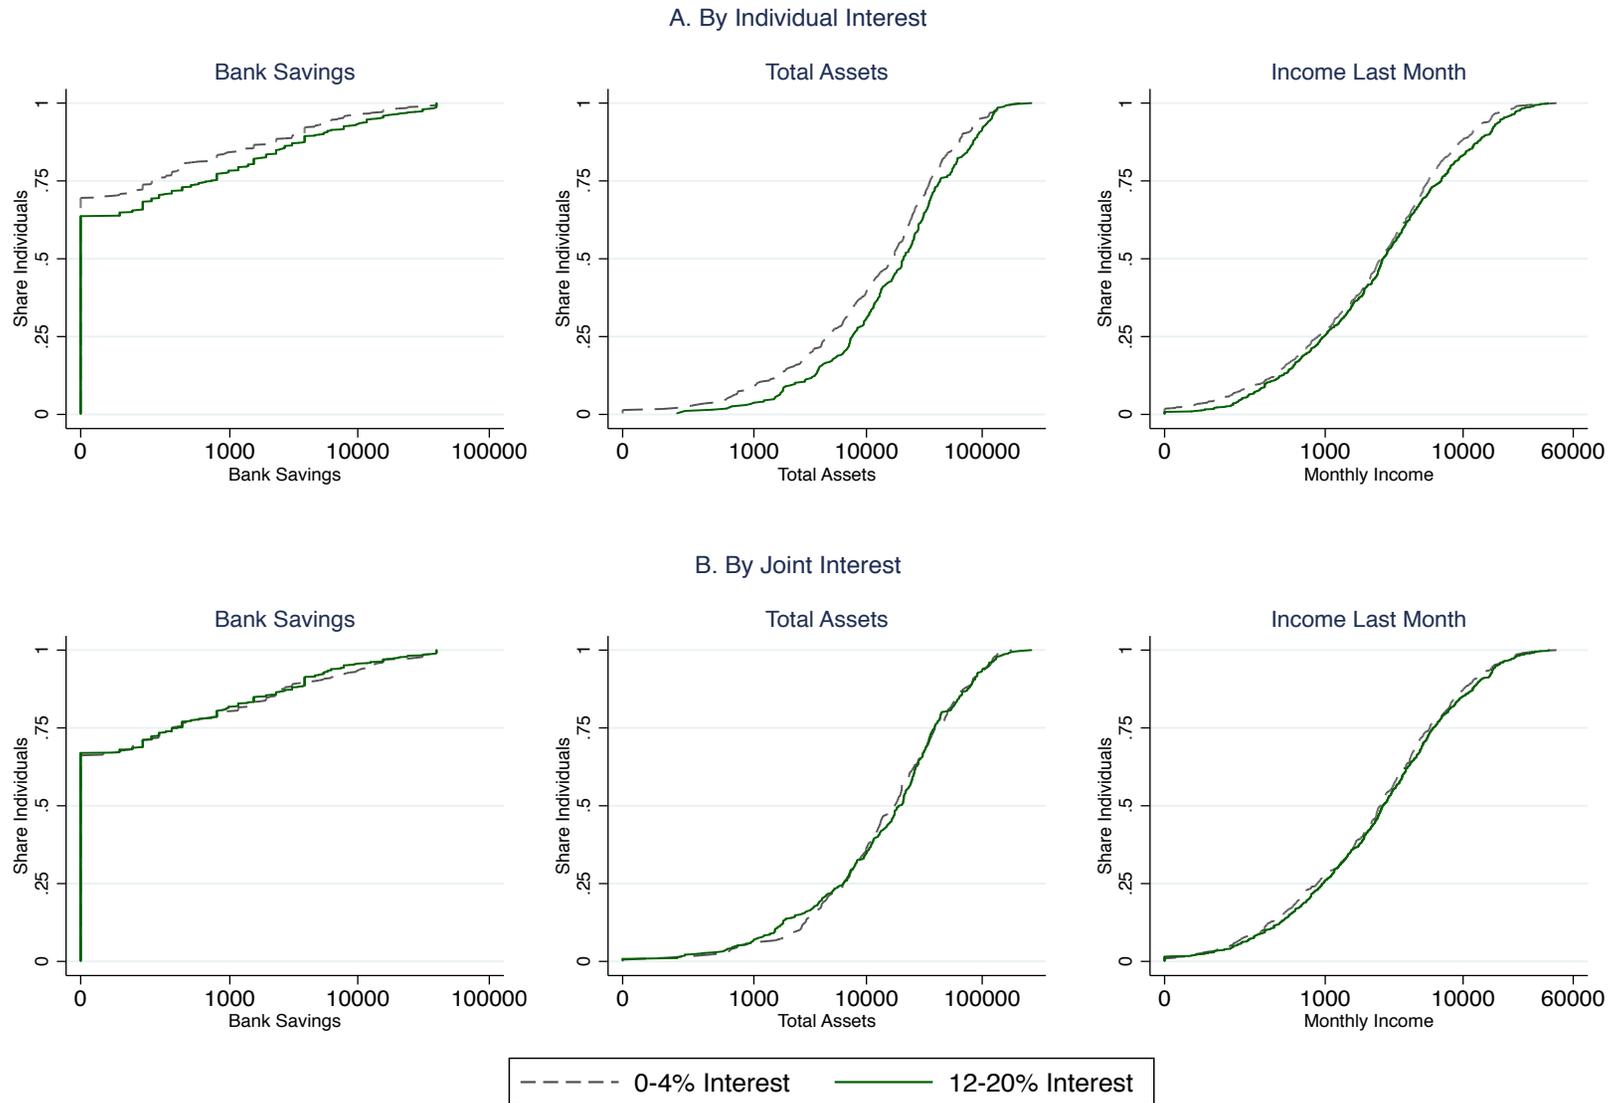
Notes: The figure plots the share of "potential" individual accounts (panel A.i and B.i) and joint accounts (Panels A.ii and B.ii) that received either a deposit (Panel A) or withdrawal (Panel B) in the relevant three month period. The first three month period drops all cash-prize eligible accounts to reflect transaction rates independent of the cash prize. All unopened accounts are coded as zeros.

Figure 2. Impact of Interest Subsidies on Short-Run Experimental Account Deposits and Balances



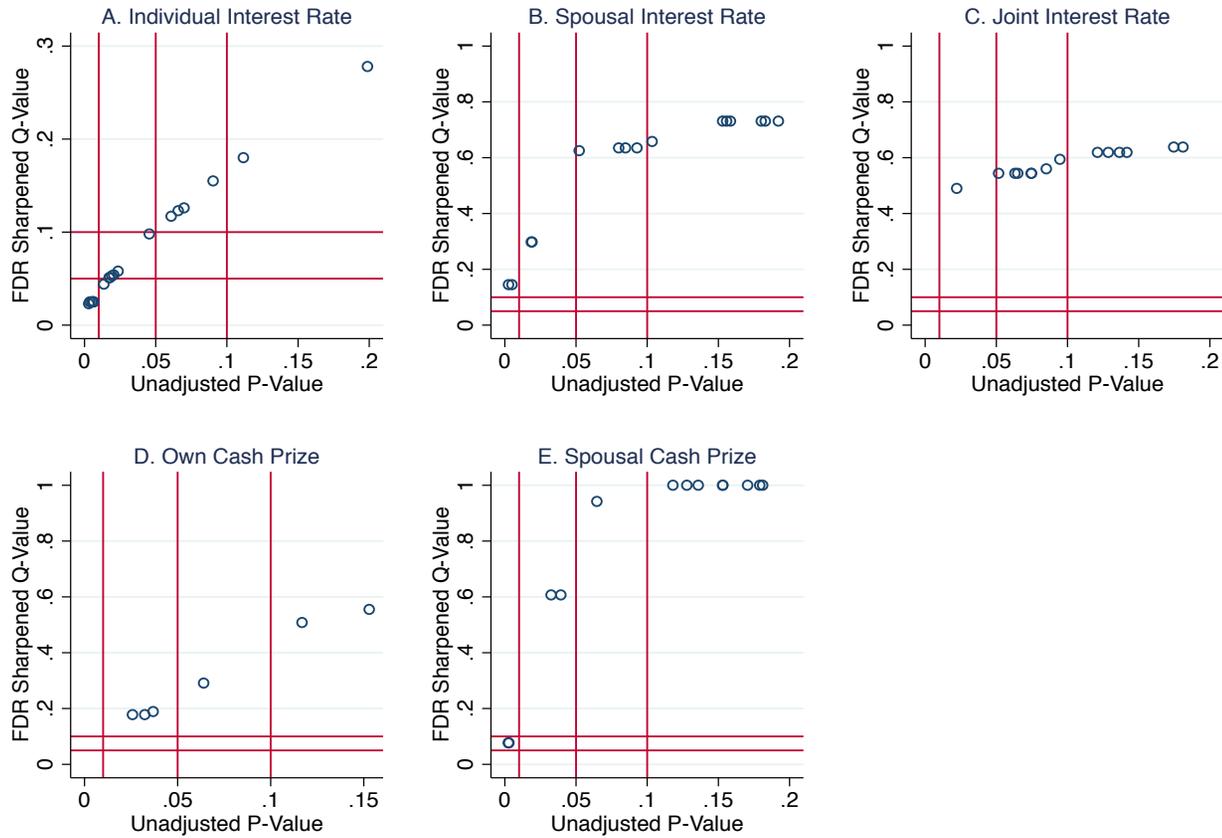
Notes: This figure provides a graphical representation of estimates in Table 2. The darkest bars give the direct impact of the interest rate on the account attached to that interest rate. The lightest bars give the effect of a given interest rate on other potential accounts. The mid-toned bar reflects the total effect of a given interest rate on deposits/balances across all experimental accounts. Whiskers give 90 and 95 percent confidence intervals based on robust standard errors clustered at the couple level when relevant. See notes to Table 2 for additional detail.

Figure 3. Distributional Impact of Interest Rates on Long-Run Economic Outcomes



Notes: This chart graphs the CDF of individual-level income and assets variables by interest rate treatment. All income and asset variables are top-coded at the 99th percentile.

Figure 4. Sharpened Q-Values to Control the False Discovery Rate by Treatment



Notes: This figure graphs traditional p-values on the x-axis and sharpened q-values that control the false discovery rate on the y-axis. The q-value represents the smallest possible false discovery rate at which a hypothesis would still be rejected. In other words, a q-value of 0.05 implies that a hypothesis would be rejected when the false discovery rate is set to 5 percent (i.e. 5 percent of rejected hypotheses are in fact true). I use the two-stage procedure proposed by Benjamini et al. (2006) to calculate q-values.

A Robustness of Main Results

Reporting Bias A first concern is that my results are simply driven by reporting bias. For reporting bias to rationalize the results, participants would have had to show their “appreciation” of higher individual interest rates by reporting more individual economic activity/increased rates of entrepreneurship, while showing appreciation of higher joint interest rates by reporting greater levels of spousal agreement. Although this type of systematic bias does not seem particularly likely *ex-ante*, the experimental design permits two additional tests of reporting bias. A first indication that reporting bias is likely not a concern is the fact that the cash prize treatment had no appreciable impact on long-run outcomes (see Tables 3-7). The cash awards were substantial, averaging Ksh 247 (\$3.09) for those receiving a payout. In contrast, the value of the interest payments were much smaller – the average payout (for those who received a positive payment) was just Ksh 26. One would therefore expect reporting bias to be most pronounced for the cash payments, which is not the case in practice.

Second, as part of pilot research for a separate project, a subset of individuals were asked whether they would be interested in a mobile money based commitment savings device (these questions were administered after completing the endline survey). This product was presented as a potential offering from Innovations for Poverty Action-Kenya, just as the initial interest rates were. The product was not actually on offer at endline – participants were simply asked to indicate on a scale of 0-10 how helpful they thought the product would be and how likely they would be to sign up if offered the product. Appendix Table R1 tests whether the interest rates or the cash prize are related to participants’ ratings of the pilot savings product.³⁴ There is no systematic evidence that individuals who received more favorable interest rate treatments rated the hypothetical product more favorably.

Top-Coding and Trimming Another concern is that the results are driven by a small number of observations in the upper tails of the income and assets distribution. Appendix Table R2 presents results for total and business income and assets where outcomes are top-coded at the 98th-90th percentiles or trimmed at the 99th-95th percentiles. I also present raw results with no top-coding. As expected, the point estimates using level outcomes decline with increased top-coding and trimming. However, the results are generally quite robust to additional top-coding and trimming, with many results surviving top-coding up to the 90th percentile and trimming through the 98th percentile. Thus, although right-tail observations

³⁴I do note that given that the interest rates impacted participants’ financial lives, it is not totally clear that the subsidies should be unrelated to the product rating, even absent reporting bias.

are clearly important for my results, estimates are relatively robust and are not exclusively driven by a small number of extreme outliers.

Imputing Missing Values Finally, it is important to note that I am missing endline data on outcomes due to both survey attrition and the fact that respondents occasionally reported that they did not know the answer to a question, or did not wish to provide an answer to a question. This latter source of missing data is amplified for variables that aggregate responses to a number of questions, such as my measures of total income and assets. Although attrition is not systematically correlated with treatments (recall Appendix Table A2), I still ask whether the results are robust to alternative imputation schemes in Appendix Table R3.

The methods in the first two panels of Table R3 are designed to be “treatment neutral”. In Panel A, all missing values are replaced with the sample mean among individuals with non-missing data. In Panel B, I regress each outcome of interest on the baseline control set and use the predicted values from that regression to impute missing values. The last two panels take a decidedly less neutral approach to further test the robustness of the individual interest rate results (I focus on just the individual interest rate here, since the results for the joint interest rate are more modest in magnitude and less statistically robust). In Panel C missing values in the 0 percent individual interest rate group are replaced with the mean outcome for the 20 percent interest group, while the 4 percent interest group is imputed with the mean for the 12 percent interest group, the 12 percent interest group is imputed with the mean for the 4 percent interest group, and the 20 percent interest group is imputed with the mean for the 0 percent interest group. In Panel D I replace missing values in the 0 percent interest group with the 80th percentile among non-missing values. The 4 percent interest group is imputed with the 60th percentile, the 12 percent interest group with the 40th percentile, and the 20 percent interest group with the 20th percentile. I perform the imputations using raw values, and then top-code the imputed variables at the 99th percentile. Note that all imputation and top-coding is performed on variable sub-components, which are then aggregated up. Thus, the “total income” variable is the sum of imputed/top-coded measures of farm, wage, business, and other income. (Results are very similar when aggregates are imputed/top-coded instead).

Appendix Table R3 illustrates that the results for business outcomes are robust to all four imputation schemes, while the results for overall income and assets only withstand “treatment-neutral” imputations. Yet the impacts on business income and assets are very similar in magnitude to the overall treatment effects on income and assets. This further underscores the notion that the primary mechanism by which the individual interest rate

increased income and assets was via microenterprise growth.

B Data Appendix

This paper uses four different sources of data: information from a baseline survey conducted from July-September 2009, information from a wave 1 endline survey conducted from August-November 2012, information from a wave 2 endline conducted from July-August 2013, and administrative data from Family Bank that covers the first three years of experimental account activity. All analysis is conducted at an individual level, using a core sample of 1,417 men and women who participated in the initial experiment and the first endline survey. This section provides additional detail on the construction of key variables used in the analysis. Note that top-coded variables that are an aggregate of multiple components (such as total endline assets) are the sum of top-coded components rather than the top-code of the sum.

Selected Baseline Variables

- *Income last month* – based on survey question “what was your income from all your income-generating activities last week?”. For comparability with the endline income measure, the weekly variable is multiplied by 52/12 in the analysis to construct a monthly measure.
- *Time preference questions* – participants were administered 10 tables of monetary choices, with each table consisting of 5 separate choices between a larger $x^{t+\tau} =$ Ksh 300 and a smaller Ksh $x^t \in \{290, 220, 150, 80, 10\}$. The $(t, t + \tau)$ pairs (where t and τ are expressed in weeks) were $(\frac{1}{7}, 1)$, $(\frac{1}{7}, 2)$, $(\frac{1}{7}, 3)$, $(\frac{1}{7}, 4)$, $(\frac{1}{7}, 8)$, $(\frac{1}{7}, 12)$, $(2, 3)$, $(2, 4)$, $(4, 8)$, $(4, 12)$.
 - *Time inconsistency variables* – the baseline survey included 4 different opportunities for an individual to display time inconsistent preferences. A respondent is coded as “impatient now-patient later” if this is the most prevalent form of time inconsistency in her answers. Similarly a respondent is coded as “patient now-impatient later” if this is the most prevalent form of time inconsistency in her answers.
 - *Weekly discount factor* – For each of the 10 time preference tables I assume that if an individual switches from “earlier” to “later”, that she is indifferent between the smaller and larger of the two x^t amounts and estimate the implied discount factor over monetary amounts. I then take the simple average of the estimates for each of the 10 tables as an estimate of the weekly discount factor.

Selected Wave 1 Endline Variables

- *Measures of income* – each respondent was asked about five different income sources: harvests (earnings in past 12 months), horticultural crops (earnings in past month), wage labor (earnings in past month), non-farm business (earnings in past month), and other income (earnings in past month). For each category, an individual was asked to specify how much income they themselves earned, and how much was earned jointly with his/her spouse. Individual income for each subcategory is defined as all individually earned income plus half of jointly earned income. The total income measure takes individual harvest income divided by 12 and adds the monthly individual income measures for all other subcategories together.
- *Measures of assets* – each respondent was asked about 8 different types of assets: bank accounts, SACCO accounts, money stored at home, ROSCA contributions, mobile money accounts, inventories and assets used to run a small business, the value of livestock, and other forms of savings. Respondents were asked about both individual assets and joint assets for bank accounts, business assets, and other savings. As with income, I construct total assets by adding up all individually owned assets and half of all joint assets. I assume that all livestock is jointly held (thus, this variable is always divided by two for all analysis in the paper). *Debt* – respondents were separately asked about both individually and jointly held debt to family and friends, formal/village banks, microfinance lenders, local moneylenders, shops, and other debt. Individual debt is the sum of all individually held amounts plus half of all jointly held amounts.

Selected Wave 2 Endline Variables

- *Has budget account for item X (e.g. business, home improvement)* – all respondents were first asked if they keep a budget in mind when making financial decisions, or if they simply pooled funds and met expenses as they came up. Respondents were then asked to list out all items in their budget, and specify how much money was devoted to that item. An individual is coded as having an account for item *X* if the individual reports keeping a budget and lists item *X* when giving budget detail.
- *Has “inflexible” account for item X* – this variable indicates that a respondent had a budget item for *X* and stated that they would not reduce spending on item *X* in order to meet a Ksh 1,000 expense.

Selected Variables from Administrative Bank Data

- *Used individual/joint account in first 6 months* – a respondent is coded as “used an individual account” if any transactions (other than the initial Ksh 100 opening balance) were posted to his or her account in the first 6 months of account activity (before interest payments were made). A respondent is coded as “used a joint account” if any transaction was posted to the joint account.
- *Individual/joint deposits* – equal to total deposits (excluding the Ksh 100 opening balance) posted to the account in the first 6 months of account activity.
- *Individual/joint withdrawals* – constructed the same way as total deposits. Excludes fees.
- *Individual/joint average balance* – the average daily balance (excluding the Ksh 100 opening balance) on an account for the first 6 months of account activity.
- *Used individual/joint account in year 3* – equal to 1 if any transaction was posted to the account in the final 12 months of the three year account activity observation window.
- *Used any account in the first 6 months* – equal to 1 if there was any transaction posted to any account owned by the couple, 0 otherwise.
- *Used any account in year 3* – equal to 1 if any transaction was posted to an experimental account owned by the couple in the final 12 months of the three year account activity observation window.

Appendix Table A1. Account Opening Decisions

	Share Couples	N
Joint Account Only	0.546	425
Two Individual Accounts Only	0.302	235
All Three Accounts	0.050	39
Joint and Husband's Account	0.042	33
Joint and Wife's Account	0.035	27
Husband's Account Only	0.015	12
Wife's Account Only	0.010	8
Declined to Open Any Account	0	0
Total	1	779

Appendix Table A2. Endline Attrition: Correlation with Treatment and Differential Selection

	(1)	(2)	(3)	(4)	(5)	(6)
	P-value: Treatment Correlated with Missing			P-value: Differential Selection on Observables		
	Individual Interest	Spousal Interest	Joint Interest	Individual Interest	Spousal Interest	Joint Interest
<i>A. Wave 1 Endline Outcomes</i>						
In Wave 1 Endline	0.280	0.558	0.743	0.714	0.539	0.392
Has Individual Bank Account	0.280	0.558	0.743	0.714	0.539	0.392
Has Joint Bank Account	0.335	0.743	0.219	0.687	0.074*	0.793
Individual Bank Savings	0.685	0.813	0.896	0.828	0.073*	0.373
Joint Bank Savings	0.735	0.785	0.850	0.666	0.662	0.793
Total Assets	0.342	0.272	0.609	0.265	0.228	0.339
Net Assets	0.217	0.269	0.542	0.231	0.168	0.344
Monthly Income	0.901	0.619	0.410	0.347	0.694	0.149
Main Occ. Entrepreneur	0.280	0.558	0.743	0.714	0.539	0.392
Has Business Profits/Assets	0.685	0.877	0.809	0.893	0.137	0.447
Business Assets	0.513	0.615	0.780	0.838	0.423	0.946
Business Income	0.614	0.740	0.750	0.636	0.104	0.210
Value Livestock	0.044**	0.143	0.490	0.942	0.605	0.571
Home Reno. Last Year	0.105	0.164	0.811	0.825	0.413	0.363
Has Permanent Roof	0.330	0.593	0.890	0.603	0.360	0.264
Agree: Spend Money	0.414	0.540	0.871	0.437	0.641	0.500
Agree: How Much to Save	0.221	0.948	0.945	0.575	0.582	0.702
Savings: Decide Together	0.366	0.626	0.738	0.712	0.421	0.365
Net Spousal Transfer	0.320	0.281	0.807	0.659	0.629	0.192
<i>B. Wave 2 Endline Outcomes</i>						
In Wave 2 Endline	0.168	0.101	0.736	0.776	0.796	0.685
No Budget	0.239	0.051*	0.962	0.886	0.882	0.397
Fraction Budget Inflexible	0.076*	0.017**	0.555	0.884	0.920	0.387
Has Business Account	0.139	0.056*	0.832	0.924	0.896	0.640
How Helped	0.253	0.220	0.592	0.836	0.287	0.784
N	1558	1558	1558	1558	1558	1558

Notes: P-values are from F-tests of coefficient restrictions on regressions with robust standard errors clustered at the couple level. The underlying regression in the first three columns regresses an indicator for whether the specified outcome is missing on the treatment of interest. The table reports p-values for whether attrition is significantly correlated with treatment. The underlying regression in the last three columns regresses the missing indicator on the treatment of interest, all demographic controls listed in Table 1, and interactions between the demographic variables and the treatment of interest. The table reports the p-value from a joint test of whether the interaction terms are equal to zero. The individual interest rate is renormalized to run from 0-1, while the joint interest rate is renormalized to run from 0.2-1. ***, **, and * indicate significance at the 1, 5, and 10 percent levels respectively.

Appendix Table A3. Impacts of the Maximum Interest Rate on Experimental Bank Account Use

	(1)	(2)	(3)	(4)	(5)	(6)
	In First 6 Months:					
	Used	Total Deposits	Total Withdrawals	Ending Balance	Avg. Daily Balance	Used - Year 3
Maximum Interest 12 Percent	0.102 (0.070)	641 (532)	241 (453)	204*** (85.7)	250*** (76.2)	0.039 (0.034)
Maximum Interest 20 Percent	0.151** (0.066)	1444*** (511)	910** (463)	314*** (75.5)	315*** (59.9)	0.087*** (0.033)
Husband's Account Has Max. Int.	0.008 (0.037)	70.2 (438)	47.7 (363)	-26.2 (78.0)	-22.2 (75.0)	0.041* (0.024)
Wife's Account Has Max. Int.	0.037 (0.037)	-328 (418)	-291 (348)	-38.8 (79.7)	-93.6 (73.0)	0.013 (0.024)
DV Mean (4% Max. Int.)	0.339	893	710	137	123	0.051
N	779	779	779	779	779	779

Notes: Heteroskedasticity robust standard errors in parentheses. Variables denominated in Kenyan Shillings are top-coded at the 99th percentile. ***, **, and * indicate significance at the 1, 5, and 10 percent levels respectively.

Appendix Table A4. Impact of Interest Rates on Other Types of Assets (Excluding Bank, Business Assets)

	(1)	(2)	(3)	(4)	(5)	(6)
	Value Livestock	ROSCA Savings	SACCO Savings	Mobile Money Savings	Home Savings	Other Savings
Individual Interest	175 (1179)	321 (261)	404 (601)	27.3 (61.8)	85.9 (73.2)	138 (916)
Spouse's Interest	-831 (1196)	87.3 (249)	-1678*** (597)	67.1 (60.1)	109 (81.6)	-989 (757)
Joint Interest	2582 (1732)	231 (320)	53.1 (836)	-9.79 (71.6)	47.3 (91.9)	69.3 (1054)
Cash Prize - Self	-1114 (1106)	-170 (233)	235 (682)	-60.6 (55.5)	-20.4 (75.0)	-856 (702)
Cash Prize - Spouse	-1080 (1189)	-119 (250)	81.9 (672)	90.6 (63.4)	-25.0 (72.0)	-846 (790)
P-value: Ind+Spouse=0	0.768	0.252	0.074*	0.284	0.094*	0.523
P-value: (Ind+Spouse)=Joint	0.227	0.705	0.253	0.315	0.306	0.647
P-value: Cash Self+Spouse=0	0.291	0.420	0.717	0.716	0.661	0.116
P-value: Net Ind.=Net Cash	0.617	0.153	0.128	0.591	0.120	0.613
DV Mean (0% Ind.)	12458	1823	1454	269	359	2375
DV Mean (4% Joint)	11657	1946	1780	284	432	2196
N	1366	1343	1398	1388	1387	1398

Notes: Standard errors clustered at the couple level in parentheses. The individual interest rate is renormalized to run from 0-1, while the joint interest rate is renormalized to run from 0.2-1. On the endline survey each individual was asked to report total livestock owned by the household. The individual was then assigned half this value. All variables are denominated in Kenyan Shillings and top-coded at the 99th percentile. ***, **, and * indicate significance at the 1, 5, and 10 percent levels respectively.

Appendix Table A5. Impact of Interest Rates on Income, by Source (Excluding Business Income)

	(1)	(2)	(3)
	Farm	Wage	Other
	Income	Income	Income
Individual Interest	-47.2 (128)	466 (293)	-34.8 (149)
Spouse's Interest	-21.2 (127)	-199 (275)	26.0 (176)
Joint Interest	-35.2 (171)	-24.9 (331)	217 (198)
Cash Prize - Self	25.7 (134)	361 (288)	-130 (134)
Cash Prize - Spouse	204 (137)	316 (287)	54.3 (153)
P-value: Ind+Spouse=0	0.730	0.501	0.966
P-value: (Ind+Spouse)=Joint	0.894	0.574	0.419
P-value: Cash Self+Spouse=0	0.314	0.083*	0.696
P-value: Net Ind.=Net Cash	0.347	0.442	0.768
DV Mean (0% Ind.)	1102	1705	605
DV Mean (4% Joint)	1080	1869	635
N	1354	1390	1403

Notes: Standard errors clustered at the couple level in parentheses. The individual interest rate is renormalized to run from 0-1, while the joint interest rate is renormalized to run from 0.2-1. All variables are denominated in Kenyan Shillings and top-coded at the 99th percentile. ***, **, and * indicate significance at the 1, 5, and 10 percent levels respectively.

Appendix Table A6. Impact of Interest Rates on Debt, by Source

	(1)	(2)	(3)	(4)	(5)	(6)
	Debt From:					
	Banks	MFIs	Money-lenders	Shops	Family Members	Other Sources
Individual Interest	-780 (1368)	-80.9 (252)	201 (299)	-23.7 (20.7)	18.4 (19.8)	654* (348)
Spouse's Interest	445 (1298)	-41.7 (219)	883*** (292)	13.3 (20.6)	25.4 (17.7)	696* (358)
Joint Interest	2994 (2036)	478 (357)	294 (361)	6.59 (23.1)	6.16 (26.9)	542 (399)
Cash Prize - Self	-646 (1469)	77.8 (267)	-190 (250)	-3.31 (20.2)	0.846 (20.3)	-402 (281)
Cash Prize - Spouse	1494 (1834)	1.27 (253)	436 (318)	-19.7 (15.8)	-41.4*** (13.8)	-91.7 (327)
P-value: Ind+Spouse=0	0.860	0.716	0.008***	0.735	0.113	0.006***
P-value: (Ind+Spouse)=Joint	0.245	0.284	0.165	0.626	0.382	0.216
P-value: Cash Self+Spouse=0	0.716	0.833	0.568	0.356	0.123	0.244
P-value: Net Ind.=Net Cash	0.677	0.650	0.145	0.752	0.020**	0.004***
DV Mean (0% Ind.)	4118	705	1564	52.5	68.7	1201
DV Mean (4% Joint)	2234	582	1531	41.5	93.8	1359
N	1410	1413	1407	1414	1414	1406

Notes: Standard errors clustered at the couple level in parentheses. The individual interest rate is renormalized to run from 0-1, while the joint interest rate is renormalized to run from 0.2-1. All variables are denominated in Kenyan Shillings and top-coded at the 99th percentile. ***, **, and * indicate significance at the 1, 5, and 10 percent levels respectively.

Appendix Table A7. Heterogeneity in Impacts on Business Outcomes by Baseline Entrepreneurship

	(1)	(2)	(3)	(4)	(5)	(6)
	Main Occupation Entrepreneur	Any Business Profits or Capital	Business Capital	Monthly Business Profits	Has Business Budget	Downwardly Rigid Business Budget
<i>Main Occupation Entrepreneurial at Baseline</i>						
Individual Interest	0.109** (0.054)	0.201*** (0.055)	5355*** (2278)	1108*** (397)	0.114** (0.051)	0.112*** (0.041)
Spouse's Interest	0.046 (0.054)	0.109* (0.056)	-1321 (2052)	-151 (358)	0.046 (0.051)	0.046 (0.040)
Joint Interest	0.008 (0.064)	-0.024 (0.065)	1471 (2320)	452 (389)	0.007 (0.057)	-0.039 (0.046)
DV Mean	0.498	0.588	7118	1641	0.297	0.164
N	590	548	572	560	542	542
<i>Other Main Occupation at Baseline</i>						
Individual Interest	0.053 (0.037)	0.033 (0.045)	736 (975)	146 (173)	0.013 (0.032)	0.020 (0.023)
Spouse's Interest	-0.009 (0.038)	-0.002 (0.045)	-319 (971)	537** (235)	0.029 (0.036)	-0.011 (0.025)
Joint Interest	0.037 (0.048)	0.064 (0.055)	688 (1481)	137 (232)	0.043 (0.043)	0.015 (0.033)
DV Mean	0.224	0.338	2818	725	0.145	0.076
N	822	787	803	804	759	759
<i>Tests of Equality (Ent.=Non-Ent.)</i>						
P-Value: Indiv. Interest	0.402	0.017**	0.059*	0.025**	0.096*	0.050*
P-Value: Spousal Interest	0.415	0.124	0.656	0.105	0.782	0.225
P-Value: Joint Interest	0.714	0.291	0.775	0.483	0.613	0.338

Notes: Standard errors clustered at the couple level in parentheses. The individual interest rate is renormalized to run from 0-1, while the joint interest rate is renormalized to run from 0.2-1. Business profits and capital are top-coded at the 99th percentile. All regressions also control for own and spousal cash prize selection. Regressions for entrepreneurs and non-entrepreneurs are jointly estimated using seemingly unrelated regression in order to perform tests of equality. ***, **, and * indicate significance at the 1, 5, and 10 percent levels respectively.

Appendix Table A8. Returns to Capital and External Validity

A. Returns to Capital and Time to Effect Size from Individual Interest Subsidies

	Δ Monthly Income/ Δ Capital	Initial Investment (in Ksh) Needed to Meet Effect Sizes in 32 Months When Marginal Propensity to Reinvest is:	
		0.5	1
Business Assets and Income	0.207	114	6
Total Assets and Income	0.190	330	23

Panel B. Existing Evidence on Returns to Capital in Microenterprises

Paper	Δ Monthly Profit/ Δ Capital	Monthly Return to Capital	Maximum Follow-Up Period
		(i) 0.17-0.25,	
Udry and Anagol (2006) ^a	--	(ii) 0.05	N/A
de Mel et al. (2008)	0.05-0.06	0.05	2 years
McKenzie and Woodruff (2008)	0.20-0.33	--	1.25 years
de Mel et al. (2012)	0.06-0.12	0.11	6 years
Field et al. (2013)	0.10-0.14	0.11-0.13	3 years
Fafchamps et al. (2014) ^b	0.21-0.29	--	1 year
Blattman et al. (2014)	0.04-0.07	--	4 years

Notes: ^aEstimate (i) is for return to pineapple cultivation, (ii) is return to automotive capital in Accra. ^bThese estimates divide the impact of an in-kind cash grant on monthly profits (Table 3) by the value of the grant (150 cedis).

Appendix Table A9. Heterogeneity in Impacts by Baseline Bank Account Ownership

	(1)	(2)	(3)	(4)	(5)	(6)
	Total Assets (Hypersine)	Total Income Last Month (Hypersine)	Main Occupation Entrepreneur	Any Business Profits or Capital	Business Capital	Monthly Business Profits
<i>Had Bank Account at Baseline</i>						
Individual Interest	0.594** (0.293)	0.227 (0.205)	0.167*** (0.071)	0.120 (0.078)	5963* (3524)	659 (588)
Spouse's Interest	-0.431 (0.292)	-0.239 (0.198)	0.061 (0.070)	-0.002 (0.077)	-6642** (3189)	-539 (613)
Joint Interest	0.568* (0.330)	0.147 (0.246)	-0.013 (0.084)	-0.025 (0.088)	6145* (3616)	494 (641)
DV Mean	10.7	9.25	0.390	0.524	10020	1885
N	225	278	308	286	296	295
<i>Did Not Have Bank Account at Baseline</i>						
Individual Interest	0.425*** (0.167)	0.276** (0.137)	0.051 (0.037)	0.094** (0.041)	1547 (1077)	413** (192)
Spouse's Interest	-0.017 (0.157)	0.187 (0.133)	0.027 (0.037)	0.078* (0.042)	1204 (996)	548*** (178)
Joint Interest	-0.025 (0.224)	-0.005 (0.171)	0.028 (0.047)	0.014 (0.050)	-789 (1224)	98.0 (204)
DV Mean	9.74	8.12	0.326	0.419	3142	888
N	828	1001	1109	1053	1084	1073
<i>Tests of Equality (Banked=Non-Banked)</i>						
P-Value: Individual Interest	0.615	0.841	0.157	0.774	0.233	0.690
P-Value: Spousal Interest	0.207	0.068*	0.672	0.357	0.020**	0.087*
P-Value: Joint Interest	0.122	0.607	0.665	0.695	0.060*	0.552

Notes: Standard errors clustered at the couple level in parentheses. The individual interest rate is renormalized to run from 0-1, while the joint interest rate is renormalized to run from 0.2-1. Business profits and capital are top-coded at the 99th percentile. All regressions control for individual and spousal cash prize selection. Regressions for the banked and non-banked are jointly estimated using seemingly unrelated regression in order to perform tests of equality. ***, **, and * indicate significance at the 1 %, 5 %, and 10 percent levels respectively.

Appendix Table R1. Reporting Bias - Impact of Treatments on Pilot Product Ratings

	(1)	(2)	(3)	(4)
	Helpful: Scale	Sign Up: Scale	Helpful: 10/10	Sign Up: 10/10
Individual Interest	0.269 (0.199)	0.023 (0.214)	-0.011 (0.049)	0.047 (0.051)
Spouse's Interest	-0.148 (0.201)	-0.173 (0.226)	0.053 (0.049)	0.011 (0.051)
Joint Interest	0.309 (0.252)	0.303 (0.294)	0.092 (0.064)	0.042 (0.062)
Cash Prize - Self	0.232 (0.177)	0.289 (0.216)	0.044 (0.049)	0.076 (0.047)
Cash Prize - Spouse	-0.186 (0.189)	-0.436* (0.246)	-0.082* (0.048)	-0.075 (0.049)
P-value: Ind+Spouse=0	0.662	0.649	0.568	0.431
P-value: (Ind+Spouse)=Joint	0.609	0.296	0.588	0.864
P-value: Cash Self+Spouse=0	0.856	0.670	0.602	0.989
P-value: Net Ind.=Net Cash	0.849	0.994	0.449	0.590
DV Mean (0% Ind.)	8.36	8.18	0.440	0.464
DV Mean (4% Joint)	8.44	8.03	0.423	0.514
N	651	650	650	651

Notes: Standard errors clustered at the couple level in parentheses. The variables in the first two columns run from 0 (least helpful/least likely to sign up) to 10 (most helpful/would definitely sign up). The individual interest rate is renormalized to run from 0-1, while the joint interest rate is renormalized to run from 0.2-1. The sample sizes in this table are small because these questions were only asked to a subset of respondents. ***, **, and * indicate significance at the 1, 5, and 10 percent levels respectively.

Appendix Table R2. Robustness of Main Results to Additional Top-Coding and Trimming

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Raw	Top-Coded, 98th Percentile	Top-Coded, 95th Percentile	Top-Coded, 90th Percentile	Trimmed, 99th Percentile	Trimmed, 98th Percentile	Trimmed, 95th Percentile
<i>Panel A - Total Assets</i>							
Individual Interest	10927*** (4608)	4455** (2112)	2013 (1502)	1124 (1039)	6436** (2932)	5028** (2464)	2587 (1875)
Spousal Interest	-5652 (5142)	-4746** (2058)	-1962 (1480)	-1062 (1024)	-6845*** (2799)	-4287* (2347)	-3089* (1839)
Joint Interest	9797 (6549)	4171 (2746)	3411* (2016)	2706* (1432)	6254* (3530)	3554 (3017)	2546 (2411)
N	1053	1053	1053	1053	1043	1032	1001
<i>Panel B - Monthly Income</i>							
Individual Interest	1247** (584)	961*** (393)	630** (293)	342* (188)	1143*** (471)	836** (390)	362 (305)
Spousal Interest	-277 (590)	175 (377)	121 (281)	90.5 (181)	105 (480)	120 (399)	468 (299)
Joint Interest	680 (697)	515 (443)	392 (340)	322 (223)	377 (550)	372 (454)	627 (386)
N	1279	1279	1279	1279	1267	1254	1216
<i>Panel C - Business Capital</i>							
Individual Interest	4533** (2034)	2096** (969)	1219* (647)	459** (229)	1875* (977)	1048 (765)	368 (378)
Spousal Interest	-1429 (1437)	-64.9 (894)	119 (615)	138 (224)	890 (892)	215 (706)	334 (363)
Joint Interest	-17.2 (2301)	623 (1138)	301 (775)	89.7 (283)	390 (1108)	239 (903)	276 (441)
N	1380	1380	1380	1380	1368	1353	1311
<i>Panel D - Monthly Business Profit</i>							
Individual Interest	576** (271)	482*** (176)	350*** (126)	248*** (94.6)	437*** (169)	315** (145)	154 (94.9)
Spousal Interest	228 (260)	346* (178)	239* (126)	199** (96.5)	393** (170)	346*** (142)	123 (96.4)
Joint Interest	178 (293)	197 (199)	114 (151)	102 (115)	301 (205)	41.7 (165)	-6.26 (115)
N	1368	1368	1368	1368	1356	1345	1304

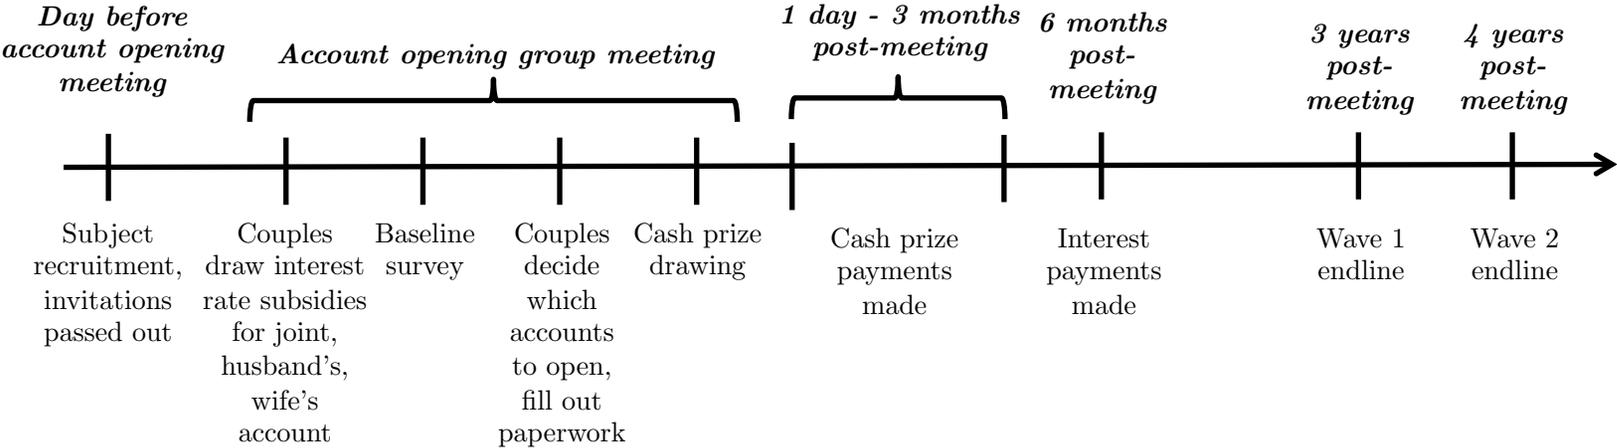
Notes: Standard errors clustered at the couple level in parentheses. The individual interest rate is renormalized to run from 0-1, while the joint interest rate is renormalized to run from 0.2-1. All regressions control for individual and spousal cash prize selection. ***, **, and * indicate significance at the 1, 5, and 10 percent levels respectively.

Appendix Table R3. Robustness of Main Results to Imputing Missing Values

	(1)	(2)	(3)	(4)
	Total Assets	Monthly Income	Business Assets	Business Profits
<i>Panel A - Impute Mean Values</i>				
Individual Interest	3737* (1983)	795** (391)	2361*** (1007)	477*** (175)
Spousal Interest	-3437* (1930)	140 (379)	-106 (931)	298* (178)
Joint Interest	2485 (2559)	361 (462)	451 (1187)	219 (194)
<i>Panel B - Predict Missing Values Using Baseline Covariates</i>				
Individual Interest	3578* (2093)	725* (404)	2207** (1017)	469*** (178)
Spousal Interest	-3194 (2074)	195 (399)	97.0 (944)	328* (189)
Joint Interest	3295 (2723)	253 (477)	314 (1211)	221 (203)
<i>Panel C - Reverse Imputation by Individual Interest Rate</i>				
Individual Interest	3028 (1985)	678* (391)	1955* (1010)	411*** (175)
Spousal Interest	-3512* (1931)	135 (379)	-147 (934)	297* (178)
Joint Interest	2379 (2559)	353 (462)	371 (1187)	210 (194)
<i>Panel D - Percentile Imputation by Individual Interest Rate</i>				
Individual Interest	1181 (2026)	451 (398)	2439*** (1013)	414*** (176)
Spousal Interest	-3286* (1967)	145 (387)	-26.4 (937)	296* (179)
Joint Interest	2290 (2595)	294 (474)	484 (1192)	198 (194)
N	1558	1558	1558	1558

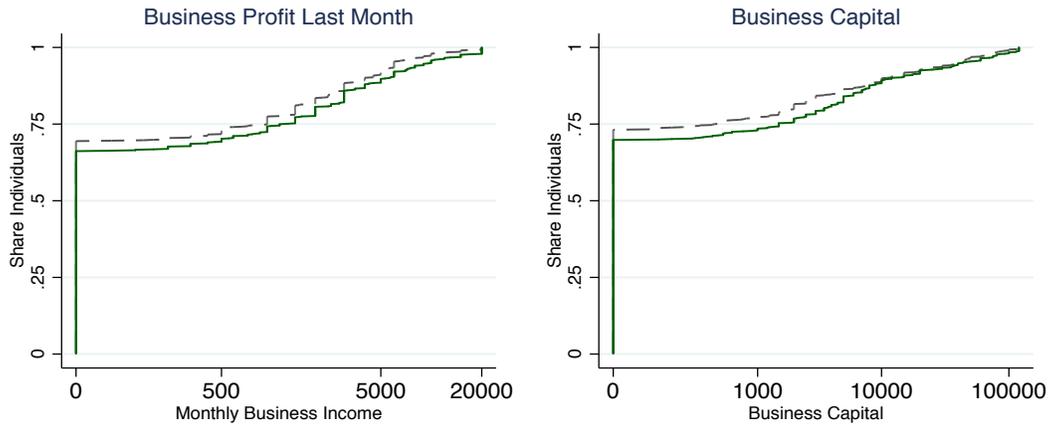
Notes: Standard errors clustered at the couple level in parentheses. All regressions control for individual and spousal cash prize selection. The individual interest rate is renormalized to run from 0-1, while the joint interest rate is renormalized to run from 0.2-1. All variables are top-coded at the 99th percentile. In Panel A, missing values of the dependent variable are replaced with overall means among non-missing observations. In Panel B, missing values are replaced by predicted values obtained by regressing the dependent variable of interest on the baseline control set listed in Table 1. In Panel C, missing values in the 0 percent individual interest rate group are replaced by the mean value in the 20 percent individual interest rate group. Missing values in the 4 percent interest group are replaced with the mean in the 12 percent interest group, missing values in the 12 percent interest group are replaced with the mean in the 4 percent interest group, and missing values in the 20 percent interest group are replaced with the mean value in the 0 percent interest group. In Panel D missing values in the 0 percent individual interest group are replaced with the 80th percentile among non-missing values. Missing values in the 4 percent interest group are replaced with the 60th percentile, missing values in the 12 percent interest group are replaced with the 40th percentile, and missing values in the 20 percent interest group are replaced with the 20th percentile. ***, **, and * indicate significance at the 1%, 5%, and 10 percent levels respectively.

Appendix Figure A1. Timeline of Experimental Activities

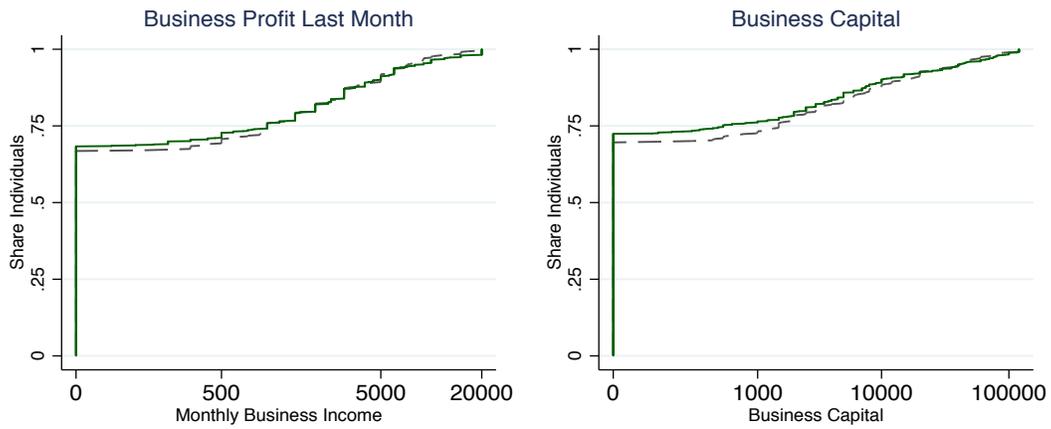


Note: Activities took place in sequential order from left to right.

A. By Individual Interest



B. By Joint Interest



Notes: This chart graphs the CDF of individual-level income and assets variables by interest rate treatment. All income and asset variables are top-coded at the 99th percentile.