4-2015

Do Opposites Detract? Intrahousehold Preference Heterogeneity and Inefficient Strategic Savings

Simone Schaner
Dartmouth College

Do Opposites Detract? Intrahousehold Preference Heterogeneity and Inefficient Strategic Savings*

Simone Schaner†

July 10, 2013

Abstract

This paper uses a field experiment to test whether intrahousehold heterogeneity in discount factors leads to inefficient strategic savings behavior. I gave married couples in rural Kenya the opportunity to open both joint and individual bank accounts at randomly assigned interest rates. I also directly elicited discount factors for all individuals in the experiment. Couples who are well matched on discount factors are less likely to use costly individual accounts and respond robustly to relative rates of return between accounts, while their poorly matched peers do not. Consequently, poorly matched couples forgo significantly more potential interest earnings on their savings.

JEL Codes: C93, D13, D14, O12

Keywords: intrahousehold allocation, saving, preference heterogeneity

*I thank Abhijit Banerjee, Esther Duflo, and Tavneet Suri for invaluable advice and feedback at all stages of this project. I am also grateful to Pascaline Dupas, Eric Edmonds, Dan Keniston, Erzo Luttmer, Ben Olken, Josh Schwartzstein, Rob Townsend, Jonathan Zinman, and numerous seminar participants for many 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 Noreen Makana for her superb field management, Moses Baraza for his expert advice, and the IPA enumerators 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, and the National Science Foundation's Graduate Research Fellowship. The study protocol in this paper was approved by IRBs at MIT, the Kenya Medical Research Institute, and IPA-Kenya. All errors are my own.

†simone.schaner@dartmouth.edu

1 Introduction

Members of the same household generally have different preferences and priorities – consequently, households must decide how to aggregate these preferences when allocating resources across members and over time. Moreover, it is clear that individuals actively advocate for their own preferences during the decision making process, since shifts in individual bargaining power translate into shifts in allocations (Lundberg et al. 1997; Chiappori et al. 2002; Duflo 2003; Bobonis 2009). These observations give rise to a natural set of questions: does preference heterogeneity ever lead individuals to take costly action to manipulate intrahousehold resource allocations? If so, how big are the resulting distortions? This is especially relevant for understanding households' intertemporal choices. While most studies that focus on static choices (e.g. what to consume) cannot reject the null of efficient households, most studies that focus on dynamic choices (e.g. savings and mutual insurance) reject efficiency.[1] While this result is well established, the causes of intertemporal inefficiency remain poorly understood.

This paper investigates these causes by using a field experiment conducted in rural Kenya to test whether heterogeneity in discount factors is associated with inefficiency in a fundamental intertemporal choice: savings behavior. When designing this study I chose to focus on heterogeneity in discount factors (rather than other aspects of preference heterogeneity) for two reasons. First, the empirical literature supporting static efficiency suggests that households do not have difficulty reconciling heterogeneous preferences over what to consume. It is therefore intuitive to turn to heterogeneity in intertemporal preferences as a driver of inefficient intertemporal behavior. Second, there is a strong theoretical justification for focusing on the difficulties associated with aggregating preferences with differing discount factors.[2] However this paper is, to my knowledge, the first to assess the implications of this heterogeneity empirically (a novel feature of the experiment is that it allows me to size the magnitude of the savings inefficiency), and one of few papers that explicitly focuses on the consequences of discount factor heterogeneity within the household.[3] Moreover, the application to savings is particularly important in developing country contexts, where individuals often make use of costly, secure savings devices even when lower cost alternatives appear to

[1] This point is made by Chiappori and Donni (2009), who also provide a summary of the empirical evidence. A partial selection of papers finding evidence of efficient intrahousehold resource allocation include Browning and Chiappori (1998), Chiappori et al. (2002), Rangel and Thomas (2005), and Bobonis (2009). Studies finding evidence of inefficiency include Udry (1996), Duflo and Udry (2004), Mazzocco (2007), de Mel et al. (2009), and Robinson (2011).

[2] See, e.g. Marglin 1963; Feldstein 1964; Browning 2000; Caplin and Leahy 2004; Gollier and Zeckhauser 2005; Jackson and Yariv 2010; Zuber 2010.

[3] An important exception is Browning (2000). However, in his model allocations are always Pareto efficient because all savings devices offer the same rate of return and carry no transaction costs. Difficulties posed by intrahousehold discount factor heterogeneity are also discussed by, but not the main focus of, Bernheim (1999) and Chiappori and Donni (2009). Mazzocco (2007) writes a baseline model of intrahousehold decision making that allows for differing discount factors, but assumes this away in order to perform Euler-equation based tests of intertemporal efficiency.

be available (Collins et al. 2009; Karlan and Morduch 2010).

This paper's central premise is that when spouses do not agree about the time path of consumption, they may take costly action to manipulate outcomes in their favor. This is easy to see with a simple example: imagine the savings problem of a patient wife paired with an impatient husband. The wife can save in either a joint bank account that she shares with her spouse, or an individual account. She knows that if she deposits funds into the joint account, her husband, who only cares about the present, will simply withdraw all her savings and spend the funds on current consumption. However, if she saves in her individual account (which her husband cannot access), she is guaranteed that her savings will be preserved for the next period. In this context, the wife may choose to save individually even if the joint account offers a much higher rate of return. Or put another way, heterogeneity in intertemporal preferences may compel the wife to make costly strategic savings choices.

I formalize this idea using a simple model of non-cooperative household savings behavior that incorporates heterogeneity in time preferences. The model predicts that losses due to strategic behavior will increase with preference heterogeneity. The model also allows me to characterize efficient savings choices in an environment that mimics my experimental context, and to identify cases where individual bank account use should increase with preference heterogeneity.

I test these predictions with a field experiment, which I conducted in Western Kenya in the Summer of 2009. I gave 598 married couples the opportunity to open three savings accounts (two individual, one joint) bearing randomly assigned interest rates. These interest rates were substantially higher than those available on the market at the time. An innovative feature of the experiment is that it created random variation in relative rates of return between accounts, even conditional on an account's own interest rate. This provides a clear test of efficient investment: the household should always choose to save in the account that offers the highest rate of return. I also asked each respondent in the experiment a battery of questions to directly elicit estimates of discount factors, which I use to calculate measures of intrahousehold heterogeneity.

In practice, I find a substantial amount of heterogeneity in elicited discount factors – the median couple's estimated discount factors are 0.5 standard deviations apart. Moreover, I find that couples who are "well matched" in terms of discount factors (i.e. above-median in match quality) robustly respond to relative rates of return in a manner consistent with efficient investment. In contrast, poorly matched couples are completely insensitive to relative rates of return and tend to make especially heavy use of individual accounts. I also reject that the responses of well and poorly matched couples are the same. These differences in behavior have financial consequences for poorly matched couples – I exploit the design of the field experiment to calculate forgone interest earnings by match quality and find that poorly matched couples leave at least 58 percent more interest on the table.

The above results suggest that spousal conflict over how much to save can give rise to inefficient savings behavior. However, costly individual accounts may also be valuable because they can be used to hide resources from spouses. I designed the experiment to explicitly address this alternative theory and to assess whether the main results are robust to accounting for intrahousehold information sharing. I find evidence that information matters – households in which spouses are poorly informed about one another's finances at baseline are more likely to choose individual accounts and less likely to consent to a randomized information sharing treatment. However, these concerns are unrelated to the initial findings regarding preference heterogeneity – well matched couples have no better information flows than poorly matched couples, and the main results are unchanged when accounting for intrahousehold information sharing.

Ideally, the experiment would have randomly assigned discount factors to individuals in couples. This is, of course, not feasible, so much of the analysis relies on heterogeneous treatment effects. It is important to note that I pre-specified the theoretical framework, my focus on heterogeneity in discount factors, and the associated heterogeneous treatment effects before implementing the experiment or undertaking any data analysis.[4] That said, even though the results are very robust to including controls for observable characteristics, one can never completely rule out the hypothesis that the results are actually driven by unobservables that are correlated with discount factor heterogeneity. This paper also requires the external validity caveat common to many field experiments – the couples in my sample reside in a small number of communities in Western Kenya, so additional research would be needed to verify whether preference heterogeneity has similar consequences in other countries and contexts.

The most important contribution of this paper is to shed light on the nature of intertemporal household decision making and the barriers that households face to allocating resources efficiently. Overall, the results are inconsistent with intertemporal efficiency and support the hypothesis that individuals choose savings levels non-cooperatively. However, the results also underscore that not all couples bear substantial efficiency losses – indeed, when couples have similar rates of time preference, they respond robustly to relative rates of return on savings accounts. This suggests that some of the differences observed in household efficiency across different studies and geographical contexts could be due to marriage markets and the quality of matches that they produce.

This paper also contributes to a growing literature that studies savings and investment in developing countries. Specifically, recent research indicates that individuals prefer to use costly but secure informal savings devices over less costly, less secure alternatives (such as saving at home) in order to protect resources from other members of the household, especially spouses (Anderson

[4]Documents detailing this plan were drafted in May and August of 2009 and are archived in the J-PAL hypothesis registry. These documents can be downloaded here: `https://sites.google.com/site/sschaner/home/Preanalysis.zip?attredirects=0&d=1`

and Baland 2002).[5] My results provide a rationale for this behavior and suggest that preference heterogeneity plays a key role in determining how, and how efficiently, people save. Moreover, many households in developing countries engage in entrepreneurial activities – for these households, business investment is an important tool for transferring resources over time. Seen in this light, this paper's insights are also relevant for understanding the widely noted heterogeneity in returns to household-run microenterprises (McKenzie and Woodruff 2008; de Mel et al. 2009; Fafchamps et al. 2012).

The remainder of the paper proceeds as follows: Section 2 develops the theoretical framework that I use to structure the empirical analysis. Section 3 outlines the experimental design, Section 4 presents main results, Section 5 extends the analysis to account for hidden information, Section 6 discusses other alternative explanations, and Section 7 concludes.

2 Theoretical Framework and Testable Predictions

2.1 Economic Environment and Decision Structure

This section develops a simple model to (i) illustrate how heterogeneity in time preferences can generate inefficient strategic savings behavior and (ii) derive a series of testable implications that I take to the data in Section 4. I model a household comprised of two agents (i.e. spouses), A and B. To focus on strategic behavior stemming from differences in savings motives, I assume that both spouses have identical utility functions defined over a single public consumption good, c (this way the only choice that the household must make is *when* to consume).[6] The spouses live in a two period world and exponentially discount per-period utility. The individual discount factor for agent i is δ_i and I assume that the per-period utility function $u(\cdot)$ is continuous, increasing, and concave with $u'(c) \to \infty$ as $c \to 0$.

At the beginning of each period the household receives a deterministic endowment, y_t, which can be consumed or saved (I assume there is no borrowing in this economy).[7] Households have access to three different savings technologies: (1) a public (i.e. joint) bank account, which yields rate of return R_J, (2) a private (i.e. individual) bank account for agent A, which yields rate of return R_A, and (3) a private bank account for agent B, which yields rate of return R_B. What makes

[5] Of course, there are other reasons why individuals may value informal savings devices. Other leading drivers of demand include a need to protect savings from oneself (as in Ashraf et al. 2006) and a need to protect savings from appropriation by members of the community (as in Baland et al. 2007; Jakiela and Ozier 2012).

[6] In practice, households must contend with both heterogeneity in intertemporal preferences and heterogeneity in *what* to consume in a given period. I discuss this issue in subsection 2.3.

[7] Alternatively, one could assume that the husband and wife receive separate endowments that they have exclusive access to. As long as savings constraints do not bind (i.e. an individual never saves his or her entire endowment), the analysis would be unchanged.

the "public" account public is that any member of the household can deposit and withdraw funds. In contrast, only the owner of a "private" account can access that account to make deposits and withdrawals.

I also assume that accounts have time and travel costs associated with them, which I refer to as "banking costs". Specifically, an individual must pay a banking cost of $b > 0$ every time he or she travels to the bank to transact. This is meant to capture the fact that financial markets in developing countries are often characterized by very high transaction costs (Karlan and Morduch 2010) – this is especially true for individuals in my study sample, who live in rural areas outside the town that hosts the bank branch.

Within each period, the following sequence of events occurs:

1. The endowment, and returns on any previous-period savings are revealed.
2. Nature selects which of the two spouses will have the first opportunity to travel to the bank. This sequential ordering of trips is meant to capture the fact that the bank is located in town, and opportunities to go to town may arrive at different times for different spouses.
3. The first mover chooses whether to pay banking cost b to go to the bank. The first mover can only deposit/withdraw from his or her individual account and the joint account.
4. The second spouse observes the choices of the first mover and decides whether to pay b to go to the bank. The second mover can only deposit/withdraw from his or her individual account and the joint account.
5. Any unsaved resources are consumed and the period ends.

I assume that both spouses have perfect information about endowments, interest rates, banking costs, and one another's savings choices.

Given this setup, the objective is now to study how heterogeneity in discount factors impacts the efficiency of household savings behavior. The answer to this question will depend on how spouses make decisions. For example, if spouses bargained cooperatively with one another and were able to commit to future actions, households would never save in a bank account with a dominated interest rate. In order to explore strategic behavior, I therefore assume that agents cannot commit to future actions, and cannot commit to sanction a spouse for past behavior. Given this, I study subgame perfect Nash equilibria to the savings game outlined above.

2.2 Equilibrium Savings Strategies

In what follows, I restrict my attention to pure strategy equilibria. I further assume that if more than one pure strategy equilibrium exists, the couple will never choose a Pareto dominated equilibrium.

Note that when $\delta_A \neq \delta_B$ this refines the set of subgame perfect Nash equilibria to those that generate a unique consumption allocation $\{c_1^*, c_2^*\}$.[8]

I solve the game by working backwards. For expositional clarity, I will refer to the first mover as agent A or the husband and the second mover as agent B or the wife. The solution to the second (and final) period problem is straightforward – since there is no opportunity to consume in the future, individuals will withdraw their bank savings and consume all available resources. In the first period, the second mover will make savings choices that account for this fact, taking as given the earlier choices of the first mover. Since the second mover has the "last word" on savings allocations, she will always invest as efficiently as possible. Moreover, if the second mover saves in account a, she will always allocate savings so that $u'(c_1) = R_a \delta_B u'(c_2)$.

The first mover takes account of this behavior when forming his savings strategy. It is easy to see how preference heterogeneity can cause inefficiencies by considering an example where the second mover's individual account bears the highest rate of return ($R_B = R_{max} \equiv \max\{R_A, R_B, R_J\}$) and strictly dominates the joint account ($R_J < R_{max}$) and the first mover's individual account ($R_A < R_{max}$).

Suppose parameters are such that if the first mover does nothing, the second mover will save in her individual account so that $u'(c_1) = \delta_B R_{max} u'(c_2)$. If the couple is *perfectly matched* on discount factors ($\delta_A = \delta_B$), the first mover will be happy to let the second mover save, as the same allocation maximizes both agents' utilities. In contrast, consider the case where the first mover is *more patient* than his wife ($\delta_A > \delta_B$) – that is, the first mover wishes to save more at R_{max} than his spouse will allow. Instead of letting his wife save, the first mover could consider saving more at a dominated interest rate (R_A or R_J) – as long as this strategy increases second period consumption, a sufficiently patient first mover will find it attractive. Note that while the first mover can deposit as much as he pleases in his individual account (since his spouse cannot withdraw funds), if he deposits "too much" in the joint account his wife will pay b to go back to the bank, withdraw the joint savings, and reoptimize. Thus, while preference heterogeneity can lead to use of both dominated joint and individual accounts, the security of individual accounts make them more attractive from a strategic perspective, all else equal.

An interesting insight of the model is that the *less patient* spouse may also engage in inefficient savings behavior. Suppose instead that $\delta_A < \delta_B$. The first mover knows that if he does nothing, his wife will save in her individual account when it is her turn to go to the bank. Since the first

[8] Multiple pure strategy equilibria will exist when the first mover is indifferent between different strategies (this can occur due to the transaction costs of saving). However, if the first mover is indifferent between one strategy generating $\{c_1^*, c_2^*\}$ and another strategy generating a different consumption allocation $\{c_1^{**}, c_2^{**}\}$, the second mover must have a strict preference between the two consumption allocations, since the second mover has a different discount factor. Thus, the refinement implies that when otherwise indifferent, the first mover will select the strategy that generates the highest utility for his or her spouse.

mover does not value the second period very much, he would like to find a way to reduce first period savings. To accomplish this, he could consider saving just enough (in either the joint or his individual account) to make the second mover indifferent between going back to the bank and staying at home. The idea behind this strategy is that even though the second mover would like to save more, she also wishes to avoid incurring additional banking costs. Thus, transaction costs make it possible for the impatient spouse to preempt additional savings on the part of the more patient spouse. Appendix A describes potentially optimal savings strategies in more detail.

The preceding discussion has illustrated that the equilibrium to the savings game need not be efficient – whenever the couple makes use of a bank account with $R_a < R_{max}$, simply reallocating savings to the highest return account would lead to a Pareto improvement. Moreover, the theory makes an even stronger prediction: it is possible to obtain a comparative static where the extent of the interest rate inefficiency monotonically increases in preference heterogeneity. To see this, define the "interest loss", L, to be the difference between the highest possible interest rate and the average interest rate earned on bank savings. In other words:

$$L \equiv \begin{cases} R_{max} - \sum_a \frac{R_a s_1^a}{\sum_a s_1^a} & \text{if the couple saves} \\ 0 & \text{if the couple does not save} \end{cases}$$

Given the structure of the game, the second mover will always save as efficiently as possible (conditional on the choices of the first mover). Thus, the first mover effectively has control over the magnitude of the interest loss. The next proposition shows that the first mover will implement increasingly inefficient savings strategies as his or her discount factor diverges from that of the second mover:

Proposition 1 *Consider a couple with access to interest rates $\{R_A, R_B, R_J\}$ who are playing a pure strategy, non Pareto dominated, subgame perfect Nash equilibrium to the savings game. Fix endowments $\{y_1, y_2\}$ and transaction costs $\{b, w\}$, as well as the discount factor of the second mover. When a couple is perfectly matched on discount factors ($\delta_A = \delta_B$), $L = 0$. L must increase as δ_A monotonically diverges from δ_B.*

Proof. See Appendix A. ■

This result is intuitive – as discount factors diverge, spouses are willing to pay more to exert control over the time path of consumption. Since the discount factor of the second mover is fixed, the first mover can always choose from the same set of savings strategies. Therefore as preference heterogeneity increases the first mover is willing to use increasingly costly savings strategies to exert control, which means L must get increasingly larger.

2.3 Caveats

The model developed above is highly stylized and focuses on a single motivation for inefficient savings behavior. One key assumption is that *how* resources are saved (e.g. individually versus jointly) has no impact on the within-period consumption allocation (this is by default, since all consumption in the model is public). In practice, households allocate resources over a range of different public and private consumption goods. In this context, another reason spouses may save inefficiently is if the type of account used for saving impacts the within-period consumption allocation, even conditional on the aggregate level of consumption. This would be the case if saving individually increased individual bargaining power, or if saving individually helped spouses hide resources from their partner. These forces would give spouses in couples with perfectly matched time preferences (but imperfectly aligned preferences within-period) incentives to save in lower-return individual accounts. These forces seem plausible *ex-ante* (indeed, I explicitly designed the experiment to test for informational considerations). I therefore discuss empirical support for these alternative hypotheses in Sections 5 and 6.

Moreover, while it makes a strong prediction, one drawback of Proposition 1 is that it holds conditional on incomes, banking costs, and the discount factor of the second mover. Couples are formed by matches made on the marriage market, so it is not feasible to isolate plausibly exogenous variation in preference heterogeneity. Therefore a key concern is that preference heterogeneity may be correlated with other determinants of household savings behavior, such as incomes or overall patience. I address this in two ways: in what follows I derive one testable prediction that holds *unconditional* on the above-mentioned factors, as well as two predictions that hold conditionally. When evaluating the predictions empirically, I also progressively control for a wide range of observable characteristics of couples to test the robustness of the results.

2.4 Testable Predictions

To map the model to the experimental context, consider a sample of couples who are characterized by some distribution of income, banking costs, and spousal discount factors $\{\delta_A, \delta_B\}$. These couples are randomly assigned different sets of interest rates $\{R_A, R_B, R_J\}$.

I begin with a simple prediction that holds unconditional on income, banking costs, and other economic parameters:

T1. Perfectly matched couples ($\delta_A = \delta_B$) will always save efficiently (in the sense that $L = 0$), whereas poorly matched couples ($\delta_A \neq \delta_B$) will not.

The next testable prediction characterizes patterns of individual and joint account use with respect to preference heterogeneity. It is important to keep in mind that *both* individual and joint accounts

may be used strategically. Thus, overall patterns of individual versus joint account use are ambiguous. However, consider the subset of couples for whom the joint account offers the highest rate of return. Since the interest rate loss L is increasing in discount factor heterogeneity by Proposition 1, it must be that *for this subset of couples* individual account use increases with discount factor heterogeneity.

T2. Consider the subset of couples for whom the joint account offers the highest rate of return. Conditional on banking costs, the discount factor of the second mover, and income, rates of individual account use will increase with preference heterogeneity.

This testable prediction would also be consistent with a model where couples save according to different rules of thumb. However, this alternative can be ruled out by the final testable prediction:

T3. Conditional on banking costs, the discount factor of the second mover, and income, the interest rate loss, L, will increase in preference heterogeneity.

An attractive property of the testable predictions is that there is no mechanical reason for all three to hold simultaneously. This is especially useful since the bulk of my analysis relies on heterogeneous treatment effects – the likelihood that an alternative mechanism generates all three patterns in the data is lower than the likelihood that an alternative mechanism generates patterns consistent with a subset of the predictions. I now describe the field experiment and the data used to test these predictions.

3 Experimental Design and Data

3.1 Experimental Design

Context The experiment took place in Western Province, Kenya, in areas surrounding the town of Busia. Busia is a commercial trading center straddling the Kenya-Uganda border. The town is well served by the formal banking sector, hosting over six banks at the time of field activities. The financial partner for this study is Family Bank of Kenya. At the time of the experiment the bank had over 600,000 customers, 50 branches throughout the country, Ksh 13 billion (approximately $163 million at an exchange rate of Ksh 80 per $1) in assets, and actively targeted low and middle income individuals as clients.

All study participants were offered Family Bank's *Mwananchi* accounts. This account could be opened with any amount of money, though a minimum operating balance of Ksh 100 ($1.25) could not be withdrawn. The account paid no interest, but deposits were free of charge and there were no recurring maintenance fees. The only fees associated with the account were withdrawal

fees, which were Ksh 62 ($0.78) over the counter and Ksh 30 ($0.38) with an ATM card. Account holders could purchase an ATM card for Ksh 300 ($3.75), though this was not mandatory.

At the outset of the study, I identified communities surrounding 19 local primary schools, which would serve as group meeting grounds (all experimental activities were conducted in these group sessions). These communities were situated either on the outskirts of Busia town or in nearby rural areas. Trained field officers issued meeting invitations to married couples in these areas where (1) neither spouse had an account with Family Bank but at least one spouse was potentially interested in opening one and (2) both spouses had national ID cards and were able to attend the meeting.[9] Just 7 percent of otherwise eligible couples were excluded due to pre-existing ownership of Family Bank accounts and approximately 29 percent of issued invitations were redeemed over the course of the study. Thus while far from universal, takeup rates are high enough that the sample represents a nontrivial fraction of targeted married couples in the catchment area.

Interventions All participating couples were given the opportunity to open up to three Family Bank accounts: an individual account in the name of the husband, an individual account in the name of the wife, and a joint account. To maximize takeup, I funded each opened account with the Ksh 100 minimum operating balance (this amount could not be withdrawn by participants – it simply made opening an account costless). While participants could in principle open an account with Family Bank at any time, only those accounts opened during experimental meetings were eligible for the operating balance subsidy and the experimental interventions, which are described below.[10]

Intervention 1 - Interest Rates Each potential account was randomly assigned an interest rate (respondents drew envelopes with the interest rates from tins upon arrival at the meeting). Since prediction T2 only holds when the joint account offers the highest rate of return, I designed the experiment so that individual accounts could bear either 0, 2, 6, or 10 percent 6-month yields (with equal probability), while joint accounts could bear either 2, 6, or 10 percent 6-month yields (with equal probability). These interest rates were very high compared to market alternatives: small scale savings balances could earn at most 0.5-2.0 percentage points of interest annually given bank accounts available in Busia at the time of the experiment. The experimental interest rates were temporary, and expired after six months.[11] Since many respondents had low levels of

[9]Family Bank (and all other banks in Kenya) require that account holders have national ID cards. The majority of individuals in Kenya have a national ID card as it is legally required of all adult citizens and necessary in order to vote, buy or sell land, and seek formal employment.

[10]A subset of opened accounts were also randomly selected to receive free ATM cards. A description and analysis of this treatment is presented in Schaner (2013). I do not discuss this intervention further, as accounting for it has no impact on any of the results presented in this paper.

[11]In order to make interest rates as salient as possible, couples were given reminder cards for each account that they opened. All cards, including those given to individuals opening accounts that did not bear any interest, featured a

education, enumerators explained what an interest rate was and provided numerical examples for each interest rate that was drawn. The empirical results show that respondents understood and robustly responded to the interest rates (see Appendix Table D3).

The three interest rate draws were independent of one another, and therefore created random variation in the relative rates of return between accounts, even conditional on an account's interest rate. I use what I refer to as the "excess interest rate" to capture this variation:

$$excess_a = R_a - \max\left\{R_j : j \neq a\right\}$$

Conditional on R_a the experiment created 10 percentage points of random variation in $excess_a$ for each account type.[12] After observing their interest rates, couples were separated and each spouse was administered a baseline survey. One concern is that randomizing the interest rates before conducting the baseline influenced survey responses. However, interest rates are not systematically associated with elicited discount factors, baseline self reports of savings levels, savings device use, or self reported decision making power regarding consumption and saving. It is therefore likely that the randomization had little impact on survey responses. After the baseline, couples were reunited and decided which accounts to open.[13]

Intervention 2 - Extra Statements In order to test whether the ability to hide savings was an important driver of individual account use, 50 percent of participating couples were randomly selected for an "extra statements" offer.[14] If a selected couple decided to open an individual account for (without loss of generality) spouse A, the enumerator processing the couple's paperwork asked if the spouses would consent to allow spouse B to receive extra statement cards. The cards, if presented by spouse B at the bank, entitled him or her to learn the current balance of spouse A's account. These cards were only valid for 6 months, and were not given to couples unless both spouses gave their consent.

reminder to save. The interest payments were made by IPA-Kenya and after the six month period, balances earned no interest (respondents were informed of this *ex-ante*), which is standard for the *Mwananchi* account and other current accounts in Kenya.

[12] Appendix Figure D1 illustrates the interest rate design, including the variation in the excess interest rates.

[13] The fact that couples decided which accounts to open together could have impacted account opening choices. In particular, this could have increased the degree of cooperation between spouses – this would reduce the likelihood of observing inefficient savings behavior in the sample. The public decision making could have also nudged couples to open joint accounts over individual accounts, in which case my results will understate the preference for individual accounts in the sample.

[14] Due to delays in approvals from the bank, extra statements were not offered to the 98 couples (16 percent of the sample) in the first 6 (of 33) experimental sessions.

3.2 Data

I use two data sources in the analysis – data from one-on-one baseline surveys (spouses were separated for the interviews), and administrative data on account use from the bank. The administrative data provided by the bank includes the first six months' transaction history for all accounts opened under the auspices of the project. The baseline survey collected basic demographic information, as well as information on rates of time preference, decision making power in the household, income, current use of a variety of savings devices, and cross reports of spousal income and savings.

Measuring Rates of Time Preference The baseline elicited time preferences using choices between different amounts of money at different times, as opposed to different amounts of goods at different times. I made this choice for three reasons. First, Ashraf et al. (2006) find that while time preference parameters estimated using choices between money, rice, and ice cream were all correlated, only the parameters estimated using money choices significantly predicted takeup and use of a commitment savings product. Second, even though discount rates estimated using money choices should theoretically reflect external interest rates, in practice respondents do not appear to take account of this when making choices (see Andreoni and Sprenger 2010 for a summary). Finally, cash lotteries made intuitive sense to respondents given that the group sessions revolved around bank accounts and savings.

All questions were framed as a choice between a smaller amount of money at a nearer time t (x^t) and a larger amount of money at a farther time $t+\tau$ ($x^{t+\tau}$).[15] In total, participants responded to 10 tables of monetary choices, with each table consisting of 5 separate choices between a smaller Ksh $x^t \in \{290, 220, 150, 80, 10\}$ and larger $x^{t+\tau} =$ Ksh 300.[16] In order to make decisions salient, respondents were given a 1 in 5 chance of winning one of their choices (the choice was also selected at random). The questions involved sizable amounts of cash relative to respondents' incomes – for comparison, median reported weekly earnings in the sample were Ksh 700 for men and Ksh 300 for women. If a respondent won one of her choices, she had the option of having the funds deposited directly in her bank account, or picking the cash up at the field office, also located in Busia town.[17]

[15] This method is common to most empirical studies that attempt to measure rates of time preference in developing countries. Examples include Ashraf et al. (2006), Bauer and Chytilova (2009), Shapiro (2010), Tanaka et al. (2010), and Dupas and Robinson (2013).

[16] The 10 $(t, t+\tau)$ pairs were: $\left(\frac{1}{7}, 1\right)$, $\left(\frac{1}{7}, 2\right)$, $\left(\frac{1}{7}, 3\right)$, $\left(\frac{1}{7}, 4\right)$, $\left(\frac{1}{7}, 8\right)$, $\left(\frac{1}{7}, 12\right)$, $(2,3)$, $(2,4)$, $(4,8)$, and $(4,12)$ weeks. I chose to set the lowest near term t to "tomorrow" $\left(\frac{1}{7}\right)$ instead of "today" (0) to avoid confounding the discount factor estimates with differences in transaction costs of obtaining the funds in the near versus far term, or degrees of trust as to whether the money would be delivered (Harrison et al. 2004). I also assume that all respondents would prefer Ksh 300 in the future to Ksh 0 sooner, and that all respondents would prefer Ksh 300 sooner to Ksh 300 in the future. Adding these imputed responses leaves 70 choices for each individual.

[17] The majority of cash winners (79 percent) chose to have their payments deposited in a bank account. The bank account may have been attractive because the respondents did not have to remember to pick up the funds at any specific time, because the bank was more conveniently located, or because the individuals intended to use their new accounts

As in Tanaka et al. (2010), I use nonlinear least squares to estimate the discount factors. For each individual I assume that utility is linear in money amounts over the range Ksh 0-300. Then the utility gains of the near and far amounts for person i considering choice q can be expressed as $\Delta U_i\left(x_q^t\right) = \delta_i^t x_q^t$ and $\Delta U_i\left(x_q^{t+\tau}\right) = \delta_i^{t+\tau} x_q^{t+\tau} + \varepsilon_{iq}$ where $\varepsilon_{iq} \sim Logistic\left(0, \mu_i\right)$. Define the dummy variable $now_{iq} = 1\left(x_q^t \succ x_q^{t+\tau}\right)$. Nonlinear least squares solves

$$\left(\hat{\delta}_i, \hat{\mu}_i\right) = \arg\min_{\delta_i, \mu_i} \sum_{q=1}^{70} \left(now_{iq} - \frac{1}{1+\exp\left(-\mu_i\left(\delta_i^t x_q^t - \delta_i^{t+\tau} 300\right)\right)} \right)^2$$

I topcoded $\hat{\delta}_i$ at $\overline{\delta}$, the value of $\hat{\delta}_i$ obtained via nonlinear least squares for always-patient responses and bottomcoded $\hat{\delta}_i$ at $\underline{\delta}$, the value of $\hat{\delta}_i$ for always-impatient responses.[18]

Panels A and B of Figure 1 graph the distribution of estimated discount factors for men and women. Discount factors span a wide range of values, but on average study participants appear to be very impatient, with weekly discount factors averaging 0.72 for men and 0.70 for women. These discount factor estimates are lower than estimates in studies of individuals in developing countries in Asia, but consistent with other studies in Africa that have found very high rates of impatience regarding the timing of cash payments.[19]

The histograms also illustrate the extent of censoring in the sample. First, 14 percent of individuals were "always impatient", and preferred Ksh 10 sooner to Ksh 300 in the future in all tables. Nonlinear least squares converges to a discount factor estimate very close to zero for this group. Another 24 percent of individuals were "always patient" and preferred Ksh 300 in the future to all sooner amounts. In general, this measurement error will lead me to overestimate match quality in couples with censored discount factors, which should bias the empirical results away from the testable predictions.

My baseline measure of intrahousehold preference heterogeneity is simply the difference between the male and female estimated discount factors: $het_c \equiv \hat{\delta}_{Mc} - \hat{\delta}_{Fc}$. While 13 percent of couples had identical discount factor estimates, many couples had estimates that differ substantially. In the analysis I frequently compare the behavior of well and poorly matched couples. To do this I label the 50 percent of the sample with the most closely aligned discount factors as "well matched" and refer to the remaining couples as "poorly matched". This corresponds to $\left|\hat{\delta}_{Mc} - \hat{\delta}_{Fc}\right| \leq 0.193$ and is equivalent to the couple's discount factors being within 0.528 standard deviations of one another.

for saving anyway.

[18]This led to the censoring of 17 estimated discount factors from below and 37 estimated discount factors from above.

[19]For Asia see Ashraf et al. (2006) Bauer and Chytilova (2009), Shapiro (2010), and Tanaka et al. (2010)). For Africa see Dupas and Robinson (2013) and Giné et al. (2011).

Panel C of Figure 1 shows a weighted scatter plot of $\hat{\delta}_M$ and $\hat{\delta}_F$, with well matched couples indicated by darker shading. The figure shows that there is only weak assortative matching in my sample: the correlation coefficient between spousal discount factors is 0.08. The figure also shows that individuals in well matched couples are more patient on average (also see Appendix Table D1). Since a couple's overall patience level could be an important determinant of savings behavior, in what follows I explicitly show that controlling for the level discount factors of the husband and wife does not substantively impact my empirical results.

It is likely that my discount factor estimates reflect multiple determinants of savings motives, including discount factors, risk aversion, state dependence, and so on. Given this, I prefer to interpret the results in terms of heterogeneity in savings motives more broadly, rather than heterogeneity in discount factors specifically. A bigger concern would be if the measure of heterogeneity is correlated with other characteristics of couples that determine savings behavior for reasons unrelated to differences in savings motives (e.g. a general ability of spouses to work well together). To address this issue I present specifications that control for spousal *levels of* and *heterogeneity in* observable characteristics. Also note that regardless of interpretation, my measure of preference heterogeneity is still useful for testing a null hypothesis of efficient households: if all households make efficient investment choices, I should not observe differences in efficiency by match quality.

Sample Characteristics and Randomization Verification The sample consists of 598 non-polygamous married couples.[20] Table 1 presents baseline characteristics of the sample and checks that these characteristics are uncorrelated with treatments. Respondents are of relatively low socioeconomic status, but almost all respondents reported using at least one savings device at baseline, with saving at home and saving with ROSCAs (rotating savings and credit associations – a type of informal savings group found in much of the developing world) most common. Formal bank and mobile money accounts were less prevalent, particularly among women − while 30 percent of men reported owning a bank account at baseline, just 12 percent of women reported owning a bank account. Over half of men and women also reported other informal savings (this was mostly investment in livestock).

I check randomization by running individual level regressions of each characteristic listed in Table 1 on five treatments of interest: the excess interest rate on the husband's, wife's, and joint account, as well as the extra statements treatment and a dummy for whether the individual was selected for a cash payment.[21] Columns 3-7 of Table 1 report regression coefficients and standard

[20] I dropped 179 polygamous couples from the sample since strategic behavior may be very different in households with more than one female head. However, the results are robust (though in some cases somewhat attenuated) to including them.

[21] I present randomization verification results for excess interest rates rather than the interest rates themselves because I focus on excess interest rates in the analysis. Randomization verification results for the interest rates themselves

errors for each treatment (rows correspond to a single regression). I also estimate all equations jointly via seemingly unrelated regression to test whether each treatment is significant across all equations. The p-values from these joint tests appear in braces in the last row of Table 1.

Overall, the randomization appears to have functioned well. The number of coefficients significant at the 90 percent level or higher is approximately equal to the expected number due to chance, and there are no systematic patterns across the different treatments. Importantly, excess interest rates are uncorrelated with discount factor heterogeneity.

The joint test does suggest that balance is off for cash payments, however. Additional randomization verification (see Appendix Table D2) shows that significantly *fewer* cash payments were awarded than expected (17.5 percent instead of the expected 20 percent). It is important to note that all randomization was conducted in the field, by allowing respondents to draw folded envelopes from tins. Since fewer cash payments were awarded than expected, this suggests that the lack of balance is due to chance rather than enumerator deviations from the experimental protocol. (Enumerators conducting randomization were supervised throughout the study and carefully trained not to allow a participant a second draw from the tin if the participant was unhappy with his or her result.) Fortunately, cash prize receipt is uncorrelated with preference heterogeneity and the excess interest rates. Thus there is no mechanical reason why these payments should impact the main testable implications. However, in the interest of caution I control for husband and wife cash payment selection throughout the analysis. As expected, the results are extremely similar if I omit these controls. Furthermore, results are similar if I simply drop couples selected for at least one cash payment or exclude the cash payment deposit from measures of account use.

Basic Overview of Account Use Before beginning the main analysis, Table 2 summarizes well and poorly matched couples' use of the experimental bank accounts. The first panel describes account opening choices. While all couples opened at least one account, very few couples opened all three. Instead the most popular choices were either only opening a joint account (56 percent of couples) or opening two individual accounts (29 percent of couples). Even though there was no monetary cost to open all three accounts, the additional time spent doing paperwork may have been enough to dissuade couples from opening accounts that they were very certain they would never use – since enumerators explained that the Ksh 100 opening balance could not be withdrawn from the accounts, there was no strategic reason to open all three accounts in order to earn additional cash.

I exclude the Ksh 100 minimum balance from all measures of account use – thus, unopened and opened but unused accounts are treated equivalently. As motivated by the theoretical framework,

are very similar. The extra statements treatment dummy is coded to zero for the individuals in the first six sessions who were not eligible for the treatment. I therefore include an additional dummy variable that identifies these individuals in all regressions.

a key outcome is whether an account was used for saving – I define a couple to have saved in account a if at least one deposit (other than the initial minimum balance subsidy) was made in the first 6 months following account opening. The next two panels show that 43 percent of couples saved in at least one account. This figure drops to 27 percent when deposits for the discount factor elicitation payoffs are excluded.[22] The 43 percent of couples who do save make an average of 3 deposits in the six months following account opening and hold an average daily balance of just under Ksh 1,000 ($12.50) in their experimental accounts.

Based on Table 2, the savings behavior of well and poorly matched couples appears to be quite similar. However, these summary statistics are not very useful for evaluating the theory in Section 2. First, recall that individuals in poorly matched couples may make strategic use of both joint and individual accounts – so simply comparing overall rates of joint and individual account use by match quality is not instructive. Moreover, since both more and less patient spouses may save strategically, there are no testable predictions regarding average balances or number of deposits in accounts. Rather, it is necessary to evaluate the theory by focusing on how efficiently resources are saved. The next section therefore turns to the testable predictions to assess support for the the ideas in Section 2.

4 Main Results

4.1 Unconditional Testable Predictions

I begin the empirical analysis with the simplest testable prediction, which holds unconditional on spousal discount factors and other determinants of savings propensity. This first prediction is that perfectly matched couples always save efficiently, while poorly matched couples do not. Before moving to the analysis, it is necessary to address one complication in assessing the efficiency of couples' savings choices. To simplify the theory, I assumed that banking costs were nonstochastic and the same for all accounts. In practice, the marginal cost of going to the bank is low when an individual is in town for another reason, but high when an individual must make a trip to town specifically to go to the bank. In such a context, a joint bank account offers an important advantage: the couple can always send the spouse with the lowest cost of going to town to the bank. Consequently it may be efficient for a couple to choose a joint account even when the joint interest rate is dominated by that of an individual account.

A simple way to address this issue is to study the subsample of couples where the joint account offers the highest rate of return ($R_J = R_{max}$) – in this case use of an account with a dominated

[22] For the main analysis I include the cash payments in measures of account use, since efficient households should always invest these payments in the highest return account. The results are, however, robust to using measures of account use that ignore these payments.

interest rate is clear evidence of inefficient savings behavior. In practice, I see that just 3.9 percent (7 of 180) of well matched couples in this subsample incur a positive interest rate loss, whereas 12.6 percent (23 of 183) of poorly matched couples incur a positive interest rate loss. The difference in these proportions is significant at the 99 percent confidence level.

This result is also consistent with the hypothesis that well and poorly matched couples use different rules of thumb for saving (i.e. "save jointly" vs. "save individually"). Fortunately, one can use the fact that efficiency requires couples to respond to relative rates of return between accounts (as measured by $excess_a$) to derive a more novel way to assess savings efficiency. Specifically, imagine a population of couples with heterogeneous banking costs. Fix individual interest rate R_i. Panel A of Figure 2 graphs the share of this population that would efficiently save in individual account i at different levels of $excess_i$ (note that increases in $excess_i$ correspond to decreases in R_{-i} and R_J). Since the individual account has higher transaction costs it will never be efficient to save in account i if $excess_i < 0$. If there is a mass of households who have equivalent individual interest rates that dominate the joint interest rate, there could be a discrete jump up in the savings rate at an excess rate of zero. Since the field experiment generated lumpy variation in the set of interest rates presented to couples, the existence of such a mass seems reasonable. As the excess interest rate increases beyond zero, the share of households saving in account i will increase, as the individual account comes to dominate more and more joint accounts. At some point, the excess rate on the individual account will become so high (i.e. R_J will become so low) that account i dominates the joint account for all couples who are willing to save at R_i. On the graph, this occurs at $excess_i = E_i(R_i)$.

Panel B of Figure 2 illustrates efficient savings responses for the joint account. The shape is similar to that for individual accounts – however, since the joint account has lower banking costs than individual accounts, it is efficient for couples to begin to save jointly at negative excess interest rates. On the graph, the lowest excess rate at which any couple is willing to save jointly is $excess_J = -E_J(R_J)$, and all households will choose the joint account once $excess_J$ reaches zero. Notice that there is a positive slope only for *positive* excess rates when considering individual accounts and a positive slope only for *negative* excess rates when considering joint accounts. This asymmetry is a striking implication of efficient investment in the presence of heterogeneous banking costs.

I generate the empirical analog of Figure 2 by running the following regression separately by account type (individual versus joint) and match quality:

$$saved_{ac} = \beta_0 + ex'_{ac}\delta + z'_{ac}\lambda + \varepsilon_{ac} \tag{1}$$

Where $saved_{ac}$ indicates that couple c saved in account a, ex_{ac} is a vector of dummy variables for the excess interest rate on account a, and z_{ac} is a vector of dummy variables for account a's interest

rate.[23] I then calculate predicted values of $saved_{ac}$ for each value of the excess rate, assuming equal distribution of the sample at each interest rate (0, 2, 6, and 10 percent for individual accounts; 2, 6, and 10 percent for joint accounts). Figure 3 presents the result of this exercise. The dashed lines are regression lines fit to the point estimates, where each point is weighted by the inverse of its standard error. Recall from Figure 2 that individual account use by well matched couples could jump discretely up at $excess_{ac} = 0$. Since -2 is the largest negative value of the excess interest rate in the sample, I therefore fit separate lines for $excess_{ac} \leq -2$ (this slope should be zero) and $excess_{ac} \geq -2$ (this slope should be positive). In contrast, the slope for joint accounts should be positive below an excess rate of zero and flat thereafter, so the lines are drawn above and below $excess_{ac} = 0$.

Column A graphs account use for well matched couples. The results for individual and joint accounts are both consistent with efficient savings behavior. Specifically, the share of couples using individual accounts is essentially flat and not statistically different from zero until the excess interest rate reaches zero, where the share saving jumps up. Beyond this point the share increases, though there is no subsequent plateau. For joint accounts, the share of well matched couples saving increases when the excess interest rate is negative and then plateaus after an excess rate of zero. However, there is no initial plateau at a low savings rate when the excess interest rate is negative.

These "missing" plateaus suggest that the experimental variation in the excess interest rate was not large enough to reach the thresholds $E_i(R_i)$ and $-E_J(R_J)$ illustrated in Figure 2; in other words, the excess interest rate was never large enough to dominate banking cost concerns for the majority of couples in the sample. This is conceivable given the temporary nature of the interest rates and the low savings levels of most couples – the 75$^{\text{th}}$ percentile total average daily balance *among savers* was Ksh 782. Even foregoing 10 percentage points of interest would only cost the 75$^{\text{th}}$ percentile household Ksh 78. In comparison, the cost of round trip travel to and from the bank meets or exceeds Ksh 100 for over half the sample. I therefore conjecture that relative rates of return only trumped banking cost concerns for couples with small banking cost differentials.

Column B of Figure 3 illustrates the behavior of poorly matched couples. Here rates of saving are completely insensitive to the excess interest rate. It is possible that the variation in the excess interest rate was simply not large enough to incite a substantial change in behavior for these couples. Since interest rates were temporary and average balances small, private savings deviations were "cheap" in terms of interest rate losses.

[23] As a result of the experimental design, some values of the excess interest rate were only realized for a very small number of accounts: 14 accounts had an excess interest rate of 2, 12 accounts had an excess interest rate of 6, and 13 accounts had an excess interest rate of 10. For each of these values, I downcode the excess interest rate by two percentage points (results are invariant to simply dropping these accounts). Similarly, I pool $excess_{ac} = -10$ and $excess_{ac} = -8$ as the omitted category in the regressions. I do this in order to identify all interest rate dummy variables, as accounts with zero percent interest had excess interest rates unique to them.

To test the significance of the patterns in Figure 3, I generate splines in the excess interest rate. To match the theoretical shifts in slope in Figure 2, I place a knot at $excess_{ac} = -2$ for individual accounts and a knot at $excess_{ac} = 0$ for joint accounts. I then run the following account level regression separately for well matched and poorly matched couples:

$$\begin{aligned} saved_{ac} &= \beta_0 + \beta_1 (below \times indiv)_{ac} + \beta_2 (above \times indiv)_{ac} + \\ &\quad \beta_3 (below \times joint)_{ac} + \beta_4 (above \times joint)_{ac} + z'_{ac}\lambda + x'_c\alpha + \varepsilon_{ac} \end{aligned} \tag{2}$$

where $below_{ac}$ is the spline capturing the slope below the knot, $above_{ac}$ is the spline capturing the slope above the knot, $indiv_{ac}$ is an individual account indicator, $joint_{ac}$ is a joint account indicator, z_{ac} is a vector including interest rate dummies, the joint account dummy, and joint$\times$interest rate interactions, and x_c is a vector of additional controls.

Panels A and B of Table 3 present results of this specification first for well matched couples, then for poorly matched couples. The first column only includes additional controls for husband and wife cash prize selection – this column essentially mirrors the results in Figure 3 and confirms that the positive responses exhibited by well matched couples are significantly different from zero. Furthermore, the magnitude of the slopes is large. These results imply that increasing the individual excess interest rate from 0 to 10 would result in a 16 percentage point increase in the savings rate, while increasing the joint excess interest rate from -10 to 0 would result in a 21 percentage point increase in the savings rate. These impacts are very large when compared to dependent variable means (10 percent for individual accounts and 31 percent for joint accounts). In contrast, I cannot reject the null that poorly matched couples did not respond to the excess interest rate.

I also test the robustness of the results in Table 3 to the inclusion of a variety of controls for observable characteristics of couples. First, the second column adds controls related to rates of time preference. To account for general differences in patience levels between well and poorly matched couples I control for $\hat{\delta}_{Mc}$ and $\hat{\delta}_{Fc}$ linearly (the addition of these controls is also motivated by the conditional predictions, which hold given the discount factor of the second mover spouse). I also include two dummy variables to identify censoring of the discount factor for each spouse – one indicating upper censoring and one indicating lower censoring. The second column adds additional controls for demographic characteristics and the final column adds further controls for economic characteristics.[24] For non-binary characteristics I include the linear and squared terms for both husband and wife, as well as the interaction between the linear values for husband and wife. For

[24]Demographic controls include age, years of education, a literacy dummy, number of children, and experimental session fixed effects, which capture distance from the bank and area of residence. The economic controls include individual income, a dummy for mobile phone ownership, and a dummy indicating that an individual is either a subsistence farmer or has no job. When the value of a control variable is missing, I recode it to zero and generate a separate dummy variable to identify these observations. I therefore also include interactions between husband and wife missing dummies in all specifications. This convention is held throughout the analysis.

binary variables, I include the dummy variable for both husband and wife as well as the interaction. This allows me to control for both levels of and *heterogeneity in* observable characteristics. The results in Panels A and B are very robust to adding these controls.

Panel C of Table 3 formally tests whether the excess interest rate responses of well and poorly matched couples differ. This panel also progressively allows for heterogeneous treatment effects with respect to the time preference, demographic, and economic control sets. To do this, I include main effects for each covariate and interactions of the covariate with the excess interest rate splines. There are 68 main effects in the economic control set. If I were to separately interact each main effect with the four excess interest rate splines, this would lead to 340 covariates in the economic control set alone. To reduce the number of covariates, I impose the restrictions $\beta_1 = \beta_4$ ("excess zero" in Panel C) and $\beta_2 = \beta_3$ ("excess positive" in Panel C). This seems reasonable, as I cannot reject these restrictions in Panels A and B (p-values for the relevant F-tests range from 0.51-0.92).

I estimate the restricted version of equation 2 separately by match quality. I demean each included covariate using means *among well matched couples* before generating the excess interest rate interaction terms and running the regressions. When demeaned this way, the coefficients on the splines in Panel C of Table 3 reflect the response to the excess interest rate at the average value of included demographic characteristics observed among well matched couples. Thus if the heterogeneous responses observed in Panels A and B were driven by some other characteristic in the control sets, then the results for poorly matched couples would mirror those for well matched couples once the excess interest rate response is allowed to vary with that control. I use seemingly unrelated regression to test equality of the excess interest rate responses for well and poorly matched couples.

As expected given random assignment of interest rates, the results for well matched couples do not substantially change with the addition of controls. In contrast, the excess interest rate responses for poorly matched couples become more *negative* after allowing for heterogeneous treatment effects with respect to controls. This suggests that differences in observables between well and poorly matched couples are not driving the initial heterogeneous treatment effects. Furthermore, I consistently reject equality of responses to the excess interest rate between well and poorly matched couples across all specifications.

I now turn to conditional predictions T2 and T3 to explore how patterns of account use and actual interest rate losses vary with match quality.

4.2 Conditional Testable Predictions

The second testable prediction states that when $R_J = R_{max}$, couples' use of individual accounts should increase as match quality decreases. (That is, there should be a U-shaped relationship

between individual account use and $\delta_M - \delta_F$). Figure 4 tests this prediction graphically. The figure presents results of the following local linear regression

$$y_c = g\left(het_c\right) + \varepsilon_c \tag{3}$$

where y_c is the outcome of interest and $het_c = \hat{\delta}_{Mc} - \hat{\delta}_{Fc}$. The sample is limited to the subset of couples for whom the joint account offers the highest rate of return. The solid line in Panel A graphs savings rates in any individual account (i.e. a couple is coded as saving individually if either the husband's or the wife's individual account was used for saving). As predicted, rates of individual account use follow a striking U shape – well matched couples (who are demarcated by gray vertical lines in each panel) are least likely to save individually and rates of individual account use increase in preference heterogeneity. This is not just because poorly matched couples are more likely to save – Panel B presents the same graph when the sample is limited to those couples who saved in at least one account. The U-shaped pattern remains, and shows that well matched savers choose to save jointly (the efficient choice), whereas poorly matched savers are much more likely to make use of inefficient individual accounts.

The following regression tests the significance of the graphical results while controlling for determinants of savings propensity and other potentially confounding factors:

$$y_c = \beta_0 + \beta_1 badmatch_c + int_c'\delta + x_c'\lambda + \varepsilon_c \tag{4}$$

where y_c is the outcome of interest and $badmatch_c$ indicates poorly matched couples. The regressions also include a vector of dummy variables for each experimental interest rate (int_c) and a vector of additional controls (x_c).[25]

Table 4 presents estimates of β_1 ("Poorly Matched") and tests robustness of the results to adding additional controls. The "basic" control set, in Panel A, only includes the interest rate dummy variables and dummies for husband and wife cash payment selection. This regression essentially mirrors the results of Figure 4. I then progressively add the same set of time preference, demographic, and economic controls used in Table 3.

As expected given the graphical results, poorly matched couples are significantly more likely to save in individual accounts, even when limiting the sample to savers. Moreover, the differences by match quality are substantial – while just 7.8 percent of well matched couples save individually when $R_J = R_{max}$, over 20 percent of their poorly matched peers choose to save individually. Panel

[25]Note that although $badmatch_c$ is a generated regressor, under the null hypothesis $\beta_1 = 0$. In this case, traditional standard errors are consistent (Newey and McFadden 1994). Since the unit of randomization is the couple, I therefore present either heteroskedasticity robust standard errors (for couple level regressions) or standard errors clustered at the couple level (for account level regressions).

B shows that this pattern persists when limiting the sample to couples who saved. Just 18.2 percent of well matched savers make use of an individual account, whereas over 40 percent of their badly matched peers save individually. These results are quite robust to adding additional controls – even though results for savers are only marginally significant after adding the demographic control set, the point estimates are consistent from specification to specification and large in magnitude.

The results so far strongly suggest that poorly matched couples invest less efficiently than well matched couples. But how big are these distortions in economic terms? I now exploit the experimental design to estimate the magnitude of savings misallocation by match quality. This also permits a test of the final prediction – that the interest rate loss, L, increases with preference heterogeneity. To estimate L, I calculate the actual interest rate that each couple earned on experimental savings balances and subtract it from the maximum interest rate. I code interest rate losses to zero for all non-savers.

Since lower return joint accounts may be more efficient than higher return individual accounts when banking cost differentials are large, I also present results where I discount individual interest rates to reflect higher banking costs. I adjust individual interest rates downward to account for banking costs in two different ways. First, I attempt to proxy banking costs using observables. I conjecture that those couples who travel to Busia town frequently for non-bank related reasons and those couples who have low travel costs to town will have smaller differential banking costs (i.e. the joint account offers less in transaction cost savings).[26] Appendix B provides additional detail on how I used principal components analysis to construct a "banking costs index", which runs from zero (lowest hypothesized banking costs) to one (highest hypothesized banking costs).[27] To discount individual interest rates I multiply the cost index by an assumed maximum interest discount and adjust individual interest rates by the resulting product. As a further robustness check, I present a set of specifications where I discount all individual interest rates uniformly. While this method cannot capture heterogeneity in banking costs within the population, it does capture the fact that individual accounts incur higher banking costs without inducing a correlation between the size of the discount and other determinants of savings propensity correlated with the banking cost index.

I first study interest rate losses graphically. Figure 5 presents the results of local linear regressions of interest rate losses on preference heterogeneity. Consistent with the theory, losses follow a U-shape, with the lowest values observed among well matched couples. This pattern holds with and without adjustments to individual interest rates. (The figure illustrates results using adjust-

[26] In practice I assume that subsistence farmers and the unemployed have higher banking costs. I also assume that couples who live closer to the bank and couples with pre-existing formal savings accounts have lower differential banking costs. Here I categorize both bank accounts and SACCO accounts as formal accounts.

[27] Appendix B also shows that savers with higher proxied costs are more likely to use joint accounts, and that excess interest rate responses are concentrated among couples with low proxied costs.

ments made with proxied banking costs. Results are very similar when individual interest rates are adjusted uniformly).

Table 5 quantifies interest rate losses by match quality and presents regression results that include additional controls for observables. The first column does not discount individual interest rates at all. If poorly matched couples had always chosen the highest return account available, the average couple would have earned 8.15 percentage points of interest. In practice, these couples averaged 7.10 percentage points of interest, leading to an interest loss of 1.05 percentage points. In contrast, well matched couples could have earned a maximum of 8.17 percentage points of interest and actually earned 7.50 percentage points. Therefore, the "loss gap" between poorly and well matched couples is 0.39 percentage points of interest, which is significantly different from zero. Even without accounting for differential banking costs, poorly matched couples appear to suffer from greater savings misallocation – their losses are 58 percent larger than those of their well matched peers. Subsequent rows of Table 5 test robustness of this result by first controlling for cash prize selection and account specific interest rates. I then include the time preference, demographic, and economic control sets respectively. The loss gap remains significant and its magnitude grows in size after including additional controls.

The remaining columns repeat this analysis using banking cost adjusted individual interest rates. Columns 2-4 use the banking cost index and a maximum individual interest discount ranging from 5 to 15 percentage points (recall that well matched couples' response to the excess interest rate in Figure 3 suggests that the maximum discount exceeds 8). The estimated loss gap increases as discounting increases and is robust to including additional controls. The final three columns discount all individual interest rates uniformly (I subtract the enumerated discount from the interest rate for all individual accounts, regardless of proxied banking costs). These results are quite similar to the results incorporating proxied banking costs.

Overall, the losses in Table 5 are large in percentage terms, but small in absolute terms. For example, a loss of 3 percentage points in interest amounts to just Ksh 24 ($0.30) for the 75th percentile saving couple. On the other hand, banking cost differentials persist for the life of the account, so long run absolute losses due to inefficient individual account use could be much larger.

The core empirical results fit the predictions of the theoretical framework very well. While well matched couples do a good job of saving efficiently, poorly matched couples do not account for rates of return between accounts and tend to make use of inefficient individual accounts. These patterns appear even though badly matched couples understand the interest rates – the results in Panel C of Appendix Table D3 confirm that the interest rate treatments significantly impacted the account opening and use choices of badly matched couples.

A central idea behind the private savings theory is that inefficient accounts are used to gain *control of resources* rather than obscure *information about resources*. However, individual ac-

counts may also enable agents to hide savings from their spouses. Hidden savings could be used to manipulate the time path of consumption, or to finance hidden consumption, thereby changing the composition of consumption in a given period. If the benefit of hiding savings is correlated with preference heterogeneity, this could generate the patterns I observe in the data. The next section tests this hypothesis.

5 Hidden Information and Account Use

Hidden information appears to be important in households in developing countries.[28] Moreover, there is evidence of hidden savings in my data. For example, when considering spousal cross reports for saving at home, which is the most common method of saving among couples, 52 percent of individuals in the sample asserted that they either did not know if their spouse saved at home or did not know how much money their spouse saved at home. Among those individuals who had a spouse guess his or her savings amount, 48 percent of the spouses underestimated relative to the individual's self report, while 25 percent of the spousal reports were overestimates (the other 27 percent matched their spouse's report – these were mostly instances where the individual and spouse reported no savings). Similar patterns appear when considering weekly income and savings in other devices, such as bank accounts and ROSCAs.

If the return to hiding savings from a spouse is increasing in discount factor heterogeneity, then hidden savings concerns could be responsible for the main empirical results (though it is not obvious that such a correlation should exist). I added the "extra statements" treatment to the field experiment to assess the overall importance of hidden savings in the study population. To gauge if hidden savings concerns are correlated with preference heterogeneity, I use baseline spousal cross reports of income and savings device use to construct an "information sharing index" which ranges from zero (worst informed couples) to one (best informed couples).

Appendix C presents the hidden savings analysis in detail. While the extra statements intervention is somewhat underpowered, I find evidence that hidden savings concerns are important: just 60 percent of individuals who were presented with the extra statement offer consented.[29] Moreover, the information index appears to identify couples for whom hiding information is particularly valuable. Poorly informed couples were less likely to consent to the extra statements intervention

[28]For example, Anderson and Baland (2002) find that women's use of ROSCAs in Kenya is consistent with a model of hidden information. Boozer, Goldstein, and Suri (2009) analyze spousal cross reports of food expenditure in Ghana and find evidence of hidden consumption. Ashraf (2009) finds evidence that the informational environment has significant impact on the investment decisions of spouses with low levels of financial control in the Philippines and de Laat (2008) finds that individuals in split migrant couples in Kenya are willing to expend considerable resources to acquire information about one another.

[29]Consent rates were nearly identical by gender – 59.2 percent of men and 60.9 percent of women consented to the extra statements.

and also respond more adversely to the extra statements treatment.[30]

If the earlier results were driven by hidden savings concerns, one would expect to see two things: first, the information sharing index should be lower among poorly matched couples. Second, accounting for information sharing in the previous analyses should dilute the impact of preference heterogeneity. However, the information index is totally uncorrelated with preference heterogeneity (the correlation coefficient between the index and the absolute value of preference heterogeneity is 0.017). This foreshadows that accounting for information sharing has no impact on the main results.

First I check to see if accounting for information flows changes the findings regarding patterns of account use. Table 6 presents the regression described by equation 4 augmented to include a dummy variable identifying "poorly informed" couples (those who have a below median information sharing index). The initial results pertaining to preference heterogeneity are very robust to controlling for information flows. At the same time, I find evidence that individual accounts may be valued for informational purposes: poorly informed households are significantly more likely to use individual accounts, even conditional on saving. Additional specification testing shown in Appendix Tables D5 and D6 confirms that information sharing is not driving the results regarding responses to the excess interest rate and interest rate losses.

Overall, these results are compatible with the theory that hidden information concerns impact household savings decisions. However, to the extent that these concerns are important, they appear to be largely orthogonal to preference heterogeneity. This is plausible – hiding savings is likely valuable because it allows individuals to increase their share of consumption, or tilt consumption towards goods that they favor. If the benefit of doing so is equally large for individuals in well and poorly matched households, accounting for it should leave the core results unchanged, which is what I observe. To complete the discussion of robustness, the next section considers other alternative explanations for the main results.

6 Alternative Explanations

In setting up the model, the choice to feature only a public consumption good obviated the possibility that private accounts could be strategically used to change *what* individuals consume (rather, the only margin of influence was *when* consumption took place). In practice, individuals may use private accounts as a tool to change the composition of consumption allocations. One possibility is that spouses strategically save in individual accounts in order to increase bargaining power. Alternatively, individual accounts could change the composition of future consumption if there

[30] I code a couple as "poorly informed" if their information index is below the sample median.

are mental accounting norms in the household. These concerns would be particularly important if heterogeneity in time preferences is correlated with heterogeneity in other preferences.

When saving privately impacts bargaining power or the within-period composition of consumption, both spouses will have incentives to save simultaneously in their individual accounts (this result is established formally in the context of labor supply by Basu 2006 and Browning et al. 2011). In the experiment, just 4 percent of couples saved in both individual accounts. Even among those couples who opened *both* individual accounts *and* saved in at least one account, just 27 percent saved in both individual accounts. This suggests that these concerns are not a major driver of the use of the individual accounts in this study.

Another way to test whether inefficient account use is driven by a desire to manipulate within-period consumption allocations is to ask whether the results are robust to including measures of consumption decision making. I purposely excluded spousal reports of consumption decision making from the control sets in the main analysis because these reports could be endogenous to the degree of strategic savings behavior in the household. As a further robustness check, Appendix Tables D4-D6 present results that include these decision making controls. As it turns out, the results are very robust to including these measures of household functioning in addition to the core time preference, demographic, and economic control sets.

Another possibility is that poorly matched couples choose savings accounts based on rules of thumb, while well matched couples optimally choose accounts taking account of relative rates of return. One model that could generate such behavior is one where household bargaining is costly, and this cost increases as the preferences of household members diverge. If costs are large enough, households could develop rules of thumb for how to manage savings in order to avoid repeated bargaining costs. However, poorly matched couples' lack of response to the excess interest rate is still somewhat of a puzzle in this case – if savings management were tasked to a single individual, he or she should still optimally take account of excess interest rates when deciding between his or her individual account and the joint account.

A final alternative is that poorly matched couples are simply less financially sophisticated than well matched couples. This theory could rationalize both poorly matched couples' overuse of dominated accounts and their lack of sensitivity to the excess interest rate. However, individuals in well and poorly matched couples have very similar levels of education and literacy (Appendix Table D1) and both types of couples respond to interest rate *levels* (Appendix Table D3). Moreover, all the main results are robust to allowing the response to the excess interest rate to vary with spousal education and literacy (as well as within-couple heterogeneity in education and literacy). Thus, there is no evidence that poorly matched couples made less efficient decisions due to cognitive constraints.

7 Conclusion

This paper sheds light on the underlying drivers of inefficient intertemporal resource allocation by households. I structured the analysis by first specifying a model in which heterogeneity in rates of time preference creates incentives for individuals to save strategically, even when doing so is costly. I then derived three testable implications of the model: (1) perfectly matched couples will always save efficiently, whereas poorly matched couples will not, (2) as long as $R_J = R_{max}$ individual account use will increase in preference heterogeneity, and (3) interest rate losses on experimental bank accounts will increase in preference heterogeneity.

The empirical results are consistent with all of these predictions. This is, of course, subject to the caveat that I cannot completely rule out the hypothesis that the results are driven by some other omitted characteristic that is correlated with my estimates of preference heterogeneity. However, the stability of the results to the inclusion of flexible demographic and economic controls, and to the inclusion of measures of household information flows and decision making, is comforting and suggests that the results are indeed driven by inefficiencies arising from conflicting savings motives. I also find evidence that informational concerns impact account use and account choice in the sample. Couples who are poorly informed about one another's financial activities respond most adversely to the "extra statements" treatment and are also more likely to gravitate towards individual accounts over joint accounts.

An innovative feature of the experimental design is that it allows me to quantify investment efficiency in terms of interest rates. However, the experimental interest rates were temporary – it is therefore important to ask whether match quality has broader implications for households' investment choices. I do observe that well matched couples are significantly more likely to invest in livestock and the family farm, which are inherently joint methods of saving that likely bear a higher rate of return than more private savings devices like ROSCAs. While this finding is suggestive, the baseline data lack detailed information on the costs of and returns to different savings devices, so it is difficult to precisely assess how this translates into actual interest rate losses for couples. Incorporating measures of preference heterogeneity into future data collection efforts designed to measure returns to savings and investment would help shed further light on this issue.

My results add to a growing body of literature that rejects dynamic household efficiency, while presenting evidence that heterogeneity in discount factors drives inefficient behavior. A novel feature of this idea is that it provides a mechanism for why some households function well while others do not: when preferences are well aligned there are no incentives to behave strategically and therefore no barriers to attaining an efficient outcome. Although this paper studies strategic savings behavior, the applications are more general. For example, many households in developing countries either engage in home production (such as farming or animal husbandry) or run small

businesses. Investment in these activities is an important way of transferring resources across periods. The insights in this paper suggest that when preferences in the household differ, capital for these activities will not always be allocated to the most efficient user. This mechanism may therefore help account for some of the heterogeneity in plot yields (Udry 1996) or microenterprise returns (de Mel et al. 2009) observed in the developing world. A broader implication of this mechanism is that greater marriage market frictions could lead to lower quality matches in terms of preferences, which could give rise to geographical and cultural variation in household efficiency.

References

Anderson, S. and J.-M. Baland (2002). The Economics of Roscas and Intrahousehold Resource Allocation. *Quarterly Journal of Economics 117*(3), 963–995.

Andreoni, J. and C. Sprenger (2010). Estimating Time Preferences from Convex Budgets. Working Paper 16347, National Bureau of Economic Research.

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 (2006). Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines. *Quarterly Journal of Economics 121*(2), 635–672.

Baland, J.-M., C. Guirkinger, and C. Mali (2007). Pretending to be Poor: Borrowing to Escape Forced Solidarity in Credit Cooperatives in Cameroon. Mimeo.

Basu, K. (2006). Gender and Say: a Model of Household Behaviour with Endogenously Determined Balance of Power. *The Economic Journal 116*(511), 558–580.

Bauer, M. and J. Chytilova (2009). Women, Children and Patience: Experimental Evidence from Indian Villages. IZA Discussion Paper 4241, Institute for the Study of Labor.

Bernheim, B. D. (1999). Comment on 'Family Bargaining and Retirement Behavior'. In H. Aaron (Ed.), *Behavioral Economics and Retirement Policy*, pp. 273–281. Brookings Institution Press.

Bobonis, G. J. (2009). Is the Allocation of Resources within the Household Efficient? New Evidence from a Randomized Experiment. *Journal of Political Economy 117*(3), 453–503.

Boozer, M. A., M. Goldstein, and T. Suri (2009). Household Information: Implications for Poverty Measurement and Dynamics. Mimeo.

Browning, M. (2000). The Saving Behaviour of a Two-person Household. *Scandinavian Journal of Economics 102*(2), 235–251.

Browning, M. and P. A. Chiappori (1998). Efficient Intra-Household Allocations: A General Characterization and Empirical Tests. *Econometrica 66*(6), 1241–1278.

Browning, M., P.-A. Chiappori, and Y. Weiss (2011, Forthcoming). *Family Economics*, Chapter 6: Uncertainty and Dynamics in the Collective Model, pp. 243–288. Cambridge University Press.

Caplin, A. and J. Leahy (2004). The Social Discount Rate. *The Journal of Political Economy 112*(6), 1257–1268.

Chiappori, P.-A. and O. Donni (2009). Non-unitary Models of Household Behavior: A Survey of the Literature.

Chiappori, P.-A., B. Fortin, and G. Lacroix (2002). Marriage Market, Divorce Legislation, and Household Labor Supply. *The Journal of Political Economy 110*(1), 37–72.

Collins, D., J. Morduch, S. Rutherford, and O. Ruthven (2009). *Portfolios of the Poor: How the World's Poor Live on Two Dollars a Day*. Princeton University Press.

de Laat, J. (2008). Household Allocations and Endogenous Information. Mimeo.

de Mel, S., D. McKenzie, and C. Woodruff (2009). Are Women More Credit Constrained? Experimental Evidence on Gender and Microenterprise Returns. *American Economic Journal: Applied Economics 1*(3), 1–32.

Duflo, E. (2003). Grandmothers and Granddaughters: Old-Age Pensions and Intrahousehold Allocation in South Africa. *World Bank Economic Review 17*(1), 1–25.

Duflo, E. and C. Udry (2004). Intrahousehold Resource Allocation in Cote d'Ivoire: Social Norms, Separate Accounts and Consumption Choices. NBER Working Papers 10498.

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.

Fafchamps, M., D. McKenzie, S. R. Quinn, and C. Woodruff (2012). Female Microenterprises and the Fly-paper Effect: Evidence from a Randomized Experiment in Ghana.

Feldstein, M. S. (1964). The Social Time Preference Discount Rate in Cost Benefit Analysis. *The Economic Journal 74*(294), 360–379.

Giné, X., J. Goldberg, D. Silverman, and D. Yang (2011). Revising Commitments: Time Preference and Time Inconsistency in the Field. Mimeo.

Gollier, C. and R. Zeckhauser (2005). Aggregation of Heterogeneous Time Preferences. *Journal of Political Economy 113*(4), 878–896.

Harrison, G. W., M. I. Lau, E. E. Rutstrom, and M. B. Sullivan (2004). *Research in Experimental Economics*, Volume 10, Chapter 5, Eliciting Risk and Time Preferences Using Field Experiments: Some Methodological Issues, pp. 125–218. Emerald Group Publishing Limited.

Jackson, M. O. and L. Yariv (2010). Collective Dynamic Choice: The Necessity of Time Inconsistency. Mimeo.

Jakiela, P. and O. Ozier (2012). Does Africa Need a Rotten Kin Theorem? Experimental Evidence from Village Economics. Mimeo.

Karlan, D. and J. Morduch (2010). *The Handbook of Development Economics*, Chapter 5, Access to Finance. North-Holland.

Lundberg, S. J., R. A. Pollak, and T. J. Wales (1997). Do Husbands and Wives Pool Their Resources? Evidence from the United Kingdom Child Benefit. *The Journal of Human Resources 32*(3), 463–480.

Marglin, S. A. (1963). The Social Rate of Discount and The Optimal Rate of Investment. *The Quarterly Journal of Economics 77*(1), 95–111.

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.

Meier, S. and C. Sprenger (2010). Present-Biased Preferences and Credit Card Borrowing. *American Economic Journal: Applied Economics 2*(1), 193–210.

Newey, W. K. and D. McFadden (1994). Large Sample Estimation and Hypothesis Testing. Volume 4 of *Handbook of Econometrics*, Chapter 36, pp. 2111 – 2245. Elsevier.

Rangel, M. A. and D. Thomas (2005). Out of West Africa: Evidence on the Efficient Allocation of Resources within Farm Households.

Robinson, J. (2011). Limited Insurance Within the Household: Evidence from a Field Experiment in Kenya. MPRA Paper 8314.

Schaner, S. (2013). The Cost of Convenience? Transaction Costs, Bargaining Power, and Savings Account Use in Kenya. (Mimeo).

Shapiro, J. (2010). Discounting for You, Me, and We: Time Preference in Groups and Pairs. Mimeo.

Tanaka, T., C. F. Camerer, and Q. Nguyen (2010). Risk and Time Preferences: Linking Experimental and Household Survey Data from Vietnam. *American Economic Review 100*(1), 557–571.

Udry, C. (1996). Gender, Agricultural Production, and the Theory of the Household. *The Journal of Political Economy 104*(5), 1010–1046.

Zuber, S. (2010). The Aggregation of Preferences: Can We Ignore the Past? *Theory and Decision*, 1–18.

Table 1. Demographic Characteristics and Balance Check

	Summary Statistics		Balance Check					
			Excess Interest Rate			Extra	Cash	
	Husbands	Wives	Husband	Wife	Joint	Statements	Payment	N
Age	41.8	35.2	0.172	0.073	0.295	-1.04	1.16	1196
	[13.3]	[11.8]	(0.188)	(0.194)	(0.201)	(1.08)	(0.988)	
Education	8.07	5.92	-0.035	-0.080	-0.126**	0.073	-0.213	1191
	[3.57]	[4.03]	(0.049)	(0.053)	(0.056)	(0.298)	(0.295)	
Literate	0.865	0.669	-0.007	-0.007	-0.014***	0.004	-0.017	1196
	[0.342]	[0.471]	(0.005)	(0.005)	(0.005)	(0.028)	(0.033)	
Number Children	4.85	4.49	0.047	0.032	0.077*	0.069	0.253	1196
	[2.92]	[2.54]	(0.037)	(0.040)	(0.043)	(0.225)	(0.219)	
Subsistence Farmer/No Job	0.407	0.466	0.008	0.010*	0.016***	-0.025	-0.029	1191
	[0.492]	[0.499]	(0.006)	(0.006)	(0.007)	(0.035)	(0.038)	
Income Last Week	1346	797	-18.7	-28.4	-28.8	-20.0	-357***	1164
	[2747]	[1617]	(28.5)	(25.0)	(23.4)	(152)	(116)	
Owns Mobile Phone	0.487	0.425	0.000	0.000	0.001	0.072**	0.034	1191
	[0.500]	[0.495]	(0.006)	(0.006)	(0.007)	(0.036)	(0.037)	
Participates in ROSCA	0.492	0.657	-0.002	0.004	0.005	0.010	0.004	1196
	[0.500]	[0.475]	(0.006)	(0.006)	(0.006)	(0.033)	(0.037)	
Has Bank Account	0.303	0.119	-0.002	-0.008	-0.008	-0.023	0.023	1196
	[0.460]	[0.324]	(0.005)	(0.005)	(0.005)	(0.028)	(0.032)	
Has a SACCO Account	0.062	0.010	-0.001	-0.003	-0.002	-0.010	-0.003	1193
	[0.241]	[0.100]	(0.002)	(0.002)	(0.002)	(0.012)	(0.014)	
Saves at Home	0.851	0.898	0.001	0.005	-0.001	0.005	0.047**	1195
	[0.356]	[0.303]	(0.003)	(0.004)	(0.004)	(0.021)	(0.022)	
Has Mobile Money Account	0.309	0.132	-0.007	-0.006	-0.014***	0.004	-0.020	998
	[0.462]	[0.339]	(0.005)	(0.005)	(0.006)	(0.028)	(0.033)	
Saves Other Ways	0.582	0.536	0.010	0.007	0.001	0.070**	-0.038	998
	[0.494]	[0.499]	(0.006)	(0.007)	(0.007)	(0.032)	(0.042)	
Total Reported Savings	19874	5266	409	-54.9	-304	963	5690	923
	[56668]	[13119]	(428)	(491)	(401)	(3067)	(6970)	
Consumption - Husband Decides	0.471	0.377	-0.002	-0.010*	-0.004	-0.024	0.002	1189
	[0.500]	[0.485]	(0.005)	(0.006)	(0.006)	(0.032)	(0.038)	
Consumption - Wife Decides	0.076	0.170	0.001	0.003	-0.001	-0.009	0.010	1189
	[0.265]	[0.376]	(0.004)	(0.004)	(0.004)	(0.021)	(0.025)	
Consumption - Both Decide	0.403	0.411	0.000	0.005	0.004	0.012	-0.012	1189
	[0.491]	[0.492]	(0.006)	(0.006)	(0.006)	(0.033)	(0.037)	
Savings - Husband Decides	0.354	0.282	0.001	0.001	0.000	-0.029	-0.048	1189
	[0.479]	[0.450]	(0.005)	(0.005)	(0.006)	(0.030)	(0.035)	
Savings - Wife Decides	0.440	0.491	-0.001	-0.007	-0.005	0.012	0.046	1189
	[0.497]	[0.500]	(0.006)	(0.006)	(0.006)	(0.034)	(0.039)	
Savings - Both Decide	0.183	0.204	0.001	0.007	0.006	0.014	-0.004	1189
	[0.387]	[0.403]	(0.005)	(0.005)	(0.005)	(0.028)	(0.031)	
Distance from Bank (Miles)	3.78	3.78	0.035	0.023	0.054	0.006	0.041	1196
	[2.21]	[2.21]	(0.034)	(0.035)	(0.039)	(0.211)	(0.152)	
Weekly Discount Factor	0.718	0.704	-0.003	-0.004	-0.006	0.059***	-0.072***	1196
	[0.368]	[0.363]	(0.004)	(0.004)	(0.004)	(0.023)	(0.031)	
Well Matched Couple	0.500	0.500	-0.003	-0.009	-0.007	-0.012	0.055	1196
	[0.500]	[0.500]	(0.007)	(0.008)	(0.008)	(0.045)	(0.037)	
Joint Test (Cross Equation)			0.958	0.510	0.588	0.253	{0.011}**	

Notes: Standard deviations in brackets, standard errors in parentheses, p-values in braces. Columns 3-7 present results from regressions of observable characteristics on listed treatments of interest. All standard errors clustered at the couple level. Mobile money and other savings data not available for the 98 couples in the first 6 experimental sessions. Variables recoded to missing if response was don't know/refused. P-values from the joint test are calculated by jointly estimating equations by seemingly unrelated regression. ***, **, and * indicate significance at the 99, 95, and 90 percent confidence levels respectively.

Table 2. Summary of Bank Account Use

	Well Matched	Poorly Matched
Couple Chose to Open:		
All Three Accounts	0.050	0.057
Joint Account Only	0.562	0.562
Both Individual Accounts	0.274	0.298
One Individual, One Joint Account	0.090	0.054
One Individual Account	0.023	0.030
Couple Saved In:		
Any Account	0.421	0.438
Joint Account	0.308	0.274
Individual Account	0.124	0.184
Excluding Cash Payments, Couple Saved In:		
Any Account	0.268	0.278
Joint Account	0.174	0.171
Individual Account	0.104	0.127
If Saved, Number Deposits In:		
All Accounts	2.53	2.92
	[2.82]	[3.74]
Joint Account	2.19	2.55
	[2.46]	[3.84]
Individual Accounts	3.12	3.19
	[3.24]	[3.04]
If Saved, Average Daily Balance In:		
All Accounts	991	998
	[2069]	[1794]
Joint Account	719	897
	[1710]	[1695]
Individual Accounts	1587	1040
	[2591]	[1664]

Notes: Standard deviations in brackets.

Table 3. Responses to the Excess Interest Rate by Match Quality

Panel A. Accounts Owned by Well Matched Couples Only				
Excess Low×Indiv	0.000	0.000	0.001	0.000
	(0.005)	(0.005)	(0.006)	(0.006)
Excess High×Indiv	0.016**	0.017**	0.014*	0.014*
	(0.008)	(0.008)	(0.008)	(0.008)
Excess Low×Joint	0.021*	0.022*	0.023**	0.023*
	(0.012)	(0.012)	(0.012)	(0.012)
Excess High×Joint	0.006	0.006	0.005	0.005
	(0.011)	(0.012)	(0.012)	(0.012)
DV Mean	0.164	0.164	0.164	0.164
N	897	897	897	897
Panel B. Accounts Owned by Poorly Matched Couples Only				
Excess Low×Indiv	-0.004	-0.003	-0.008	-0.009
	(0.007)	(0.007)	(0.008)	(0.008)
Excess High×Indiv	-0.003	-0.003	-0.002	-0.002
	(0.007)	(0.007)	(0.007)	(0.007)
Excess Low×Joint	-0.006	-0.005	-0.008	-0.008
	(0.010)	(0.010)	(0.011)	(0.011)
Excess High×Joint	-0.008	-0.008	-0.006	-0.005
	(0.011)	(0.011)	(0.011)	(0.012)
DV Mean	0.175	0.175	0.175	0.175
N	897	897	897	897
Panel C. Tests of Equality Across Equations				
Well Matched Couples				
Excess Positive	0.018***	0.018***	0.021***	0.022***
	(0.007)	(0.007)	(0.007)	(0.007)
Excess Zero	0.004	0.004	0.006	0.006
	(0.005)	(0.006)	(0.005)	(0.006)
Poorly Matched Couples				
Excess Positive	-0.004	-0.022***	-0.020**	-0.020*
	(0.006)	(0.008)	(0.010)	(0.010)
Excess Zero	-0.004	-0.013	-0.013	-0.014
	(0.006)	(0.009)	(0.008)	(0.009)
Tests of Equality of Coefficients Across Equations				
Excess Positive	{0.024}**	{0.000}***	{0.001}***	{0.001}***
Excess Zero	{0.290}	{0.101}	{0.061}*	{0.071}*
Joint Test - Both Splines	{0.014}**	{0.000}***	{0.000}***	{0.000}***
Control Set	Basic	+Time Pref	+Demo.	+Economic

Notes: Robust standard errors clustered at the couple level in parentheses, p-values in braces. All regressions include fixed effects that saturate interest rate×joint account and dummies for husband and wife cash payment selection. Time preference controls include separate dummies for upper/lower censoring of the discount factors of each spouse and the estimated discount factor of each spouse. The demographic control set adds controls for experimental session fixed effects, spousal heterogeneity in age, education, number of children, and literacy. The economic control set adds controls for heterogeneity in income, an indicator for subsistence farmers or the unemployed, and phone ownership. All controls are demeaned to the value among well matched couples. Specifications in Panel C also include interactions between the excess interest splines and the included control variables. Cross-equation F-tests calculated using seemingly unrelated regression. ***, **, and * indicate significance at the 99, 95, and 90 percent confidence levels respectively.

Table 4. Preference Heterogeneity and Use of Dominated Individual Accounts

Panel A. All Couples With Dominated Individual Accounts				
Poorly Matched	0.125***	0.177***	0.151***	0.175***
	(0.035)	(0.038)	(0.042)	(0.045)
DV Mean (Well Matched)	0.078	0.078	0.078	0.078
N	363	363	363	363
Panel B. Subset of Couples Who Saved in at Least One Account				
Poorly Matched	0.244***	0.376***	0.224*	0.240*
	(0.078)	(0.085)	(0.117)	(0.143)
DV Mean (Well Matched)	0.182	0.182	0.182	0.182
N	157	157	157	157
Control Set	Basic	+Time Pref	+Demo.	+Economic

Notes: The sample is limited to the subset of couples for whom the joint account bears the highest interest rate. The dependent variable is a dummy indicating that a couple saved in any individual account. Heteroskedasticity robust standard errors in parentheses. All regressions include dummy variables for each account's interest rate and dummies for husband and wife cash payment selection. Time preference controls include separate dummies for upper/lower censoring of the discount factors of each spouse and the estimated discount factor of each spouse. The demographic control set adds controls for experimental session fixed effects, spousal heterogeneity in age, education, number of children, and literacy. The economic control set adds controls for heterogeneity in income, an indicator for subsistence farmers or the unemployed, and phone ownership. ***, **, and * indicate significance at the 99, 95, and 90 percent confidence levels respectively.

Table 5. Interest Rate Losses by Match Quality

	No	Proxied Banking Cost Discounting			Uniform Discounting		
Maximum Individual Discount	Discounting	5	10	15	5	10	15
Poorly Matched Couples							
Maximum Interest Earnings	8.15	6.45	6.01	6.01	6.74	6.16	6.04
Actual Interest Earnings	7.10	5.28	4.17	3.29	5.66	4.72	4.03
Loss	1.05	1.18	1.85	2.72	1.08	1.44	2.01
Well Matched Couples							
Maximum Interest Earnings	8.17	6.41	5.93	5.93	6.73	6.11	5.96
Actual Interest Earnings	7.50	5.85	4.92	4.31	6.24	5.46	5.00
Loss	0.664	0.566	1.02	1.63	0.486	0.655	0.965
Loss Gap							
A. No Controls	0.389**	0.613***	0.830***	1.10***	0.591***	0.784***	1.04***
	(0.172)	(0.205)	(0.316)	(0.447)	(0.185)	(0.253)	(0.338)
B. + Basic Controls	0.447***	0.694***	0.968***	1.29***	0.660***	0.894***	1.19***
	(0.163)	(0.202)	(0.308)	(0.435)	(0.183)	(0.248)	(0.332)
C. + Time Preference Controls	0.486**	0.917***	1.41***	1.92***	0.798***	1.21***	1.66***
	(0.219)	(0.218)	(0.326)	(0.464)	(0.209)	(0.266)	(0.355)
D. + Demographic Controls	0.462**	0.855***	1.36***	1.90***	0.744***	1.14***	1.61***
	(0.218)	(0.223)	(0.342)	(0.490)	(0.211)	(0.274)	(0.369)
E. + Economic Controls	0.462**	0.868***	1.41***	2.00***	0.747***	1.14***	1.63***
	(0.222)	(0.228)	(0.348)	(0.496)	(0.216)	(0.282)	(0.378)
N	598	598	598	598	598	598	598

Notes: Robust standard errors in parentheses. Basic controls include dummy variables for each account's interest rate and dummies for husband and wife cash payment selection. Time preference controls include separate dummies for upper/lower censoring of the discount factors of each spouse and the estimated discount factor of each spouse. The demographic control set adds controls for experimental session fixed effects, spousal heterogeneity in age, education, number of children, and literacy. The economic control set adds controls for heterogeneity in income, an indicator for subsistence farmers or the unemployed, and phone ownership. Proxied banking costs constructed using a principal components index of distance from the bank, spouse specific indicators for subsistence farmers/the unemployed, and spouse specific indicators for baseline bank account ownership and SACCO membership. ***, **, and * indicate significance at the 99, 95, and 90 percent confidence levels respectively.

Table 6. Preference Heterogeneity, Information Sharing, and Use of Dominated Individual Accounts

Panel A. All Couples With Dominated Individual Accounts				
Poorly Matched	0.130***	0.184***	0.161***	0.181***
	(0.035)	(0.037)	(0.042)	(0.045)
Poorly Informed	0.113***	0.122***	0.119***	0.096**
	(0.039)	(0.038)	(0.042)	(0.043)
DV Mean (Well Matched)	0.078	0.078	0.078	0.078
N	363	363	363	363
Panel B. Subset of Couples Who Saved in at Least One Account				
Poorly Matched	0.217***	0.335***	0.169	0.216
	(0.078)	(0.087)	(0.122)	(0.157)
Poorly Informed	0.217***	0.206***	0.232**	0.096
	(0.085)	(0.085)	(0.108)	(0.131)
DV Mean (Well Matched)	0.182	0.182	0.182	0.182
N	157	157	157	157
Control Set	Basic	+Time Pref	+Demo.	+Economic

Notes: The sample is limited to the subset of couples for whom the joint account bears the highest interest rate. The dependent variable is a dummy indicating that a couple saved in any individual account. Heteroskedasticity robust standard errors in parentheses. All regressions include dummy variables for each account's interest rate and dummies for husband and wife cash payment selection. Time preference controls include separate dummies for upper/lower censoring of the discount factors of each spouse and the estimated discount factor of each spouse. The demographic control set adds controls for experimental session fixed effects, spousal heterogeneity in age, education, number of children, and literacy. The economic control set adds controls for heterogeneity in income, an indicator for subsistence farmers or the unemployed, and phone ownership. ***, **, and * indicate significance at the 99, 95, and 90 percent confidence levels respectively.

Figure 1. Distributions of Estimated Discount Factors and Discount Factor Heterogeneity

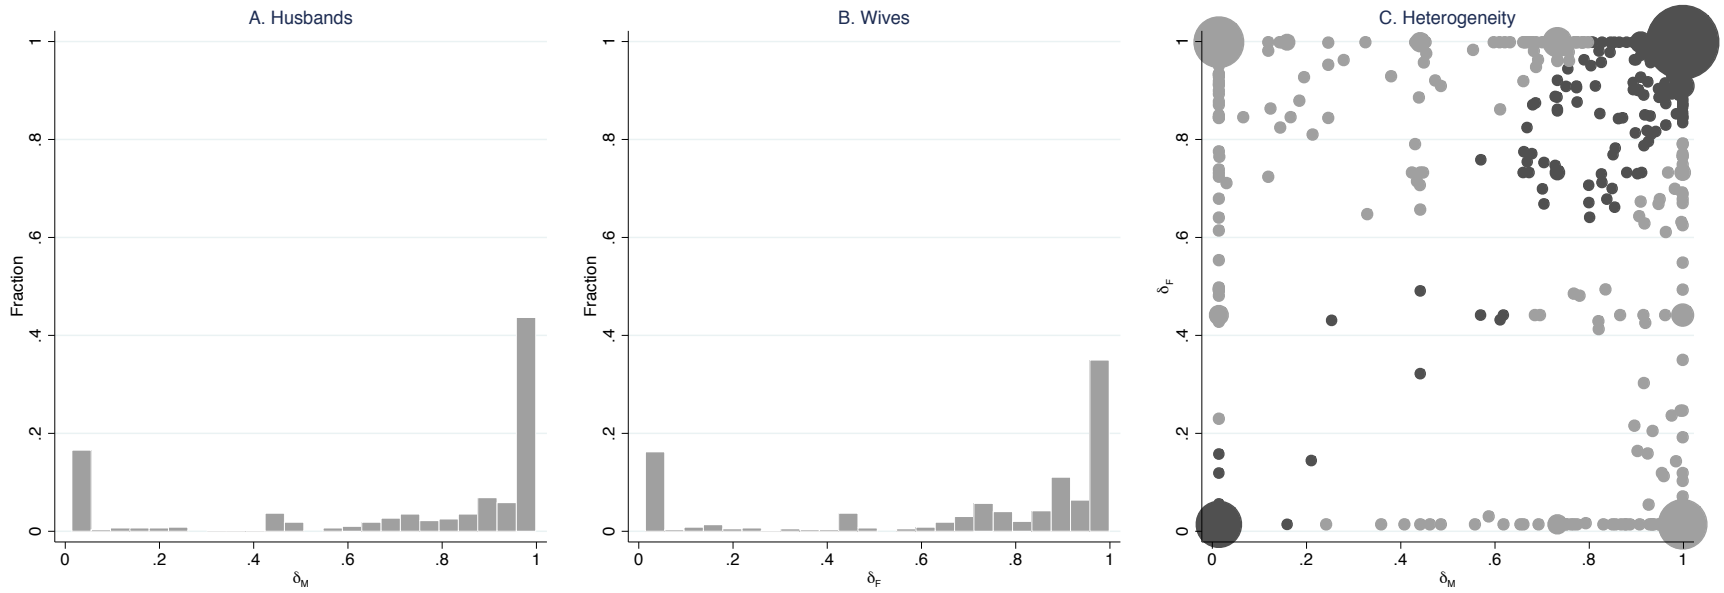

Note: Panels A and B are histograms of husbands' and wives' discount factors respectively. Panel C is a weighted scatterplot of husbands' discount factors (x axis) and wives' discount factors (y axis). The size of each circle in Panel C is proportional to the number of couples with the relevant discount factor combination. Well matched couples in Panel C are demarcated with darker shading.

Figure 2. Efficient Responses to the Excess Interest Rate by Account Type

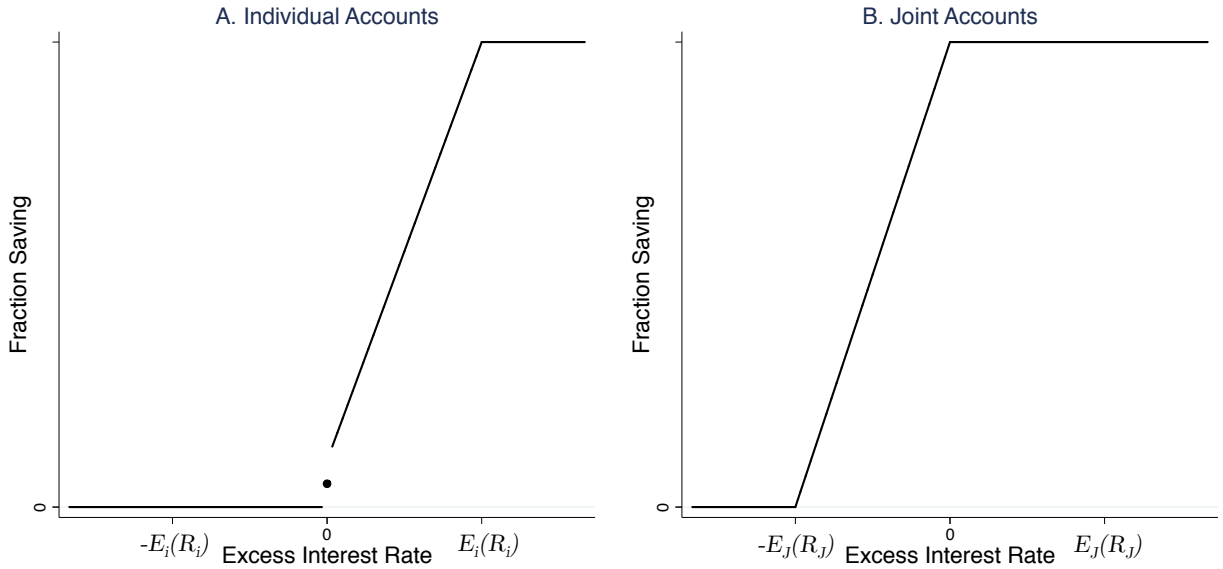

Figure 3. Savings Response to Excess Interest Rate by Match Quality and Account Type

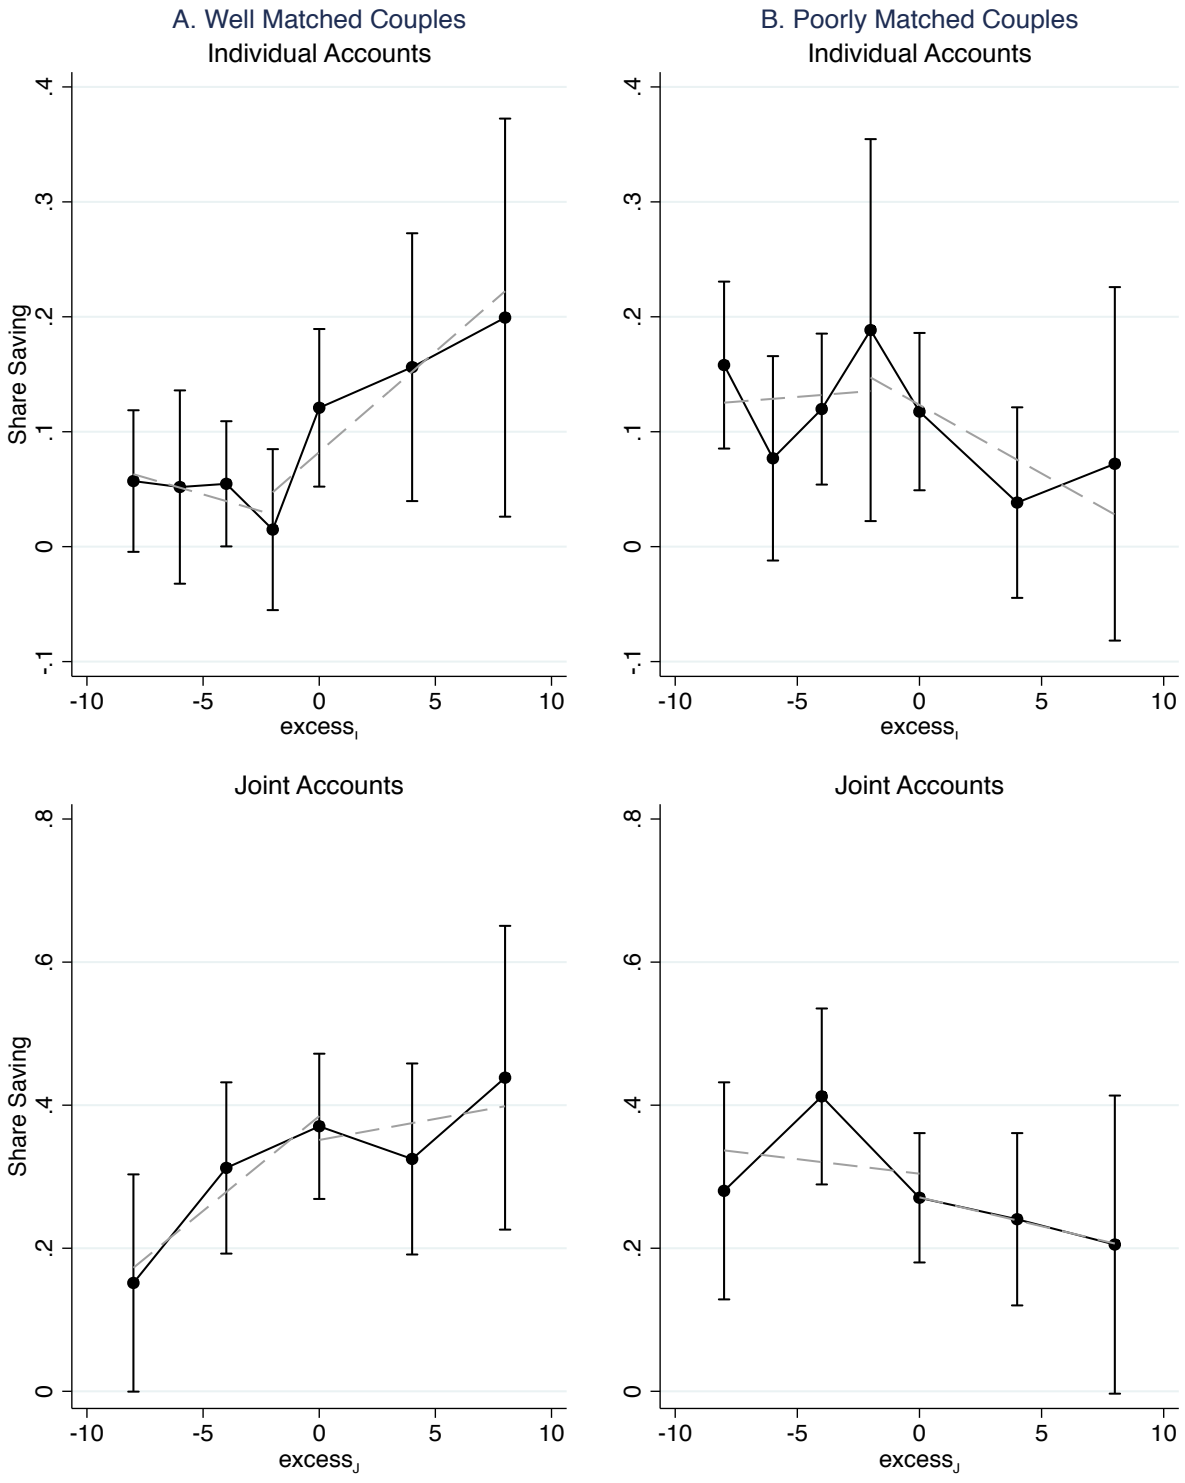

Note: Predicted values calculated from regressions with standard errors clustered at the couple level. Whiskers indicate 95 percent confidence intervals for each predicted value.

Figure 4. Use of Dominated Individual Accounts and Heterogeneity in Estimated Discount Factors

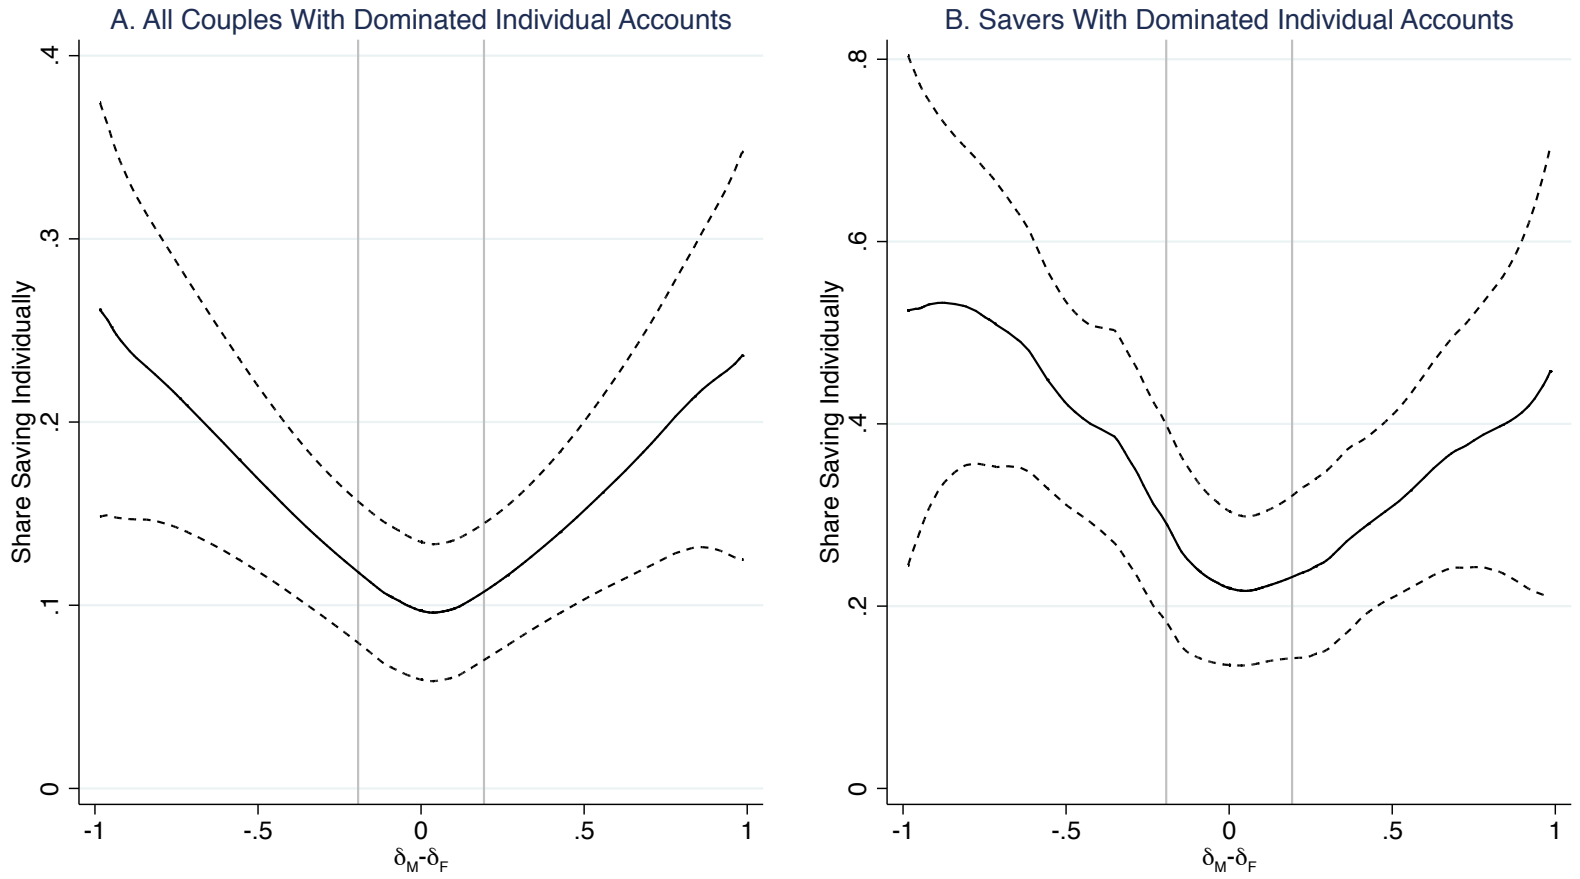

Notes: Local linear regression results. The sample is limited to the subset of couples for whom the joint account bears the highest interest rate. The dependent variable is a dummy indicating that a couple saved in any individual account. Dashed lines give 95 percent confidence intervals. Well matched couples are delineated by gray vertical lines.

Figure 5. Interest Rate Losses by Match Quality and Individual Interest Discounting

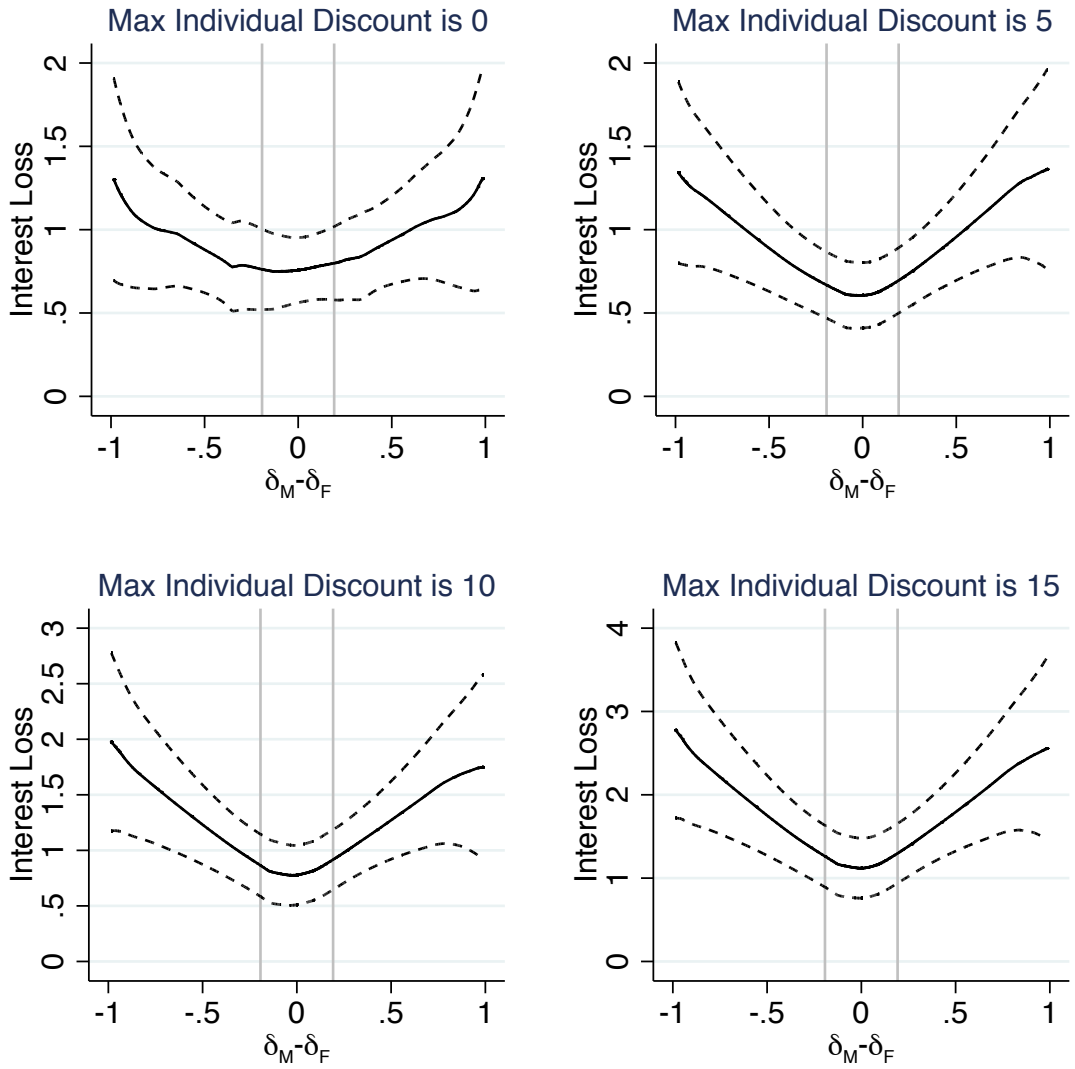

Notes: Local linear regression results. Dashed lines give 95 percent confidence intervals. Gray vertical lines demarcate well matched couples. Individual interest rate discounting performed using proxied banking costs.

A Proofs

Before proving Proposition 1, it is useful to describe the nature of potentially optimal savings strategies more explicitly. The following two lemmas do this.

Lemma 1 *At most one spouse will travel to the bank to transact in a given period.*

Proof. The proof will show that any strategy that results in both spouses going to the bank cannot be optimal. Note that if the second mover saves in account a, she will always save so that $u'(c_1) = R_a\delta_B u'(c_2)$ (and must therefore only be saving at interest rate R_a). Now consider the options of the first mover.

First suppose the second mover only deposits savings in the joint account. Then the first mover could simply place additional savings in the joint account to satisfy $u'\left(y_1 - s_1^A - s_1^J - b\right) = R_J\delta_B u'\left(y_2 + R_A s_1^A + R_J s_1^J - b\right)$. The second mover will prefer not to go back to the bank, and consumption in both periods will be higher, since at least b in banking costs are saved. This is a profitable deviation.

Instead suppose the second mover deposits at least some savings in her individual account (this requires $R_B \geq R_J$). If $R_B \geq R_A$, the first mover could deviate and simply not go to the bank. The second mover would continue to save at interest rate R_B – since total resources increase by at least b, the only way $u'(c_1) = R_a\delta_B u'(c_2)$ will continue to hold is if consumption in both periods increases, which is a profitable deviation. Suppose instead that $R_B < R_A$. In this case the first mover could deviate by placing the value of both his and the first mover's savings, as well as b, in his individual account. The second mover's best response to this strategy would be to stay at home – thus, first period consumption would be unchanged while second period consumption would strictly increase. This is a profitable deviation.

Combined, these observations show that no strategy in which both spouses go to the bank can be optimal. ■

Lemma 2 *Consider the case where $\delta_A \neq \delta_B$. Any pure strategy, subgame perfect Nash equilibrium (PSPNE) to the savings game must involve one of the following strategies:*

i. Neither spouse saves.

ii. The first mover does not save and the second mover saves in his/her individual account with $R_B = R_{max}$ such that $u'(c_1) = R_{max}\delta_B u'(c_2)$. This strategy will only be used when $R_J < R_{max}$ and $R_A < R_{max}$.

iii. The first mover saves in his/her individual account such that $u'(c_1) = R_A\delta_A u'(c_2)$. The second mover strictly prefers to stay at home. This strategy will only be used when $R_A \geq R_J$.

iv. The first mover saves in the joint account such that $u'(c_1) = R_J\delta_A u'(c_2)$. The second mover strictly prefers to stay at home. This strategy will only be used when $R_J \geq R_A$.

v. The first mover saves to the point that the second mover is indifferent between staying home and going back to the bank to reoptimize savings. When $R_A > R_J$ the first mover will only use his/her individual account. When $R_J \geq R_A$ the first mover may use just the individual account, just the joint account, or a mixture of the two.

Proof. The proof will show that any savings strategy other than those listed above cannot be optimal. First consider strategies in which the second mover saves. By Lemma 1, the first mover will not go to the bank. It follows that the second mover will make use of the highest return account available to her. Suppose that both the first and second mover have access to an account bearing R_{max}. In this case the first mover could go to the bank and deposit the second mover's optimal savings plus (if $\delta_A > \delta_B$) or minus (if $\delta_A < \delta_B$) ε into the high-return account. Since $b > 0$, the second mover will prefer to stay at home for sufficiently small ε and the first mover will be strictly better off. This rules out the possibility that $R_J = R_{max}$, or that $R_A = R_B = R_{max}$.

Suppose instead that $R_A = R_{max}$, but $R_J < R_{max}$, $R_B < R_{max}$. Then the first mover could go to the bank and deposit the value of the second mover's optimal savings into the high-return individual account. The second mover will optimally stay home, first period consumption will stay unchanged, and second period consumption will strictly increase. Thus, this is a profitable deviation. By process of elimination, case (ii) is the only remaining scenario in which the second mover saves.

Now consider cases where only the first mover saves. Suppose that the second mover strictly prefers to stay at home. Then it must be that the first mover is saving in the account with the highest possible return (call this return $R^* = \max\{R_A, R_J\}$) such that $u'(c_1^*) = R^* \delta_A u'(c_2^*)$ – otherwise the first mover could reallocate ε of his savings in such a way to become strictly better off. The only scenarios satisfying these requirements are (iii) and (iv).

Finally, suppose the first mover saves to the point that the second mover is indifferent between going to the bank to reoptimize savings and staying at home. This strategy must solve the following optimization problem:

$$\max_{s_1^A, s_1^J} \quad u\left(y_1 - s_1^A - s_1^J - b\right) + \delta_A u\left(y_2 + R_A s_1^A + R_J s_1^J - b\right) \quad \text{subject to} \quad (1)$$

$$u\left(y_1 - s_1^A - s_1^J - b\right) + \delta_B u\left(y_2 + R_A s_1^A + R_J s_1^J - b\right) \geq V_2\left(s_1^A\right)$$

$$s_1^A \geq 0,\ s_1^J \geq 0$$

where $V_2\left(s_1^A\right)$ is the utility that agent two would get from paying b and reoptimizing savings at the bank. $V_2(\cdot)$ depends on s_1^A, since savings in the first mover's individual account cannot be withdrawn by the second mover. In contrast, $V_2(\cdot)$ does not depend on s_1^J, since the second mover can freely access these funds once she is at the bank.

When $R_A > R_J$, the first mover will only use the individual account for an optimal indifference strategy. To see this, suppose instead that $s_1^J > 0$. The first mover could simply reallocate joint savings to his individual account – first period consumption would stay unchanged, second period consumption would increase, and the second mover would strictly prefer to stay at home. Thus, the only remaining strategies involving indifference of the second mover are those listed in (v). ■

We are now prepared to move on to Proposition 1.

Proof of Proposition 1. Suppose $\delta_A = \delta_B$. There will be a PSPNE that implements the (c_1^*, c_2^*) that maximizes both agents' utility (neither spouse would have an incentive to deviate from the strategy that generates this allocation). Since we have assumed that couples never play Pareto-dominated equilibria, it must be the case that $L = 0$ when a couple is perfectly matched.

To complete the proof, it is sufficient to show that L will never strictly decrease as preference heterogeneity increases. The proof proceeds by contradiction: suppose that when $\Delta' = \left|\delta_A' - \delta_B\right|$ increases to $\Delta'' = \left|\delta_A'' - \delta_B\right|$ the interest rate loss decreases from $L' > 0$ to L''. Denote the equilibrium consumption allocation at Δ' as (c_1', c_2') and the consumption allocation at Δ'' as (c_1'', c_2'').

Since the preferences of the second mover are constant, optimality of the first mover's strategy requires that

$$\delta_A' \left(u\left(c_2''\right) - u\left(c_2'\right)\right) \leq u\left(c_1'\right) - u\left(c_1''\right) \leq \delta_A'' \left(u\left(c_2''\right) - u\left(c_2'\right)\right)$$

with at least one of the inequalities being strict. (This observation uses the fact that monotonic increases in heterogeneity require either $\delta_A^{''} < \delta_A^{'} < \delta_B$ or $\delta_A^{''} > \delta_A^{'} > \delta_B$). This implies that the second mover's preference ordering is $(c_1^*, c_2^*) \underset{B}{\succ} (c_1', c_2') \underset{B}{\succ} (c_1'', c_2'')$. The remainder of the argument will show that any savings strategy that generates (c_1', c_2') cannot be part of a PSPNE.

Since $L' > 0$, (c_1', c_2') must be generated by the use of at least one account with $R_a < R_{max}$. Lemma 2 rules out the possibility that the dominated account is that of the second mover, as the second mover's account will only be used when $R_B = R_{max}$.

A second possibility is that (c_1', c_2') is generated by saving only in the dominated joint account. Lemma 2 shows that the first mover will only save in the joint account when $R_A \leq R_J$. In this case it must be that $R_{max} = R_B$ and (c_1'', c_2'') must involve either no saving at all or the second mover optimally saving in her individual account. But this would imply that $(c_1^*, c_2^*) \underset{B}{\sim} (c_1'', c_2'')$, which is a contradiction.

The final possibility is that (c_1', c_2') involves the use of the first mover's dominated individual account. It must be the case that $R_J > R_A$ and (c_1'', c_2'') involves use of the joint account – otherwise (c_1'', c_2'') would have to involve either no saving or the second mover saving in her individual account, which we have seen is impossible. Lemma 2 also shows that the second mover must be indifferent between staying at home and going to the bank at $\left(c_1^{'}, c_2^{'}\right)$. There are two possibilities to consider:

i. The second mover is indifferent at (c_1', c_2') but not (c_1'', c_2'').

ii. The second mover is indifferent at both (c_1', c_2') and (c_1'', c_2'').

First consider subcase (i). If the second mover is not indifferent at (c_1'', c_2''), then by Lemma 2 (c_1'', c_2'') must involve just joint savings. This implies that

$$V_2\left(s_1^{A'}\right) = u\left(c_1'\right) + \delta_B u\left(c_2'\right) > u\left(c_1''\right) + \delta_B u\left(c_2''\right) > V_2\left(0\right)$$

where $V_2\left(s_1^{A'}\right)$ is the utility agent 2 would receive from going back to the bank and reoptimizing savings given $s_1^{A'}$. However, if $R_A < R_J$ it must be that $V_2\left(0\right) > V_2\left(s_1^{A}\right)$ – thus, this subcase is not possible.

Finally consider subcase (ii). This requires that

$$V_2\left(s_1^{A'}\right) = u\left(c_1'\right) + \delta_B u\left(c_2'\right) > u\left(c_1''\right) + \delta_B u\left(c_2''\right) = V_2\left(s_1^{A''}\right)$$

These inequalities will only be satisfied if $s_1^{A'} < s_1^{A''}$. If $L'' < L'$ this must mean that both joint and individual savings are higher at (c_1'', c_2'') (so $\delta_A'' > \delta_A' > \delta_B$). However, for $\left(s_1^{A'}, s_1^{J'}\right)$ to be a solution to (1), it must be that agent B prefers to deviate for any allocation involving $s_1^A > s_1^{A'}$ and $s_1^J > s_1^{J'}$. This however, would imply that (c_1'', c_2'') is not feasible.

Thus, there is no possible scenario in which $L'' < L'$. This completes the proof. ■

B Proxying Banking Costs

I conjecture that those couples who travel to Busia town very frequently for non-bank related reasons and those couples who have low travel costs to town will have smaller differential banking costs (i.e. the joint account offers less in transaction cost savings). This should be negatively correlated with distance from the bank and economic activity (here I assume that subsistence farmers and the unemployed are less likely to take frequent trips to town). Moreover, pre-existing use of a formal savings account should signal lower differential banking costs. Here I include both bank accounts and SACCO accounts as formal accounts.[31] To aggregate these measures, I used principal components analysis to extract the first principal component of the data matrix formed by the above-listed variables (results are similar if I convert the measures into standard deviation units and equally weight them). I then normalized this component to construct a "banking costs index", which runs from zero (lowest hypothesized banking costs) to one (highest hypothesized banking costs). Appendix Table B1 checks to see if the cost index is correlated with account behavior in ways predicted by the theory. The first three columns present couple-level regressions of the following form:

$$y_c = \beta_0 + \beta_1 index_c + x_c'\delta + \varepsilon_c \tag{2}$$

where y_c is the outcome of interest, $index_c$ is the banking cost index, and x_c is a vector of controls. I use the same control sets as in previous regressions, but I exclude controls for occupation (since this is included in the index) as well as experimental session fixed effects (which absorb distance from the bank, which is also included in the index).

It is likely that the cost index is correlated with absolute banking costs and socioeconomic status in addition to differential banking costs. To test the former hypothesis, the first column of Table B1 examines the correlation between the banking cost index and a dummy variable for whether or not a couple saved in any bank account. Indeed, couples with the highest proxied costs are significantly less likely to save in any bank account, which suggests that they have fewer resources and/or face

[31] SACCO stands for "savings and credit cooperative". SACCOs function like credit unions, and are generally organized around higher paying professions, such as teaching and commercial farming. There is an argument for excluding the account ownership variables from the index, since this may be determined in part by match quality. In practice, the results are very similar if I exclude the account ownership variables from the banking cost index.

greater costs of using the accounts. The next three columns limit the population to couples who saved in at least one account, and examine the correlation between use of joint versus individual accounts and the cost index.[32] Higher cost savers are significantly less likely to save individually and more likely to save jointly as compared to lower cost savers. These results are relatively robust to adding observable controls that capture socioeconomic status and other characteristics – though the results examining account use are attenuated upon including economic controls, the coefficients are still large in magnitude and mostly significant.

Earlier, I argued that well matched couples' response to (and poorly matched couples' lack of response to) the excess interest rate suggest that interest rate losses only impacted the decisions of well matched savers with relatively low differential banking costs. I use the proxied banking costs to test this hypothesis directly: in this case well matched couples with low proxied banking costs should respond robustly to the excess interest rate, while well matched couples with large proxied banking costs should be much less responsive to the excess rate. To test this, I limit the sample to well matched couples, define a couple to have high banking costs if their index value is above the sample median, and run the following regression:

$$\begin{aligned} saved_{ac} &= \beta_0 + \beta_1 high_c + \beta_2 excess_{ac} + \beta_3 high_c \times excess_{ac} + \beta_4 joint_{ac} + \\ &\quad \beta_5 \left(joint \times high\right)_{ac} + z'_{ac}\lambda + x'_c\delta + \varepsilon_{ac} \end{aligned} \tag{3}$$

where $high_c$ is a dummy variable indicating high banking costs, $excess_{ac}$ is the excess interest rate on account a for couple c, $joint_{ac}$ is a joint account indicator, z_{ac} is a vector of interest rate dummies, their interactions with $joint_{ac}$, and further interactions with the high cost dummy. Finally, x_c is a vector of the same control sets used in the earlier specifications in Appendix Table B1.

The fourth column of Table B1 presents results of this regression. As expected, the response to the excess interest rate is positive and highly significant for couples with lower proxied banking costs and very close to zero for couples with higher proxied banking costs. Overall, these results suggest that the index is a reasonable proxy for differential banking costs, subject to the caveat that it is correlated with other determinants of account use, such as socioeconomic status and absolute banking costs.

[32] In the absence of a selection effect (i.e. conditional on b^J/savings capacity), larger differential banking costs will push couples to make more intensive use of joint accounts. However, larger absolute banking costs/lower savings capacities will select out smaller scale savers. Since the hypothesized costs are fixed and not proportional to balances, these savers will be more likely to opt for joint accounts, all else equal. Therefore this selection effect would bias me away from finding a positive correlation between joint account use and the cost index.

C Extra Statements Sample and Information Index

As described in Section 3, 50 percent of couples (who attended the seventh experimental session or above) were sampled for an "extra statements" offer. In order to keep selection into individual account opening constant between treatment and control, the experimental protocol dictated that the extra statement offer only be made to participants *after* they decided which accounts to open. However, extra statement provision is significantly, negatively correlated with the probability of opening an individual account in the overall sample. The enumerators never reported informing couples about the extra statements before account choice, nor did they report any cases where couples changed their minds about opening individual accounts after getting an extra statements offer. However, in a subset of the sessions the enumerators were able to observe couples' extra statements treatment status before couples made account opening choices. The correlation is only significant in this subset of sessions, so I conjecture that some enumerators guided selected couples to joint accounts, as filling out the extra statement cards involved time consuming paperwork. To address this concern, I ran the following regression among all individual accounts in the suspect sessions enumerator by enumerator:

$$open_{ic} = \beta_0 + \beta_1 es_c + \gamma_e + \varepsilon_{ic}$$

Where i indexes the individual, es_c indicates extra statement selection and γ_e are co-enumerator fixed effects (the enumerators worked in teams of two).

Out of 14 enumerators, the coefficient on es_c was negative and significant for just four enumerators. I dropped observations for these four enumerators in the sessions where enumerators could observe extra statement selection prior to the couple's account opening decision. All told, I dropped 366 of 1,000 individual account observations. Columns 1 and 2 of Appendix Table C1 verify that account opening is uncorrelated with extra statement selection once I drop the suspect enumerators.

Before dropping suspect observations, I created an information sharing index. The index inputs are responses to baseline survey questions addressing income earned last week, bank accounts, savings at home, SACCOs, and ROSCAs, which I use to create five subindices. The subindices range from 0 (perfect information) to 1 (most misinformation). If cross reports exactly matched own reports, I coded the subindex to 0. If an individual reported that they did (did not) use a device, but the spouse reported that they did not (did), I coded the subindex to 1. If a spouse asserted that they did not know if an individual used a given device, or if they did not know how much savings was in the device, I also coded the subindex to 1. For other instances where I had an own report and a cross report of the amount (or in the case of ROSCAs, the number of ROSCAs), I coded the

subindex to equal $\min\left\{\frac{|own_{ic}-cross_{ic}|}{own_{ic}}, 1\right\}$. I then created a household-level information index equal to:

$$index_c^{info} = 1 - \frac{1}{2}\sum_{i\in\{M,F\}} \frac{index_{ic}^{inc} + index_{ic}^{hh} + index_{ic}^{bank} + index_{ic}^{sacco} + index_{ic}^{rosca}}{5}$$

where a value of 1 represents a perfectly informed household and a value of 0 represents a poorly informed household.

Appendix Table C1 presents the results of the extra statements intervention. All regressions are of the following form:

$$y_{ac} = \beta_0 + \beta_1 es_c + ht_c'\lambda + x_c'\delta + \varepsilon_{ac} \tag{4}$$

where y_{ac} is the outcome of interest, es_c indicates that the couple was selected for (or, in some specifications, consented to) extra statements, and x_c is a vector of additional controls. To examine treatment effects by preference heterogeneity and household information sharing, in some specifications I also include the vector ht_c, which includes a dummy for poorly matched couples, a dummy for poorly informed couples, and the interaction of these variables with the extra statements indicator.

The first two columns of Table C1 verify that, as per experimental protocol, the probability of opening an individual account is uncorrelated with extra statement selection. This is in fact the case, and for the remaining specifications I limit the sample to opened individual accounts. The next two specifications examine extra statement consent rates. Panel A, Column 4 reveals substantial differences in consent rates by household information sharing. Poorly informed households were 28 percentage points less likely to consent to extra statements, though additional controls reduce the magnitude and significance of this estimate.

Columns 5-8 of Table C1 examine the reduced form impact of the extra statements offer on savings rates and average daily balances of open individual accounts. The overall impact of extra statements (columns 5 and 7) is relatively small, insignificant, and actually positive in sign. Given the low consent rate, this is not very surprising – individuals who would be most adversely impacted by the extra statements should have been most likely to refuse consent. However, columns 6 and 8 suggest that this aggregate zero impact may mask differences by information sharing. In particular, conditional on demographic and economic characteristics the extra statements intervention reduced savings rates for couples who are both poorly informed and poorly matched (I reject that the main effect plus the interaction effects for the poorly informed and matched are equal to zero at the 90 percent level or better in the last two panels). In contrast, there is some evidence that the extra statements intervention actually had a *positive* impact on well informed and well matched couples. It is possible that these couples viewed the extra statements as a commitment device to stick to cooperative savings plans and were therefore more likely to use their experimental

accounts.

D Supplementary Tables and Robustness Checks

Appendix Table D1 summarizes individual level demographic characteristics for well and poorly matched couples and tests whether characteristics differ by match quality. Appendix Table D2 presents additional randomization verification results. Panel A tests whether the empirical distributions of treatments differ from their theoretical analogs. This panel displays the share of individuals exposed to each treatment and the p-value from a binomial test that this share is equal to the theoretically expected share. Panel B tests whether treatments are correlated with one another. Each column corresponds to a regression where the treatment of interest is regressed on all other treatments.

Appendix Table D3 verifies that respondents understood and responded to the experimental interest rates. This table displays the results of account level regressions of the outcome of interest on interest rate dummy variables, a joint account dummy, and dummy variables for husband and wife cash payment selection.

Appendix Tables D4-D6 test the robustness of the main results to (1) alternative measures of preference heterogeneity, (2) including additional controls, and (3) alternative samples. Table D4 tests the robustness of the results in Table 2, while Tables D5 and D6 test the robustness of the results in Tables 3 and 4 respectively.

In addition to the main discount factor estimates, I estimate discount factors via nonlinear least squares with normally distributed errors, and via an *ad hoc* bounding strategy similar to that in Meier and Sprenger (2010). Specifically, suppose that for individual i, Ksh X at time t is preferred to Ksh 300 at time $t+\tau$, but Ksh 300 at time $t+\tau$ is preferred to Ksh $Y < X$ at time t. I assume that the individual is indifferent between Ksh 300 at time $t+\tau$ and the midpoint of the two amounts $\left(\frac{X+Y}{2}\right)$ at time t. I calculate the implied discount factor using this midpoint. I do this for each table of monetary choices, obtaining 10 discount factor estimates, and take the simple average of them.

The tables also present results using alternative rules of thumb for match quality, where I define either $1/3$ or $2/3$ of couples to be "well matched" based on discount factors. Table D5 also presents results where I use the absolute value of preference heterogeneity, $|het_c|$, as a measure of match quality. Finally, the robustness check tables include specifications where I drop all couples who were randomly selected to receive at least one cash payment, and specifications where I include controls for information sharing and consumption decision making.

Appendix Table B1. Proxied Banking Costs and Savings Behavior

	Couple Level			Account Level
	Saved Any	Saved Indiv.	Saved Joint	Saved
Panel A. Basic Controls				
Cost Index/High Cost	-0.381***	-0.870***	0.582***	-0.170***
	(0.113)	(0.167)	(0.177)	(0.063)
Excess				0.015***
				(0.005)
Excess×High Cost				-0.013*
				(0.007)
Panel B. + Time Preference Controls				
Cost Index/High Cost	-0.381***	-0.886***	0.588***	-0.181***
	(0.115)	(0.173)	(0.181)	(0.065)
Excess				0.016***
				(0.006)
Excess×High Cost				-0.015*
				(0.007)
Panel C. + Demographic Controls				
Cost Index/High Cost	-0.289***	-0.912***	0.546***	-0.199***
	(0.119)	(0.202)	(0.217)	(0.065)
Excess				0.018***
				(0.006)
Excess×High Cost				-0.017**
				(0.007)
Panel D. + Economic Controls				
Cost Index/High Cost	-0.326***	-0.789***	0.348	-0.213***
	(0.134)	(0.247)	(0.242)	(0.066)
Excess				0.020***
				(0.006)
Excess×High Cost				-0.018***
				(0.008)
DV Mean	0.430	0.677	0.358	0.164
N	598	257	257	897
Sample	All	Savers	Savers	Well Matched

Notes: Robust standard errors (clustered at the couple level when relevant) in parentheses. All regressions include dummies for husband and wife cash payment selection. Time preference controls include separate dummies for upper/lower censoring of the discount factors of each spouse and the estimated discount factor of each spouse. The demographic control set adds controls for spousal heterogeneity in age, education, number of children, and literacy. The economic control set adds controls for heterogeneity in income and phone ownership. The final column includes dummy variables that fully saturate interest rate×joint account, as well as a high cost×joint account, and interactions between the high cost dummy and the interest rate×joint account interaction set. ***, **, and * indicate significance at the 99, 95, and 90 percent confidence levels respectively.

Appendix Table C1. Impact of Extra Statements on Savings and Average Balances of Individual Accounts

	Protocol Check		First Stage		Reduced Forms				Two Stage Least Squares			
	Opened		Consented to ES		Saved		Average Balance		Saved		Average Balance	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel A. Basic Controls												
Extra Statement	-0.023	0.043	0.579***	0.765***	0.037	0.191*	230	529	0.065	0.280	397	805
	(0.052)	(0.085)	(0.067)	(0.109)	(0.056)	(0.105)	(180)	(451)	(0.097)	(0.171)	(312)	(700)
x Poorly Matched		-0.043		-0.129		-0.151		-362		-0.248		-682
		(0.104)		(0.137)		(0.111)		(335)		(0.238)		(674)
x Poorly Informed		-0.056		-0.280*		-0.187		-155		-0.342		-20.5
		(0.115)		(0.145)		(0.123)		(415)		(0.254)		(772)
Panel B. + Time Preference Controls												
Extra Statement	-0.030	0.038	0.560***	0.758***	0.055	0.193*	280	548	0.098	0.253	500	774
	(0.054)	(0.087)	(0.066)	(0.111)	(0.054)	(0.102)	(182)	(462)	(0.097)	(0.162)	(327)	(706)
x Poorly Matched		-0.039		-0.169		-0.093		-141		-0.113		-175
		(0.106)		(0.135)		(0.109)		(337)		(0.242)		(735)
x Poorly Informed		-0.059		-0.252*		-0.219*		-297		-0.382		-264
		(0.116)		(0.148)		(0.129)		(441)		(0.252)		(773)
Panel C. + Demographic Controls												
Extra Statement	-0.012	0.017	0.587***	0.798***	0.047	0.158*	155	1031**	0.081	0.160	264	1410**
	(0.054)	(0.095)	(0.063)	(0.096)	(0.050)	(0.095)	(120)	(461)	(0.087)	(0.135)	(206)	(637)
x Poorly Matched		-0.019		-0.224*		-0.147		-755**		-0.267		-1669*
		(0.110)		(0.123)		(0.115)		(379)		(0.231)		(870)
x Poorly Informed		-0.003		-0.153		-0.244*		-948*		-0.371*		-1330*
		(0.133)		(0.137)		(0.138)		(555)		(0.209)		(798)
Panel D. + Economic Controls												
Extra Statement	-0.009	0.024	0.593***	0.732***	-0.007	0.070	81.6	1007**	-0.011	0.053	138	1468**
	(0.055)	(0.098)	(0.061)	(0.107)	(0.056)	(0.097)	(151)	(441)	(0.094)	(0.149)	(254)	(643)
x Poorly Matched		-0.018		-0.147		-0.143		-828***		-0.366		-2013***
		(0.112)		(0.141)		(0.123)		(301)		(0.232)		(834)
x Poorly Informed		0.020		-0.107		-0.294**		-862*		-0.509**		-1452*
		(0.136)		(0.151)		(0.148)		(521)		(0.224)		(877)
DV Mean (ES=0)	0.395		0		0.273		241		0.273		241	
N	634		244		244		244		244		244	
Sample	All Accts.		Open Accts.		Open Accts.		Open Accts.		Open Accts.		Open Accts.	

Notes: Robust standard errors clustered at the couple level in parentheses. All regressions except for columns (1) and (2) limited to open individual accounts. All regressions include dummies for husband and wife cash payment selection. Time preference controls include separate dummies for upper/lower censoring of the discount factors of each spouse and the estimated discount factor of each spouse. The demographic control set adds controls for experimental session fixed effects, spousal heterogeneity in age, education, number of children, and literacy. The economic control set adds controls for heterogeneity in income, an indicator for subsistence farmers or the unemployed, and phone ownership. ***, **, and * indicate significance at the 99, 95, and 90 percent confidence levels respectively.

Appendix Table D1. Demographic Characteristics of Study Sample by Match Quality

	Well Matched	Badly Matched	Difference	N
Age	38.4	38.6	-0.211	1196
	[13.2]	[12.7]	(0.748)	
Education	7.01	6.98	0.025	1191
	[4.00]	[3.90]	(0.229)	
Literate	0.753	0.781	-0.028	1196
	[0.432]	[0.414]	(0.024)	
Number Children	4.61	4.72	-0.109	1196
	[2.76]	[2.73]	(0.159)	
Subsistence Farmer/No Job	0.416	0.457	-0.041	1191
	[0.493]	[0.499]	(0.029)	
Income Last Week	1044	1101	-56.8	1164
	[1936]	[2564]	(133)	
Owns Mobile Phone	0.497	0.415	0.083***	1191
	[0.500]	[0.493]	(0.029)	
Participates in ROSCA	0.582	0.567	0.015	1196
	[0.494]	[0.496]	(0.029)	
Has Bank Account	0.227	0.194	0.033	1196
	[0.420]	[0.396]	(0.024)	
Has a SACCO Account	0.039	0.034	0.005	1193
	[0.193]	[0.180]	(0.011)	
Saves at Home	0.889	0.860	0.030	1195
	[0.314]	[0.348]	(0.019)	
Has Mobile Money Account	0.207	0.234	-0.026	998
	[0.406]	[0.424]	(0.026)	
Saves Other Ways	0.608	0.511	0.097***	998
	[0.489]	[0.500]	(0.031)	
Total Reported Savings	14261	10768	3493	923
	[53389]	[25265]	(2766)	
Consumption - Husband Decides	0.396	0.452	-0.056**	1189
	[0.489]	[0.498]	(0.029)	
Consumption - Wife Decides	0.131	0.114	0.017	1189
	[0.338]	[0.318]	(0.019)	
Consumption - Both Decide	0.418	0.397	0.021	1189
	[0.494]	[0.490]	(0.029)	
Savings - Husband Decides	0.324	0.312	0.012	1189
	[0.468]	[0.464]	(0.027)	
Savings - Wife Decides	0.471	0.459	0.013	1189
	[0.500]	[0.499]	(0.029)	
Savings - Both Decide	0.178	0.209	-0.031	1189
	[0.383]	[0.407]	(0.023)	
Distance from Bank (Miles)	3.97	3.59	0.378***	1196
	[2.16]	[2.24]	(0.127)	
Weekly Discount Factor	0.844	0.579	0.265***	1196
	[0.274]	[0.396]	(0.020)	

Notes: Standard deviations in brackets, standard errors in parentheses. Mobile money and other savings data not available for the 98 couples in the first 6 experimental sessions. Variable recoded to missing if response was don't know/refused. ***, **, and * indicate significance at the 99, 95, and 90 percent confidence levels respectively.

Appendix Table D2. Additional Randomization Verification

Panel A. Do Realized Treatment Proportions Differ From Theoretical Treatment Proportions?

	Level Interest Rate on:				
	Husband's Account	Wife's Account	Joint Account	Extra Statements	Cash Payment
10% Interest Rate/Extra Stmts/Cash	0.232	0.239	0.319	0.456*	0.175**
	{0.345}	{0.571}	{0.488}	{0.054}	{0.030}
6% Interest Rate	0.242	0.264	0.355		
	{0.706}	{0.422}	{0.278}		
2% Interest Rate	0.256	0.266	0.326		
	{0.741}	{0.370}	{0.729}		
0% Interest Rate	0.269	0.231			
	{0.278}	{0.299}			

Panel B. Are Treatments Correlated With One Another?

	Excess Interest Rate on:				
	Husband's Account	Wife's Account	Joint Account	Extra Statements	Cash Payment
Excess Interest - Husband's Account		-0.736***	-0.714***	-0.001	0.002
		(0.024)	(0.022)	(0.008)	(0.004)
Excess Interest - Wife's Account	-0.811***		-0.767***	-0.003	0.001
	(0.025)		(0.021)	(0.008)	(0.004)
Excess Interest - Joint Account	-0.917***	-0.894***		0.001	0.003
	(0.021)	(0.019)		(0.009)	(0.004)
Extra Statements	-0.035	-0.075	0.015		-0.014
	(0.249)	(0.237)	(0.221)		(0.023)
Cash Payment	0.115	0.070	0.127	-0.026	
	(0.204)	(0.196)	(0.178)	(0.040)	
N	1196	1196	1196	1000	1196

Notes: Panel A shows the proportion of the sample selected for each treatment, as well as p-values (in braces) from binomial tests that the realized proportion is equal to the theoretical proportion. Theoretical probabilities are 1/4 for the first two columns, 1/3 for the third column, 1/2 for the fourth column, and 1/5 for the final column. Each column in panel B presents results of regressing the treatment of interest (column heading) on the remaining treatments. Regressions in columns 1-4 include an additional dummy variable for couples who were not eligible for the extra statements treatment. Standard errors clustered at the couple level are in parentheses. ***, **, and * indicate significance at the 99, 95, and 90 percent confidence levels respectively.

Appendix Table D3. Impact of Interest Rates on Account Takeup and Use

	Opened	Saved	Number Deposits	Average Balance
Panel A. All Accounts Offered to All Couples				
2% Interest	0.063*	0.010	0.023	9.11
	(0.035)	(0.021)	(0.085)	(37.7)
6% Interest	0.143***	0.065***	0.034	60.6
	(0.035)	(0.021)	(0.081)	(51.1)
10% Interest	0.209***	0.094***	0.305***	119**
	(0.037)	(0.022)	(0.104)	(51.7)
Joint	0.262***	0.190***	0.433***	120***
	(0.037)	(0.024)	(0.095)	(46.6)
DV Mean (Omitted)	0.288	0.054	0.164	49.9
N	1794	1794	1794	1794
Panel B. All Accounts Offered to Well Matched Couples				
2% Interest	0.032	0.015	0.009	27.6
	(0.047)	(0.026)	(0.123)	(38.8)
6% Interest	0.134***	0.065***	-0.027	99.8
	(0.050)	(0.027)	(0.105)	(75.4)
10% Interest	0.225***	0.122***	0.288**	191***
	(0.050)	(0.031)	(0.142)	(67.2)
Joint	0.287***	0.213***	0.463***	95.6
	(0.052)	(0.033)	(0.115)	(71.6)
DV Mean (Omitted)	0.288	0.033	0.144	20.1
N	897	897	897	897
Panel C. All Accounts Offered to Poorly Matched Couples				
2% Interest	0.094*	0.004	0.037	-10.3
	(0.051)	(0.032)	(0.118)	(64.8)
6% Interest	0.152***	0.062*	0.101	21.0
	(0.050)	(0.032)	(0.122)	(70.0)
10% Interest	0.194***	0.065**	0.324**	49.4
	(0.054)	(0.031)	(0.151)	(80.6)
Joint	0.239***	0.167***	0.403***	145***
	(0.053)	(0.034)	(0.151)	(60.9)
DV Mean (Omitted)	0.288	0.075	0.185	81.0
N	897	897	897	897

Notes: Omitted group is individual accounts, zero percent interest. Robust standard errors clustered at the couple level in parentheses. All regressions include separate dummies for husband and wife cash payment selection. ***, **, and * indicate significance at the 99, 95, and 90 percent confidence levels respectively.

Appendix Table D4. Robustness Checks - Responses to the Excess Interest Rate by Match Quality

Panel A. Well Matched Couples Only							
Excess Low×Indiv	-0.012	-0.003	-0.004	-0.003	0.007	-0.003	-0.006
	(0.008)	(0.006)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
Excess High×Indiv	0.010	0.010	0.009	0.019***	0.019***	0.016*	0.015*
	(0.009)	(0.007)	(0.009)	(0.008)	(0.008)	(0.009)	(0.009)
Excess Low×Joint	0.005	0.013	0.012	0.002	0.019	0.020	0.017
	(0.016)	(0.011)	(0.013)	(0.014)	(0.013)	(0.013)	(0.013)
Excess High×Joint	0.008	-0.001	0.002	0.011	0.013	0.005	0.005
	(0.015)	(0.010)	(0.013)	(0.012)	(0.013)	(0.012)	(0.013)
DV Mean	0.161	0.162	0.164	0.163	0.092	0.164	0.164
N	597	1197	897	897	588	897	897
Panel B. Poorly Matched Couples Only							
Excess Low×Indiv	-0.003	0.006	-0.003	0.001	-0.003	-0.006	-0.003
	(0.007)	(0.012)	(0.009)	(0.008)	(0.008)	(0.009)	(0.009)
Excess High×Indiv	0.006	0.005	0.007	0.000	-0.002	0.002	0.004
	(0.008)	(0.015)	(0.011)	(0.008)	(0.010)	(0.009)	(0.010)
Excess Low×Joint	-0.002	0.021	0.005	0.004	-0.037***	-0.006	-0.005
	(0.011)	(0.016)	(0.013)	(0.012)	(0.014)	(0.012)	(0.013)
Excess High×Joint	0.008	-0.041***	-0.021*	-0.009	-0.005	-0.017	-0.017
	(0.012)	(0.017)	(0.012)	(0.012)	(0.014)	(0.013)	(0.012)
DV Mean	0.174	0.184	0.175	0.176	0.107	0.175	0.175
N	1197	597	897	897	618	897	897
Panel C. Allowing for Heterogeneity Across Demographic Controls							
Well Matched Couples							
Excess Positive	0.013	0.014**	0.019***	0.016**	0.015**	0.025***	0.024***
	(0.009)	(0.006)	(0.008)	(0.007)	(0.007)	(0.007)	(0.007)
Excess Zero	-0.003	-0.001	-0.002	0.004	-0.008	0.005	0.003
	(0.006)	(0.005)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
Poorly Matched Couples							
Excess Positive	0.015*	0.001	-0.023**	-0.016	-0.003	-0.022**	-0.010
	(0.008)	(0.028)	(0.010)	(0.010)	(0.010)	(0.011)	(0.011)
Excess Zero	-0.010	-0.053***	-0.013	0.000	-0.021***	-0.013	-0.009
	(0.007)	(0.015)	(0.009)	(0.008)	(0.008)	(0.010)	(0.009)
Tests of Equality of Coefficients Across Equations							
Excess Positive	{0.883}	{0.651}	{0.001}***	{0.010}***	{0.104}	{0.000}***	{0.010}**
Excess Zero	{0.479}	{0.001}***	{0.293}	{0.694}	{0.107}	{0.109}	{0.253}
Joint Test - Both Splines	{0.778}	{0.003}***	{0.001}***	{0.020}**	{0.056}*	{0.000}***	{0.008}***
Specification	Logit - 33%	Logit - 67%	Probit - 50%	Ad Hoc - 50%	No Cash Payments	Info Sharing	Decn Making

Notes: Robust standard errors clustered at the couple level in parentheses, p-values in braces. All specifications include fixed effects that saturate interest rate×joint account and dummies for husband and wife cash payment selection, as well as the time preference, demographic, and economic control sets described in Table 5. The information sharing specification also includes a dummy for well informed households. The consumption decision making controls in column 7 include dummy variables for individual reports of whether the husband decides, the wife decides, both decide together, both decide independently, or someone else decides, as well as interactions between husband and wife reports. All controls are demeaned to the value among well matched couples. Specifications in Panel C also include interactions between the excess interest splines and the included control variables. Cross-equation F-tests performed using seemingly unrelated regression. ***, **, and * indicate significance at the 99, 95, and 90 percent confidence levels respectively.

Appendix Table D5. Robustness Checks - Preference Heterogeneity and Account Use

	Couples with Dominated Individual Account	
	All	Savers
Panel A. Logit NLLS, Well Matched - 33%		
Poorly Matched	0.159***	0.233
	(0.041)	(0.203)
DV Mean (Well Matched)	0.050	0.118
Panel B. Logit NLLS, Well Matched - 67%		
Poorly Matched	0.130***	0.175
	(0.049)	(0.201)
DV Mean (Well Matched)	0.097	0.230
Panel C.Probit NLLS, Well Matched - 50%		
Poorly Matched	0.155***	0.193
	(0.048)	(0.193)
DV Mean (Well Matched)	0.082	0.195
Panel D. Ad Hoc Discount Factors, Well Matched - 50%		
Poorly Matched	0.175***	0.280*
	(0.042)	(0.148)
DV Mean (Well Matched)	0.079	0.184
Panel E. Absolute Value Preference Heterogeneity		
Absolute Value of Heterogeneity	0.251***	0.388
	(0.066)	(0.264)
DV Mean (All)	0.132	0.306
Panel F. Drop Cash Payment Recipients		
Poorly Matched	0.068**	0.248**
	(0.032)	(0.115)
DV Mean (Well Matched)	0.036	0.167
Panel G. Control for Consumption Decision Making		
Poorly Matched	0.166***	0.308
	(0.048)	(0.189)
DV Mean (Well Matched)	0.078	0.182

Notes: Heteroskedasticity robust standard errors in parentheses. Sample size is N=363 and N=157 for columns 1 and 2 respectively, expect for Panel F, which has sample sizes of N=238 and N=57 respectively. All regressions include time preference, demographic, and economic controls sets except Panel F. This panel only includes the time preference control set because the small sample size cannot support larger control sets in the "savers only" specification. See notes to Table 3 for time preference, demographic, and economic control sets. The consumption decision making controls in Panel G include dummy variables for individual reports of whether the husband decides, the wife decides, both decide together, both decide independently, or someone else decides, as well as interactions between husband and wife reports. All regressions include dummy variables for each account's interest rate and dummies for husband and wife cash payment selection. ***, **, and * indicate significance at the 99, 95, and 90 percent confidence levels respectively.

Appendix Table D6. Robustness Checks - Interest Rate Losses by Match Quality (No Individual Interest Discounting)

	Logit 50%	Logit 33%	Logit 67%	Probit 50%	Ad Hoc 50%	No Cash Payment Recipients
Poorly Matched Couples						
Maximum Interest Earnings	8.15	8.07	8.29	8.18	8.11	8.04
Actual Interest Earnings	7.10	7.08	7.32	7.18	7.12	7.07
Loss	1.05	0.982	0.973	1.00	0.998	1.11
Well Matched Couples						
Maximum Interest Earnings	8.17	8.35	8.10	8.14	8.21	8.12
Actual Interest Earnings	7.50	7.74	7.29	7.43	7.49	7.26
Loss	0.664	0.610	0.801	0.713	0.719	0.960
Loss Gap						
A. No Controls	0.389**	0.372**	0.172	0.291*	0.279	0.315*
	(0.172)	(0.167)	(0.188)	(0.172)	(0.172)	(0.168)
B. + Interest Rate/Cash Prize Controls	0.541***	0.486***	0.311*	0.394***	0.450***	0.366**
	(0.157)	(0.154)	(0.179)	(0.159)	(0.162)	(0.168)
C. + Time Preference Controls	0.508***	0.481***	0.103	0.309	0.375**	0.265
	(0.211)	(0.180)	(0.240)	(0.206)	(0.190)	(0.286)
D. + Demographic Controls	0.486***	0.452***	0.087	0.363*	0.412**	0.248
	(0.208)	(0.177)	(0.249)	(0.203)	(0.189)	(0.303)
E. + Economic Controls	0.486**	0.440***	0.101	0.365*	0.415**	0.268
	(0.213)	(0.181)	(0.252)	(0.208)	(0.193)	(0.318)
F. + Information Sharing Controls	0.481**	0.433***	0.095	0.356*	0.404**	0.279
	(0.215)	(0.181)	(0.251)	(0.209)	(0.192)	(0.325)
G. + Decision Making Controls	0.477**	0.379**	0.098	0.356	0.391**	0.242
	(0.220)	(0.187)	(0.258)	(0.217)	(0.196)	(0.320)
N	598	598	598	598	598	402

Notes: Robust standard errors in parentheses. Interest rate controls include a set of dummy variables that fully saturate the set of experimental interest rates and dummies for husband and wife cash payment selection. See notes to Table 7 for time preference, demographic, and economic control sets. The information sharing control set includes a dummy for well informed households. The consumption decision making controls include dummy variables for individual reports of whether the husband decides, the wife decides, both decide together, both decide independently, or someone else decides, as well as interactions between husband and wife reports. Proxied banking costs constructed using a principal components index of distance from the bank, spouse specific indicators for subsistence farmers/the unemployed, and spouse specific indicators for baseline bank account ownership and SACCO membership. ***, **, and * indicate significance at the 99, 95, and 90 percent confidence levels respectively.

Appendix Figure D1. Interest Rate Design

$R_J=2$				
	$R_M=0$	$R_M=2$	$R_M=6$	$R_M=10$
$R_F=0$	2, -2,-2	0, 0,-2	-4, 4,-6	-8, 8,-10
$R_F=2$	0, -2,0	0, 0,0	-4, 4,-4	-8, 8,-8
$R_F=6$	-4, -6,4	-4, -4,4	-4, 0,0	-8, 4,-4
$R_F=10$	-8, -10,8	-8, -8,8	-8, -4,4	-8, 0,0

$R_J=6$				
	$R_M=0$	$R_M=2$	$R_M=6$	$R_M=10$
$R_F=0$	6, -6,-6	4, -4,-6	0, 0,-6	-4, 4,-10
$R_F=2$	4, -6,-4	4, -4,-4	0, 0,-4	-4, 4,-8
$R_F=6$	0, -6,6	0, -4,6	0, 0,0	-4, 4,-4
$R_F=10$	-4, -10,4	-4, -8,4	-4, -4,4	-4, 0,0

$R_J=10$				
	$R_M=0$	$R_M=2$	$R_M=6$	$R_M=10$
$R_F=0$	10,-10, -10	8, -8,-10	4, -4,-10	0, 0,-10
$R_F=2$	8, -10,-8	8, -8,-8	4, -4,-8	0, 0,-8
$R_F=6$	4, -10,-4	4, -8,-4	4, -4,-4	0, 0,-4
$R_F=10$	0, -10,0	0, -8,0	0, -4,0	0, 0,0

Notes: The first number in interior cells is the excess interest rate on the joint account. The excess interest rate on the husband's and wife's account follow respectively.