Intrahousehold Preference Heterogeneity, Commitment, and Strategic Savings*

Simone G. Schaner†

December 31, 2011

Abstract

This paper develops and tests a model of household bargaining where heterogeneity in discount factors drives inefficient savings behavior. The model highlights that when discount factors differ and individuals cannot commit intertemporally, the availability of a "private" savings technology (a device that is only accessible by a single owner) may incite individuals to take costly strategic savings action in order to manipulate the time path of consumption. To test this model and the efficient alternative, we designed and implemented a field experiment where married couples in rural Kenya were given the opportunity to open joint and individual bank accounts at randomly assigned interest rates. The behavior of individuals in the experiment is inconsistent with *ex-ante* Pareto efficiency and a variety of alternative models of intrahousehold resource allocation, but consistent with the proposed model of strategic savings. Savings misallocation due to strategic behavior is substantial: couples who are poorly matched on discount factors forgo at least 58 percent more potential interest earnings when compared to their well matched peers.

*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 Dan Keniston, Cynthia Kinnan, 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. All errors are my own.

†simone.schaner@dartmouth.edu

1 Introduction

Despite a great deal of research, there remains vigorous debate among economists as to whether households allocate resources in a Pareto efficient manner. While a substantial body of empirical work supports the efficient households hypothesis (a partial selection includes Bobonis 2009; Browning and Chiappori 1998; Chiappori et al. 2002; Rangel and Thomas 2005), a countering body of work rejects efficiency (see, for example, de Mel et al. 2009; Duflo and Udry 2004; Mazzocco 2007; Robinson 2011; Udry 1996).[1] In comparing these two groups of results, Chiappori and Donni (2009) note a striking pattern: most studies that support efficient households test for *static* efficiency (i.e. whether, in a given time period, households allocate consumption in a Pareto efficient manner) while most studies that reject the efficient model test for *dynamic* efficiency (this includes testing for efficient insurance and production arrangements within the household, as well as testing whether the household can *ex-ante* commit to bargaining weights assigned to its members). These results suggest that although households are capable of acting cooperatively, they have trouble enforcing the types of binding contracts that are generally needed for intertemporal efficiency.

However, the drivers of these intertemporal distortions remain poorly understood, as little research to date has moved beyond testing the efficient households paradigm to testing alternative models of intertemporal resource allocation.[2] This paper contributes to filling this knowledge gap by testing whether heterogeneity in intertemporal preferences incites individuals to take strategic action to influence the time path of household consumption. The intuition behind this premise is simple: suppose household members can commit to binding contracts and bargain cooperatively to set a savings and consumption plan. When members have different rates of time preference, they will be willing to pay to alter this plan – specifically, the more patient member would like to increase savings (to increase consumption in the future), while the less patient member would like to decrease savings (to increase consumption in the present). Absent commitment, individuals may therefore prefer to exploit a secure, "private" savings technology to strategically alter savings levels, even when saving this way is costly. For example, a patient wife could deposit a large share of her income in an individual bank account and refuse her husband's requests to make a withdrawal. Her impatient husband may also be able to save strategically: he could deposit less than his wife would like in a secure account. If there are transaction costs associated

[1] The workhorse model of the efficient household is the collective model, in which spouses with distinct utility functions bargain cooperatively with one another to arrive at a Pareto efficient allocation of resources (Chiappori 1988; Chiappori 1992).

[2] A notable exception is Duflo and Udry (2004), who argue that their results with respect to incomplete insurance are consistent with mental accounting.

with banking, it may not be worth the wife's while to make an additional deposit to increase savings levels.

We structure the analysis by first developing a model of bargaining in a two person household that formalizes the above intuition. Individuals in the model have access to two classes of savings devices: a private savings device and a "public" savings device (such as a joint bank account or home savings), which can be accessed by any member of the household. These savings devices differ in terms of rates of return and transaction costs. Differences in individual savings preferences are captured by differences in discount factors. The central theoretical result is that willingness to pay to deviate from the efficient cooperatively bargained consumption-savings allocation increases with discount factor heterogeneity. This implies that while couples who are well matched in terms of discount factors will invest their resources efficiently, poorly matched couples' savings decisions may be distorted by strategic action. This result gives rise to three testable predictions. First, we use the theory to identify cases in which households with greater preference heterogeneity will be more likely to save individually and less likely to save jointly, all else equal. Second, we characterize how efficient households respond to relative rates of return between different account types in an environment with heterogeneity in transaction costs. Finally, the model predicts that interest rate losses on savings deposits increase with preference heterogeneity.

The model informed the design of a field experiment, which we conducted in Western Kenya in the Summer of 2009. We gave 598 married couples the opportunity to open three savings accounts (two individual, one joint) bearing randomly assigned interest rates. 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 simple test of efficient investment: the household should always choose to save in the account that offers the highest rate of return. We also asked each respondent in the experiment a battery of questions to directly elicit estimates of discount factors, which we use to calculate measures of intrahousehold heterogeneity. We then combine these data with administrative data on bank account use to test the implications of the theory. We pre-specified our focus on heterogeneity in discount factors (and the associated heterogeneous treatment effects) before implementing the experiment or undertaking any data analysis.[3]

The empirical results are consistent with the testable theoretical predictions. We find that couples who are well matched in terms of discount factors make more intensive use of joint accounts, less intensive use of individual accounts, and respond to relative rates of return in a manner consistent with efficient investment. In contrast, poorly matched couples

[3]Pre-analysis documents detailing this plan were drafted in May and August of 2009 and are archived in the J-PAL hypothesis registry.

are completely insensitive to relative rates of return (we 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 – we 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.

We also test whether individuals make use of private accounts to hide savings from other members of the household. We find suggestive evidence that information sharing matters – households in which spouses are poorly informed about one another's finances are more likely to choose individual accounts, less likely to choose joint accounts, and more likely to reduce savings in response 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 empirical results are unchanged when accounting for intrahousehold information sharing. Furthermore, the main findings are not driven by a correlation between discount factor heterogeneity and other observable characteristics of couples. Finally, the patterns in the data are not consistent with other theories of household saving such as mental accounting or rules of thumb.

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.

The paper also contributes to a growing literature that studies savings and investment in developing countries (for a comprehensive review see Karlan and Morduch 2010). Specifically, individuals in developing countries often appear to favor secure but costly individual savings devices over more public forms of saving. We provide evidence that conflicting time preferences within the household is an important motivation for this behavior. Since many households in developing countries engage in entrepreneurial activities, this paper's insights are also relevant for understanding the widely noted heterogeneity in returns to household-run microenterprises (Fafchamps et al. 2011; de Mel et al. 2009; McKenzie and Woodruff 2008).

Finally, our paper relates to an established theoretical literature emphasizing the difficulties associated with aggregating preferences with differing discount factors (e.g. Browning

2000; Caplin and Leahy 2004; Feldstein 1964; Gollier and Zeckhauser 2005; Jackson and Yariv 2010; Marglin 1963; Zuber 2010). Here the contribution is twofold: first, very little of the existing literature explicitly focuses on the consequences of discount factor heterogeneity within the household.[4] Second, this paper is, to our knowledge, the first to assess the implications of this heterogeneity empirically.

The remainder of the paper is structured as follows: Section 2 presents the model of strategic savings behavior, Section 3 outlines the experimental design and derives testable implications of the theory, 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 A Model of Strategic Savings

The primary goal of the model is to illustrate how heterogeneity in time preferences generates incentives to behave strategically. Since the focus of inquiry is intertemporal efficiency, we present a stylized model with a single public consumption good (so the only relevant choice is *when* to consume) in which discount factors are the only source of strategic incentives.[5] In this section we set up the model and characterize the *ex-ante* Pareto efficient outcome that would obtain if agents could commit to binding intertemporal contracts and bargain cooperatively with one another. We then characterize how heterogeneity in time preferences impacts agents' willingness to pay to deviate from the cooperative solution. After describing the experiment in Section 3, we feed the theoretical results through the experimental design to derive testable predictions that can be taken to the data.

2.1 Model Setup

Economic Environment The household consists of two agents (i.e. spouses), A and B. They live in a two period world with one public consumption good, c. Individuals have deterministic income streams $\left\{y_t^A, y_t^B\right\}_{t=1}^{T=2}$ and must decide whether or not to save any of their income for the second period. Though individuals can save for the future,

[4] 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, Chiappori and Donni (2009) and Bernheim (1999). 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.

[5] The choice of a public good is for ease of exposition, not necessity – if we impose more structure on preferences and decision making, the model can be written with just private or both public and private consumption goods.

there is no borrowing in this economy.[6] Agents have perfect information regarding own and spousal income streams, preferences, savings strategies, and rates of return earned on savings. Households have access to three different savings technologies:

- A public/joint bank account, which yields rate of return R_J
- A private/individual bank account for agent A, which yields rate of return R_A
- A private/individual bank account for agent B, which yields rate of return R_B

What makes the "public" account public is that any member of the household can deposit and withdraw funds. In contrast, "private" accounts can only be accessed by a single owner (though balances are known to all members of the household).

We associate two types of costs with the savings technologies, to capture the fact that financial markets in developing countries are often characterized by very high transaction costs (Karlan and Morduch 2010). First, while it is free to make deposits into all accounts, withdrawals incur a fee, $w \geq 0$. Second, accounts have time and travel costs associated with them, which we refer to as "banking costs". The idea here is that the bank is located in town, while most individuals live outside of town. If individual i is going to town for some other reason in period t, the cost of travelling to the bank is low. However, if the individual must travel to town specifically to go to the bank, the cost is high. An important advantage of a joint account is that the household can always send the agent with the lowest travel cost to the bank. To capture this intuition with minimal complexity, we assume that travel costs are nonstochastic, but that the cost of travel for an individual account weakly exceeds the cost of travel for a joint account $\left(b^i \geq b^J \geq 0 \text{ for } i = A, B\right)$.

Preferences Both members of the household $i \in \{A, B\}$ exponentially discount per-period utility of the public consumption good c:

$$U_t^i = \sum_{\tau=t}^{T=2} \delta_i^{\tau-t} u\left(c_\tau\right)$$

where we assume the per-period utility function $u\left(\cdot\right)$ is continuous, increasing, and concave with $u'\left(c\right) \to \infty$ as $c \to 0$. Without loss of generality, we assume that agent B is more patient than agent A ($\delta_B \geq \delta_A$) for the rest of the theoretical discussion.

[6] A perfect savings and credit market without transaction costs would eliminate all scope for strategic behavior. However, the model would generalize to an environment with an imperfect credit market – this would just put constraints on some types of strategic behavior.

2.2 The Cooperative Outcome

Before discussing strategic savings behavior, we characterize the consumption allocation that would occur if agents bargained cooperatively with one another. We follow the collective approach of Browning and Chiappori (1998) and assume that whatever the nature of the household's cooperative bargaining process, the outcome is efficient. We can therefore represent the cooperative bargaining problem by writing the optimization problem of a social planner who puts weight μ on A's utility and weight $(1-\mu)$ on B's utility, where μ is A's relative bargaining power.[7] Then the collective (or cooperative) bargaining problem is:

$$\max_{\left\{c_t s_t^A, s_t^B, s_t^J\right\}_{t=1}^{T=2}} \sum_{t=1}^{T=2} \left[\mu \delta_A^{t-1} u\left(c_t\right) + (1-\mu)\, \delta_B^{t-1} u\left(c_t\right)\right] \tag{CB}$$

subject to

$$y_t^A + y_t^B + \sum_{a \in \{J,A,B\}} \max\left\{R_a s_{t-1}^a - w - b^a, 0\right\} \geq c_t + \sum_{a \in \{J,A,B\}} 1\left(s_t^a > 0\right)\left(s_t^a + b^a\right) \quad t = 1, 2$$

$$s_t^a \geq 0 \ \forall a \in \{J, A, B\}$$

where $1\left(\cdot\right)$ is the indicator function. Due to nonzero transaction costs, this problem is not convex. However, we can imagine the planner solving for the optimal savings allocation *conditional* on paying each relevant combination of transaction costs, and then selecting the plan that generates the highest utility. Since each conditional problem is convex, first order conditions will be necessary for an interior optimum. Then if savings is placed in account a:

$$u'\left(c_1\right) = R_a \Omega u'\left(c_2\right) \tag{1}$$

where the effective "household" discount factor, $\Omega \equiv \mu\delta_A + (1-\mu)\,\delta_B$, is just a bargaining power weighted average of individual discount factors.[8]

[7] In practice, bargaining weights will depend on "distribution factors" (in the parlance of Browning and Chiappori (1998)) that reflect individual bargaining positions. The model presented here abstracts away from the possibility that private savings accounts are used strategically to influence bargaining weights. Since this is an important alternative motivation for costly private savings, we discuss the possibility that this concern is generating the patterns in our data in Section 6.

[8] Note that in the case where $T > 2$, the effective discount factor is time varying and equal to $\rho_t \delta_A + (1-\rho_t)\,\delta_B$ where $\rho_t \equiv \mu\delta_A^{t-1} / \left[\mu\delta_A^{t-1} + (1-\mu)\delta_B^{t-1}\right]$. So as $t \to \infty$ the effective discount factor approaches that of the more patient agent. In other words, when individual discount factors differ, collective preferences are discounted hyperbolically. This observation reflects broader issues associated with aggregating individual preferences with differing discount factors. In particular, aggregated preferences will be time inconsistent as long as positive weight is placed on at least two agents with $\delta_j \neq \delta_i$ (Jackson and Yariv 2010).

2.3 Preference Heterogeneity and Strategic Behavior

How Much are Agents Willing to Pay to Deviate? Now suppose agent i is given a choice between accepting the collective solution to the savings problem or paying κ (and thereby reducing first period resources to $y_1^A + y_1^B - \kappa$) to select a savings and consumption path that maximizes his or her own utility. Clearly, when a couple is perfectly matched ($\delta_A = \delta_B = \Omega$) neither agent will be willing to pay to deviate, so there are no barriers to efficient savings among perfectly matched couples. However, as preference heterogeneity increases conditional on Ω, the maximum κ that an agent is willing to pay will increase. Note that when Ω is unchanged, the collectively optimal allocation also remains unchanged – thus we obtain a comparative static holding the cooperative optimum constant.[9] The following proposition formally states this result:

Proposition 1 *Fix incomes, Ω, $\mathbf{b^a}$, and w. The κ^* that makes an individual indifferent between accepting the collectively optimal consumption and savings allocation and paying κ^* to choose an individually optimal consumption and savings allocation is increasing in $|\delta_i - \Omega|$. Furthermore, when the collectively optimal allocation involves savings, κ^* is strictly increasing in preference heterogeneity for all $\delta_i \neq \Omega$.*

Proof. See Appendix A. ■

The intuition behind this result is very straightforward. When the couple bargains cooperatively, consumption and savings is dictated by the household discount factor, Ω. As δ_i moves away from Ω, agent i's most preferred allocation moves further and further away from the cooperative outcome – agent i is therefore increasingly willing to pay to deviate. Very similar arguments show that as preference heterogeneity increases, agents will be willing to accept lower interest rates and greater transaction costs on bank accounts that help them control the consumption allocation (these extensions are also established formally in Appendix A). This is intuitive – as agents become increasingly willing to pay to deviate from the collective allocation, they will be more willing to accept less favorable account terms in order to deviate.

The Nature of Strategic Savings Behavior The above analysis makes it clear that when discount factors in the household differ, agents will have incentives to take strategic action to deviate from the cooperative outcome. What might these deviations look like?

[9] This concordance does not hold when $T > 2$ – in this case the collectively optimal savings path will differ for couples with the same Ω but different degrees of preference heterogeneity. However, it is straightforward to show that conditional on δ_i, agent i's willingness to pay to change the cooperative allocation increases as δ_{-i} diverges from δ_i.

Private accounts offer a particularly attractive way to manipulate consumption and savings allocations – since agents have proprietary access to private accounts, they can choose savings levels in these accounts to maximize their own utility without regard for one another's welfare. When there are transaction costs, agents could also potentially make use of the joint account for strategic purposes. However, joint bank accounts can be accessed by either member of the household, and therefore hold less strategic value (e.g. if B makes a joint savings deposit that is very unappealing to A, A could always pay the cost to go back to the bank to make a withdrawal). To focus the discussion on the strategic benefit of private accounts, we therefore make the simplifying assumption that individuals do not use joint accounts strategically (in other words, agents always bargain cooperatively over joint savings levels).[10]

Given this assumption, all strategic deviations will involve either (1) depositing savings in a private account or (2) refusing to deposit savings in a private account when it is collectively optimal to save in that account. First consider the options of agent B, who is more patient. She wishes to save more than the collective optimum. She could achieve this goal by placing a large amount of savings in her individual account – as long as this increases second period consumption relative to the collective optimum, this will be appealing when she is sufficiently patient (this is easy to see in the extreme case where $\delta_B \to \infty$ so agent B only cares about the second period). When she has the most attractive account from the collective perspective, she could also deviate by *refusing* to save in her account. Since $\delta_B \geq \Omega$ this would only make sense if her refusal to save individually pushed the couple to save in the joint account. However, it is straightforward to show that this type of deviation will never be profitable.[11] Consequently, strategic behavior on the part of the more patient agent will always result in increased use of that agent's private account.

Now consider agent A, who is less patient. First suppose he has access to the most attractive savings device and the collective optimum entails saving in his private account. He could simply choose to save less than the collective optimum in his individual account, even to the point of refusing to save at all. When this deviation increases first period consumption A will find it attractive when he is sufficiently impatient (consider $\delta_A \to 0$). Even when A's account is dominated by the joint and/or B's account he may still be able to manipulate consumption streams in his favor. For example, suppose the household would

[10] A previous version of this paper (Schaner 2011a) formally solves a strategic savings game implied by this structure and characterizes the subgame perfect pure strategy Nash equilibria. Here we simply discuss the types of strategic deviations that *may* be optimal, as this is sufficient for the derivation of the testable predictions.

[11] If the couple collectively prefers B's account to the joint account, this implies $R_B \geq R_J$. Euler equations then imply that the collective optimum using B's account will entail greater second period consumption than the joint collective optimum. Since $\delta_B \geq \Omega$, this implies that B will prefer to save in her individual account over forcing the couple to save jointly.

collectively choose to save in the joint account if agent A took no action. Then A could rush to the bank and preemptively save "just enough" in his account to prevent the couple from travelling to the bank again to save jointly. If this "just enough" amount results in increased first period consumption, a sufficiently impatient agent will find this deviation profitable. In this case, the presence of transaction costs is essential – in their absence, this type of deviation would always result in decreased consumption in both periods.

Intuitively, one may expect that the use of individual accounts increases with preference heterogeneity. Yet the discussion above underscores that strategic behavior on the part of the less patient spouse may sometimes involve refusing to use an individual account – in some cases this will work to increase rates of joint account use.[12] However, we can obtain an unambiguous result by limiting attention to cases in which $R_J \geq R_A$: when this condition holds and the joint account offers banking cost savings ($b^J < b^A$), the joint account will always be strictly preferred from a collective perspective – in this case, all strategic savings deviations will work to increase rates of individual account use.

This analysis has generated two key insights that underpin the testable predictions that we will take to the data. First, as preference heterogeneity increases, individuals are willing to pay more to influence the consumption-savings allocation (Proposition 1). In other words, as match quality declines, individuals are willing to make more inefficient strategic savings choices. Second, as long as the joint interest rate exceeds the interest rate on the less patient spouse's account, this insight also implies that individual account use is increasing in preference heterogeneity, all else equal.

The ideal test of this theory would require randomly assigning discount factors to individual household members and evaluating the impact on savings behavior – this, of course, is not feasible. However, we can randomly assign interest rates to bank accounts, and compare the relative efficiency of investment choices made by couples with well versus poorly aligned time preferences. This idea inspired the design of a field experiment we conducted in Western Kenya, where married couples were given the opportunity to open three bank accounts (two individual accounts and one joint account) with randomly assigned interest rates. The next section describes this experiment in detail and then derives testable implications of the theory by overlaying the above insights with the experimental design.

[12]Specifically, when there are differential banking costs ($b^J < b^A$) and $\delta_A < \Omega$ agent A may prefer the collectively bargained joint savings allocation to any allocation involving use of this individual account.

3 Experimental Design and Testable Implications

3.1 Experimental Context

The experiment was conducted 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. It is only recently, however, that major banks have begun to offer products suitable for low income individuals. Traditionally, Kenyan bank accounts required minimum opening balances upwards of Ksh 1,000 (approximately equal to $12.50 at an exchange rate of Ksh 80 per $1, or $19.23 using a PPP exchange rate of Ksh 52 per $1) and charged monthly maintenance fees around Ksh 50 ($0.63).[13] However, recently banks have begun to target lower income individuals, and several banks currently offer lower fee alternatives to traditional bank accounts.

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 ($167 million) in assets, and actively targeted low and middle income individuals as part of its corporate strategy. All study participants were offered Family Bank's *Mwananchi* accounts. This account can be opened with any amount of money, though a minimum operating balance of Ksh 100 (approximately $1.25) cannot be withdrawn. The account pays no interest, but deposits are free of charge and there are no recurring maintenance fees. The only fees associated with the account are withdrawal fees, which are Ksh 62 ($0.78) over the counter and Ksh 30 ($0.38) with an ATM card. Account holders may purchase an ATM card for Ksh 300 ($3.75), though this is not mandatory.

3.2 Experimental Design

Targeted Population At the outset of the study, we 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. In these areas trained field officers issued meeting invitations to married couples 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.[14] Approximately 29

[13]For comparison, the median couple in the sample reported Ksh 1,200 in combined income in the week before the baseline survey.

[14]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

percent of issued invitations were redeemed over the course of the study. 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, we funded each opened account with the Ksh 100 ($1.25) minimum operating balance (this amount could not be withdrawn by participants – it simply made opening an account costless). Participants were randomly allocated into two core interventions, which are described below.[15]

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 some of the theoretical results are only unambiguous when the joint interest rate dominates the individual interest rate of the less patient spouse, we designed the experiment so the joint account had the greatest probability of bearing the highest rate of return: individual accounts could bear either 0, 2, 6, or 10 percent 6-month yields (with equal probability), and joint accounts could bear either 2, 6, or 10 percent 6-month yields (with equal probability). Figure 1 illustrates the design of the interest rate intervention. These interest rates were temporary, and expired after six months.[16] Since many respondents had low levels of education, enumerators explained what an interest rate was and provided numerical examples for each interest rate that was drawn.

The three interest rate draws were independent of one another, and therefore created random variation in the return of account a relative to its alternatives, even conditional on account a's interest rate. We define the "excess" interest rate on account a for couple c to be

$$excess_{ac} = R_{ac} - \max\{R_{\hat{a}c} : \hat{a} \neq a\} \tag{2}$$

In other words, the excess interest rate is the difference between account a's rate and the best alternative rate. The interior cells of Figure 1 illustrate the variation in the excess rate for joint, husbands' and wives' accounts respectively.

After observing their interest rates, couples were separated and each spouse was ad-

necessary in order to vote, buy or sell land, and seek formal employment.

[15] 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 (2011b). We do not discuss this intervention here, as accounting for it has no impact on any of the results presented in this paper.

[16] 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. 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.

ministered 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 baseline self reports of savings levels or savings device use, nor are they systematically associated with self reported decision making power regarding consumption and saving. We therefore expect that the randomization had little impact on survey responses. After the baseline, couples were reunited and decided which accounts they wanted to open.[17]

Intervention 2 - Extra Statements In order to test whether the ability to hide savings was an important driver of individual account use, we randomly selected 50 percent of participating couples for an "extra statements" offer.[18] If a selected couple decided to open an individual account for (without loss of generality) the wife, the enumerator processing the couple's paperwork asked if the spouses would consent to allow the husband to receive extra statement cards. The cards, if presented by the husband at the bank, entitled him to learn the current balance of his wife's account. These cards were only valid for 6 months, and were not given to couples unless both spouses gave their consent.

3.3 Testable Implications

We now merge the theoretical results and the experimental design to generate three key testable implications, which we take to the data in Section 4. To mimic the experimental context, consider a finite population of households where there is some nondegenerate distribution of preference parameters, bargaining weights, income, and banking costs.[19] As in the field experiment, all households in this population have access to three bank accounts with randomly assigned interest rates: R_A (the individual rate for the less patient spouse), R_B (the individual rate for the more patient spouse), and R_J (the joint rate).

The first testable implication characterizes patterns of overall account use by match quality. The central insight of the private savings theory is that as match quality decreases, individuals become increasingly willing to engage in costly strategic savings behavior. In many instances this strategic behavior will lead to increased use of individual accounts, but not always: the theoretical discussion highlighted that in some cases the less patient spouse in a badly matched couple may choose to *refuse* to use his or her account when it offers the

[17] 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 reminder to save.

[18] Due to delays in approvals from the bank, extra statements were not offered to the 97 couples (16 percent of our sample) in the first 6 (of 33) experimental sessions.

[19] Variation in banking costs is due to, for example, variation in distance from the bank and occupational mobility.

highest rate of return. However, if we limit the sample to couples for whom $R_A \leq R_J$, this type of strategic deviation is not relevant (since in this case it is never efficient to use the less patient spouse's account).

For this subsample, condition on the average household discount factor (Ω), incomes ($\mathbf{y}$), and banking costs ($\mathbf{b^a}$) so all households would choose the same allocation if they bargained cooperatively. Individuals in perfectly matched couples are not willing to pay to deviate from the collective optimum, so perfectly matched couples will always invest their savings efficiently (or not save at all, depending on parameters). Proposition 1 implies that as preference heterogeneity increases, individuals will be more willing to save strategically, which implies increased use of individual accounts since $R_A \leq R_J$.[20] Strategic savings will generate the reverse pattern with respect to joint accounts – rates of joint account use should be highest among perfectly matched couples and decline as match quality worsens. To summarize:

T1. Consider the population of couples for whom $R_A \leq R_J$. Conditional on Ω, $\mathbf{y}$, and $\mathbf{b^a}$, rates of individual account use will increase with preference heterogeneity, while rates of joint account use will decrease with preference heterogeneity.

The next testable implication characterizes the efficient response to relative rates of return on bank accounts. Here we exploit the fact that conditional on an account's interest rate, R_a, the experiment generated random variation in its relative rate of return, captured by $excess_a$. However, when there are differential banking costs (which is certainly the case in our rural Kenyan context), it will not always be efficient to save in the account with the highest interest rate: in some cases, the joint account will offer the highest return net of transaction costs, even when $excess_J < 0$. Specifically, for a given Ω, $\mathbf{y}$, $\mathbf{b^a}$, and individual interest rate R_i, suppose that a cooperatively bargaining couple would choose to save in individual account i. Then there exists a joint interest rate $\widetilde{R}_J(R_i) \leq R_i$ that would make this couple indifferent between the joint account and account i. We refer to $R_i - \widetilde{R}_J(R_i)$ as the *individual interest discount.* Adjusting individual interest rates downward by this discount is a simple way to put individual accounts on parity with joint accounts in terms of rates of return. We are now prepared to characterize how efficient savings choices change over a range of excess interest rates for a population with heterogeneous banking costs (and therefore heterogeneous individual interest discounts for a given R_i).

Panel A of Figure 2 graphs how the fraction of couples saving in account i varies with

[20] In theory, strategic behavior could lead some badly matched couples to save jointly if the couple is implementing a mixed strategy equilibrium to the savings game. We assume that the share of couples implementing mixed strategies is sufficiently small that it does not impact our overall results. In practice, the empirical results support this.

$excess_i$ conditional on R_i. In this context, increases in $excess_i$ correspond to decreases in R_{-i} and R_J. When $b^A = b^B > b^J$ it will never be efficient to save in account i if $excess_i < 0$.[21] At an excess rate of zero, if there is a mass of households who have equivalent individual interest rates that dominate the joint interest rate, we will observe a discrete jump up in the savings rate. Since the field experiment generated lumpy variation in the set of interest rates presented to a couple, 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, perfectly matched couples begin to save jointly at negative excess interest rates. On the graph, the lowest excess rate at which a 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 provided $b^J < b^i$. 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. To summarize, the second testable implication is:

T2. If couples save efficiently, use of individual accounts will begin to increase at an excess interest rate of zero and continue to increase until a plateau is reached at a positive excess interest rate. The use of joint accounts should increase over negative excess interest rates and then plateau at an excess interest rate of zero.

An attractive property of this implication is that unlike T1, it holds unconditional on Ω, $\mathbf{y}$, and $\mathbf{b^a}$. T2 therefore provides a useful benchmark to assess whether the savings behavior of well and poorly matched couples is consistent with efficient investment without having to directly address unobservable banking costs.

Suppose that in addition to actual interest rates, we were able to observe banking costs and could adjust individual rates by the relevant discount to obtain $\tilde{R}_J(R_i)$. Perfectly matched savers would always choose to save in the account associated with the highest return: $\max\left\{R_J, \tilde{R}_J(R_A), \tilde{R}_J(R_B)\right\}$ – if we could calculate the "banking cost adjusted"

[21] If we allow for $b^A \neq b^B$, then the share of couples using account i may increase as $excess_i$ increases even when $excess_i$ is negative. However, the share of couples using account i when $excess_i < 0$ should be small, as in many cases account i will be dominated by the joint account.

rate of return on their savings, they would bear no interest loss. In contrast, poorly matched couples would not always choose the account with the highest effective rate of return – as a group they would bear a positive interest rate loss. Moreover, conditional on $\mathbf{b^a}$, $\mathbf{y}$, and Ω, Proposition 1 implies that the maximum acceptable interest rate loss increases as match quality decreases. The final testable implication is therefore:

T3. Poorly matched couples will bear larger banking cost adjusted interest rate losses than perfectly matched couples. Conditional on $\mathbf{b^a}$, $\mathbf{y}$, and Ω interest rate losses will monotonically increase with preference heterogeneity.

An attractive property of these predictions is that there is no mechanical reason for all three to hold simultaneously. For example, rates of joint account use could be higher among well matched couples due to differing rules of thumb for how to save – but in this case, we would not find support for T2 and T3. Similarly, even if we find support for T3, there is no reason we need to observe the particular pattern of savings among well matched couples implied by T2. This is especially useful since we must rely 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 just one or two of the predictions. We now describe the data used to test these predictions.

3.4 Data

We use two data sources for this project – survey data from one-on-one baseline questionnaires (spouses were separated for the interviews), and administrative data on account use from the bank. 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 use of savings devices. The administrative data provided by the bank includes the first six months' transaction history of all accounts opened under the auspices of the project.

Measuring Rates of Time Preference We chose to elicit time preferences using choices between different amounts of money at different times, as opposed to different amounts of goods at different times. We made this choice for two reasons. First, Ashraf, Karlan, and Yin (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, cash lotteries made intuitive sense to respondents given that the group sessions revolved around bank accounts and savings.

We framed all questions 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 $t+\tau$ ($x^{t+\tau}$).[22] 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.[23] In order to make decisions salient, respondents were given a 1 in 5 chance of winning one of their choices. The questions involved sizable amounts of cash relative to respondents' incomes (for comparison, median reported weekly earnings in our 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 our field office, also located in Busia town.[24]

A substantial mass of men and women either always chose the nearer option or always chose Ksh 300 in the future.[25] For these individuals, the discount factor is not identified. To obtain an estimate, we 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 two imputed questions to the tables leaves us with 70 choices for each individual.

While these data allow us to calculate $(\beta - \delta)$ quasi-hyperbolic preference parameters, we focus on calculating a single exponential discount factor, as motivated by the two period model. As in Tanaka, Camerer, and Nguyen (2010), we use nonlinear least squares to estimate the discount factors. For each individual we 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 we assume $\varepsilon_{iq} \sim Logistic\left(0, \mu_i\right)$. Then

$$\Pr\left(x_q^t \succ x_q^{t+\tau}\right) = \Pr\left(\delta_i^t x_q^t > \delta_i^{t+\tau} x_q^{t+\tau} + \varepsilon_{iq}\right) = \frac{1}{1 + \exp\left(-\mu_i\left(\delta_i^t x_q^t - \delta_i^{t+\tau} 300\right)\right)}$$

[22] This method is common to most empirical studies that attempt to measure rates of time preference in developing countries. Examples include Ashraf, Karlan, and Yin (2006), Bauer and Chytilova (2009), Dupas and Robinson (2011), Shapiro (2010), and Tanaka, Camerer, and Nguyen (2010).

[23] 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. We 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, Lau, Rutstrom, and Sullivan 2004).

[24] 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 for saving anyway.

[25] 26 percent of men and 22 percent of women always waited for Ksh 300 in the future, while 15 percent of men and 14 percent of women always wanted the smaller, nearer amount.

Define the dummy variable $now_{iq} = 1\left(x_q^t \succ x_q^{t+\tau}\right)$. Nonlinear least squares solves

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

We topcoded $\hat{\delta}_i$ at $\overline{\delta}$, the value of $\hat{\delta}_i$ obtained via nonlinear least squares when $now_{iq} = 0$ for every question, and bottomcoded $\hat{\delta}_i$ at $\underline{\delta}$, the value of $\hat{\delta}_i$ obtained when $now_{iq} = 1$ for every question.[26]

As robustness checks, we also computed discount factors assuming normally distributed errors, and computed discount factors using a simple *ad hoc* bounding strategy similar to that in Meier and Sprenger (2010).[27] The correlation coefficient between the two sets of nonlinear least squares estimates is 0.98. The correlation between the logistic estimates and the *ad hoc* estimates is 0.96. As expected given such high correlations, results in this paper are very similar when using these alternative methods.

One concern with the discount factor estimates is that responses could reflect collective rather than individual preferences. Even though respondents answered questions separately from their spouses, their decisions may reflect the intrahousehold resource allocation process. If *ex-ante* commitment describes the household, spouses should make identical choices. Even when action is strategic, if the monetary amounts were inframarginal to savings decisions then choices should reflect the discount factor used to determine savings allocations. If respondents answered questions in this way, rather than according to their own preferences, the majority of intracouple variation in discount factors would be due to measurement error. This would generally bias us towards rejecting T1-T3.[28]

It is likely that the discount factor estimates reflect heterogeneity in multiple determinants of savings motives, including discount factors, risk aversion, state dependence, and so on.[29] We therefore prefer to interpret our empirical results by match quality as reflecting the impact of heterogeneity in savings motives more broadly, rather than heterogeneity in discount factors specifically (indeed, at the expense of simplicity and intuition, it is possible to rewrite

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

[27] Specifically, suppose that for individual i, $x_q^t \succ 300^{t+\tau}$, but that $300^{t+\tau} \succ x_{q+1}^t$ (where $x_q^t > x_{q+1}^t$). We then assume that the individual is indifferent between Ksh 300 at time $t+\tau$ and the midpoint of the two amounts at time t. Using this midpoint, we can then calculate the implied discount factor. We do this for each table, obtaining 10 discount factor estimates, and take the simple average of them.

[28] A notable exception is if measurement error were correlated with cognitive ability. This could potentially generate patterns in the data consistent with T1-T3. To test for this, we present regression specifications that control for spousal levels of and heterogeneity in education and literacy.

[29] We do note however, that Dupas and Robinson (2011) elicit estimates of discount factors and risk aversion in a Kenyan sample very similar to our own and find very low correlations between the two measures (0.12 for men and 0.06 for women) (Pascaline Dupas, personal communication).

the model in this way). A bigger concern would be if our measure of heterogeneity is correlated with other characteristics of couples that determine savings behavior for reasons unrelated to differences in savings motives (e.g. financial literacy, education, a general ability of spouses to work well together). To address this issue we present specifications that control for spousal *levels of* and *heterogeneity in* observable characteristics. Of course, we cannot account for unobservable correlates with elicited preference heterogeneity.

Sample Characteristics The sample consists of 598 non-polygamous married couples.[30] Table 1 presents baseline characteristics of the sample. Respondents are of relatively low socioeconomic status – husbands average 8 years of schooling, and their wives average just under 6 years. While most men are literate (87 percent), one third of women cannot read and write. On average, men reported earning Ksh 1,346 (about $17) in the past week, while women reported earning Ksh 797 ($10). However, median reported weekly earnings are substantially lower, at Ksh 700 and Ksh 300 for husbands and wives respectively.

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. Savings accounts with formal banks were less prevalent, particularly for women – while 30 percent of men reported owning a savings account at baseline (and those men reported substantial savings in their accounts), just 12 percent of women reported owning a savings account. All of the savings accounts owned at baseline were individual accounts. At baseline both spouses were most likely to report that the husband was in charge of deciding what to consume, while the wife was most likely to be reported as responsible for deciding how much to save.

Finally, Table 1 reports the average estimated discount factor (the distribution of these discount factors is illustrated in Panels A and B of Figure 3). 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 (Tanaka, Camerer, and Nguyen 2010; Ashraf, Karlan, and Yin 2006; Bauer and Chytilova 2009; Shapiro 2010). However, Dupas and Robinson (2011) and Giné et al. (2011) find very impatient preferences among samples of individuals in Kenya and Malawi respectively.[31]

Panel C of Figure 3 graphs a histogram of intrahousehold discount factor heterogeneity,

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

[31] Note that it is not essential that our estimated discount factors reflect the actual *level* values that govern most intertemporal decisions. What we require is that the intracouple heterogeneity in our estimates reflects true heterogeneity in the sample.

defined as the difference between the male and female estimated discount factors: $het_c = \hat{\delta}_c^M - \hat{\delta}_c^F$. While 13 percent of couples have identical discount factor estimates, many couples have estimates that differ substantially. When making comparisons between well and poorly matched couples we choose a rule of thumb that divides the sample into two equal sized groups: we label the 50 percent of the sample with the most closely aligned discount factors as well matched and refer to the remaining couples as badly matched.[32] This corresponds to $\left|\hat{\delta}_c^M - \hat{\delta}_c^F\right| \leq 0.193$ (this is equivalent to the couple's discount factors being within 0.528 standard deviations of one another). These well matched couples are delineated by dashed lines in Panel C of Figure 3.[33]

Table 2 checks for correlations between estimated discount factors and observables.[34] While demographic characteristics are essentially uncorrelated with estimated discount factors, variables related to savings are more significant. Most importantly, estimated discount factors are positively correlated with observed savings levels in Family Bank accounts, especially for well matched couples.[35] These correlations are reassuring and suggest that the estimated discount factors are capturing preferences over savings levels.

Randomization Verification All randomization was conducted in the field, by allowing respondents to draw folded envelopes from tins. Enumerators were carefully trained not to allow any participants a second draw from the tin if the participant was unhappy with his or her result. To ensure that this protocol was followed, we check the deviation of treatment shares from their theoretical probabilities in addition to checking to see if treatments are correlated with observables.

Table 3 presents the randomization verification exercise for interest rates, extra statements, and selection for cash payments for discount rate elicitation. Panel A presents the actual proportions of each treatment observed in the sample. P-values from a binomial test that the observed proportion equals the theoretical proportion ($1/4$ for each interest rate on individual accounts, $1/3$ for interest rates on joint accounts, $1/2$ for extra statements, and $1/5$ for cash payments) are presented in braces. While couples were somewhat less likely to be selected for extra statements and cash prizes than expected, the results suggest that

[32] Table A1 illustrates demographic differences between well matched and poorly matched couples.

[33] Results are similar when using a variety of alternative measures, including defining well matched couples to be the best matched 1/3 or 1/4 of couples, and to using the absolute value of the difference in spousal discount factors as a measure of preference heterogeneity.

[34] Missing demographic characteristics were recoded to zero and dummied out. This convention is held for the remainder of the paper.

[35] One concern is that the act of giving cash for discount rate elicitation choices, which could be deposited into Family Bank accounts, would mechanically cause this correlation. However, the correlation for well matched couples remains statistically significant at the 95 percent level when dropping all couples where at least one spouse was randomly selected to receive a cash payment.

enumerators and participants respected the experimental protocol.

Panel B tests for balance across observables. Here we regress treatments on observables and present p-values from F-tests that the coefficients on observables are equal to zero. We also check that the treatments are uncorrelated with one another. 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. The exception is that balance is somewhat off for the cash payments awarded to incentivize the discount factor elicitation. Most importantly, cash payment selection is significantly correlated with match quality – individuals in well matched couples were slightly more likely to be selected for a cash payment than their counterparts in poorly matched couples (19 percent versus 16 percent). Second, cash payment selection is also marginally significantly correlated with the joint interest rate (cash payment winners had a 0.5 percentage point higher joint rate on average). Fortunately, there is no mechanical reason why these payments should impact our main testable implications, which center about investment efficiency and account use patterns. However, in the interest of caution we control for husband and wife cash payment selection throughout the analysis (as expected the results are extremely similar if we omit these controls).

4 Results

4.1 Account Takeup and Interest Rate Responses

Despite the low cost to couples of opening all three accounts, only 5 percent of couples chose to do so: instead, 56 percent of couples chose to open only the joint account, 3 percent of couples only opened one individual account, 29 percent of couples opened two individual accounts, and the remaining 7 percent of couples opened a joint account and one individual account. When we informally asked couples why they decided not to open all the accounts, most couples stated that they did not have the resources to maintain three accounts, or that they only had use for the account(s) that they chose to open. 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 we 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.

The relative popularity of joint accounts in the sample could reflect large differential banking costs. In addition, couples may have been biased towards joint accounts because they decided which accounts to open together. However, since use of joint accounts is higher

even conditional on opening (42 percent of open joint accounts were used for saving while 25 percent of opened individual accounts were used for saving), it seems unlikely that social pressure at the meeting is entirely responsible for the high rates of joint account use.

Table 4 examines the impact of interest rates on account opening and use. It presents results from regressions of the following form:

$$y_{ac} = \alpha_0 + int'_{ac}\beta + x'_c\delta + \varepsilon_{ac} \tag{3}$$

where y_{ac} is the outcome of interest, int_{ac} is a vector including interest rate dummy variables, a joint account dummy, and interactions between the joint account and interest rate dummies, and x_c includes dummy variables for husband and wife cash payment selection. The table examines several measures of account use. Using the Family Bank transaction data, we calculated the average daily balance of account a owned by couple c using 6 months of data from the day of account opening (this covers the period where the account was eligible to earn interest). The average daily balance does not include the Ksh 100 minimum operating balance deposited into every open account, since this amount could not be accessed and was deposited on behalf of the respondents. All unopened accounts were assigned an average daily balance of zero. In addition to the average daily balance, we also created a savings indicator for accounts with positive average daily balances, and calculated total deposits and withdrawals over the 6-month study period, the total number of account transactions, and the total amount of fees charged to the account.

Though the impact of the interest rates on savings balances is theoretically ambiguous, the income effect should be small since the interest rates were temporary. Table 4 confirms that higher interest rates are associated with significantly higher rates of account opening and use. Indeed, while 43 percent of the couples in our sample saved in at least one account, the second column illustrates that only 5 percent of individual accounts with no interest were used for saving. In contrast, 13 percent of individual accounts with 10 percent interest were used for savings. The table also illustrates that joint accounts were much more popular among the study couples than individual accounts – for example, 34 percent of joint accounts with 10 percent interest were used for savings. However, the interest rate interaction terms reveal that while level use of joint accounts is much higher, the slope with respect to the interest rate does not differ between the two account types. The most important takeaway from Table 4 is that couples understood, and robustly responded to, interest rates – an essential precondition to analyzing responses to relative rates of return.[36]

[36]Since free ATM cards were only issued to open accounts (and are therefore correlated with the interest rate), we also checked to see if the results in Table 4 were driven by the free ATM treatment. Accounting for this treatment has little effect on estimated coefficients and does not change the result that couples robustly

4.2 Preference Heterogeneity and Account Use

The first testable prediction outlined in Section 3 (prediction T1) is that individual account use increases with preference heterogeneity, while joint account use decreases with preference heterogeneity. Recall that this prediction is only unambiguous when the less patient spouse's account bears an interest rate that is weakly lower than the joint interest rate. Figure 4 explores the relationship between account choice and preference heterogeneity graphically, using the subsample of couples who meet this criterion. The figure presents results of the following local linear regression

$$y_c = g\left(het_c\right) + \varepsilon_c \tag{4}$$

where y_c is the outcome of interest and $het_c = \hat{\delta}_c^M - \hat{\delta}_c^F$. Panel A graphs savings rates in joint accounts and 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). 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. In contrast, rates of joint account use are highest among well matched couples and lower for poorly matched couples. There does, however, appear to be a sharp increase in joint account use among couples where the wife is much more patient than the husband. Recall that the testable prediction holds conditional on savings propensity (i.e. conditional on Ω, $\mathbf{b^a}$, and $\mathbf{y}$) – it may be that these couples simply have particularly high likelihoods of saving overall due to different parameter values. To help account for this, Panel B limits the sample to savers and graphs the fraction of total savings stored in individual and joint accounts (since most couples only saved in one account, the analogous graph using savings dummies is very similar). This greatly attenuates the uptick in joint account use among very poorly matched couples where the wife is more patient. We also graph results for average daily balances in Panel C. Though there is no sharp theoretical prediction for individual accounts (as preference heterogeneity increases more patient individuals will wish to use individual accounts to store increasing amounts of savings, while less patient individuals will wish to store decreasing amounts of savings), balances in joint accounts should decrease with preference heterogeneity, all else equal. Overall, patterns for average balances are not very clear, but this is not entirely surprising. Account balances are generally low with some very large outliers (for example, among couples who saved, the median average daily balance was Ksh 290, while the 95th percentile was Ksh 4,639 and the 99th percentile was Ksh 9,794). As such, outliers have large impacts on means.

The following regression tests the significance of the graphical results while controlling

respond to the interest rate. These results are available from the author upon request.

for determinants of savings propensity and other potentially confounding factors:

$$y_c = \beta_0 + \beta_1 badmatch_c + \beta_2 imphigh \times badmatch_c + \beta_3 imphigh_c + int_c'\delta + x_c'\lambda + \varepsilon_c \quad (5)$$

where y_c is the outcome of interest, $badmatch_c$ indicates poorly matched couples, and $imphigh \times badmatch_c$ indicates poorly matched couples where the less patient spouse's account has a strictly higher return than the joint account. Less patient spouses in these households may refuse to use their individual accounts, potentially pushing the household to save jointly. In other words, these spouses may have attractive "joint deviations" that make account use patterns ambiguous. We therefore expect β_1 and β_2 to be of opposite signs. We also control for a vector of dummy variables that fully saturates the combination of experimental interest rates on each account (int_c) and a vector of additional controls (x_c). The omitted match quality group is well matched couples.[37]

Table 5 presents estimates of β_1 ("Poorly Matched") and β_2 ("Joint Deviation"), with four separate control sets. The "basic" control set, in Panel A, only includes the interest rate dummy variables, the joint deviation main effect ($imphigh_c$), and dummies for husband and wife cash payment selection. This regression essentially mirrors the results of Figure 4. Panel B then adds additional controls related to rates of time preference. To account for general differences in patience levels between well and poorly matched couples (Ω), we control for $\frac{\hat{\delta}_c^M + \hat{\delta}_c^F}{2}$ linearly.[38] The discount factor estimates are censored for some individuals (those who were always-patient and always-impatient, as well as those individuals whose estimates went outside the bounds implied by the always-patient and always-impatient response sets). This censoring would flatten the curves presented in Figure 4. To address this, we include two dummy variables for each spouse – one indicating upper censoring and one indicating lower censoring.

Panel C adds controls for a variety of demographic characteristics. Panel D adds additional controls for economic characteristics. A natural concern is that discount factor heterogeneity is correlated with other forms of intracouple heterogeneity. To control for both spousal levels of and heterogeneity in observables, for non-binary characteristics we 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 binary variables, we include the dummy variable

[37]Note that although $badmatch_c$ is a generated regressor, under the null hypothesis $\beta_1 = \beta_2 = 0$. In this case, traditional standard errors are consistent (Newey and McFadden 1994). Throughout the paper, we therefore present either heteroskedasticity robust standard errors (for couple level regressions) or standard errors clustered at the couple level (for account level regressions).

[38]Choosing a 50/50 weighting for the average discount factor in a couple is somewhat *ad hoc*, as the true weighting should be related to the individual bargaining weights. However, results are robust to a wide variety of different weighting schemes, including proxying the bargaining weight with relative income shares and using self reported savings decision making power to construct weights.

for both husband and wife as well as the interaction.[39] Therefore each non-binary characteristic contributes 5 covariates to the regression, while binary characteristics contribute 3.

Consistent with the theory, poorly matched couples are significantly more likely to save in individual accounts (17.7 percentage points/143 percent in Panel D) and marginally significantly less likely to save in joint accounts (10.4 percentage points/34 percent). Moreover, poorly matched couples with available "joint deviations" are indeed more likely to use joint accounts and less likely to use individual accounts when compared to poorly matched couples where such a deviation is not theoretically profitable. Limiting the sample to savers, poorly matched couples store a significantly higher share of balances in individual accounts and a significantly lower share of balances in joint accounts. In contrast, estimates for average balances are generally insignificant and the point estimate on the poorly matched dummy is positive for joint accounts when additional controls are added. However, this is largely due to outliers – all coefficients are of the expected sign (and significant for individual accounts) if we topcode balances at the 95th percentile.

We now turn to the second testable prediction (T2). First we generate the empirical analog of Figure 2 by running the following regression separately for well and poorly matched couples:

$$saved_{ac} = \beta_0 + ex'_{ac}\delta + z'_{ac}\lambda + \varepsilon_{ac} \tag{6}$$

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.[40] We 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 5 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

[39] Demographic controls include age, years of education, an individual literacy dummy, the number of children reported by each spouse, 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.

[40] 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, we downcode the excess interest rate by two percentage points (results are invariant to simply dropping these accounts). Similarly, we pool $excess_{ac} = -10$ and $excess_{ac} = -8$ as the omitted category in the regressions. We do this in order to identify all interest rate dummy variables, as accounts with zero percent interest had excess interest rates unique to them. Results are essentially unchanged if we drop all accounts with a zero percent interest rate.

the largest negative value of the excess interest rate in the sample, we 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 $excess_{ac} = 0$.

Panel A presents account use for well matched couples. The results for individual and joint accounts are both consistent with theoretical predictions. Specifically, the share of couples using individual accounts is essentially flat 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 slopes upward 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 $E_i\left(R_i\right)$ and $-E_J\left(R_J\right)$ illustrated in Figure 2; in other words, the excess interest rate was never large enough to dominate differential banking costs for the vast 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 75th percentile total average daily balance *among savers* was Ksh 782. Even foregoing 10 percentage points of interest would only cost the 75th percentile household Ksh 78. In comparison, the cost of round trip travel (via a mix of public and private transport) to and from the bank meets or exceeds Ksh 100 for over half the sample. We therefore conjecture that interest rate differentials only trumped banking cost concerns for couples with small banking cost differentials.

Panel B of Figure 5 illustrates the behavior of poorly matched couples. Their 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 are "cheap" in terms of interest rate losses. In this case, it could be that the majority of badly matched couples with small differential banking costs chose to save privately, while poorly matched couples with large differential banking costs chose to save jointly, and very few marginal couples were swayed by the excess interest rate.

We now test the significance of the patterns in Figure 5. To do so, we generate splines in the excess interest rate. Consistent with Figure 5, we place a knot at $excess_{ac} = -2$ for individual accounts and a knot at $excess_{ac} = 0$ for joint accounts. We then run the following

regression separately among well matched and poorly matched couples:

$$\begin{aligned} saved_{ac} &= \beta_0 + \beta_1 (below \times indiv)_{ac} + \beta_2 (above \times indiv)_{ac} + \\ &\quad \beta_3 (below \times joint)_{ac} + \beta_4 (above \times joint)_{ac} + z'_{ac}\lambda + x'_c\alpha + \gamma_s + \varepsilon_{ac} \end{aligned} \tag{7}$$

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×interest rate interactions, and x_{ac} is a vector containing additional controls. Panels A and B of Table 6 presents the results.

The first column only includes additional controls for husband and wife cash prize selection and their interactions with the excess interest rate splines – this column essentially mirrors the results in Figure 5 and confirms that the positive responses exhibited by well matched couples are indeed 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 results in a 19 percentage point increase in the savings rate, while increasing the joint excess interest rate from -10 to 0 results in a 25 percentage point increase in the savings rate. These impacts are very large when compared to dependent variable means (8 percent for individual accounts and 31 percent for joint accounts). In contrast, column 1 of Panel B illustrates that we cannot reject that poorly matched couples are completely insensitive to the excess interest rate. Panel C then uses seemingly unrelated regression to test the equality of interest rate responses among well and poorly matched couples. We reject equality of β_2 and β_3 (the splines where well matched couples exhibit a significant positive response), as well as overall equality of the interest rate responses.

An important concern is that the above differences by match quality are actually driven by differences in other couple-specific characteristics that are correlated with match quality. To address this, the last three columns of Table 6 test the robustness of our initial results by allowing the response to the excess interest rate to vary with other observable dimensions of heterogeneity. Specifically, we progressively allow for heterogeneous treatment effects with respect to the time preference, demographic, and economic characteristics included in the control sets in Table 5. To do this, we include main effects for each characteristic and their interactions with the excess interest rate. We demean each included covariate using means *among well matched couples* before generating the interaction terms and running the regression. When demeaned this way, the coefficients on the splines in Table 6 reflect the response to the excess interest rate at the average value of included demographic characteristics observed among well matched couples. If the heterogeneous responses observed in column 1

were driven by an included demographic control (for example, heterogeneity in education), then the results for poorly matched couples would mirror those for well matched couples once the control is included.

As expected given random assignment of the excess interest rate, the results for well matched couples do not substantially change with the addition of controls. The only major change observed for poorly matched couples is that the response to positive excess interest rates on the joint account becomes *negative* and significant after adding controls. This suggests that differences in observables between well and poorly matched couples are not driving the initial heterogeneous treatment effects. Furthermore, we continue to reject equality of responses to the excess interest rate after adding controls.

Taken together, the results in Tables 5 and 6 are incompatible with a model where couples can commit intertemporally. Specifically, we only observe evidence of efficient investment among well matched couples. Though badly matched couples' lack of a response to the excess interest rate could be rationalized if they had substantially larger differential banking costs, we observe them making more intensive use of individual accounts. Furthermore, it is not the case that badly matched couples do not understand the interest rates – results not shown here confirm that the interest rate treatments do significantly impact the account opening and use choices of badly matched couples. This suggests that badly matched couples' response to the excess interest rate reflects inefficient investment rather than demographic differences.

4.3 Proxying Banking Costs

The results so far suggest that banking costs are an important concern for couples in the sample – both well and poorly matched couples were much more likely to use joint accounts, and the excess interest rate responses among well matched couples in Figure 5 imply that even an excess rate of 8 (for individual accounts) or -8 (for joint accounts) did not outweigh banking cost concerns for a nontrivial share of the couples in the sample.

The unobservability of banking costs makes the final testable implication (T3), that well matched couples invest their resources more efficiently than poorly matched couples, difficult to analyze. We therefore assess investment efficiency three different ways: first, we make no adjustments to the interest rates. Second, we uniformly discount individual interest rates to reflect the fact that they have higher banking costs. Finally, we attempt to proxy banking costs with observables, and discount individual interest rates proportionally to proxied costs. For this final exercise we 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 we 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 we include both bank accounts and SACCO accounts as formal accounts.[41] To aggregate these measures, we used principal components analysis to extract the first principal component of the data matrix formed by the above-listed variables (results are similar if we convert the measures into standard deviation units and equally weight them). We then normalized this component to construct a "banking costs index", which runs from zero (lowest hypothesized banking costs) to one (highest hypothesized banking costs). Table 7 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:

$$y_c = \beta_0 + \beta_1 index_c + x_c'\delta + \varepsilon_c \tag{8}$$

where y_c is the outcome of interest, $index_c$ is the banking cost index, and x_c is a vector of controls. We use the same control sets as in previous regressions, but we 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 7 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 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.[42] As expected, 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

[41] 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.

[42] Since the cost index is correlated with absolute banking costs/socioeconomic status, the correlation between the index and joint account use is technically ambiguous. 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 us away from finding a positive correlation between joint account use and the cost index.

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, we 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. We 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, we 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:

$$\begin{aligned} saved_{ac} &= \beta_0 + \beta_1 high_c + \beta_2 excess_{ac} + \beta_3 high_c \times excess_{ac} + \beta_4 joint_{ac} + \\ &\quad \beta_5 \left(joint \times high\right)_{ac} + z'_{ac}\lambda + x'_c\delta + \varepsilon_{ac} \end{aligned} \tag{9}$$

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 Table 7.

The fourth column of Table 7 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 important determinants of account use, such as socioeconomic status and absolute banking costs.

4.4 Investment Efficiency by Match Quality

The results in Table 6 and Figure 5 imply that poorly matched couples invest less efficiently than well matched couples. But how big are these distortions in economic terms? We now exploit the experimental design to estimate the magnitude of savings misallocation among couples. To measure misallocation we calculate the maximum interest rate that each couple could have earned on their balance, compare this to the actual interest rate they earned, and then check to see if interest rate losses are greater for poorly matched couples. One issue is how to treat couples that did not save. In some cases patient individuals in poorly matched

couples may make costly savings deposits in individual accounts when the efficient response is to not save at all (due to very high banking costs, for example). To account for this, we code interest rate losses to zero for all non-savers. In practice, results are very similar when we either drop all non-savers or code the loss for non-savers to be equal to the maximum possible interest rate less the maximum interest rate across opened accounts.

Since lower return joint accounts may be more efficient than higher return individual accounts when banking cost differentials are large, we also present results where we discount individual interest rates to reflect higher banking costs. Recall that given an average household discount factor, Ω, a vector of banking costs, $\mathbf{b^a}$, an endowment $\mathbf{y}$, and an individual interest rate, R_i, we can calculate the joint interest rate that would make a well matched couple indifferent between saving jointly and individually, $\tilde{R}_J(R_i)$. The "individual interest discount", $R_i - \tilde{R}_J(R_i)$ then reflects the importance of differential banking costs.

In this spirit, we calculate "banking cost adjusted" actual and maximum interest earnings for all savers. To do so, we adjust individual interest rates downward to account for banking costs in two different ways. First, we use the cost index that we constructed in the previous subsection. To perform the adjustment we simply multiply the cost index by an assumed maximum joint interest discount and adjust individual interest rates by the resulting product. However, the cost index is an imperfect proxy and likely correlated with other important determinants of savings. As a further robustness check, we present a set of specifications where we 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.

We first study investment efficiency graphically – recall that Proposition 1 states that willingness to pay to change the cooperative allocation increases in preference heterogeneity. We therefore expect to see increasing interest rate losses as match quality degrades. Figure 6 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 adjustments made with proxied banking costs. Results are very similar when individual interest rates are adjusted uniformly).

Table 8 compares interest rate losses by match quality and presents results that include additional controls for observables. The first column does not discount individual interest rates at all. Here we see that 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 0.4 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 their well matched peers. Subsequent rows of Table 8 test robustness of this result by first controlling for cash prize selection, as well as a vector of dummies that saturate experimental interest rates on the three accounts. We then include the time preference, demographic, and economic control sets respectively. The loss gap remains significant and its magnitude grows in size upon including additional controls.

The next three columns repeat this analysis where individual interest rates are discounted using proxied banking costs and maximum interest discounts ranging from 5 to 15 percentage points (recall that well matched couples' response to the excess interest rate in Figure 5 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 by the stated value – these results are quite similar to the results incorporating proxied banking costs. Overall, the losses in Table 8 are very large in percentage terms, but relatively small in absolute terms. A loss of 3 percentage points in interest amounts to Ksh 24 ($0.30) for the 75th percentile saving couple. However, since banking cost differentials persist for the life of the account, long run absolute losses due to inefficient individual account use could be quite large.

Taken as a whole, the core results fit the predictions of the private savings theory very well. A key feature of this theory is the assumption of perfect information: individual accounts are valuable because they offer *security* of savings. However, individual accounts may also enable agents to *hide* savings from their spouses, which 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. This theory is similar to the theory of private savings, but rests on the assumption that in equilibrium, spouses are (and benefit from) concealing information from one another. If the benefit of hiding savings is correlated with preference heterogeneity, this could generate the patterns we observe in the data. The next section tests this idea.

5 Hidden Information and Account Use

Hidden information appears to be important in households in developing countries. 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.

Moreover, there is evidence of hidden savings in our data. For example, when considering spousal cross reports for saving at home (which is the most common method of saving among couples, and arguably the most observable), 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 intracouple discount factor heterogeneity, then hidden savings concerns could be responsible for the empirical results (though it is not obvious that such a correlation should exist). We added the "extra statements" treatment to the field experiment in order to assess the importance of hidden savings in the study population, and to gauge if hidden savings concerns are correlated with preference heterogeneity. We only present the extra statements analysis for a subset of couples in the data, as some enumerators did not follow the proper protocol when offering extra statements. In the analysis we also make use of baseline spousal cross reports of income and savings device use to construct an "information sharing index" which ranges from 0 (worst informed couples) to 1 (best informed couples). We map this index to responses to the extra statements treatment to assess whether couples in which individuals are poorly informed about spousal finances are more likely to value hidden information. Appendix B discusses how we identified protocol problems, the creation of the extra statements analysis sample, and the creation of the information sharing index, which measures how well informed spouses are about one another's income and savings device use.

Table 9 presents the results of the extra statements intervention. All regressions are of the following form:

$$y_{ic} = \beta_0 + \beta_1 es_c + ht_c'\lambda + x_c'\delta + \varepsilon_{ic} \tag{10}$$

where y_{ic} 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 we 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.[43]

The first two columns of Table 9 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 we limit the analysis to opened individual accounts that were eligible to be randomly selected for the extra statements offer.

The next two specifications examine extra statement consent rates among open individual accounts. Only 60 percent of individuals who were presented with the extra statement offer consented – this suggests that informational concerns were important to a sizable fraction of couples choosing individual accounts. 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 9 examine the reduced form impact of the extra statement 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 to extra statements, this is not very surprising – individuals most likely to be 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 poorly informed and poorly matched couples (we reject that the main effect plus the interaction effects for the poorly informed and matched are equal to zero at the 90 percent level 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. One possibility is that these couples considered the extra statements to be a commitment device to stick to cooperative savings plans, and were therefore more likely to use their experimental accounts.

Overall, the extra statements intervention is somewhat underpowered, so we choose to interpret the results with caution. However, the results suggest that the ability to hide savings from other household members is indeed an important concern for at least some households in the sample. Moreover, the information index appears to identify couples for whom hiding information is particularly valuable. As such, if the earlier results were driven by hidden savings concerns, we would expect to see two things: First, the information sharing index should be lowest among poorly matched couples. Second, properly accounting for informa-

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

tion sharing in our 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 the difference in individual discount factors is 0.017). This foreshadows that accounting for information sharing has no impact on the core results.

We check to see if accounting for information flows changes our findings regarding match quality and joint versus individual account use. Table 10 presents the regression described by equation 5 augmented to include the poorly informed indicator. We make two observations – first, the initial results pertaining to preference heterogeneity are robust to controlling for information flows. Second, the poorly informed indicator has predictive content for household account choice conditional on preference heterogeneity. As expected, poorly informed households are less likely to use joint accounts and more likely to use individual accounts. Moreover, additional specification testing not shown here confirms that information sharing is not driving the results regarding responses to the excess interest rate.

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 we observe. To complete our discussion of robustness, the following section considers other alternative explanations for the main results.

6 Alternative Explanations

In setting up the model, our 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. Though the analysis in the previous section directly addressed one way individual accounts could be used to change consumption (by hiding information), other channels are possible. For example, secure individual stores of savings could improve outside options and thereby increase bargaining power. The existence of other motivations does not pose an inherent problem for our results (indeed, accounting for information sharing had no impact on our core results) – the concern would be if the prevalence of these motivations were

positively correlated with discount factor heterogeneity. When saving privately impacts 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 our context, 89 percent of couples who saved only saved in one account. If we further limit the sample to couples who opened both individual accounts *and* saved in at least one account, just 28 percent saved in both individual accounts. This suggests that private consumption concerns are not a major driver of the use of the individual accounts in this study.

Furthermore, our results are not well rationalized by mental accounting, which is frequently cited as motivating intrahousehold resource allocation in the developing world. The theory of mental accounting, as described in Thaler (1990), specifies that households earmark funds saved in different accounts for different types of consumption. Duflo and Udry (2004) present evidence of such behavior among households in Côte d'Ivoire. If households have good outside options for individual investment technologies and poor outside options for joint technologies, mental accounting could help explain why joint accounts were so popular in the sample. However, mental accounting models do not easily rationalize the different excess interest rate responses observed among well and poorly matched couples.

Another related 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. An example of a 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 model – 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.

Finally, a broad concern with the analysis is that since we cannot randomly assign time preferences to individuals, we must rely on heterogeneous treatment effects to study the implications of preference heterogeneity. Consequently, we can never completely rule out that the core results are actually driven by unobservable characteristics correlated with match quality. However, the robustness of the results to controls for household information sharing and heterogeneity in observables, including education, income, and distance from the bank (through experimental session fixed effects) suggests that this is not the case.

7 Conclusion

This paper seeks to shed light on the underlying drivers of inefficient intertemporal resource allocation by households. We structure 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. In order to test the model, we conducted a field experiment in Western Kenya, where married couples were given the opportunity to open joint and individual bank accounts bearing randomly assigned interest rates. By feeding the model through the experimental design, we derive three testable implications: (1) as long as the joint interest rate exceeds the interest rate on the less patient spouse's account, individual account use will increase in preference heterogeneity while joint account use will decrease in preference heterogeneity (2) well matched couples invest savings efficiently, whereas poorly matched couples savings decisions may be distorted by strategic behavior (3) well matched couples earn higher rates of return than poorly matched savers, with interest rate losses increasing in preference heterogeneity.

The empirical results support all three of these predictions. This is, of course, subject to the caveat that we cannot randomly assign preference heterogeneity in the field. As such, we cannot completely rule out the hypothesis that the results are driven by some other omitted characteristic that is correlated with estimates of preference heterogeneity. However, the stability of the results to the inclusion of flexible demographic and economic controls, and to the inclusion of a measure of household information flows, is comforting and suggests that the results we observe are indeed driven by inefficiencies arising from conflicting savings motives. We also find some 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 information sharing treatment and are also more likely to gravitate towards individual accounts over joint accounts.

Our results add to a growing body of literature that rejects dynamic household efficiency, while proposing and testing an alternative model of household decision making. A novel feature of the model 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. When credit markets are imperfect, 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, McKenzie, and Woodruff 2009) observed in the developing world. A broader implication of our 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.

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.

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. (1988). Rational Household Labor Supply. *Econometrica 56*(1), pp. 63–90.

Chiappori, P.-A. (1992). Collective Labor Supply and Welfare. *Journal of Political Economy 100*(3), pp. 437–467.

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.

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. 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 (2011). Savings Constraints and Microenterprise Development: Evidence from a Field Experiment in Kenya. Mimeo.

Fafchamps, M., D. McKenzie, S. R. Quinn, and C. Woodruff (2011, July). When is Capital Enough to Get Female Microenterprises Growing? Evidence from a Randomized Experiment in Ghana. Working Paper 17207, National Bureau of Economic Research.

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.

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

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. G. (2011a). Intrahousehold Preference Heterogeneity, Commitment, and Strategic Savings: Theory and Evidence from Kenya. Mimeo, MIT.

Schaner, S. G. (2011b). 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.

Thaler, R. H. (1990). Anomalies: Saving, Fungibility, and Mental Accounts. *The Journal of Economic Perspectives 4*(1), 193–205.

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.

A Proofs

Proof of Proposition 1. To establish the result formally, we work with the indirect utility function that arises from agent i's utility maximization problem (this is just (CB) where only agent i's utility appears in the maximand). Denote this indirect utility function as $v(\kappa, \delta_i)$. The only complication is that due to transaction costs, this function may not be differentiable at some points. It is straightforward to show that when there are transaction costs, it is never optimal to save in more than one account. We can then solve i's utility maximization problem by taking the maximum utility over the solution to four subproblems: (1) consume endowments, (2) optimally save in the joint account conditional on paying joint transaction costs, (3) optimally save in agent A's account conditional on paying transaction costs, (4) optimally save in agent B's account conditional on paying transaction costs. Denote the conditional indirect utility functions arising from these subproblems as $v_a(\kappa, \delta_i)$, $a \in \{J, A, B, E\}$ where E denotes the allocation in which no accounts are used and consumption is equivalent to endowments. Since each subproblem is convex, each $v_a(\kappa, \delta_i)$ will be continuous and differentiable in κ and δ_i. Then the indirect utility function is the upper envelope of the four conditional indirect utility functions that arise from these subproblems: $v(\kappa, \delta_i) = \max_a \{v_a(\kappa, \delta_i)\}$. The upper envelope will be continuous, but not differentiable at places where two or more $v_a(\kappa, \delta_i)$ intersect.

First we note that the result is trivially true when $\delta_i = \Omega$. Since the collective allocation is always available to individual i, it must be that $\kappa^* \geq 0$ for all δ_i. Since $\kappa^* = 0$ when $\delta_i = \Omega$, the result follows.

Next we establish the result for each conditional indirect utility function when $\delta_i \neq \Omega$. Denote the allocation that arises when the couple bargains cooperatively as $\left(c_1^{HH}, c_2^{HH}\right)$. Then the κ_a^* that makes the agent indifferent between the collective outcome and the individual optimum conditional on exercising savings option a is

$$u\left(c_1^{HH}\right) + \delta_i u\left(c_2^{HH}\right) = v_a\left(\kappa_a^*, \delta_i\right) \tag{11}$$

We limit our attention to instances where $\kappa_a^* \geq 0$, since if $\kappa_a^* < 0$ savings option a cannot be on the upper envelope. By the Envelope Theorem $\frac{\partial v_a}{\partial \delta_i} = u\left(c_2^{ai}\right)$, where c_2^{ai} is second period consumption under agent i's most preferred allocation using savings option a. Also note that $\frac{\partial v_a}{\partial \kappa} = -\lambda_1^{ai}$, where λ_1^{ai} is the multiplier on the first period budget constraint on agent i's conditional utility maximization problem. Then we differentiate equation 11 implicitly to obtain:

$$\frac{\partial \kappa_a^*}{\partial \delta_i} = \frac{u\left(c_2^{HH}\right) - u\left(c_2^{ai}\right)}{-\lambda_1^{ai}}$$

The denominator of this equation is negative. We now show that $u\left(c_2^{aB}\right) \geq u\left(c_2^{HH}\right)$ and $u\left(c_2^{aA}\right) \leq u\left(c_2^{HH}\right)$. First consider agent B. To obtain a contradiction suppose instead that $u\left(c_2^{aB}\right) < u\left(c_2^{HH}\right)$. Then for equation 11 to hold it must be that $u\left(c_1^{aB}\right) > u\left(c_1^{HH}\right)$.

$$u\left(c_1^{aB}\right) - u\left(c_1^{HH}\right) = \delta_B\left[u\left(c_2^{HH}\right) - u\left(c_2^{aB}\right)\right] > \Omega\left[u\left(c_2^{HH}\right) - u\left(c_2^{aB}\right)\right]$$

because $\delta_B > \Omega$. This implies that $u\left(c_1^{HH}\right) + \Omega u\left(c_2^{HH}\right) < u\left(c_1^{aB}\right) + \Omega u\left(c_2^{aB}\right)$, which is impossible because $\left(c_1^{HH}, c_2^{HH}\right)$ is collectively optimal and $\kappa_a^* \geq 0$ implies that $\left(c_1^{aB}, c_2^{aB}\right)$ is

a feasible allocation. Then it must be that $u\left(c_2^{aB}\right) \geq u\left(c_2^{HH}\right)$, which implies that $\frac{\partial k_a^*}{\partial \delta_B} \geq 0$. A very similar argument for δ_A establishes that $\frac{\partial k_a^*}{\partial \delta_A} \leq 0$. Since preference heterogeneity is increasing in δ_B and decreasing in δ_A, these derivatives imply that willingness to pay is weakly increasing in preference heterogeneity.

Note that agent i's maximum willingness to pay, $\kappa^* = \max_a \{\kappa_a^*\}$, solves $u\left(c_1^{HH}\right) + \delta_i u\left(c_2^{HH}\right) = v\left(\kappa^*, \delta_i\right)$. Then $\frac{\partial k_a^*}{\partial \delta_B} \geq 0$ and $\frac{\partial k_a^*}{\partial \delta_A} \leq 0$ imply that κ^* must (potentially weakly) increase in δ_B and decrease in δ_A (and therefore increase in $|\delta_i - \Omega|$). Note that when $\delta_i \neq \Omega$ and $c_2^{HH} > y_2^A + y_2^B$ (i.e. it is collectively optimal to save), the individually optimal solution to (CB) will never coincide with the collectively optimal solution to (CB). Then it must be that $u\left(c_2^{HH}\right) \neq u\left(c_2^{ai}\right)$ for any savings option on the upper envelope. Then in this case willingness to pay is strictly increasing in preference heterogeneity. ■

It is straightforward to modify the above proof to conceptualize willingness to pay in terms of the interest rate or transaction costs on an account. For example, we can define R_a^* for $a \in \{J, A, B\}$ such that $u\left(c_1^{HH}\right) + \delta_i u\left(c_2^{HH}\right) = v_a\left(R_a^*, \delta_i\right)$. Then as long as agent i elects to save in account a at the individual optimum (which will be necessary for option a to be on the upper envelope) we have $\frac{\partial R_a^*}{\partial \delta_i} = \frac{u\left(c_2^{HH}\right) - u\left(c_2^{ai}\right)}{\lambda a_2^i s_1^{ai}}$. This implies that the minimum acceptable interest rate on account a is decreasing in preference heterogeneity. We can calculate a similar partial derivative for banking costs, b^a: $\frac{\partial b_a^*}{\partial \delta_i} = \frac{u\left(c_2^{HH}\right) - u\left(c_2^{ai}\right)}{-\lambda_1^{ai} - \lambda_2^{ai}}$ and the withdrawal fee, w: $\frac{\partial w_a^*}{\partial \delta_i} = \frac{u\left(c_2^{HH}\right) - u\left(c_2^{ai}\right)}{-\lambda_2^{ai}}$. In both these cases, the maximum acceptable transaction cost is increasing in preference heterogeneity.

B 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. While 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, they were able to observe whether or not a couple was randomly selected for extra statements in a subset of the sessions. The correlation is only significant in this subset of sessions, so we 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, we 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. We 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, we dropped 366 of 1,000 individual account observations. Columns 1 and 2 of Table 9 verify that account opening is uncorrelated with extra statement selection once we drop the suspect enumerators.

After dropping suspect observations, we created an information sharing index. The index inputs are responses to baseline survey questions addressing income earned least week, bank accounts, savings at home, SACCOs, and ROSCAs, creating five subindices. The subindices range from 0 (perfect information) to 1 (most misinformation). If cross reports exactly matched own reports, we coded the index to 0. If an individual reported that they did (did not) use a device, but the spouse reported that they did not (did), we coded the index 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, we also coded the index to 1. For other instances where we had an own report and a cross report of the amount (or in the case of ROSCAs, the number of ROSCAs), we coded the index to equal $\min\left\{\frac{|own_{ic} - cross_{ic}|}{own_{ic}}, 1\right\}$. We 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.

Table 1. Demographic Characteristics of Study Sample

	Husbands	Wives	Difference	N
Age	41.8	35.2	6.52***	1196
	[13.3]	[11.8]	(0.724)	
Education	8.07	5.92	2.14***	1191
	[3.57]	[4.03]	(0.221)	
Literate	0.865	0.669	0.196***	1196
	[0.342]	[0.471]	(0.024)	
Number Children	4.85	4.49	0.360**	1196
	[2.92]	[2.54]	(0.158)	
Subsistence Farmer or No Job	0.407	0.466	-0.060**	1191
	[0.492]	[0.499]	(0.029)	
Income Last Week	1346	797	549***	1164
	[2747]	[1617]	(132)	
Owns Mobile Phone	0.798	0.686	0.112***	1196
	[0.402]	[0.465]	(0.025)	
Participates in ROSCA	0.492	0.657	-0.166***	1196
	[0.500]	[0.475]	(0.028)	
Has Bank Account	0.303	0.119	0.184***	1196
	[0.460]	[0.324]	(0.023)	
Savings in Bank Account (Among Savers)	10217	4495	5722***	208
	[18695]	[5680]	(1702)	
Saves at Home	0.851	0.898	-0.047***	1195
	[0.356]	[0.303]	(0.019)	
Savings at Home (Among Savers)	1206	858	348**	1025
	[2549]	[2920]	(171)	
Consumption - I Decide	0.471	0.170	0.301***	1189
	[0.500]	[0.376]	(0.026)	
Consumption - Spouse Decides	0.076	0.377	-0.301***	1189
	[0.265]	[0.485]	(0.023)	
Consumption - Decide Together	0.388	0.377	0.011	1189
	[0.488]	[0.485]	(0.028)	
Consumption - Decide Alone	0.015	0.034	-0.019**	1189
	[0.122]	[0.181]	(0.009)	
Savings - I Decide	0.354	0.491	-0.137***	1189
	[0.479]	[0.500]	(0.028)	
Savings - Spouse Decides	0.440	0.282	0.158***	1189
	[0.497]	[0.450]	(0.027)	
Savings - Decide Together	0.106	0.103	0.003	1189
	[0.308]	[0.304]	(0.018)	
Savings - Decide Alone	0.077	0.101	-0.024	1189
	[0.267]	[0.302]	(0.017)	
Weekly Discount Factor	0.718	0.704	0.014	1196
	[0.368]	[0.363]	(0.021)	
Distance from Family Bank (Miles)	3.78	3.78		1196
	[2.21]	[2.21]		

Notes: Standard deviation in brackets, standard errors in parentheses. Variable recoded to missing if response was don't know/refused. ***, **, and * indicate significance at the 99, 95, and 90 percent levels respectively.

Table 2. Correlations Between Estimated Discount Factors and Socioeconomic Characteristics

	All	Well Matched	Poorly Matched
Female	-0.011	-0.029	0.005
	(0.026)	(0.028)	(0.043)
Age	-0.001	-0.001	-0.001
	(0.001)	(0.001)	(0.002)
Education	0.004	0.003	0.003
	(0.004)	(0.004)	(0.007)
Literate	-0.044	-0.013	-0.033
	(0.035)	(0.040)	(0.060)
Number Children	-0.009	-0.007	-0.012
	(0.005)	(0.005)	(0.008)
Subsistence Farmer or No Job	0.012	0.031	0.012
	(0.024)	(0.024)	(0.040)
Income Last Week^{+}	0.304	0.402	0.288
	(0.399)	(0.574)	(0.599)
Owns Phone	0.031	0.024	-0.004
	(0.028)	(0.032)	(0.045)
Consumption - I Decide	0.002	0.050**	-0.037
	(0.024)	(0.025)	(0.036)
Consumption - My Spouse Decides	-0.044	-0.010	-0.100*
	(0.033)	(0.034)	(0.051)
ROSCA Member	0.034	0.049	-0.022
	(0.034)	(0.040)	(0.051)
Has Bank Account	0.110***	0.028	0.179***
	(0.045)	(0.044)	(0.077)
Saves at Home	-0.063**	-0.021	-0.055
	(0.030)	(0.032)	(0.048)
Has SACCO Account	0.035	-0.026	0.070*
	(0.025)	(0.026)	(0.041)
Has MPESA Account	0.045*	0.016	0.100**
	(0.025)	(0.030)	(0.044)
Saves Other Ways	-0.007	0.022	-0.005
	(0.032)	(0.033)	(0.051)
Savings - I Decide	0.060*	0.092***	-0.017
	(0.032)	(0.035)	(0.048)
Savings - My Spouse Decides	0.074**	0.086***	0.039
	(0.033)	(0.036)	(0.049)
Self Reported Savings^{+}	0.012	0.015	-0.048
	(0.022)	(0.023)	(0.092)
Savings in Family Bank Accounts^{+}	1.22*	2.09***	1.71
	(0.716)	(0.715)	(1.49)
N	1196	598	598

Notes: $^{+}$Coefficients and standard errors have been scaled by 100,000 for readability. Robust standard errors clustered at the couple leve reported in parentheses. All regressions includeexperimental session fixed effects. ***, **, and * indicate significance at the 99, 95, and 90 percent levels respectively.

Table 3. Randomization Verification - Interest Rates and Extra Statements

	Individual Accounts		Joint	Extra	Cash
	Husband	Wife	Accounts	Statements	Prize
Panel A. Adherence to Theoretical Probabilities					
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. Correlation with Baseline Characteristics and Other Treatments					
Age	{0.700}	{0.542}	{0.753}	{0.202}	{0.606}
Education	{0.037}**	{0.812}	{0.136}	{0.727}	{0.465}
Literate	{0.305}	{0.621}	{0.047}**	{0.679}	{0.520}
Number Children	{0.148}	{0.778}	{0.099}*	{0.275}	{0.870}
Subsistence Farmer/Unemployed	{0.193}	{0.950}	{0.942}	{0.778}	{0.219}
Income Last Week	{0.413}	{0.675}	{0.606}	{0.914}	{0.035}**
Has Mobile Phone	{0.201}	{0.160}	{0.174}	{0.233}	{0.723}
Participates in ROSCA	{0.091}*	{0.661}	{0.217}	{0.991}	{0.895}
Has Bank Account	{0.594}	{0.397}	{0.490}	{0.292}	{0.183}
Saves at Home	{0.813}	{0.106}	{0.189}	{0.522}	{0.018}**
Has SACCO Account	{0.690}	{0.195}	{0.684}	{0.320}	{0.811}
Has MPESA Account	{0.957}	{0.377}	{0.266}	{0.977}	{0.563}
Saves Other Ways	{0.187}	{0.503}	{0.071}*	{0.049}**	{0.149}
Consumption Decision Making	{0.621}	{0.091}*	{0.385}	{0.156}	{0.945}
Saving Decision Making	{0.450}	{0.919}	{0.328}	{0.742}	{0.599}
Well Matched Couple	{0.820}	{0.340}	{0.934}	{0.173}	{0.019}**
Weekly Discount Factor	{0.468}	{0.956}	{0.130}	{0.001}***	{0.003}***
Distance from Family Bank (Miles)	{0.528}	{0.245}	{0.511}	{0.809}	{0.682}
Interest Rate - Husband's Account		{0.635}	{0.445}	{0.973}	{0.203}
Interest Rate - Wife's Account	{0.636}		{0.152}	{0.631}	{0.470}
Interest Rate - Joint Account	{0.445}	{0.152}		{0.242}	{0.057}*
Extra Statements	{0.869}	{0.597}	{0.254}		{0.658}
Cash Prize	{0.203}	{0.471}	{0.058}*	{0.704}	
Joint Test: All Baseline Characteristics	{0.221}	{0.258}	{0.159}	{0.168}	{0.099}*

Notes: P-values from binomial tests of theoretical probabilities (Panel A) or F-tests for the significants of treatments (Panel B) in braces. 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. Standard errors in Panel B clustered at the couple level. ***, **, and * indicate significance at the 99, 95, and 90 percent levels respectively.

Table 4. Impact of Interest Rates on Account Takeup and Use

	Opened	Saved	Average Balance	Total Deposits	Total Withdrawals	Number Transactions	Fees
2% Interest	0.066*	0.018	-8.63	6.29	-58.1	0.074	0.425
	(0.036)	(0.020)	(32.9)	(222)	(154)	(0.141)	(4.02)
6% Interest	0.152***	0.059***	90.0	536	439	0.177	5.61
	(0.039)	(0.022)	(58.7)	(422)	(346)	(0.138)	(4.75)
10% Interest	0.197***	0.071***	107**	689*	555	0.301*	8.04
	(0.040)	(0.023)	(50.2)	(398)	(372)	(0.163)	(5.18)
Joint	0.256***	0.150***	166***	3156	2967	0.576**	15.5*
	(0.051)	(0.033)	(66.6)	(2649)	(2638)	(0.249)	(8.22)
Joint×6% Interest	-0.015	0.030	-118	-3271	-3125	-0.331	-13.9
	(0.061)	(0.046)	(100)	(2683)	(2661)	(0.298)	(10.0)
Joint×10% Interest	0.037	0.066	-16.4	-2259	-2272	0.277	-0.889
	(0.059)	(0.049)	(118)	(2784)	(2742)	(0.449)	(13.4)
DV Mean (Individual Accounts, No Interest)	0.288	0.054	49.9	302	211	0.244	5.16
N	1794	1794	1794	1794	1794	1794	1794

Notes: Robust standard errors clustered at the couple level in parentheses. All regressions include separate dummies for husband and wife cash prize selection. ***, **, and * indicate significance at the 99, 95, and 90 percent levels respectively.

Table 5. Preference Heterogeneity and Savings Levels, By Account Type

	Individual Accounts			Joint Accounts		
	Saved	Fraction Balance	Average Balance	Saved	Fraction Balance	Average Balance
Panel A. Basic Controls						
Poorly Matched	0.108***	0.200***	18.3	-0.074*	-0.200***	-56.5
	(0.032)	(0.070)	(104)	(0.041)	(0.070)	(80.9)
Joint Deviation	-0.127*	-0.529***	-130	0.267***	0.529***	483***
	(0.073)	(0.192)	(146)	(0.080)	(0.192)	(207)
Panel B. + Time Preference Controls						
Poorly Matched	0.150***	0.286***	80.0	-0.062	-0.286***	105
	(0.035)	(0.078)	(97.4)	(0.050)	(0.078)	(84.2)
Joint Deviation	-0.138*	-0.604***	-192	0.258***	0.604***	433**
	(0.074)	(0.194)	(151)	(0.080)	(0.194)	(207)
Panel C. + Demographic Controls						
Poorly Matched	0.160***	0.315***	69.8	-0.098*	-0.315***	70.1
	(0.038)	(0.102)	(79.9)	(0.053)	(0.102)	(106)
Joint Deviation	-0.128*	-0.417**	-123	0.303***	0.417**	481**
	(0.073)	(0.202)	(146)	(0.086)	(0.202)	(220)
Panel D. +Economic Controls						
Poorly Matched	0.177***	0.322***	113	-0.104*	-0.322***	73.1
	(0.039)	(0.107)	(80.4)	(0.054)	(0.107)	(104)
Joint Deviation	-0.120*	-0.312	-135	0.301***	0.312	498**
	(0.073)	(0.222)	(150)	(0.088)	(0.222)	(224)
DV Mean (Well Matched)	0.124	0.290	196	0.308	0.710	221
N	598	257	598	598	257	598

Notes: Heteroskedasticity robust standard errors in parentheses. Time preference controls include separate dummies for upper/lower censoring of the discount factors of each spouse and the average household discount factor. The demographic control set adds controls for spousal heterogeneity in age, education, number of children, literacy, and experimental session fixed effects. The economic control set adds controls for heterogeneity in income, an indicator for subsistence farmers or the unemployed, and phone ownership. All regressions include fixed effects that fully saturate the set of experimental interest rates, dummies for husband and wife cash prize selection, and a dummy identifying households where the less patient spouse's account has a higher interest rate than the joint account. ***, **, and * indicate significance at the 99, 95, and 90 percent levels respectively.

Table 6. Responses to the Excess Interest Rate by Match Quality

Panel A. Well Matched Couples Only				
Excess Low×Indiv	0.000	0.000	-0.004	-0.007
	(0.005)	(0.005)	(0.005)	(0.006)
Excess High×Indiv	0.019***	0.020***	0.022***	0.026***
	(0.008)	(0.008)	(0.008)	(0.009)
Excess Low×Joint	0.025**	0.023**	0.024**	0.023**
	(0.011)	(0.011)	(0.011)	(0.012)
Excess High×Joint	0.008	0.008	0.007	-0.010
	(0.011)	(0.011)	(0.010)	(0.013)
DV Mean	0.155	0.155	0.155	0.155
N	897	897	897	897
Panel B. Poorly Matched Couples Only				
Excess Low×Indiv	-0.002	0.002	0.007	0.005
	(0.007)	(0.011)	(0.011)	(0.012)
Excess High×Indiv	-0.004	-0.019**	-0.009	-0.009
	(0.008)	(0.010)	(0.012)	(0.012)
Excess Low×Joint	0.000	-0.015	-0.019	-0.011
	(0.010)	(0.016)	(0.022)	(0.024)
Excess High×Joint	-0.013	-0.060***	-0.072***	-0.094***
	(0.011)	(0.017)	(0.026)	(0.027)
DV Mean	0.167	0.167	0.167	0.167
N	897	897	897	897
Panel C. Tests of Equality Across Equations				
Excess Low×Indiv	{0.826}	{0.885}	{0.373}	{0.351}
Excess High×Indiv	{0.043}**	{0.001}***	{0.033}**	{0.019}**
Excess Low×Joint	{0.097}*	{0.051}*	{0.080}*	{0.207}
Excess High×Joint	{0.185}	{0.001}***	{0.006}***	{0.006}***
Joint Test - All Splines	{0.030}**	{0.000}***	{0.001}***	{0.002}***
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. All regressions also include dummies for husband and wife cash prize selection and their interactions with the excess interest rate controls. Time preference controls include separate dummies for upper/lower censoring of the discount factors of each spouse, the average household discount factor, and interactions with the excess interest rate controls. The demographic control set adds controls for experimental session fixed effects, spousal heterogeneity in age, education, number of children, literacy, and interactions with the excess interest controls. The economic control set adds controls for heterogeneity in income, an indicator for subsistence farmers or the unemployed, phone ownership, and interactions with excess interest rate controls. All controls are demeaned to the value among well matched couples. Panel C uses seemimgly unrelated regression to test for equality of excess interest rate responses by match quality. ***, **, and * indicate significance at the 99, 95, and 90 percent levels respectively.

Table 7. Proxied Banking Costs and Savings Behavior

	Couple Level				Account Level
	Saved Any	Saved Joint	Saved Indiv.	Fraction Joint	Saved
Panel A. Basic Controls					
Cost Index/High Cost	-0.380***	0.611***	-0.900***	0.797***	-0.159***
	(0.113)	(0.179)	(0.168)	(0.163)	(0.062)
Excess					0.015***
					(0.005)
Excess×High Cost					-0.012*
					(0.007)
Panel B. + Time Preference Controls					
Cost Index/High Cost	-0.383***	0.621***	-0.919***	0.815***	-0.166***
	(0.114)	(0.181)	(0.172)	(0.166)	(0.064)
Excess					0.016***
					(0.005)
Excess×High Cost					-0.013*
					(0.007)
Panel C. + Demographic Controls					
Cost Index/High Cost	-0.287***	0.589***	-0.959***	0.817***	-0.184***
	(0.118)	(0.217)	(0.203)	(0.198)	(0.066)
Excess					0.018***
					(0.005)
Excess×High Cost					-0.015**
					(0.007)
Panel D. + Economic Controls					
Cost Index/High Cost	-0.330***	0.388	-0.837***	0.642***	-0.204***
	(0.134)	(0.248)	(0.252)	(0.237)	(0.067)
Excess					0.019***
					(0.005)
Excess×High Cost					-0.017**
					(0.007)
DV Mean	0.430	0.677	0.358	0.654	0.155
N	598	257	257	257	897
Sample	All	Savers	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 prize selection. Time preference controls include separate dummies for upper/lower censoring of the discount factor for each member of the couple and the average household discount factor. 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 two columns include dummy variables that fully saturate interest rate×joint, as well as a high cost×joint dummy. The fourth column also includes interactions between the high cost dummy and the interest rate×joint interaction set. ***, **, and * indicate significance at the 99, 95, and 90 percent levels respectively.

Table 8. 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	6.75	6.17	6.04	6.45	6.01	6.01
Actual Interest Earnings	7.10	5.68	4.75	4.06	5.28	4.17	3.29
Loss	1.05	1.07	1.42	1.98	1.18	1.85	2.72
Well Matched Couples							
Maximum Interest Earnings	8.17	6.75	6.13	5.97	6.41	5.93	5.93
Actual Interest Earnings	7.50	6.26	5.49	5.03	5.85	4.92	4.31
Loss	0.664	0.486	0.641	0.937	0.566	1.02	1.63
Loss Gap							
A. No Controls	0.388**	0.587***	0.781***	1.04***	0.612***	0.830***	1.10***
	(0.172)	(0.184)	(0.250)	(0.334)	(0.205)	(0.316)	(0.447)
B. + Interest Rate/Cash Prize Controls	0.540***	0.755***	1.02***	1.34***	0.794***	1.09***	1.45***
	(0.157)	(0.180)	(0.247)	(0.330)	(0.199)	(0.310)	(0.441)
C. + Time Preference Controls	0.510***	0.837***	1.29***	1.75***	0.963***	1.49***	2.02***
	(0.209)	(0.202)	(0.264)	(0.356)	(0.215)	(0.333)	(0.478)
D. + Demographic Controls	0.490***	0.848***	1.31***	1.82***	0.969***	1.52***	2.12***
	(0.207)	(0.216)	(0.291)	(0.393)	(0.236)	(0.371)	(0.530)
E. + Economic Controls	0.490**	0.846***	1.31***	1.84***	0.981***	1.59***	2.24***
	(0.212)	(0.222)	(0.299)	(0.402)	(0.242)	(0.376)	(0.533)
N	598	598	598	598	598	598	598

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 prize selection. Time preference controls include separate dummies for upper/lower censoring of the discount factors of each spouse, and the average household discount factor. 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 levels respectively. Principal components index constructed from distance from the bank, spouse specific indicators for subsistence farmers/the unemployed, and spouse specific indicators for baseline bank account ownership and SACCO membership.

Table 9. 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***	0.046	0.197*	230	529	0.080	0.289*	397	805
	(0.052)	(0.085)	(0.067)	(0.109)	(0.057)	(0.106)	(180)	(451)	(0.099)	(0.173)	(312)	(700)
x Poorly Matched		-0.043		-0.129		-0.159		-362		-0.261		-682
		(0.104)		(0.137)		(0.113)		(335)		(0.242)		(674)
x Poorly Informed		-0.056		-0.280*		-0.173		-155		-0.307		-20.0
		(0.115)		(0.145)		(0.126)		(415)		(0.256)		(772)
Panel B. + Time Preference Controls												
Extra Statement	-0.028	0.039	0.562***	0.765***	0.062	0.205**	282	558	0.111	0.275*	502	780
	(0.053)	(0.087)	(0.066)	(0.111)	(0.054)	(0.102)	(182)	(462)	(0.098)	(0.164)	(326)	(706)
x Poorly Matched		-0.037		-0.159		-0.101		-127		-0.139		-158
		(0.106)		(0.134)		(0.116)		(334)		(0.252)		(723)
x Poorly Informed		-0.059		-0.268*		-0.218*		-319		-0.366		-286
		(0.116)		(0.146)		(0.132)		(442)		(0.254)		(771)
Panel C. + Demographic Controls												
Extra Statement	-0.014	0.017	0.594***	0.812***	0.054	0.198**	169	1044***	0.091	0.225	285	1452**
	(0.054)	(0.095)	(0.063)	(0.095)	(0.050)	(0.097)	(122)	(447)	(0.087)	(0.137)	(209)	(637)
x Poorly Matched		-0.022		-0.223*		-0.174		-754**		-0.334		-1664*
		(0.110)		(0.125)		(0.118)		(378)		(0.231)		(870)
x Poorly Informed		-0.007		-0.169		-0.289**		-965*		-0.452**		-1398*
		(0.133)		(0.137)		(0.140)		(536)		(0.216)		(797)
Panel D. + Economic Controls												
Extra Statement	-0.009	0.024	0.601***	0.747***	0.010	0.126	129	1063***	0.017	0.143	214	1638***
	(0.055)	(0.097)	(0.060)	(0.110)	(0.058)	(0.102)	(154)	(439)	(0.096)	(0.165)	(256)	(683)
x Poorly Matched		-0.018		-0.157		-0.183		-866***		-0.475*		-2165***
		(0.112)		(0.141)		(0.122)		(304)		(0.255)		(912)
x Poorly Informed		0.019		-0.115		-0.351**		-894*		-0.611***		-1558*
		(0.136)		(0.153)		(0.151)		(521)		(0.233)		(904)
DV Mean (ES=0)	0.395		0		0.266		241		0.266		241	
N	634		244		244		244		244		244	

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 prize receipt. Time preference controls include separate dummies for upper/lower censoring of the discount factors of each spouse and the average household discount factor. The demographic control set adds controls for spousal heterogeneity in age, education, number of children, literacy, and experimental session fixed effects. 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 levels respectively.

Table 10. Preference Heterogeneity, Information Sharing, and Savings Levels, By Account Type

	Individual Accounts			Joint Accounts		
	Saved	Fraction Bal.	Average Balance	Saved	Fraction Bal.	Average Balance
Panel A. Basic Controls						
Poorly Matched	0.113***	0.180***	24.0	-0.075*	-0.180***	-55.5
	(0.032)	(0.067)	(107)	(0.041)	(0.067)	(81.9)
Joint Deviation	-0.129*	-0.486***	-137	0.274***	0.486***	481**
	(0.072)	(0.192)	(144)	(0.079)	(0.192)	(208)
Poorly Informed	0.140***	0.266***	180*	-0.069*	-0.266***	33.3
	(0.032)	(0.068)	(93.5)	(0.038)	(0.068)	(102)
Panel B. + Time Preference Controls						
Poorly Matched	0.153***	0.254***	82.5	-0.061	-0.254***	107
	(0.035)	(0.074)	(99.1)	(0.050)	(0.074)	(85.8)
Joint Deviation	-0.141*	-0.577***	-200	0.265***	0.577***	434**
	(0.073)	(0.191)	(150)	(0.080)	(0.191)	(207)
Poorly Informed	0.140***	0.284***	182*	-0.070*	-0.284***	36.9
	(0.032)	(0.067)	(95.0)	(0.039)	(0.067)	(102)
Panel C. + Demographic Controls						
Poorly Matched	0.165***	0.280***	74.3	-0.101*	-0.280***	66.7
	(0.038)	(0.102)	(81.6)	(0.053)	(0.102)	(108)
Joint Deviation	-0.134*	-0.422**	-137	0.308***	0.422**	479**
	(0.072)	(0.202)	(145)	(0.085)	(0.202)	(220)
Poorly Informed	0.107***	0.272***	140	-0.078*	-0.272***	-41.9
	(0.034)	(0.083)	(109)	(0.043)	(0.083)	(87.1)
Panel D. + Economic Controls						
Poorly Matched	0.180***	0.295***	111	-0.106**	-0.295***	67.9
	(0.038)	(0.108)	(80.6)	(0.054)	(0.108)	(105)
Joint Deviation	-0.124*	-0.303	-148	0.305***	0.303	498**
	(0.071)	(0.213)	(149)	(0.087)	(0.213)	(223)
Poorly Informed	0.089***	0.209**	69.0	-0.068	-0.209**	-91.0
	(0.035)	(0.094)	(124)	(0.044)	(0.094)	(86.7)
DV Mean (Omitted Matched)	0.124	0.290	196	0.308	0.710	221
N	598	257	598	598	257	598

Notes: Heteroskedasticity robust standard errors in parentheses. "Omitted Matched" refers to couples who are poorly matched and have no joint deviation. Time preference controls include separate dummies for upper/lower censoring of the discount factors of each spouse and the average household discount factor. The demographic control set adds controls for spousal heterogeneity in age, education, number of children, literacy, and experimental session fixed effects. The economic control set adds controls for heterogeneity in income, an indicator for subsistence farmers or the unemployed, and phone ownership. All regressions include fixed effects that fully saturate the set of experimental interest rates, dummies for husband and wife cash prize selection, and a dummy identifying households where the less patient spouse's account has a higher interest rate than the joint account. ***, **, and * indicate significance at the 99, 95, and 90 percent levels respectively.

Figure 1. 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 on the joint account. The excess interest on the husband's and wife's account follow respectively.

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

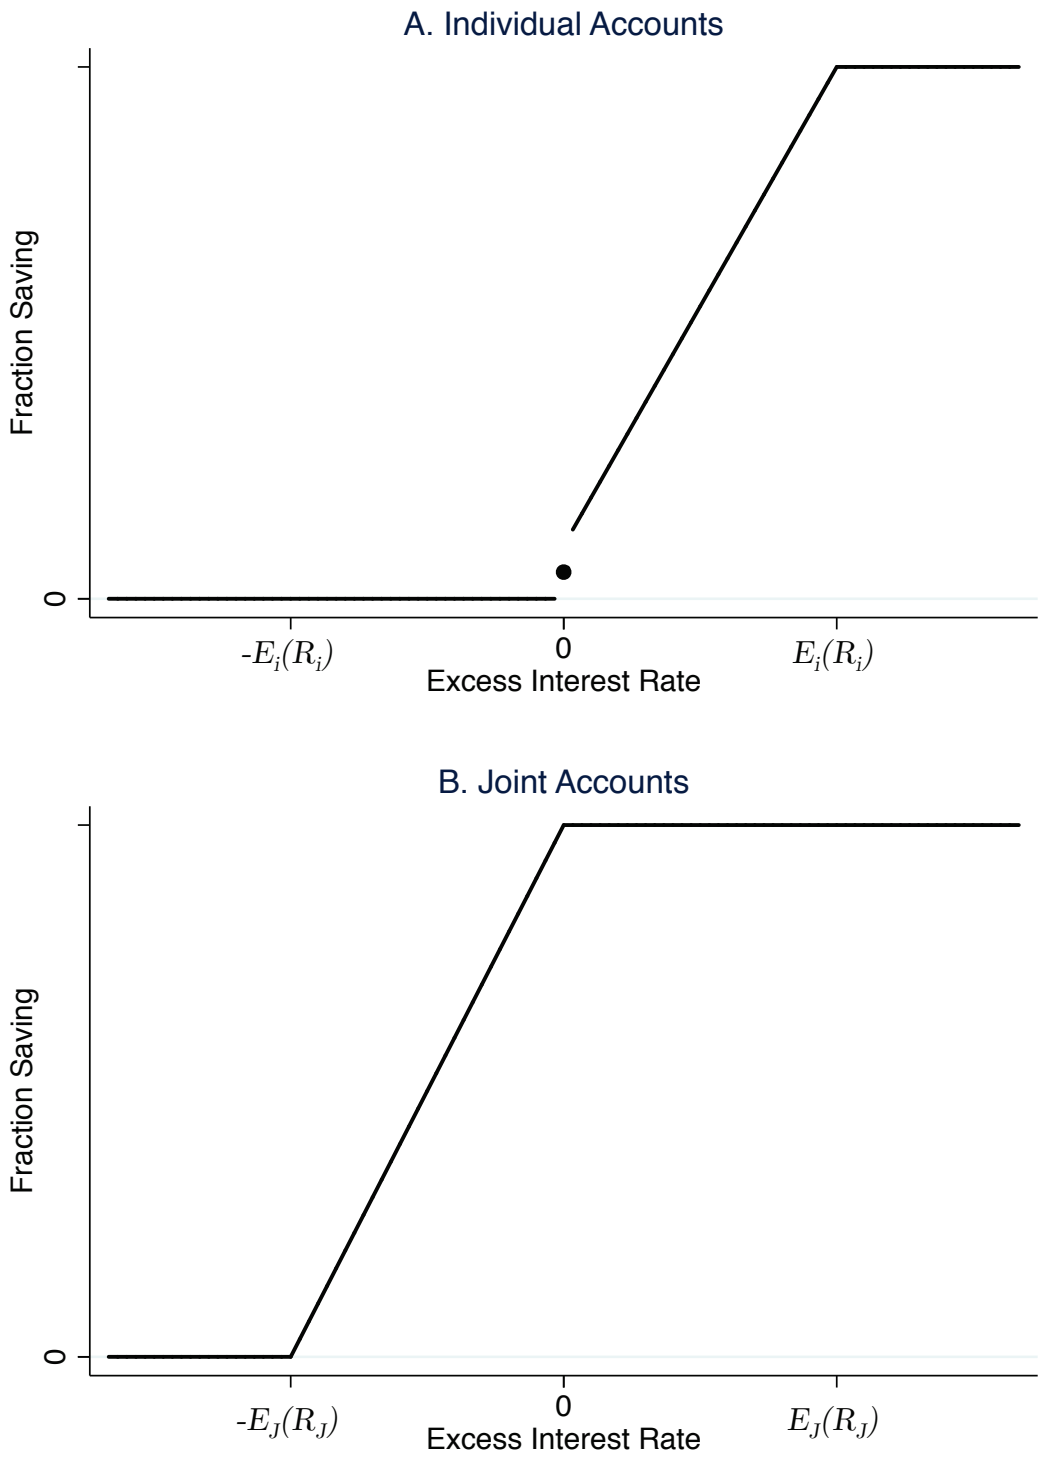

Figure 3. Distributions of Estimated Discount Factors and Discount Factor Heterogeneity

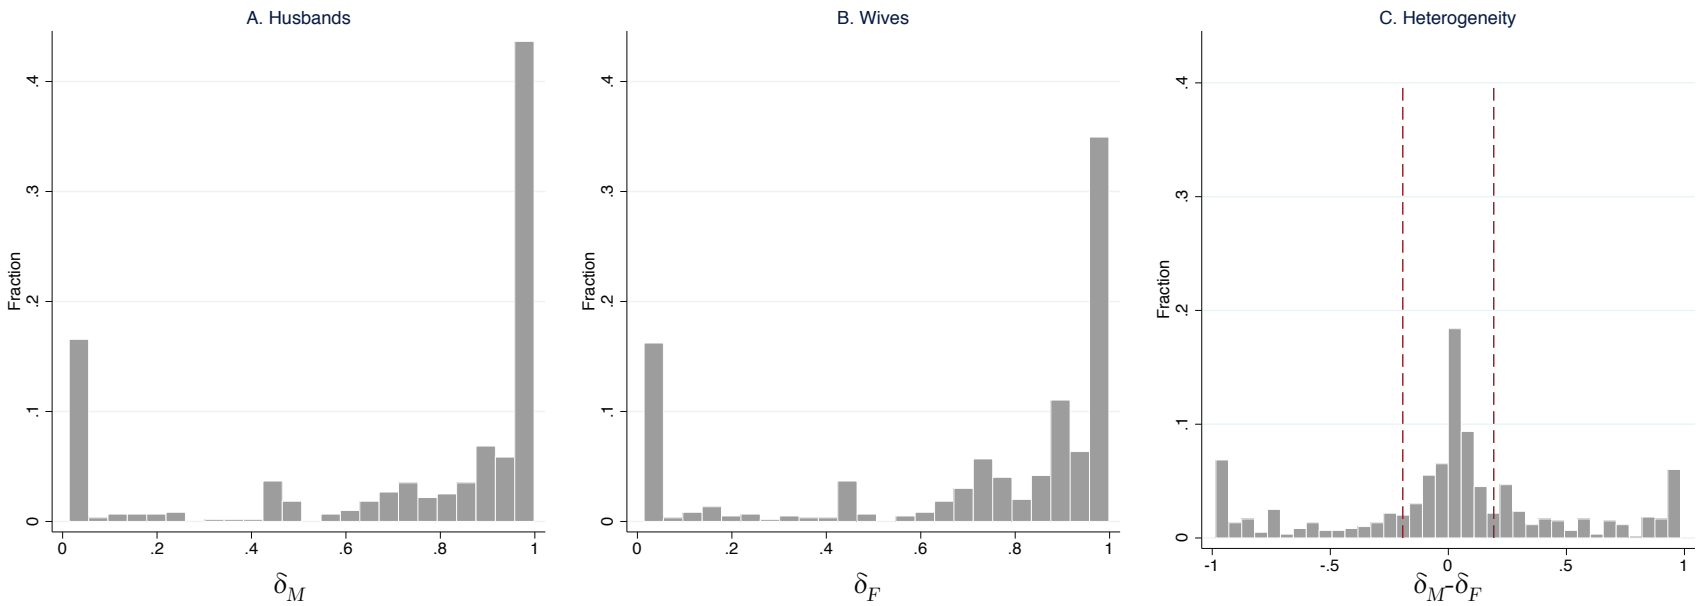

Note: Well matched couples delineated by dashed vertical lines.

Figure 4. Relationship Between Savings Balances and Heterogeneity in Estimated Discount Factors

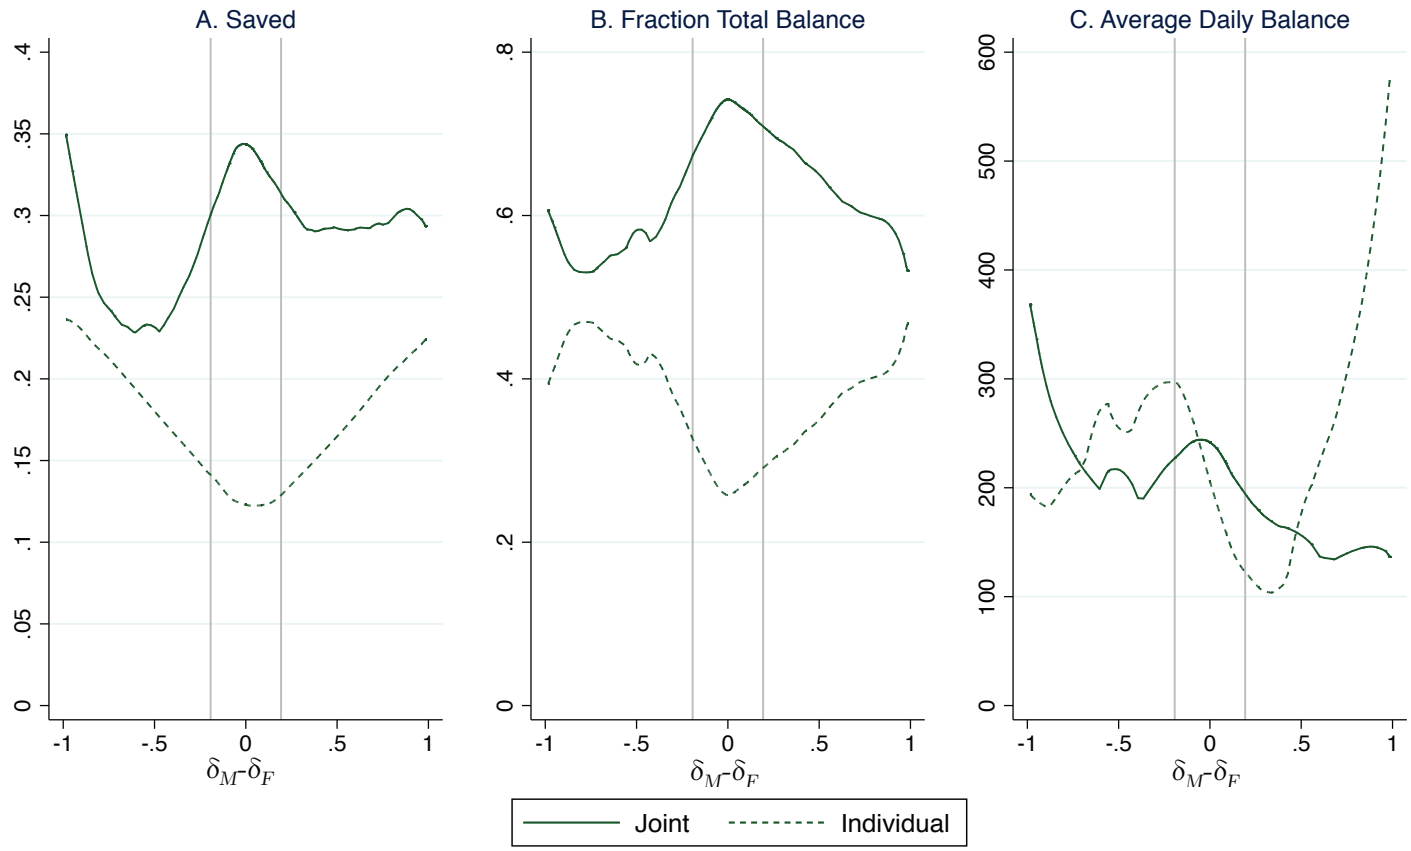

Note: Well matched couples delineated by vertical lines.

Figure 5. Savings Response to Excess Interest Rate by Match Quality and Account Type

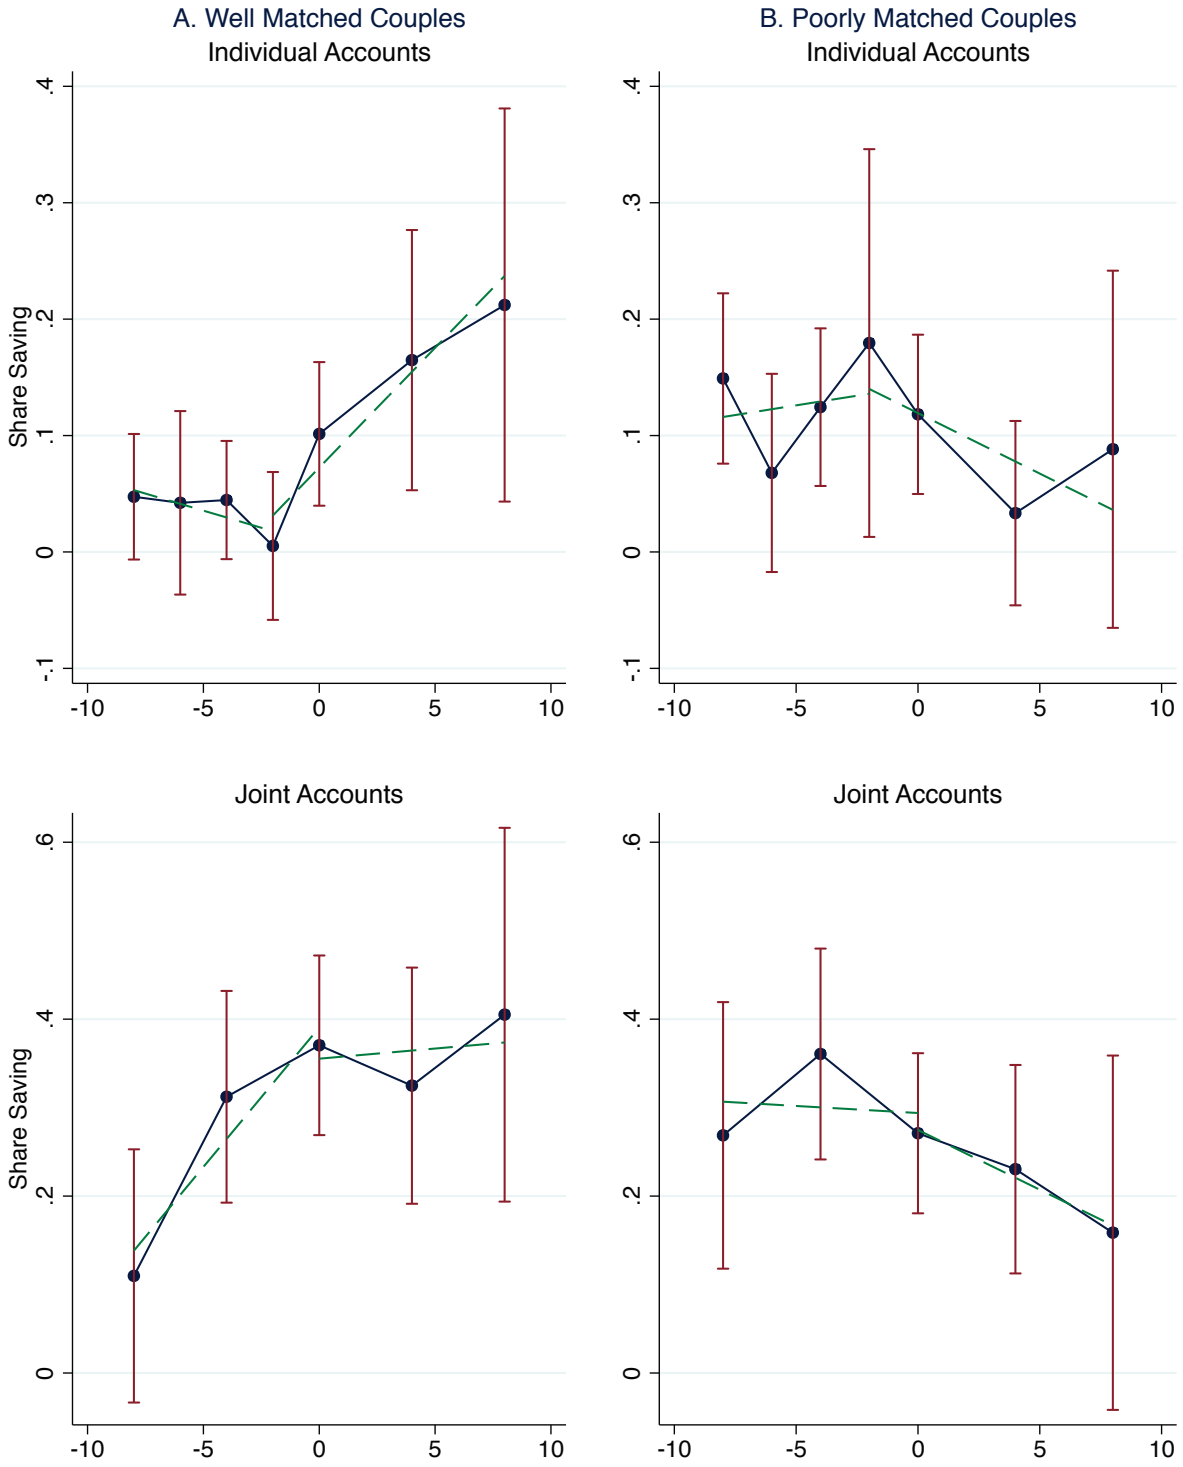

Note: Predicted values calculated from regressions with standard errors clustered at the couple level. Whiskers indicate 95 percent confidence intervals for each predicted value.

Figure 6. Interest Rate Losses by Match Quality and Individual Interest Discounting

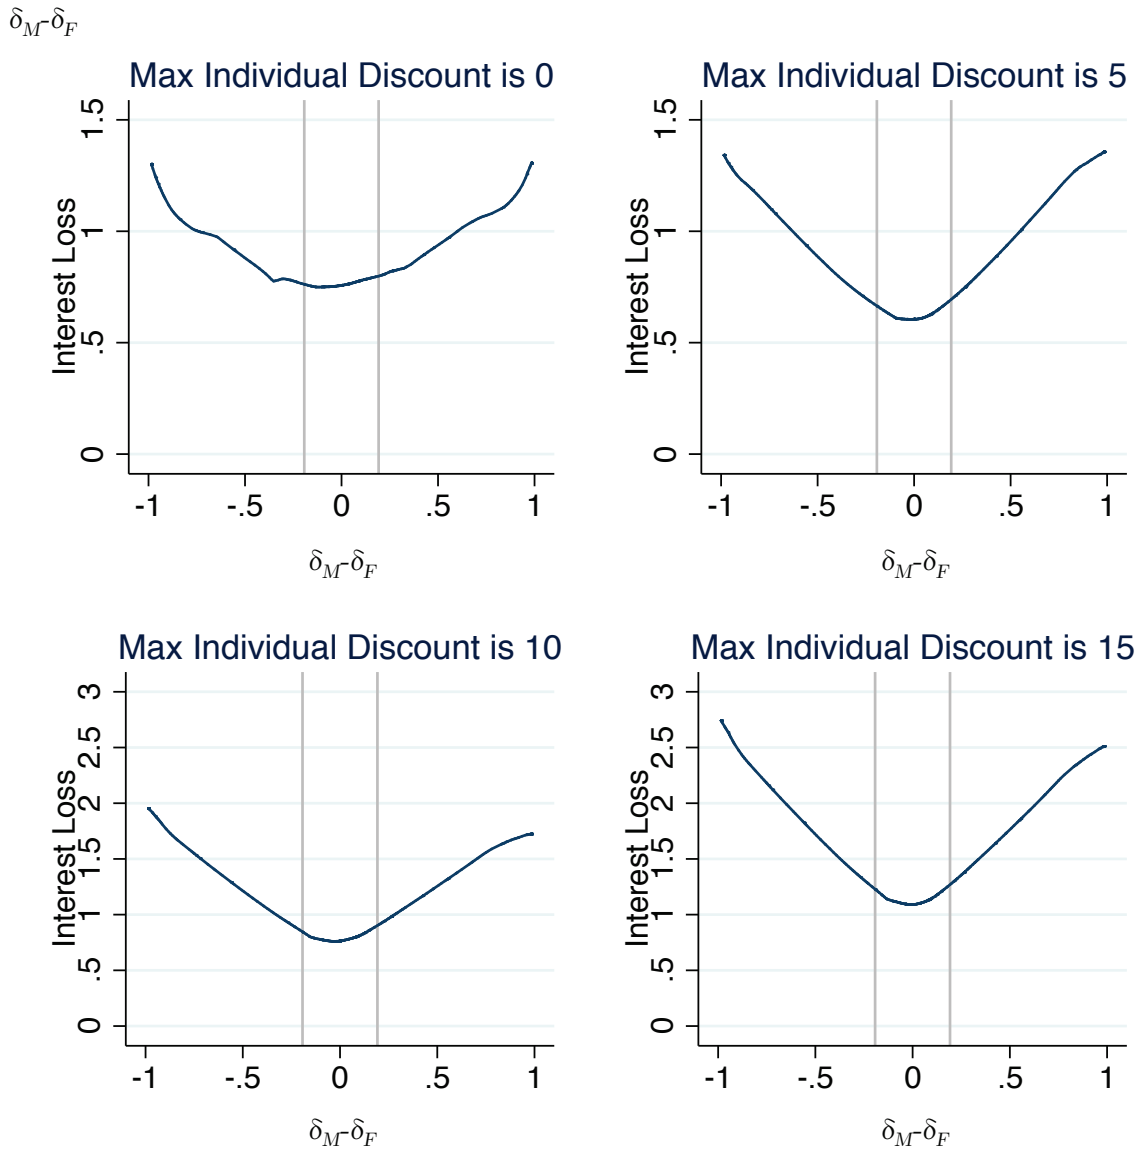

Notes: Gray lines demarcate well matched couples. Individual interest rate discounting performed using proxied banking costs.

Table A1. 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 or 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	0.769	0.714	0.055**	1196
	[0.422]	[0.452]	(0.025)	
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)	
Savings in Bank Account	9674	7283	2392	208
	[19485]	[11257]	(2163)	
Saves at Home	0.889	0.860	0.030	1195
	[0.314]	[0.348]	(0.019)	
Savings at Home	1064	988	75.4	1025
	[3151]	[2272]	(171)	
Consumption - I Decide	0.323	0.318	0.006	1189
	[0.468]	[0.466]	(0.027)	
Consumption - Spouse Decides	0.204	0.249	-0.045*	1189
	[0.403]	[0.433]	(0.024)	
Consumption - Decide Together	0.389	0.376	0.012	1189
	[0.488]	[0.485]	(0.028)	
Consumption - Decide Alone	0.029	0.020	0.008	1189
	[0.167]	[0.141]	(0.009)	
Savings - I Decide	0.451	0.393	0.058**	1189
	[0.498]	[0.489]	(0.029)	
Savings - Spouse Decides	0.344	0.378	-0.034	1189
	[0.475]	[0.485]	(0.028)	
Savings - Decide Together	0.082	0.126	-0.044***	1189
	[0.275]	[0.333]	(0.018)	
Savings - Decide Alone	0.096	0.083	0.013	1189
	[0.294]	[0.276]	(0.017)	
Weekly Discount Factor	0.844	0.579	0.265***	1196
	[0.274]	[0.396]	(0.020)	
Distance from Family Bank (Miles)	3.97	3.59	0.378***	1196
	[2.16]	[2.24]	(0.127)	

Notes: Standard deviation in brackets, standard errors in parentheses. Variable recoded to missing if response was don't know/refused. ***, **, and * indicate significance at the 99, 95, and 90 percent levels respectively.