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

Simone Schaner†

March 24, 2014

Abstract

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

JEL Codes: C93, D13, D14, O12

Keywords: intrahousehold allocation, saving, preference heterogeneity

*I thank Abhijit Banerjee, Esther Duflo, and Tavneet Suri for invaluable advice and feedback at all stages of this project. I am also grateful to Pascaline Dupas, Eric Edmonds, Dan Keniston, Erzo Luttmer, Ben Olken, Josh Schwartzstein, Rob Townsend, Jonathan Zinman, and numerous seminar participants for many useful comments. This project would not have been possible without the tireless assistance, hard work, and commitment of many employees of Family Bank. I am particularly indebted to Victor Keriri Mwangi, Steve Mararo, and Michael Aswani Were. I also thank Noreen Makana for her superb field management, Moses Baraza for his expert advice, and the IPA enumerators for their excellent assistance with the data collection. I gratefully acknowledge the financial support of the Russell Sage Foundation, the George and Obie Shultz Fund, MIT's Jameel Poverty Action Lab, and the National Science Foundation's Graduate Research Fellowship. The study protocol in this paper was approved by IRBs at MIT, the Kenya Medical Research Institute, and IPA-Kenya. All errors are my own.

†simone.schaner@dartmouth.edu

1 Introduction

Households across the world have members with differing preferences and priorities – consequently, individuals must find a way to aggregate these preferences when making collective decisions. Moreover, it is clear that individuals actively advocate for their own preferences during the decision making process, since shifts in individual bargaining power translate into shifts in allocations (Lundberg et al. 1997; Chiappori et al. 2002; Duflo 2003; Bobonis 2009). These observations give rise to a natural set of questions: does preference heterogeneity ever lead individuals to take costly action to manipulate intrahousehold resource allocations? If so, how big are the resulting distortions? These questions may be especially relevant for understanding households' intertemporal choices. In their review of the literature on intra-household decision making, Chiappori and Donni (2009) note that studies focusing on static choices (e.g. what to consume) generally fail to reject the null of efficient households, while studies focusing on dynamic choices (e.g. savings and mutual insurance) usually reject efficiency.[1] Indeed, achieving intertemporal efficiency is likely more difficult, since this requires that couples not only reach an agreement on the Pareto frontier, but that they stick to this agreement over multiple time periods.

In spite of these observations, the causes of intertemporal inefficiency remain poorly understood. To shed light on this issue, this paper uses a field experiment to test whether heterogeneity in *intertemporal preferences* is associated with inefficiency in a fundamental intertemporal choice: savings behavior. The motivating hypothesis is that when spouses do not agree about the time path of consumption, they may take costly strategic action to manipulate outcomes in their favor. This is easy to see with a simple example: imagine the savings problem of a patient wife paired with an impatient husband. The wife can save in either a joint bank account that she shares with her spouse, or an individual account. She knows that if she deposits funds into the joint account, her husband, who only cares about the present, will simply withdraw all her savings and spend the funds on current consumption. Alternatively, if she saves in her individual account (which her husband cannot access), her savings will be preserved for the next period. In this context, the wife may choose to save individually even if the joint account offers a much higher rate of return.

I formalize this idea using a simple model of non-cooperative household savings behavior that incorporates heterogeneity in time preferences. I then use a field experiment to evaluate the model empirically. This experiment, which I implemented in Western Kenya in the Summer of 2009, was specifically designed to mirror the conditions of the model: I gave 544 married couples the opportunity to open three savings accounts (two individual, one joint) bearing randomly assigned

[1] A partial selection of papers finding evidence of efficient intrahousehold resource allocation include Browning and Chiappori (1998), Chiappori et al. (2002), Rangel and Thomas (2005), and Bobonis (2009). Studies finding evidence of inefficiency include Udry (1996), Duflo and Udry (2004), Mazzocco (2007), de Mel et al. (2009), and Robinson (2011).

interest rates. These interest rates were substantially higher than those available on the market at the time. A central 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. I also asked each respondent in the experiment a battery of questions to directly elicit estimates of discount factors, which I use to calculate measures of intrahousehold heterogeneity. In practice, I find a substantial amount of heterogeneity in elicited discount factors in my study sample – the median couple's estimated discount factors are nearly 0.5 standard deviations apart.

I use the model to generate three main testable predictions, which I map to the experimental data. The first prediction characterizes patterns of joint versus individual account use. Here, the theory allows me to identify cases where inefficient individual account use should increase with discount factor heterogeneity. Indeed, in these cases I find that couples who are "well matched" in terms of discount factors (i.e. above-median in match quality) are much less likely to use dominated individual accounts. While just 8 percent of well-matched couples save in inefficient individual accounts, 19-23 percent of poorly matched couples do the same. This does not simply reflect a relative preference for joint accounts among well matched couples. I confirm this with the second testable prediction, which states that well matched couples should respond to relative rates of return on *both* individual *and* joint accounts in a manner consistent with efficient investment. This prediction is borne out in the data and provides a striking contrast to the behavior of poorly matched couples, who appear to be completely insensitive to relative rates of return.

These differences in behavior have financial consequences for poorly matched couples – I exploit the design of the field experiment to calculate forgone interest earnings by match quality and evaluate the final testable prediction – that interest rate losses should increase monotonically with preference heterogeneity. Consistent with the theory, I find that poorly matched couples leave at least 52 percent more interest on the table. While this final prediction is intuitive, the model does deliver several unexpected insights. For example, since both individual and joint accounts may be used strategically, the testable prediction with respect individual account use only holds when the joint account bears the highest rate of return. Moreover, since both the more- and less-patient spouse may save strategically, there is no clear testable prediction with respect to average balances in inefficient accounts – instead, one must focus on the average interest rate earned on savings.

My results suggest that spousal conflict over how much to save can give rise to inefficient savings behavior. Several important caveats are in order, however. First, since it was not feasible to randomly assign discount factors to individuals (the experimental ideal), the bulk of my analysis relies on heterogeneous treatment effects. It is important to note that I pre-specified the theoretical framework, my focus on heterogeneity in discount factors, and the associated heterogeneous treatment effects before implementing the experiment or undertaking any data analysis.[2] Indeed,

[2]Documents detailing this plan were drafted in May and August of 2009 and are archived in the J-PAL hypoth-

my decision to focus on heterogeneity in discount factors (rather than other aspects of preference heterogeneity) was not *ad hoc*, but rather motivated by two bodies of existing research. First, since the earlier-cited empirical literature supporting static efficiency suggests that households can reconcile heterogeneous preferences over *what* to consume, it is intuitive to turn to heterogeneity in preferences over *when* to consume as a driver of inefficient intertemporal behavior. Second, there is a well developed body of theoretical work that highlights the particular difficulties associated with aggregating preferences with differing discount factors.[3]

This does not, of course, eliminate the concern that it is not heterogeneity in discount factors, but rather some omitted characteristic correlated with this heterogeneity, that causes couples to make inefficient savings choices. For example, costly individual accounts may also be valuable because they can be used to hide resources from spouses. I designed the experiment to explicitly address this alternative theory and to assess whether the main results are robust to accounting for intrahousehold information sharing. I do find evidence that information matters – households in which spouses are poorly informed about one another's finances at baseline are more likely to choose individual accounts and there is some evidence that they also respond more adversely to a randomized information sharing treatment. However, these concerns are unrelated to the initial findings regarding preference heterogeneity – well matched couples have no better information flows than poorly matched couples, and the main results are unchanged when accounting for intrahousehold information sharing.

Another possibility is that poorly matched couples could be more prone to decision-making errors as compared to well matched couples (due to differences in financial literacy, for example). In practice, my results are robust to controlling for measures of spousal education, literacy, decision error in the discount factor elicitation questions, within-couple disagreement regarding consumption and savings decision making, self reported decision-making power, and a range of other observable characteristics recorded during the baseline survey. Even so, given the nature of the analysis I cannot completely rule out the hypothesis that my results are driven by some unobservable factor correlated with discount factor heterogeneity.

Another important caveat is that most couples in my sample have small savings balances – as a result, even though my results pertaining to interest rate losses are quite large in percentage terms, they are necessarily small in absolute terms. This makes it difficult to assess whether the efficiency losses I observe are indicative of more meaningful economic losses in other aspects of couples' lives. While I do find that poorly matched couples are somewhat less likely to remain married three years after the initial experiment, which suggests that the costs of preference heterogeneity

esis registry. These documents can be downloaded here: `https://sites.google.com/site/sschaner/home/Preanalysis.zip?attredirects=0&d=1`

[3]See, for example, Marglin (1963), Feldstein (1964), Browning (2000), Caplin and Leahy (2004), Gollier and Zeckhauser (2005), Zuber (2010), and Jackson and Yariv (2012).

can be substantial, this difference is only statistically significant when controlling for observable characteristics of couples. Finally, this paper requires the external validity caveat common to many field experiments – the couples in my sample reside in a relatively small number of communities in Western Kenya, so additional research would be needed to evaluate whether preference heterogeneity has similar consequences in other countries and contexts.

The main 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, my 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 insight could be useful for reconciling some of the differences observed in household efficiency both within and across studies and geographical contexts.

This paper also contributes to a growing literature that studies savings and investment in developing countries. Recent research indicates that individuals prefer to use costly but secure informal savings devices over less costly, less secure alternatives (such as saving at home) in order to protect resources from other members of the household, especially spouses (Anderson and Baland 2002; Collins et al. 2009; Karlan and Morduch 2010).[4] My results provide a rationale for this behavior and suggest that preference heterogeneity could play a key role in determining how, and how efficiently, people save. Moreover, many households in developing countries engage in entrepreneurial activities – for these households, business investment is an important tool for transferring resources over time. Seen in this light, this paper's insights may also be relevant for understanding the widely noted heterogeneity in returns to household-run microenterprises (McKenzie and Woodruff 2008; de Mel et al. 2009; Fafchamps et al. 2014).

The remainder of the paper proceeds as follows: Section 2 develops the theoretical framework used to structure the empirical analysis. Section 3 describes the experimental design and data, Section 4 presents main results, Section 5 discusses other alternative explanations, and Section 6 concludes.

2 Theoretical Framework and Testable Predictions

This section develops a simple model to (i) illustrate how heterogeneity in time preferences can generate inefficient strategic savings behavior and (ii) derive a series of testable implications that I

[4]There are also other reasons why individuals may value informal savings devices, such as a need to protect savings from oneself (as in Ashraf et al. 2006b) and a need to protect savings from appropriation by members of the community (as in Baland et al. 2007; Jakiela and Ozier 2012).

take to the data in Section 4.

2.1 Economic Environment and Decision Structure

I model a household comprised of two agents (i.e. spouses), A and B. To focus on strategic behavior stemming from differences in savings motives, I assume that both spouses have identical utility functions defined over a single public consumption good, c (this way the only choice that the household must make is *when* to consume).[5] The spouses live in a two period world and exponentially discount per-period utility. The individual discount factor for agent i is δ_i and I assume that the per-period utility function $u(\cdot)$ is continuous, increasing, and concave with $u'(c) \to \infty$ as $c \to 0$.

At the beginning of each period the household receives a deterministic endowment, y_t, which can be consumed or saved (I assume there is no borrowing in this economy).[6] Households have access to three different savings technologies: (1) a public (i.e. joint) bank account, which yields rate of return R_J, (2) a private (i.e. individual) bank account for agent A, which yields rate of return R_A, and (3) a private bank account for agent B, which yields rate of return R_B. What makes the "public" account public is that any member of the household can deposit and withdraw funds. In contrast, a "private" account can only be accessed by its owner.

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

Within each period, the following sequence of events occurs:

1. The endowment, and returns on any previous-period savings are revealed.

2. Nature selects which of the two spouses will have the first opportunity to travel to the bank. This sequential ordering of trips is meant to capture the fact that the bank is located in town, and opportunities to go to town may arrive at different times for different spouses.

3. The first mover chooses whether to pay banking cost b to go to the bank. The first mover can only deposit/withdraw from his or her individual account and the joint account.

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

[6] Alternatively, one could assume that the husband and wife receive separate endowments that they have exclusive access to. As long as savings constraints do not bind (i.e. an individual never saves his or her entire endowment), the analysis would be unchanged. It is also straightforward to make the second period endowment stochastic – allowing for this leaves the arguments and implications that follow essentially unchanged.

4. The second mover observes the choices of the first mover and decides whether to pay b to go to the bank. The second mover can only deposit/withdraw from his or her individual account and the joint account.

5. Any unsaved resources are consumed and the period ends.

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

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

2.2 Equilibrium Savings Strategies

In what follows, I restrict attention to pure strategy equilibria. I further assume that if more than one pure strategy equilibrium exists, the couple will never choose a Pareto dominated equilibrium. Note that when $\delta_A \neq \delta_B$ this refines the set of subgame perfect Nash equilibria to those that generate a unique consumption allocation $\{c_1^*, c_2^*\}$.[7]

I solve the game by working backwards. For expositional clarity, I will refer to the first mover as agent A or the husband and the second mover as agent B or the wife. The solution to the second (and final) period problem is straightforward – individuals will withdraw their bank savings and consume all available resources. In the first period, the second mover will make savings choices that account for this fact, taking the earlier choices of the first mover as given. The first mover will in turn make savings decisions taking account of the second mover's response. Appendix A, which characterizes the solution to the savings game in more detail, shows that in equilibrium at most one spouse will travel to the bank (see Lemma 1). The intuition for this result is that spouses plan their actions to minimize wasteful banking costs as much as possible. This means that one can think of the first mover (agent A) as effectively having two choices: he can either stay home and let the second mover (agent B) save, or he can save in his individual and/or the joint account

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

in a way that does not incite the second mover to return to the bank to further adjust the savings allocation. As a result, the first mover's decisions ultimately determine how efficiently the couple will save.

To concretely see how preference heterogeneity can generate inefficient equilibrium savings strategies, it is useful to consider an example in which δ_B is held constant and equilibrium choices are analyzed with respect to δ_A. Panel A of Figure 1 illustrates how equilibrium savings strategies change with δ_A when $u(c) = ln(c)$, $y_1 = 30$, $y_2 = 10$, $b = 0.05$, $R_A = 1.03$, $R_B = 1.05$, $R_J = 1.04$, and δ_B is fixed at 0.7. In this example, agent B has access to the highest rate of return. Thus, if agent A wishes to save efficiently, he must defer to his wife, who will deposit savings in her individual account according to her own preferences. Alternatively, agent A can consider depositing funds in either his individual or the joint account in order to manipulate the equilibrium consumption allocation. Both these actions, however, would mean that the couple's savings would not earn the highest possible rate of return.

When the couple i*s* perfectly matched on discount factors ($\delta_A = \delta_B = 0.7$), the first mover/agent A is happy to let the second mover/agent B save, as the same allocation maximizes both agents' utilities. In fact, Figure 1 shows that the couple saves efficiently (at agent B's most-preferred savings level) as long as preference heterogeneity is not too large – this is because when δ_A is close to δ_B, agent A is relatively satisfied with his wife's savings choice – it is therefore not worth it to him to "pay" the efficiency loss needed to manipulate the consumption allocation. However, when agent A becomes sufficiently more patient than his wife he begins to take strategic action to shift additional consumption to the second period (this begins at $\delta_A = 0.8$ on the graph). Notice that at first, agent A only saves in the joint account, but as preference heterogeneity increases he starts to make increasingly intensive use of his individual account, even though $R_A < R_J < R_B$. This is because if agent A were to deposit too much in the joint account, agent B would return to the bank to make a withdrawal – thus, while preference heterogeneity can lead to the use of both dominated joint and individual accounts, the security of the individual account makes it more attractive for extreme savings deviations.[8]

An interesting insight of the model is that the less patient spouse may also engage in inefficient savings behavior, which in turn implies that it need *not* be the case that savings balances in dominated accounts will increase with preference heterogeneity. Consider the part of the graph where $\delta_A < \delta_B = 0.7$. Agent A knows that if he does nothing, his wife will save in her individual account when it is her turn to go to the bank. Since agent A does not value the second period very much, he would like to find a way to reduce first period savings. To accomplish this, he could

[8]Notice that joint savings does not to drop to zero, even when $\delta_A = 1$. Since agent B would have to pay banking cost $b = 0.05$ to withdraw from the joint account, this implies that agent A will always be able to save at least 0.05 in the joint account without triggering a withdrawal.

consider saving just enough to make his wife indifferent between paying b to go back to the bank and staying at home. On the graph, notice that when δ_A is relatively close to δ_B, agent A makes use of the joint account, whereas once preference heterogeneity becomes sufficiently large he makes more intensive use of his even less efficient individual account. This is because even though the joint account offers a higher return than A's account, saving jointly makes returning to the bank relatively more attractive for the wife (since she can withdraw the joint savings and reallocate it to her own, highest return account). Thus, saving individually allows agent A to reduce the total amount of first period savings – again we see that individual accounts are more useful for extreme savings deviations due to their extra security.

The patterns in Panel A of Figure 1 suggest that the average interest rate earned on the couple's savings declines as preference heterogeneity increases. To formally capture this, I define the "interest loss", L, to be the difference between the highest possible interest rate and the average interest rate earned on bank savings. In other words:

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

where $R_{max} \equiv \max\{R_A, R_B, R_J\}$. Figure 1, Panel B shows that the interest loss on savings increases with preference heterogeneity, with the interest rate loss following a pronounced U-shape with respect to the discount factor of agent A. This result is intuitive – as discount factors diverge, spouses are willing to pay more to exert control over the time path of consumption. Since the discount factor of agent B is fixed, the agent A can always choose from the same set of savings strategies. Therefore as preference heterogeneity increases agent A is willing to use increasingly costly savings strategies to exert control, which means that L gets increasingly larger.

To recap, this example has generated four important insights: (1) perfectly matched couples will always save efficiently, (2) both more and less patient spouses may engage in strategic behavior, (3) both joint and individual accounts may be used strategically, (4) the interest rate loss on the couple's savings allocation increases with preference heterogeneity. The following proposition shows that these insights translate beyond the particular example in Figure 1:

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

Proof. See Appendix A. ■

2.3 Caveats and Extensions

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

Another related concern is that the model abstracts away from other motives for saving, such as risk preferences. This is an especially important consideration because my empirical measure of time preferences likely captures multiple savings motives, including risk aversion (Andreoni and Sprenger 2012). A model that maps more directly to the empirics would therefore be one in which couples are indexed by some generalized measure of heterogeneity in savings preferences. Appendix A shows that it is possible to derive an analog to Proposition 1 when one defines an increase in preference heterogeneity to be a perturbation in agent A's preferences that makes him generally more present/future oriented as compared to agent B.[9] Given this, I prefer to interpret the empirical results in terms of heterogeneity in savings preferences broadly defined, rather than in terms of heterogeneity in discount factors specifically.

It is also important to ask how central some of the modeling choices are for obtaining the main theoretical results. For example – instead of requiring that couples decide whether to go to the bank sequentially, a natural alternative would be to have the spouses simultaneously decide whether or not to go to the bank. In fact, it is possible to specify a simultaneous move version of the model and arrive at essentially the same results (Appendix A sketches this alternate model).[10] A second question is whether the theoretical results would be robust to allowing for infinitely repeated game

[9]The basic idea is as follows: consider the set of all non-dominated consumption pairs (c_1, c_2) and (c_1', c_2'), where $c_1 > c_1'$ and $c_2 < c_2'$. Now consider the set of pairs for which initially, $(c_1, c_2) \underset{A}{\preceq} (c_1', c_2')$. Then I define a preference perturbation to make A globally more savings oriented if $(c_1, c_2) \underset{A}{\prec} (c_1', c_2')$ after the perturbation. Similarly, consider the set of pairs for which initially $(c_1, c_2) \underset{A}{\succeq} (c_1', c_2')$. Then I define a preference perturbation to make A globally more present oriented if $(c_1, c_2) \underset{A}{\succ} (c_1', c_2')$ after the perturbation.

[10]Under the alternate model it is still important to give both spouses the chance to re-optimize savings deposits within a period. This allows the model to capture the strategic advantage of individual accounts (which are only accessible by the owner) over joint accounts.

play, as this setup would allow for the application of the Folk Theorem. In this case, one would expect couples to be able to reach the Pareto frontier provided they are sufficiently patient. This insight suggests that inefficient savings behavior should be more prevalent among couples in which at least one of the partners is very impatient, or the couple has a high risk of separation in the future. In practice, this "minimum discount factor" effect is difficult to isolate because conditional on preference heterogeneity, the minimum discount factor in the couple is directly correlated with the average level of patience (and thus the overall savings motive) of the couple. For this reason, I focus on the simpler implications of the two period model.

2.4 Testable Predictions

The goal is now to generate a set of theoretical predictions that can be evaluated in the data. To map the model to the experimental context, consider a sample of couples who are characterized by some distribution of income, banking costs, and spousal discount factors. These couples are randomly assigned different sets of interest rates $\{R_A, R_B, R_J\}$.

The first prediction characterizes patterns of joint versus individual account use. Given that individual accounts are more useful for strategic savings purposes, it is intuitive to conjecture that individual account use should increase with preference heterogeneity. In fact, since joint accounts may be used strategically as well (recall Figure 1), the model does *not* predict that rates of individual account use will monotonically increase with preference heterogeneity. However, consider the subset of couples for whom the joint account offers the highest rate of return. Since the interest rate loss L is increasing in discount factor heterogeneity by Proposition 1, it must be that *for this subset of couples* individual account use increases with discount factor heterogeneity.

T1. Consider the subset of couples for whom the joint account offers the highest rate of return. All else equal, rates of individual account use will increase with preference heterogeneity.

One concern with the first testable prediction is that it could also be consistent with a model where couples save according to different rules of thumb – that is, well matched couples may simply prefer to engage in joint endeavors, while poorly matched couples may prefer to make more independent choices. However, the theory predicts that as long as couples are sufficiently well matched on discount factors, they will save efficiently – and this entails responding to relative rates of return on *both* individual *and* joint accounts. This insight motivates the second testable prediction:

T2. Perfectly matched couples ($\delta_A = \delta_B$) will always save efficiently (in the sense that $L = 0$). Thus, perfectly matched couples will respond to relative rates of return on both individual and joint accounts in a manner consistent with efficient investment. In contrast, poorly matched couples ($\delta_A \neq \delta_B$) need not save efficiently.

Finally, I map Proposition 1 directly onto the data to form the final testable prediction:

T3. All else equal, the interest rate loss (L) will increase in preference heterogeneity.

I now describe the field experiment that I use to evaluate testable predictions T1-T3.

3 Experimental Design and Data

3.1 Experimental Design

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

All study participants were offered Family Bank *Mwananchi* accounts. This account could be opened with any amount of money, though a minimum operating balance of Ksh 100 (\$1.25) could not be withdrawn. The account paid no interest, but deposits were free of charge and there were no recurring maintenance fees. The only fees associated with the account were withdrawal fees, which were Ksh 62 (\$0.78) over the counter and Ksh 30 (\$0.38) with an ATM card. Account holders could purchase an ATM card for Ksh 300 (\$3.75), though this was not mandatory.

At the outset of the study, I identified communities surrounding 19 local primary schools, which would serve as group meeting grounds (all experimental activities were conducted in these group sessions). These communities were situated either on the outskirts of Busia town or in nearby rural areas. Trained field officers issued meeting invitations to married couples 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.[11] Just 7 percent of otherwise eligible couples were excluded due to pre-existing ownership of Family Bank accounts and approximately 29 percent of issued invitations were redeemed over the course of the study. Thus while far from universal, takeup rates are high enough that the sample represents a nontrivial fraction of targeted married couples in the catchment area.[12]

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

[12]Unfortunately I do not have data on characteristics of nonparticipants, which makes it difficult to study selection into study participation. An analysis of how demographic characteristics vary with the cost of traveling to the bank finds no evidence of differential selection by match quality, however. (Results available upon request).

Interventions All participating couples were given the opportunity to open up to three Family Bank accounts: an individual account in the name of the husband, an individual account in the name of the wife, and a joint account. To maximize takeup, I funded each opened account with the Ksh 100 minimum operating balance (this amount could not be withdrawn by participants – it simply made opening an account costless). While participants could in principle open an account with Family Bank at any time, only those accounts opened during experimental meetings were eligible for the operating balance subsidy and the experimental interventions, which are described below.[13]

Intervention 1 - Interest Rates Each potential account was randomly assigned an interest rate (respondents drew envelopes with the interest rates from tins upon arrival at the meeting). Since prediction T1 only holds when the joint account offers the highest rate of return, I designed the experiment so that individual accounts could bear either 0, 2, 6, or 10 percent 6-month yields (or 0, 4, 12, or 20 percent returns on an annual basis) with equal probability, while joint accounts could bear either 2, 6, or 10 percent 6-month yields (with equal probability). These interest rates were very high compared to market alternatives: small scale savings balances could earn at most 0.5-2.0 percentage points of interest annually given bank accounts available in Busia at the time of the experiment. The experimental interest rates were temporary, and expired after six months.[14] 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. Appendix Table D1 shows that respondents were significantly more likely to open and use accounts with higher temporary interest rates – this suggests that respondents understood the treatment.

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

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

Conditional on R_a the experiment created 10 percentage points of random variation in $excess_a$ for each account type.[15] After observing their interest rates, couples were separated and each spouse was administered a baseline survey. One concern is that randomizing the interest rates before

[13] 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 (2013b). I do not analyze this intervention in this paper, as there are no clear testable predictions with respect to the ATM treatment.

[14] 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. The interest payments were made by IPA-Kenya and after the six month period, balances earned no interest (respondents were informed of this *ex-ante*), which at the time was standard for the *Mwananchi* account and other current accounts in Kenya.

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

conducting the baseline influenced survey responses. However, interest rates are not systematically associated with elicited discount factors, baseline self reports of savings levels, savings device use, or self reported decision making power regarding consumption and saving. It is therefore likely that the randomization had little impact on survey responses.

After the baseline, couples were reunited and decided which accounts to open. The fact that couples decided which accounts to open together could have impacted account opening choices. For example, if joint decision making increased the degree of cooperation between spouses, this would reduce the likelihood of observing inefficient savings behavior in the sample. The public decision making could have also nudged couples to open joint accounts over individual accounts, in which case my results will understate the preference for individual accounts in the sample.

Intervention 2 - Extra Statements Hidden information appears to be important in households in developing countries.[16] In order to test whether the ability to hide savings was an important driver of individual account use, 50 percent of participating couples were randomly selected for an "extra statements" offer.[17] If a selected couple decided to open an individual account for (without loss of generality) spouse A, the enumerator processing the couple's paperwork asked if the spouses would consent to allow spouse B to receive extra statement cards. The cards, if presented by spouse B at the bank, entitled him or her to learn the current balance of spouse A's account. These cards were only valid for 6 months, and were not given to couples unless both spouses gave their consent.

Marital Status Verification All participating couples were re-contacted approximately 3 years after the initial intervention for the purposes of a follow-up study (see Schaner 2013c for details). As part of this follow up, enumerators updated the marital status of each study participant and found that 32 (5 percent) of the couples in the core sample were not actually married at the time of initial field activities. Moreover, these “false” couples were somewhat more likely to be poorly matched on time preferences (see Appendix Table D2), which would be expected provided couples match assortatively on time preferences. One would also expect false couples to favor individual accounts and ignore excess interest rates. To avoid the risk of biasing my results in favor of the testable predictions, I therefore limit the sample in this paper to those couples whose marital status

[16]For example, Anderson and Baland (2002) find that women's use of ROSCAs in Kenya is consistent with a model of hidden information. Boozer et al. (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 a 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.

[17]Due to delays in approvals from the bank, extra statements were not offered to the 84 couples (15 percent of the sample) in the first 6 (of 33) experimental sessions.

was successfully re-confirmed at follow up.[18]

3.2 Data

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

Measuring Rates of Time Preference The baseline elicited time preferences using choices between different amounts of money at different times, as opposed to different amounts of goods at different times. I made this choice for three reasons. First, Ashraf et al. (2006b) find that while time preference parameters estimated using choices between money, rice, and ice cream were all correlated, only the parameters estimated using money choices significantly predicted takeup and use of a commitment savings product. Second, even though discount rates estimated using money choices should theoretically reflect external interest rates, in practice respondents do not appear to take account of this when making choices (see Andreoni and Sprenger 2012 for a summary). Finally, cash lotteries made intuitive sense to respondents given that the group sessions revolved around bank accounts and savings.

All questions were framed as a choice between a smaller amount of money at a nearer time (x^t) and a larger amount of money at a farther time $(x^{t+\tau})$.[19] In total, participants responded to 10 tables of monetary choices, with each table consisting of 5 separate choices between a smaller $x^t \in \{\text{Ksh } 290, 220, 150, 80, 10\}$ and larger $x^{t+\tau} = \text{Ksh } 300$.[20] The questions involved sizable amounts of cash relative to respondents' incomes – for comparison, median reported weekly earnings in the sample were Ksh 700 for men and Ksh 300 for women. In order to make decisions salient, respondents were given a 1 in 5 chance of winning one of their choices (the choice was also

[18]Over 96 percent of couples were re-contacted for followup. All told I drop 32 "false" couples and 22 unconfirmed couples from the core sample of 598, resulting in a final sample size of 544 couples. Follow up is not correlated with baseline match quality (Appendix Table D2) or any of the treatments (Table 2, row 1).

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

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

selected at random). If a respondent won one of her choices, she had the option of having the funds deposited directly in her bank account, or picking the cash up at the field office, also located in Busia town.[21]

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

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

I topcoded $\hat{\delta}_i$ at $\overline{\delta}$, the value of $\hat{\delta}_i$ obtained via nonlinear least squares for always-patient responses and bottomcoded $\hat{\delta}_i$ at $\underline{\delta}$, the value of $\hat{\delta}_i$ for always-impatient responses.[22]

Panels A and B of Figure 2 graph the distribution of estimated discount factors for men and women. Discount factors span a wide range of values, but on average study participants appear to be very impatient, with weekly discount factors averaging 0.72 for men and 0.70 for women. These discount factor estimates are lower than estimates in studies of individuals in developing countries in Asia, but consistent with other studies in Africa that have found very high rates of impatience regarding the timing of cash payments.[23]

The histograms also illustrate the extent of censoring in the sample. First, 15 percent of individuals were "always impatient", and preferred Ksh 10 sooner to Ksh 300 in the future in all tables. Nonlinear least squares converges to a discount factor estimate very close to zero for this group. Another 25 percent of individuals were "always patient" and preferred Ksh 300 in the future to all sooner amounts. In general, this measurement error will lead me to overestimate match quality in couples with censored discount factors, which should bias the empirical results away from the testable predictions.

My baseline measure of intrahousehold preference heterogeneity is simply the difference between the male and female estimated discount factors: $het_c\equiv\hat{\delta}_{Mc}-\hat{\delta}_{Fc}$. While 12 percent of couples had identical discount factor estimates, many couples had estimates that differ substantially.

[21]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.

[22]This led to the censoring of 16 estimated discount factors from below and 35 estimated discount factors from above.

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

In the analysis I frequently compare the behavior of well and poorly matched couples. To do this I label the 50 percent of the sample with the most closely aligned discount factors as "well matched" and refer to the remaining couples as "poorly matched". This corresponds to $\left|\hat{\delta}_{Mc} - \hat{\delta}_{Fc}\right| \leq 0.21$ and is equivalent to the couple's discount factors being within 0.49 standard deviations of one another.

Panel C of Figure 2 shows a weighted scatter plot of $\hat{\delta}_M$ and $\hat{\delta}_F$, with well matched couples indicated by darker shading. The figure shows that there is only weak assortative matching in my sample: the correlation coefficient between spousal discount factors is 0.09. As mentioned earlier, the central concern with my measure of heterogeneity is whether it is correlated with other characteristics of couples that determine savings behavior for reasons unrelated to differences in savings motives. To explore observable differences between well and poorly matched couples, I begin the empirical analysis with a comparison of demographic characteristics by match quality.

3.3 Background Results

Sample Characteristics by Match Quality The sample consists of 544 non-polygamous married couples.[24] Table 1 presents baseline characteristics of the sample by match quality. In terms of demographic and economic characteristics, respondents are of relatively low socioeconomic status. Yet almost all respondents reported using at least one savings device at baseline, with saving at home and saving with ROSCAs (rotating savings and credit associations – a type of informal savings group found in much of the developing world) most common. Formal bank and mobile money accounts were less prevalent, with each being used by approximately 21 percent of individuals. Well and poorly matched couples make similar use of all savings devices except for "other" ways of saving, which is more common among well matched couples.[25] At baseline individuals in well matched couples were also significantly more likely to own a mobile phone and lived further away from the bank branch. These two variables are straightforward to include as covariates in the main analysis.

The next rows of Table 1 study differences in individual time preferences by match quality. Importantly, well matched couples have significantly higher discount factors. To the extent that well matched couples are more likely to save overall, this should bias me away from finding evidence consistent with the testable predictions (since, from a theoretical perspective, not saving incurs no efficiency loss).

The remaining variables in Table 1 attempt to capture various aspects of individual and house-

[24] I dropped 161 polygamous couples from the sample since strategic behavior may be very different in households with more than one female head. However, as Appendix Tables D4-D6 show, the results are robust to including them.

[25] This category mostly consists of investments in livestock (arguably a "joint" method of saving, since farm animals are easily accessible to both spouses).

hold decision-making. These variables are especially important, since they can shed light on alternative reasons why poorly matched couples might save inefficiently. The first difference of note is that women in well matched couples had somewhat more say in household decisions: individuals in these couples were less likely to report that the husband is the primary decision maker and slightly more likely to report that the wife made most decisions about how the household spends money. In contrast, there were no significant differences by match quality in terms of which spouse did most of the *saving* in the household. It is possible that imbalances in bargaining power could also lead to inefficient strategic behavior, especially if spouses with low bargaining power (which usually means women in my rural Kenyan context) take strategic action to shift the balance of power in their favor (indeed, this mechanism would arguably exacerbate incentives for strategic behavior stemming from preference heterogeneity). On the other hand, one could also argue that differences in decision-making power by match quality are to be expected, since reported power should reflect the resource allocation process in the household (individuals who engage in strategic behavior may feel that they have greater control over the intrahousehold resource allocation process, for example).

Another possibility is that couples who are better matched on time preferences could be better aligned in general and less prone to decision-making error. For example, suppose that couples match assortatively on the marriage market, and there is no actual discount factor heterogeneity in my sample. Then all *observed* preference heterogeneity would be due to measurement error, which would likely be correlated with financial literacy and cognitive ability. To the extent that couples with more measurement error in their discount factors make less efficient savings choices, this would bias the results in favor of the testable predictions. The remaining variables in Table 1 are intended to address these concerns.

First, an attractive property of the nonlinear least squares procedure that I use to estimate discount factors is that it not only delivers individual-specific estimates of the discount factor – it also delivers individual-specific estimates of the standard deviation of the choice error (specifically, I use $\hat{\mu}_i$ as my measure of "decision error"). Surprisingly, decision error appears to be greater among well matched couples, which would bias me away from finding results consistent with the testable predictions. It is important to note, however, that my measure of decision error is highly correlated with the level estimate of the discount factor (more patient individuals tend to have greater estimated decision error). Conditional on the level discount factor, there is no significant difference in decision error between well and poorly matched couples.

The next two variables use self-reported decision making power to construct estimates of how well aligned spouses are in their perception of who decides about consumption and saving. Overall, there is a substantial degree of misalignment, with spouses in over half of couples giving conflicting reports regarding decision-making.[26] While there is no difference between well and poorly matched couples in terms of misalignment over how to spend money, poorly matched couples are significantly more likely to disagree about who does most of the saving in the household. I interpret this with result caution: while it is conceivable that strategic savings behavior could increase savings misalignment (especially if control over saving shifts from period to period), it is also possible that poorly matched couples could be misaligned for other reasons, which would bias the results in favor of the testable predictions.

Finally, to test the robustness of my results to hidden information, I use spousal cross reports of income and savings device use to construct an "information sharing index" which ranges from zero (worst informed couples) to one (best informed couples). I then define a couple to be "poorly informed" if their information sharing index is below the sample median.[27] If the return to hiding savings from a spouse is increasing in discount factor heterogeneity, then hidden savings concerns could be responsible for the main empirical results (though it is not obvious that such a correlation should exist).

I added the "extra statements" treatment to the field experiment to validate my measure of information sharing and assess the overall importance of hidden savings in the study population. Appendix C presents the hidden savings analysis in detail. While the extra statements intervention is somewhat underpowered, there is some suggestive evidence that hidden savings concerns are important: just 69 percent of individuals who were presented with the extra statement offer consented, and poorly informed and poorly matched couples appear to respond more adversely to the extra statements treatment.[28] Table 1 suggests that informational concerns may largely be orthogonal to match quality, as there is no significant difference in information sharing by match quality.

Overall, Table 1 makes it clear that there are some important differences between well and poorly matched couples. However, I find no systematic patterns suggesting that poorly matched couples should save less efficiently than well matched couples. To evaluate the robustness of my results, I present my main estimates both with and without controls for all time preference, demographic, economic, and decision making controls in Table 1 except those capturing savings device use, overall savings levels, and reported savings behavior of couples since these could be outcomes of strategic behavior. (I do, however, control for intra-couple disagreement over who does most of the saving, as there is no explicit link between this variable and the model.)

[26]There are two primary types of disagreements: one spouse reports that one individual decides while the other reports that both decide, or each spouse reports that a different individual decides. About three quarters of disagreements regarding consumption are of the first type, while savings disagreements are evenly split between the two types.

[27]Slightly over half of couples are coded as badly informed due to a mass of couples with the same index value around the median. See Appendix C for more detail on the information sharing index.

[28]Consent rates were nearly identical by gender – 68.2 percent of men and 69.2 percent of women consented to the extra statements.

Randomization Verification I check randomization by running individual level regressions of each characteristic listed in Table 1 on five treatments of interest: the excess interest rate on the husband's, wife's, and joint account, as well as the extra statements treatment and a dummy for whether the individual was selected for a cash payment.[29] Columns 1-5 of Table 2 report regression coefficients and standard errors for each treatment (rows correspond to a single regression). I also estimate all equations jointly via seemingly unrelated regression to test whether each treatment is significant across all equations. The p-values from the joint tests are in the last row of Table 2.

Overall, the randomization appears to have functioned well. With one exception the joint tests fail to reject the null of no correlation, and there are no systematic patterns across the different treatments. Importantly, excess interest rates are uncorrelated with discount factor heterogeneity. Although the joint significance test does not suggest a lack of balance, column 3 does show that couples who received higher excess interest rates on the joint account are of lower socioeconomic status (in particular, they have less education, are less likely to be literate, are more likely to be subsistence farmers, and have more children). As the main analysis will show, my results are robust to including controls for these observable characteristics.[30]

The joint significance test does suggest that balance is off for cash payments, however. Additional randomization verification (see Appendix Table D3) shows that significantly *fewer* cash payments were awarded than expected (17.7 percent instead of the expected 20 percent). All randomization was conducted in the field, by allowing respondents to draw folded envelopes from tins. Since fewer cash payments were awarded than expected, this suggests that the lack of balance is due to chance rather than enumerator deviations from the experimental protocol. While cash prize receipt is uncorrelated with the excess interest rates, it is correlated with preference heterogeneity – this could serve to elevate overall savings rates among well matched couples. Since I treat not saving as efficient, this should bias the results away from the testable predictions. To address this I control for husband and wife cash payment selection throughout the analysis.

Basic Overview of Account Use Before beginning the main analysis, Table 3 summarizes well and poorly matched couples' use of the experimental bank accounts. The first panel describes account opening choices. While all couples opened at least one account, very few couples opened all three. Instead the most popular choices were either only opening a joint account (58 percent of couples) or opening two individual accounts (26 percent of couples). Even though there was no monetary cost to open all three accounts, the additional time spent doing paperwork may have been

[29]The extra statements treatment dummy is coded to zero for the individuals in the first six sessions who were not eligible for the treatment. I therefore include an additional dummy variable that identifies these individuals in all regressions and follow this convention throughout the paper.

[30]As an additional robustness check, Appendix Tables D9-D12 show that the main results are robust to reweighting the sample to eliminate this lack of balance.

enough to dissuade couples from opening accounts that they were very certain they would never use – since enumerators explained that the Ksh 100 opening balance could not be withdrawn from the accounts, there was no strategic reason to open all three accounts in order to earn additional cash.

I exclude the Ksh 100 minimum balance from all measures of account use – thus, unopened and opened but unused accounts are treated equivalently. As motivated by the theoretical framework, a key outcome is whether an account was used for saving – I define a couple to have saved in account *a* if at least one deposit (other than the initial minimum balance subsidy) was made in the first 6 months following account opening. The next two panels show that 44 percent of couples saved in at least one account. This figure drops to 27 percent when deposits for the discount factor elicitation payoffs are excluded.[31] The 44 percent of couples who do save make an average of 3 deposits in the six months following account opening and hold an average daily balance of just over Ksh 1,000 ($12.50) in their experimental accounts.

Based on Table 3, the savings patterns of well and poorly matched couples appear to be quite similar. However, these summary statistics are not very useful for evaluating the theory in Section 2. First, recall that individuals in poorly matched couples may make strategic use of both joint and individual accounts – so simply comparing overall rates of joint and individual account use by match quality is not instructive. Moreover, since both more and less patient spouses may save strategically, there are no testable predictions regarding average balances or number of deposits in accounts. The next section therefore turns to the testable predictions to assess support for the the ideas in Section 2.[32]

4 Main Results

Testable Prediction T1 The first testable prediction states that when $R_J = R_{max}$, couples' use of individual accounts should increase as match quality decreases. (That is, there should be a U-shaped relationship between individual account use and $\delta_M - \delta_F$). Figure 3 tests this prediction graphically. The figure presents results of the following local linear regression

$$y_c = g\left(het_c\right) + \varepsilon_c \tag{1}$$

[31]For the main analysis I include the cash payments in measures of account use, since efficient households should always invest these payments in the highest return account. The results are, however, generally robust to using measures of account use that ignore these payments, or simply dropping cash payment recipients (see Appendix D).

[32]Although there are no direct testable predictions regarding other measures of account use like the average daily balance and number of deposits, results using these outcomes are generally similar. Appendix D presents these results in the interest of completeness.

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

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

$$y_c = \beta_0 + \beta_1 badmatch_c + w_c'\gamma + x_c'\lambda + \varepsilon_c \qquad (2)$$

where y_c is the outcome of interest and $badmatch_c$ indicates poorly matched couples. The regressions also include a vector of dummy variables for the interest rate on each experimental account, controls for husband and wife cash prize selection, extra statement selection, and the free ATM status of every account (w_c).[33,34] I also evaluate the robustness of the results to controls for time preference, demographic, economic, and decision-making characteristics (x_c).

Table 4 presents estimates of β_1 ("Poorly Matched") and tests robustness of the results to adding additional controls. The "basic" control set, in column 1, only includes w_c. This regression essentially mirrors the results of Figure 3. As expected given the graphical results, poorly matched couples are significantly more likely to save in individual accounts, even when limiting the sample to savers. Moreover, the differences by match quality are substantial – while just 7.9 percent of well matched couples save individually when $R_J = R_{max}$, nearly 19 percent of their poorly matched peers choose to save individually. Panel B shows that this pattern persists when limiting the sample to couples who saved. Just 18.1 percent of well matched savers make use of an individual account, whereas almost 40 percent of their badly matched peers save individually.

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

[34] Note that the free ATM treatment was randomized conditional on an account being opened, but the present analysis does not condition on account opening. To address this I control for "ex-ante ATM" selection – this is equal to actual ATM treatment status for open accounts and is equal to 1 for a randomly selected subset of unopened accounts, where the ex-ante ATM selection probability for unopened accounts is set to the theoretical free ATM selection probability a couple would have faced had they opened a given account.

The next four columns test the robustness of the results to the inclusion of additional controls. Column 2 adds controls for features of couples' time preferences. To account for general differences in patience between well and poorly matched couples I control for $\hat{\delta}_{Mc}$ and $\hat{\delta}_{Fc}$ linearly. I also include dummies indicating that each spouse is either impatient now-patient later or patient now-impatient later, and dummy variables to identify upper and lower censoring of the discount factor for each spouse. The third column adds additional controls for demographic characteristics and the fourth column adds further controls for economic characteristics. The demographic and economic controls sets also include controls to capture intra-couple heterogeneity: for non-binary characteristics I include the linear and squared terms for both husband and wife, as well as the interaction between the linear values for husband and wife. For binary variables, I include the dummy variable for both husband and wife as well as the interaction.[35]

To evaluate alternative explanations for savings inefficiency, the final column adds a variety of controls related to household decision making. First, to account for hidden information, I include a dummy variable identifying poorly informed couples. Hidden savings does appear to be important – the last column of Table 4, Panel A shows that poorly informed couples are significantly more likely to save in dominated individual accounts. Next, I add controls for both husbands' and wives' decision error in the time preference questions. Both these variables are unrelated to account use. I also control for whether the couple disagrees about who in the household decides how money is spent and who in the household does most of the saving – again, neither of these factors is significantly related to individual account use. Finally, the decision making control set also includes both husbands' and wives' reports of decision making power regarding how money is spent in the household. Coefficient estimates for these variables are omitted for space constraints, but there is no evidence that self-reported decision making power is related to the use of dominated individual accounts.

Testable Prediction T2 While the results in Table 4 are striking, they are also compatible with an alternative theory in which couples with similar preferences prefer to save jointly while couples with different preferences prefer to save individually. The second prediction – that perfectly matched couples always save efficiently – can be used to rule out this possibility. Before moving to the analysis, it is necessary to address one complication in assessing the efficiency of couples' savings choices. To simplify the theory, I assumed that banking costs were nonstochastic and the same for all accounts. In practice, the marginal cost of going to the bank is low when an individual

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

is in town for another reason, but high when an individual must make a trip to town specifically to go to the bank. In such a context, a joint bank account offers an important advantage: the couple can always send the spouse with the lowest cost of going to town to the bank. To capture this idea as simply as possible, imagine that the banking cost on a joint account is always less than the banking cost on an individual account ($b_J < b_i$).

In this case it may be efficient to save in a joint account even when $excess_J < 0$, and the joint account will be the clear choice when $excess_J \geq 0$. In contrast, it will never be efficient to save in an individual account i with $excess_i < 0$, and it may (or may not) be efficient to save in an individual account with $excess_i \geq 0$. Now, suppose there is a distribution of banking costs in the population and that all couples save efficiently. Figure 4 uses a specific example, in which each interest rate R_a is evenly distributed across the range $[0, 20]$ in 2 percentage point intervals, to illustrate how savings rates vary with the excess interest rate for individual accounts (Panel A) and joint accounts (Panel B).[36]

Panel A confirms that no couples save in the individual account when $excess_i < 0$. There is a discrete jump up in the savings rate at $excess_i = 0$ – this corresponds to cases where the interest rate on the joint account is very unattractive, and the interest rate on both individual accounts is the same.[37] As the excess interest rate increases beyond zero, the share of households saving in account i increases. Note that the savings rate plateaus around $excess_i = 14$ – at this point the joint interest rate is so much lower than R_i that all couples in the sample prefer to save individually regardless of their banking costs.[38] Panel B illustrates a contrasting pattern for joint accounts. Here, couples begin to save jointly around $excess_J = -14$ (these are the couples with the most extreme gap between b_J and b_i) and all couples prefer the joint account once $excess_J = 0$.

Note that Figure 4 is just a numerical example – in practice the shape of the positive sloped parts of the graphs will depend on a number of factors, including the distribution of banking costs in the population. The universal insight is that when considering individual accounts there should only be a positive slope for *positive* excess rates, while when considering joint accounts there should only be a positive slope for *negative* excess rates. This asymmetry is a striking implication of efficient investment in the presence of heterogeneous banking costs.

I generate the empirical analog of Figure 4 by running the following regression separately by

[36] In this example, I also assume that $y_1 = 30$, $y_2 = 10$, $\delta_A = \delta_B = 0.7$. Each couple in the $n = 100,000$ population draws two banking costs from a $U[0, 0.5]$ distribution – the larger draw is the banking cost on the individual account, the smaller draw the banking cost on the joint account.

[37] Since the field experiment generated lumpy variation in the set of interest rates presented to couples, the existence of such a mass in practice seems reasonable.

[38] In the example, all households save at this upper plateau. In practice, it is likely that the share saving at the upper plateau will be less than one, since not all households will want to save at interest rate R_i.

account type (individual versus joint) and match quality:

$$saved_{ac} = \beta_0 + ex'_{ac}\delta + int'_{ac}\lambda + \varepsilon_{ac} \tag{3}$$

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 int_{ac} is a vector of dummy variables for account a's interest rate.[39] I then calculate predicted values of $saved_{ac}$ for each value of the excess rate, assuming equal distribution of the sample at each interest rate (0, 2, 6, and 10 percent for individual accounts; 2, 6, and 10 percent for joint accounts). Figure 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 4 that individual account use by well matched couples could jump discretely up at $excess_i = 0$. Since -2 is the largest negative value of the excess interest rate in the sample, I therefore fit separate lines for $excess_i \leq -2$ (this slope should be zero) and $excess_i \geq -2$ (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 $excess_J = 0$.

Column A graphs account use for well matched couples, while column B graphs account use for poorly matched couples. The results are strikingly different: well matched couples appear to respond to the excess interest rate when predicted by the theory, whereas if anything, poorly matched couples appear to save *less* at higher excess interest rates. While the patterns in Figure 5 are suggestive, some of the point estimates have very large standard errors (indeed, I cannot reject that any of the positive slopes in Panel A are equal to zero). To create a higher powered test, I generate splines in the excess interest rate. To match the theoretical shifts in slope illustrated in Figure 4, I place a knot at $excess_{ac} = -2$ for individual accounts and a knot at $excess_{ac} = 0$ for joint accounts. The splines therefore have two components: $exyes_{ac}$ (this captures the slope on the excess interest rate that is expected to be positive provided couples save efficiently) and $exno_{ac}$ (this captures the slope on the excess interest rate that is expected to be zero). I then pool both joint and individual accounts and run the following account level regression:

$$\begin{aligned} saved_{ac} &= \beta_0 + \beta_1 badmatch_c + (ex \times match)'_{ac}\delta + (ex \times badmatch)'_{ac}\gamma + \\ &\quad z'_{ac}\lambda + (z \times badmatch)'_{ac}\eta + x'_c\alpha + (ex \times x)'_{ac}\zeta + (z \times x)'_{ac}\psi + w'_{ac}\phi + \varepsilon_{ac} \end{aligned} \tag{4}$$

[39] As a result of the experimental design, some values of the excess interest rate were only realized for a very small number of accounts: 13 accounts had an excess interest rate of 2, 11 accounts had an excess interest rate of 6, and 11 accounts had an excess interest rate of 10. For each of these values, I downcode the excess interest rate by two percentage points (results are invariant to simply dropping these accounts). Similarly, I pool $excess_{ac} = -10$ and $excess_{ac} = -8$ as the omitted category in the regressions. I do this in order to identify all interest rate dummy variables, as accounts with zero percent interest had excess interest rates unique to them.

where ex_{ac} is vector containing the two excess interest rate splines ($exyes_{ac}$ and $exno_{ac}$), $badmatch_c$ is a dummy variable identifying poorly matched couples, $match_c$ identifies well matched couples, z_{ac} is a vector containing interest rate dummies and the joint account dummy, x_c is a vector of additional controls, and w_{ac} includes controls for husband and wife cash prize selection, extra statement, and free ATM selection. This specification allows me to estimate separate responses to the excess interest rate for well and poorly matched couples, while also allowing the response to the excess interest rate to vary with respect to other observable characteristics.[40]

The primary coefficients of interest are included in δ, which measures the response of well matched couples to the excess interest rate, and γ, which provides analogous estimates for poorly matched couples. The first column of Table 5 presents results with only basic controls (that is, x_c and its interactions are omitted from the regression). The results mirror the patterns in Figure 5 – well matched couples only respond to the excess interest rate when predicted by the theory (though power is still limited, as this is only significant at the 10 percent level). The estimated response is quite large in magnitude, however – the coefficient on $matched \times exyes$ implies that increasing the excess interest rate on an individual account from -2 to 8 would increase well matched couples' savings rate in that account by 13 percentage points. This represents a substantial increase given that overall, only 16 percent of accounts offered to well matched couples are used for saving. In contrast, poorly matched couples appear to be completely insensitive to the excess interest rate. Moreover, I am able to reject that the overall responses of well and poorly matched couples differ at the 5 percent level.

The next three columns test the robustness of the results to the inclusion of controls. Since I also include interactions between these controls and the excess interest rate treatments, I demean all included controls using the mean *among well matched couples*. When demeaned this way, the coefficients on the excess interest rate splines reflect the response to the excess interest rate at the average value of included demographic characteristics observed among well matched couples. Thus if the heterogeneous responses observed in column 1 were driven by some other characteristic in the control sets, then the results for poorly matched couples should mirror those for well matched couples once the excess interest rate response is allowed to vary with that control. Columns 2-5 progressively add the same time preference, economic, demographic, and decision making controls sets included in Table 4.

The results are remarkably stable, which suggests that differences in observables between well and poorly matched couples are not driving the initial heterogeneous treatment effects. However, it is important to caveat that the specifications in Table 5 suffer from limited power: I cannot formally

[40] Note that the excess interest rate is only random conditional on the interest rate and type of account. Therefore I interact z_{ac} with match quality (and the additional controls in x_c) – this is to avoid attributing any heterogeneity in responses to the interest rate or account type to heterogeneity in responses to the excess interest rate.

reject that the responses to *exyes* are the same for well and poorly matched couples, and I am only able to reject that the overall responses of well and poorly matched couples are the same under the basic control set.

Testable Prediction T3 The results so far suggest that poorly matched couples invest less efficiently than well matched couples. But how big are these distortions in economic terms? I now exploit the experimental design to estimate the magnitude of savings misallocation by match quality. This also permits a test of the final prediction – that the interest rate loss, L, increases with preference heterogeneity. I construct L according to the definition put forth in the theory section: I calculate the actual interest rate that each couple earned on experimental savings balances and subtract it from the maximum interest rate, and set interest rate losses to zero for all non-savers.

Since lower return joint accounts may be more efficient than higher return individual accounts when banking cost differentials are large, I also present results where I discount individual interest rates to reflect higher banking costs. I adjust individual interest rates downward to account for banking costs in two different ways. First, I attempt to proxy banking costs using observables. I conjecture that those couples who travel to Busia town frequently for non-bank related reasons and those couples who have low travel costs to town will have smaller differential banking costs (i.e. the joint account offers less in transaction cost savings).[41] Appendix B provides additional detail on how I used principal components analysis to construct a "banking costs index", which runs from zero (lowest hypothesized banking costs) to one (highest hypothesized banking costs).[42] To discount individual interest rates I multiply the cost index by an assumed maximum interest discount and adjust individual interest rates by the resulting product. As a further robustness check, I present a set of specifications where I discount all individual interest rates uniformly. While this method cannot capture heterogeneity in banking costs within the population, it does capture the fact that individual accounts incur higher banking costs without inducing a correlation between the size of the discount and observables correlated with the banking cost index.

I first study interest rate losses graphically. 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 (recall Panel B of Figure 1), with the lowest values observed among well matched couples. This pattern holds with and without banking cost adjustments to individual interest rates.[43]

Table 6 summarizes interest rate losses by match quality and presents regression results that

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

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

[43] The figure illustrates results using adjustments made with proxied banking costs. Results are very similar when individual interest rates are adjusted uniformly.

include additional controls for observables. The first column does not discount individual interest rates at all. If poorly matched couples had always chosen the highest return account available, the average couple would have earned 8.13 percentage points of interest. In practice, these couples averaged 7.11 percentage points of interest, leading to an interest loss of 1.03 percentage points. In contrast, well matched couples could have earned a maximum of 8.19 percentage points of interest and actually earned 7.51 percentage points. Therefore, the "loss gap" between poorly and well matched couples is 0.35 percentage points of interest, which is significantly different from zero. Even without accounting for differential banking costs, poorly matched couples suffer from greater savings misallocation – their losses are 52 percent larger than those of their well matched peers. Subsequent rows of Table 6 test robustness of this result by first controlling for account specific interest rates, cash prize, extra statement, and free ATM selection (here I use the regression specification described by equation 2). I then include the time preference, demographic, economic, and decision making control sets respectively. The loss gap remains significant and its magnitude grows in size after including additional controls, although in several cases the results are just marginally significant due to lost precision.

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

It is important to note that while the losses in Table 6 are large in percentage terms, they are small in absolute terms. For example, a loss of 3 percentage points in interest amounts to just Ksh 24 (\$0.30) for the 75$^{\text{th}}$ percentile saving couple. On the other hand, banking cost differentials persist for the life of the account, so long run absolute losses due to inefficient individual account use could be much larger. Moreover, results from the long-run follow up in Appendix Table D2 offer some suggestive evidence that, all else equal, the costs of a poor match are economically meaningful. Overall, poorly matched couples are 3 percentage points less likely to still be married at the 3 year follow up. While this difference is not statistically significant, it grows in magnitude (to over 9 percentage points) and becomes highly significant after including the observable control sets. This is consistent with the idea that it is costly for a couple to be mis-matched on time preferences, but that poorly matched couples are at least partially compensated for this cost on other dimensions of match quality.

Overall, my results fit the predictions of the theoretical framework very well. While well matched couples appear to save efficiently, poorly matched couples do not account for rates of return between accounts and tend to make use of inefficient individual accounts. The next section discusses alternative theories that could generate the patterns that I observe in the data.

5 Alternative Explanations

As mentioned earlier, a leading alternative hypothesis is that poorly matched couples simply make noisier or less efficient financial decisions as compared to well matched couples. This theory could rationalize both poorly matched couples' overuse of dominated accounts and their lack of sensitivity to the excess interest rate. This was the motivation behind including controls for education, literacy, decision error, and spousal alignment over decision making in the main results. Overall, I do not find any compelling evidence that the "noisiness hypothesis" is driving my results. First, recall from Table 1 that there are no systematic differences in these characteristics by match quality, and Tables 4-6 show that including these controls generally strengthens, rather than attenuates, the results. Second, both well and poorly matched couples respond robustly to interest rate *levels* (Appendix Table D3), which suggests that on average, both groups understood and made decisions informed by the interest rates. Of course, I am not able to speak to unobservable aspects of sophistication or decision noise, but it is comforting that the results are robust to a range of different observable proxies of these characteristics.

In setting up the model, the choice to feature only a public consumption good obviated the possibility that private accounts could be strategically used to change *what* individuals consume (rather, the only margin of influence was *when* consumption took place). In practice, individuals may use private accounts as a tool to change the composition of consumption allocations. One possibility is that spouses strategically save in individual accounts in order to increase bargaining power. Alternatively, individual accounts could be used to change the composition of future consumption if there are mental accounting norms in the household or if spouses use individual accounts to hide resources from their partners. This would be particularly important if heterogeneity in time preferences is correlated with heterogeneity in other preferences, or if individuals with low levels of bargaining power engage in strategic behavior to alter consumption allocations.

When saving privately impacts bargaining power, both spouses will have incentives to save simultaneously in their individual accounts (this result is established formally in the context of labor supply by Basu 2006 and Browning et al. 2011). In the experiment, just 3.5 percent of couples saved in both individual accounts. Even among those couples who opened *both* individual accounts *and* saved in at least one account, just 29 percent saved in both individual accounts. This suggests that these concerns are not a major driver of the use of the individual accounts in this study. Moreover, my results are robust to controlling for self-reported consumption decision making power.

I do, however, find some evidence that hidden savings concerns are relevant in my sample – Table 4 shows that spouses who are poorly informed about one another's finances are significantly more likely to make use of dominated individual accounts. But to the extent that these concerns are important, they appear to be largely orthogonal to preference heterogeneity. This is plausible – hiding savings is likely valuable because it allows individuals to increase their share of consumption, or tilt consumption towards goods that they favor. If the benefit of doing so is equally large for individuals in well and poorly matched households, accounting for it should leave the core results unchanged, which is what I observe.

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

6 Conclusion

This paper sheds light on the underlying drivers of inefficient intertemporal resource allocation by households. I structured the analysis by first specifying a model in which heterogeneity in rates of time preference creates incentives for individuals to save strategically, even when doing so is costly. I then derived three testable implications of the model: (1) as long as $R_J = R_{max}$ individual account use will increase in preference heterogeneity, (2) perfectly matched couples will respond to positive excess interest rates on individual accounts and negative excess interest rates on joint accounts, and (3) interest rate losses on experimental bank accounts will increase in preference heterogeneity.

The empirical results are consistent with all these predictions. This is, of course, subject to the caveat that I cannot completely rule out the hypothesis that the results are driven by some other omitted characteristic that is correlated with my estimates of preference heterogeneity. However, the stability of the results to the inclusion of a wide range of observable controls, including measures of other aspects of the household decision-making process, is consistent with the idea that the results are indeed driven by inefficiencies arising from conflicting savings motives.

An innovative feature of the experimental design is that it allows me to quantify investment efficiency in terms of interest rates. However, the experimental interest rates were temporary –

it is therefore important to ask whether match quality has broader implications for households' investment choices. I do observe that well matched couples are significantly more likely to invest in livestock and the family farm, which are inherently joint methods of saving that likely bear a higher rate of return than more private savings devices like ROSCAs. While this finding is suggestive, the baseline data lack detailed information on the costs of and returns to different savings devices, so it is difficult to precisely assess how this translates into actual interest rate losses for couples. It is also interesting to note that well matched couples were more likely to still be married at the post-intervention marital status check. Although this difference is only statistically significant after adding controls for observable characteristics, this suggests that welfare losses due to poor match quality could extend beyond savings and investment decisions.

My results add to a growing body of literature that rejects dynamic household efficiency, while presenting evidence that heterogeneity in intertemporal preferences drives inefficient savings behavior. A novel feature of this idea is that it provides a mechanism for why some households function well while others do not: when preferences are well-aligned there are no incentives to behave strategically and therefore no barriers to attaining an efficient outcome. Although this paper studies strategic savings behavior, the applications are more general. For example, many households in developing countries either engage in home production (such as farming or animal husbandry) or run small businesses. Investment in these activities is an important way of transferring resources across periods. The insights in this paper suggest that when preferences in the household differ, capital for these activities will not always be allocated to the most efficient user. This mechanism may therefore help account for some of the heterogeneity in plot yields (Udry 1996) or microenterprise returns (de Mel et al. 2009) observed in the developing world. A broader (and more speculative) implication of this mechanism is that greater marriage market frictions could lead to lower quality matches in terms of preferences, which could give rise to geographical variation in household efficiency.

References

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

Andreoni, J. and C. Sprenger (2012). Estimating Time Preferences from Convex Budgets. *American Economic Review 102*(7), 3333–3356.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Dupas, P. and J. Robinson (2013). Savings Constraints and Microenterprise Development: Evidence from a Field Experiment in Kenya. *American Economic Journal: Applied Economics 5*(1), 163–192.

Fafchamps, M., D. McKenzie, S. R. Quinn, and C. Woodruff (2014). Microenterprise Growth and the Fly-paper Effect: Evidence from a Randomized Experiment in Ghana. *Journal of De 106*, 211–226.

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 (2012). Collective Dynamic Choice: The Necessity of Time Inconsistency. Mimeo.

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

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

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

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

Mazzocco, M. (2007). Household Intertemporal Behaviour: A Collective Characterization and a Test of Commitment. *Review of Economic Studies 74*(3), 857 – 895.

McKenzie, D. and C. Woodruff (2008). Experimental Evidence on Returns to Capital and Access to Finance in Mexico. *World Bank Economic Review 22*(3), 457–482.

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

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

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

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

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

Schaner, S. (2013c). The Persistent Power of Behavioral Change: Long-Run Impacts of Temporary Savings Subsidies for the Poor. Dartmouth College Working Paper.

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

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

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

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

Table 1. Observable Characteristics of Study Sample by Match Quality

	Well Matched	Badly Matched	Difference	N
Age	38.9	39.5	-0.572	1088
	[13.1]	[12.8]	(1.04)	
Education	6.87	6.84	0.025	1084
	[4.04]	[3.93]	(0.287)	
Literate	0.737	0.776	-0.039	1088
	[0.441]	[0.417]	(0.028)	
Number Children	4.74	4.88	-0.138	1088
	[2.77]	[2.70]	(0.222)	
Subsistence Farmer/No Job	0.430	0.472	-0.042	1084
	[0.496]	[0.500]	(0.034)	
Distance from Bank (Miles)	4.04	3.71	0.326*	1088
	[2.14]	[2.14]	(0.184)	
Income Last Week	960	1177	-217	1057
	[1759]	[2819]	(148)	
Owns Mobile Phone	0.495	0.423	0.073**	1083
	[0.500]	[0.494]	(0.034)	
Participates in ROSCA	0.574	0.564	0.009	1088
	[0.495]	[0.496]	(0.032)	
Has Bank Account	0.219	0.202	0.017	1088
	[0.414]	[0.402]	(0.027)	
Has a SACCO Account	0.041	0.031	0.009	1086
	[0.197]	[0.174]	(0.011)	
Saves at Home	0.890	0.858	0.031	1087
	[0.314]	[0.349]	(0.020)	
Has Mobile Money Account	0.200	0.218	-0.018	918
	[0.400]	[0.413]	(0.029)	
Saves Other Ways	0.624	0.529	0.095***	918
	[0.485]	[0.500]	(0.034)	
Total Reported Savings	14644	11410	3234	851
	[54989]	[26733]	(3000)	
Weekly Discount Factor	0.844	0.577	0.268***	1088
	[0.280]	[0.397]	(0.020)	
Impatient Now-Patient Later	0.209	0.215	-0.006	1070
	[0.407]	[0.411]	(0.025)	
Patient Now-Impatient Later	0.295	0.281	0.014	1070
	[0.456]	[0.450]	(0.029)	
Consumption - Husband Decides	0.389	0.444	-0.055*	1083
	[0.488]	[0.497]	(0.031)	
Consumption - Wife Decides	0.137	0.101	0.036*	1083
	[0.344]	[0.302]	(0.020)	
Consumption - Both Decide	0.419	0.413	0.006	1083
	[0.494]	[0.493]	(0.031)	
Saving - Husband Mostly Saves	0.323	0.322	0.001	1082
	[0.468]	[0.468]	(0.029)	
Saving - Wife Mostly Saves	0.474	0.450	0.024	1082
	[0.500]	[0.498]	(0.033)	
Saving - Both Save	0.177	0.204	-0.027	1082
	[0.382]	[0.403]	(0.026)	
Decision Error: Standard Deviation	19.7	13.9	5.84***	1088
	[37.8]	[33.7]	(2.20)	
Spouses Disagree: Consumption	0.590	0.583	0.007	1078
	[0.492]	[0.494]	(0.042)	
Spouses Disagree: Saving	0.519	0.619	-0.101***	1076
	[0.500]	[0.486]	(0.043)	
Poorly Informed Couple	0.506	0.523	-0.016	944
	[0.500]	[0.500]	(0.046)	

Notes: Standard deviations in brackets, heteroskedasticity robust standard errors in parentheses. Mobile money and other savings data not available for the 84 couples in the first 6 experimental sessions. Variables are recoded to missing if response was don't know/refused. ***, **, and * indicate significance at the 99, 95, and 90 percent confidence levels respectively.

Table 2. Balance Check

	Treatment					
	Excess Interest Rate			Extra	Cash	
	Husband	Wife	Joint	Statements	Payment	N
Marital Status Confirmed	0.001	-0.001	-0.001	-0.002	0.010	1196
	(0.003)	(0.003)	(0.003)	(0.016)	(0.012)	
Age	0.185	0.068	0.313	-0.684	0.946	1088
	(0.198)	(0.203)	(0.213)	(1.13)	(1.03)	
Education	-0.035	-0.090	-0.116**	0.051	-0.232	1084
	(0.052)	(0.055)	(0.059)	(0.313)	(0.313)	
Literate	-0.006	-0.006	-0.012**	0.003	-0.016	1088
	(0.005)	(0.005)	(0.005)	(0.030)	(0.035)	
Number Children	0.047	0.028	0.078*	0.143	0.221	1088
	(0.038)	(0.041)	(0.045)	(0.233)	(0.231)	
Subsistence Farmer/No Job	0.007	0.010	0.016**	-0.024	-0.018	1084
	(0.006)	(0.007)	(0.007)	(0.037)	(0.040)	
Distance from Bank (Miles)	0.076**	0.051	0.092**	-0.038	0.086	1088
	(0.034)	(0.035)	(0.039)	(0.211)	(0.154)	
Income Last Week	-21.7	-26.5	-24.7	48.1	-433***	1057
	(30.5)	(26.2)	(24.3)	(163)	(121)	
Owns Mobile Phone	0.002	0.002	0.003	0.073*	0.041	1083
	(0.006)	(0.007)	(0.007)	(0.038)	(0.039)	
Participates in ROSCA	-0.002	0.004	0.006	0.005	0.024	1088
	(0.006)	(0.006)	(0.007)	(0.035)	(0.039)	
Has Bank Account	0.000	-0.006	-0.004	-0.022	0.021	1088
	(0.005)	(0.005)	(0.006)	(0.029)	(0.034)	
Has a SACCO Account	-0.001	-0.003	-0.003	-0.006	0.001	1086
	(0.002)	(0.002)	(0.002)	(0.013)	(0.015)	
Saves at Home	0.001	0.005	-0.002	0.005	0.047**	1087
	(0.003)	(0.004)	(0.004)	(0.022)	(0.024)	
Has Mobile Money Account	-0.005	-0.004	-0.012*	-0.012	-0.013	918
	(0.005)	(0.006)	(0.006)	(0.029)	(0.034)	
Saves Other Ways	0.011*	0.008	0.002	0.079***	-0.034	918
	(0.006)	(0.007)	(0.007)	(0.034)	(0.043)	
Total Reported Savings	423	-44.0	-315	835	5432	851
	(441)	(514)	(423)	(3277)	(7348)	
Weekly Discount Factor	-0.002	-0.003	-0.006	0.064***	-0.053*	1088
	(0.004)	(0.005)	(0.005)	(0.024)	(0.032)	
Impatient Now-Patient Later	-0.002	0.000	0.001	-0.024	-0.062**	1070
	(0.005)	(0.004)	(0.005)	(0.027)	(0.030)	
Patient Now-Impatient Later	0.002	0.002	-0.002	0.050	0.018	1070
	(0.005)	(0.005)	(0.006)	(0.030)	(0.037)	
Well Matched Couple	-0.004	-0.012	-0.011	-0.014	0.080**	1088
	(0.008)	(0.008)	(0.009)	(0.047)	(0.037)	
Consumption - Husband Decides	-0.002	-0.009	-0.004	-0.016	0.010	1083
	(0.006)	(0.006)	(0.007)	(0.033)	(0.040)	
Consumption - Wife Decides	0.001	0.003	-0.001	-0.024	0.014	1083
	(0.004)	(0.004)	(0.004)	(0.022)	(0.026)	
Consumption - Both Decide	-0.001	0.004	0.005	0.016	-0.020	1083
	(0.006)	(0.006)	(0.007)	(0.034)	(0.039)	
Saving - Husband Mostly Saves	-0.001	0.000	-0.001	-0.041	-0.038	1082
	(0.005)	(0.005)	(0.006)	(0.032)	(0.037)	
Saving - Wife Mostly Saves	-0.002	-0.008	-0.008	0.028	0.034	1082
	(0.006)	(0.006)	(0.007)	(0.035)	(0.041)	
Saving - Both Save	0.003	0.008	0.009*	0.009	-0.003	1082
	(0.005)	(0.005)	(0.005)	(0.029)	(0.033)	
Decision Error: Standard Deviation	-0.199	-0.379	-0.572	0.278	1.71	1088
	(0.362)	(0.418)	(0.455)	(2.46)	(3.06)	
Spouses Disagree: Consumption	-0.016**	-0.014*	-0.008	-0.036	0.019	1078
	(0.008)	(0.008)	(0.009)	(0.046)	(0.038)	
Spouses Disagree: Saving	0.007	0.008	0.011	-0.061	0.070*	1076
	(0.008)	(0.008)	(0.009)	(0.047)	(0.038)	
Poorly Informed Couple	0.000	0.000	0.005	-0.058	-0.095**	944
	(0.008)	(0.009)	(0.009)	(0.050)	(0.041)	
P-value: Joint Test (Cross Equation)	0.654	0.227	0.394	0.250	0.001***	

Notes: Standard errors clustered at the couple level are in parentheses. Each row presents the results of a regression of observable characteristics on all listed treatments of interest (rows correspond to a single regression). All regressions also include a dummy variable identifying couples ineligible for the extra statements treatment. Excess interest rate variables range from -10 to 10. Mobile money and other savings data not available for the 84 couples in the first 6 experimental sessions. Variables are recoded to missing if response was don't know/refused. P-values from the joint test are calculated by jointly estimating equations by seemingly unrelated regression. ***, **, and * indicate significance at the 99, 95, and 90 percent confidence levels respectively.

Table 3. Summary of Bank Account Use

	Well Matched	Poorly Matched	Difference	N
Couple Chose to Open:				
All Three Accounts	0.051	0.055	-0.004	544
	[0.221]	[0.229]	(0.019)	
Joint Account Only	0.574	0.592	-0.018	544
	[0.495]	[0.492]	(0.042)	
Both Individual Accounts	0.261	0.257	0.004	544
	[0.440]	[0.438]	(0.038)	
One Individual, One Joint Account	0.088	0.066	0.022	544
	[0.284]	[0.249]	(0.023)	
One Individual Account	0.026	0.029	-0.004	544
	[0.159]	[0.169]	(0.014)	
Couple Saved In:				
Any Account	0.423	0.449	-0.026	544
	[0.495]	[0.498]	(0.043)	
Joint Account	0.320	0.298	0.022	544
	[0.467]	[0.458]	(0.040)	
Individual Account	0.114	0.169	-0.055*	544
	[0.318]	[0.376]	(0.030)	
Excluding Cash Payments, Couple Saved In:				
Any Account	0.257	0.287	-0.029	544
	[0.438]	[0.453]	(0.038)	
Joint Account	0.176	0.180	-0.004	544
	[0.382]	[0.385]	(0.033)	
Individual Account	0.092	0.125	-0.033	544
	[0.289]	[0.331]	(0.027)	
If Saved, Average Daily Balance In:				
All Accounts	1024	1051	-26.8	237
	[2160]	[1841]	(261)	
Joint Account	731	912	-181	168
	[1757]	[1702]	(267)	
Individual Accounts	1748	1182	566	77
	[2803]	[1772]	(566)	
If Saved, Number Deposits In:				
All Accounts	2.49	3.02	-0.538	237
	[2.82]	[3.87]	(0.438)	
Joint Account	2.16	2.65	-0.493	168
	[2.46]	[3.99]	(0.516)	
Individual Accounts	3.16	3.35	-0.187	77
	[3.29]	[3.18]	(0.754)	

Notes: Standard deviations are in brackets, heteroskedasticity robust standard errors in parentheses. The first two columns show means for well and poorly matched couples respectively. Column 3 shows the difference between well and poorly matched couples. ***, **, and * indicate significance at the 99, 95, and 90 percent confidence levels respectively.

Table 4. Preference Heterogeneity and Use of Dominated Individual Accounts

Panel A. All Couples With Dominated Individual Accounts					
Poorly Matched	0.106***	0.145***	0.127***	0.142***	0.146***
	(0.036)	(0.039)	(0.042)	(0.046)	(0.048)
Poorly Informed					0.093***
					(0.040)
Decision Error: Husband					0.000
					(0.001)
Decision Error: Wife					0.000
					(0.001)
Spouses Disagree: Consumption					-0.011
					(0.045)
Spouses Disagree: Saving					-0.015
					(0.043)
DV Mean (Well Matched)	0.079	0.079	0.079	0.079	0.079
N	331	331	331	331	331
Panel B. Subset of Couples Who Saved in at Least One Account					
Poorly Matched	0.221***	0.299***	0.315***	0.321***	0.274*
	(0.077)	(0.093)	(0.122)	(0.132)	(0.165)
Poorly Informed					0.199
					(0.171)
Decision Error: Husband					0.000
					(0.001)
Decision Error: Wife					0.003
					(0.002)
Spouses Disagree: Consumption					-0.113
					(0.137)
Spouses Disagree: Saving					-0.028
					(0.120)
DV Mean (Well Matched)	0.181	0.181	0.181	0.181	0.181
N	147	147	147	147	147
Control Set	Basic	+Time Pref	+Demo.	+Economic	+Decisions

Notes: The sample is limited to the subset of couples for whom the joint account bears the highest interest rate. The dependent variable is a dummy indicating that a couple saved in any individual account. Heteroskedasticity robust standard errors in parentheses. Basic controls include dummy variables for each account's interest rate and dummies for husband and wife cash payment selection, free ATM status for each of the three accounts a couple could open, and extra statement selection. Time preference controls include separate dummies for upper/lower censoring of the discount factors of each spouse, the estimated discount factor of each spouse, and patient now-impatient later and impatient now-patient later dummies for each spouse. The demographic control set adds controls for village fixed effects, spousal heterogeneity in age, education, number of children, and literacy. The economic control set adds controls for heterogeneity in income, for being a subsistence farmer or unemployed, and mobile phone ownership. The decision making controls set includes controls for both spouses' self-reports of consumption decision-making, the estimated time preference decision error of each spouse, dummies for whether a couple disagrees about consumption and savings decision making, and a dummy identifying poorly informed couples. When controlling for intra-couple heterogeneity in a characteristic I include separate controls for the husband's and wife's value, the interaction between these values and (when the characteristic is not binary) the square of the husband's and wife's value. ***, **, and * indicate significance at the 99, 95, and 90 percent confidence levels respectively.

Table 5. Responses to the Excess Interest Rate by Match Quality

δ_1: Well Matched×Excess Yes	0.013*	0.016***	0.017**	0.016*	0.016*
	(0.007)	(0.006)	(0.007)	(0.008)	(0.008)
δ_2: Well Matched×Excess No	0.005	0.004	0.007	0.006	0.003
	(0.006)	(0.006)	(0.006)	(0.007)	(0.008)
γ_1: Poorly Matched×Excess Yes	0.000	0.005	0.007	0.001	-0.002
	(0.007)	(0.009)	(0.012)	(0.013)	(0.014)
γ_2: Poorly Matched×Excess No	-0.008	-0.011	-0.009	-0.003	-0.005
	(0.006)	(0.008)	(0.010)	(0.012)	(0.013)
β_1: Poorly Matched	-0.035	-0.024	-0.029	-0.001	0.014
	(0.064)	(0.073)	(0.088)	(0.105)	(0.118)
P-value from F-test: δ_1=γ_1	0.174	0.340	0.453	0.322	0.237
P-value from F-test: δ_2=γ_2	0.126	0.142	0.211	0.534	0.610
P-value from F-test: δ_1=γ_1 & δ_2=γ_2	0.048**	0.124	0.250	0.416	0.371
DV Mean (Well Matched)	0.163	0.163	0.163	0.163	0.163
N	1632	1632	1632	1632	1632
Control Set	Basic	+Time Pref	+Demo.	+Economic	+Decisions

Notes: Data is at the account level. The dependent variable is a dummy indicating that an account received a savings deposit. Robust standard errors clustered at the couple level in parentheses. Basic controls include a dummy identifying joint accounts, dummies for the account's interest rate, and dummies for free ATM selection, extra statements selection, and husband and wife cash payment selection. The joint account and interest rate dummies are also interacted with the poorly matched dummy. See notes to Table 4 for the time preference, demographic, economic, and decision making controls sets. All time preference, demographic, economic, and decision making controls are demeaned to the value among well matched couples and interacted with the two excess interest rate splines and the interest rate and account type dummy variables. ***, **, and * indicate significance at the 99, 95, and 90 percent confidence levels respectively.

Table 6. 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	8.13	6.45	6.04	6.04	6.73	6.19	6.08
Actual Interest Earnings	7.11	5.37	4.33	3.52	5.74	4.88	4.25
Loss	1.03	1.08	1.72	2.52	0.990	1.31	1.83
Well Matched Couples							
Maximum Interest Earnings	8.19	6.44	5.96	5.96	6.76	6.14	5.99
Actual Interest Earnings	7.51	5.88	4.98	4.42	6.28	5.51	5.08
Loss	0.676	0.557	0.975	1.54	0.477	0.631	0.911
Loss Gap							
A. No Controls	0.350**	0.527***	0.742**	0.989**	0.513***	0.680***	0.916***
	(0.178)	(0.209)	(0.324)	(0.457)	(0.188)	(0.256)	(0.342)
B. + Basic Controls	0.421***	0.598***	0.873***	1.18***	0.574***	0.781***	1.06***
	(0.168)	(0.208)	(0.320)	(0.451)	(0.187)	(0.254)	(0.340)
C. + Time Preference Controls	0.432*	0.754***	1.23***	1.69***	0.653***	1.01***	1.41***
	(0.231)	(0.233)	(0.348)	(0.492)	(0.222)	(0.282)	(0.374)
D. + Demographic Controls	0.442*	0.775***	1.29***	1.81***	0.669***	1.05***	1.50***
	(0.232)	(0.243)	(0.366)	(0.518)	(0.229)	(0.296)	(0.392)
E. + Economic Controls	0.438*	0.781***	1.29***	1.82***	0.679***	1.05***	1.49***
	(0.238)	(0.248)	(0.375)	(0.529)	(0.233)	(0.303)	(0.402)
E. + Decision Making Controls	0.474*	0.790***	1.27***	1.79***	0.690***	1.04***	1.47***
	(0.244)	(0.255)	(0.384)	(0.539)	(0.240)	(0.312)	(0.412)
N	544	544	544	544	544	544	544

Notes: Heteroskedasticity robust standard errors in parentheses. See notes to Table 4 for the basic, time preference, demographic, and economic controls sets. For columns 2-4 I proxy banking costs by extracting the first principal component of distance from the bank, spouse-specific indicators for subsistence farmers/the unemployed, and spouse-specific indicators for baseline bank account ownership and SACCO membership. This index is re-normalized to range from 0 (lowest proxied costs) to 1 (highest proxied costs). Individual accounts are then discounted by the product of this index and the maximum individual discount specified on the table. For columns 5-7 all individual accounts are discounted by the same maximum discount listed on the table. ***, **, and * indicate significance at the 99, 95, and 90 percent confidence levels respectively.

Figure 1. Preference Heterogeneity and Equilibrium Savings Strategies

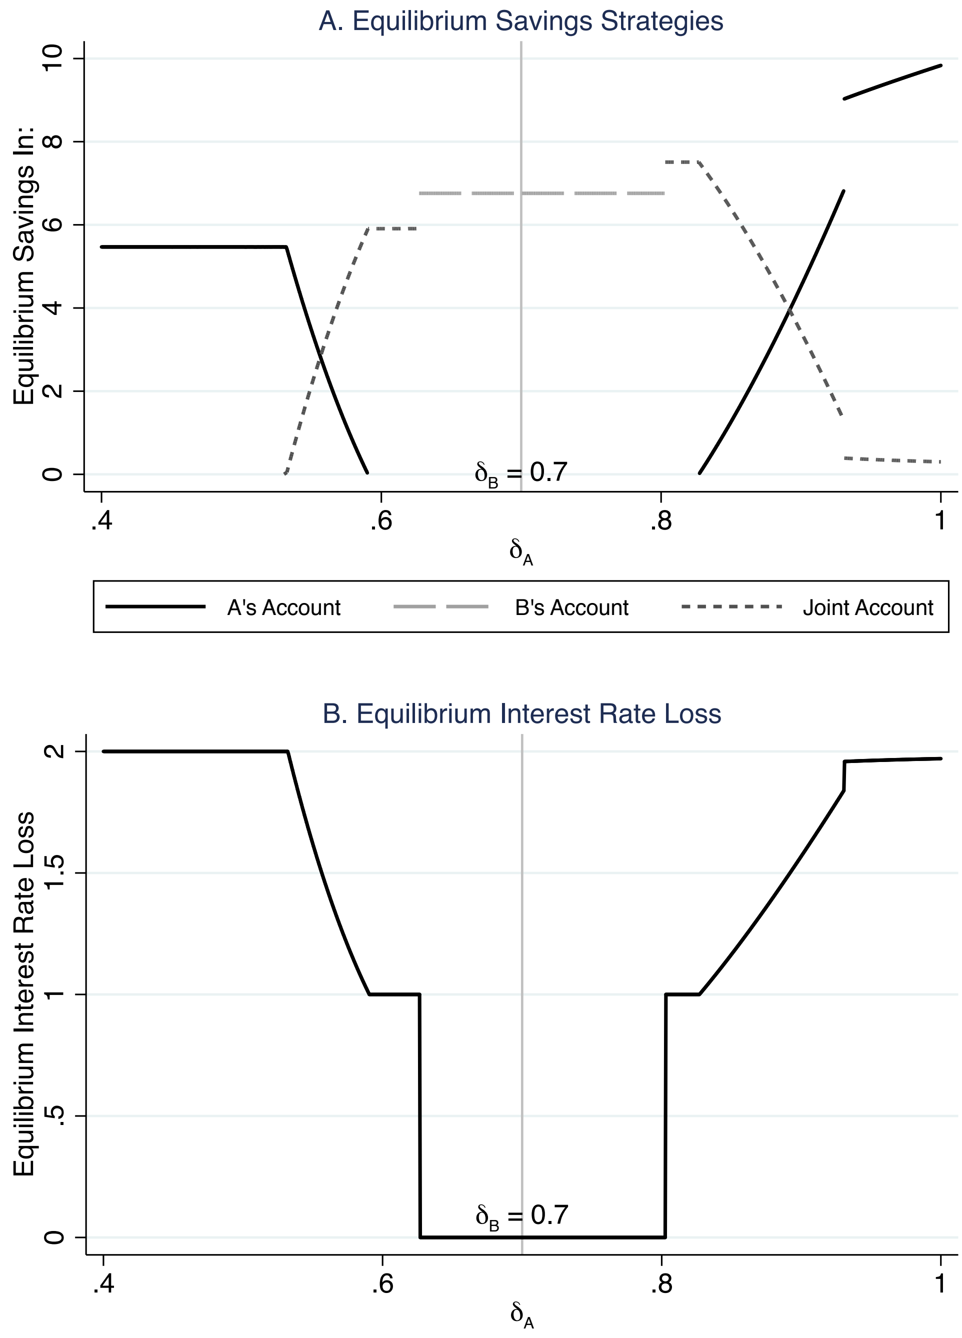

Notes: This figure provides a numerical example of how equilibrium savings strategies (Panel A) and the interest rate loss (Panel B) change with preference heterogeneity. In this example, y_1=30, y_2=10, b=0.05, R_A=1.03, R_B=1.05, R_J=1.04, δ_B=0.7, and the utility of per-period consumtion is u(c)=ln(c).

Figure 2. Distribution of Estimated Discount Factors and Discount Factor Heterogeneity

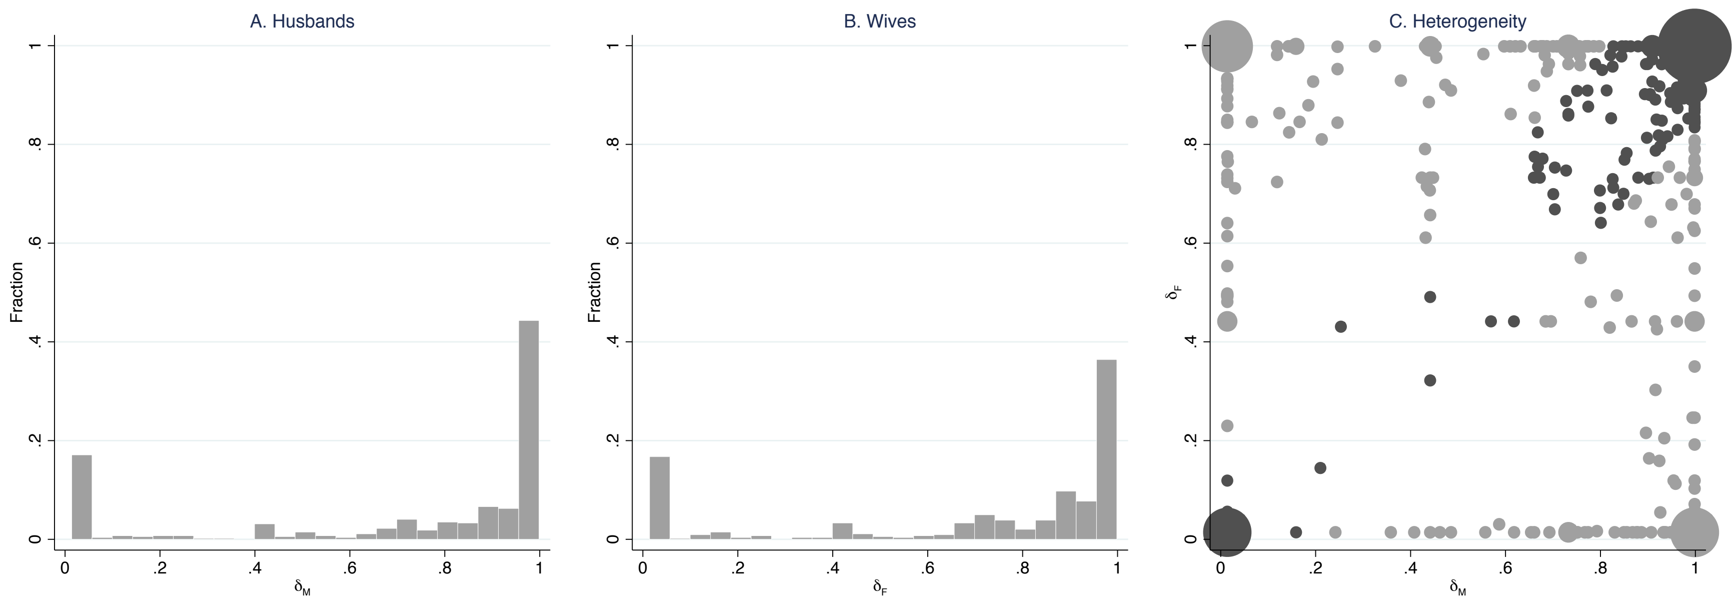

Notes: Data for N=544 couples. Panels A and B are histograms of husbands' and wives' discount factors respectively. Panel C is a weighted scatterplot of husbands' discount factors (x axis) and wives' discount factors (y axis). The size of each circle in Panel C is proportional to the number of couples with the relevant discount factor combination. Well matched couples in Panel C are demarcated with darker shading.

Figure 3. Heterogeneity in Estimated Discount Factors and Use of Dominated Individual Accounts

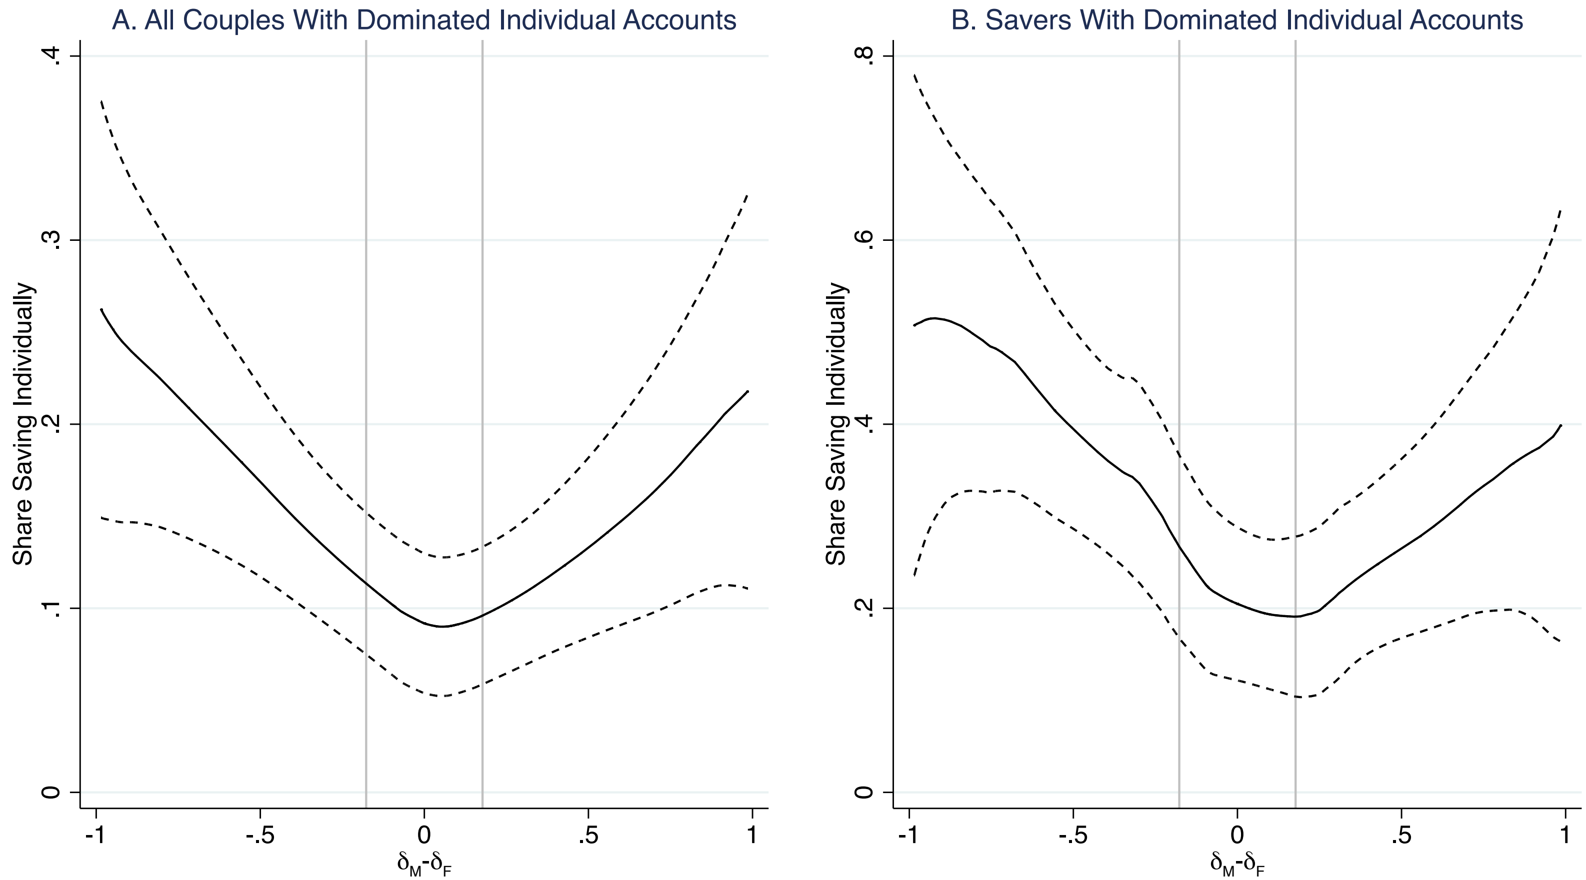

Notes: Local linear regression results. The sample is limited to the subset of couples for whom the joint account bears the highest interest rate. The dependent variable is a dummy indicating that a couple saved in any individual account. Dashed lines give 95 percent confidence intervals. Well matched couples are delineated by gray vertical lines. Sample size in Panel A is N=331 couples, sample size in Panel B is N=147 couples.

Figure 4. Efficient Responses to the Excess Interest Rate by Account Type

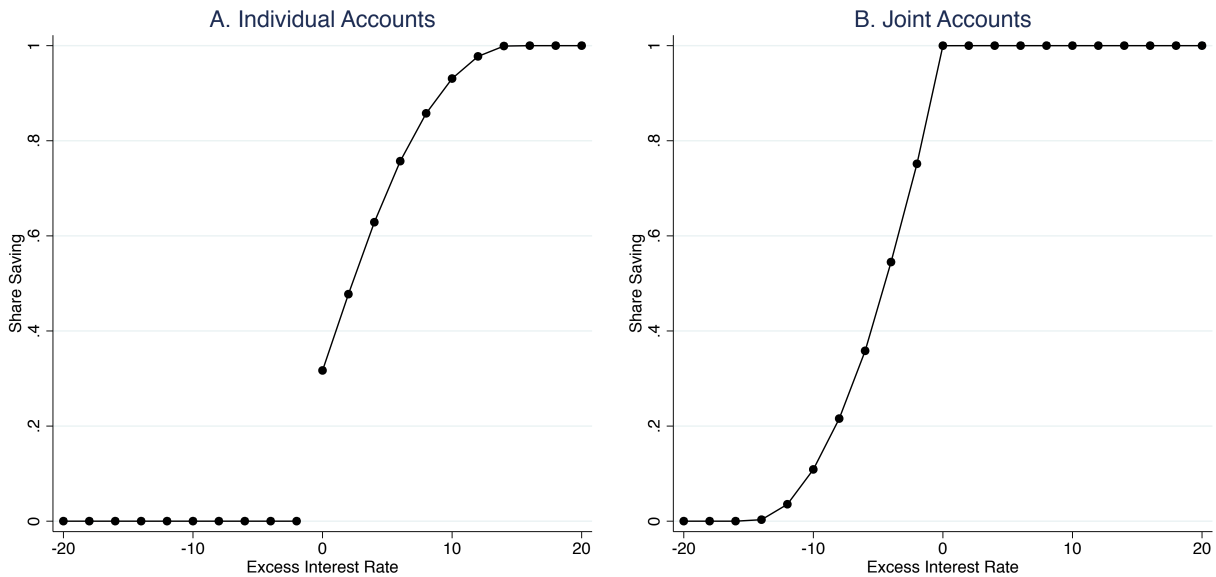

Notes: This figure use a simulated numerical example among N=100,000 couples to show how the savings rate changes with the excess interest rate, provided couples save efficiently. Here, interest rates are evenly distributed over the range 0-20 in 2 percentage point increments, $\delta_A=\delta_B=0.7$, $y_1=30$, $y_2=10$, and utility of per-period consumtion is u(c)=ln(c). Each couple receives 2 draws from a U[0.0.5] distribution - the minimum draw is the joint banking cost and the maximum draw is the individual banking cost. The first panel graphs the share of couples who save in agent A's bank account at each excess interest rate, the second panel graphs the share of couples who save in the joint account. In practice, the precise shape of the response will depend on a number of factors, including the distribution of banking costs in the population - the general result reflected in the figure is that savings rates will only increase over positive excess interest rates for individual accounts and negative excess interest rates for joint accounts.

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

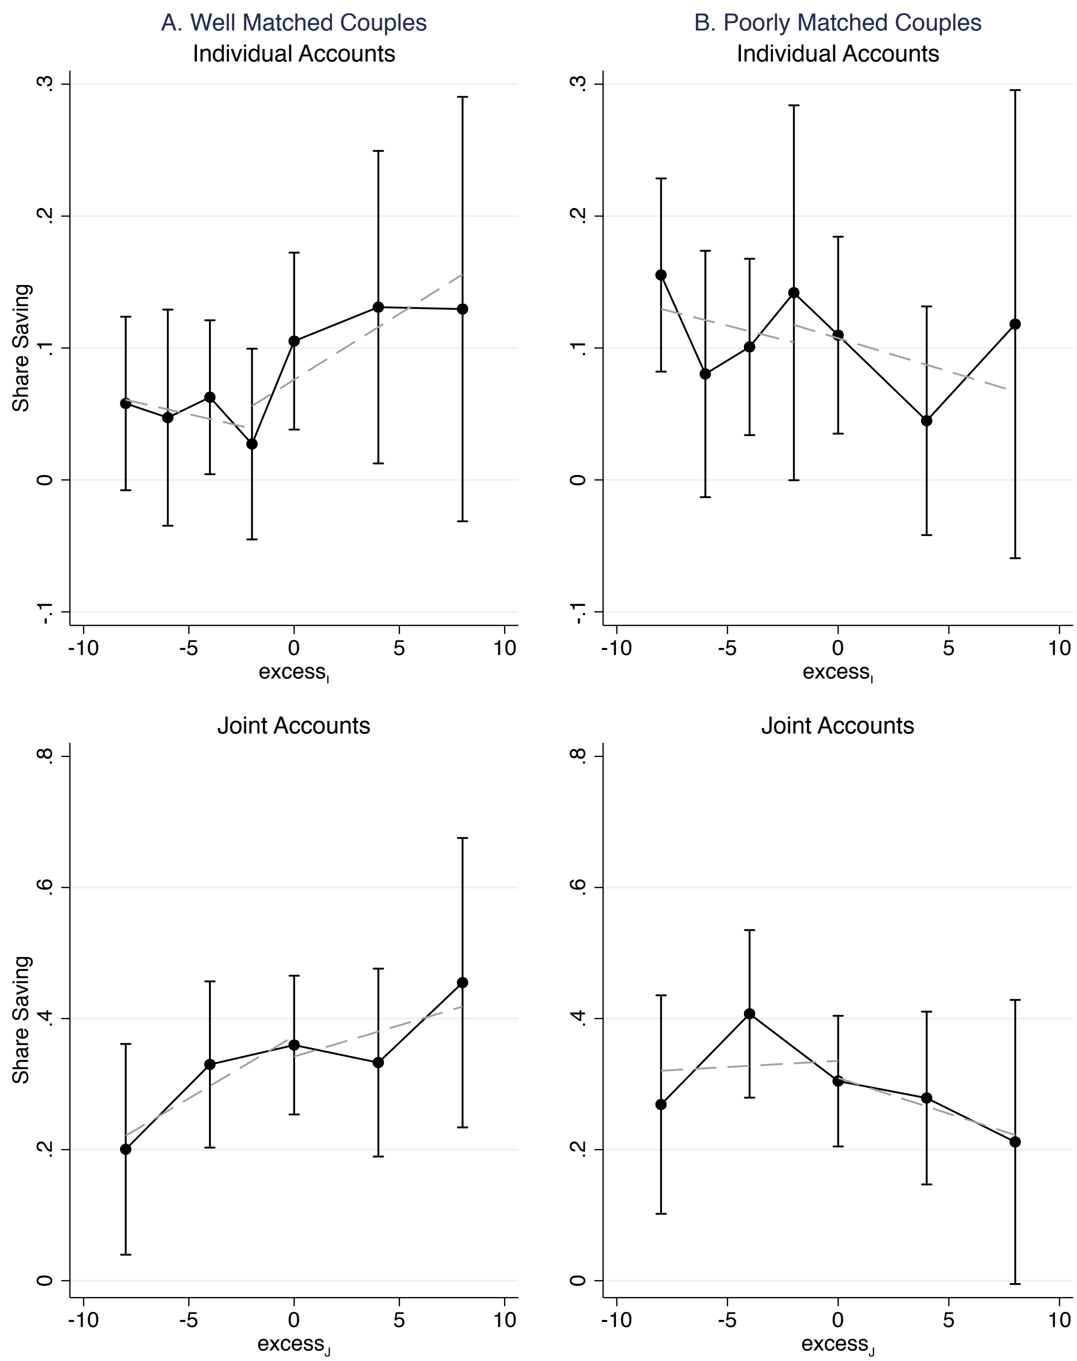

Note: Each panel plots predicted values and confidence intervals from account-level regressions of a savings dummy variable on a set of dummy variables for the excess interest rate and dummies for the level interest rate. Panel A limits the sample to all well matched couples and plots predicted values separately for all (open and unopened) individual accounts (top row, N=544) and all joint accounts (bottom row, N=272). Panel B repeats this exercise for poorly matched couples, plotting results for 544 individual and 292 joint accounts. Each predicted value assumes an equal distribution of accounts at each possible level interest rate. The 95 percent confidence intervals on predicted values are based on robust standard errors clustered at the couple level. The gray dotted lines show best fit lines for the predicted values, where each predicted value is weighted by one over its standard error.

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

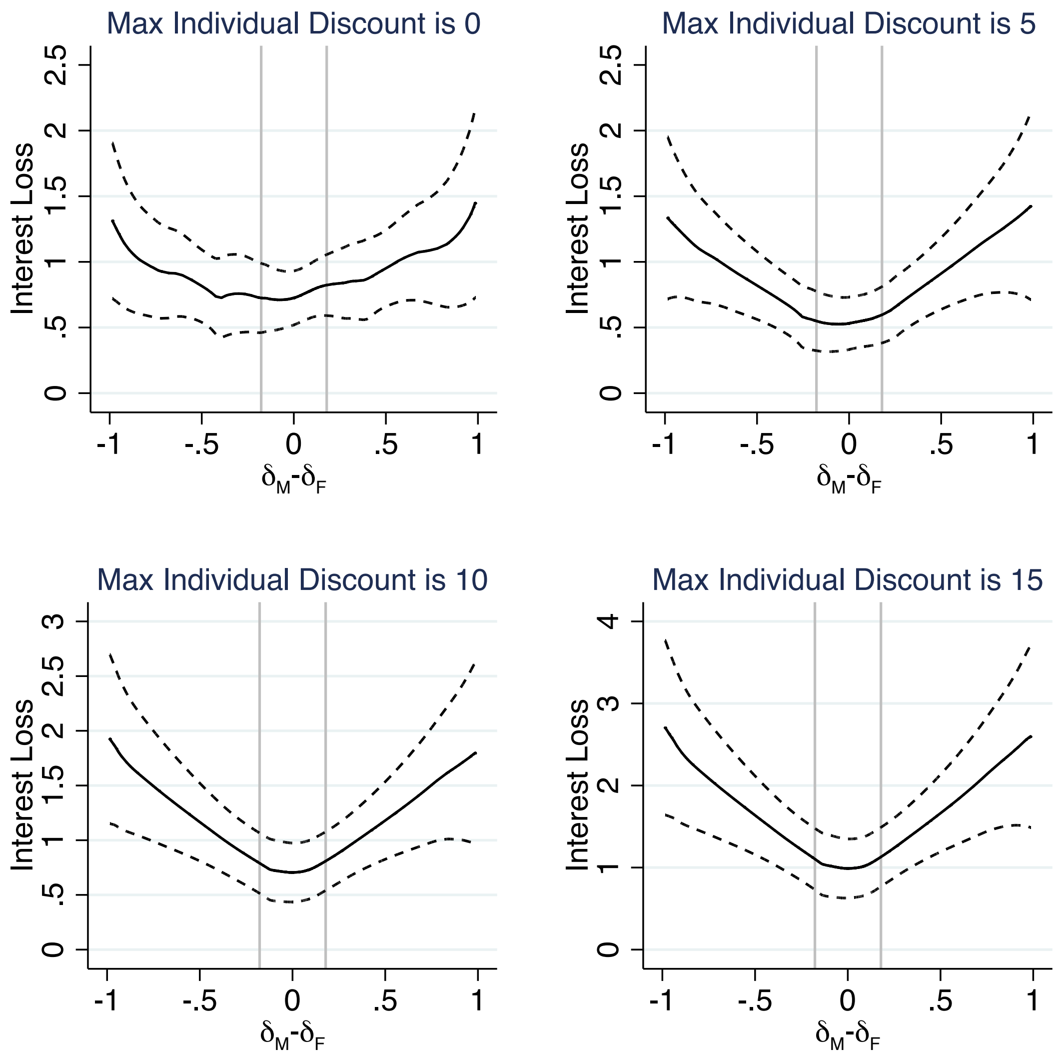

Notes: Local linear regression results for N=544 couples. The dependent variable is the interest rate loss. Dashed lines give 95 percent confidence intervals. Gray vertical lines demarcate well matched couples. Individual interest rate discounting is performed using proxied banking costs (see notes to Table 6 for additional detail).