

**Barriers to the Adoption and Optimal Use of Savings and Health Technologies**

by

Simone Gabrielle Schaner

A.B. Economics, Princeton University (2003)

Submitted to the Department of Economics  
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy


at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2011

© 2011 Simone Gabrielle Schaner. All rights reserved.

The author hereby grants to MIT permission to reproduce and distribute publicly paper and electronic copies of this thesis document in whole or in part in any medium now known or hereafter created.



**ARCHIVES**

Author .....

.....  
Department of Economics  
May 13, 2011

Certified by

.....  
Abhijit Banerjee  
and International Professor of Economics  
Thesis Supervisor

Certified by

.....  
Abdul Latif Jameel Professor of Poverty Alleviation and Development Economics  
Esther Duflo  
Thesis Supervisor

Certified by..

.....  
Assistant Professor of Applied Economics, MIT Sloan School of Management  
Tavneet Suri  
Thesis Supervisor

Accepted by

.....  
Esther Duflo  
Abdul Latif Jameel Professor of Poverty Alleviation and Development Economics  
Chairman, Departmental Committee on Graduate Studies

# Barriers to the Adoption and Optimal Use of Savings and Health Technologies

by

Simone Gabrielle Schaner

Submitted to the Department of Economics  
on May 13, 2011, in partial fulfillment of the  
requirements for the degree of  
Doctor of Philosophy

## Abstract

This thesis studies Kenyan households' use of savings accounts and malaria testing and treatment technologies.

The first chapter studies whether or not married couples use savings accounts strategically. In the absence of commitment, the availability of a "private" savings technology (a device that is only accessible by a single owner) may incite individuals to take costly strategic savings action in order to manipulate the time path of consumption. This chapter presents a model that formalizes this idea and derives several testable theoretical implications. In particular, households where husbands and wives are well matched in terms of time preference should make greater use of joint (public) accounts, less use of individual (private) accounts, and make more efficient investment choices as compared to their poorly matched peers. The model informed the design of a field experiment where married couples in rural Kenya were given the opportunity to open joint and individual bank accounts at randomly assigned interest rates. The behavior of individuals in the experiment is inconsistent with ex-ante Pareto efficiency and a variety of alternative models of intrahousehold resource allocation, but consistent with the proposed model of strategic savings. Savings misallocation due to strategic behavior may be substantial: in the experiment poorly matched couples forgo at least 64 percent more interest than well matched couples.

The second chapter studies the impact of reducing bank account transaction costs. Free ATM cards were offered to a randomly selected subset of newly opened formal bank accounts in Western Kenya. The ATM card reduced withdrawal fees by over 50 percent (from \$0.78 to \$0.38) and enabled account holders to make withdrawals from their accounts at any time of the day. The cards also enabled accounts to be accessed without the in-person verification of a national identity card. Targeting ATM cards to joint accounts and accounts owned by men substantially increased savings rates (by 39 percent) and average daily balances (by 16 percent) in the bank accounts. In contrast, the intervention had a negative impact when targeted to individual accounts owned by women. This gender difference appears to be driven by differences in bargaining power within the household: the positive treatment effect for men is concentrated in households where men have above median bargaining power, whereas the negative treatment effect for women is concentrated in households where women have below median bargaining power.

The final chapter (co-authored with Jessica Cohen and Pascaline Dupas) uses data from a randomized controlled trial conducted with over 2,900 households in rural Kenya to study the tradeoffs between the affordability of effective antimalarials (ACTs) and overuse. We compare a 95-percent ACT subsidy (currently under consideration by the global health community) to an alternative policy regime that explicitly acknowledges the problem of overuse by providing access to a subsidized rapid diagnostic test for malaria (RDT) in tandem with subsidized ACTs. We find that ACT access increases by 60 percent in the presence of an ACT subsidy of 80 percent or more. Under the proposed 95-percent ACT subsidy, however, only 56 percent of those buying an ACT at

the drug shop test positive for malaria. We show that targeting could be substantially increased (without compromising access) when the ACT subsidy is reduced to 80 percent but accompanied by an RDT subsidy.

Thesis Supervisor: Abhijit Banerjee

Title: Ford International Professor of Economics

Thesis Supervisor: Esther Duflo

Title: Abdul Latif Jameel Professor of Poverty Alleviation and Development Economics

Thesis Supervisor: Tavneet Suri

Title: Assistant Professor of Applied Economics, MIT Sloan School of Management

## Acknowledgments

I owe a debt of gratitude to many, and this list is sure to be incomplete. First, I thank my advisors, Abhijit Banerjee, Esther Duflo, and Tavneet Suri, for invaluable advice and feedback at all stages of my thesis research. They have taught me much, and it has been a great pleasure working with them. Over my graduate career I have benefited from the advice and tutelage of many more MIT faculty members – Ben Olken, Rob Townsend, and Josh Angrist deserve special mention. I also thank Garance Genicot and Dean Karlan, who nurtured my interest in development economics as an undergraduate, and my colleagues at the Urban Institute (especially Pam Loprest), who inspired me to get a Ph.D. in economics.

Identification issues aside, I can confidently say that the MIT peer effect is large, positive, and highly significant. My classmates have been great teachers, colleagues, and friends. Special thanks go to current and past office mates Dan Keniston, Emily Breza, Cynthia Kinnan, Nick Ryan, Rick Hornbeck, and Jeremy Shapiro.

My first two chapters 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. Andrey Vakhovskiy provided much appreciated support and advice during field activities.

It has been an utter pleasure working with Jessica Cohen and Pascaline Dupas on the third chapter of this thesis. They have been wonderful collaborators and have taught me a great deal about both field work and economics.

I gratefully acknowledge the financial support of the Russell Sage Foundation, the George and Obie Shultz Fund, and MIT's Jameel Poverty Action Lab (Chapters 1 and 2), the Clinton Foundation Health Access Initiative and Novartis Pharmaceuticals (Chapter 3), and the and the National Science Foundation's Graduate Research Fellowship.

Most of all, I thank my wonderful parents, Peter and Nancy Schaner – their unwavering love, support, and belief in me has sustained me not just through graduate school, but throughout my life.



# Contents

<b>1</b>	<b>Intrahousehold Preference Heterogeneity, Commitment, and Strategic Savings: Theory and Evidence from Kenya</b>	<b>9</b>
1.1	Introduction . . . . .	9
1.2	A Model of Strategic Savings . . . . .	11
1.2.1	Model Setup . . . . .	12
1.2.1.1	General Economic Environment . . . . .	12
1.2.1.2	Savings Technologies . . . . .	12
1.2.1.3	Preferences . . . . .	13
1.2.1.4	Timing and Strategic Actions . . . . .	14
1.2.1.5	Solution Concept . . . . .	14
1.2.2	Efficient Bargaining . . . . .	15
1.2.3	Incentives to Deviate from the Efficient Allocation . . . . .	16
1.2.4	The Strategic Solution . . . . .	18
1.2.4.1	Incorporating the Sharing Rule . . . . .	18
1.2.4.2	The Collective Allocation Problem . . . . .	18
1.2.4.3	Characterization of Optimal Strategies . . . . .	19
1.2.4.4	Preference Heterogeneity and Account Choice . . . . .	20
1.3	Experimental Design and Testable Implications . . . . .	25
1.3.1	Experimental Context . . . . .	25
1.3.2	Experimental Design . . . . .	25
1.3.2.1	Targeted Population . . . . .	25
1.3.2.2	Interventions . . . . .	26
1.3.3	Testable Implications . . . . .	28
1.3.4	Data . . . . .	31
1.3.4.1	Measuring Rates of Time Preference . . . . .	32
1.3.4.2	Sample Characteristics . . . . .	34
1.3.4.3	Randomization Verification . . . . .	35
1.4	Results . . . . .	36
1.4.1	Account Takeup and Interest Rate Responses . . . . .	36

1.4.2	Discount Factor Heterogeneity and Account Use . . . . .	37
1.4.3	Proxying Banking Costs . . . . .	42
1.4.4	Investment Efficiency by Match Quality . . . . .	44
1.5	Hidden Information and Account Use . . . . .	46
1.6	Alternative Explanations . . . . .	49
1.7	Conclusion . . . . .	50
1.A	Appendix: Proofs . . . . .	52
1.B	Appendix: Extra Statements Sample . . . . .	58
1.C	Appendix: Tables and Figures . . . . .	59
<b>2</b>	<b>Cost and Convenience: The Impact of ATM Card Provision on Formal Savings</b>	
	<b>Account Use in Kenya</b>	<b>81</b>
2.1	Introduction . . . . .	81
2.2	Theoretical Framework . . . . .	83
2.3	Experimental Design and Data . . . . .	86
2.3.1	Experimental Context . . . . .	86
2.3.2	Experimental Design . . . . .	86
2.3.2.1	Targeted Population . . . . .	86
2.3.2.2	Interest Rates . . . . .	87
2.3.2.3	ATM Cards . . . . .	88
2.3.3	Data . . . . .	88
2.3.4	Sample Characteristics and Randomization Verification . . . . .	89
2.4	Results . . . . .	91
2.4.1	Summary of Account Use . . . . .	91
2.4.2	Impact of ATM Card Provision . . . . .	93
2.4.3	Bargaining Power and ATM Card Treatment Effects . . . . .	98
2.5	Conclusion . . . . .	101
2.A	Appendix: Additional Time Preference Analysis . . . . .	103
2.A.1	Survey Questions on Rates of Time Preference . . . . .	103
2.A.2	Time Inconsistency and ATM Card Treatment Effects . . . . .	104
2.B	Appendix: Tables and Figures . . . . .	105
<b>3</b>	<b>Price Subsidies, Diagnostic Tests, and Targeting of Malaria Treatment: Evidence</b>	
	<b>from a Randomized Controlled Trial*</b>	<b>123</b>
3.1	Introduction . . . . .	123
3.2	Background . . . . .	128
3.2.1	Background on Malaria . . . . .	128
3.2.2	The Affordable Medicines Facility for Malaria (AMFm) . . . . .	129

---

\*This chapter is co-authored with Jessica Cohen and Pascaline Dupas.

3.2.3	Health Providers and Health Treatment Seeking in Rural Kenya . . . . .	129
3.3	Model . . . . .	131
3.3.1	Model Setup . . . . .	132
3.3.2	Impact of an ACT Subsidy at the Drug Shop . . . . .	132
3.3.3	Impact of Adding an RDT Subsidy at the Drug Shop . . . . .	134
3.4	Study Design and Data . . . . .	135
3.4.1	Experimental Design . . . . .	135
3.4.2	Data . . . . .	137
3.4.3	Characteristics of Study Sample . . . . .	138
3.5	Results . . . . .	139
3.5.1	Predicting Malaria Positivity . . . . .	139
3.5.2	Status Quo Treatment-Seeking Behavior . . . . .	140
3.5.3	The Impact of Subsidies on Treatment Seeking Behavior . . . . .	141
3.5.4	Within-Subsidy Price Variation, ACT Access, and ACT Targeting . . . . .	146
3.6	Cost Effectiveness . . . . .	149
3.6.1	Methodology . . . . .	149
3.6.2	Results . . . . .	150
3.7	Conclusion . . . . .	152
3.A	Appendix: Predicted Positivity and Regression Bias . . . . .	153
3.B	Appendix: Tables and Figures . . . . .	156

**Bibliography**

**176**



## Chapter 1

# Intrahousehold Preference Heterogeneity, Commitment, and Strategic Savings: Theory and Evidence from Kenya

### 1.1 Introduction

Informal and semi-formal savings devices abound in the developing world, even though they are generally characterized by high costs, illiquidity, and substantial risk (Rutherford 1999; Rutherford 2000).<sup>1</sup> Moreover, households often use a variety of such arrangements (Collins, Morduch, Rutherford, and Ruthven 2009), even though storing savings at home should be essentially costless in the absence of complications. As such, the popularity of these devices presents a puzzle – what constraints make these costly practices attractive? Recent research reveals three central themes: the need to protect savings from oneself (as in Ashraf, Karlan, and Yin 2006), the need to protect savings from appropriation by members of the community (as in Baland, Guirking, and Mali 2007), and the need to protect savings from other members of the household, especially one’s spouse (as in Anderson and Baland 2002). This final theme presents a particular challenge to traditional economic representations of the household – since members share a unified budget constraint and interact repeatedly, they should be able to contract with one another in a Pareto efficient manner, even when they have different preferences (Browning and Chiappori 1998). While a growing number of papers have documented evidence of households behaving in ways incompatible with Pareto effi-

---

<sup>1</sup>A stark example is that of deposit collectors. Deposit collectors regularly visit their clients to take savings deposits. They are free to do what they wish with the deposits while they store them, and they often charge fees for the service. Steel and Aryeetey (1994) document that in Ghana, these fees amount to an annual return of *negative* 54 percent. Another popular informal device is rotating savings and credit associations (ROSCAs). ROSCAs consist of a group of individuals who meet at predetermined intervals (e.g. weekly, monthly) to put a fixed amount of money into a common pot. At each meeting, a different member of the group receives the pot. ROSCAs are by nature illiquid and often quite risky, as group members can defect before the ROSCA cycle is completed.

ciency (see, for example, Ashraf 2009; de Mel, McKenzie, and Woodruff 2009; Duflo and Udry 2004; Robinson 2008; Udry 1996), the underlying *causes* of these inefficiencies remain poorly understood.

To make progress on this front, this paper abstracts from all but one potential driver of inefficient household savings behavior: intrahousehold heterogeneity in rates of time preference, coupled with an inability to commit to binding contracts. The idea here is that when one household member is very impatient, he may be tempted to spend any readily accessible savings, even if he promises not to do so. In this case, other household members may resort to saving in a "private" device (such as an individual bank account) that cannot be accessed by the less patient individual, even if using this device is very costly (because it offers a negative rate of return, for instance). To formally study this mechanism, we develop a model of a two person household where individuals have access to two classes of savings devices: a private savings device and a "public" savings device (such as a joint bank account), which can be accessed by any member of the household. These savings devices may also differ in terms of rates of return and transaction costs.

We show that when discount factors within the household differ, implementing the *ex-ante* Pareto efficient consumption allocation requires the ability to commit to binding intertemporal contracts. When individuals cannot commit, they may be tempted to make strategic use of private accounts in order to manipulate the time path of consumption. The model underscores that *both* agents in the household may privately benefit from saving strategically. The more patient agent would like to push additional consumption into the future, while the less patient agent would like to push additional consumption into the present – in both cases it may be possible to achieve these goals through the use of private accounts. Moreover, the model illustrates that impatient individuals may be willing to take costly action to deny the household higher return savings devices. In this context households may make intensive use of lower return, higher transaction cost savings devices, even when more attractive (from a rate of return perspective) alternatives are readily available.

This model informed the design of a field experiment, which was conducted in Western Kenya in the Summer of 2009. We gave 597 married couples the opportunity to open three savings accounts at randomly assigned interest rates: an individual account for the husband, an individual account for the wife, and a joint account. We also asked each respondent in the experiment a battery of questions designed to elicit discount factors, which are used to calculate measures of intrahousehold discount factor heterogeneity. By applying our theoretical results to the context of the experiment, we are able to derive a series of testable predictions of the private savings model. A central theoretical result is that couples who are well matched in terms of discount factors invest their resources efficiently, while poorly matched couples savings decisions are distorted by strategic action. A key feature of the experimental design is that it created random variation in relative rates of return, even conditional on an account's own interest rate. Our theory has sharp predictions regarding patterns of account use by match quality and the shape of well matched couples' response to relative rates of return, both of which we test in the data.

Consistent with our theory of strategic savings, we find that couples who are well matched in

terms of rates of time preference make more intensive use of joint accounts, less intensive use of individual accounts, and respond to relative rates of return in a manner consistent with investment efficiency. In contrast, poorly matched couples are completely insensitive to relative rates of return. These differences in behavior have financial consequences for poorly matched couples – interest rate losses on this group’s newly opened bank account deposits were 64 percent larger than those of their well matched peers. The empirical results also suggest that transaction costs play a very important role in determining bank account choice and use in our sample.

Overall, our results are consistent with our theory of strategic savings behavior and inconsistent with *ex-ante* Pareto efficient bargaining. The results do not appear to be driven by a correlation between discount factor heterogeneity and other characteristics of couples. Moreover, the patterns in the data are not consistent with other theories of household saving such as mental accounting or rules of thumb.

We also investigate an alternative theory that could generate the patterns we observe in the data: hidden savings. While our model presumes complete information, agents may use private accounts to systematically hide resources from other household members (as in Anderson and Baland 2002). We make use of a randomized information treatment that we implemented as part of the field experiment, as well as spousal cross reports of income and savings device use to examine the role that hidden information plays in household savings decisions. We find evidence that households with poor information flows are more likely to choose individual accounts, less likely to choose joint accounts, and more likely to reduce savings in response to the information treatment. However, these concerns are unrelated to our initial findings regarding preference heterogeneity; well matched couples have no better information flows than poorly matched couples, and the empirical results are unchanged when accounting for intrahousehold information sharing.

The remainder of the paper is structured as follows: Section 1.2 presents our model of strategic savings behavior, Section 1.3 outlines the experimental design and derives testable implications of the theory, Section 1.4 presents main results, Section 1.5 extends the analysis to account for hidden information, Section 1.6 discusses alternative explanations, and Section 1.7 concludes.

## 1.2 A Model of Strategic Savings

We are interested in understanding how heterogeneity in discount factors impacts individuals’ incentives to engage in costly savings behavior. To do so, we develop a strategic model of a household consisting of two agents with potentially differing discount factors, who must decide how much to consume and how much to save in a portfolio of public and private savings devices. In this section we set up the model and characterize the *ex-ante* Pareto efficient savings allocation. We then show that when discount factors differ, individuals will have incentives to deviate from this allocation. Given this observation, we then characterize the equilibrium of the strategic model and close the section by deriving comparative statics with respect to discount factor heterogeneity and interest

rates. Then in Section 1.3, we feed these results through the experimental design to derive testable predictions that can be taken to the data.

## 1.2.1 Model Setup

### 1.2.1.1 General Economic Environment

The household consists of a husband ( $M$ ) and a wife ( $F$ ). They live in a two period world with one personal consumption good,  $c$ . Individuals have deterministic income streams  $\{y_t^M, y_t^F\}_{t=1}^{T=2}$  and must decide whether or not to save any of their income for the second period. Though individuals can save for the future, there is no borrowing in this economy.<sup>2</sup> Agents have perfect information regarding own and spousal income streams, preferences, savings strategies, and rates of return earned on savings.

### 1.2.1.2 Savings Technologies

Households have access to three different savings technologies:

- A public/joint bank account, which yields rate of return  $R_J > 1$
- A private/individual husband's bank account, which yields rate of return  $R_M > 1$
- A private/individual wife's bank account, which yields rate of return  $R_F > 1$

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

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

---

<sup>2</sup>A perfect savings and credit market without transaction costs would eliminate all scope for strategic behavior. However, our model would generalize to an environment with an imperfect credit market – this would just put constraints on some types of strategic behavior.



### 1.2.1.3 Preferences

Both members of the household  $i \in \{M, F\}$  have CRRA preferences over the personal consumption good  $c_t^i$  (we note that the results would be unchanged if we generalized preferences to be a CES aggregate of a personal consumption good and a public, nonrival consumption good):

$$U_t^i = E_t \left[ \sum_{\tau=t}^T \delta_i^{\tau-t} \frac{(c_\tau^i)^{1-\sigma}}{1-\sigma} \right]$$

Without loss of generality, we will assume that the wife is more patient than the husband ( $\delta_F \geq \delta_M$ ) for the rest of the theoretical discussion.

The agents in the model act in self interested ways whenever possible. When an individual has proprietary access to a resource, we assume that he can make a unilateral decision regarding that resource – as a result, he will take the decision that maximizes his own utility without regard for spousal welfare. We refer to these choices as private decisions. In the context of the model, saving in individual accounts is a private decision, as resources stored in these accounts cannot be accessed without the consent of the owner.

However, some decisions in the household need to be made collectively (we refer to these as public decisions). If both members of the household can freely access a resource, then it cannot be distributed unless both spouses agree on the allocation. In order to reach a consensus, we assume that spouses bargain cooperatively with one another. We assume that the husband’s bargaining power can be represented by  $\mu \in (0, 1)$  and is a function of a variety of time invariant distributions factors (as in Browning and Chiappori 1998).<sup>3</sup>

Savings allocated to the joint account is a public decision. This is because either spouse can access the joint account at any time; in order for funds to be deposited and remain in the account, the deposit must be determined by consensus. Similarly, the distribution of consumption between husband and wife is a public decision. In other words, we assume that the majority of consumption is akin to food eaten at home – since any household member can put more food on his or her plate, the final allocation must be determined by collective agreement. Note that this holds regardless of the *source* of the resources used to finance consumption (current income vs. joint saving vs. individual saving). Even when consumption is financed out of private savings, the act of transforming financial resources into the consumable good makes the resources appropriable by both spouses and therefore subject to collective bargaining.

---

<sup>3</sup>One potential issue with our model is the assumption that distribution factors do not change over time. Indeed, the observation that distribution factors can shift unexpectedly has inspired a large body of empirical work (e.g. Angrist 2002; Bobonis 2009; Chiappori, Fortin, and Lacroix 2002; Duffo 2003; Lafortune 2010; Lundberg, Pollak, and Wales 1997). We could easily expand our framework to accommodate unexpected innovations in bargaining weights. A bigger issue is if the act of allocating savings to individual or joint accounts alters the bargaining weight (presumably by shifting outside options). We discuss whether the availability of such deviations could be driving our empirical results in Section 1.6.

This assumption is important, as it eliminates any scope for private accounts to be used to increase individual shares of aggregate per-period consumption. In practice, not all consumption choices are public decisions. However, while some consumption goods are undoubtedly best thought of as private (many "vice" goods such as alcohol and cigarettes have this property), these goods account for a small share of total expenditures of poor households in developing countries, particularly when compared to general food expenditures, which make up around two-thirds of all spending (Banerjee and Duflo 2007). In this context, private consumption concerns may be inframarginal to the savings motive and therefore ignorable from a modeling perspective. Moreover, ruling out private consumption motives allows us to focus on strategic action driven by differences in rates of time preference, which is the primary goal of our model. That said, private consumption deviations may well be important determinants of savings behavior, so we discuss whether such concerns could be generating our results in Section 1.6.

#### 1.2.1.4 Timing and Strategic Actions

Within a given period, the model proceeds as follows:

1. Incomes  $(y_t^M, y_t^F)$  and returns from any previous period's savings are realized.
2. The husband and the wife simultaneously make private savings decisions. Denote private savings by individual  $i \in \{M, F\}$  at time  $t$  as  $s_t^i$ . An individual cannot save more than  $y_t^i + R_i s_{t-1}^i - b^i$  in any period (resources he or she has proprietary access to, less the cost of going to the bank).
3. The husband and wife observe total resources available, as well as resources saved privately, and jointly decide how much to consume ( $c_t$ ). Any additional savings is placed in the joint account – denote this "household" saving as  $s_t^J$ . The spouses also decide how to apportion consumption between husband and wife subject to  $c_t^M + c_t^F = c_t$ .
4. Consumption takes place and the period ends.

#### 1.2.1.5 Solution Concept

We solve for subgame perfect Nash equilibria of the private savings game. The strategy set for spouse  $i$  at time  $t$  is  $S_t^i = [0, y_t^i + R_i s_{t-1}^i - b^i]$ . The strategy set for the couple at time  $t$  is given by  $S_t^J = [0, Y_t(\mathbf{s}_t^i) - b^J]$ , where  $Y_t(\mathbf{s}_t^i)$  denotes period  $t$  resources net of private savings. A strategy for actor  $a \in \{J, M, F\}$  is given by a probability distribution  $\varphi_t^a$  over  $S_t^a$ .

In some cases, the private savings game will have multiple equilibria. To refine the set of equilibria that we need to consider, we make the following assumptions:

- A1. Pure strategy equilibria will always be chosen over mixed strategy equilibria.

A2. If there exists more than one Nash equilibrium of a given type (pure/mixed), but one equilibrium Pareto dominates the others, then the couple will always choose the dominant equilibrium.

When there are no transaction costs and more than one account bears the highest rate of return, there will be some cases in which a continuum of pure strategy equilibria exists, with each involving saving the same aggregate amount in different accounts. This result is not very robust – as we will see later on, as long as there is an arbitrarily small transaction cost associated with each account, at most one account will be used in any implemented pure strategy Nash equilibrium (this is a result of imposing Assumption A2). To eliminate this knife edge multiple account case, we make the following assumption:

A3. There is always some (potentially arbitrarily small) transaction cost associated with banking: either  $w > 0$  or  $b^a > 0 \forall a \in \{J, M, F\}$ .

### 1.2.2 Efficient Bargaining

Before studying the strategic solution, we establish an efficient benchmark by which to measure the behavior of households in our model. Since this is a multiperiod setting, we adopt the standard of *ex-ante* Pareto efficiency. Any *ex-ante* Pareto efficient allocation can be captured by writing the optimization problem of a social planner who puts weight  $\eta$  on the husband's utility and weight  $(1 - \eta)$  on the wife's utility. By varying  $\eta$  over  $[0, 1]$ , we trace out the Pareto frontier. Here we set  $\eta$  equal to  $\mu$  (the husband's bargaining power in the collective allocation problem) – this solves for the allocation that would result if the husband and wife could write intertemporally binding contracts with one another. Then the planner's problem (or the "efficient bargaining" problem) is:

$$\max_{\{c_t^M, c_t^F, s_t^M, s_t^F, s_t^J\}_{t=1}^{T=2}} \mu \sum_{t=1}^{T=2} \delta_M^{t-1} \frac{(c_t^M)^{1-\sigma}}{1-\sigma} + (1-\mu) \sum_{t=1}^{T=2} \delta_F^{t-1} \frac{(c_t^F)^{1-\sigma}}{1-\sigma} \quad (\text{EB})$$

subject to

$$y_t^M + y_t^F + \sum_{a \in \{J, M, F\}} \max \{R_a s_{t-1}^a - w - b^a, 0\} \geq c_t^M + c_t^F + \sum_{a \in \{J, M, F\}} 1(s_t^a > 0)(s_t^a + b^a)$$

$$s_t^a \geq 0 \forall a \in \{J, M, F\}$$

where  $1(\cdot)$  is the indicator function. Due to nonzero transaction costs, this problem is not convex. However, we can imagine the planner solving for the optimal savings allocation *conditional* on paying each relevant combination of transaction costs, and then selecting the plan that generates the highest utility. Since each conditional problem is convex, first order conditions will be necessary for an interior optimum. Taking first order conditions with respect to  $c_t^M$  and  $c_t^F$  we see that

$$\frac{u'_M(c_t^M)}{u'_F(c_t^F)} = \left(\frac{c_t^M}{c_t^F}\right)^{-\sigma} = \frac{1-\mu}{\mu} \left(\frac{\delta_F}{\delta_M}\right)^{t-1} \quad (1.1)$$

so when  $\delta_F > \delta_M$ , the *ex-ante* Pareto efficient sharing rule necessitates that for a given  $c_t = c_t^M + c_t^F$ ,  $\frac{c_t^F}{c_t}$  increase over time.<sup>4</sup> We can use equation 1.1 to solve for the *ex-ante* Pareto efficient consumption sharing rule in each period:

$$c_t^M = \rho_t c_t \text{ and } c_t^F = (1 - \rho_t) c_t \text{ where } \rho_t \equiv \frac{(\mu \delta_M^{t-1})^{\frac{1}{\sigma}}}{(\mu \delta_M^{t-1})^{\frac{1}{\sigma}} + ((1 - \mu) \delta_F^{t-1})^{\frac{1}{\sigma}}}$$

Note that  $\rho_t$  monotonically increases with  $\mu$ , the husband's bargaining power. When  $\delta_M < \delta_F$ ,  $\rho_1 > \rho_2$ . In contrast, when  $\delta_M = \delta_F$ ,  $\rho_t$  is time invariant and equal to  $\frac{(\mu)^{\frac{1}{\sigma}}}{(\mu)^{\frac{1}{\sigma}} + (1 - \mu)^{\frac{1}{\sigma}}}$ .

### 1.2.3 Incentives to Deviate from the Efficient Allocation

It may be difficult to enforce the *ex-ante* efficient allocation when  $\rho_t$  evolves over time. In fact, when discount factors differ, there are incentives to deviate at both the private savings and collective bargaining stages of the game. First we show that individuals could make themselves better off by deviating from the efficient savings path. When  $s_1^a > 0$ , first order conditions from the efficient bargaining problem require that  $(c_1^i)^{-\sigma} = \delta_i R_a (c_2^i)^{-\sigma}$ . But when  $\delta_M < \delta_F$ , the wife's marginal utility of an  $\varepsilon$  increase in  $s_t^a$  is

$$\left[ - (c_1^F)^{-\sigma} (1 - \rho_1) + \delta_F R_a (c_2^F)^{-\sigma} (1 - \rho_2) \right] \varepsilon > 0$$

and the husband's marginal utility of an  $\varepsilon$  decrease in  $s_t^a$  is

$$\left[ (c_1^M)^{-\sigma} \rho_1 - \delta_M R_a (c_2^M)^{-\sigma} \rho_2 \right] \varepsilon > 0$$

where both inequalities use  $\rho_1 > \rho_2$ . In contrast, if  $\delta_M = \delta_F$  then  $\rho_1 = \rho_2$  and there are no individual incentives to deviate. A linear consumption sharing rule is essential for this result. If the sharing rule also depended on the *level* of consumption, there could be incentives to deviate in the absence of discount factor heterogeneity. We purposefully abstract away from this complication in order to focus on discount factors. However, it is important to note that assuming a linear sharing rule puts strong restrictions on individual utility functions: the per-period sharing rule will be linear if and only if  $\mu u_M(c_t^M) + (1 - \mu) u_F(c_t^F)$  is homothetic (indeed, we could easily rewrite our model with more general utility functions under this assumption).<sup>5</sup>

---

<sup>4</sup>This observation reflects broader issues associated with aggregating individual preferences with differing discount factors. In particular, aggregated preferences will be time inconsistent as long as positive weight is placed on at least two agents with  $\delta_j \neq \delta_i$  (Jackson and Yariv 2010). In our context, it is straightforward to show that when  $\delta_F > \delta_M$  the effective discount factor governing the *ex-ante* efficient allocation asymptotes to  $\delta_F$  as  $T \rightarrow \infty$ . Other studies that address this aggregation problem include Feldstein (1964), Marglin (1963), Caplin and Leahy (2004), Gollier and Zeckhauser (2005), Weitzman (2001), and Zuber (2010).

<sup>5</sup>This highlights an important point: in practice intrahousehold heterogeneity in discount factors may be correlated with heterogeneity in utility functions. As such, we will not be able to rule out that some part of the patterns we

Second, the couple may have difficulty enforcing the time varying consumption sharing rule. Imagine the spouses collectively decide upon a consumption path at  $t = 1$ . Then if they were given the opportunity to reoptimize at time  $\tau > 1$ , they would choose to do so whenever  $\delta_M \neq \delta_F$ , and would reallocate a larger share of time  $\tau$  consumption to the less patient spouse such that  $\rho_\tau = \rho_1$ .<sup>6</sup> Since only the less patient spouse stands to benefit from this reallocation, there may be some natural barriers to renegotiation (Ligon 2002); however, in the long run it seems implausible that households would be able to enforce allocations where one member gets an ever shrinking share of aggregate consumption. Indeed, the results in Duflo and Udry (2004), Mazzocco (2007), and Robinson (2008) all suggest that couples cannot commit intertemporally.

We therefore assume that all agents "live in the moment", in that if reoptimization is attractive in period  $t$ , they will reoptimize. Then the sharing rule will be governed by  $\rho = \rho_1$ . This assumption also implies that if an individual can make himself better off by deviating from the allocation that he would collectively choose with his spouse, he will do so. To see if this has bite, we solve (EB) for the optimal savings path *imposing* a time invariant, linear sharing rule. Denote  $c^M(c_t) = \rho c_t$  and  $c^F(c_t) = (1 - \rho) c_t$ . Then if the couple were to collectively choose  $s_1^a > 0$ , the following equality would be satisfied for individual  $i$

$$(c_1^i)^{-\sigma} = R_a \delta_i (c_2^i)^{-\sigma} + R_a (1 - c^i(c_2)) (\delta_{-i} - \delta_i) (c_2^i)^{-\sigma} \quad (1.2)$$

Note that when  $\delta_F > \delta_M$ , the remainder term on the right hand side is negative for the wife and positive for the husband, which implies that  $(c_1^F)^{-\sigma} < R_a \delta_F (c_2^F)^{-\sigma}$  and  $(c_1^M)^{-\sigma} > R_a \delta_M (c_2^M)^{-\sigma}$ . In this case, if individuals do not take strategic action and the couple saves, the collective outcome will leave the wife feeling savings constrained and the husband feeling borrowing constrained – both would like to alter the time path of consumption if possible. It therefore seems likely that households with discount factor heterogeneity will exhibit strategic behavior. The next subsection characterizes this behavior and derives comparative statics with respect to preference heterogeneity and interest rates.

---

observe in the data are driven by more general preference heterogeneity. Section 1.6 discusses whether it seems likely that other forms of preference heterogeneity are driving our results.

<sup>6</sup>This temptation to reoptimize reflects the fact that when discount factors differ, the household has time inconsistent collective preferences (Jackson and Yariv 2010). This type of temptation problem is similar to the internal temptation problems studied by a large literature where either time inconsistent preferences or differential preferences between different "selves" lead to distorted consumption and savings behavior (examples include Banerjee and Mullainathan 2010, Fudenberg and Levine 2006, Gul and Pesendorfer 2004, Harris and Laibson 2001, Laibson 1997, and O'Donoghue and Rabin 1999). We also note that heterogeneity in discount factors is not a necessary condition for a time inconsistent household: Hertzberg (2010) proposes a non-cooperative model of the household where two agents with identical exponential discount factors behave like a single time inconsistent agent.

## 1.2.4 The Strategic Solution

### 1.2.4.1 Incorporating the Sharing Rule

It is convenient to use the sharing rule to rewrite individual utility functions so that they are defined over aggregate per-period consumption  $c_t$ :

$$\begin{aligned}\tilde{U}_t^M &= \frac{U_t^M}{\rho^{1-\sigma}} = E_t \left[ \sum_{\tau=t}^{T=2} \delta_M^{\tau-t} \frac{c_\tau^{1-\sigma}}{1-\sigma} \right] \\ \tilde{U}_t^F &= \frac{U_t^F}{(1-\rho)^{1-\sigma}} = E_t \left[ \sum_{\tau=t}^{T=2} \delta_F^{\tau-t} \frac{c_\tau^{1-\sigma}}{1-\sigma} \right]\end{aligned}$$

Since the couple bargains cooperatively over joint savings, it will choose a Pareto efficient allocation (subject to the time invariant  $\rho$ ). This choice can be represented by the maximization of a "household" utility function

$$\tilde{U}_t^H = \frac{\mu U_t^M + (1-\mu) U_t^F}{\Omega_0} = E_t \left[ \sum_{\tau=t}^{T=2} \frac{\Omega_{\tau-t}}{\Omega_0} \frac{c_\tau^{1-\sigma}}{1-\sigma} \right]$$

where  $\Omega_t = \mu \rho^{1-\sigma} \delta_M^t + (1-\mu)(1-\rho)^{1-\sigma} \delta_F^t$ . Note that

$$\frac{\Omega_1}{\Omega_0} = \Omega = \frac{\mu \rho^{1-\sigma} \delta_M + (1-\mu)(1-\rho)^{1-\sigma} \delta_F}{\mu \rho^{1-\sigma} + (1-\mu)(1-\rho)^{1-\sigma}} = \rho \delta_M + (1-\rho) \delta_F$$

The "household" discount factor,  $\Omega$ , is just a weighted average of individual discount factors, where the weights are given by each individual's share of aggregate consumption (also recall that  $\rho$  is just a rescaling of the bargaining weight,  $\mu$ ).

### 1.2.4.2 The Collective Allocation Problem

We solve the model by working backwards. In the second (final) period, all parties would like to maximize consumption. Therefore agents optimally set  $s_2^{M*} = s_2^{F*} = s_2^{J*} = 0$ . Then the collective  $t = 1$  saving problem is given by:

$$\begin{aligned}\arg \max_{s_1^J} & \frac{\left( y_1 - \sum_{a \in \{J, M, F\}} 1(s_1^a > 0) (s_1^a + b_1^a) \right)^{1-\sigma}}{1-\sigma} + \\ & \Omega \frac{\left( y_2 + \sum_{a \in \{J, M, F\}} 1(s_1^a > 0) \max \{ R_a s_1^a - w - b_2^a, 0 \} \right)^{1-\sigma}}{1-\sigma} \\ & \text{subject to } s_1^J \geq 0\end{aligned}$$

where  $s_1^M$  and  $s_1^F$  are taken as given. Again, this problem is not convex due to transaction costs. As before, we can examine the convex subproblem by assuming that joint banking and withdrawal

costs are paid, solve for the optimal  $s_1^J$  conditional on the costs being paid, and compare the utility of this allocation to the utility of setting  $s_1^J = 0$ . Then if the couple saves in the joint account, the following household Euler condition holds:

$$c_1^{-\sigma} = \Omega R_J c_2^{-\sigma} \quad (1.3)$$

### 1.2.4.3 Characterization of Optimal Strategies

Both individuals know the cooperative outcome given any private savings allocation  $\mathbf{s}_1^i = [s_1^M, s_1^F]'$  and endowment  $\mathbf{y} = [y_1, y_2]'$ . Since we do not focus on changes in the endowment, we refer to the couple's optimal joint savings choice given a private savings allocation  $\mathbf{s}_1^i$  as  $s_1^J(\mathbf{s}_1^i)$ .<sup>7</sup> Before characterizing individual strategic behavior, it is useful to consider how each spouse can use his or her private account to manipulate consumption streams. The wife's goal is to increase consumption in the second period. She can achieve this by "oversaving" in her individual account – as long as she has sufficiently large  $y_1^F$ , this will be a viable strategy even when her individual account is dominated (in terms of rate of return) by the joint and/or husband's account.

In contrast, the husband's goal is to increase consumption in the first period. Here, his optimal strategy will depend on the context. Suppose that in the absence of strategic behavior, the couple would collectively choose to save in his account because he has access to the best interest rate. In order to increase first period consumption he could either save less than the desired collective amount in his account, even to the point of refusing to save at all. Even when his account is dominated by the joint and/or wife's account, he may still be able to manipulate consumption streams in his favor. For example, suppose the couple would collectively choose to save in the joint account if he took no action. Then the husband could rush to the bank and preemptively save "just enough" in his account to prevent the couple from travelling to the bank again to save jointly. If this "just enough" amount results in increased first period consumption, a sufficiently impatient husband will find this deviation profitable. Note that in this case, the presence of transaction costs is essential – in their absence, this type of deviation would always result in decreased consumption in both periods. This is an important insight, particularly given our focus on developing countries: transaction costs greatly expand the scope of private savings deviations available to the less patient spouse.

We begin our characterization of optimal strategies by showing that at most one bank account will be in use in any implemented pure strategy Nash equilibrium. This will give us a very simple way to determine which households will save privately and which will not. First, we establish the following lemma:

**Lemma 1** *Let  $\mathbf{s}_1^i = [s_1^M, s_1^F]'$  be a pure strategy private savings allocation. If  $\mathbf{s}_1^i$  is part of a pure strategy Nash equilibrium and either  $s_1^M > 0$  or  $s_1^F > 0$ , then  $s_1^J(\mathbf{s}_1^i) = 0$ .*

---

<sup>7</sup>When convenient, we will write  $s_1^J(\mathbf{s}_1^i)$  when considering only one private savings choice.

**Proof.** See Appendix A. ■

This lemma highlights that it is never optimal to save individually such that the couple travels back to the bank at the collective allocation stage. With this result in hand, we are prepared to show the following:

**Proposition 1** *No couple will choose a pure strategy Nash equilibrium to the private savings game where more than one account is in use. Moreover, when  $R_M\delta_M \neq R_F\delta_F$ , no such pure strategy equilibrium exists.*

**Proof.** See Appendix A. ■

Given this result, it is straightforward to determine which households will save privately. For each individual  $i$ , we need only check if he or she would find a pure strategy private savings choice profitable when  $s_1^{-i} = 0$ . If neither member does, Proposition 1 and Assumption A1 imply that households default to  $s_1^M = s_1^F = 0$ . If one or both members strictly prefer to save privately assuming their spouse does not, then any Nash equilibrium must involve the use of individual accounts with positive probability. If a pure strategy Nash equilibrium can be constructed where just one spouse saves, the household will choose this over any mixed strategy by Assumption A1. If such an equilibrium cannot be constructed, both individuals in the household will randomize over private savings choices.

#### 1.2.4.4 Preference Heterogeneity and Account Choice

We are now prepared to analyze how heterogeneity in rates of time preference impacts the efficiency of household investment choices. As a baseline, we show that perfectly matched households always invest their savings efficiently:

**Proposition 2** *If  $\delta_M = \delta_F$ , then a solution to the ex-ante Pareto Efficient planning problem (EB) will always be chosen by the household in the private savings game.*

**Proof.** See Appendix A. ■

To explore how preference heterogeneity impacts savings behavior, we fix incomes, the average household discount factor ( $\Omega$ ), banking costs ( $\mathbf{b}^a$ ), and interest rates. This lets us consider a subset of households who would all choose the same allocation in the absence of individual strategic behavior, and study how the chosen allocation changes with preference heterogeneity. The benevolent planner would also choose the same savings allocation for all these households if he were restricted to the time invariant consumption rule given by  $\rho$ . This allocation also corresponds to the outcome of the private savings game when  $\delta_M = \delta_F = \Omega$  (Proposition 2). This concordance is a useful result of choosing utility functions that result in a linear consumption sharing rule, and provides a natural benchmark for efficient behavior.<sup>8</sup>

---

<sup>8</sup>However, note that these households would *not* all choose the same allocation if they could implement time varying consumption rules. This is easily seen by comparing the solution to (EB) for a household where  $\delta_M = \delta_F = \Omega$



In order to derive comparative statics with respect to discount factor heterogeneity and account use, we need a way to index couples according to match quality. To do so, we define  $\gamma \geq 1$  such that  $\delta_M(\gamma) = \frac{\Omega}{\gamma}$  and  $\delta_F(\gamma; \rho) = \frac{\Omega}{\gamma} \left( \frac{\gamma - \rho}{1 - \rho} \right)$ . When  $\gamma = 1$ ,  $\delta_M = \delta_F = \Omega$  - a couple is perfectly matched. By varying  $\gamma$  and  $\rho$  conditional on  $\Omega$ ,  $\delta_F - \delta_M$  can be made arbitrarily large (note that  $\delta_F - \delta_M = \frac{\Omega\gamma(\gamma-1)}{1-\rho}$ , which is strictly increasing in  $\gamma$  when  $\gamma \geq 1$ ), all while ensuring that  $\delta_M \in (0, \Omega]$  and  $\delta_F \in \left[ \Omega, \frac{\Omega}{1-\rho} \right)$ .

To move forward, we need to put a bit more structure on the types of private savings choices that might be profitable for an individual. To do so, take all feasible private savings choices (the set  $S_1^i$ ), discard choices that lead to a strictly Pareto dominated consumption allocation (as compared to  $s_1^i = 0$ ) and assign the remaining choices to two sets, depending on how the choices change the time path of consumption relative to the "household alternative", which prevails when  $s_1^i = 0$ :

**Definition 1** Given  $s_1^{-i}$ , a *temporally advantaged* private savings choice for individual  $i$  is any  $s_1^i \in S_1^i$  that results in increased consumption in the first period when  $i = M$  or the second period when  $i = F$ , relative to the household alternative.

**Definition 2** Given  $s_1^{-i}$ , a *temporally disadvantaged* private savings choice for individual  $i$  is any  $s_1^i \in S_1^i$  that results in increased consumption in the second period when  $i = M$  or the first period when  $i = F$ , relative to the household alternative.

Note that temporally advantaged choices favor an individual's relatively preferred period ( $t = 2$  for the wife and  $t = 1$  for the husband). In contrast, temporally disadvantaged choices favor the opposite of the relatively preferred period. Denote individual  $i$ 's set of temporally advantaged private savings choices given  $s_1^{-i}$  as  $A_1^i(s_1^{-i})$  and the analogous set of temporally disadvantaged private savings choices as  $D_1^i(s_1^{-i})$ . When it does not cause confusion, we drop the dependence on  $s_1^{-i}$  and simply refer to  $A_1^i$  and  $D_1^i$ . When  $\hat{s}_1^i > 0$  results in the same consumption allocation as  $s_1^i = 0$ , we assign  $\hat{s}_1^i$  to both  $A_1^i$  and  $D_1^i$ . Given this setup, we are now prepared to characterize how the profitability of private savings changes with preference heterogeneity.

To do this, we establish "preference heterogeneity thresholds" for private savings. Specifically, we show that when temporally advantaged deviations are available, there exists a level of preference heterogeneity (indexed by  $\gamma$ ) *above* which couples will always save privately. We also show that there exists a separate threshold for temporally disadvantaged deviations *below* which couples will always save privately. The following proposition formalizes this result:

**Proposition 3** Fix  $\Omega$ ,  $\mathbf{y}$ ,  $\mathbf{b}^a$ , and interest rates. Then the following preference heterogeneity thresholds for private savings obtain:

---

to a household where  $\delta_M = 0$  and  $\delta_F = \frac{\Omega}{1-\rho}$ . The household with greater preference heterogeneity will save more, all else equal.

1. Suppose  $A_1^F(0)$  is nonempty. Then  $\exists \rho^* \in [0, 1)$  s.t.  $\forall \rho \in (\rho^*, 1) \exists \gamma_F^*(\rho) \in [1, \infty)$  s.t. all households with sharing rule  $\rho$  and  $\gamma > \gamma_F^*(\rho)$  will exhibit private savings in any Nash equilibrium.
2. Suppose  $A_1^M(0)$  is nonempty. Then  $\exists \gamma_M^* \in [1, \infty)$  s.t. all households with  $\gamma > \gamma_M^*$  will exhibit private savings in any Nash equilibrium.
3. Suppose  $D_1^F(0)$  is nonempty. Then  $\forall \rho \in (0, 1) \exists \gamma_F^{**}(\rho) \geq 1$  s.t. all households with  $\gamma < \gamma_F^{**}(\rho)$  will exhibit private savings in any Nash equilibrium. It may be that  $\gamma_F^{**}(\rho) = 1$ , in which case no element of  $D_1^F(0)$  will ever be profitable for the wife.
4. Suppose  $D_1^M(0)$  is nonempty. Then  $\exists \gamma_M^{**} \geq 1$  s.t. all households with  $\gamma < \gamma_M^{**}$  will exhibit private savings in any Nash equilibrium. It may be that  $\gamma_M^{**} = 1$ , in which case no element of  $D_1^M(0)$  will ever be profitable for the husband.

**Proof.** See Appendix A. ■

The first two thresholds apply to temporally advantaged choices. A key insight of this proposition is that no matter how wasteful a temporally advantaged private savings choice may be, we can always find a preference heterogeneity threshold beyond which an agent will find this choice profitable. This is intuitive: as  $\delta_M \rightarrow 0$  ( $\delta_F \rightarrow \infty$ ), the agent only cares about the first (second) period. Since temporally advantaged choices favor this relatively preferred period by definition, in the limit agents will prefer them to  $s_1^i = 0$  even when these choices incur higher transaction costs and/or require the use of substantially lower interest rates. At the same time, it is important to note that private savings choices need not be inefficient – this will depend on the parameter values under consideration.

The final two thresholds apply to temporally disadvantaged choices. Here, the sign of the inequalities are reversed: since these choices favor the *opposite* of the agent’s preferred period, they become less attractive as preference heterogeneity increases. Intuitively, one may expect that the use of individual accounts increases with preference heterogeneity, all else equal. However, the different thresholds necessary for temporally advantaged and disadvantaged deviations highlight that this may not always be the case. Indeed, the next proposition illustrates that individual account use will only be increasing in preference heterogeneity when the less patient spouse’s bank account is dominated by the joint account.

Before presenting that result, we must introduce some notation. When there are differential banking costs (i.e. it is more costly to travel to the bank to use an individual account than a joint account), a perfectly matched couple ( $\gamma = 1$ ) will always strictly prefer the joint account to individual account  $i$  bearing the same interest rate. To size the banking cost gap, define  $\tilde{R}_i(R_J; \Omega, \mathbf{y}, \mathbf{b}^a)$  to be the individual interest rate that makes a couple composed of two agents with  $\delta_i = \Omega$ , endowment  $\mathbf{y}$ , and banking cost vector  $\mathbf{b}^a$  indifferent between individual account  $i$  and the joint account. When it does not cause confusion, we shorten the notation to  $\tilde{R}_i(R_J)$ . If parameters are such that this couple would not save given  $R_J$ , let  $\tilde{R}_i(R_J) \equiv -\infty$ .

**Proposition 4** Fix  $\Omega$ ,  $\rho$ ,  $\mathbf{y}$ ,  $\mathbf{b}^a$ , and interest rates. Then the following hold:

- (a) If private savings is weakly preferred by the wife given  $s_1^M = 0$  and discount factor  $\delta_F(\underline{\gamma})$ , then it will be strictly preferred by wives in all households with  $\gamma > \underline{\gamma}$ .
- (b) Suppose  $R_M < \tilde{R}_M(R_J)$ . Then if private savings is weakly preferred by the husband given  $s_1^F = 0$  and discount factor  $\delta_M(\underline{\gamma})$ , it will be strictly preferred by husbands in all households with  $\gamma > \underline{\gamma}$ .
- (c) Suppose  $R_M \geq \tilde{R}_M(R_J)$ . The relationship between the husband's preference for private savings and  $\gamma$  given  $s_1^F = 0$  need not display upward monotonicity.

**Proof.** See Appendix A. ■

The most important result here is that whenever  $R_M < \tilde{R}_M(R_J)$  (i.e. a perfectly matched couple would prefer the joint account to the husband's account), households' use of individual accounts will be upwardly monotonic in preference heterogeneity – this follows directly from parts (a) and (b) of Proposition 4. For these households, strategic private savings action will always entail saving in an individual account: wives will "oversave", forcing more consumption into the future, while husbands exploit banking costs by running to the bank and saving in order to prevent the couple from returning to the bank to save in the joint account. Since the attractiveness of these deviations is increasing in discount factor heterogeneity, individual account use is also increasing in heterogeneity.

In contrast, part (c) of the proposition highlights that when the husband's account dominates the joint account, account use patterns need not be monotonic in preference heterogeneity. Recall that when the husband has the most attractive account, his best strategy may be to undersave (relative to the collective optimum) in his account. As preference heterogeneity increases, he may wish to undersave to the point of not saving at all, so this can lead to a negative correlation between individual account use and preference heterogeneity. Moreover, for some parameter values it may be that as preference heterogeneity increases, first the husband undersaves, then he refuses to save (forcing the couple to save jointly), and then makes use of the "save just enough" deviation. This would lead to a nonmonotonic relationship between individual account use and match quality, as indexed by  $\gamma$ .

Also note that in the proof, we show that temporally disadvantaged choices are never optimal for wives. As such, we do not analyze  $\gamma_F^{**}$  going forward.

Propositions 3 and 4 characterize the prevalence of private savings over a range of preference heterogeneity *given* interest rates, and identify conditions under which individual account use will be increasing in preference heterogeneity. We now study how *changes* in interest rates impact the incidence of private savings. We do so by studying the impact of one interest rate on the heterogeneity thresholds established in Proposition 3, conditional on the relevant alternative interest rate. First consider changes in  $R_i$ . Denote individual  $i$ 's preference heterogeneity thresholds (conditional on  $\Omega$ ,  $\rho$ ,  $\mathbf{y}$ , and  $R_J$ ) for private savings given  $R_i$  as  $\gamma_i^*(R_i)$  and  $\gamma_i^{**}(R_i)$ .

**Proposition 5** Fix  $\Omega$ ,  $\rho$ ,  $\mathbf{y}$ ,  $\mathbf{b}^a$ , and  $R_J$ . Suppose the private savings threshold for individual  $i$  is s.t.  $\gamma_i^*(R_i) \in (1, \infty)$ . Then  $\gamma_i^*(R'_i) < \gamma_i^*(R_i) \forall R'_i > R_i$ . Suppose  $\gamma_M^{**}(R_M) \in (1, \infty)$ . Then  $\gamma_M^{**}(R'_M) > \gamma_M^{**}(R_M) \forall R'_M > R_M$ .

**Proof.** See Appendix A. ■

The movements of  $\gamma_i^*$  and  $\gamma_M^{**}$  illustrate that increasing the individual interest rate (conditional on  $R_J$ ) makes private savings more attractive. Moreover, if we limit our attention to cases where temporally disadvantaged deviations are never optimal (i.e.  $R_M < \tilde{R}_M(R_J)$ ), then Proposition 5 implies that individuals in households with greater preference heterogeneity are willing to accept lower rates of return on individual accounts than their counterparts in better matched households.

We now perform the analogous comparative statics exercise, this time fixing  $R_i$  and changing the joint rate. Since we now vary  $R_J$ , we refer to preference heterogeneity thresholds as  $\gamma_i^*(R_J)$  and  $\gamma_i^{**}(R_J)$ .

**Proposition 6** Fix  $\Omega$ ,  $\rho$ ,  $\mathbf{y}$ ,  $\mathbf{b}^a$  and  $R_i$ . Suppose  $\gamma_F^*(R_J) \in (1, \infty)$ . Then  $\gamma_F^*(R'_J) < \gamma_F^*(R_J) \forall R'_J < R_J$ . In contrast, the comovement of  $\gamma_M^*(R_J)$  and  $\gamma_M^{**}(R_J)$  with  $R_J$  is ambiguous, unless  $\frac{\partial s_1^j}{\partial R_J} < 0$ .

**Proof.** See Appendix A. ■

For the wife, the result here mirrors that of Proposition 5. An increase in  $R_J$  always increases second period consumption under the household alternative, even when  $\frac{\partial s_1^j}{\partial R_J} < 0$ . So as  $R_J$  increases, more wives will find the household alternative attractive (and therefore  $\gamma_F^*$  increases). The key difference for husbands is that increasing the joint interest rate could actually make the household alternative *less* attractive when  $\delta_M < \Omega$ , since first period consumption decreases whenever  $\frac{\partial s_1^j}{\partial R_J} > 0$ . An interesting implication of Proposition 6 is that when given a menu of joint account interest rates, wives would always choose the highest rate. In contrast, a husband may prefer a lower joint rate if this serves to reduce the couple's savings. This is closely related to our earlier observation that a husband may sometimes refuse to save in his private account when it bears the highest rate of return. These results imply that less patient spouses may sometimes be willing to take costly action to block a household's access to higher return savings devices.

Our analysis has characterized how preference heterogeneity impacts account use given a set of interest rates, and how the prevalence of private savings changes given changes in individual and joint interest rates. While we cannot randomly assign preference heterogeneity to couples to test our theory, we can randomly assign interest rates to bank accounts. This observation inspired the design of a field experiment we conducted in Western Kenya, where married couples were given the opportunity to open three bank accounts (two individual accounts and one joint account) with randomly assigned interest rates. The following section describes this experiment in detail, and then derives testable implications of the theory by overlaying the above propositions with the experimental design.

## 1.3 Experimental Design and Testable Implications

### 1.3.1 Experimental Context

Our experiment was conducted in Western Province, Kenya, in areas surrounding the town of Busia. Busia is a commercial trading center straddling the Kenya-Uganda border. The town is well served by the formal banking sector, hosting over six banks at the time of field activities. It is only recently, however, that major banks have begun to offer products suitable for low income individuals. Traditionally, Kenyan bank accounts required opening balances upwards of Ksh 1,000 (approximately equal to \$12.50 at an exchange rate of Ksh 80 per \$1, or \$19.23 using a PPP exchange rate of Ksh 52 per \$1) and charged monthly maintenance fees around Ksh 50 (\$0.63).<sup>9</sup> However, recently banks have begun to target lower income individuals, and several banks currently offer lower fee alternatives to traditional bank accounts.

The financial partner for this study is Family Bank of Kenya. The bank currently has over 600,000 customers, 50 branches throughout the country, Ksh 13 billion (\$167 million) in assets, and actively targets low and middle income individuals as part of its corporate strategy. All study participants were offered Family Bank's *Mwananchi* accounts. This account can be opened with any amount of money, though a minimum operating balance of Ksh 100 (approximately \$1.25) cannot be withdrawn. The account pays no interest, but deposits are free of charge and there are no recurring maintenance fees. The only fees associated with the account are withdrawal fees, which are Ksh 62 (\$0.78) over the counter and Ksh 30 (\$0.38) with an ATM card. Account holders may purchase an ATM card for Ksh 300 (\$3.75), though this is not mandatory.

### 1.3.2 Experimental Design

#### 1.3.2.1 Targeted Population

The experiment targeted low income married couples who did not currently have an account with Family Bank but were potentially interested in opening one. At the outset of the study, we identified communities surrounding 19 local primary schools, which would serve as group meeting grounds. These schools were located between 0.2 and 7.7 miles from Family Bank's Busia branch, which is located in the town's commercial center. Targeted communities were situated either on the outskirts of Busia town or in nearby rural areas. Figure 1 illustrates the location of the schools relative to Family Bank and our field office on a map.

All experimental activities were conducted in group meetings. Trained field officers recruited households in communities surrounding a study school the day before each meeting. With the help of a local guide, they made door-to-door visits to households headed by married couples and issued meeting invitations to eligible households. To be eligible for invitation, a household had to be

---

<sup>9</sup>For comparison, the median household in our sample reported Ksh 1,200 in combined income in the week before the survey.

headed by a married couple, with both spouses present and able to attend the meeting. In addition, only households where both spouses had a valid Kenyan national ID card were admitted to the meetings, as Family Bank requires this document of all account holders.<sup>10</sup>

In order to compensate respondents for their time and to provide an additional incentive to attend the meetings, each individual who participated in the study received Ksh 100 in cash at the end of the meeting. Approximately 29 percent of issued invitations were redeemed over the course of the study. While far from universal, takeup rates are high enough that our sample represents a nontrivial fraction of targeted married couples in our catchment area.

### 1.3.2.2 Interventions

All couples attending our group meetings were given the opportunity to open up to three Family Bank accounts: an individual account in the name of the husband, an individual account in the name of the wife, and a joint account. To maximize takeup, we funded each opened account with the Ksh 100 (\$1.25) minimum operating balance (this amount could not be withdrawn by participants – it simply made opening an account costless). Participants were eligible to be randomly allocated into two core interventions, which are described below.<sup>11</sup>

**Intervention 1 - Interest Rates** Figure 2 illustrates the design of the interest rate intervention. Before deciding which accounts they wanted to open, couples drew three random interest rates (one rate per potential account). All interest rates were quoted to respondents as the 6-month yield on the average daily account balance. These interest rates were temporary, and expired after six months – respondents were told that the interest was a special promotion to help them save.<sup>12</sup> 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. Since some of our theoretical results are only unambiguous when the joint interest rate dominates the individual interest rate of the less patient spouse, we designed the experiment so the joint account had a greater probability of bearing the highest rate of return than an individual account: individual accounts could bear either 0, 2, 6, or 10 percent 6-month yields (with equal probability), and joint accounts could bear either 2, 6, or 10 percent 6-month yields (with equal probability).

---

<sup>10</sup>This requirement is common to all banks in Kenya. 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.

<sup>11</sup>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 (2010). We do not discuss this intervention here, as accounting for it has no impact on any of the results presented in this paper.

<sup>12</sup>After the six month period, balances earned no interest, which is standard for the *Mwananchi* account and other current accounts in Kenya. These interest rates were very high compared to market alternatives: small scale savings balances could earn at most 0.5-2.0 percentage points of interest annually given bank accounts available in Busia at the time of the experiment.

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

$$excess_{ac} = R_{ac} - \max \{R_{\hat{a}c} : \hat{a} \neq a\} \quad (1.4)$$

That is, the excess interest rate is the difference between account  $a$ 's rate and the best alternative rate. The interior cells of Figure 2 illustrate the variation in the excess rate for joint, husbands' and wives' accounts respectively.

After observing their interest rates, couples were separated and each spouse was administered a baseline survey. One concern is that randomizing the interest rates before conducting the baseline influenced survey responses. However, interest rates are not systematically associated with baseline self reports of savings levels or savings device use, nor are they associated with self reported decision making power regarding consumption and saving. We therefore expect that the randomization had little impact on survey responses. After the baseline, couples were reunited and decided which accounts they wanted to open.

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. Cards given for accounts bearing interest prominently featured the interest rate and several numerical examples of the 6-month return on different balances. If couples opened more than one account, enumerators marked the cards with the account type so couples could remember which account bore which rate of return. All cards were translated into Swahili, since individuals in the study area are generally more proficient in Swahili than English. Figure A1 illustrates the interest rate reminder cards and a translation of the text.

**Intervention 2 - Extra Statements** While not captured by our theory, a leading alternative model for why costly individual savings devices are popular in the developing world is that these devices enable individuals to hide resources from other household members, which in turn increases individual utility (Anderson and Baland 2002). In order to test whether the ability to hide savings was an important driver of individual account use in our sample, we randomly selected 50 percent of participating couples for an "extra statements" offer.<sup>13</sup> If a selected couple decided to open an individual account for (without loss of generality) the wife, the enumerator processing the couple's paperwork asked if they would consent to allow the husband to receive extra statement cards. The cards, if presented by the husband at the bank, entitled him to learn the current balance of his wife's account. These cards were only valid for 6 months, and were not given to couples unless both spouses gave their consent.

---

<sup>13</sup>Extra statements were not offered to the 97 couples (16 percent of our sample) in the first 6 experimental sessions.

### 1.3.3 Testable Implications

We now apply our theoretical results to the experimental design to generate three key testable implications, which we will take to the data in Section 1.4. To mimic our experimental context, consider a finite population of households where there is some nondegenerate distribution of preference parameters, bargaining weights, income, and banking costs, given by  $f(\gamma, \Omega, \rho, \mathbf{y}, \mathbf{b}^a)$ .<sup>14</sup> As in the field experiment, all households in this population have access to three bank accounts with randomly assigned interest rates:  $R_M$ ,  $R_F$ , and  $R_J$ . Let all interest rates vary over the range  $[\underline{R}, \overline{R}]$ .

Our theoretical discussion highlighted that certain types of inefficient private savings choices will always be profitable as long as preference heterogeneity in the household is *large enough*. In practice,  $f(\cdot)$  could be such that inefficient choices are never profitable for any individual in the population. Since we are interested in deriving predictions for the case where inefficient strategic behavior is observed among a nontrivial share of couples, we make the following assumption:

- A4. The distribution of preference heterogeneity, as governed by  $f(\cdot)$ , is such that for any interest rate draw  $[R_J, R_M, R_F]$  (where each  $R_a \in [\underline{R}, \overline{R}]$ ), there exists a positive mass of couples with  $\gamma > 1$  who save in individual accounts in any Nash Equilibrium.

For convenience, we also assume the following:

- A5.  $b^M = b^F > b^J$

Given our rural Kenyan context, it is essential that we account for banking costs when deriving our testable implications. Recall our definition of  $\tilde{R}_i(R_J; \Omega, \mathbf{y}, \mathbf{b}^a)$ : it is the individual interest rate that would make a couple composed of two agents with  $\delta_i = \Omega$ , endowment  $\mathbf{y}$ , and banking cost vector  $\mathbf{b}^a$  indifferent between individual account  $i$  and the joint account. If parameters are such that the couple would not save given  $R_J$ ,  $\tilde{R}_i(R_J; \Omega, \mathbf{y}, \mathbf{b}^a) \equiv -\infty$ . Similarly, define  $\tilde{R}_J(R_i; \Omega, \mathbf{y}, \mathbf{b}^a)$  to be the joint interest rate that would make a perfectly matched couple indifferent between individual account  $i$  and the joint account. If parameters are such that the couple would not save given  $R_i$ , let  $\tilde{R}_J(R_i; \Omega, \mathbf{y}, \mathbf{b}^a) \equiv \infty$ .

Since we consider a finite population, we can bound banking costs, incomes, and the average household discount factor. Then for a given  $R_J$  we can define the *maximum individual interest premium* in the population to be:

$$E_J(R_J) = \max \left\{ \tilde{R}_i(R_J, \Omega, \mathbf{y}, \mathbf{b}^a) \mid R_J \right\} - R_J$$

Similarly, for a given  $R_i$  we can define the *maximum joint interest discount* in the population to be

$$E_i(R_i) = R_i - \min \left\{ \tilde{R}_J(R_i, \Omega, \mathbf{y}, \mathbf{b}^a) \mid R_i \right\}$$

---

<sup>14</sup>Variation in banking costs is due to, for example, variation in distance from the bank and occupational mobility.



Given this setup, we analyze the savings choices of two groups of households – one consisting of a population of perfectly matched couples given by the distribution  $f(\gamma, \Omega, \rho, \mathbf{y}, \mathbf{b}^a \mid \gamma = 1)$ , and another population of poorly matched households given by the distribution  $f(\gamma, \Omega, \rho, \mathbf{y}, \mathbf{b}^a \mid \gamma > 1)$ .

Our first testable implication characterizes patterns of overall account use by match quality. The central idea behind the private savings theory is that badly matched couples make strategic use of individual accounts to manipulate consumption streams. Intuitively, one would therefore expect that perfectly matched couples make more intensive use of joint accounts, while poorly matched couples make more intensive use of individual accounts. This is true in many instances, but not always: the theoretical discussion highlighted that in some cases a husband in a badly matched couple may find it profitable to *refuse* to use his account when it bears the highest rate of return.<sup>15</sup> Fortunately, our formal characterization of private savings behavior enables us to identify a subset of the population for which the initial intuition always holds.

To do so, condition on  $\Omega$ ,  $\mathbf{y}$ , and banking costs,  $\mathbf{b}^a$ . Proposition 2 states that perfectly matched couples will always invest their savings optimally. Consider interest rate draws where a perfectly matched couple would save jointly. Propositions 3, 4, and Assumption A4 imply that a positive mass of poorly matched couples will save in individual accounts. The same argument applies to interest rate draws where a perfectly matched couple would not save at all, subject to the caveat that some poorly matched couples will be observed saving jointly due to mixed strategy equilibria. To the extent that the share of couples who must implement a mixed strategy is small (in practice, our empirical results support this) mixed strategies will not substantively impact our results. Next consider interest rate draws where a perfectly matched couple would prefer to save in an individual account. Proposition 4 implies that when  $R_M < \tilde{R}_M(R_J)$ , all poorly matched couples would also save in an individual account with positive probability (again subject to the mixed strategy caveat). Though  $\tilde{R}_M(R_J)$  is unobservable, limiting the sample to those couples for whom  $R_M \leq R_J$  will ensure that well matched couples are more likely to use the joint account, since  $R_J$  is always less than  $\tilde{R}_M(R_J)$  by Assumption A5. Then our first testable implication follows:

- T1. Consider the population of couples for whom  $R_M \leq R_J$ . Conditional on  $\Omega$ ,  $\mathbf{y}$ , and  $\mathbf{b}^a$ , well matched couples will be more likely to save in the joint account, while poorly matched couples will be more likely to save in the individual account.

Our next testable implication characterizes couples' responses to relative rates of return on their bank accounts. Here we exploit the fact that conditional on an account's interest rate,  $R_a$ , our experiment generated random variation in its relative rate of return, captured by  $excess_a$ . Since perfectly matched couples always invest efficiently (Proposition 2), they will always save in the highest return account available, while poorly matched couples will sometimes make inefficient choices. However, due to differential banking costs, we cannot simply compare savings rates in

---

<sup>15</sup>Here we refer to the account of the less patient spouse as the husband's account for rhetorical convenience. In practice, this account could belong to either the husband or the wife.

accounts where  $excess_a \geq 0$  and  $excess_a < 0$  by match quality: in some cases, the joint account will offer the highest return net of transaction costs, even when  $excess_J < 0$ . Instead, we characterize how savings rates change over a range of excess interest rates for a population with heterogeneous banking costs. We first describe the efficient response of perfectly matched couples, and then the response of poorly matched couples.

Panel A of Figure 3 graphs the average response to  $excess_i$  conditional on  $R_i$ . Note that perfectly matched households will never save in account  $i$  if  $excess_i < 0$  (here we use Assumption A5).<sup>16</sup> At an excess rate of zero, if there is a mass of households who have  $R_M = R_F$  but  $R_J < \tilde{R}_J(R_i)$ , we will observe a discrete jump up in the savings rate. Since our field experiment generated lumpy variation in the set of interest rates presented to a couple, the existence of such a mass seems reasonable. As the excess interest rate increases beyond zero, the share of households saving in account  $i$  will increase, until  $excess_i = E_i(R_i)$ . By definition, at this point account  $i$  will dominate account  $-i$  and the joint account for all perfectly matched households who are willing to save at interest rate  $R_i$ . Note that conditional on  $(\Omega, \mathbf{y})$ , the slope of the graph over  $(0, E_i(R_i))$  will be driven by the CDF of  $\tilde{R}_J(R_i)$ ; perfectly matched couples' response to the excess interest rate reflects the distribution of banking costs in the population. However, this result is only conditional – integrating over the distribution of  $\Omega$  and  $\mathbf{y}$  will also impact the shape of the graph.

Panel B of Figure 3 illustrates average savings responses for the joint account. The shape is essentially the same for perfectly matched couples – however, since the joint account has lower banking costs than individual accounts, perfectly matched couples begin to save at  $excess_J = -E_J(R_J)$ , and all choose the joint account once  $excess_J$  reaches zero. Notice that we should observe a positive slope only for *positive* excess rates when considering individual accounts and a positive slope only for *negative* excess rates when considering joint accounts. This asymmetry is a striking implication of efficient investment in the presence of heterogeneous banking costs.

The behavior of poorly matched households is more difficult to characterize. By Proposition 3 and Assumption A4, strategic behavior will lead a mass of poorly matched households to save in individual account  $i$  even when  $excess_i < 0$ . We also expect to see a positive mass of joint accounts in use when  $excess_J < 0$ . (This will include a mix of couples who save jointly because individual banking costs are too high, and couples where the less patient spouse has the highest return individual account but refuses to use it). It is also clear that a positive mass of couples will save in each type of account when the excess interest rate is positive.

It is not clear that the share of badly matched couples saving in account  $a$  will monotonically increase as  $excess_a$  increases. One issue is how the excess interest rate effects the prevalence of mixed strategy equilibria. The impact that this has on the graph in Figure 3 will depend on the shape of  $f(\cdot)$ . If we assume that mixed strategy equilibria do not meaningfully impact our predictions,

---

<sup>16</sup>If we allow for  $b^M \neq b^F$ , then the share of perfectly matched couples using account  $i$  may increase as  $excess_i$  increases even when  $excess_i$  is negative. However, the share of couples using account  $i$  when  $excess_i < 0$  should be small, as in many cases account  $i$  will be dominated by the joint account.

then we still have an ambiguous result for individual accounts, as illustrated by Proposition 6.<sup>17</sup> In contrast, Proposition 5 implies that the share of joint accounts in use should monotonically increase as  $excess_J$  increases.

We also note that we do not have a clear prediction for when or whether the share of couples saving in account  $a$  will be higher for perfectly matched as compared to poorly matched couples, as this will depend on  $f(\cdot)$ . Ultimately, the response of poorly matched couples to the excess interest rate is largely an empirical question. We are now prepared to state our second testable implication:

- T2. As illustrated by Figure 3, the share of well matched couples saving in individual accounts will begin to increase around an excess interest rate of zero and continue to increase until a plateau is reached at a positive excess interest rate. The share of well matched couples saving in the joint account should peak, then plateau, around an excess interest rate of zero. In contrast, poorly matched couples will not exhibit such breaks in behavior around an excess interest rate of zero.

Our final implication uses the result that perfectly matched couples invest efficiently. Suppose that in addition to actual interest rates, we were able to observe  $\tilde{R}_J(R_i)$ . Then perfectly matched savers would always choose to save in the account associated with  $\max\{R_J, \tilde{R}_J(R_M), \tilde{R}_J(R_F)\}$  – if we could calculate the "banking cost adjusted" rate of return on their savings, they would bear no interest loss. In contrast, poorly matched couples would not always choose the account with the highest effective rate of return – as a group they would bear a positive interest rate loss. This leads to our final testable prediction:

- T3. Poorly matched couples will bear larger banking cost adjusted interest rate losses than perfectly matched couples.

We now describe the data used to test T1-T3.

### 1.3.4 Data

We use two data sources for this project – survey data from one-on-one baseline questionnaires administered during the group sessions (spouses were separated for the interviews), and administrative data on account use from the bank. The baseline survey collected basic demographic information, as well as information on rates of time preference, decision making power in the household, income, current use of a variety of savings devices, and cross reports of spousal income and use of savings

---

<sup>17</sup>Increases in the excess rate driven by decreases in  $R_J$  may make private savings *less* attractive to the less patient spouse, which could push the share of households saving privately down as the excess interest rate increases. However, this would be counteracted by the effect for the more patient spouse, which would push the share of badly matched households saving privately up as  $excess_i$  increases. Changes in the excess interest rate due to changes in  $R_{-i}$  would also work to generate a positive slope. Since more forces favor a positive slope, we illustrate this in Figure 3.

devices. The administrative data provided by the bank includes the first six months' transaction history of all accounts opened under the auspices of the project. Each entry in the transaction history includes the deposit or withdrawal amount, any fees, and the type and time of the transaction.

### 1.3.4.1 Measuring Rates of Time Preference

We chose to elicit time preferences using choices between different amounts of money at different times, as opposed to different amounts of goods at different times. We made this choice for two reasons. First, Ashraf, Karlan, and Yin (2006) find that while time preference parameters estimated using choices between money, rice, and ice cream were all correlated, only the parameters estimated using money choices significantly predicted takeup and use of a commitment savings product. Second, cash lotteries made intuitive sense to respondents given that the group meetings revolved around bank accounts and savings.

We framed all questions as a choice between a smaller amount of money at a nearer time  $t$  ( $x^t$ ) and a larger amount of money at a farther time  $t + \tau$  ( $x^{t+\tau}$ ).<sup>18</sup> In order to make choices salient, respondents were given a 1 in 5 chance of winning one of their choices. Enumerators also used calendars to visually show respondents the number of days they would have to wait for both the smaller and larger amount of money.

In total, participants responded to 10 tables of monetary choices, with each table consisting of 5 separate choices between a smaller Ksh  $x^t \in \{290, 220, 150, 80, 10\}$  and larger  $x^{t+\tau} = \text{Ksh } 300$ . This was a sizeable amount of cash for the study participants. (For comparison, median reported daily earnings in our sample were Ksh 100 for men and Ksh 43 for women). The 10  $(t, t + \tau)$  pairs were:  $(\frac{1}{7}, 1)$ ,  $(\frac{1}{7}, 2)$ ,  $(\frac{1}{7}, 3)$ ,  $(\frac{1}{7}, 4)$ ,  $(\frac{1}{7}, 8)$ ,  $(\frac{1}{7}, 12)$ ,  $(2, 3)$ ,  $(2, 4)$ ,  $(4, 8)$ , and  $(4, 12)$  weeks. We chose to set the lowest near term  $t$  to "tomorrow" ( $\frac{1}{7}$ ) instead of "today" (0) to avoid confounding our discount factor estimates with differences in transaction costs of obtaining the funds in the near versus far term, or degrees of trust as to whether the money would be delivered (Harrison, Lau, Rutstrom, and Sullivan 2004). If a respondent won one of her choices, she had the option of having the funds deposited directly in her bank account, or picking the cash up at our field office, also located in Busia town.<sup>19</sup>

While our data allow us to calculate  $(\beta - \delta)$  quasi-hyperbolic preference parameters, we focus on calculating a single exponential discount factor, as motivated by our two period model. As

---

<sup>18</sup>This method is common to most empirical studies that attempt to measure rates of time preference in developing countries. Examples include Ashraf, Karlan, and Yin (2006), Bauer and Chytilová (2009), Dupas and Robinson (2011), Shapiro (2010), and Tanaka, Camerer, and Nguyen (2010).

<sup>19</sup>Despite the fact that the field office and Family Bank were proximately located, and that accessing cash deposited in an account would entail paying a withdrawal fee, 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 (in the commercial center of town), because the withdrawal fee was seen as a commitment device not to spend the money frivolously, or because the individuals intended to use their new accounts for saving anyway.

in Tanaka, Camerer, and Nguyen (2010), we use nonlinear least squares to estimate the discount factors. For each individual we assume that utility is linear in money amounts over the range Ksh 0–Ksh 300. Then the utility gains of the near and far amounts for person  $i$  considering choice  $q$  can be expressed as  $\Delta U_i(x_q^t) = \delta_i^t x_q^t$  and  $\Delta U_i(x_q^{t+\tau}) = \delta_i^{t+\tau} x_q^{t+\tau} + \varepsilon_{iq}$  where we assume  $\varepsilon_{iq} \sim \text{Logistic}(0, \mu_i)$ .<sup>20</sup> Then

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

A substantial mass of men and women either always chose the nearer option or always chose Ksh 300 in the future.<sup>21</sup> For these individuals, the discount factor is not identified. To obtain an estimate, we assume that all respondents would prefer Ksh 300 in the future to Ksh 0 sooner, and that all respondents would prefer Ksh 300 sooner to Ksh 300 in the future. Adding these two imputed questions to our tables leaves us with 70 responses for each individual. Define the dummy variable  $now_{iq} = 1(x_q^t \succ x_q^{t+\tau})$ . Nonlinear least squares solves

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

The nonlinear least squares algorithm converged for all but 1 individual in the sample. We dropped the couple where the algorithm failed for one of the spouses. We also topcoded  $\hat{\delta}_i$  at  $\bar{\delta}$ , the value of  $\hat{\delta}_i$  obtained via nonlinear least squares when  $now_{iq} = 0$  for every question, and bottomcoded  $\hat{\delta}_i$  at  $\underline{\delta}$ , the value of  $\hat{\delta}_i$  obtained when  $now_{iq} = 1$  for every question.<sup>22</sup>

As a robustness check, we also computed discount factors using a simple *ad hoc* bounding strategy similar to that found in Meier and Sprenger (2010). Specifically, suppose that for individual  $i$ ,  $x_q^t \succ 300^{t+\tau}$ , but that  $300^{t+\tau} \succ x_{q+1}^t$  (where  $x_q^t > x_{q+1}^t$ ). We then assume that the individual is indifferent between Ksh 300 at time  $t + \tau$  and the midpoint of the two amounts at time  $t$ ,  $\bar{x}_{q,q+1}^t = \frac{x_q^t + x_{q+1}^t}{2}$ . Using this midpoint, we can then calculate the implied discount factor  $\hat{\delta}_i^{q,q+1} = \left(\frac{\bar{x}_{q,q+1}^t}{300}\right)^{\frac{1}{\tau}}$ . We do this for each table, obtaining 10 discount factor estimates, and take the simple average of them. The correlation coefficient between the nonlinear least squares estimates and the *ad hoc* estimates is 0.97 and results in this paper are generally very robust to using this alternative method.

One concern with the discount factor estimates is that responses could reflect collective rather than individual preferences. Even though respondents answered questions separately from their spouses, their decisions may reflect the intrahousehold resource allocation process. If *ex-ante* com-

<sup>20</sup>Results assuming normally distributed  $\varepsilon_{iq}$  are essentially unchanged. Results using alternative discount factor estimates are available from the author upon request.

<sup>21</sup>26 percent of men and 22 percent of women always waited for Ksh 300 in the future, while 15 percent of men and 14 percent of women always wanted the smaller, nearer amount.

<sup>22</sup>This led to the censoring of 19 estimated discount factors from below and 53 estimated discount factors from above.

mitment describes the household, spouses should make identical choices. Even when action is strategic, if our monetary amounts are inframarginal to the savings decision then choices should reflect the discount factor used to determine savings allocations. If our respondents answered questions in this way, the majority of intracouple variation in discount factors would be due to measurement error. This would generally bias us towards rejecting T1-T3.<sup>23</sup> It is also possible that our discount factor estimates capture other aspects of individual preferences, such as risk aversion – unfortunately, while we can control for heterogeneity in observable characteristics, we cannot account for heterogeneity in discount factors driven by unobservable aspects of preferences.

#### 1.3.4.2 Sample Characteristics

Our baseline sample consists of 597 non-polygamous married couples where both spouses had valid national ID cards, valid estimates of  $\delta_i$ , and no pre-existing accounts with Family Bank.<sup>24</sup> Table 1 presents baseline characteristics of the sample. Respondents are of relatively low socioeconomic status – husbands average 8 years of schooling, and their wives average just under 6 years. While most men are literate (86 percent), one third of women cannot read and write. On average, men reported earning Ksh 1,348 (about \$17) in the past week, while women reported earning Ksh 798 (\$10). However, median reported weekly incomes are substantially lower, at Ksh 700 and Ksh 300 for husbands and wives respectively.

Almost all respondents reported using at least one savings device at baseline. Most common was saving cash at home, reported by 85 and 90 percent of husbands and wives respectively. Reported savings levels at home were substantial and approximately equal to average weekly earnings. Half of men and two-thirds of women reported participating in at least one ROSCA, corroborating earlier studies demonstrating that ROSCAs are generally more popular among women than men in East Africa (Anderson and Baland 2002; Dupas and Robinson 2011). Savings accounts with formal banks were less common, particularly for women – while 30 percent of men reported owning a savings account (and those men reported substantial savings in their accounts), just 12 percent of women reported owning a savings account.

Pronounced gender patterns are also present in responses to questions regarding decision making power within the household. The majority of individuals reported that either the husband made most consumption decisions (47 percent of men, 38 percent of women) or that both spouses jointly made consumption decisions (39 percent of men and 38 percent of women). In contrast, individuals were most likely to report that the *wife* was primarily responsible for making savings decisions in the household. In addition, only 10-11 percent of individuals reported that they jointly made savings

---

<sup>23</sup>A notable exception is if measurement error were correlated with cognitive ability. This could potentially generate patterns in the data consistent with T1-T3. To test for this, we will present regression specifications that control for spousal levels of and heterogeneity in education and literacy.

<sup>24</sup>We dropped 179 polygamous couples from our sample since strategic behavior may be very different in households with more than one wife. However, our results are robust (though in some cases somewhat attenuated) to including them.

decisions with their spouse, substantially less than the share of individuals reporting joint decision making regarding consumption.

Finally, Table 1 reports the average estimated discount factor (the distribution of these discount factors can be seen in Figure 4). On average, study participants appear to be very impatient, with weekly discount factors averaging 0.80 for men and 0.79 for women. These discount factor estimates are lower than estimates in studies of individuals in developing countries in Asia (Tanaka, Camerer, and Nguyen 2010; Ashraf, Karlan, and Yin 2006; Bauer and Chytlovic̃ 2009; Shapiro 2010). However, Dupas and Robinson (2011) find very impatient preferences among a Kenyan sample of low-income entrepreneurs.<sup>25</sup>

Next we examine the correlation between the discount factor estimates and observables. Table 2 presents results of regressing estimated discount factors on demographic and economic characteristics.<sup>26</sup> While demographic characteristics are essentially uncorrelated with estimated discount factors, variables related to savings are more significant. Most importantly, estimated discount factors are positively correlated with observed savings levels in Family Bank accounts, and significantly so for women.<sup>27</sup> These correlations are reassuring and suggest that the estimated discount factors are capturing preferences over savings levels. While the correlations are generally insignificant for men, this may be due to distortions in savings behavior stemming from strategic action – the correlation between discount factors and Family Bank savings deposits is positive and statistically significant for both genders when we limit the sample to well matched couples.

#### 1.3.4.3 Randomization Verification

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

Table 3 presents the randomization verification exercise for interest rates, extra statements, and cash prizes. Panel A presents the actual proportions of each treatment observed in the sample. P-values from a binomial test that the observed proportion equals the theoretical proportion (1/4 for individual accounts, 1/3 for joint accounts, 1/2 for extra statements, and 1/5 for cash prizes)

---

<sup>25</sup>It is not essential that our estimated discount factors reflect the actual level values that govern individuals' intertemporal decisions. What we require is that the intracouple heterogeneity in our estimates reflects true heterogeneity in the sample.

<sup>26</sup>Missing demographic characteristics were recoded to zero and dummied out. This convention is held for the remainder of the paper.

<sup>27</sup>One concern is that the act of giving cash for discount rate elicitation choices, which could be deposited into Family Bank accounts, would mechanically cause this correlation. However, the correlation for women remains statistically significant at the 95 percent level when dropping all couples where at least one spouse was randomly selected to receive a cash prize.

are presented in braces. While couples were somewhat less likely to be selected for extra statements and cash prizes than expected, the results suggest that enumerators and participants respected the experimental protocol.

Panel B presents chi-square test statistics and associated p-values from tests for the equality of treatment distributions across demographic characteristics. The first two columns are limited to either the husband or the wife, whereas the last three columns use both husband and wife demographic characteristics. Overall, the randomization appears to have functioned well. The number of coefficients significant at the 90 and 95 percent level are approximately equal to the expected number due to chance, and there are no systematic patterns across the different treatments.

## 1.4 Results

### 1.4.1 Account Takeup and Interest Rate Responses

While couples had the option of opening all three bank accounts and incurred no costs in doing so (aside from the cost of additional time spent filling out paperwork), most couples chose to open only one or two accounts. Although all couples opened at least one account, only 5 percent of couples in our sample opened all three accounts. Fifty-six percent of couples chose to open only the joint account, 3 percent of couples only opened one individual account, 29 percent of couples opened two individual accounts, and the remaining 7 percent of couples opened a joint account and one individual account. When we informally asked couples why they decided not to open all the accounts, most couples stated that they did not have the resources to maintain three accounts, or that they only had use for the account(s) that they chose to open. The additional time spent doing paperwork may have been enough to dissuade couples from opening accounts that they were very certain they would never use – since we explained that our 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.

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

$$y_{ac} = \alpha_0 + int'_{ac}\beta + \gamma_s + \varepsilon_{ac} \quad (1.5)$$

where  $y_{ac}$  is the outcome of interest,  $int_{ac}$  is a vector including interest rate dummy variables, a joint account dummy, and interactions between the joint account and interest rate dummies, and  $\gamma_s$  are experimental session fixed effects. (We include experimental session fixed effects for continuity with the results involving preference heterogeneity – however, results are essentially unchanged if we do not include them). We constructed several measures of account use for Table 4. Using the Family Bank transaction data, we calculated the average daily balance of account  $a$  owned by couple  $c$  using 6 months of data from the day of account opening (this covers the period where the



account was eligible to earn interest). The average daily balance nets out the Ksh 100 minimum operating balance deposited into every open account, since this amount could not be accessed and was deposited on behalf of the respondents. All unopened accounts were assigned an average daily balance of zero. In addition to the average daily balance, we also created a savings indicator for accounts with positive average daily balances, and calculated total deposits and withdrawals over the 6-month study period, the total number of account transactions, and the total amount of fees charged to the account.

Though the impact of the interest rate on saving is theoretically ambiguous, the income effect should be very small since our interest rates were temporary. We therefore expect that higher interest rates will be associated with greater savings. Table 4 confirms that higher interest rates are associated with significantly higher rates of account opening and use. For example, individual accounts with 10 percent interest were 66 percent more likely to be opened, more than twice as likely to receive savings deposits, and had 3.3 times the average daily balances of individual accounts that bore no interest. The table also illustrates that joint accounts were much more popular among our study couples than individual accounts. For example, at 2 percent interest, 60 percent of the couples chose to open a joint account, while only 34 percent of couples chose to open an individual account. This pattern carries over to account use as well. The interest rate interaction terms reveal that while level use of joint accounts is much higher, the slope with respect to the interest rate does not differ between the two account types.<sup>28</sup>

The relative popularity of joint accounts in our sample could reflect large differential banking costs. In addition, couples may have been biased towards joint accounts because they decided which accounts to open together. However, since use of joint accounts is higher even conditional on opening (42 percent of open joint accounts were used for saving while 25 percent of opened individual accounts were used for saving), it seems unlikely that social pressure at the meeting is entirely responsible for the high rates of joint account use.

#### 1.4.2 Discount Factor Heterogeneity and Account Use

As illustrated by Figure 4, there is substantial variation in estimated discount factors within our sample. For couple  $c$ , define preference heterogeneity to be the difference between the male and female estimated discount factors:  $het_c = \hat{\delta}_c^M - \hat{\delta}_c^F$ . Figure 5 presents the histogram of  $het_c$ . While 13 percent of couples have identical discount factor estimates, many couples have estimates that differ substantially.

Our three testable implications involve comparing the behavior of couples who are perfectly and poorly matched in terms of discount factors. To approximate this empirically, we choose a rule of

---

<sup>28</sup>Since free ATM cards were only issued to open accounts (and are therefore correlated with the interest rate), we also checked to see if the results in Table 4 were driven by the free ATM treatment. Accounting for this treatment has little effect on estimated coefficients and does not change the result that couples robustly respond to the interest rate. These results are available from the author upon request.

thumb that divides the sample into three equal sized groups: "wife more patient", "well matched", and "husband more patient". We label the 1/3 of the sample with the most closely aligned discount factors as well matched and refer to the remaining couples as badly matched.<sup>29</sup> This corresponds to  $\left| \hat{\delta}_c^M - \hat{\delta}_c^F \right| \leq 0.081$  (this is equivalent to the couple's discount factors being within 0.366 standard deviations of one another). These well matched couples are delineated by dashed lines in Figure 5. Our results are similar when using a variety of alternative measures, including defining well matched couples to be the best matched 1/2 or 1/4 of couples, and to using  $\left| \hat{\delta}_c^M - \hat{\delta}_c^F \right|$  as a measure of preference heterogeneity.

Our first testable prediction outlined in Section 1.3 (prediction T1) is that well matched couples make more intensive use of joint accounts, while poorly matched couples make more intensive use of individual accounts. Recall that this prediction is only unambiguous when the less patient spouse's account bears an interest rate that is weakly lower than the joint interest rate. Figure 6 explores the relationship between account choice and preference heterogeneity, using the subsample of couples who meet this criterion. The figure presents results of the following local linear regression

$$y_c = g(het_c) + \varepsilon_c \quad (1.6)$$

where  $y_c$  is the outcome of interest and  $het_c = \hat{\delta}_c^M - \hat{\delta}_c^F$ . Panel A graphs savings rates by account type. Average joint account savings rates are higher for well matched couples, who are demarcated by the grey vertical lines in each panel. In contrast, the savings rates for both individual accounts are U-shaped, with higher savings rates among badly matched couples. This pattern also emerges when limiting the sample to savers and examining the fraction of total savings that are stored in each account (Panel C). Patterns for average daily balances (Panel B) are similar though somewhat less clear, but this is not entirely surprising. Account balances are generally low with some very large outliers (for example, among couples who saved, the median average daily balance deposited in all accounts was Ksh 290, while the 95<sup>th</sup> percentile was Ksh 4,639 and the 99<sup>th</sup> percentile was Ksh 9,794). As such, outliers have large impacts on means.

In order to test the significance of these results and to control for potentially confounding factors, we run the following regression:

$$y_c = \beta_0 + \beta_1 match_c + \beta_2 joint\_dev_c + int'_c \delta + x'_c \lambda + \gamma_s + \varepsilon_c \quad (1.7)$$

where  $y_c$  is the outcome of interest,  $match_c$  indicates well matched couples,  $joint\_dev_c$ , indicates badly matched couples where the less patient spouse's account has a strictly higher return than the joint account,  $int_c$  is a vector of dummy variables for each account's interest rate,  $x_c$  is a vector of additional controls, and  $\gamma_s$  are session fixed effects. The omitted match quality group is badly matched couples where the less patient spouse's individual account does not dominate the joint account in terms of the interest rate. The private savings theory predicts that this group will have

---

<sup>29</sup>Table A1 illustrates demographic differences between well matched and poorly matched couples.

the highest rate of individual account use and the lowest rate of joint account use. Note that although  $match_c$  and  $joint\_dev_c$  are generated regressors, the null hypothesis specifies that  $\beta_1 = \beta_2 = 0$ . In this case, traditional standard errors are consistent (Newey and McFadden 1994). Throughout the paper, we therefore present either heteroskedasticity robust standard errors (for couple level regressions) or standard errors clustered at the couple level (for account level regressions).

Table 5 presents estimates of  $\beta_1$  ("Well Matched") and  $\beta_2$  ("Joint Deviation"), with four separate control sets. Our "basic" control set, in Panel A, only includes interest rate dummy variables and session fixed effects. This regression essentially mirrors the results of Figure 6. Panel B then adds additional controls related to rates of time preference. To account for general differences in patience levels between well and poorly matched couples, we control for  $\frac{\delta_c^M + \delta_c^F}{2}$  linearly.<sup>30</sup> Our discount factor estimates are censored for all couples who answered that they either always preferred the later, larger amount (they always waited) or that they always preferred the nearer, smaller amount (they always wanted the funds as soon as possible). This censoring would flatten the curves presented in Figure 6. To address this, we include dummy variables for each spouse indicating that he or she always chose to wait, analogous dummies for spouses who always wanted the sooner amount, and separate dummy variables indicating that spouse  $i$  in couple  $c$  had his or her discount factor top or bottom-coded because the nonlinear least squares estimate went outside the bounds implied by always-patient and always-impatient responses.

In Panel C we add controls for a variety of demographic characteristics. Panel D adds additional controls for economic characteristics. To ensure flexible accounting for both levels and intracouple heterogeneity, for non-binary characteristics we include the linear and squared terms for both husband and wife, as well as the interaction between the linear values for husband and wife. For binary variables, we include the dummy variable for both husband and wife as well as the interaction.<sup>31</sup>

Since adding demographic and economic covariates has little impact on our results, we take the estimates in Panel B as our baseline specification. As predicted, well matched couples are significantly less likely to save in individual accounts (75 percent, 8.6 percentage points) and significantly more likely to save in joint accounts (41 percent/11.2 percentage points) when compared to their poorly matched counterparts. Moreover, we observe that poorly matched couples with available "joint deviations" are indeed more likely to use joint accounts and less likely to use individual accounts when compared to poorly matched couples where such a deviation is never profitable. (However, these differences are generally insignificant and our results are substantively unchanged if we pool all poorly matched couples). Limiting the couples to savers, we also see that well matched

---

<sup>30</sup>Choosing a 50/50 weighting for the average discount factor in a couple is somewhat *ad hoc*, as the true weighting should be related to the individual bargaining weights. However, results are robust to a wide variety of different weighting schemes, including proxying the bargaining weight with relative income shares and using self reported savings decision making power to construct weights.

<sup>31</sup>Demographic controls include age, years of education, an individual literacy dummy, and the number of children reported by each spouse. The economic controls include individual income, dummies for mobile phone ownership, and dummies indicating that an individual is either a subsistence farmer or has no job.

couples store a lower share of balances in individual accounts (58 percent/11.6 percentage points) and a greater share of balances in joint accounts (40 percent/24.2 percentage points). Estimates for average balances are also large and directionally correct, though generally insignificant.

We now consider our second testable prediction (T2). First we generate the empirical analog of Figure 3 by running the following regression separately for well and poorly matched couples:

$$saved_{ac} = \beta_0 + ex'_{ac}\delta + z'_{ac}\lambda + \varepsilon_{ac} \quad (1.8)$$

Where  $saved_{ac}$  indicates that couple  $c$  saved in account  $a$ ,  $ex_{ac}$  is a vector of dummy variables for the excess interest rate on account  $a$ , and  $z_{ac}$  is a vector of dummy variables for account  $a$ 's interest rate.<sup>32</sup> We then calculate predicted values of  $saved_{ac}$  for each value of the excess rate, assuming equal distribution of the sample at each interest rate (0, 2, 6, and 10 percent for individual accounts; 2, 6, and 10 percent for joint accounts). Figure 7 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 3 that individual account use by well matched couples could jump discretely up at  $excess_{ac} = 0$ . Since  $-2$  is the largest negative value of the excess interest rate in our sample, we therefore fit separate lines for  $excess_{ac} \leq -2$  (this slope should be zero) and  $excess_{ac} \geq -2$  (this slope should be positive). In contrast, the slope for joint accounts should be positive below an excess rate of zero and flat thereafter, so our lines are drawn above and below  $excess_{ac} = 0$ .

Panel A presents account use for well matched couples. The results for individual and joint accounts are both consistent with theoretical predictions. Specifically, the share of couples using individual accounts is essentially flat until the excess interest rate reaches zero, where the share saving jumps up. Beyond this point the share increases, though we do not observe a subsequent plateau. For joint accounts, the share of well matched couples saving slopes upward when the excess interest rate is negative and then plateaus after an excess rate of zero. We do not, however, observe an initial plateau at a low savings rate when the excess interest rate is negative.

These "missing" plateaus suggest that the variation we created in the excess interest rate was not large enough to reach the thresholds  $E_i(R_i)$  and  $-E_J(R_J)$  illustrated on Figure 3; in other words, the excess interest rate was never large enough to dominate differential banking costs for the vast majority of couples in our sample. This is conceivable given the temporary nature of the interest rates and the low savings levels of our couples – the 75<sup>th</sup> percentile total average daily

---

<sup>32</sup>As a result of the experimental design, some values of the excess interest rate were only realized for a very small number of accounts: 14 accounts had an excess interest rate of 2, 12 accounts had an excess interest rate of 6, and 13 accounts had an excess interest rate of 10. For each of these values, we downcode the excess interest rate by two percentage points (results are invariant to simply dropping these accounts). Similarly, we pool  $excess_{ac} = -10$  and  $excess_{ac} = -8$  as the omitted category in our regressions. We do this in order to identify all interest rate dummy variables, as accounts with zero percent interest had excess interest rates unique to them. Results are essentially unchanged if we drop all accounts with a zero percent interest rate.

balance *among savers* was Ksh 789. Even foregoing 10 percentage points of interest would only cost the 75<sup>th</sup> percentile household Ksh 79. In comparison, the cost of round trip travel to and from the bank meets or exceeds Ksh 100 for over half of our study households. Given this observation and the patterns in the data, we infer that interest rate differentials only trumped banking cost concerns for couples with small banking cost differentials.

Panel B of Figure 7 examines the behavior of poorly matched couples. Their rates of saving are completely insensitive to the excess interest rate. While the response to the excess interest rate is theoretically ambiguous for individual accounts, our theory predicts a positive slope for joint accounts. However, it is possible that the variation in the excess interest rate was simply not large enough to incite a substantial change in behavior for these couples. Since interest rates were temporary and average balances small, private savings deviations are "cheap" in terms of interest rate losses. In this case, it could be that the majority of badly matched couples with small differential banking costs chose to save privately, while poorly matched couples with large differential banking costs chose to save jointly, and very few marginal couples were swayed by the excess interest rate.

We now test the significance of the patterns exhibited by well matched couples in Figure 7. To do so, we generate splines in the excess interest rate. Consistent with Figure 7, we place a knot at  $excess_{ac} = -2$  for individual accounts and a knot at  $excess_{ac} = 0$  for joint accounts. We then run the following regression among well matched couples:

$$\begin{aligned}
 saved_{ac} = & \beta_0 + \beta_1 (below \times indiv)_{ac} + \beta_2 (above \times indiv)_{ac} + \\
 & \beta_3 (below \times joint)_{ac} + \beta_4 (above \times joint)_{ac} + z'_{ac}\lambda + x'_c\alpha + \gamma_s + \varepsilon_{ac}
 \end{aligned} \tag{1.9}$$

where  $below_{ac}$  is the spline capturing the slope below the knot,  $above_{ac}$  is the spline capturing the slope above the knot,  $indiv_{ac}$  is an individual account indicator,  $joint_{ac}$  is a joint account indicator  $z_{ac}$  is a vector including interest rate dummies, the joint account dummy, and joint×interest rate interactions,  $x_{ac}$  is a vector containing the control sets used in Table 5, and  $\gamma_s$  are experimental session fixed effects. To test whether well matched couples' responses to the excess interest rate are significantly different from poorly matched couples, we also present a specification where we include both types of couples, include a badly matched indicator, estimate  $\beta_1$ - $\beta_4$ , and  $\lambda$  separately for well and poorly matched couples (by interacting the relevant variables in equation 1.9 with the badly matched indicator), and constrain the remaining coefficients ( $\alpha$  and  $\gamma_s$ ) to be the same for all couples. Table 6 presents the results.

Panel A examines well matched couples. The first two coefficients report the response to excess interest rates on individual accounts. As expected, the coefficient on the lower spline in the excess interest rate is very close to zero and insignificant. In contrast, the coefficient on the upper spline is much larger, positive, and significant. The pattern for joint accounts is, as expected, reversed. The coefficient on the lower spline is large, positive, and significant, while the point estimates for the upper spline are much smaller in magnitude and insignificant. The coefficients reveal very

large responses to the excess interest rate. For example, the final column of panel A implies that increasing the individual excess interest rate from 0 to 10 results in an 18.2 percentage point increase in the savings rate, while increasing the joint excess interest rate from  $-10$  to 0 results in a 38.4 percentage point increase in the savings rate. These impacts are very large when compared to dependent variable means (6 percent for individual accounts and 34 percent for joint accounts).

Panel B presents the results for all couples pooled together. The responses for well matched couples can be read off of the main effects – we see that pooling all couples has little impact on the original estimates. Examining differential responses for poorly matched couples reveals that their aggregate responses (which are given by the sum of the main effect and the badly matched interaction) are very close to (and not significantly different from) zero. Furthermore, the responses of well and poorly matched couples differ significantly where expected based on the patterns in Figure 7 (the upper spline of individual accounts and the lower spline of joint accounts).

Taken together, the results in Tables 5 and 6 are incompatible with a model where couples can commit intertemporally. Specifically, we only observe evidence of efficient investment among well matched couples. Though badly matched couples' lack of a response to the excess interest rate could be rationalized if they had substantially larger differential banking costs, we observe them making more intensive use of individual accounts. Moreover, Table A1, which illustrates individual demographic characteristics by match quality, reveals that these two groups of couples are very similar with respect to most observables, including distance to the bank and baseline bank account ownership. This suggests that badly matched couples' response to the excess interest rate reflects inefficient investment rather than demographic differences. Thus, our results add to a body of literature demonstrating that households in developing countries often make decisions that are not Pareto efficient.<sup>33</sup>

### 1.4.3 Proxying Banking Costs

Our results suggest that banking costs are an important concern for couples in our study – both well and poorly matched couples were much more likely to use joint accounts, and the excess interest rate responses among well matched couples in Figure 7 imply that even an excess rate of 8 (for individual accounts) or  $-8$  (for joint accounts) did not outweigh banking cost concerns for a nontrivial share of the couples in our sample. Moreover, the positive slopes in Figure 7 imply a distribution of banking costs in the population; if all couples had identical banking costs, we would observe an initial flat plateau, then a discrete jump up followed by another flat plateau at a positive savings rate.

The unobservability of differential banking costs makes our final testable implication (T3), that well matched couples invest their resources more efficiently than poorly matched couples, difficult to analyze. In order to make some adjustment for unobservable banking costs, we attempt to proxy them with observables. We conjecture that those couples who travel to Busia town very frequently

---

<sup>33</sup>Studies documenting inefficient behavior include Ashraf (2009), de Mel, McKenzie, and Woodruff (2009), Duflo and Udry (2004), Robinson (2008), and Udry (1996).

for non-bank related reasons and those couples who have low travel costs to town will have smaller differential banking costs. This should be negatively correlated with distance from the bank and economic activity (here we assume that subsistence farmers and the unemployed are less likely to take frequent trips to town). Moreover, pre-existing use of a formal savings account should signal lower differential banking costs. Here we include both bank accounts and SACCO accounts as formal accounts.<sup>34</sup> To aggregate these measures, we used principal components analysis to extract the first principal component of the data matrix formed by the above-listed variables. We then normalized this component to construct a "banking costs index", which runs from zero (lowest hypothesized banking costs) to one (highest hypothesized banking costs). Table 7 checks to see if the cost index is correlated with account behavior in ways predicted by our theory. The first three columns present couple-level regressions of the following form:

$$y_c = \beta_0 + \beta_1 index_c + x'_c \delta + \varepsilon_c \quad (1.10)$$

where  $y_c$  is the outcome of interest,  $index_c$  is the banking cost index, and  $x_c$  is a vector of controls. We use the same control sets as in previous regressions, but we exclude the subsistence farming indicators from the economic control set, as this variable is used in the creation of the cost index. Since session fixed effects absorb distance from the bank, we exclude these as well.

It is likely that the cost index is correlated with both absolute banking costs and differential banking costs. To test the former hypothesis, the first column of Table 7 examines the correlation between the banking cost index and a dummy variable for whether or not a couple saved in any bank account. Indeed, couples with the highest costs are significantly less likely to save in any bank account. The next two columns limit the population to couples who saved in at least one account, and examine the correlation between use of the joint account and the cost index.<sup>35</sup> Here, higher cost savers are significantly more likely to save jointly as compared to lower cost savers. These results are relatively robust to adding observable controls – though the results examining joint account use are substantially attenuated upon including economic controls, the coefficients are still large in magnitude. This is encouraging, and suggests that the cost index is in part capturing differential banking costs.

Earlier, we argued that well matched couples' response to (and poorly matched couples' lack of response to) the excess interest rate suggest that interest rate losses only impacted the decisions of well matched savers with relatively low differential banking costs. We can use our proxied banking

---

<sup>34</sup>SACCO stands for "savings and credit cooperative". SACCOs function like credit unions, and are generally organized around higher paying professions, such as teaching and commercial farming.

<sup>35</sup>Since the cost index is correlated with absolute banking costs, the correlation between the index and joint account use is technically ambiguous. In the absence of a selection effect (i.e. conditional on  $b^j$ ), larger differential banking costs will push couples to make more intensive use of joint accounts. However, larger absolute banking costs will select out smaller scale savers. Since the hypothesized costs are fixed and not proportional to balances, these savers will be more likely to opt for joint accounts, all else equal. Therefore this selection effect would bias us away from finding a positive correlation between joint account use and the cost index.

costs to test this hypothesis directly: in this case well matched couples with low proxied banking costs should respond robustly to the excess interest rate, while well matched couples with large proxied banking costs should be much less responsive to the excess rate. To test this, we limit the sample to well matched couples, define a couple to have high banking costs if their index value is above the sample median, and run the following regression:

$$\begin{aligned} saved_{ac} = & \beta_0 + \beta_1 high_c + \beta_2 excess_{ac} + \beta_3 high_c \times excess_{ac} + \beta_4 joint_{ac} + \\ & \beta_5 (joint \times high)_{ac} + z'_{ac} \lambda + x'_c \delta + \varepsilon_{ac} \end{aligned} \quad (1.11)$$

where  $high_c$  is a dummy variable indicating high banking costs,  $excess_{ac}$  is the excess interest rate on account  $a$  for couple  $c$ ,  $joint_{ac}$  is a joint account indicator,  $z_{ac}$  is a vector of interest rate dummies and their interactions with  $joint_{ac}$ , and  $x_c$  is a vector of the same control sets used in the earlier specifications in Table 7.

The fourth column of Table 7 presents results of this regression. As expected, the response to the excess interest rate is positive and highly significant for couples with lower proxied banking costs and very close to zero for couples with higher proxied banking costs. One issue with this regression is that we constrain  $\lambda$  (the vector capturing interest rate responses) to be the same for high and low cost couples. The final column of Table 7 presents an alternative specification where we add interactions between  $z_{ac}$  and the high cost dummy. Although this decreases precision somewhat, the point estimates are essentially unchanged. We also note that the p-values on an F-test that the interest rate-high cost interactions are equal to zero range from 0.76-0.97, so imposing the restriction seems reasonable. Overall, these results suggest that our index is a reasonable proxy for differential banking costs, and that couples' responses to the excess interest rate are in fact being driven by banking costs.

#### 1.4.4 Investment Efficiency by Match Quality

Our results suggest that poorly matched couples invest less efficiently than well matched couples, as per implication T3. We now use our proxied banking costs to estimate the magnitude of misallocation by poorly matched couples. A simple technique to compare the investment efficiency of well and poorly matched couples would be to calculate the maximum interest rate that each saving couple could have earned on their balance, compare this to the actual interest rate they earned, and then check to see if interest rate losses are greater for poorly matched couples. However, since lower return joint accounts may be more efficient than higher return individual accounts when banking cost differentials are large, this comparison may overstate the degree of savings misallocation among well matched couples. Our theoretical discussion highlighted that given an average household discount factor,  $\Omega$ , a vector of banking costs,  $\mathbf{b}^a$ , an endowment  $\mathbf{y}$ , and an individual interest rate,  $R_i$ , we can calculate the joint interest rate that would make a well matched couple indifferent between saving jointly and individually,  $\tilde{R}_J(R_i)$ . The "joint interest discount",  $R_i - \tilde{R}_J(R_i)$  then reflects



the importance of differential banking costs.

In this spirit, we calculate "banking cost adjusted" actual and maximum interest earnings for all savers. To do so, we adjust individual interest rates downward to account for banking costs using the cost index that we constructed in the previous subsection. To perform the adjustment we simply multiply the cost index by a maximum joint interest discount ( $E_i$ ), and adjust individual interest rates by the resulting product.

While this method is somewhat *ad hoc*, conducting the analysis for a range of  $E_i$  should give a general idea of differences in savings misallocation between well and poorly matched couples. Table 8 presents the results of this exercise. The first column sets  $E_i$  to zero – this column makes no banking costs adjustment whatsoever. Here we see that if poorly matched savers had always chosen the highest return account available, the average saving couple would have earned 8.3 percentage points of interest. In practice, these couples averaged 6.0 percentage points of interest, leading to an interest loss of 2.3 percentage points. In contrast, well matched couples could have earned a maximum of 8.3 percentage points of interest and actually earned 6.9 percentage points. Therefore, the "loss gap" between poorly and well matched couples is 0.9 percentage points of interest, which is significantly different from zero. Even without accounting for differential banking costs, poorly matched couples appear to suffer from greater savings misallocation – their losses are 64 percent larger than their well matched peers.

The next four columns repeat this analysis with different values of  $E_i$ , ranging from 5 to 20 percentage points (recall that well matched couples' response to the excess interest rate in Figure 7 suggests that  $E_i$  exceeds 8). The estimated loss gap increases as  $E_i$  increases, though poorly matched losses as a percent of well matched losses peak at  $E_i = 5$ , where they are 162 percent larger. These losses are very large in percentage terms, but relatively small in absolute terms. A loss of 3 percentage points in interest amounts to Ksh 24 for the 75<sup>th</sup> percentile saving couple. However, since banking cost differentials persist for the life of the account, long run absolute losses due to inefficient individual account use could be large.

Taken as a whole, our results fit the predictions of the private savings model very well. A key feature of this model is the assumption of perfect information: individual accounts are valuable because they offer *security* of savings. However, individual accounts may also enable agents to *hide* savings from their spouses, which could be used to manipulate the time path of consumption, or to finance hidden consumption. This theory is similar to the theory of private savings, but rests on the assumption that in equilibrium, spouses are (and benefit from) concealing information from one another. If the benefit of hiding savings is correlated with preference heterogeneity, this could generate the patterns we observe in the data. The next section presents evidence of hidden savings in our data and demonstrates that the empirical results are robust to accounting for hidden savings concerns.

## 1.5 Hidden Information and Account Use

Hidden information appears to be important in households in developing countries. For example, Anderson and Baland (2002) find that women's use of ROSCAs in Kenya is consistent with a model of hidden information. Boozer, Goldstein, and Suri (2009) analyze spousal cross reports of food expenditure in Ghana and find evidence of hidden consumption. Ashraf (2009) finds evidence that the informational environment has significant impact on the investment decisions of spouses with low levels of financial control in the Philippines and de Laat (2008) finds that individuals in split migrant couples in Kenya are willing to expend considerable resources to acquire information about one another. In contrast, Mani (2010) finds that changing the information environment in a lab experiment in India had little impact on spousal investment decisions.

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

If the return to hiding savings from a spouse is increasing in intracouple discount factor heterogeneity, then hidden savings concerns could be responsible for our empirical results (though it is not obvious that such a correlation should exist). We added the "extra statements" treatment to our field experiment in order to assess the importance of hidden savings in our study population, and to gauge if hidden savings concerns are correlated with preference heterogeneity. We only present the extra statements analysis for a subset of couples in the data, as some enumerators did not follow the proper protocol when offering extra statements. Appendix B discusses how we identified protocol problems and the creation of the extra statements analysis sample.

We also make use of baseline spousal cross reports of income and savings device use to identify couples who are most likely to find hiding financial resources beneficial. The assumption here is that couples in which individuals are poorly informed about spousal finances are more likely to value hidden information. To construct an "information sharing" index, we considered responses to baseline survey questions addressing income earned last week, bank accounts, savings at home, SACCOs, and ROSCAs, creating five subindices. The subindices range from 0 (perfect information) to 1 (most misinformation). If cross reports exactly matched own reports, we coded the index to 0. If an individual reported that they did (did not) use a device, but the spouse reported that they did not (did), we coded the index to 1. If a spouse asserted that they did not know if an individual

used a given device, or if they did not know how much savings was in the device, we also coded the index to 1. For other instances where we had an own report and a cross report of the amount (or in the case of ROSCAs, the number of ROSCAs), we coded the index to equal  $\min \left\{ \frac{|own_{ic} - cross_{ic}|}{own_{ic}}, 1 \right\}$ . We then created a household-level information index equal to:

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

where a value of 1 represents a perfectly informed household and a value of 0 represents a poorly informed household. Figure 8 graphs the distribution of the index among the 516 couples for whom we have nonmissing values.<sup>36</sup> We then code a couple as "well informed" if their information index is above the sample median.

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

$$y_{ic} = \beta_0 + \beta_1 es_c + ht'_c \lambda + x'_c \delta + \gamma_s + \varepsilon_{ic} \quad (1.12)$$

where  $y_{ic}$  is the outcome of interest,  $es_c$  indicates that the couple was selected for (or, in some specifications, consented to) extra statements,  $x_c$  is a vector of time preference, demographic, and economic characteristics, and  $\gamma_s$  are session fixed effects. To examine treatment effects by preference heterogeneity and household information sharing, in some specifications we also include the vector  $ht_c$ , which includes a dummy for well matched couples, a dummy for well informed couples, and the interaction of these variables with the extra statements indicator.<sup>37</sup>

The first two columns of Table 9 verify that, as per experimental protocol, the probability of opening an individual account is uncorrelated with extra statement selection. We see that this is in fact the case, and for the remaining specifications we limit our analysis to opened individual accounts that were eligible to be randomly selected for the extra statements offer.

The next two specifications examine extra statement consent rates among open individual accounts. We see that only 60 percent of individuals who were presented with the extra statement offer consented – this suggests that informational concerns are important to a sizeable fraction of couples choosing individual accounts. Column 4 reveals interesting differences in consent rates by household information sharing. Well informed households are 26 percentage points more likely to consent to extra statements, which is statistically significant in Panel A, though additional controls reduce precision.

Columns 5-8 of Table 9 examine the reduced form impact of the extra statement offer on savings rates and average daily balances of open individual accounts. The aggregate estimated impact of ex-

---

<sup>36</sup>We could not construct the index for couples where at least one spouse refused to answer one or more of the relevant questions, or had a missing response for other reasons.

<sup>37</sup>When the information sharing index is missing, we recode the well informed dummy to zero and include a dummy variable indicating that the index is missing, as well an interaction set that parallels any interactions included for the well informed dummy. Simply dropping these couples does not change the results.

tra statements (columns 5 and 7) is relatively small, insignificant, and actually positive upon adding controls. Given the low consent rate to extra statements, this is not very surprising – although the standard errors are large enough that we cannot rule out nontrivial negative impacts of extra statements, these results suggest that the extra statement treatment had little overall impact on savings behavior. However, columns 6 and 8 suggest that this aggregate zero impact may mask differences by information sharing. In particular, the extra statements intervention significantly reduced savings rates for poorly informed and poorly matched couples, while having no impact on well informed and well matched couples. Patterns for average balances are similar, though more imprecise. Since consent was partial, we also present two stage least squares estimates where extra statement consent (and its interactions with the well matched and well informed dummies) are instrumented with random assignment to the extra statement treatment and the relevant interactions. These IV results are presented in columns 9-12. We note that the estimated impacts on poorly matched, poorly informed households are substantial – the baseline specification suggests that consenting to extra statements reduces the probability of saving by 83 percentage points. (This point estimate is implausibly large when compared to the dependent variable mean and is most likely driven by our relatively small sample size. However, the 95 percent confidence interval contains a range of more reasonable estimates).

Overall, the extra statements intervention suggests that the ability to hide savings from other household members is indeed an important concern for at least some households in our sample. Moreover, our information index appears to successfully identify couples for whom hiding information is more valuable. As such, if our earlier results were driven by hidden savings concerns, we would expect to see two things: First, the information sharing index should be lowest among poorly matched couples. Second, properly accounting for information sharing in our previous analyses should dilute the impact of preference heterogeneity. However, the information index is totally uncorrelated with preference heterogeneity (the correlation coefficient between the information index and the absolute value of the difference in individual discount factors is 0.015). This foreshadows our finding that accounting for information sharing has no impact on our results.

We check to see if accounting for information flows changes our findings regarding match quality and joint versus individual account use. Table 10 presents the regression described by equation 1.7 augmented to include the well informed indicator. We make two observations – first, our initial results pertaining to preference heterogeneity are robust to controlling for information flows. Second, the well informed indicator has predictive content for household account choice conditional on preference heterogeneity. As we would expect, better informed households are more likely to use joint accounts and less likely to use individual accounts, though these results are attenuated by additional controls. Moreover, additional specification testing not shown here confirms that information sharing is not driving our results regarding responses to the excess interest rate.

Overall, our results are compatible with the theory that hidden information concerns impact household savings decisions. However, to the extent that these concerns are important, they appear

to be largely orthogonal to preference heterogeneity. This is plausible – hiding savings is likely valuable because it allows individuals to increase their share of consumption, or tilt consumption towards goods that they favor. If the benefit of doing so is equally large for individuals in well matched and poorly matched households, accounting for it should leave our results unchanged, which is what we observe. To complete our discussion of robustness, the following section considers other alternative explanations for the patterns in our data.

## 1.6 Alternative Explanations

Perhaps the most obvious concern with our results is that discount factor heterogeneity is irrelevant for behavior, but correlated with other unobservables (in terms of demographic characteristics and preferences) that drive account choice. Indeed, our theoretical analysis suggests that heterogeneity in other aspects of preferences could generate incentives for strategic behavior, even in the absence of discount factor heterogeneity. While we cannot address unobservables, the robustness of our results to the inclusion of controls accounting for both levels of and intracouple heterogeneity in observables, including education, income, and distance from the bank (through session fixed effects), suggests that this is not the case. Furthermore, it is not clear that discount factors should be strongly correlated with other important aspects of preferences: Dupas and Robinson (2011) elicit estimates of discount factors and risk aversion in a Kenyan sample very similar to our own and find very low correlations between the two measures (0.12 for men and 0.06 for women) (Pascaline Dupas, personal communication).

Another possibility is that individual accounts are desirable because large pools of individual savings can be used to shift the bargaining weight, allowing an agent to secure a greater share of consumption for herself in the future. Similarly, individual accounts may allow agents to make unilateral private consumption choices (or restrict spousal private consumption). If the benefit of individual accounts is correlated with discount factor heterogeneity, this could explain the patterns in our data. Both cases suggest that individual accounts would be more popular than joint accounts, and that both spouses would save simultaneously in their individual accounts. However, we do not find either of these patterns in our data – 29 percent of joint accounts were used for saving during our study horizon, as compared to 11 percent of individual accounts with an interest rate of 2 percent or better. Moreover, only 28 percent of households who opened both individual accounts *and* saved in at least one account saved in both individual accounts. In contrast, the private savings theory predicts that most couples will only make use of one bank account (Proposition 1). Indeed, 89 percent of couples who saved only saved in one account. On the other hand, individual accounts could be valuable even if ownership (without use) increases an individual’s option value outside the marriage. However, in this case it is not clear that savings decisions should be distorted as long as couples remain together.

Furthermore, our results are not well rationalized by mental accounting, which is frequently

cited as motivating intrahousehold resource allocation in the developing world. The theory of mental accounting, as described in Thaler (1990), specifies that households earmark funds saved in different accounts for different types of consumption. Duflo and Udry (2004) present evidence of such behavior among households in Côte d’Ivoire. If households have good outside options for individual investment technologies and poor outside options for joint technologies, mental accounting could help explain why joint accounts were more popular among our study households. However, there is no clear reason why mental accounting would generate differential behavior by preference heterogeneity, particularly with respect to the excess interest rate and investment efficiency.

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

## 1.7 Conclusion

We develop a model of household decision making in which agents can strategically use private savings accounts to manipulate consumption streams. We show that costly strategic action can be beneficial only when individuals in the household have differing rates of time preference. In order to test our model, we conducted a field experiment in Western Kenya, where married couples were recruited to group meetings and given the opportunity to open joint and individual bank accounts with randomly assigned interest rates. By overlaying our model with our experimental design, we derived three testable implications: (1) well matched couples will make more intensive use of joint accounts while poorly matched couples will make more intensive use of individual accounts, (2) well matched couples will take account of relative rates of return on accounts in accordance with Figure 3, and (3) well matched savers earn higher rates of return on their savings than poorly matched savers.

Our results support all three of these predictions. This is, of course, subject to the caveat that we cannot randomly assign preference heterogeneity in the field. As such, we cannot completely rule out the hypothesis that our results are driven by some other omitted characteristic that is correlated with estimates of preference heterogeneity. However, the stability of our results to the inclusion of flexible demographic and economic controls, and to the inclusion of a measure of household information flows, is comforting and suggests that the results we observe are indeed driven by inefficiencies due to conflicting rates of time preference.

We also find some evidence that informational concerns impact account use and account choice among our couples. Couples who are poorly informed about one another's financial activities respond most adversely to a randomized information treatment and are also more likely to gravitate towards individual bank accounts. However, our treatment of hidden information is largely a robustness check. Understanding how information sharing impacts household decision making, particularly from a theoretical perspective, is an area in need of additional research.

Our results add to a growing body of literature that rejects a Pareto efficient model of the household, while presenting a theory that sheds light on specific mechanisms that could be driving this inefficiency. Several recent papers find that individuals make strategic investment decisions when they report little control over finances in the household (Ashraf, Aycinena, Martinez A., and Yang 2010; Chin, Karkoviata, and Wilcox 2010; Ashraf 2009). Our results dovetail with these papers, as perceived "control" could be a function of how closely an agent's preferences align with collective decisions. (Indeed, Table A1 shows that individuals in well matched couples are significantly more likely to state that they make decisions regarding saving – here, perceived decision making power may reflect how closely actual outcomes are aligned with an individual's preferences.)

The private savings theory has interesting implications for financial product design for the poor. In order to achieve the *ex-ante* Pareto efficient allocation, households need a way to commit themselves to binding intertemporal contracts. A savings device that committed future consumption to specific individuals in the household could help enforce such contracts. While designing such a product does not seem particularly feasible, a savings account that committed resources to particular *types* of consumption could approximate this product if members of the household favor different goods. Indeed, women and men in developing countries do appear to have differential preferences (Bobonis 2009; Duflo 2003; Thomas 1994). As such, savings accounts that constrain balances to be used for certain goods such as education, health, or home improvement could be useful tools for improving intrahousehold contracting.

## 1.A Appendix: Proofs

**Proof of Lemma 1.** It is given that  $s_1^i$  is part of a pure strategy Nash equilibrium and that  $s_1^i > 0$ . Now assume that  $s_1^J(s_1^i) > 0$  and that this allocation is a PSNE. We show that this always leads to a contradiction.

If  $s_1^J(s_1^i) > 0$  then the household Euler relation given by equation 1.3 must hold. First consider the case where  $R_i s_1^i \leq w + b^i$ . Consider the deviation  $\hat{s}_1^i = 0$ . Since this increases  $t = 1$  resources, the couple will always save in the joint account after this deviation, so equation 1.3 will continue to hold. In order for the Euler to hold before and after the deviation  $t = 1$  and  $t = 2$  consumption must both be higher when  $\hat{s}_1^i = 0$ . In this case, the original  $s_1^i$  cannot be part of a PSNE because  $\hat{s}_1^i = 0$  makes individual  $i$  strictly better off.

Next consider the case where  $R_i s_1^i > w + b^i$ . We consider the subcases  $R_J \geq R_i$  and  $R_J < R_i$  in turn. When  $R_J \geq R_i$ , consider the deviation  $\hat{s}_1^i = 0$ . The household will continue to save, so equation 1.3 will continue to hold. Since  $R_J \geq R_i$ , this is only possible if consumption in both periods strictly increases (we use Assumption A3 when  $R_J = R_i$ ). Then  $\hat{s}_1^i = 0$  is a profitable deviation, which implies the original  $s_1^i$  was not a Nash best response.

When  $R_i > R_J$ , consider the deviation  $\hat{s}_1^i = s_1^J + s_1^i + b^J$ . It is clear that the couple would choose  $s_1^J(\hat{s}_1^i) = 0$ , so  $t = 1$  consumption would stay the same, while  $t = 2$  consumption would strictly increase. This is a profitable deviation, which implies that the original  $s_1^i$  was not a best response. ■

In order to prove Proposition 1, we first establish two useful lemmas.

**Lemma 2** *Suppose  $\exists$  a pure strategy Nash equilibrium where  $s_1^M > 0$  and  $s_1^F > 0$ . Then  $\delta_M R_M = \delta_F R_F$  and  $c_1^{-\sigma} = \delta_i R_i c_2^{-\sigma}$  for  $i \in \{M, F\}$  must hold in the equilibrium.*

**Proof of Lemma 2.** We will establish the result by showing that for each individual  $i$ , it must be that  $c_1^{-\sigma} = R_i \delta_i c_2^{-\sigma}$ . First, note that for any pure strategy Nash equilibrium where both spouses save, Lemma 1 implies that  $s_1^J(s_1^i) = 0$ .

Suppose instead that  $c_1^{-\sigma} < R_i \delta_i c_2^{-\sigma}$ . Since  $s_1^J(s_1^i) = 0$ , the Euler inequality implies that individual  $i$  would be strictly better off increasing  $s_1^i$ . This contradicts the definition of a PSNE and therefore cannot be the case. Next suppose that  $c_1^{-\sigma} > R_i \delta_i c_2^{-\sigma}$  and the preference to set  $s_1^J = 0$  at the household allocation stage is strict. Then the Euler inequality implies that individual  $i$  would be strictly better off decreasing  $s_1^i$ , which cannot be the case if  $s_1^i$  is part of a PSNE.

Next suppose that  $c_1^{-\sigma} > R_i \delta_i c_2^{-\sigma}$  and the preference to set  $s_1^J = 0$  at the household allocation stage is weak (the couple is just indifferent between  $s_1^J > 0$  and  $s_1^J = 0$ ). Given Assumption A3 it must be that  $c_1^{-\sigma} < R_J \Omega c_2^{-\sigma} \Rightarrow R_J \Omega > R_i \delta_i$ . For this case, we consider  $i \in \{M, F\}$  separately. First consider  $i = F$ . Note  $\delta_F \geq \Omega \Rightarrow R_J > R_F$ . If the wife were to reduce  $s_1^F$  by  $\varepsilon$ , then the couple would strictly prefer to save jointly (and  $t = 2$  consumption would increase). Since  $R_J > R_F$ , the couple would be better off under the deviation. Since  $\delta_F \geq \Omega$  and  $t = 2$  consumption increases, this would be profitable for the wife. Since  $s_1^F$  is part of a PSNE, this cannot be the case, so it must be that  $c_1^{-\sigma} = \delta_F R_F c_2^{-\sigma}$ . Next consider  $i = M$ . Using the result we just established for the wife, it must be that  $c_1^{-\sigma} = \delta_F R_F c_2^{-\sigma} < R_J \Omega c_2^{-\sigma} \Rightarrow R_J > R_F$ . By the same logic we used for the wife's case, it would again be profitable for the wife to reduce  $s_1^F$  by  $\varepsilon$  to induce joint savings. Since  $s_1^F$  is part of a PSNE, this cannot be the case. Then it must be that  $c_1^{-\sigma} = \delta_M R_M c_2^{-\sigma}$ . Combining the result for the husband and the wife we have  $c_1^{-\sigma} = \delta_M R_M c_2^{-\sigma} = \delta_F R_F c_2^{-\sigma}$ , which implies  $\delta_M R_M = \delta_F R_F$ . This completes the proof. ■



**Lemma 3** Suppose  $\Omega R_J < \delta_F R_F = \delta_M R_M$  and  $\tilde{s}_1^F > 0$ . If  $\tilde{s}_1^M > 0$  is a best response given  $\tilde{s}_1^F$ , then the  $\hat{s}_1^M > 0$  that satisfies  $\hat{c}_1^{-\sigma} = R_M \delta_M \hat{c}_2^{-\sigma}$  will be a best response given  $\hat{s}_1^F = 0$ .

**Proof of Lemma 3.** Optimality of  $\tilde{s}_1^M$ , combined with  $\Omega R_J < \delta_M R_M$  implies  $\tilde{c}_1^{-\sigma} = R_M \delta_M \tilde{c}_2^{-\sigma}$  (otherwise the husband would have a profitable deviation available to him). Then when  $\hat{s}_1^F = 0 \exists \hat{s}_1^M > 0$  that satisfies

$$(y_1 - \hat{s}_1^M - b^M)^{-\sigma} = \delta_M R_M (y_2 + R_M \hat{s}_1^M - w - b^M)^{-\sigma} \quad (1.13)$$

It is clear that  $\hat{s}_1^M$  is strictly preferred to any other  $\hat{s}_1^M > 0$ . Then to prove that  $\hat{s}_1^M > 0$  is a best response to  $\hat{s}_1^F = 0$ , we just need to show that  $\hat{s}_1^M > 0$  is strictly preferred to  $s_1^M = 0$ . For the remainder of the proof, we refer to the allocation that results when  $s_1^M = s_1^F = 0$  as the joint alternative/regime, and the allocation that results given  $\hat{s}_1^M > 0$  as the husband's allocation/regime.

Now we establish some notation. Denote period  $t$  consumption under the husband's regime given endowment  $\mathbf{y}$  as  $c_t^M(\mathbf{y})$ . Similarly denote period  $t$  consumption under the joint regime given  $\mathbf{y}$  as  $c_t^H(\mathbf{y})$ . Then the difference in husband's utility under the two regimes can be written as

$$\Delta_M(\mathbf{y}) \equiv \frac{c_1^M(\mathbf{y})^{1-\sigma}}{1-\sigma} + \delta_M \frac{c_2^M(\mathbf{y})^{1-\sigma}}{1-\sigma} - \frac{c_1^H(\mathbf{y})^{1-\sigma}}{1-\sigma} - \delta_M \frac{c_2^H(\mathbf{y})^{1-\sigma}}{1-\sigma}$$

Next, note that changing the wife's strategy from  $\tilde{s}_1^F$  to  $\hat{s}_1^F = 0$  is isomorphic (from the perspective of the husband) to switching from endowment  $\mathbf{y}^0 = [y_1^0, y_2^0]'$  (where  $y_1^0 = y_1 - s_1^F - b^F$  and  $y_2^0 = y_2 + \max\{R_F s_1^F - b^F - w, 0\}$ ) to endowment  $\mathbf{y}^1 = [y_1^1, y_2^1]'$  (where  $y_1^1 = y_1$  and  $y_2^1 = y_2$ ). Note  $y_1^1 > y_1^0$  and  $y_2^1 \leq y_2^0$ . It is clear that  $\frac{\partial s_1^F}{\partial y_1} \geq 0$  and  $\frac{\partial s_1^F}{\partial y_2} \leq 0$ : this implies that saving under both the husband's and joint regimes must be higher at  $\mathbf{y}^1$  as compared to  $\mathbf{y}^0$  (weakly so, for the joint regime).

Since  $\tilde{s}_1^M$  is a best response,  $\Delta_M(\mathbf{y}^0) \geq 0$ . To complete the proof, we will show that  $\Delta_M(\mathbf{y}^0) \geq 0 \Rightarrow \Delta_M(\mathbf{y}^1) > 0$ . When  $\Delta_M(\mathbf{y}^0) \geq 0$  there are two cases: (a) the husband's allocation generates weakly higher consumption in both periods at  $\mathbf{y}^0$ , and (b) the husband's allocation generates strictly higher consumption in one period and strictly lower consumption in the other period at  $\mathbf{y}^0$ .

First consider case (a). Since desired savings under both regimes is higher given  $\mathbf{y}^1$  and since  $\Omega R_J < \delta_M R_M \Rightarrow R_M > R_J$ , it must be that  $\Delta_M(\mathbf{y}^1) > 0$ . (Noting that increasing  $s_1^M$  by the increase in joint savings would yield more consumption in both periods relative to the joint alternative makes this clear). This implies  $\Delta_M(\mathbf{y}^1) > 0$  in case (a). Next consider case (b). It must be that  $c_1^M(\mathbf{y}^0) < c_1^H(\mathbf{y}^0)$  and  $c_2^M(\mathbf{y}^0) > c_2^H(\mathbf{y}^0)$  (Euler equations cannot support the other possibility). Then (with some abuse of notation, as  $\Delta_M(\mathbf{y})$  will jump discontinuously at one point if  $\delta_M < \Omega$  and  $s_1^J(\mathbf{y}^0) = 0$  but  $s_1^J(\mathbf{y}^1) > 0$ ):

$$\Delta_M(\mathbf{y}^1) = \Delta_M(\mathbf{y}^0) + \int_{y_1^0}^{y_1^1} \frac{\partial}{\partial y_1} \Delta_M(y_1, y_2^0) dy_1 - \int_{y_2^1}^{y_2^0} \frac{\partial}{\partial y_2} \Delta_M(y_1^1, y_2) dy_2$$

Suppose the couple switches from  $s_1^J = 0$  to  $s_1^J > 0$  at some  $\mathbf{y}^*$ . If  $\delta_M < \Omega$  then  $\Delta_M(\mathbf{y})$  will discontinuously increase at  $\mathbf{y}^*$ . Since the presence of this point only helps our argument, we ignore it going forward. Then for all non-discontinuous points we have (using the Envelope Theorem for

the husband's private savings response):

$$\begin{aligned}\frac{\partial}{\partial y_1} \Delta_M(y_1, y_2^0) &= [c_1^M(y_1)^{-\sigma} - c_1^H(y_1)^{-\sigma}] + \frac{\partial s_1^J}{\partial y_1} [c_1^H(y_1)^{-\sigma} - \delta_M R_J c_2^H(y_1)^{-\sigma}] \\ -\frac{\partial}{\partial y_2} \Delta_M(y_1^1, y_2) &= \delta_M [c_2^H(y_2)^{-\sigma} - c_2^M(y_2)^{-\sigma}] - \frac{\partial s_1^J}{\partial y_2} [c_1^H(y_2)^{-\sigma} - \delta_M R_J c_2^H(y_2)^{-\sigma}]\end{aligned}$$

Whenever  $s_1^J > 0$ ,  $\delta_M \leq \Omega \Rightarrow (c_1^H)^{-\sigma} \geq \delta_M R_J (c_2^H)^{-\sigma}$ . Then the second bracketed terms in each equation will always work to make private savings more attractive, so we ignore them going forward.

Begin at  $\mathbf{y}^0$  and consider increasing  $y_1$  from  $y_1^0$  to  $y_1^1$ . Since  $c_1^M(\mathbf{y}^0) < c_1^H(\mathbf{y}^0)$ ,  $\frac{\partial}{\partial y_1} \Delta_M(y_1^0, y_2^0) > 0$ . If  $c_1^M(\mathbf{y}) \leq c_1^H(\mathbf{y}) \forall y_1 \in [y_1^0, y_1^1]$  then  $\Delta_M(y_1^1, y_2^0) > 0$ . Suppose instead that at some  $y_1^*$   $c_1^M(y_1^*, y_2^0) > c_1^H(y_1^*, y_2^0)$ . For this to be the case, the joint alternative must involve  $s_1^J > 0$ , which implies

$$\begin{aligned}c_1^M(y_1^*, y_2^0)^{-\sigma} &= \delta_M R_M c_2^M(y_1^*, y_2^0)^{-\sigma} < c_1^H(y_1^*, y_2^0)^{-\sigma} = \Omega R_M c_2^H(y_1^*, y_2^0)^{-\sigma} \\ &\Rightarrow c_2^M(y_1^*, y_2^0) > c_2^H(y_1^*, y_2^0) \text{ and } c_1^M(y_1^*, y_2^0) > c_1^H(y_1^*, y_2^0)\end{aligned}$$

But if consumption is higher in both periods at  $(y_1^*, y_2^0)$  under the husband's allocation, he will prefer saving privately at  $(y_1^1, y_2^0)$  and  $\mathbf{y}^1$  as well. So it must be that  $\Delta_M(y_1^1, y_2^0) > 0$ . Consider the case where at  $(y_1^1, y_2^0)$   $c_1^M < c_1^H$  and  $c_2^M > c_2^H$  (the only other possibility is  $c_1^M > c_1^H$  and  $c_2^M > c_2^H \Rightarrow \Delta_M(\mathbf{y}^1) > 0$  in which case we are done). Then  $-\frac{\partial}{\partial y_2} \Delta_M(y_1^1, y_2^0) > 0$ . This derivative will remain positive unless there exists  $y_2^*$  such that  $c_2^M(y_2^*) < c_2^H(y_2^*)$ . In this case Eulers imply:

$$c_1^M(y_2^*)^{-\sigma} = \delta_M R_M c_2^M(y_2^*)^{-\sigma} > c_1^H(y_2^*)^{-\sigma} = \Omega R_J c_2^H(y_2^*)^{-\sigma}$$

which implies  $\Delta_M(y_1^1, y_2^*) < 0$ . Let  $\hat{y}_2 = \inf\{y_2^* : c_2^M(y_2^*) < c_2^H(y_2^*)\}$ . Recall that if the couple switches from  $s_1^J = 0$  to  $s_1^J > 0$  at  $\hat{y}_2$ ,  $\Delta_M(\mathbf{y})$  will discontinuously increase. Otherwise,  $\Delta_M(\mathbf{y})$  will be continuous at  $\hat{y}_2$ . In the continuous case, note that  $\delta_M R_M > \Omega R_J$  and  $c_2^M(\hat{y}_2) = c_2^H(\hat{y}_2) \Rightarrow c_1^M(\hat{y}_2) < c_1^H(\hat{y}_2)$ . In either case this implies that  $\exists \varepsilon > 0$  such that  $\Delta_M(y_1^1, \hat{y}_2 - \varepsilon) < 0$ , which is impossible. Then  $\nexists y_2^*$  as above, which implies that  $\Delta_M(\mathbf{y}^0) \geq 0 \Rightarrow \Delta_M(\mathbf{y}^1) > 0$  in cases (a) and (b). ■

We are now prepared to establish Proposition 1.

**Proof of Proposition 1.** By Lemma 1, no PSNE will involve  $s_1^J > 0$  and at least one  $s_1^i > 0$ . Then the second part of the proposition follows immediately from Lemma 2.

We now establish the first part. Suppose  $\exists$  a PSNE in which  $s_1^M > 0$  and  $s_1^F > 0$ . By Lemma 1  $\delta_M R_M = \delta_F R_M$  and  $c_1^{-\sigma} = R_i \delta_i c_2^{-\sigma}$  for  $i \in \{M, F\}$ . By Assumption A2, this equilibrium will never be chosen by any couple as long as  $\exists$  some other Pareto dominant PSNE. We divide the problem into 4 cases: (a)  $\Omega = \delta_M = \delta_F$ , (b)  $\Omega R_J = \delta_i R_i$  and  $\delta_M < \delta_F$ , (c)  $\Omega R_J > \delta_i R_i$ , and (d)  $\Omega R_J < \delta_i R_i$  and show that in each case a Pareto dominant PSNE exists.

In case (a)  $R_M = R_J = R_F$ . It is clear that  $s_1^M = s_1^F = 0$  and  $s_1^J = s_1^J(\mathbf{0}) > 0$  will be a Pareto dominant PSNE (since  $b^J \leq b^i$  strictly fewer transaction costs are paid when saving in just the joint account so consumption will be higher in both periods).

In cases (b) and (c),  $R_F < R_M$  and  $R_F < R_J$ . Note that in order for the original  $s_1^F$  to be a Nash best response, it must be that the couple would choose  $s_1^J = 0$  if the wife deviated to  $s_1^F = 0$ . (Since  $\Omega R_J \geq \delta_F R_F$  this would have to be due to transaction costs. If the couple chose  $s_1^J > 0$  in response to the deviation, second period consumption would increase and the wife would be strictly

better off). Now consider the alternative allocation  $\hat{s}_1^F = 0$  and  $\hat{s}_1^M$  set according to equation 1.13. Since equation 1.13 holds under the original and the proposed alternative, and since  $R_M > R_F$ , it must be that  $\hat{c}_1 > c_1$  and  $\hat{c}_2 > c_2$ . If  $\hat{s}_1^M$  is a best response for the husband, then  $s_1^F = 0$ ,  $s_1^M = \hat{s}_1^M$ ,  $s_1^J = 0$  will be a Pareto dominant PSNE (it is clear that  $s_1^J = 0$  will be a best response for the couple given  $\hat{s}_1^M$  since the original PSNE implies  $s_1^J(s_1^M, 0) = 0$ ). Clearly  $\hat{s}_1^M$  is preferred to any other positive  $s_1^M$ , so if  $\hat{s}_1^M \succ_M s_1^M = 0$  we are done with cases (b) and (c). Suppose instead that  $s_1^M = 0 \succ_M \hat{s}_1^M$ . Then it must be that  $s_1^J(\mathbf{0}) > 0$  – denote the resulting allocation  $(\tilde{c}_1, \tilde{c}_2)$ . Since  $\Omega R_J \geq \delta_i R_i$ , Eulers show that  $\tilde{c}_2 < \hat{c}_2 \Rightarrow \tilde{c}_1 < \hat{c}_1$ , so  $(\tilde{c}_1, \tilde{c}_2) \succ_M (\hat{c}_1, \hat{c}_2) \Rightarrow \tilde{c}_2 > \hat{c}_2$ . Then  $\delta_F \geq \delta_F \Rightarrow (\tilde{c}_1, \tilde{c}_2) \succ_F (\hat{c}_1, \hat{c}_2)$ , which implies  $s_1^M = s_1^F = 0$  and  $s_1^J = s_1^J(\mathbf{0})$  is a Pareto dominant PSNE.

For case (d) take the alternative allocation  $\hat{s}_1^F = \hat{s}_1^J = 0$ ,  $\hat{s}_1^M > 0$  according to equation 1.13. Then Assumption A3 and  $R_M \geq R_F$  imply that consumption in both periods must be higher under this alternative allocation (otherwise the husband's Euler relation would not be maintained). Lemma 3 implies that  $\hat{s}_1^M$  is a best response for the husband. By Lemma 2, the alternative allocation results in  $\hat{c}_1^{-\sigma} = \delta_F R_F \hat{c}_2^{-\sigma} > \Omega R_J \hat{c}_2^{-\sigma}$ . This implies that  $\hat{s}_1^F = 0$  is a best response for the wife, and that  $\hat{s}_1^J = 0$  is optimal for the couple. Then the alternative allocation is a Pareto dominant PSNE, and involves only one account in use. ■

**Proof of Proposition 2.** When  $\delta_M = \delta_F = \Omega$ , the *ex-ante* Pareto efficient consumption sharing rules will be time invariant and given by  $c_t^M = \rho c_t$  and  $c_t^F = (1 - \rho) c_t$ . Substituting this into the efficient bargaining problem (EB) yields the following maximand:

$$\left( \mu \rho^{1-\sigma} + (1 - \mu) (1 - \rho)^{1-\sigma} \right) \sum_{t=1}^T \delta^{t-1} \left( \frac{c_t^{1-\sigma}}{1 - \sigma} \right)$$

Note that any savings allocation  $\{s_t^M, s_t^F, s_t^J\}_{t=1}^{T-2}$  that maximizes this term *also* maximizes  $\tilde{U}_t^M$  and  $\tilde{U}_t^F$  in both  $t = 1$  and  $t = 2$ . Then by definition of an optimum, there is no deviation from this savings plan that makes either individual strictly better off, so the solution to (EB) will be sustainable as a pure strategy Nash equilibrium to the private savings game. By definition of an optimum, it also follows that any other pure strategy equilibrium that does not solve (EB) will result in strictly less utility for both agents. Then by Assumption A2, the allocation chosen by the household in the private savings game must correspond to an *ex-ante* Pareto efficient allocation. ■

**Proof of Proposition 3.** Suppose  $\tilde{s}_1^i$  results in the same consumption allocation as  $s_1^i = 0$ . For the proof, we only consider  $\tilde{A}_1^i = A_1^i \setminus \tilde{s}_1^i$  and  $\tilde{D}_1^i = D_1^i \setminus \tilde{s}_1^i$ , as a private savings choice that is equivalent to  $s_1^i = 0$  will never lead to a strict preference for private savings, so we need not consider these choices. Note that the presence of transaction costs implies that if such an  $\tilde{s}_1^i$  exists,  $\exists \tilde{s}_1^i$  that results in strictly more (or less) consumption in agent  $i$ 's preferred period, so  $A_1^i \neq \emptyset \Rightarrow \tilde{A}_1^i \neq \emptyset$  and  $D_1^i \neq \emptyset \Rightarrow \tilde{D}_1^i \neq \emptyset$ . Denote the  $t$  period consumption allocation that arises given  $s_1^i$  and  $s_1^{-i} = 0$  as  $c_t(s_1^i)$ . Then define  $\tilde{U}_i(s_1^i, \delta_i) \equiv u_1(c_1(s_1^i)) + \delta_i u_2(c_2(s_1^i))$ . Then  $\tilde{U}_i(0, \delta_i)$  is the utility individual  $i$  would receive in the absence of private savings. Then  $\forall s_1^i \in \tilde{A}_1^i$

$$\frac{\partial \Delta_i(s_1^i, \delta_i)}{\partial \delta_i} \equiv \frac{\partial [\tilde{U}_i(s_1^i, \delta_i) - \tilde{U}_i(0, \delta_i)]}{\partial \delta_i} = u_2(c_2(s_1^i)) - u_2(c_2(0)) \quad (1.14)$$

This term is always positive for elements of  $\tilde{A}_1^F$  and negative for elements of  $\tilde{A}_1^M$ , unless  $s_1^M$  results in more consumption in both periods, in which case  $\Delta_M(s_1^M, \delta_i) > 0 \forall \gamma \geq 1$ . Then either

$\tilde{U}_i(s_1^i, \Omega) - \tilde{U}_i(0, \Omega) \geq 0$  or  $\tilde{U}_i(s_1^i, \Omega) - \tilde{U}_i(0, \Omega) < 0$ . In the former case, this implies that  $s_1^i$  is preferred to  $\hat{s}_1^i = 0 \forall \delta_F > \Omega$  ( $i = F$ ) or  $\delta_M < \Omega$  ( $i = M$ ). Then for the wife, our thresholds are  $\rho^* = 0$  and  $\gamma_F^*(\rho) = 1 \forall \rho$ . For the husband, the threshold is  $\gamma_M^* = 1$ .

In the latter case  $\exists \tilde{\delta}_i(s_1^i)$  that solves  $\tilde{U}_i(s_1^i, \tilde{\delta}_i) - \tilde{U}_i(0, \tilde{\delta}_i) = 0$ :  $\tilde{\delta}_i(s_1^i) = \frac{u_1(c_1(0)) - u_1(c_1(s_1^i))}{u_2(c_2(s_1^i)) - u_2(c_2(0))}$ . Since  $\tilde{U}_i(s_1^i, \Omega) - \tilde{U}_i(0, \Omega) < 0$ , by the definition of  $\tilde{A}_1^i$  it must be that the numerator and denominator are of the same sign (otherwise the private savings choice would be strictly preferred for all  $\delta_i$ ). Then our derivatives and initial conditions imply that for husbands  $0 < \tilde{\delta}_M(s_1^i) < \Omega$  and for wives  $\tilde{\delta}_F(s_1^i) > \Omega$ . Then define:  $\delta_M^* \equiv \sup \{ \tilde{\delta}_M(s_1^M) : s_1^M \in \tilde{A}_1^M \}$  and  $\delta_F^* \equiv \inf \{ \tilde{\delta}_M(s_1^M) : s_1^F \in \tilde{A}_1^F \}$ . For husbands, we then set  $\gamma_M^* = \frac{\Omega}{\delta_M^*}$ . For wives we must address two parameters. The lower bound  $\rho^*$  solves  $\frac{\Omega}{1-\rho^*} = \delta_M^*$ . Then for each  $\rho > \rho^*$ ,  $\gamma_F^*(\rho) = \frac{\Omega\rho}{(\Omega - \delta_M^*(1-\rho))}$ .

We can repeat this exercise with the set  $\tilde{D}_1^i/\tilde{A}_1^i$ . Note that if  $\tilde{U}_i(s_1^i, \Omega) - \tilde{U}_i(0, \Omega) \leq 0$ ,  $s_1^i$  will never be profitable for any individual where  $\delta_i \neq \Omega$ . In this case, define  $\gamma_M^{**} = 1$  or  $\gamma_F^{**}(\rho) = 1 \forall \rho$ . When  $\tilde{U}_i(s_1^i, \Omega) - \tilde{U}_i(0, \Omega) > 0$ , the mechanics are essentially the same as the argument above, except the signs are reversed. Since  $s_1^i$  is preferred at  $\delta_i = \Omega$ , we can find  $\gamma_F^{**}(\rho) \forall \rho \in (0, 1)$ . When  $s_1^i \in \tilde{A}_1^i \cap \tilde{D}_1^i$ , private savings increases consumption in both periods and will be preferred to the household alternative. In this case, let  $\gamma_M^{**} \equiv \infty$  or  $\gamma_F^{**}(\rho) \equiv \infty \forall \rho \in (0, 1)$ . These thresholds meet the initial requirements of the proposition, which therefore completes the proof. ■

**Proof of Proposition 4.** For parts (a) and (b), there are two possibilities, either the Nash best response at  $\delta_i(\underline{\gamma})$ ,  $s_1^{i*} > 0$ , results in the same consumption allocation as  $s_1^i = 0$ , or it does not. When it does, the result follows trivially: individual  $i$  could always save  $s_1^{F*} + \varepsilon$  (for the wife) or  $s_1^{M*} - \varepsilon$  (for the husband). Since there is some transaction cost associated with the joint account (Assumption A3), we can always find  $\varepsilon$  small enough such that the resulting allocation is strictly preferred to  $s_1^i = 0 \forall \delta_i(\underline{\gamma})$  with  $\gamma > \underline{\gamma}$ . To complete the proofs of (a) and (b) we just establish the result for the latter possibility ( $s_1^{i*}$  results in a different allocation).

First consider (a). We will show that any  $s_1^{F*} > 0$  that is a Nash best response of the wife, is an element of  $A_1^F$  by contradiction. Suppose instead that  $s_1^{F*} \in D_1^F \setminus A_1^F$ . Then  $c_1(s_1^{F*}) > c_1(s_1^J(0)) \Rightarrow s_1^J(0) > 0$ . Since  $\delta_F \geq \Omega$ , an agent with discount factor  $\Omega$  would also prefer  $s_1^{F*}$  to  $s_1^F = 0$ . Since  $b^F \geq b^J$  we have  $R_F \geq \tilde{R}_F(R_J) \geq R_J$ . This implies the following Euler relationship:

$$c_1(s_1^{F*})^{-\sigma} < c_1(s_1^J(0))^{-\sigma} = R_J \Omega c_2(s_1^J(0))^{-\sigma} \leq R_F \delta_F(\gamma) c_2(s_1^{F*})^{-\sigma}$$

which is not possible if  $s_1^{F*}$  is a best response: Lemma 1 implies that  $s_1^J(s_1^{F*}) = 0$ , so the above inequality implies that the wife could make herself even better off by increasing  $s_1^F$ . This is a contradiction, so we conclude that any optimal  $s_1^{F*}$  must be such that  $c_2(s_1^{F*}) > c_2(s_1^J(0))$ . Then Proposition 3 implies the result.

Now consider (b). Suppose  $s_1^{M*} \in D_1^M \setminus A_1^M$ . Since  $R_M < \tilde{R}_M(R_J)$ , it must be that an agent with discount factor  $\Omega$  strictly prefers the allocation under the joint alternative. This implies that

$$\frac{c_1(s_1^J(0))^{1-\sigma}}{1-\sigma} - \frac{c_1(s_1^{M*})^{1-\sigma}}{1-\sigma} > \Omega \left[ \frac{c_2(s_1^{M*})^{1-\sigma}}{1-\sigma} - \frac{c_2(s_1^J(0))^{1-\sigma}}{1-\sigma} \right]$$

but since  $\delta_M \leq \Omega$ , it is clear that  $s_1^{M*}$  would not be a Nash best response. Since all profitable  $s_1^{M*} \in A_1^M$ , Proposition 3 implies the result.

Finally consider (c). It suffices to illustrate a case where use of the individual account is non-monotonic in  $\gamma$ . Suppose that  $R_M > \tilde{R}_M(R_J)$  but there is no  $s_1^M$  that can deliver weakly more

consumption in both periods relative to the joint alternative (i.e.  $A_1^M \cap D_1^M = \emptyset$ ) – note that when  $b^M = b^J$  this scenario is never possible. This implies that if the husband sets  $s_1^M$  such that first period consumption were equal to that under the joint alternative, second period consumption under  $s_1^M$  would be less than that under the joint alternative:

$$R_M (s_1^J(0) - b^M + b^J) - b^M < R_J s_1^J(0) - b^J \quad (1.15)$$

When  $\delta_M = \Omega$ , the husband will strictly prefer saving privately to  $s_1^M = 0$  (this follows from  $R_M > \tilde{R}_M(R_J)$ ), and it must be that  $c_2(s_1^{M*}) > c_2(s_1^J(0))$  under the Nash best response,  $s_1^{M*}$ . Otherwise we would have

$$(c_1(s_1^{M*}))^{-\sigma} < (c_1(s_1^J(0)))^{-\sigma} = \Omega R_J (c_2(s_1^J(0)))^{-\sigma} < \Omega R_M (c_2(s_1^{M*}))^{-\sigma}$$

which contradicts optimality of  $s_1^{M*}$ . So when  $\delta_M = \Omega$ ,  $s_1^{M*} \in D_1^M$ .

Since  $A_1^M \cap D_1^M = \emptyset$  and saving privately is strictly profitable at  $\gamma = 1$ , Proposition 3 implies that  $\exists$  finite  $\underline{\gamma}$  below which the husband prefers to save privately. If  $A_1^M = \emptyset$  we will see downward monotonicity. However, if  $A_1^M$  is not empty, Proposition 3 (part 2) implies that private savings is again preferred by the husband for all  $\gamma > \tilde{\gamma}$ .

It is straightforward to construct an example where  $\tilde{\gamma} > \underline{\gamma}$ . Note that when  $\gamma = 1$ ,  $s_1^{M*} > s_1^J(0)$  and  $s_1^{M*}$  must satisfy the husband's Euler relation. As  $\gamma$  increases,  $\delta_M$  will decrease, so the  $s_1^{M*}$  that satisfies the husband's Euler will decrease continuously. Suppose parameters are such that at  $\tilde{\gamma}$ , the  $s_1^{M*}$  satisfying the husband's Euler is such that  $s_1^{M*} = s_1^J(0)$ , and that he would prefer the savings allocation  $\{s_1^J = 0, s_1^F = 0, s_1^M = s_1^{M*}\}$  to consuming endowments. Then equation 1.15 implies that the husband would be even better off under the household allocation, so at  $\tilde{\gamma}$  his best response is  $s_1^M = 0$ . Then it must be that  $\underline{\gamma} < \tilde{\gamma} < \bar{\gamma}$ . ■

**Proof of Proposition 5.** Note  $\gamma_i^*(R_i) > 1 \Rightarrow A_1^i(R_i) \cap D_1^i(R_i) = \emptyset$ . By construction, individual  $i$  is indifferent between saving privately and the household alternative at  $\gamma_i^*(R_i)$  (this follows immediately from continuity of utility in  $\delta_i$ ). Then Lemma 1 implies that the best possible private savings choice at  $\gamma_i^*(R_i)$ ,  $s_1^i(\gamma_i^*(R_i))$ , must be such that  $s_1^J(s_1^i(\gamma_i^*(R_i))) = 0$ . Now increase  $R_i$  to  $R'_i$  and leave  $s_1^i$  unchanged. Then we still have  $s_1^J(s_1^i) = 0$ , since increasing  $R_i$  decreases second period marginal utility. Furthermore, individual  $i$ 's utility of  $s_1^i(\gamma_i^*(R_i))$  must strictly increase (since second period consumption increases and first period consumption remains the same), while the utility of the joint alternative remains the same. This implies  $\gamma_i^*(R'_i) < \gamma_i^*(R_i)$ .

Now consider  $\gamma_M^{**}(R_M)$ . The proof of Proposition 4 illustrates that finite  $\gamma_M^{**}(R_M)$  implies  $R_M \geq \tilde{R}_M(R_J) \geq R_J$  and  $A_1^M(R_M) \cap D_1^M(R_M) = \emptyset$ . Then we can apply the same logic as above to show that  $\gamma_M^{**}(R'_M) > \gamma_M^{**}(R_M)$ . ■

**Proof of Proposition 6.** First consider  $\gamma_F^*(R_J)$ . Note  $\gamma_F^*(R_J) > 1 \Rightarrow A_1^F(R_J) \cap D_1^F(R_J) = \emptyset$ . By construction, the wife is indifferent between saving privately and the household alternative at  $\gamma_F^*(R_J)$ . By Lemma 1,  $s_1^J(s_1^F(\gamma_F^*(R_J))) = 0$  at  $R_J$ . Then  $s_1^J(s_1^F(\gamma_F^*(R_J))) = 0 \forall R'_J < R_J$ . It follows that  $\forall \gamma \geq \gamma_F^*(R_J)$ , the wife's utility under the private savings regime must be weakly increasing as  $R_J$  decreases. Moreover, the wife's utility under the household alternative will always fall when  $R_J$  falls. When  $\frac{\partial s_1^J}{\partial R_J} > 0$ ,  $\delta_F \geq \Omega$  and the fact that household utility strictly increases is sufficient to imply this. When  $\frac{\partial s_1^J}{\partial R_J} < 0$ , joint savings must change so that consumption in each period decreases as  $R_J$  decreases (otherwise the joint Euler would not hold), which is sufficient for the implication. This implies that a wife who is indifferent between saving privately and the household allocation at  $R_J$  will strictly prefer to save privately at  $R'_J$ . This implies that  $\gamma_F^*(R_J) > \gamma_F^*(R'_J)$ .

Now consider the husband. It suffices to illustrate that  $\gamma_M^*(R_J)$  and  $\gamma_M^{**}(R_J)$  can increase or decrease with  $R_J$  by example. By the same argument used for the wife, the husband's utility of saving privately must be weakly increasing as  $R_J$  declines. Assume the couple strictly prefers to save in the joint account when  $s_1^M = 0$ . Then decreasing  $R_J$  has the following impact on the husband's utility of the joint savings allocation:

$$-\frac{\partial s_1^J}{\partial R_J} \left[ - (c_1 (s_1^J (0)))^{-\sigma} + \delta_M R_J (c_2 (s_1^J (0)))^{-\sigma} \right] - s_1^J \delta_M (c_2 (s_1^J (0)))^{-\sigma} \quad (1.16)$$

When  $\frac{\partial s_1^J}{\partial R_J} \leq 0$  this is always negative. However, when  $\frac{\partial s_1^J}{\partial R_J} > 0$  this will be positive for sufficiently small  $\delta_M$ . By Proposition 5,  $\gamma_M^*$  is decreasing with  $R_M$  given  $R_J$ . Then we can choose parameters so that equation 1.16 is either positive or negative at the  $\delta_M$  implied by  $\gamma_M^*(R_J)$ . We can also choose parameters so that the  $s_1^M > 0$  that is optimal at  $\gamma_M^*(R_J)$  and  $R_J$  satisfies the husband's Euler equation. Then the movement of  $\gamma_M^*(R_J)$  will be given by the sign of equation 1.16. We can perform an analogous exercise with  $\gamma_M^{**}(R_J)$ . ■

## 1.B Appendix: Extra Statements Sample

As described in Section 1.3, 50 percent of couples (who attended our seventh experimental session or above) were sampled for an "extra statements" offer. In order to keep selection into individual account opening constant between treatment and control, our experimental protocol dictated that the extra statement offer only be made to participants *after* they decided which accounts to open. However, extra statement provision is significantly, negatively correlated with the probability of opening an individual account in our sample. While our enumerators never reported informing couples about the extra statements before account choice, nor did they report any cases where couples changed their minds about opening individual accounts after getting an extra statements offer, they were able to observe whether or not a couple was randomly selected for extra statements in a subset of our sessions. The correlation is only significant in this subset of sessions, so we conjecture that some enumerators guided selected couples to joint accounts, as filling out the extra statement cards involved time consuming paperwork. To address this concern, we ran the following regression among all individual accounts in the suspect sessions enumerator by enumerator:

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

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

Out of 14 enumerators, the coefficient on  $es_c$  was negative and significant for just four enumerators. We dropped observations for these four enumerators in the sessions where enumerators could observe extra statement selection prior to the couple's account opening decision. All told, we dropped 366 of 1,000 individual account observations. Columns 1 and 2 of Table 9 verify that account opening is uncorrelated with extra statement selection once we drop the suspect enumerators.

**1.C Appendix: Tables and Figures**

Table 1-1: Demographic Characteristics of Study Sample

	Husbands	Wives	Difference	N
Age	41.8 [13.3]	35.2 [11.8]	6.53*** (0.726)	1194
Education	8.06 [3.57]	5.91 [4.02]	2.14*** (0.220)	1189
Literate	0.864 [0.343]	0.668 [0.471]	0.196*** (0.0238)	1194
Number Children	4.84 [2.93]	4.48 [2.54]	0.358** (0.159)	1194
Subsistence Farmer or No Job	0.406 [0.491]	0.467 [0.499]	-0.0615** (0.0287)	1189
Income Last Week	1348 [2749]	798 [1618]	549*** (132)	1162
Owns Mobile Phone	0.487 [0.500]	0.426 [0.495]	0.0615** (0.0289)	1189
Participates in ROSCA	0.491 [0.500]	0.657 [0.475]	-0.166*** (0.0282)	1194
Has Bank Account	0.303 [0.460]	0.119 [0.324]	0.184*** (0.0230)	1194
Savings in Bank Account (Among Savers)	10217 [18695]	4495 [5680]	5722*** (1702)	208
Saves at Home	0.851 [0.357]	0.898 [0.303]	-0.0472*** (0.0192)	1193
Savings at Home (Among Savers)	1209 [2551]	858 [2923]	351** (171)	1023
Consumption - I Decide	0.470 [0.500]	0.170 [0.376]	0.299*** (0.0257)	1187
Consumption - Spouse Decides	0.0758 [0.265]	0.378 [0.485]	-0.302*** (0.0227)	1187
Consumption - Decide Together	0.389 [0.488]	0.376 [0.485]	0.0128 (0.0282)	1187
Consumption - Decide Alone	0.0152 [0.122]	0.0337 [0.181]	-0.0186** (0.00896)	1187
Savings - I Decide	0.355 [0.479]	0.492 [0.500]	-0.137*** (0.0284)	1187
Savings - Spouse Decides	0.439 [0.497]	0.282 [0.450]	0.157*** (0.0275)	1187
Savings - Decide Together	0.106 [0.308]	0.103 [0.304]	0.00284 (0.0178)	1187
Savings - Decide Alone	0.0773 [0.267]	0.0997 [0.300]	-0.0224 (0.0165)	1187
Weekly Discount Factor	0.801 [0.227]	0.793 [0.215]	0.00776 (0.0128)	1194
Distance from Family Bank (Miles)	3.78 [2.21]	3.78 [2.21]		1194

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



Table 1-2: Correlations Between Estimated Discount Factors and Socioeconomic Characteristics

	Nonlinear Least Squares			Ad Hoc
	All	Husbands	Wives	All
Female	-0.0133 (0.0156)			-0.0205 (0.0207)
Age	-0.000853 (0.000704)	-0.000798 (0.00100)	-0.000975 (0.000962)	-0.00113 (0.000933)
Education	0.00155 (0.00235)	0.00175 (0.00337)	0.00264 (0.00354)	0.00351 (0.00300)
Literate	-0.0268 (0.0212)	-0.0287 (0.0350)	-0.0196 (0.0294)	-0.0348 (0.0280)
Number Children	-0.00531 (0.00326)	-0.0105** (0.00463)	0.000503 (0.00429)	-0.00617 (0.00432)
Subsistence Farmer or No Job	0.00718 (0.0144)	0.0142 (0.0229)	0.0144 (0.0191)	0.00798 (0.0187)
Income Last Week <sup>+</sup>	0.0854 (0.259)	-0.00133 (0.251)	0.412 (0.585)	0.210 (0.326)
Owns Phone	-0.000766 (0.0142)	0.0177 (0.0206)	-0.00447 (0.0188)	0.00115 (0.0184)
Consumption - I Decide	0.0236 (0.0186)	0.0378 (0.0392)	0.0290 (0.0247)	0.0334 (0.0248)
Consumption - Decide Together	0.00132 (0.0194)	0.0317 (0.0420)	-0.0112 (0.0232)	0.00321 (0.0259)
Consumption - Decide Alone	-0.0140 (0.0438)	-0.00205 (0.0921)	-0.0458 (0.0549)	-0.0138 (0.0564)
Consumption - Other	-0.00926 (0.0376)	-0.00565 (0.0614)	-0.00121 (0.0536)	-0.00939 (0.0498)
ROSCA Member	0.00342 (0.0145)	-0.000996 (0.0201)	0.00833 (0.0197)	0.00296 (0.0188)
Has Bank Account	-0.0160 (0.0195)	0.00525 (0.0249)	-0.0661* (0.0348)	-0.0354 (0.0258)
Saves at Home	0.0384* (0.0205)	0.0429 (0.0289)	0.0356 (0.0320)	0.0393 (0.0270)
Has SACCO Account	0.0483 (0.0317)	0.0381 (0.0355)	0.232*** (0.0848)	0.0696* (0.0404)
Has MPESA Account	-0.0489*** (0.0190)	-0.0677*** (0.0254)	-0.0293 (0.0330)	-0.0459* (0.0237)
Saves Other Ways	0.0166 (0.0153)	0.0300 (0.0234)	-0.00463 (0.0214)	0.0300 (0.0199)
Savings - My Spouse Decides	0.0107 (0.0147)	0.0207 (0.0221)	-0.00173 (0.0220)	0.00815 (0.0192)
Savings - Decide Together	-0.0288 (0.0249)	-0.0146 (0.0377)	-0.0425 (0.0326)	-0.0341 (0.0324)
Savings - Decide Alone	-0.0549** (0.0264)	-0.0447 (0.0427)	-0.0646* (0.0343)	-0.0740** (0.0336)
Savings - Other	-0.0594 (0.0457)	-0.0804 (0.0747)	-0.0332 (0.0641)	-0.0403 (0.0584)
Self Reported Savings <sup>+</sup>	0.0136 (0.0144)	0.00440 (0.0347)	0.0562 (0.0796)	0.0211 (0.0188)
Savings in Family Bank Accounts <sup>+</sup>	0.716 (0.470)	0.138 (0.622)	1.60** (0.757)	1.25** (0.603)
N	1194	597	597	1194

Notes: <sup>+</sup>Coefficients and standard errors have been scaled by 100,000 for readability. Robust standard errors (clustered at the couple level for pooled regressions) reported in parentheses. All regressions include session fixed effects. \*\*\*, \*\*, and \* indicate significance at the 99, 95, and 90 percent levels respectively.

Table 1-3: Randomization Verification - Interest Rates and Extra Statements

	Individual Accounts		Joint Accounts	Extra Statements	Cash Prize
	Husband	Wife			
<i>Panel A. Adherence to Theoretical Probabilities</i>					
10% Interest Rate/Extra Stmt/Cash	0.233 {0.345}	0.240 {0.571}	0.320 {0.515}	0.456* {0.0544}	0.175** {0.0327}
6% Interest Rate	0.243 {0.706}	0.263 {0.478}	0.353 {0.298}		
2% Interest Rate	0.256 {0.741}	0.266 {0.369}	0.327 {0.761}		
0% Interest Rate	0.268 {0.321}	0.231 {0.299}			
<i>Panel B. Correlation with Demographic Characteristics</i>					
Age	170 {0.745}	156 {0.416}	61.5 {1.000}	24.7 {1.000}	37.3 {0.995}
Education	42.1 {0.468}	35.1 {0.766}	15.2 {0.977}	7.60 {0.909}	8.91 {0.837}
Literate	1.50 {0.682}	1.21 {0.752}	4.20 {0.122}	0.00986 {0.921}	0.276 {0.599}
Number Children	51.3 {0.577}	37.4 {0.543}	33.9 {0.567}	18.4 {0.303}	13.0 {0.792}
Occupation	19.6 {0.721}	36.3 {0.109}	21.0 {0.281}	20.8** {0.0136}	4.65 {0.864}
Income Last Week (Decile)	24.5 {0.605}	33.9* {0.0865}	26.1* {0.0986}	3.21 {0.956}	11.7 {0.228}
Has Mobile Phone	2.54 {0.468}	1.69 {0.640}	0.455 {0.797}	1.42 {0.233}	0.229 {0.632}
Participates in ROSCA	3.44 {0.329}	1.61 {0.657}	2.40 {0.301}	0.0954 {0.757}	0.0259 {0.872}
Has Bank Account	2.75 {0.432}	2.67 {0.445}	1.31 {0.520}	0.697 {0.404}	0.456 {0.500}
Saves at Home	0.667 {0.881}	5.51 {0.138}	2.23 {0.328}	0.0156 {0.900}	3.52* {0.0606}
Has SACCO Account	1.61 {0.657}	3.82 {0.282}	0.186 {0.911}	0.588 {0.443}	0.0492 {0.824}
Has MPESA Account	4.45 {0.217}	2.63 {0.452}	9.75*** {0.00763}	0.0101 {0.920}	0.438 {0.508}
Saves Other Ways	6.37* {0.0950}	0.381 {0.944}	4.18 {0.124}	4.42** {0.0355}	0.979 {0.323}
Self Reported Savings (Decile)	24.6 {0.596}	27.5 {0.439}	20.1 {0.330}	12.2 {0.203}	17.7** {0.0388}
Consumption Decision Making	17.9 {0.269}	23.4* {0.0752}	6.52 {0.770}	5.03 {0.284}	1.87 {0.867}
Savings Decision Making	13.5 {0.562}	15.9 {0.388}	5.12 {0.883}	3.22 {0.666}	2.11 {0.834}
Estimated Discount Factor (Decile)	21.1 {0.272}	15.5 {0.799}	23.1* {0.0590}	8.55 {0.286}	15.5** {0.0296}
Well Matched Couple	1.72 {0.633}	1.55 {0.671}	0.0798 {0.961}	1.42 {0.233}	0.290 {0.590}
Distance from Bank (Decile)	13.0 {0.966}	31.5 {0.139}	10.1 {0.862}	6.25 {0.619}	16.4** {0.0370}

Notes: P-values from binomial tests of theoretical probabilities (Panel A) or chi-squared tests for the equality of treatment distributions across demographic characteristics (Panel B) in braces. Theoretical probabilities are 1/4 for the first two columns, 1/3 for the third column, 1/2 for the fourth column, and 1/5 for the final column. The last three columns of Panel B test the equality of treatment distributions over both husband and wife's characteristics. These tests are adjusted for clustering at the couple level. \*\*\*, \*\*, and \* indicate significance at the 99, 95, and 90 percent levels respectively.

Table 1-4: Impact of Interest Rates on Account Takeup and Use

	Opened	Saved	Average Balance	Total Deposits	Total Withdrawals	Number Transactions	Fees
2% Interest	0.0524 (0.0354)	0.0182 (0.0197)	-21.0 (34.0)	-24.9 (293)	-82.0 (244)	0.0421 (0.135)	-0.360 (3.94)
6% Interest	0.144*** (0.0377)	0.0635*** (0.0218)	96.5 (61.5)	538 (451)	437 (373)	0.167 (0.136)	5.31 (4.79)
10% Interest	0.191*** (0.0399)	0.0851*** (0.0233)	114** (50.2)	794* (439)	644 (415)	0.319* (0.166)	8.91* (5.19)
Joint	0.263*** (0.0524)	0.138*** (0.0342)	166*** (65.3)	3058 (2470)	2871 (2461)	0.551** (0.244)	15.1* (8.04)
Joint×6% Interest	-0.0208 (0.0637)	0.0432 (0.0491)	-124 (103)	-3218 (2498)	-3066 (2474)	-0.297 (0.302)	-13.3 (10.1)
Joint×10% Interest	0.0282 (0.0612)	0.0802 (0.0516)	-10.9 (111)	-2041 (2484)	-2061 (2447)	0.310 (0.432)	-0.899 (12.7)
DV Mean (Individual Accounts, No Interest)	0.289	0.0537	50.0	303	212	0.245	5.18
N	1791	1791	1791	1791	1791	1791	1791

Notes: Robust standard errors clustered at the couple level in parentheses. All regressions include session fixed effects. \*\*\*, \*\*, and \* indicate significance at the 99, 95, and 90 percent levels respectively.

Table 1-5: Preference Heterogeneity and Savings Levels, By Account Type

	Individual Accounts			Joint Accounts		
	Saved	Average Balance	Fraction Balance	Saved	Average Balance	Fraction Balance
<i>Panel A. Basic Controls</i>						
Well Matched	-0.0675*** (0.0206)	-57.2 (37.0)	-0.0818*** (0.0310)	0.100** (0.0452)	208 (133)	0.173*** (0.0652)
Joint Deviation	-0.0523 (0.0330)	-69.1 (50.9)	-0.0722 (0.0492)	0.0708 (0.0652)	409** (197)	0.146 (0.103)
<i>Panel B. + Time Preference Controls</i>						
Well Matched	-0.0858*** (0.0223)	-90.0* (52.2)	-0.116*** (0.0325)	0.112** (0.0500)	97.3 (96.8)	0.242*** (0.0696)
Joint Deviation	-0.0468 (0.0323)	-65.3 (51.0)	-0.0778 (0.0491)	0.0736 (0.0661)	414** (193)	0.154 (0.104)
<i>Panel C. + Demographic Controls</i>						
Well Matched	-0.0885*** (0.0229)	-80.9 (54.6)	-0.122*** (0.0328)	0.117** (0.0506)	92.2 (97.6)	0.247*** (0.0730)
Joint Deviation	-0.0611* (0.0320)	-58.8 (53.0)	-0.0617 (0.0500)	0.0811 (0.0685)	392** (195)	0.123 (0.109)
<i>Panel D. + Economic Controls</i>						
Well Matched	-0.0870*** (0.0228)	-84.2 (56.2)	-0.119*** (0.0324)	0.109** (0.0518)	95.6 (103)	0.241*** (0.0740)
Joint Deviation	-0.0610* (0.0318)	-70.2 (52.3)	-0.0693 (0.0546)	0.0845 (0.0693)	392** (197)	0.137 (0.122)
DV Mean (Omitted)	0.114	126	0.200	0.271	174	0.601
N	1194	1194	512	597	597	256

Notes: Robust standard errors (clustered at the couple level for individual accounts) in parentheses. Time preference controls include separate dummies for upper/lower censoring and top/bottomcoding of the discount factors of each spouse and the average household discount factor. The demographic control set adds controls for spousal heterogeneity in age, education, number of children, and literacy. The economic control set adds controls for heterogeneity in income, an indicator for subsistence farmers or the unemployed, and phone ownership. All regressions include session fixed effects and fixed effects for each account's interest rate. \*\*\*, \*\*, and \* indicate significance at the 99, 95, and 90 percent levels respectively.

Table 1-6: Responses to the Excess Interest Rate by Match Quality

<i>Panel A. Well Matched Couples Only</i>				
Excess Low×Indiv.	0.00285 (0.00518)	0.00349 (0.00517)	0.00316 (0.00623)	0.00708 (0.00682)
Excess High×Indiv.	0.0210** (0.00939)	0.0225*** (0.00957)	0.0222*** (0.00897)	0.0182** (0.00846)
Excess Low×Joint	0.0346** (0.0152)	0.0348** (0.0154)	0.0365*** (0.0156)	0.0384*** (0.0163)
Excess High×Joint	-0.00564 (0.0162)	-0.00601 (0.0159)	-0.00674 (0.0155)	-0.00634 (0.0159)
DV Mean	0.152	0.152	0.152	0.152
N	597	597	597	597
<i>Panel B. Impacts by Match Quality</i>				
Excess Low×Indiv.	0.00192 (0.00498)	0.00110 (0.00508)	0.00293 (0.00534)	0.00345 (0.00552)
Excess High×Indiv.	0.0202** (0.00946)	0.0223*** (0.00935)	0.0212** (0.00943)	0.0199** (0.00909)
Excess Low×Joint	0.0347*** (0.0148)	0.0346*** (0.0148)	0.0363*** (0.0147)	0.0367*** (0.0148)
Excess High×Joint	-0.00625 (0.0164)	-0.00667 (0.0164)	-0.00666 (0.0161)	-0.00864 (0.0161)
Excess Low×Indiv.×Bad Match	-0.00648 (0.00755)	-0.00457 (0.00765)	-0.00672 (0.00790)	-0.00798 (0.00808)
Excess High×Indiv.×Bad Match	-0.0279*** (0.0116)	-0.0304*** (0.0115)	-0.0299*** (0.0116)	-0.0268*** (0.0114)
Excess Low×Joint×Bad Match	-0.0323* (0.0178)	-0.0313* (0.0178)	-0.0322* (0.0177)	-0.0331* (0.0178)
Excess High×Joint×Bad Match	-0.00356 (0.0191)	-0.00249 (0.0191)	-0.00219 (0.0189)	0.000438 (0.0189)
DV Mean	.161	.161	.161	.161
N	1791	1791	1791	1791
Control Set	Basic	+Time Pref	+Demo.	+Economic

Notes: Robust standard errors clustered at the couple level in parentheses. All regressions include session fixed effects and fixed effects that saturate interest rate×joint account (×badly matched for panel B). Time preference controls include separate dummies for upper/lower censoring and top/bottomcoding of the discount factors of each spouse and the average household discount factor. The demographic control set adds controls for spousal heterogeneity in age, education, number of children, and literacy. The economic control set adds controls for heterogeneity in income, an indicator for subsistence farmers or the unemployed, and phone ownership. \*\*\*, \*\*, and \* indicate significance at the 99, 95, and 90 percent levels respectively.

Table 1-7: Proxied Banking Costs and Savings Behavior

	Couple Level			Account Level	
	Saved Any	Saved Joint	Frac. Joint	Saved	
<i>Panel A. No Controls</i>					
Cost Index/High Cost	-0.377*** (0.127)	0.571*** (0.180)	0.758*** (0.164)	-0.125*** (0.0351)	-0.158* (0.0852)
Excess				0.0169*** (0.00579)	0.0178** (0.00785)
Excess×Cost Index/High Cost				-0.0141*** (0.00542)	-0.0157 (0.0100)
<i>Panel B. + Time Preference Controls</i>					
Cost Index/High Cost	-0.397*** (0.130)	0.588*** (0.187)	0.784*** (0.169)	-0.131*** (0.0359)	-0.143 (0.0901)
Excess				0.0174*** (0.00583)	0.0176** (0.00777)
Excess×Cost Index/High Cost				-0.0138*** (0.00547)	-0.0141 (0.0103)
<i>Panel C. + Demographic Controls</i>					
Cost Index/High Cost	-0.363*** (0.145)	0.539*** (0.226)	0.766*** (0.205)	-0.113*** (0.0358)	-0.159* (0.0928)
Excess				0.0162*** (0.00589)	0.0174*** (0.00728)
Excess×Cost Index/High Cost				-0.0135*** (0.00559)	-0.0158 (0.0100)
<i>Panel D. + Economic Controls</i>					
Cost Index/High Cost	-0.426*** (0.178)	0.222 (0.257)	0.523** (0.248)	-0.0966** (0.0459)	-0.152* (0.0913)
Excess				0.0163*** (0.00602)	0.0177*** (0.00746)
Excess×Cost Index/High Cost				-0.0136*** (0.00567)	-0.0163 (0.0101)
DV Mean	0.429	0.676	0.652	0.152	
N	597	256	256	597	
Sample	All	Savers	Savers	Matched	Matched
Interest Rate Restrictions?	N/A	N/A	N/A	Yes	No

Notes: Robust standard errors (clustered at the couple level when relevant) in parentheses. Time preference controls include separate dummies for upper/lower censoring and top/bottom coding of the discount factor for each member of the couple and the average household discount factor. The demographic control set adds controls for spousal heterogeneity in age, education, number of children, and literacy. The economic control set adds controls for heterogeneity in income and phone ownership. The final two columns include dummy variables that fully saturate interest rate×joint, as well as a high cost×joint dummy. The fourth column also includes interactions between the high cost dummy and the interest rate×joint interaction set. \*\*\*, \*\*, and \* indicate significance at the 99, 95, and 90 percent levels respectively.

Table 1-8: Interest Rate Losses Borne Among Savers, by Match Quality

Assumed $E_i$	0	5	10	15	20
<i>Poorly Matched Couples</i>					
Maximum Interest Earnings	8.30	6.90	6.25	6.09	6.08
Actual Interest Earnings	6.00	4.61	3.23	1.84	0.452
Loss	2.30	2.29	3.03	4.25	5.62
<i>Well Matched Couples</i>					
Maximum Interest Earnings	8.29	7.01	6.62	6.47	6.45
Actual Interest Earnings	6.89	6.14	5.39	4.64	3.89
Loss	1.40	0.874	1.23	1.83	2.55
<i>Loss Gap</i>	0.900***	1.42***	1.80***	2.42***	3.07***
	(0.344)	(0.364)	(0.494)	(0.660)	(0.843)
As Share of Well Matched Loss	0.643	1.62	1.47	1.33	1.20
N	256	256	256	256	256

Notes: Robust standard errors in parentheses. \*\*\*, \*\*, and \* indicate significance at the 99, 95, and 90 percent levels respectively. Principal components index constructed from distance from the bank, spouse specific indicators for subsistence farmers/the unemployed, and spouse specific indicators for baseline bank account ownership and SACCO membership.

Table 1-9: Impact of Extra Statements on Savings and Average Balances of Individual Accounts

	Protocol Check		First Stage		Reduced Forms				Two Stage Least Squares			
	Opened		Consented to ES		Saved		Average Balance		Saved		Average Balance	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
<i>Panel A. Basic Controls</i>												
Extra Statement	-0.00985 (0.0523)	-0.0641 (0.109)	0.598*** (0.0694)	0.474*** (0.118)	-0.00683 (0.0629)	-0.356*** (0.125)	27.3 (123)	-417 (296)	-0.0114 (0.106)	-0.832*** (0.356)	45.6 (204)	-958 (737)
x Well Matched		0.0118 (0.107)		0.00375 (0.133)		0.264** (0.118)		422** (214)		0.607*** (0.247)		897** (417)
x Well Informed		0.0967 (0.122)		0.257** (0.127)		0.347** (0.156)		512 (517)		0.755** (0.346)		995 (919)
<i>Panel B. + Time Preference Controls</i>												
Extra Statement	-0.0112 (0.0542)	-0.0617 (0.110)	0.575*** (0.0753)	0.472*** (0.130)	0.00513 (0.0573)	-0.291*** (0.121)	91.3 (141)	-316 (345)	0.00894 (0.0996)	-0.771** (0.356)	159 (244)	-881 (873)
x Well Matched		0.0267 (0.110)		-0.0238 (0.147)		0.133 (0.124)		207 (240)		0.452* (0.259)		660 (524)
x Well Informed		0.0752 (0.125)		0.237 (0.146)		0.350** (0.157)		620 (542)		0.743** (0.338)		1138 (969)
<i>Panel C. + Demographic Controls</i>												
Extra Statement	-0.00810 (0.0541)	-0.0114 (0.118)	0.601*** (0.0644)	0.517*** (0.139)	0.0425 (0.0549)	-0.265* (0.149)	179 (146)	-429 (318)	0.0707 (0.0911)	-0.625* (0.334)	298 (244)	-1161 (806)
x Well Matched		0.0218 (0.113)		0.0864 (0.141)		0.180 (0.135)		690*** (279)		0.453* (0.262)		1483** (688)
x Well Informed		0.000729 (0.137)		0.185 (0.148)		0.307* (0.182)		952* (517)		0.548* (0.299)		1602* (894)
<i>Panel D. + Economic Controls</i>												
Extra Statement	-0.00784 (0.0563)	-0.0205 (0.117)	0.610*** (0.0632)	0.515*** (0.141)	0.0101 (0.0610)	-0.433*** (0.158)	192 (215)	-353 (352)	0.0166 (0.0997)	-0.971*** (0.370)	315 (350)	-1135 (906)
x Well Matched		0.0231 (0.115)		0.0285 (0.155)		0.173 (0.151)		513* (279)		0.517* (0.292)		1117* (647)
x Well Informed		0.00172 (0.139)		0.230 (0.175)		0.420** (0.184)		867 (534)		0.799*** (0.329)		1617 (1009)
DV Mean (ES=0)		0.395		0		0.266		241		0.266		241
N		634		244		244		244		244		244

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

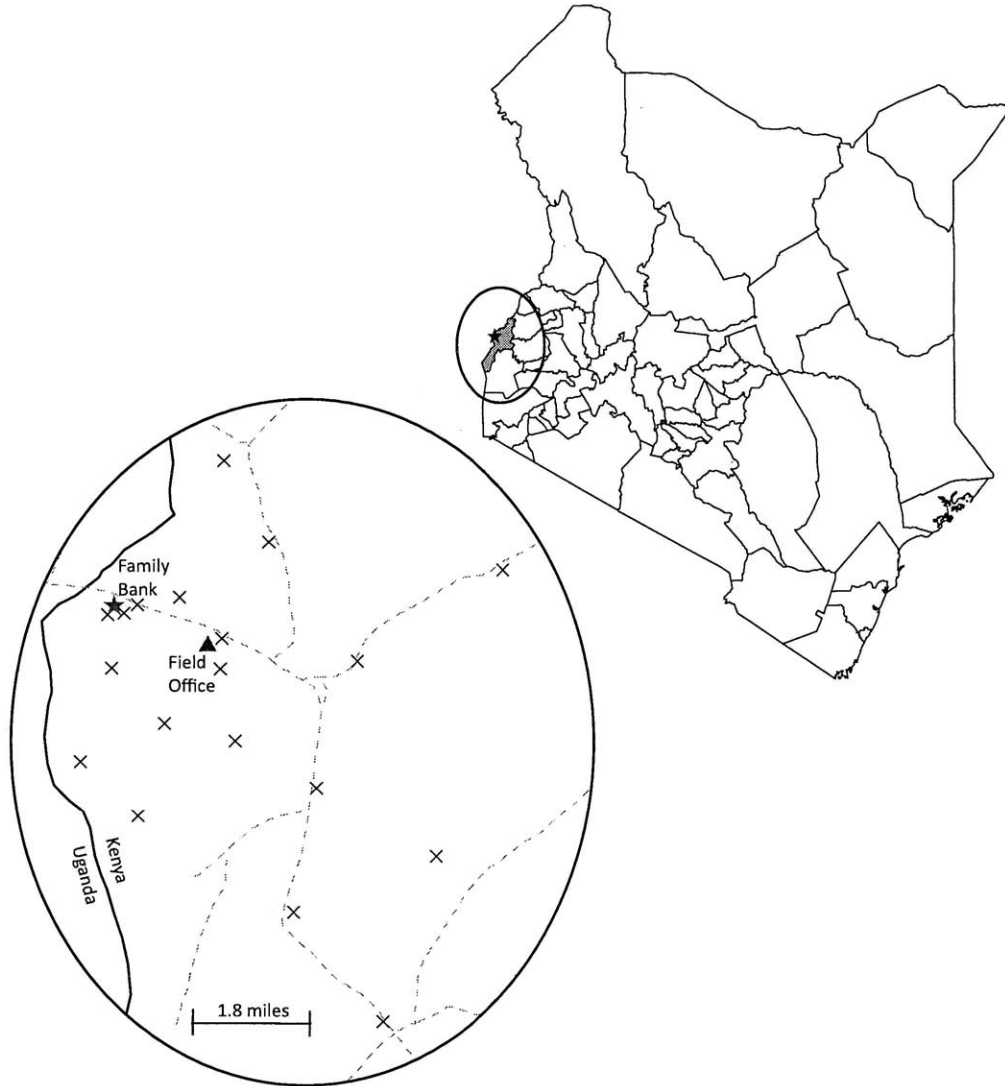


Table 1-10: Preference Heterogeneity, Information Sharing, and Savings Levels, By Account Type

	Individual Accounts			Joint Accounts		
	Saved	Average Balance	Fraction Balance	Saved	Average Balance	Fraction Balance
<i>Panel A. Basic Controls</i>						
Well Matched	-0.0689*** (0.0205)	-58.9 (37.1)	-0.0860*** (0.0299)	0.103** (0.0451)	204 (132)	0.179*** (0.0634)
Joint Deviation	-0.0478 (0.0326)	-63.2 (50.6)	-0.0506 (0.0490)	0.0616 (0.0654)	415** (201)	0.102 (0.103)
Well Informed	-0.0454** (0.0231)	-66.4 (44.6)	-0.107*** (0.0318)	0.0918** (0.0407)	-15.7 (100)	0.205*** (0.0678)
<i>Panel B. + Time Preference Controls</i>						
Well Matched	-0.0857*** (0.0222)	-89.8* (52.4)	-0.114*** (0.0323)	0.112** (0.0498)	97.2 (97.0)	0.237*** (0.0695)
Joint Deviation	-0.0458 (0.0321)	-63.6 (51.1)	-0.0812* (0.0492)	0.0718 (0.0662)	413** (194)	0.161 (0.104)
Well Informed	-0.0345* (0.0207)	-56.1 (41.9)	-0.0666** (0.0299)	0.0613 (0.0386)	24.9 (79.0)	0.125* (0.0650)
<i>Panel C. + Demographic Controls</i>						
Well Matched	-0.0886*** (0.0228)	-81.1 (54.7)	-0.120*** (0.0327)	0.118*** (0.0502)	92.5 (97.6)	0.243*** (0.0730)
Joint Deviation	-0.0597* (0.0319)	-56.7 (53.7)	-0.0669 (0.0499)	0.0778 (0.0688)	388** (195)	0.133 (0.109)
Well Informed	-0.0270 (0.0209)	-40.6 (50.3)	-0.0499 (0.0320)	0.0640 (0.0410)	65.0 (68.0)	0.0924 (0.0715)
<i>Panel D. + Economic Controls</i>						
Well Matched	-0.0872*** (0.0228)	-84.2 (55.7)	-0.118*** (0.0326)	0.110** (0.0516)	97.1 (103)	0.239*** (0.0747)
Joint Deviation	-0.0600* (0.0318)	-70.2 (53.4)	-0.0701 (0.0545)	0.0815 (0.0695)	387** (196)	0.137 (0.122)
Well Informed	-0.0161 (0.0219)	0.191 (57.1)	-0.0235 (0.0368)	0.0544 (0.0423)	89.5 (70.3)	0.0289 (0.0853)
DV Mean (Omitted-Match)	0.114	126	0.200	0.271	174	0.601
N	1194	1194	512	597	597	256

Notes: Robust standard errors (clustered at the couple level for individual accounts) in parentheses. "Omitted-Match" refers to couples who are poorly matched and have no joint deviation. Time preference controls include separate dummies for upper/lower censoring and top/bottomcoding of the discount factors of each spouse and the average household discount factor. The demographic control set adds controls for spousal heterogeneity in age, education, number of children, and literacy. The economic control set adds controls for heterogeneity in income, an indicator for subsistence farmers or the unemployed, and phone ownership. All regressions include session fixed effects and fixed effects for each account's interest rate. \*\*\*, \*\*, and \* indicate significance at the 99, 95, and 90 percent levels respectively.

Figure 1-1: Map of Study Locations



Notes: The shaded area on the map of Kenya indicates Busia district, with Busia town marked by a star. The local map shows detail surrounding Busia town, including major roads (dashed lines), boundaries (solid lines) and study locations. Group meeting locations are demarcated by an x.

Figure 1-2: Interest Rate Design

<b>R<sub>J</sub>=2</b>					<b>R<sub>J</sub>=6</b>					<b>R<sub>J</sub>=10</b>				
	R <sub>M</sub> =0	R <sub>M</sub> =2	R <sub>M</sub> =6	R <sub>M</sub> =10		R <sub>M</sub> =0	R <sub>M</sub> =2	R <sub>M</sub> =6	R <sub>M</sub> =10		R <sub>M</sub> =0	R <sub>M</sub> =2	R <sub>M</sub> =6	R <sub>M</sub> =10
R <sub>F</sub> =0	2, -2,-2	0, 0,-2	-4, 4,-6	-8, 8,-10	R <sub>F</sub> =0	6, -6,-6	4, -4,-6	0, 0,-6	-4, 4,-10	R <sub>F</sub> =0	10,-10, -10	8, -8,-10	4, -4,-10	0, 0,-10
R <sub>F</sub> =2	0, -2,0	0, 0,0	-4, 4,-4	-8, 8,-8	R <sub>F</sub> =2	4, -6,-4	4, -4,-4	0, 0,-4	-4, 4,-8	R <sub>F</sub> =2	8, -10,-8	8, -8,-8	4, -4,-8	0, 0,-8
R <sub>F</sub> =6	-4, -6,4	-4, -4,4	-4, 0,0	-8, 4,-4	R <sub>F</sub> =6	0, -6,6	0, -4,6	0, 0,0	-4, 4,-4	R <sub>F</sub> =6	4, -10,-4	4, -8,-4	4, -4,-4	0, 0,-4
R <sub>F</sub> =10	-8, -10,8	-8, -8,8	-8, -4,4	-8, 0,0	R <sub>F</sub> =10	-4, -10,4	-4, -8,4	-4, -4,4	-4, 0,0	R <sub>F</sub> =10	0, -10,0	0, -8,0	0, -4,0	0, 0,0

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

Figure 1-3: Theoretical Responses to the Excess Interest Rate by Match Quality and Account

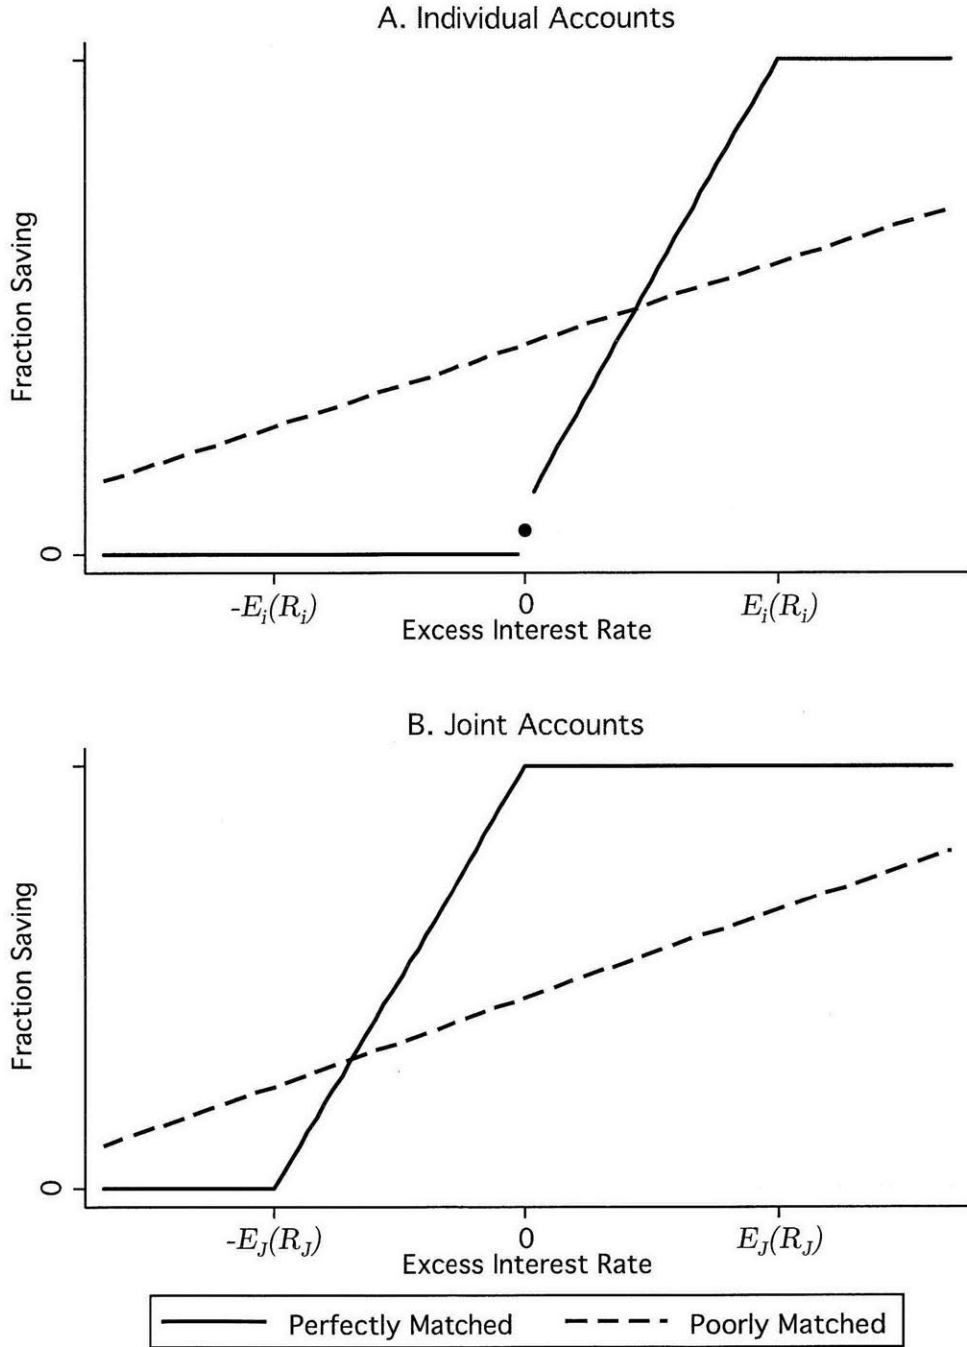
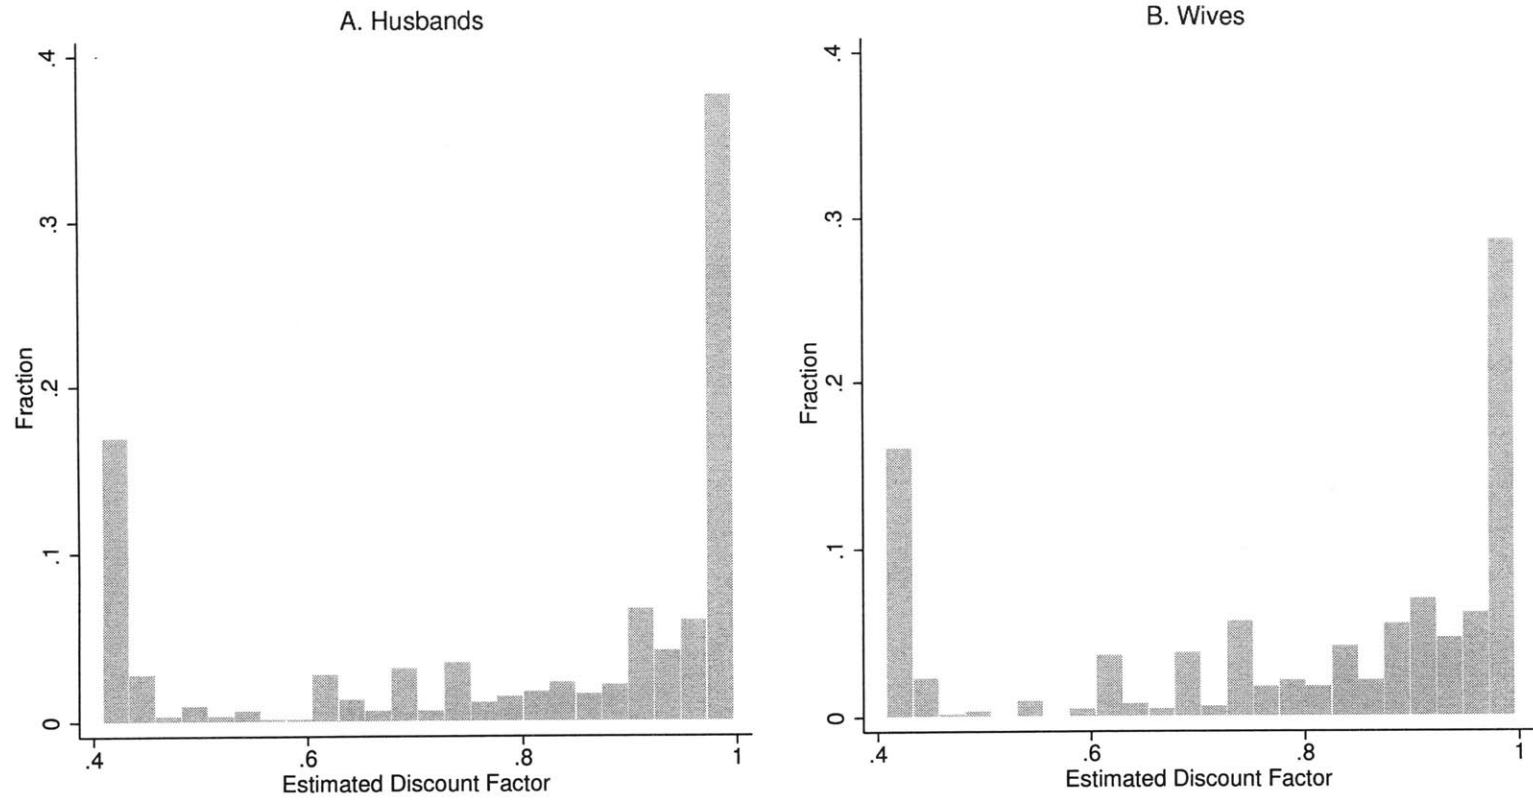


Figure 1-4: Distribution of Estimated Discount Factors by Gender


73



---

Figure 1-5: Intra-household Heterogeneity in Estimated Discount Factors

---

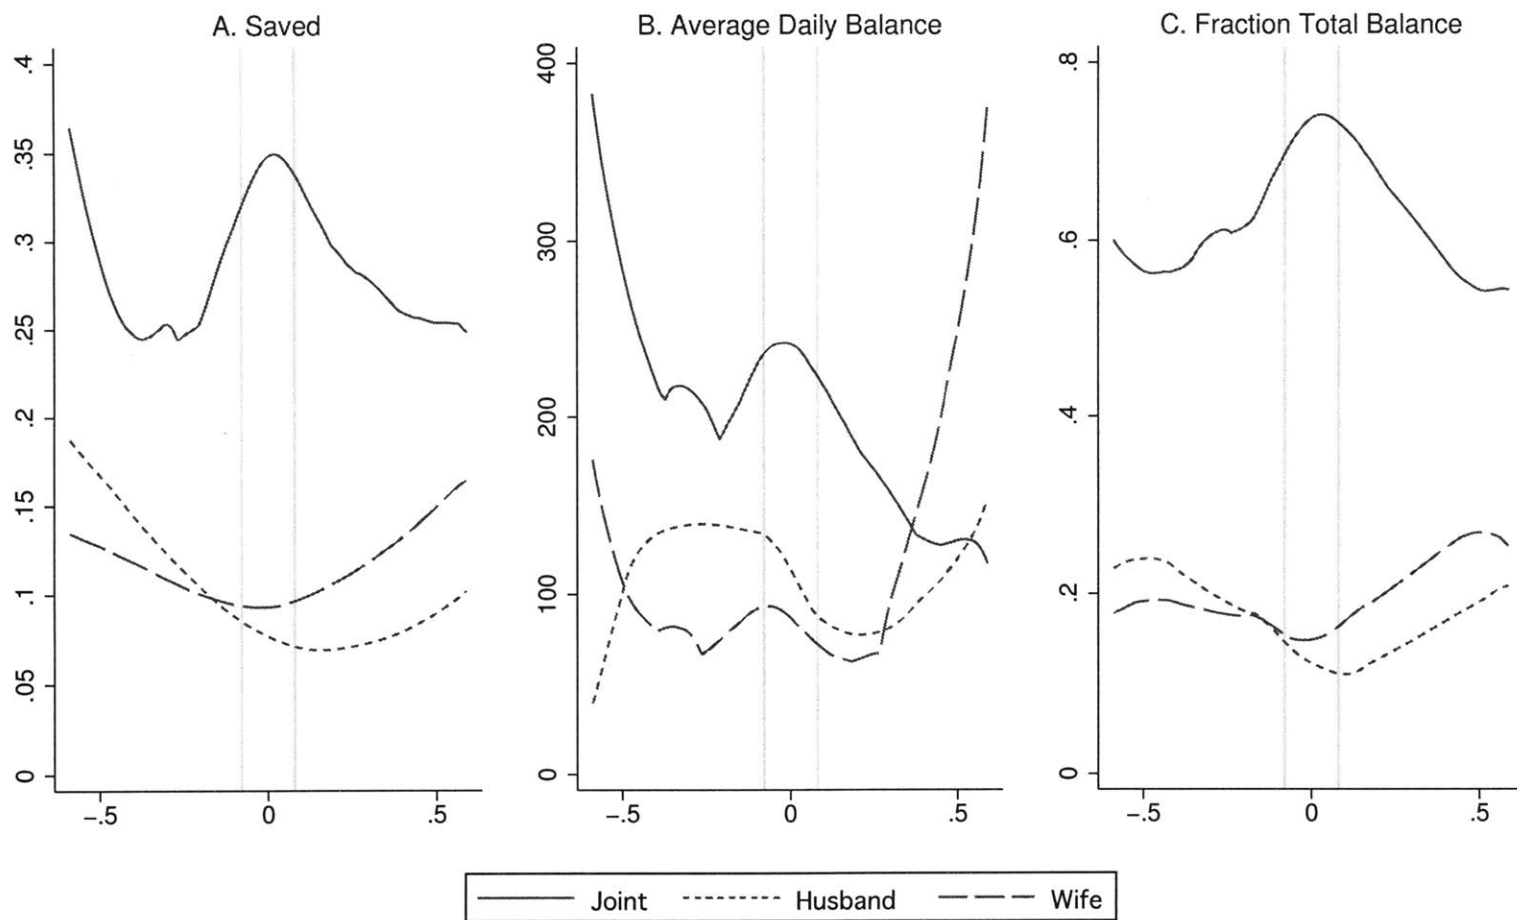


---

Note: Well matched couples delineated by dashed lines.

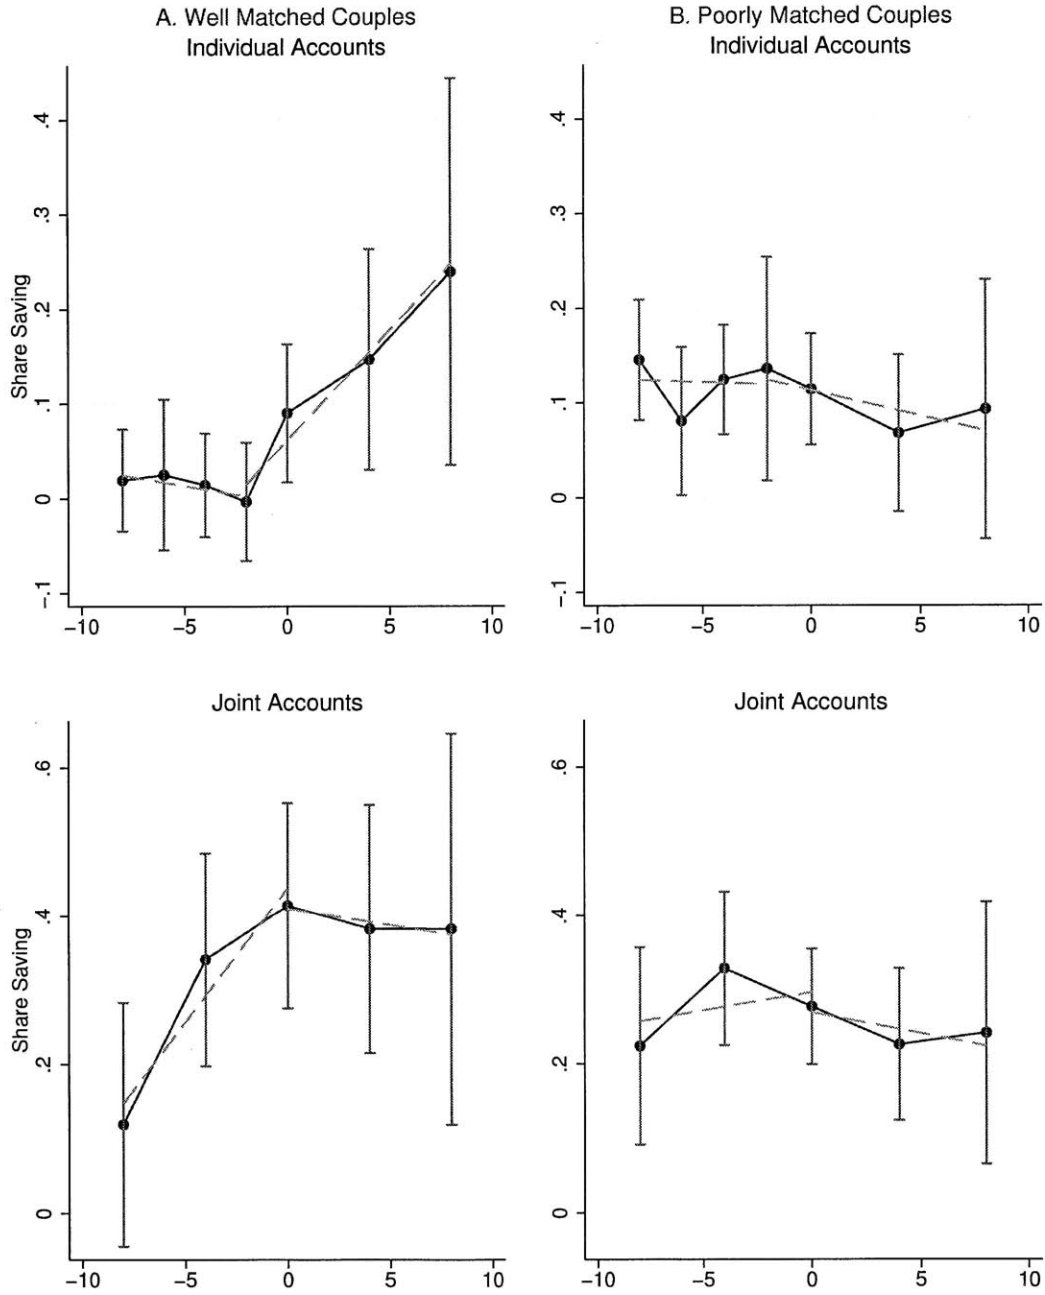
Figure 1-6: Relationship Between Savings Balances and Heterogeneity in Estimated Discount Factors

75



Note: Well matched couples delineated by gray vertical lines.

Figure 1-7: Savings Response to Excess Interest Rate by Match Quality and Account Type



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



Figure 1-8: Distribution of Household Information Index


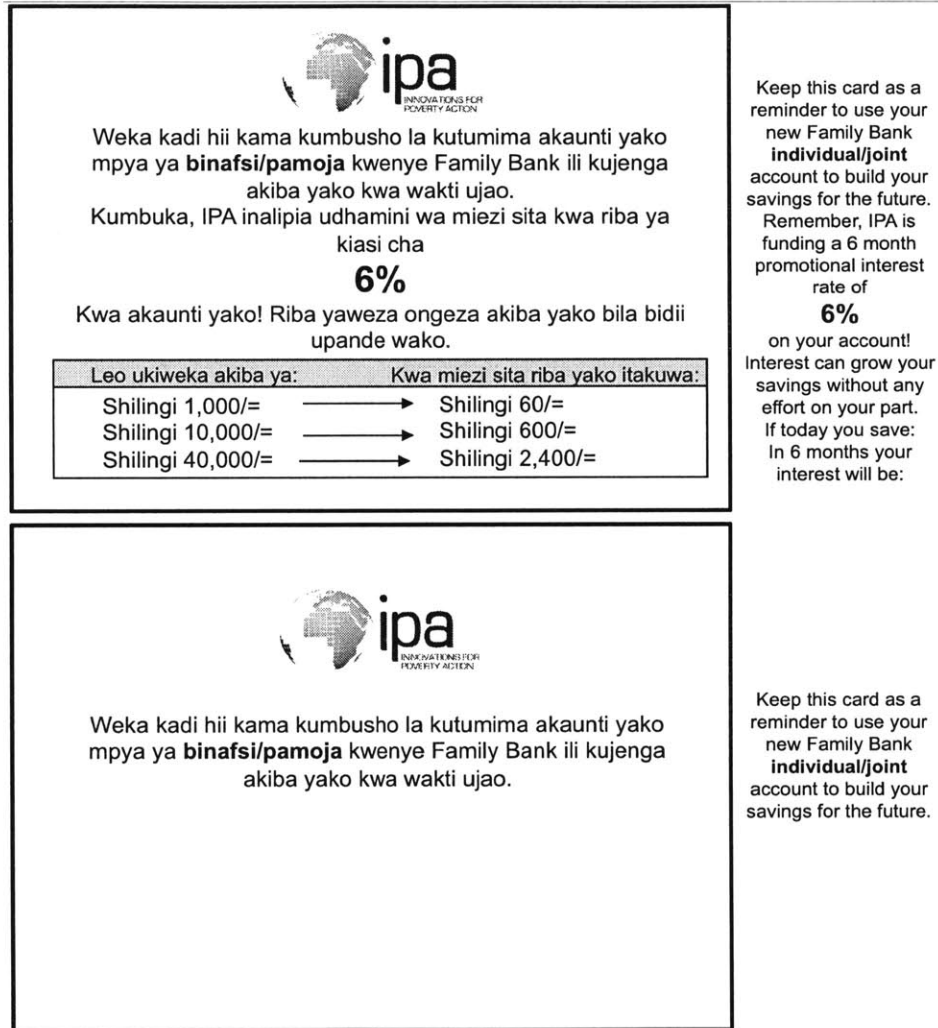


Table 1-A1: Demographic Characteristics of Study Sample by Match Quality

	Well Matched	Badly Matched	Difference	N
Age	38.1 [13.2]	38.7 [12.8]	-0.595 (0.801)	1194
Education	6.77 [4.01]	7.09 [3.91]	-0.319 (0.245)	1189
Literate	0.731 [0.444]	0.784 [0.412]	-0.0528** (0.0266)	1194
Number Children	4.60 [2.86]	4.70 [2.69]	-0.0980 (0.172)	1194
Subsistence Farmer or No Job	0.429 [0.496]	0.440 [0.497]	-0.0113 (0.0305)	1189
Income Last Week	895 [1361]	1164 [2612]	-269** (117)	1162
Owns Mobile Phone	0.496 [0.501]	0.437 [0.496]	0.0592* (0.0307)	1189
Participates in ROSCA	0.565 [0.496]	0.578 [0.494]	-0.0126 (0.0304)	1194
Has Bank Account	0.188 [0.392]	0.222 [0.416]	-0.0339 (0.0245)	1194
Savings in Bank Account	8626 [12046]	8580 [17870]	46.4 (2121)	208
Saves at Home	0.917 [0.276]	0.853 [0.354]	0.0639*** (0.0187)	1193
Savings at Home	887 [2130]	1103 [3034]	-216 (163)	1023
Consumption - I Decide	0.322 [0.468]	0.319 [0.466]	0.00329 (0.0288)	1187
Consumption - Spouse Decides	0.234 [0.424]	0.223 [0.417]	0.0103 (0.0260)	1187
Consumption - Decide Together	0.378 [0.486]	0.385 [0.487]	-0.00644 (0.0299)	1187
Consumption - Decide Alone	0.0203 [0.141]	0.0265 [0.161]	-0.00618 (0.00912)	1187
Savings - I Decide	0.482 [0.500]	0.393 [0.489]	0.0892*** (0.0306)	1187
Savings - Spouse Decides	0.338 [0.474]	0.372 [0.484]	-0.0333 (0.0294)	1187
Savings - Decide Together	0.0808 [0.273]	0.116 [0.321]	-0.0355** (0.0178)	1187
Savings - Decide Alone	0.0783 [0.269]	0.0936 [0.291]	-0.0153 (0.0170)	1187
Weekly Discount Factor	0.876 [0.197]	0.758 [0.222]	0.119*** (0.0126)	1194
Distance from Family Bank (Miles)	3.93 [2.25]	3.71 [2.19]	0.224 (0.137)	1194

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

Figure 1-A1: Interest Rate Reminder Cards and Translation



Note: Figure illustrates reminder cards given for 6 percent and 0 percent interest.



## Chapter 2

# Cost and Convenience: The Impact of ATM Card Provision on Formal Savings Account Use in Kenya

### 2.1 Introduction

The vast majority of the world's poor do not have access to formal financial services. Recent estimates suggest that nearly three quarters of individuals in developing countries are unbanked (Kendall et al. 2010), and in Sub-Saharan Africa, this estimate reaches 80 percent (Chaia et al. 2009). This lack of access does not reflect an inability or unwillingness to save. Indeed, the ability to save is essential for the poor – stores of resources are needed to smooth unexpected shocks, to boost investment in income generating activities, and to purchase valuable durable goods. Accordingly, Collins et al. (2009) document that low-income households in developing countries save resources in a wide variety of informal and semi-formal savings devices, even though these devices are often quite costly. Policymakers have taken note of the costs and risks associated with informal savings, as well as evidence from bank expansions in developing countries suggesting that increasing access to the formal financial sector increases savings, investment, and income.<sup>1</sup> As a result, there is a growing movement to connect the unbanked to formal financial services. Yet at the micro-level, relatively little is known about the poor's use of formal savings accounts. The first part of this paper contributes to filling this knowledge gap by presenting the results of a field experiment designed to answer the following questions: Would the poor save with formal institutions if given the opportunity? Would making formal sector accounts cheaper (by reducing fees) and more convenient (by reducing non-fee transaction costs) substantially increase use of these accounts?

The answer to the above is not obvious. Traditional economic theory suggests that making formal

---

<sup>1</sup>See, for example, Aportela (1999), Bruhn and Love (2009), Burgess and Pande (2005), and Kaboski and Townsend (2005).

accounts cheap and readily available would result in increased adoption by low income households. However, a growing body of literature documents that individuals, particularly in developing countries, face important internal and external constraints to building savings balances. Furthermore, these constraints are often such that fees or restrictions on access to liquidity may actually help *increase* stores of savings. First, individuals may have to contend with time inconsistent preferences – if the temptation to spend out of readily accessible savings is too great, individuals may prefer to store resources in an account that is costly (in terms of time or money) to access (Banerjee and Mullainathan 2010; Laibson 1997). Second, individuals in developing countries face frequent demands on their resources from the community and extended family members. In such a context, individuals may prefer to store savings in technologies that are costly to access in order to protect their resources from these demands (Baland et al. 2007; Jakiela and Ozier 2011). Finally, demands on savings may come from within the household. When the household is not Pareto efficient, individuals may value savings technologies that make resources difficult to access or observe (Anderson and Baland 2002; Schaner 2011). For example, a wife may have difficulty denying her husband a few shillings for a drink if he knows she has a store of savings under the mattress. However, if the funds are stored at the bank and withdrawals incur a fee, she may be able to avoid having to make such a transfer.

Moreover, these concerns are economically relevant. Many informal savings devices are characterized by commitment and/or security features, and a growing body of empirical evidence suggests that individuals in developing countries behave in ways consistent with these concerns (for a review, see Karlan and Morduch (2010)). In light of these results, the second part of this paper asks the following question: Is there evidence that the above-described internal and external constraints limit the benefits of reduced fees and transaction costs? This question is of particular importance given a recent shift in policymaker interest away from microcredit towards microsavings. For example, in November 2010 the Bill and Melinda Gates Foundation announced a \$500 million pledge to expand access to formal savings accounts to the world's poor, with an emphasis on transactions cost reducing technologies such as mobile money (Bill and Melinda Gates Foundation 2010).

The field experiment that we designed to answer these questions was conducted in rural Kenya from June 2009-February 2010. Seven hundred forty-eight married couples were given the opportunity to open up to three accounts with a formal bank: a joint account, an individual account for the husband, and an individual account for the wife. Each account was randomly assigned a temporary 6-month interest rate, which ranged from zero percent to 10 percent. Altogether, these couples opened 1,121 accounts. As is typical in Kenya, the bank accounts featured nontrivial withdrawal fees of \$0.78, and accounts were only accessible during bank hours. However, the bank also offered ATM cards for the accounts: these cards reduced withdrawal fees by over 50 percent (to \$0.38) and also enabled card holders to make withdrawals outside of bank hours. These cards were costly to obtain (ordinarily, account holders would have to pay \$3.75 to acquire an ATM card). We randomly selected a subset of opened accounts to receive an ATM card for free – given their high cost, the

intervention increased ATM card takeup by nearly 89 percentage points.

Overall, we find relatively low use of formal accounts. Even though all couples included in the study reported that they were interested in opening a savings account with the bank, just 27 percent of the couples had saved in at least one of their new accounts after six months. ATM card provision had a significant impact on account use – overall, aggregate account use increased by 0.15 standard deviations relative to the control, savings rates increased by 28 percent (7 percentage points), and average daily account balances increased by 9 percent. However, this net positive impact masks a striking heterogeneous treatment effect by account type. In particular, ATM card provision to joint and men’s accounts substantially increased usage (aggregate account use, savings rates, and average daily balances increased by 0.24 standard deviations, 39 percent, and 16 percent respectively). This overall impact is equivalent to making 8 additional percentage points of interest available to couples. In contrast, providing ATM cards to women’s accounts actually had a negative impact on overall account use.

We find evidence that this heterogeneous treatment effect is driven by the intrahousehold resource allocation concerns described above. An ATM card makes an account cheaper to access and less secure (if a husband knows his wife’s passcode and can obtain the ATM card, he can access her account without her consent). We proxy the relative bargaining power of husbands and wives by using demographic characteristics collected during our baseline survey and find that women with below-median bargaining power had a large and significantly negative response to the ATM card treatment. In contrast, women with above-median power exhibited a small positive response that is not statistically significant from zero. Furthermore, men with below-median bargaining power did not respond to the ATM treatment, while men with above-median bargaining power exhibited a very large, significantly positive response to the ATM card treatment. These results suggest that both transaction costs and security are important determinants of formal account adoption and use. Furthermore, transaction cost saving technologies that also make account balances easier to view and access may differentially favor individuals who have better bargaining positions within the household.

The remainder of the paper is structured as follows: Section 2.2 outlines a simple model of individual savings behavior that highlights the role of transaction costs and security in the decision to make use of formal bank accounts. Section 2.3 describes the experimental design and the data, Section 3.5 presents the results, and Section 3.7 concludes.

## 2.2 Theoretical Framework

This section sets up a model of savings behavior to motivate the empirical analysis. The primary goal of the model is to highlight two mechanisms by which ATM cards might impact savings behavior: First, the cards reduce transaction costs associated with formal accounts by reducing withdrawal fees and enabling account owners to make withdrawals outside of usual bank hours. Second, ATM cards

may make formal accounts less secure, for reasons discussed in the introduction. These concerns may be particularly salient in developing countries, where a growing body of evidence suggests that individuals strategically use financial services to sequester resources from other members of the household and the community (Anderson and Baland 2002; Ashraf 2009; Ashraf et al. 2010; Baland et al. 2007; Brune et al. 2010; Jakiela and Ozier 2011; Platteau 2000; Schaner 2011).

Individuals in the model exponentially discount utility, where the per-period utility function is given by  $u(c_t)$ . We assume that  $u(\cdot)$  is well behaved in that it is concave, twice continuously differentiable, and  $u'(c) \rightarrow \infty$  as  $c \rightarrow 0$ . There are  $T$  periods in the world, so at time  $t$ , individual utility is expressed as  $U_t = \sum_{\tau=t}^T \delta^{(\tau-t)} u(c_\tau)$ . For simplicity, we assume that there is no uncertainty, and that in each period  $t$  individuals receive an endowment,  $y_t$ . In order to capture intrahousehold bargaining, we can think of  $y_t$  as the share of resources allocated to a given individual after transfers to and from others and the bargaining process. This imposes an important restriction – as long as the vector of endowments is independent of consumption and savings decisions, then this implies that transfers *from* others do not respond to individual savings choices. We will, however, allow that transfers *to* others out of savings will depend on how individuals save.

Individuals cannot borrow, but they can save. Specifically, there are two different savings technologies available at any time. First, agents may store resources at home (denoted by  $h_t$ ). Saving at home has the advantage of having no transaction costs, but it also makes cash more easily appropriated by other members of the household and the community – we capture this by denoting the return (net of transfers) on home savings as  $R_h$ , where  $R_h < 1$ . Alternatively, individuals may save at the bank (denoted by  $b_t$ ). The key advantage of saving with the bank is that fewer resources are appropriated by outside agents, so the rate of return on bank savings exceeds the rate of return on home savings:  $R_b > R_s$ . However, bank accounts also have transaction costs – in particular, an individual must pay a fee  $w > 0$  every time he or she makes a withdrawal.

Then the the individual's optimal consumption and savings allocation is given by the solution to the following constrained maximization problem:

$$\begin{aligned} & \arg \max_{\{c_t, b_t, h_t\}_{t=0}^T} \sum_{t=0}^T \delta^t u(c_t) \\ & \text{subject to} \\ & c_t + h_t + s_t \leq y_t + R_h h_{t-1} + R_b b_{t-1} + 1(b_t < R_b b_{t-1}) w \quad \forall t \\ & b_t \geq 0, h_t \geq 0 \quad \forall t \end{aligned}$$

where  $1(\cdot)$  is the indicator function. As a result of the lumpy withdrawal fees, this problem is not convex – some individuals will save at home rather than with the bank to avoid withdrawal fees. Individuals will be particularly averse to saving formally when desired savings levels are small (this is due to the fact that  $w$  is fixed in absolute terms) and when  $R_h$  and  $R_b$  are not very different.

The tradeoff between a higher rate of return and lumpy fees versus a lower rate of return and



no fee is the same tradeoff highlighted by the canonical work of Baumol (1952) and Tobin (1956). Consider a population of agents with different income streams. When the withdrawal fee is reduced, the number of deposits and withdrawals into formal bank accounts will increase. Individuals who were already using bank accounts will make more deposits and withdrawals, and other individuals who were not using bank accounts will start to use them. However, the impact on balances in bank accounts is ambiguous, as there are different effects on the intensive and extensive margins. This is easily seen by studying the  $T = 2$  case. Consider an agent who was already saving at the bank. Since  $R_b > R_h$  she will not save at home, and period 1 and 2 consumption will be governed by

$$u'(y_1 - b_1) = R_b \delta u'(y_2 + R_b b_1 - w)$$

differentiating implicitly, we see that  $\frac{\partial b_1}{\partial w} = \frac{R_b^2 \delta u''(c_2) + u''(c_1)}{R_b \delta u''(c_2)} > 0$  (so decreasing  $w$  will decrease the amount deposited). The intuition is straightforward – to an inframarginal saver, reducing  $w$  is akin to increasing income in future periods. The agent spreads this increase over both periods when allocating consumption, so  $b_1$  must decrease for first period consumption to increase. However, the reduction in the balance must necessarily be small, since consumption in the second period must also go up, which implies that  $\Delta b_1 < \frac{1}{R_b} \Delta w$ . Therefore, particularly in a multiperiod setting, it seems likely that this effect will be outweighed by an extensive margin effect – pre-existing savers will deposit and withdraw from their accounts more frequently (increasing the average daily balance) and individuals who did not save at all given the higher withdrawal fee will begin to save. As such, we expect a decrease in withdrawal fees to increase formal account deposits, withdrawals, and balances. However, the impact of the withdrawal fee on total fees paid is ambiguous. Although the fee goes down, more transactions will occur (among both pre-existing savers and new savers who are brought into the formal sector by the fee reduction) – as such, the net effect will depend on the elasticity of the number of withdrawals with respect to the withdrawal fee. If the elasticity is less than  $-1$ , then total fee revenue will increase.

However, as discussed earlier, ATM cards may also make bank accounts less secure. We capture this by assuming that an ATM card reduces  $R_b$  (with an ATM card saved resources are more easily appropriated by others, so the return on savings net of transfers goes down). Holding  $w$  constant, reducing  $R_b$  will reduce the number of deposits and withdrawals. As before, the net impact on the average balance is ambiguous – if the income effect is very large, it could theoretically outweigh the substitution effect *and* the extensive margin effect (though reducing security to the point that  $R_b = R_h$  will unambiguously reduce formal account use to zero). In practice, it seems likely that average balances would decline, all else equal.

Overall, the net effect of ATM card provision is an empirical question – if the impact of the fee reduction dominates, account use will increase, but if the security effect dominates, account use will decrease. This analysis also suggests the possibility of heterogeneous treatment effects: the impact of ATM cards is more likely to be negative when security effects are large. Our field

experiment enables us to empirically estimate the impact of ATM card provision on bank account use. Furthermore, we will exploit features of the experimental design and baseline data to study the empirical relevance of the security effect and related heterogeneous treatment effects. The next section describes the experimental context and design, as well as our data.

## 2.3 Experimental Design and Data

### 2.3.1 Experimental Context

The experiment was conducted in Western Province, Kenya, in areas surrounding the town of Busia. Busia is a commercial trading center straddling the Kenya-Uganda border. The town is well served by the formal banking sector, hosting over six banks at the time of field activities. It is only recently, however, that formal banks have begun to offer products suitable for low income individuals. Traditionally, Kenyan bank accounts required opening balances upwards of Ksh 1,000 (approximately equal to \$12.50 at an exchange rate of Ksh 80 per \$1, or \$19.23 using a PPP exchange rate of Ksh 52 per \$1) and charged monthly maintenance fees around Ksh 50 (\$0.63).<sup>2</sup> However, recently banks have begun to target lower income individuals, and several banks currently offer lower fee alternatives to traditional bank accounts.

The financial partner for this study is Family Bank of Kenya. At the time of field activities the bank had over 600,000 customers, 50 branches throughout the country, Ksh 13 billion (\$167 million) in assets, and actively targeted low and middle income individuals as part of its corporate strategy. All study participants were offered Family Bank's *Mwananchi* accounts. This account can be opened with any amount of money, though a minimum operating balance of Ksh 100 (approximately \$1.25) cannot be withdrawn. The account pays no interest, but deposits are free of charge and there are no recurring maintenance fees. The only fees associated with the account are withdrawal fees, which are Ksh 62 (\$0.78) over the counter and Ksh 30 (\$0.38) with an ATM card.<sup>3</sup> Account holders may purchase an ATM card for Ksh 300 (\$3.75), though this is not mandatory.

### 2.3.2 Experimental Design

#### 2.3.2.1 Targeted Population

We examine the impact of providing free ATM cards to a randomly selected subset of 1,113 newly opened Family Bank accounts. These accounts included both joint and individual accounts opened by 748 married couples living in the vicinity of Busia town who did not have a pre-existing account

---

<sup>2</sup>For comparison, the median household in our sample reported Ksh 1,200 in combined income in the week before the survey.

<sup>3</sup>These accounts therefore offered both a negative nominal and real rate of return on savings. Year-on-year inflation averaged around 9 percent for the first 6 months of the study period (July-December 2009) and dropped to around 5.5 percent for the remaining 3 months covered by the study (January-March 2010) (Central Bank of Kenya 2009; Central Bank of Kenya 2010).

with Family Bank but stated that they were potentially interested in opening one. At the outset of the study, we identified communities surrounding 19 local primary schools, which served as group meeting grounds to implement baseline surveys, complete account opening paperwork, and conduct randomization. These schools were located between 0.2 and 7.7 miles from Family Bank's Busia branch, which is situated in the town's commercial center. Targeted communities were located either on the outskirts of Busia town or in nearby rural areas. Figure 1 illustrates the location of the schools relative to Family Bank and our field office on a map.

Trained field officers recruited households in communities surrounding a study school the day before each meeting. With the help of a local guide, they made door-to-door visits to homes in the area and issued meeting invitations to eligible households. To be eligible for invitation, a household had to be headed by a married couple, with both spouses present and able to attend the meeting. Couples were targeted jointly in order to study strategic savings behavior in the household (see Schaner (2011) for details) and to study gender differences in savings behavior and savings account use absent a selection effect. (I.e, it is difficult to answer the question "would targeting men or women for savings accounts lead to greater total savings in formal accounts?" if selection into account ownership is different for men and women. Our sampling strategy eliminates this selection problem, so that we can look at gender differences within a single population of households.) In addition, only households where both spouses had a valid Kenyan national ID card were admitted to the meetings, as Family Bank requires this document of all account holders.<sup>4</sup>

In order to compensate respondents for their time and to provide an additional incentive to attend the meetings, each individual who participated in the study received Ksh 100 in cash at the end of the meeting. Approximately 29 percent of issued invitations were redeemed over the course of the study. While far from universal, takeup rates are high enough that our sample represents a nontrivial fraction of targeted married couples in our catchment area.

### 2.3.2.2 Interest Rates

All couples attending the group meetings were given the opportunity to open up to three Family Bank accounts: an individual account in the name of the husband, an individual account in the name of the wife, and a joint account. To maximize takeup, we funded each opened account with the Ksh 100 (\$1.25) minimum operating balance (this amount could not be withdrawn by participants – it simply made opening an account costless). Each potential account was randomly assigned a temporary 6-month interest rate of either 0, 2, 6, or 10 percent.<sup>5</sup> Joint accounts could earn 2, 6, or 10 percent interest with equal probability while individual accounts could earn 0, 2, 6, or 10 percent

---

<sup>4</sup>This requirement is common to all banks in Kenya. 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.

<sup>5</sup>These percentages are 6-month yields. Annualized yields are approximately double the quoted rates. However, interest was paid on the six-month average daily balance and balances earned no interest thereafter. Respondents were aware of the temporary nature of these interest rates from the outset of the study.

interest with equal probability. This design, illustrated in Figure 2, created random variation in the maximum interest rate available to couples participating in the study. In particular, 7 percent of couples had a maximum interest rate of 2 percent, 31 percent of couples had a maximum interest rate of 6 percent, and 62 percent of couples had a maximum interest rate of 10 percent. We will exploit this variation to calculate the interest value of the free ATM card intervention in the results section.

### 2.3.2.3 ATM Cards

Each newly opened, ATM-eligible account was randomly allocated to either the ATM treatment group (in which case the account received an ATM card for free) or the control group.<sup>6</sup> The free ATM selection probability was 0.15 for the first six experimental sessions (193 open accounts, or 17 percent of all open accounts) and 0.25 for the remaining 27 experimental sessions (920 accounts). Making ATM cards free could impact observed account use through both a direct effect (account use with the card differs relative to the counterfactual) and a composition effect (the pool of open accounts changes). In order to study the direct effect absent the composition effect, we designed the experimental protocol so that respondents did not know if a given account would be selected for an ATM card when they decided which accounts to open.

Since the majority of respondents did not have prior experience with bank accounts (or ATM cards), enumerators carefully explained how the bank accounts and ATM cards worked, as well as the withdrawal fees associated with the accounts and cards. When an opened account was randomly chosen to receive a free ATM card, respondents were informed that the Ksh 300 ATM card fee would be paid on their behalf, and that they could retrieve their card at the bank branch. Due to technical constraints on the part of Family Bank, only one free ATM card was issued for both individual and joint accounts. In the case of joint accounts, it was up to the couple to decide how to allocate the card between spouses.

### 2.3.3 Data

We use two data sources for this project – survey data from one-on-one baseline questionnaires administered during the group sessions (spouses were separated for the interviews) and administrative

---

<sup>6</sup> Respondents were given the choice between two types of joint accounts. The first, dubbed “either to sign” required the consent of either spouse to make withdrawals. These accounts were eligible for ATM cards. The second, dubbed “both to sign” required that both spouses appear in person, together at the bank in order to make withdrawals – as such, these accounts were not eligible for ATM cards. Overall, “either to sign” accounts were much more popular with respondents – 93 percent of couples opening a joint account opted for this type of joint account. We exclude all couples who only opened a “both to sign” joint account from the analysis, and we exclude 8 “both to sign” joint accounts from our account-level analysis of the impact of ATM cards. Results are unchanged if we simply drop these 8 couples altogether. A subset of individual accounts were also randomly selected to be eligible for an information sharing intervention. The details of this intervention are described in Schaner (2011). We do not discuss this intervention further here, as it has no impact on our results.

data on account use from the bank. The baseline survey collected basic demographic information, as well as information on rates of time preference and time inconsistency, decision making power in the household, income, and current use of a variety of savings devices (for details on the time preference questions, see Appendix section A.1. As detailed in the Appendix, the baseline survey asked individuals to choose between different amounts of money at different times in order to directly elicit time preferences. To incentivize the questions, each respondent was given a 1 in 5 chance of winning one of his choices. The majority of respondents chose to have their cash prizes deposited into their newly opened bank accounts). The administrative data provided by the bank includes the first six months' transaction history of all accounts opened under the auspices of the project. Each entry in the transaction history includes the deposit or withdrawal amount, any fees, and the type and time of the transaction.

### 2.3.4 Sample Characteristics and Randomization Verification

The sample consists of 1,113 ATM card-eligible joint and individual bank accounts opened by 748 married couples who had valid national ID cards and no pre-existing accounts with Family Bank. Table 1 presents summary statistics for these couples. Respondents are of relatively low socioeconomic status – husbands average 8 years of schooling, and their wives average just under 6 years. While most men are literate (85 percent), one third of women cannot read and write. On average, men reported earning Ksh 1,662 (about \$21) in the past week, while women reported earning Ksh 815 (\$10). However, median reported weekly incomes are substantially lower, at Ksh 700 and Ksh 300 for husbands and wives respectively. Forty-six percent of men and 41 percent of women reported that they own their own phone. Overall 47 percent of couples reporting that either spouse owns a phone – this rate is low relative to the whole of Kenya, where 75 percent of households report owning at least one phone (Jack and Suri 2011).

Almost all respondents reported using at least one savings device at baseline, with use of informal devices much more common than use of formal or semi-formal devices. Most common was saving cash at home, reported by 85 and 90 percent of husbands and wives respectively. Reported savings levels at home were substantial and approximately equal to average weekly earnings. ROSCAs were also very popular, with 49 percent of men and 66 percent of women reporting belonging to at least one group.<sup>7</sup> Savings accounts with formal banks were less common, particularly for women – while 32 percent of men reported owning a savings account (and those men reported substantial savings in their accounts), just 12 percent of women reported owning a savings account. These numbers are very similar to reported use of mobile phone money storage technologies, though bank accounts have much larger balances. Least common was ownership of a SACCO account – just 7 percent of men

---

<sup>7</sup>ROSCA stands for “rotating savings and credit association”. ROSCAs consist of a group of individuals who meet at predetermined intervals (e.g. weekly, monthly) to put a fixed amount of money into a common pot. At each meeting, a different member of the group receives the pot. ROSCAs are by nature illiquid and often quite risky, as group members can defect before the ROSCA cycle is completed.

and 1 percent of women reported belonging to a SACCO at baseline.<sup>8</sup> The large savings balances in bank accounts and SACCO accounts in part reflect higher incomes of these account owners. However, the secure and formal nature of these devices may also make them more appealing for storing large sums of money: the average home saver stored 1.5 weeks of his income at home, while bank account owners stored 7.8 weeks of their income at the bank and SACCO accounts were used to store 38 weeks of income.

Although women in developing countries are generally thought to have less bargaining power than their husbands, questions regarding savings decision making power reveal that both genders frequently reported that women made savings decisions (49 percent of women and 43 percent of men). This may also reflect social norms: women in developing countries are often tasked with investing in the needs of the family and household (Bruce 1989). Still, a substantial share of respondents reported that men were responsible for deciding how much to save (28 percent of women and 37 percent of men). In contrast, few individuals reported joint or independent decision making regarding savings.

Table 2 checks that randomization of free ATM cards, of cash payments made to incentivize discount rate elicitation, and of the maximum interest rate available to the couples was successful. Since randomization was conducted in the field, with respondents drawing folded envelopes from tins, we check that (1) proportions treated do not differ from their theoretical selection probabilities and (2) treatment status is uncorrelated with observable demographic characteristics. Panel A of each table displays the results of the first exercise. P-values from a binomial test that realized proportions are equal to theoretical proportions are reported in braces. Overall, realized probabilities for ATM cards and cash prizes are slightly lower than, though not significantly different from, theoretical probabilities. Theoretical and actual probabilities are also very close for the maximum interest rate.

Panel B presents the results of separate regressions of demographic characteristics on treatment indicators. The first two columns are limited to characteristics of either the husband or the wife, whereas the last three columns use both husband and wife demographic characteristics. Since ATM card and cash prize receipt are binary treatments, we report coefficients and standard errors on treatment indicators in the first four columns. Since the maximum interest rate could take on values of 2, 6, or 10 percent, we regress demographic characteristics on dummy variables for 6 and 10 percent interest, and then present results of an F-test that these dummy variables are jointly equal to zero.

Overall, the randomization appears to have functioned well, with significant differences appearing at a rate approximately equal to that which would appear due to chance. We do note that cash prize provision for women is significantly (and negatively) correlated with ATM card provision for

---

<sup>8</sup>This is not surprising given the low incidence of formal sector employment in our sample. SACCO stands for "Savings and Credit Co-Operative". In Kenya they function much like credit unions and are organized around higher paying formal occupations such as teaching and commercial farming.

women's accounts. Since cash prizes could be deposited into bank accounts, and since the majority of individuals chose to do so, prize provision significantly increased measures of bank account use. For this reason, we control for cash prize receipt throughout our analysis (results are also robust to controlling for a laundry list of demographic characteristics).

## 2.4 Results

### 2.4.1 Summary of Account Use

Table 3 presents account use summary statistics. To give a sense of rates of account use absent the cash deposits from time preference elicitation, we present summary statistics for the entire sample (Panel A) and for the subset of couples where neither spouse was randomly selected to receive a cash prize (Panel B). Columns 1-3 present means and standard deviations of a variety of account use measures by account type. In these columns we drop individual accounts randomly selected to receive no interest to ensure that the interest rate compositions of the three account types are comparable. For each account type, we present the share of potential accounts that were actually opened, and then limit the sample to open accounts that were not randomly selected for the ATM treatment and present averages for measures of account use including savings rates, the average daily balance, the number of deposits and withdrawals, and the amount of transaction fees paid. Columns 4-6 present differences (and associated standard errors) in account use measures between account types. Finally, in order to give a picture of *overall* use of Family Bank accounts by couples in the study, the last column presents a summary of aggregate account use at the couple level. This column includes all couples and accounts, and displays the average number of accounts opened by a couple, the share of couples who saved in any account at all, the average daily balance stored in all accounts, and so on.

Overall, joint accounts were much more popular than individual accounts – the first row of each panel illustrates that two thirds of couples opened a joint account, whereas 45-47 percent of couples opened individual accounts, with no significant difference in the rate of account opening for men's and women's accounts. Even though we targeted a sample of individuals who stated that they were interested in opening a bank account, rates of actual account use were relatively low. Column 7 illustrates that just 45 percent of couples saved in any one of the three accounts, and this number drops to 27 percent when we exclude couples who won at least one cash prize.

These usage rates are notably lower than those documented by Prina (2011) (among a sample of Nepalese women in an urban slums) and Dupas and Robinson (2011), among a sample of Kenyan small-scale entrepreneurs.<sup>9</sup> Inspection of columns 4-6 of Table 3 also reveals that even conditional on

---

<sup>9</sup>The contrast with Prina (2011) likely reflects very different contexts and savings products (the accounts in her study bore 10 percent nominal interest, had no use fees, and had lower time and travel costs to go to the bank branch). In contrast, Dupas and Robinson (2011) conducted their study in a similar part of Kenya, and the bank accounts offered to both study populations had similar fees. In this case the differences in usage rates are likely

opening, joint accounts were substantially more likely to be used for saving than individual accounts. Joint accounts also appear to be used by more small-scale savers – despite the higher savings rates, joint accounts do not have higher average daily balances when compared to individual accounts. Men and women are just as likely to save in their individual accounts, and also have very similar average daily balances. Although differences are not significant, men do appear to transact more than women – they make more withdrawals (though the number of deposits are very similar), and consequently pay more fees. This finding also contrasts sharply with Dupas and Robinson (2011), who find that the women in their sample made much more use of accounts as compared to men. This is likely due to a selection effect – women and men in their sample engage in different income generating activities and may well have belonged to very different types of households (indeed, the women in their sample have higher incomes than the men on average). In contrast, we compare the behavior of men and women in the same set of households.

Figure 3 explores the dynamics of account use, exploiting the fact that we have time series data on account activity. For each account in the control group (i.e. not selected for an ATM card) that was assigned an interest rate of 2 percent or better, we calculate whether an account was used for saving, the closing balance, the average daily balance, the number of deposits, and the number of withdrawals at the end of each week from the first week to the 26th week following account opening, which marks the end of our study period. We exclude couples who were randomly selected for cash prizes, as the cash prize deposits could lead us to falsely conclude that couples make sustained use of their accounts. The time series reveals that account activity does not trail off after the first few weeks of account opening – the number of deposits and withdrawals continue to grow as time progresses. Furthermore, a nontrivial share of savers take a great deal of time to make their first deposit into an account – though 19.6 percent of joint accounts received a deposit within the first 5 weeks of account opening, another 5 percent of accounts received a deposit over the ensuing 19 weeks.

Though many savers use their accounts infrequently and store small sums, a small number of the couples in our sample used their new accounts very intensively. Table 4 limits the sample to those couples who saved in at least one account and presents quantiles of aggregate savings measures. While the median couple who saved (in the absence of cash prizes) accrued an average daily balance of just Ksh 844 (\$10.55) in all its Family Bank accounts, the 75th percentile couple saved Ksh 1,584 (\$19.80), and the 99th percentile couple saved Ksh 14,439 (\$180.49). This skewness is also present in transaction volume – the median saving couple made four deposits (including the opening deposit made on their behalf at the beginning of the study) and one withdrawal, while the 99th percentiles are 24 deposits and 21 withdrawals.

It is possible that this dispersion in account use is driven by a small number of couples who have

---

because Dupas and Robinson exclusively targeted small-scale entrepreneurs, who may have had a greater need for formal bank accounts. We also offered each individual Ksh 100 in compensation for participating in our baseline sessions. This may have selected some couples who only had limited interest in bank accounts but low opportunity costs of time.



very high incomes. In fact, this is not the case. Table A1 regresses aggregate couple-level account use (among couples who did not win cash prizes) on demographic and economic characteristics. Conditional on a variety of demographic characteristics, men’s income does not significantly predict account use, and women’s reported income is actually *negatively* correlated with account use (though other markers associated with women’s socioeconomic status, such as education, phone ownership, and bank account ownership are all positively and significantly correlated with aggregate savings activity). The  $R^2$  on the regressions in table A1 range from 0.13-0.15. It therefore appears that key determinants of intensive account use in our sample are largely unobservable.

### 2.4.2 Impact of ATM Card Provision

It is not entirely unexpected that the use of opened accounts is so low in the control group – transaction costs associated with the accounts are large, particularly given that most savers do not save very much (for example, a single withdrawal fee represents 7.34 percent of the average daily balance of the median saving couple in panel B of Table 4). We now ask whether substantially reducing these costs through ATM card provision can meaningfully increase account activity. Table 5 presents estimates of the impact of ATM card provision on a variety of account use measures. All regressions are of the following form:

$$y_{ac} = \beta_0 + \beta_1 \text{freeatm}_{ac} + x'_{ac} \delta + \epsilon_{ac} \quad (2.1)$$

where  $y_{ac}$  is the outcome of interest, as measured 26 weeks after account opening,  $\text{freeatm}_{ac}$  is a dummy variable indicating that account  $a$  owned by couple  $c$  was selected to receive a free ATM card, and  $x_{ac}$  is a vector of controls. These controls include a dummy variable for the first six experimental sessions (since ATM selection probability was lower for these sessions), account type dummy variables (when relevant), separate dummy variables for husband and wife’s cash prize receipt, and the interaction of these dummies with the account type dummies. In all regressions we limit our attention to open accounts, since ATM cards were randomly allocated conditional on account opening. We examine the impact of ATM card receipt by account type (Panels A-C) and pooled for all accounts (Panels D and E).

The first column of Table 5 reports the first stage – the impact of the free card treatment on whether or not an account had an active ATM card. ATM card fees are quite significant given the low incomes of our study population: in the control group, respondents purchased ATM cards for just 11 percent of their accounts (Panel D). Since the free card treatment ensured that an account received an ATM card, the first stages are very substantial. As such, we focus on the reduced form impact of free ATM card provision for the remainder of the analysis.

The subsequent columns of Table 5 examine the impact of ATM card provision on several measures of account use that, in the model absent security concerns presented in Section 2.2, should increase when withdrawal fees fall – these include a dummy variable indicating that an

account was used for saving, the average daily balance in the account, the number of deposits, and the number of withdrawals.<sup>10</sup> Given the large number of outcomes, we also present the mean effect for these outcomes in the second column, where we follow the methodology of Kling, Liebman, and Katz (2007) to construct a measure of standardized account activity. We standardize each variable relative to the subset of the control group that did not receive a cash prize (pooling all accounts). This way the magnitude of the treatment effects are in the same standard deviation units for all accounts types (as such, the dependent variable mean is only equal to 0 for the omitted group in panels that pool all accounts – results are very similar when we standardize outcomes relative to the control group mean for each account type separately). We also separately present estimates of the impact of ATM card provisions on fees paid. As highlighted by the theory, reducing the withdrawal fee will mechanically reduce total fees paid, but total withdrawals will increase (absent security concerns) – so the net effect is ambiguous.

Inspection of Panels A-C of Table 5 reveals positive impacts of ATM card provision for joint and men’s accounts. These impacts are large relative to dependent variable means – in particular, column 2 illustrates that ATM card provision resulted in a 0.305 standard deviation increase in joint account use and a 0.266 standard deviation increase in husbands’ account use. ATM cards had a substantially larger impact on withdrawals as compared to deposits (relative to dependent variable means) – the ATM treatment increased the number of deposits into joint accounts by 11 percent, but increased the number of withdrawals by 196 percent. The analogous point estimates for husbands’ accounts imply increases of 18 percent and 223 percent for deposits and withdrawals respectively.<sup>11</sup> Given the very large elasticities of withdrawals with respect to the fee (-3.8 for joint accounts and -4.3 for husbands’ accounts), total fees paid on both types of accounts more than doubles, though all of these estimates are somewhat imprecise. In contrast, point estimates for wives’ accounts are much smaller and almost always negative, though we cannot reject that any of them are equal to zero (or positive).

Despite the negative point estimates for wives’ accounts, Panel D illustrates that when pooling all accounts, the ATM card treatment significantly increased account use by 0.191 standard deviations. The negative estimates for women are quite striking, especially given that the patterns and magnitudes are so different from those observed for joint and men’s accounts. Panel E pools all accounts and tests whether impacts for women are significantly different than those for other accounts by including an interaction between the free ATM indicator and the wife’s account indicator.

---

<sup>10</sup>Since all open accounts had a positive average daily balance and a small number of accounts had very large average daily balances, as documented in Table 4, we use the log of the average daily balance in this and all subsequent analysis. Results are similar, though less precisely estimated, if we use the level instead.

<sup>11</sup>One complication associated with comparing these treatment effects relative to dependent variable means is that the mean for deposits includes one deposit made on behalf of respondents for each open account. This will push the mean of deposits in the control group up while leaving the treatment-control difference unchanged, thereby pushing the percentage increase down. However, if we omit these opening deposits, we still estimate that ATM card provision increased deposits into joint accounts by 17 percent and deposits into husband’s accounts by 31 percent – in both cases these impacts are substantially smaller than the percentage impacts for withdrawals.

When pooling accounts this way, we estimate large, significant impacts for joint/husbands accounts (a mean impact of 0.29 standard deviations) and impacts that do not significantly differ from zero for women. We also strongly reject that the treatment effect for women’s accounts is equal to the treatment effect for other accounts.

The estimated impacts on the use of joint and men’s accounts are large – however, an account-level analysis cannot determine whether the effect reflects an *aggregate* increase in couples’ use of bank accounts, or if the impact is driven by *substitution* between multiple accounts owned by the same couple. To test between these two alternatives, Table 6 examines the effect of the ATM treatment on couple-level aggregates of account use. This analysis also allows us to compare the impact of ATM card provision to the impact of interest rates by estimating how the maximum interest rate available to the couple affected aggregate account use. Panel A of Table 6 presents results from the following regression:

$$y_c = \beta_0 + \beta_1 freeJ_c + \beta_2 freeH_c + \beta_3 freeW_c + x_c' \delta + \epsilon_c \quad (2.2)$$

where  $y_c$  is the outcome of interest,  $freeJ_c$ ,  $freeH_c$ , and  $freeW_c$  indicate selection for the ATM treatment for the joint, husband’s and wife’s account respectively, and  $x_c$  is a vector of controls. Since the likelihood that a couple received an ATM card depended on whether or not an account was opened,  $x_c$  includes dummy variables that fully saturate every combination of account opening choices as well as a dummy variable for the first 6 experimental sessions and two separate dummy variables indicating husband and wife cash prize selection.<sup>12</sup>

Panel B presents results of a specification where the account-specific treatment indicators are replaced with a dummy variable that is equal to one if the couple was selected for *any* free ATM card. To test whether the treatment effect on wives’ accounts differs from the treatment effect of other accounts, Panel C presents a specification where two treatment dummy variables are included – one indicating that a couple received a free ATM card on the joint or husband’s account, and one indicating that a free ATM card was given to the wife’s account. We standardize aggregate account activity relative to couples who were not selected for any free ATM cards and did not receive any cash prizes.

Panels A-C confirm that the substantial account level impacts in Table 5 reflect increases in aggregate account use at the couple level. The first column of Panel C illustrates that receipt of a joint or husband’s ATM card results in an increase in aggregate account use of 0.239 standard deviations. As before, the impact on withdrawals is much larger, relative to control group dependent variable means, when compared to the impact on deposits.<sup>13</sup> The F-tests reported in Panel C

---

<sup>12</sup>The combination of account opening choices separately accounts for the 8 joint accounts that were opened but not eligible to receive a free ATM card because they were “both to sign”. In total, 10 different account opening combinations were realized, though most couples either opened just an “either to sign” joint account (53 percent of the sample) or two individual accounts (31 percent of the sample).

<sup>13</sup>As before, this result holds even when we omit opening deposits made on behalf of respondents, which reduces

clearly reject that the impact of wives' ATM cards and joint/husbands' ATM cards are the same. However, an important caveat is in order – although we can state whether or not the ATM treatment increased saving in Family Bank accounts, we cannot tell if this savings represents crowd-out from other savings devices, or if this represents mostly new savings. This is a limitation of our data collection strategy, which did not include an endline survey.

Panel D studies the impact of the maximum interest rate available to the couple on aggregate account use. Note that in this case, it is not appropriate to control for combinations of open accounts (since account-specific interest rates had a robust impact on the decision to open a given account), so we estimate the impact of interest rates in a separate regression of the following form:

$$y_{ac} = \beta_0 + \beta_1 max6_c + \beta_2 max10_c + x'_c \delta + \epsilon_c \quad (2.3)$$

where  $max6_c$  and  $max10_c$  indicate that the maximum interest rate available to couple  $c$  was 6 percent and 10 percent respectively and  $x_c$  includes a dummy variable for the first 6 experimental sessions and separate dummy variables for cash prize selection of men and women. For the mean effect, we use the same standardization (treating couples not selected for any free ATM card or cash prize as the reference group) in order to ensure that magnitudes are comparable across panels in the table.

Panel D illustrates that aggregate account use responded robustly to the maximum interest rate – in particular, couples who received a maximum rate of 10 percent had aggregate account activity 0.232 standard deviations above that of those couples who received a maximum rate of 2 percent. This impact is nearly identical to the impact of providing a free ATM card to the joint or husband's account. However, from the standpoint of a policymaker, increasing account use by providing ATM cards is notably more cost effective. In particular, consider the policy of purchasing a free ATM card for all joint and husbands' accounts. In the absence of the free ATM treatment, couples purchased ATM cards for just 7.7 percent of joint accounts and 7.6 percent of husband's accounts. Given the opening rates in our sample, free provision would therefore result in lost ATM card fee revenue of Ksh 25.8 per couple, while increasing transaction fee revenue by Ksh 26.0 (here we use the couple-level estimate in Table 6). The subsidy therefore *increases* revenue by Ksh 0.2. Furthermore, examination of the ATM treatment effect week by week (not presented in this paper) reveals that the treatment effect on fees for husband's and joint accounts grows over time.<sup>14</sup> It therefore seems likely that the subsidy would result in more revenue gains as the time horizon expands.<sup>15</sup> In contrast, the average interest payout to couples with a maximum interest rate of 2 percent (and no cash prizes) was Ksh 3.73, while the average interest payout to couples with a maximum interest rate of 10 percent was Ksh 48.5. At the same time, couples with a maximum interest rate of 10 percent paid an average of Ksh 25.2 more in fees, so the net cost of the interest

---

the mean of deposits in the control group (with no cash prizes) to 0.85.

<sup>14</sup>These results are available upon request from the author.

<sup>15</sup>Shifting transactions to the ATM would also reduce the need for tellers at the bank, reducing labor costs.

rate subsidy would be Ksh 19.6 per account.<sup>16</sup>

The average treatment effects presented in Tables 5 and 6 are large. Given the very skewed distribution of account use illustrated by Table 4, it is natural to ask whether this only reflects a very large impact at upper quantiles of the distribution of account use, or if the treatment effect is more diffuse. Figure 4 graphs the CDFs of couple-level aggregate account use for six groups. Panel A limits the sample to couples who opened a joint account and plots the CDF of standardized aggregate account use for couples who did and did not receive the free ATM treatment on the joint account. Panels B and C repeat the exercise for men's and women's individual accounts, in each case limiting the sample to couples who opened the relevant account. For both joint and husband's accounts, the CDF for couples who received a free ATM card is almost everywhere below the control group CDF, suggesting that the ATM card had a substantial impact at a variety of quantiles (of course, given the large point masses of couples who did not save at all, the treatment effect must necessarily be concentrated at upper quantiles). In contrast, inspection of Panel C reveals that the CDF for couples in which women's accounts received the free ATM treatment is almost everywhere above the CDF of the control group.

To test for significance of these CDF differences while simultaneously controlling for cash prize receipt, Figure 5 plots point estimates and 95 percent confidence intervals from quantile regressions following the specification given by equation 2.2. Panel A graphs quantile treatment effects for a free ATM card to the joint account, while Panels B and C present analogous results for husbands' and wives' accounts. The figure only shows estimates for quantiles 70-99 – since so many couples did not save at all, the estimated treatment effect is identically zero at almost all quantiles below 70 in each panel.

Over quantiles 70-92, the treatment effects for joint accounts average 0.19, while the treatment effects for husbands' accounts average 0.27. At quantiles 93-99, the estimated treatment effects become very large – particularly for joint accounts. Under the assumption of rank invariance, this suggests that the couples who would have made the most intensive use of bank accounts in the absence of the treatment also benefit disproportionately from ATM cards.

As expected, a distinctly different pattern emerges when studying the quantile treatment effects for wives' accounts in Panel C. Nearly all estimated treatment effects are negative, with an average of -0.395 over quantiles 70-99, and many estimates are statistically different from zero at the 90 or 95 percent level. This suggests that giving free ATM cards to women did not just result in no change in savings behavior, it may have actually *negatively* impacted account use for some women.

This heterogeneous treatment effect is striking: while women's accounts respond to the ATM treatment quite differently than joint/men's accounts, both men and women have remarkably similar rates of saving and average balances in the control group (recall Table 3). The gender difference is unlikely to be driven by selection on *couple* level unobservables, since couples generally opened both

---

<sup>16</sup>The external validity of this comparison is questionable, however, since ATM card provision was done conditional on account opening, while interest rates were randomized unconditional on account opening.

individual accounts together: 86 percent of couples who opened an account for the husband also opened an account for the wife and 89 percent of couples who opened an account for the wife also opened an account for the husband. The question then becomes, what is it about women that makes their response to the ATM treatment so different? Some explanations, such as financial literacy or prior exposure to the formal sector, could explain a zero impact on account use but cannot explain a negative impact. Although women are more likely than men to be time inconsistent (Table 1), we do not have any strong evidence that this is a key factor governing differential responses to the ATM treatment (see Appendix section A.2 for supporting analysis). This leaves differences in bargaining power between husbands and wives as a leading explanation. One may ask “if the ATM card harmed the security of a wife’s account, why wouldn’t she just throw it away?”. Our ATM card randomization was conducted when the couples were sitting together, so card receipt was public information. This may have made it difficult for a wife to simply dispose of the card if it was of interest to her spouse. The next section investigates the bargaining power hypothesis.

### **2.4.3 Bargaining Power and ATM Card Treatment Effects**

A growing body of literature documents that individuals make financial decisions strategically in order to manipulate intrahousehold resource allocation in their favor. In a lab experiment in the Philippines, Ashraf (2009) documents that spouses who report that they do not have control over savings decisions allocate experimental winnings so as to increase their own personal consumption. In India, Mani (2010) finds that individuals in married couples are willing to sacrifice experimental earnings in order to ensure that these earnings are deposited in their own individual bank accounts rather than a spouse’s account. In Kenya, Anderson and Baland (2002) show that ROSCAs are most popular among women with intermediate levels of bargaining power – the authors hypothesize that at these levels of bargaining power, women can use ROSCAs to tilt household consumption towards goods that they favor. Finally, using the same sample of couples considered in this paper, Schaner (2011) presents evidence that couples who are badly matched in terms of rates of time preference strategically use bank accounts to manipulate household savings levels.

In all the studies listed above, choices with strategic value enable the decision maker to securely sequester or hide resources from his or her spouse. In our context, individual bank accounts are useful for this reason because they can only be accessed by the account owner. However, ATM cards may dilute the strategic value of individual accounts – suppose a husband learned his wife’s ATM passcode – then he would be able to make withdrawals from (and learn the balance of) her account through the ATM. Absent the ATM card, he would have to force the wife to make withdrawals herself, in person, at the bank. If spouses make use of individual accounts strategically and value these accounts for security, then making an ATM card available could make the individual account less attractive. Furthermore, when the withdrawal cost is reduced, it may be more difficult for an account holder to refuse to make a withdrawal for her spouse (or for another member of her family or the community). ATM cards may be particularly unattractive to spouses with low bargaining

power, since they already have difficulty resisting demands made by their partners.<sup>17</sup> Since women generally have less bargaining power than men in Kenya, this could explain why a negative impact is only visible for women. To test this hypothesis, we study how responses to individual ATM cards differ between households in which the wife has higher versus lower levels of bargaining power.

We use intrahousehold differences in demographic characteristics to measure women’s relative power. In particular, we assume that having higher income, having more years of education, being more literate, and being older than a spouse correlate with greater relative bargaining power. First, we standardize each of these four variables at the individual level by subtracting the sample mean and dividing by the sample standard deviation. We then proxy the wife’s relative bargaining power by the average difference between her values for these variables and her husband’s values for these variables:

$$power_c = \frac{1}{4} \sum_{x \in X} (x_c^w - x_c^h)$$

Figure 6 plots the histogram of  $power_c$  among the 698 couples where both spouses had nonmissing values for the four demographic characteristics included in the index. As expected, husbands have more proxied bargaining power than wives – just 17 percent of women have greater proxied power than their husbands and the median difference between wives and husbands is -0.42 standard deviation units.

To check that this index has informative content about household savings behavior, Figure 7 presents results of local linear regressions of baseline self-reported account use on the bargaining power index. Indeed, savings device use for husbands (panel A) and wives (panel B) is generally correlated with the index and results symmetric across genders. When men have relatively more bargaining power, they are less likely to save at home and more likely to save at the bank/SACCO and on the phone. In contrast, women are more likely to save in the bank/SACCO and on the phone and less likely to save at home when *they* have more bargaining power. These correlations are consistent with the idea that spouses with more bargaining power are more economically empowered. On the other hand, this is not the only way to interpret the data – if individuals with less bargaining power differentially value security, it is surprising that they are also more likely to save at home.<sup>18</sup>

Since so few women have greater absolute proxied bargaining power when compared to their husbands, we define a woman to be “relatively advantaged” in terms of bargaining power if  $power_c$  is above the sample median. The following account-level specification studies heterogeneous impacts

---

<sup>17</sup>It is not obvious that this relationship should be linear. As highlighted by Anderson and Baland (2002), women with very low bargaining power may simply forfeit control of accounts to their husbands regardless of the ATM card. However, there is no evidence of such a nonlinear relationship in our data, so we focus on the distinction between high and low bargaining power.

<sup>18</sup>We also note that we do not find the patterns with respect to bargaining power and ROSCA use that are highlighted by Anderson and Baland (2002), even though relative income is an important input into our indicator.

of ATM cards by bargaining power:

$$y_{ac} = \beta_0 + \beta_1 \text{freeatm}_{ac} + \beta_2 \text{freeatm}_{ac} \times \text{wifeadv}_c + \beta_3 \text{wifeadv}_c + x'_{ac} \delta + \epsilon_{ac} \quad (2.4)$$

where all variables are as defined in equation 2.1 and  $\text{wifeadv}_c$  is either an indicator that the wife's relative bargaining power in couple  $c$  is above the sample median, or the wife's decision making index.<sup>19</sup>

Table 8 presents the results of this specification for men's and women's accounts separately (Panels A and B). Panel C presents results from a couple-level specification where equation 2.2 is augmented to include interactions of each ATM card treatment with the bargaining power variable (coefficients for the joint account are omitted from the table for clarity). We also limit the sample in this specification to those couples who opened at least one individual account, since we are interested in studying the impact of ATM cards on individual accounts. Since the bargaining power index is correlated with baseline bank account use, we present a specification in which we only control for "basic controls" (those control sets used in Table 5 and 6), as well as two specifications in which we progressively control for baseline savings device use and its interaction with the ATM treatments. The second specification adds controls for formal account ownership (here, we define both bank accounts and SACCO accounts as formal accounts, since they have similar features and are often used to store large amounts of savings at baseline) and interactions with the ATM treatments. We then add controls for mobile money, ROSCA, and home savings (and their interactions with the ATM treatment indicators) in the third specification. Note that the interpretation of the free ATM main effect therefore changes from specification to specification. In the first column, it is the ATM treatment effect for households in which the husband has the most bargaining power. In the second specification, it is households in which the husband has the most bargaining power *and* neither spouse owned a formal account at baseline. In the third specification, the main effect is for households in which neither spouse had a formal account, a mobile money account, any home savings, or any ROSCA memberships. No couples actually meet this criterion (recall the very high rates of home saving at baseline), so we do not focus on the main effect in the latter two specifications.

Panel A reveals that men's positive response to free ATM cards is driven entirely by households where men have more proxied bargaining power, even though these two groups of households did not make significantly different use of men's accounts in the absence of the ATM card treatment. The coefficient on the ATM  $\times$  bargaining power interaction is relatively robust to adding additional baseline savings controls, though it is no longer significant when we include the full baseline control set.

---

<sup>19</sup>Results are similar using an absolute threshold of  $\text{power}_c \geq 0$  or the continuous index value, though they are less precise. Proxied bargaining power is missing for 50 of 748 couples due to missing input variables. For these couples, we set  $\text{wifeadv}_c = 0$ , dummy them out separately, and include an interaction of the missing dummy with the free ATM card dummy.



In contrast, Panel B illustrates the opposite pattern for wives' accounts. In particular, women in households where they have relatively little bargaining power displayed a sizable negative response to ATM cards, while women who have more bargaining power responded positively to the ATM treatment (however, the sum of the main effect and the interaction term is not significantly different from zero for any of the specifications). Again, the interaction term is robust to adding baseline account controls. Another notable pattern present in Panel B is that in the absence of ATM cards, women with relatively more proxied bargaining power actually make less use of their individual accounts. This may be because women with relatively little bargaining power make excessive use of their secure individual accounts in order to manipulate household consumption and/or savings allocations, or because these women were more likely to have a bank account at baseline. Panel C demonstrates that these patterns persist when examining aggregate couple-level account use, though significance of the results is somewhat attenuated. Overall, our results are very consistent with the hypothesis that security concerns are an important mediator of individuals' responses to ATM cards.

## 2.5 Conclusion

We present results from a field experiment conducted in Western Kenya with a low-income, mostly rural sample of married couples. A subset of newly opened formal bank accounts belonging to these couples were randomly selected to receive a free ATM card. ATM cards reduced withdrawal fees by over 50 percent and also made the accounts easier to access and less secure. Overall formal account use in our sample is relatively low – 27 percent of couples used at least one newly opened account for saving and the median saving couple had a 6-month average daily balance of \$10.55. The free ATM treatment significantly increased overall account use. However, inspection of impacts by account type reveals a striking heterogeneous treatment effect by gender. When joint or husband's accounts were selected for the free ATM treatment, aggregate account use by couples increased by 0.24 standard deviations. This impact is substantial, and equivalent to offering the couple 8 more percentage points of interest on their accounts. When the wife's account was selected for the free ATM treatment, aggregate account use *decreased* by 0.13 standard deviations. Furthermore, quantile treatment effects for free ATM provision on wives' accounts are often large, negative, and significantly different from zero.

We hypothesize that this heterogeneous treatment effect is largely driven by a security effect: when individuals in the household have weak bargaining power, the ATM treatment may do more harm than good because the card makes it more difficult for individuals to guard their savings from appropriation by other members of the household. When women have less bargaining power than men on average, this could generate the patterns that we observe in the data. To test this, we proxy relative bargaining power between men and women by intracouple differences in demographic characteristics. We find that the positive response to men's ATM cards is concentrated among

couples where men have above median bargaining power, while the negative response to women's ATM cards is concentrated among couples where women have below median bargaining power. These results suggest that security of savings is very important to individuals with poor bargaining positions within the household. However, our results are not robust to proxying bargaining power within the household with spousal self reports of decision making power over savings. This could be because self reports of decision making power reflect not just bargaining power, but satisfaction with household savings allocations (which could be driven by successful strategic savings behavior on the part of individuals or because both members of the couple have similar preferences). Given the different results associated with these two measures, additional research is needed to clarify the relationship between relative bargaining power and a preference for secure savings devices.

Our results add to a small but growing literature that studies the use of formal bank accounts by low-income individuals in developing countries (Dupas and Robinson 2011; Prina 2011; Ashraf et al. 2006; Ashraf et al. 2010; Brune et al. 2010). They also add to a literature demonstrating that issues of control mediate the use of different savings technologies in large and important ways (Anderson and Baland 2002; Ashraf 2009; Ashraf et al. 2010; Chin et al. 2010; Schaner 2011). Finally, our results have implications for the design of formal savings products targeted to poor households. First, reducing transaction costs substantially increases the use of formal sector bank accounts. At the same time, making formal accounts more easily accessible may actually decrease use of formal accounts by individuals with poor bargaining positions within the household. In our sample, and in developing countries more generally, this often means women. This is particularly important to keep in mind given the recent policy interest in mobile money. In many ways, this technology is easily appropriated by other members of the household – a husband may be able to learn a wife's savings balance by simply catching a glimpse of her phone. Our study cannot determine whether women's account use was depressed due to the withdrawal fee reduction, or due to the fact that the ATM card made the account easier to access and possible to access remotely, or both. Better understanding how these factors mediate formal account use is an important area for future research and will provide useful insights for designing formal accounts that best meet the needs of the poor.

## 2.A Appendix: Additional Time Preference Analysis

### 2.A.1 Survey Questions on Rates of Time Preference

As part of the baseline, each respondent was asked a series of questions designed to elicit discount factors. We chose to elicit time preferences using choices between different amounts of money at different times, as opposed to different amounts of goods at different times. We made this choice for two reasons. First, Ashraf, Karlan, and Yin (2006) find that while time preference parameters estimated using choices between money, rice, and ice cream were all correlated, only the parameters estimated using money choices significantly predicted takeup and use of a commitment savings product. Second, cash lotteries made intuitive sense to respondents given that the group meetings revolved around bank accounts and savings.

We framed all questions as a choice between a smaller amount of money at a nearer time  $t$  ( $x^t$ ) and a larger amount of money at a farther time  $t + \tau$  ( $x^{t+\tau}$ ).<sup>20</sup> In order to make choices salient, respondents were given a 1 in 5 chance of winning one of their choices. Enumerators also used calendars to visually show respondents the number of days they would have to wait for both the smaller and larger amount of money.

In total, participants responded to 10 tables of monetary choices, with each table consisting of 5 separate choices between a smaller Ksh  $x^t \in \{290, 220, 150, 80, 10\}$  and larger  $x^{t+\tau} = \text{Ksh } 300$ . This was a sizable amount of cash for the study participants. (For comparison, median reported daily earnings in our sample were Ksh 100 for men and Ksh 43 for women). The 10  $(t, t + \tau)$  pairs were:  $(\frac{1}{7}, 1)$ ,  $(\frac{1}{7}, 2)$ ,  $(\frac{1}{7}, 3)$ ,  $(\frac{1}{7}, 4)$ ,  $(\frac{1}{7}, 8)$ ,  $(\frac{1}{7}, 12)$ ,  $(2, 3)$ ,  $(2, 4)$ ,  $(4, 8)$ , and  $(4, 12)$  weeks. We chose to set the lowest near term  $t$  to "tomorrow" ( $\frac{1}{7}$ ) instead of "today" (0) to avoid confounding our discount factor estimates with differences in transaction costs of obtaining the funds in the near versus far term, or degrees of trust as to whether the money would be delivered (Harrison, Lau, Rutstrom, and Sullivan 2004).

We can measure preference reversals (of both the hyperbolic, impatient-now, patient-later type, as well as the anti-hyperbolic patient-now impatient-later type) by comparing responses to the last four tables of questions to their analogues that involves choices between cash tomorrow and cash at a later date. (An important drawback of using "tomorrow" instead of "today" as the nearest choice is that we will not be able to detect hyperbolic discounting that discounts all future consumption relative to immediate consumption – this will lead us to underestimate the degree of hyperbolic discounting in our sample). If a respondent won one of her choices, she had the option of having the funds deposited directly in her bank account, or picking the cash up at our field office, also located in Busia town.<sup>21</sup>

---

<sup>20</sup>This method is common to most empirical studies that attempt to measure rates of time preference in developing countries. Examples include Ashraf, Karlan, and Yin (2006), Bauer and Chytilová (2009), Dupas and Robinson (2011), Shapiro (2010), and Tanaka, Camerer, and Nguyen (2010).

<sup>21</sup>Despite the fact that the field office and Family Bank were proximately located, and that accessing cash deposited in an account would entail paying a withdrawal fee, the majority of cash winners (77 percent) chose to have their

For the purposes of this study, we define an individual to have “hyperbolic” preferences if he or she exhibited impatient-now, patient-later preference reversals on at least 2 out of 4 of the relevant pairs of tables. Similarly, we define an individual to have “anti-hyperbolic” preferences if he or she exhibited patient-now, impatient later preference reversals on at least 2 out of 4 of the relevant pairs of tables (as a result of this definition, 32 individuals are coded as both hyperbolic and anti-hyperbolic). We computed discount factors using a simple *ad hoc* bounding strategy similar to that found in Meier and Sprenger (2010). Specifically, suppose that for individual  $i$ ,  $x_q^t \succ 300^{t+\tau}$ , but that  $300^{t+\tau} \succ x_{q+1}^t$  (where  $x_q^t > x_{q+1}^t$ ). We then assume that the individual is indifferent between Ksh 300 at time  $t + \tau$  and the midpoint of the two amounts at time  $t$ ,  $\bar{x}_{q,q+1}^t = \frac{x_q^t + x_{q+1}^t}{2}$ . Using this midpoint, we can then calculate the implied discount factor  $\hat{\delta}_i^{q,q+1} = \left(\frac{\bar{x}_{q,q+1}^t}{300}\right)^{\frac{1}{\tau}}$ . We do this for each table, obtaining 10 discount factor estimates, and take the simple average of them.

## 2.A.2 Time Inconsistency and ATM Card Treatment Effects

One reason why women may respond negatively to the ATM card treatment is time inconsistency. As illustrated by Table 1, women were more likely than men to make choices between amounts of money in a manner consistent with hyperbolic discounting. If ATM cards increase the temptation to withdraw funds, sophisticated individuals may respond negatively to the ATM card treatment. (In contrast, the impact for naive individuals is ambiguous – they may initially use the account intensively, only to withdraw more than they like). Table A2 investigates heterogeneous treatment effects with respect to time inconsistency by running the specification described by equation 2.2, augmented to include interactions between each account specific free ATM treatment indicator and separate dummy variables indicating that husband and wife gave answers to baseline time preference questions in a way consistent with hyperbolic discounting. Overall, there is little evidence that time inconsistency matters for heterogeneous treatment effects – though point estimates are sometimes large and negative for individual accounts, they are positive for joint accounts and imprecisely estimated on all accounts.<sup>22</sup>

---

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 (in the commercial center of town), because the withdrawal fee was seen as a commitment device not to spend the money frivolously, or because the individuals intended to use their new accounts for saving anyway.

<sup>22</sup>This may be due in part to our choice to use a front end delay method in estimating discount factors, particularly if the relevant tradeoff for most hyperbolic discounters is “today” versus the future.

## 2.B Appendix: Tables and Figures

Table 2-1: Demographic Characteristics of Study Sample

	Husbands	Wives	Difference	N
Age	44.0 [14.1]	36.9 [12.1]	7.10*** (0.678)	1496
Education	7.88 [3.70]	5.82 [3.99]	2.06*** (0.199)	1489
Literate	0.845 [0.362]	0.660 [0.474]	0.184*** (0.0218)	1496
Polygamous	0.192 [0.395]	0.209 [0.407]	-0.0162 (0.0208)	1486
Number Children	5.82 [4.13]	4.58 [2.47]	1.24*** (0.176)	1493
Subsistence Farmer or No Job	0.415 [0.493]	0.462 [0.499]	-0.0477* (0.0257)	1491
Income Last Week	1662 [5474]	815 [1781]	847*** (213)	1452
Owns Mobile Phone	0.464 [0.499]	0.413 [0.493]	0.0512** (0.0257)	1490
Participates in ROSCA	0.487 [0.500]	0.664 [0.473]	-0.178*** (0.0252)	1496
Has Bank Account	0.318 [0.466]	0.120 [0.326]	0.198*** (0.0208)	1496
Savings in Bank Account (Among Savers)	10853 [17994]	5967 [14629]	4886** (2134)	271
Has SACCO Account	0.0668 [0.250]	0.0121 [0.109]	0.0547*** (0.00998)	1492
Savings in SACCO Account (Among Savers)	54706 [53736]	44444 [64293]	10261 (22039)	56
Saves at Home	0.845 [0.362]	0.896 [0.306]	-0.0509*** (0.0174)	1494
Savings at Home (Among Savers)	1344 [2993]	887 [2761]	457*** (162)	1266
Saves on Mobile Phone	0.304 [0.460]	0.142 [0.350]	0.162*** (0.0231)	1251
Mobile Phone Savings (Among Savers)	581 [1670]	557 [1286]	23.9 (187)	266
Savings - I Decide	0.367 [0.482]	0.487 [0.500]	-0.119*** (0.0255)	1488
Savings - Spouse Decides	0.430 [0.495]	0.275 [0.447]	0.155*** (0.0245)	1488
Savings - Decide Together	0.101 [0.301]	0.0957 [0.294]	0.00485 (0.0154)	1488
Savings - Decide Alone	0.0791 [0.270]	0.120 [0.325]	-0.0409*** (0.0155)	1488
Impatient Now-Patient Later	0.149 [0.356]	0.191 [0.394]	-0.0423** (0.0196)	1475
Patient Now-Impatient Later	0.167 [0.373]	0.198 [0.399]	-0.0314 (0.0201)	1475
Distance from Family Bank (Miles)	3.73 [2.19]	3.73 [2.19]	- -	1496

Notes: Standard deviations in brackets, robust standard errors in parentheses. \*\*\*, \*\*, and \* indicate significance at the 99, 95, and 90 percent levels respectively.

Table 2-2: Randomization Verification

	Free ATM Card			Cash Prize	Maximum Interest Rate
	Husband	Wife	Joint		
Panel A. Adherence to Theoretical Probabilities					
Free ATM/Cash Prize/2 Percent	0.190	0.217	0.214	0.183	0.0749
	{0.115}	{0.636}	{0.309}	{0.106}	{0.428}
6 Percent					0.307
					{0.355}
10 Percent					0.618
					{0.678}
Panel B. Correlation with Demographic Characteristics					
Age	-1.02	-0.676	0.993	0.676	0.576
	(1.99)	(1.57)	(1.41)	(0.902)	{0.562}
Education	0.164	-0.111	0.294	-0.0507	1.59
	(0.548)	(0.546)	(0.371)	(0.256)	{0.205}
Literate	0.0558	-0.135**	0.0335	-0.000500	0.912
	(0.0502)	(0.0650)	(0.0371)	(0.0289)	{0.402}
Polygamous	0.0249	-0.0148	-0.0276	0.0232	0.407
	(0.0661)	(0.0590)	(0.0396)	(0.0276)	{0.666}
Number Children	0.126	-0.0587	-0.0325	-0.0618	0.326
	(0.703)	(0.332)	(0.322)	(0.211)	{0.722}
Subsistence Farmer/Unemployed	-0.133*	-0.111*	0.00809	-0.0683**	1.89
	(0.0680)	(0.0642)	(0.0432)	(0.0324)	{0.152}
Income Last Week	177	-187	27.8	-367**	0.00225
	(557)	(180)	(427)	(168)	{0.998}
Owns Mobile Phone	0.0866	0.110	-0.0334	0.0289	3.22**
	(0.0764)	(0.0678)	(0.0460)	(0.0327)	{0.0403}
Participates in ROSCA	-0.112	-0.0248	-0.0244	0.0143	1.06
	(0.0755)	(0.0653)	(0.0418)	(0.0330)	{0.346}
Has Bank Account	-0.0126	-0.0287	-0.0160	0.00437	1.89
	(0.0737)	(0.0498)	(0.0333)	(0.0278)	{0.151}
Has SACCO Account	0.0298	0.0134	-0.0166	-0.00370	0.813
	(0.0440)	(0.0225)	(0.0142)	(0.0127)	{0.444}
Saves at Home	-0.0427	0.0460	0.0185	0.0370*	0.666
	(0.0613)	(0.0397)	(0.0238)	(0.0203)	{0.514}
Saves on Mobile Phone	0.0174	-0.113**	-0.00823	-0.0303	4.15**
	(0.0813)	(0.0491)	(0.0343)	(0.0290)	{0.0162}
Savings - I Decide	0.0952	0.000876	-0.0366	0.0157	0.783
	(0.0756)	(0.0700)	(0.0381)	(0.0325)	{0.457}
Savings - Spouse Decides	-0.0505	0.000620	0.0306	-0.0204	0.124
	(0.0746)	(0.0597)	(0.0364)	(0.0318)	{0.883}
Savings - Decide Together	-0.0375	-0.0381	-0.00246	-0.0209	1.30
	(0.0371)	(0.0310)	(0.0263)	(0.0195)	{0.273}
Savings - Decide Alone	-0.00370	0.0139	0.0222	0.0184	0.531
	(0.0447)	(0.0538)	(0.0239)	(0.0209)	{0.588}
Impatient Now-Patient Later	-0.0312	0.0908	-0.0311	-0.0197	0.990
	(0.0498)	(0.0590)	(0.0311)	(0.0239)	{0.372}
Patient Now-Impatient Later	-0.0565	-0.0425	0.00875	0.0324	0.121
	(0.0508)	(0.0514)	(0.0323)	(0.0264)	{0.886}
Distance from Family Bank (Miles)	-0.622**	-0.00445	-0.0553	0.00389	1.51
	(0.290)	(0.296)	(0.250)	(0.137)	{0.221}
Interest Rate	0.771	-0.464	0.284	0.221	
	(0.600)	(0.549)	(0.354)	(0.228)	
Cash Prize - Husband	0.0276	-0.0496	0.0299		1.35
	(0.0594)	(0.0494)	(0.0440)		{0.259}
Cash Prize - Wife	-0.0194	-0.101**	0.0419		3.35**
	(0.0604)	(0.0497)	(0.0440)		{0.0357}

Notes: P-values in braces, robust standard errors (clustered at the couple level in columns 3-5) in parentheses. The first four columns present regression coefficients and standard errors on treatment dummies. The last column presents the test statistic and p-value of an F-test that demographics are equal across all treatments. For the cash prize, the interest rate is the individual interest rate when open, and the joint interest rate otherwise. \*\*\*, \*\*, and \* indicate significance at the 99, 95, and 90 percent levels respectively.

Table 2-3: Account Use Summary Statistics

	Joint	Husband	Wife	Husband-Joint	Wife-Joint	Husband-Wife	All
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A. All Couples							
Opened	0.659 [0.474]	0.463 [0.499]	0.446 [0.497]	-0.197*** (0.0353)	-0.213*** (0.0351)	0.0167 (0.0194)	1.50 [0.596]
Saved	0.404 [0.491]	0.260 [0.440]	0.308 [0.463]	-0.144*** (0.0390)	-0.0955*** (0.0409)	-0.0481 (0.0415)	0.448 [0.498]
Average Balance	338 [1364]	341 [1352]	236 [724]	2.91 (118)	-102 (86.1)	105 (109)	569 [2871]
Number Deposits	1.96 [2.34]	1.75 [2.06]	1.65 [1.57]	-0.214 (0.187)	-0.312* (0.162)	0.0985 (0.171)	1.24 [2.70]
Number Withdrawals	0.352 [1.43]	0.530 [2.30]	0.308 [1.12]	0.178 (0.178)	-0.0441 (0.107)	0.222 (0.163)	0.698 [2.61]
Fees	25.7 [110]	32.8 [134]	19.0 [73.6]	7.08 (11.0)	-6.71 (7.61)	13.8 (9.82)	43.0 [155]
N (Open Accounts)	389	200	198	589	587	398	748
B. Couples Without Cash Prizes (N=495)							
Opened	0.665 [0.473]	0.472 [0.500]	0.455 [0.499]	-0.193*** (0.0434)	-0.210*** (0.0432)	0.0172 (0.0253)	1.51 [0.603]
Saved	0.243 [0.430]	0.163 [0.371]	0.144 [0.353]	-0.0805** (0.0402)	-0.0994*** (0.0408)	0.0190 (0.0413)	0.271 [0.445]
Average Balance	269 [1158]	223 [1271]	208 [758]	-46.0 (130)	-61.8 (97.9)	15.8 (128)	409 [1532]
Number Deposits	1.80 [2.57]	1.48 [1.51]	1.49 [1.67]	-0.320 (0.203)	-0.313 (0.217)	-0.00652 (0.178)	1.01 [2.75]
Number Withdrawals	0.356 [1.57]	0.348 [1.37]	0.232 [0.824]	-0.00766 (0.152)	-0.124 (0.121)	0.116 (0.134)	0.655 [2.60]
Fees	26.1 [120]	23.4 [90.8]	14.9 [51.9]	-2.65 (10.7)	-11.1 (8.68)	8.49 (8.80)	41.2 [153]
N (Open Accounts)	267	135	125	402	392	260	495

Notes: Robust standard errors in parentheses, standard deviations in brackets. Columns 1-6 limited to accounts with 2 percent interest or better and no free ATM card. All outcomes (except for account opening) are presented conditional on account opening. Final column reports couples' aggregate activity across all three accounts.



Table 2-4: Quantiles of Account Use – Couples Who Saved

	Average			
	Balance	Deposits	Withdrawals	Fees
A. All Couples				
Minimum	105	1.00	0	0
1st Percentile	110	1.00	0	0
25th Percentile	320	2.00	0	0
50th Percentile	463	3.00	0	0
75th Percentile	1098	5.00	1.00	92.0
90th Percentile	2968	8.00	5.00	310
95th Percentile	4871	11.0	8.00	534
99th Percentile	14439	21.0	20.0	1188
Maximum	66410	30.0	31.0	1732
Mean	1421	4.24	1.56	96.1
N	335	335	335	335
B. Couples Without Cash Prizes				
Minimum	108	2.00	0	0
1st Percentile	113	2.00	0	0
25th Percentile	327	3.00	0	0
50th Percentile	844	4.00	1.00	60.0
75th Percentile	1582	6.00	2.00	162
90th Percentile	4053	11.0	8.00	496
95th Percentile	8143	14.0	10.0	664
99th Percentile	14439	24.0	21.0	1345
Maximum	15119	30.0	30.0	1404
Mean	1663	5.23	2.42	152
N	134	134	134	134

Table 2-5: Impact of Free ATM Cards on Account Use

	Has ATM Card	Mean Effect	Saved	Average Balance	Number Deposits	Number Withdrawals	Fees
Panel A. Joint Accounts							
Free ATM	0.860*** (0.0191)	0.305** (0.153)	0.0737 (0.0513)	0.301** (0.143)	0.311 (0.323)	0.715* (0.407)	32.2 (21.4)
DV Mean (No ATM, No Cash Prize)	0.0769	0.0920	0.246	5.03	2.82	0.365	26.8
N	485	485	485	485	485	485	485
Panel B. Husbands' Accounts							
Free ATM	0.884*** (0.0219)	0.266* (0.157)	0.0853 (0.0546)	0.126 (0.126)	0.440 (0.311)	0.651 (0.422)	31.6 (23.1)
DV Mean (No ATM, No Cash Prize)	0.0760	-0.0779	0.152	4.85	2.43	0.292	19.4
N	319	319	319	319	319	319	319
Panel C. Wives' Accounts							
Free ATM	0.934*** (0.0173)	-0.0672 (0.0960)	-0.0393 (0.0508)	-0.0589 (0.120)	-0.219 (0.149)	0.00297 (0.138)	-0.884 (8.39)
DV Mean (No ATM, No Cash Prize)	0.0446	-0.0674	0.153	4.92	2.49	0.204	13.1
N	309	309	309	309	309	309	309
Panel D. All Accounts							
Free ATM	0.887*** (0.0115)	0.191** (0.0849)	0.0453 (0.0312)	0.153* (0.0808)	0.202 (0.171)	0.501** (0.216)	22.9** (11.6)
DV Mean (No ATM, No Cash Prize)	0.0680	0.000	0.194	4.95	2.62	0.301	21.0
N	1113	1113	1113	1113	1113	1113	1113
Panel E. All Accounts - Is Impact for Wives Different?							
Free ATM	0.871*** (0.0144)	0.288*** (0.112)	0.0768** (0.0378)	0.230** (0.0998)	0.358 (0.232)	0.691** (0.299)	32.0** (16.0)
Free ATM×Wife	0.0594*** (0.0214)	-0.344*** (0.147)	-0.112* (0.0614)	-0.274* (0.150)	-0.555** (0.275)	-0.680** (0.338)	-32.5* (18.4)
DV Mean (No ATM, No Cash Prize)	0.0680	0.000	0.194	4.95	2.62	0.301	21.0
N	1113	1113	1113	1113	1113	1113	1113

Notes: Robust standard errors (clustered at the couple level when relevant) in parentheses. All regressions include dummy variables for the first 6 experimental sessions, cash prize receipt for each spouse, and account type and cash prize×account type interactions when relevant. \*\*\*, \*\*, and \* indicate significance at the 99, 95, and 90 percent levels respectively.

Table 2-6: Impact of Free ATM Card Provision on Total Savings by Couples

	Mean Effect	Saved (Any Account)	Average Balance	Number Deposits	Number Withdrawals	Fees
<b>A. Impact of ATM Card by Type</b>						
Joint ATM Card	0.194 (0.131)	0.0740 (0.0494)	0.233* (0.135)	0.0810 (0.338)	0.546 (0.417)	20.9 (22.3)
Husband's ATM Card	0.262* (0.158)	0.137*** (0.0574)	0.0141 (0.111)	0.951* (0.529)	0.527 (0.521)	26.8 (27.9)
Wife's ATM Card	-0.127 (0.122)	-0.0479 (0.0557)	-0.151 (0.117)	-0.0552 (0.418)	-0.363 (0.294)	-22.6 (16.8)
DV Mean (No ATM, No Cash Prize)	0.000	0.255	5.33	2.27	0.440	30.6
<b>B. Pooled Impact of ATM Cards</b>						
Any ATM Card	0.152* (0.0863)	0.0716** (0.0359)	0.0944 (0.0873)	0.261 (0.241)	0.386 (0.253)	16.1 (14.4)
DV Mean (No ATM, No Cash Prize)	0.000	0.255	5.33	2.27	0.440	30.6
<b>C. Impact by Card Type - Is Impact for Wives Different?</b>						
Joint or Husband's ATM Card	0.239*** (0.0987)	0.101*** (0.0387)	0.161* (0.0964)	0.464* (0.282)	0.597** (0.304)	26.0 (17.0)
Wife's ATM Card	-0.126 (0.124)	-0.0449 (0.0556)	-0.166 (0.118)	-0.00888 (0.435)	-0.371 (0.289)	-22.6 (16.6)
F Test - Joint/Husband=Wife	5.13** {0.0238}	4.46** {0.0351}	4.27** {0.0391}	1.01 {0.316}	4.30** {0.0384}	3.43* {0.0643}
DV Mean (No ATM, No Cash Prize)	0.000	0.255	5.33	2.27	0.440	30.6
<b>D. Impact of Interest Rates</b>						
Max Interest is 6 Percent	0.160 (0.109)	0.0479 (0.0604)	0.274*** (0.110)	0.276 (0.302)	0.221 (0.248)	13.5 (13.6)
Max Interest is 10 Percent	0.232** (0.100)	0.0589 (0.0567)	0.316*** (0.0991)	0.578** (0.288)	0.378* (0.202)	25.2** (12.2)
DV Mean (2 Percent Interest, No Cash Prize)	-0.132	0.190	5.06	2.12	0.357	18.5
N	748	748	748	748	748	748

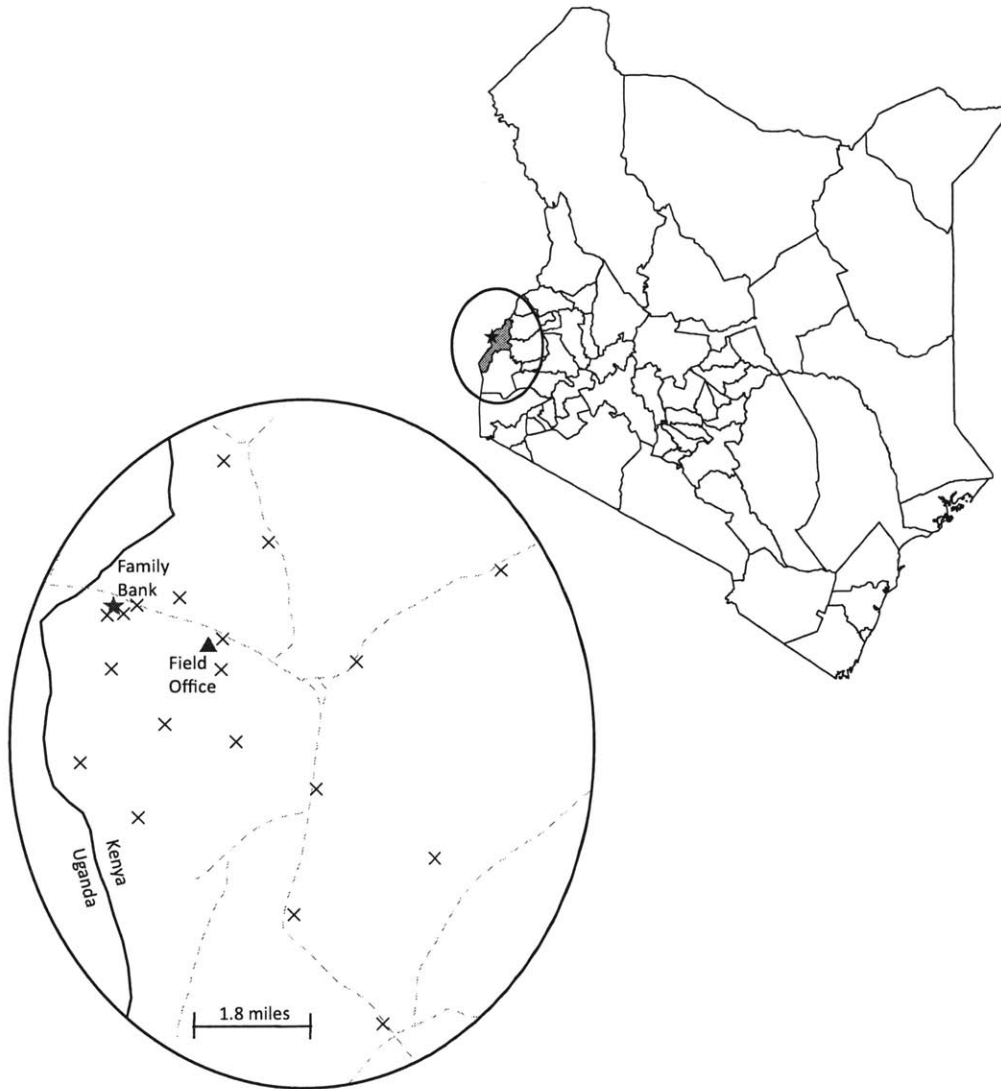
Notes: Robust standard errors in parentheses, p-values in braces. Additional controls (all panels) include cash prize dummies for both the husband and and a dummy for the first 6 experimental sessions. The first three panels also include a set of dummy variables that saturate possible combinations of open accounts. \*\*\*, \*\*, and \* indicate significance at the 99, 95, and 90 percent levels respectively.

Table 2-7: Impact of ATM Cards Interacted with Bargaining Power in the Household

Panel A. Husbands' Accounts			
Free ATM	0.569**	0.120	-0.749
	(0.266)	(0.220)	(0.794)
Free ATM×Wife Advantaged	-0.618*	-0.584**	-0.508
	(0.322)	(0.286)	(0.326)
Wife Advantaged	0.0922	0.0881	0.0763
	(0.118)	(0.118)	(0.127)
DV Mean (No ATM, No Cash Prize)	-0.0779	-0.0779	-0.0779
N	319	319	319
Panel B. Wives' Accounts			
Free ATM	-0.279***	-0.245***	-0.633*
	(0.113)	(0.105)	(0.337)
Free ATM×Wife Advantaged	0.350*	0.364**	0.420***
	(0.181)	(0.172)	(0.175)
Wife Advantaged	-0.229**	-0.223**	-0.228**
	(0.101)	(0.0991)	(0.102)
DV Mean (No ATM, No Cash Prize)	-0.0674	-0.0674	-0.0674
N	309	309	309
Panel C. Couple Level Impact			
Husband's ATM	0.407*	0.0492	-0.645
	(0.243)	(0.204)	(0.786)
Husband's ATM×Wife Advantaged	-0.295	-0.338	-0.530
	(0.334)	(0.289)	(0.333)
Wife's ATM	-0.343**	-0.308**	-0.686
	(0.159)	(0.148)	(0.538)
Wife's ATM×Wife Advantaged	0.412	0.461**	0.495**
	(0.258)	(0.231)	(0.234)
Wife Advantaged	-0.0816	-0.0678	-0.0623
	(0.133)	(0.129)	(0.144)
DV Mean (No ATM, No Cash Prize)	0.000	0.000	0.000
N	354	354	354
Baseline Account Ownership Controls?	No	Formal	All

Notes: Robust standard errors in parentheses. Baseline control sets for panels A and B match those in Table 5. Baseline control sets for panel C include those for couple-level ATM regressions in Table 6. "Formal" controls add a dummy variable for formal account ownership and its interaction with the relevant ATM treatments. "All" controls add additional dummy variables for mobile phone savings use, ROSCA membership, and a home savings dummy, as well as the relevant interactions with the ATM treatments. \*\*\*, \*\*, and \* indicate significance at the 99, 95, and 90 percent levels respectively.

Figure 2-1: Map of Study Locations



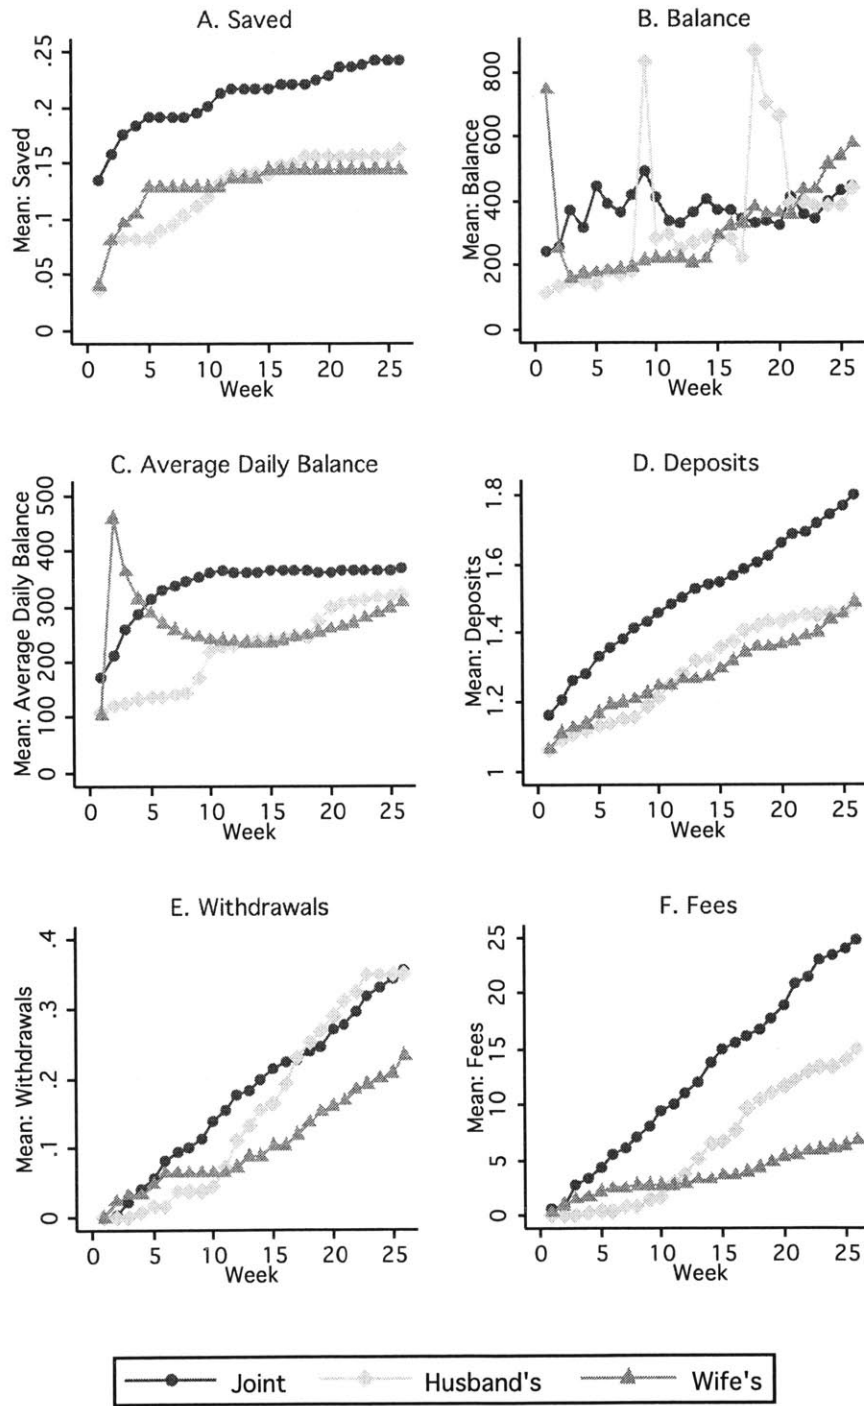
Notes: The shaded area on the map of Kenya indicates Busia district, with Busia town marked by a star. The local map shows detail surrounding Busia town, including major roads (dashed lines), boundaries (solid lines) and study locations. Group meeting locations are demarcated by an x.

Figure 2-2: Interest Rate Design

$R_J=2$					$R_J=6$					$R_J=10$				
	$R_M=0$	$R_M=2$	$R_M=6$	$R_M=10$		$R_M=0$	$R_M=2$	$R_M=6$	$R_M=10$		$R_M=0$	$R_M=2$	$R_M=6$	$R_M=10$
$R_F=0$	2	2	6	10	$R_F=0$	6	6	6	10	$R_F=0$	10	10	10	10
$R_F=2$	2	2	6	10	$R_F=2$	6	6	6	10	$R_F=2$	10	10	10	10
$R_F=6$	6	6	6	10	$R_F=6$	6	6	6	10	$R_F=6$	10	10	10	10
$R_F=10$	10	10	10	10	$R_F=10$	10	10	10	10	$R_F=10$	10	10	10	10

Notes: The maximum interest rate available to the couple is illustrated in interior cells.

Figure 2-3: Account Use Over Time - Couples Without Cash Prizes



Notes: Figures limited to opened accounts with 2 percent interest or better that were not selected for a free ATM card.

Figure 2-4: CDFs of Standardized Aggregate Account Use by Opened Account Type and ATM Treatment

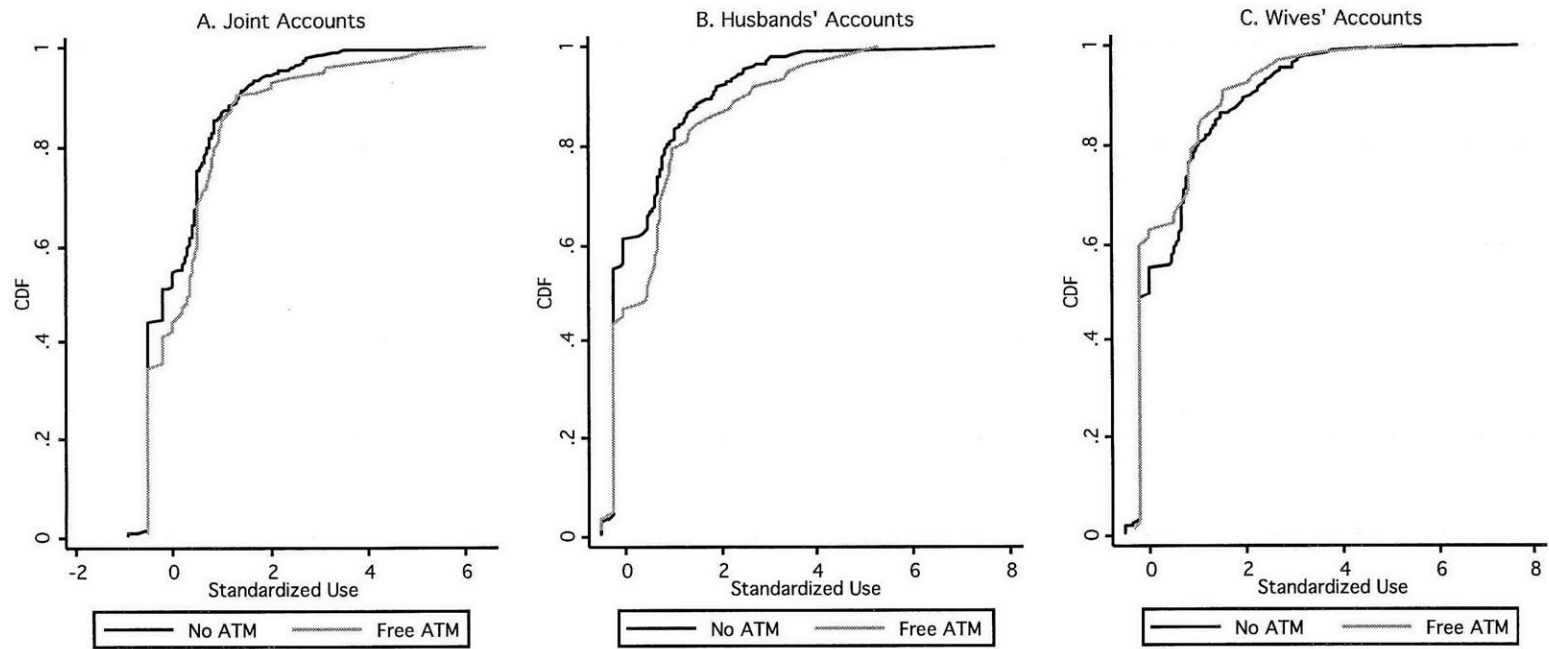
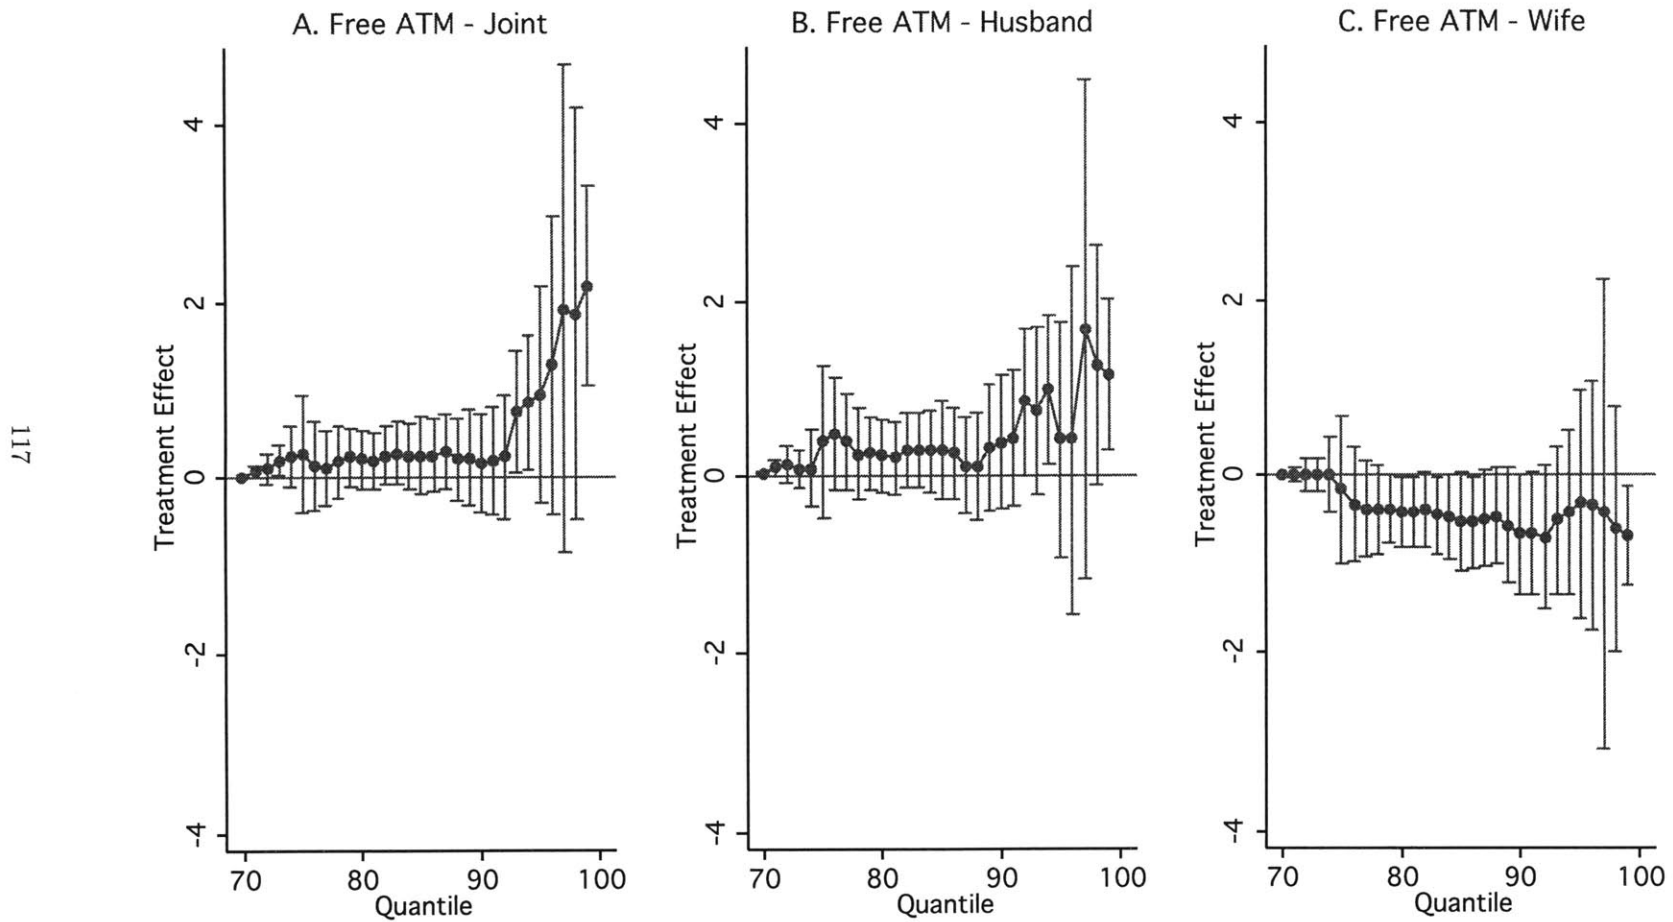




Figure 2-5: Quantile Treatment Effects of ATM Cards and Interest on Aggregate Standardized Account Use



Notes: Whiskers indicate 95 percent confidence intervals of quantile regression coefficients.

Figure 2-6: Distribution of Husband's Relative Bargaining Power

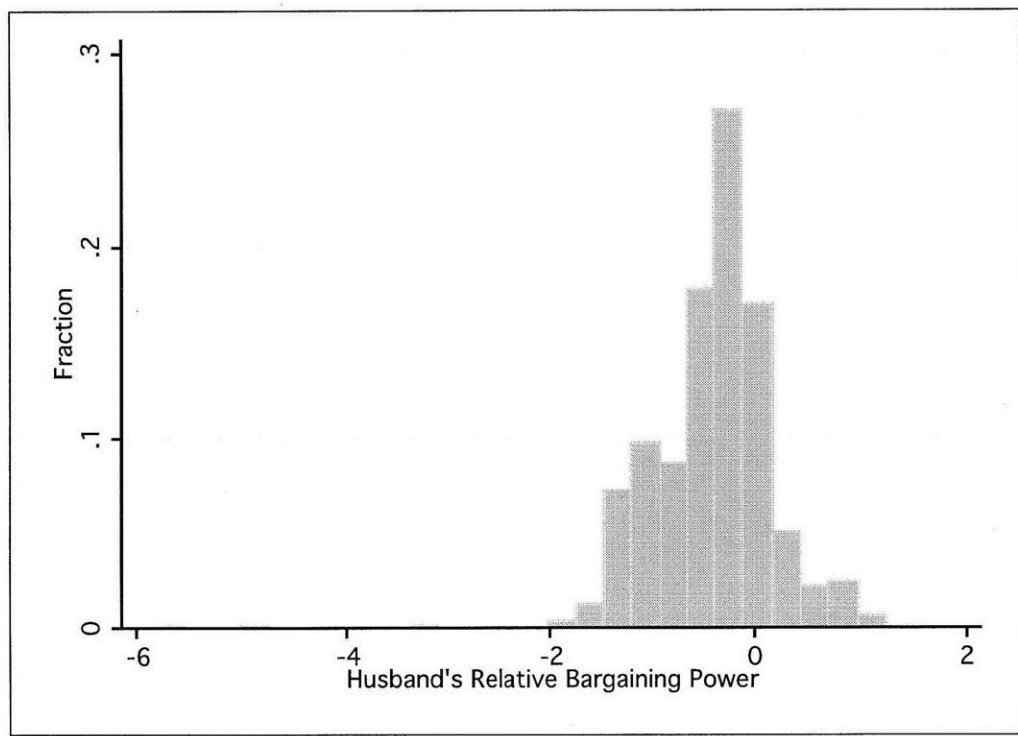


Figure 2-7: Baseline Savings Device Use and Proxied Bargaining Power

119

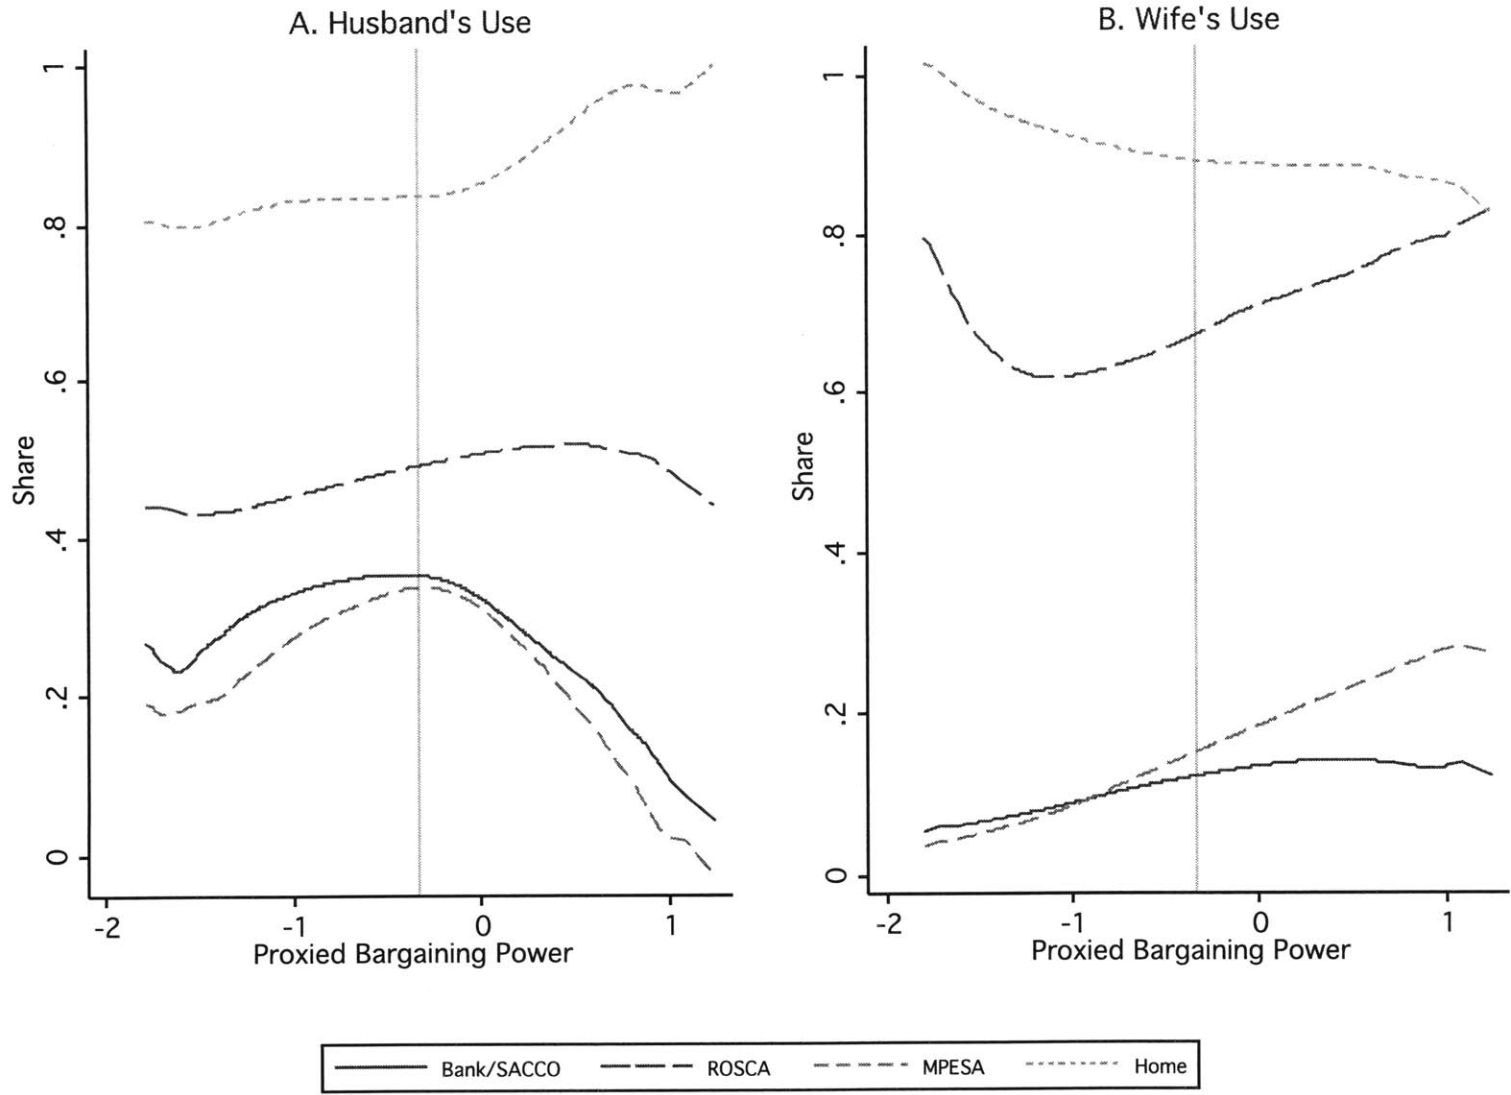


Table 2-A1: What Predicts Savings Account Use? (Couples Without Cash Prizes)

	Standardized Use		Saved		Average Daily Balance	
	Coefficient	Std. Error	Coefficient	Std. Error	Coefficient	Std. Error
Age (H)	0.00239	(0.00476)	0.00287	(0.00279)	-0.000559	(0.00570)
Age (W)	-0.000418	(0.00579)	0.0000649	(0.00323)	-0.0000938	(0.00727)
Education (H)	-0.0127	(0.0125)	-0.00606	(0.00663)	-0.00870	(0.0152)
Education (W)	0.0317***	(0.0126)	0.0121*	(0.00654)	0.0295*	(0.0152)
Children (H)	0.0278	(0.0185)	0.0152	(0.00971)	0.0281	(0.0226)
Children (W)	0.0127	(0.0248)	-0.000941	(0.0125)	0.0239	(0.0325)
Subsistence Farmer/Unemployed (H)	-0.107	(0.0869)	-0.0297	(0.0487)	-0.0894	(0.109)
Subsistence Farmer/Unemployed (W)	-0.0909	(0.0813)	-0.0177	(0.0462)	-0.0865	(0.100)
Income/1,000 (H)	0.00498	(0.00650)	0.00649	(0.00429)	0.00230	(0.00784)
Income/1,000 (W)	-0.0498***	(0.0153)	-0.0270***	(0.00750)	-0.0488***	(0.0202)
Owns Mobile Phone (H)	-0.0503	(0.0752)	-0.0304	(0.0425)	-0.0939	(0.0931)
Owns Mobile Phone (W)	0.130*	(0.0768)	0.0951**	(0.0434)	0.168*	(0.0947)
Saves in ROSCA (H)	0.0233	(0.0745)	0.0212	(0.0408)	0.0251	(0.0917)
Saves in ROSCA (W)	0.0293	(0.0761)	0.0409	(0.0434)	0.0526	(0.0935)
Owns Bank Account (H)	-0.0290	(0.102)	0.0162	(0.0481)	0.00411	(0.120)
Owns Bank Account (W)	0.373**	(0.168)	0.186***	(0.0755)	0.372*	(0.196)
Owns SACCO Account (H)	0.534**	(0.237)	0.123	(0.0928)	0.550**	(0.259)
Owns SACCO Account (W)	0.297	(0.448)	0.0942	(0.231)	0.624	(0.717)
Saves at Home (H)	0.0121	(0.106)	-0.0584	(0.0581)	-0.0511	(0.142)
Saves at Home (W)	0.0601	(0.142)	0.0610	(0.0689)	0.107	(0.153)
Saves on Mobile Phone (H)	0.143	(0.104)	0.0865	(0.0555)	0.175	(0.134)
Saves on Mobile Phone (W)	0.0986	(0.139)	0.0334	(0.0701)	0.265	(0.184)
Hyperbolic Discounter (H)	-0.0289	(0.0947)	0.0324	(0.0598)	0.0136	(0.131)
Hyperbolic Discounter (W)	-0.0993	(0.0934)	0.0173	(0.0540)	-0.151	(0.118)
Anti-Hyperbolic Discounter (H)	0.152	(0.119)	0.0810	(0.0592)	0.130	(0.139)
Anti-Hyperbolic Discounter (W)	-0.0936	(0.0811)	-0.0225	(0.0471)	-0.129	(0.107)
Discount Factor (H)	0.000132	(0.158)	0.0199	(0.0838)	0.132	(0.185)
Discount Factor (W)	0.00391	(0.152)	-0.0826	(0.0889)	0.175	(0.185)
Polygamous	-0.0775	(0.124)	-0.0521	(0.0611)	-0.0652	(0.141)
Distance from Bank	-0.0269	(0.0196)	-0.0115	(0.00938)	-0.0305	(0.0230)
DV Mean / R2	0.000	.154	0.271	.127	5.41	.146

Notes: Robust standard errors in parentheses. Time preference controls are defined in the Appendix. \*\*\*, \*\*, and \* indicate significance at the 99, 95, and 90 percent levels respectively.

Table 2-A2: Treatment Effects by Time Inconsistency

	Mean Effect	Saved	Average Balance	Number Deposits	Number Withdrawals	Fees
Joint ATM	0.177 (0.159)	0.0351 (0.0553)	0.199 (0.158)	0.0290 (0.410)	0.671 (0.518)	28.0 (27.7)
Joint ATM×Hyperbolic (H)	-0.0430 (0.236)	0.181 (0.144)	0.0429 (0.321)	-0.275 (0.571)	-0.837* (0.464)	-44.2 (26.9)
Joint ATM×Hyperbolic (W)	0.106 (0.243)	0.0794 (0.133)	0.183 (0.335)	0.445 (0.640)	-0.210 (0.533)	-15.4 (27.5)
Husband's ATM	0.312 (0.200)	0.126* (0.0677)	0.00924 (0.135)	1.28* (0.668)	0.676 (0.664)	33.1 (35.7)
Husband's ATM×Hyperbolic (H)	-0.253 (0.292)	-0.0875 (0.181)	0.178 (0.276)	-0.939 (0.720)	-0.977 (0.656)	-30.9 (45.2)
Husband's ATM×Hyperbolic (W)	-0.113 (0.250)	0.140 (0.148)	-0.0446 (0.213)	-1.33** (0.662)	-0.273 (0.824)	-17.0 (40.9)
Wife's ATM	-0.0983 (0.169)	-0.0627 (0.0635)	-0.0238 (0.162)	-0.00402 (0.612)	-0.364 (0.379)	-20.1 (22.4)
Wife's ATM×Hyperbolic (H)	-0.231 (0.262)	-0.0156 (0.177)	-0.505** (0.248)	-0.725 (0.767)	-0.121 (0.451)	-21.0 (30.2)
Wife's ATM×Hyperbolic (W)	0.0755 (0.234)	0.0942 (0.134)	-0.269 (0.241)	0.358 (0.772)	0.340 (0.574)	14.5 (30.7)
Hyperbolic (H)	-0.00724 (0.103)	-0.0291 (0.0553)	0.0644 (0.125)	0.0418 (0.287)	-0.0735 (0.193)	-3.27 (14.8)
Hyperbolic (W)	-0.103 (0.0895)	-0.0298 (0.0492)	-0.0786 (0.102)	-0.302 (0.239)	-0.222 (0.180)	-15.0 (11.5)
DV Mean	9.91e-09	0.255	5.33	2.27	0.440	30.6
N	748	748	748	748	748	748

Notes: Robust standard errors (clustered at the couple level when relevant) in parentheses. \*\*\*, \*\*, and \* indicate significance at the 99, 95, and 90 percent levels respectively.



## Chapter 3

# Price Subsidies, Diagnostic Tests, and Targeting of Malaria Treatment: Evidence from a Randomized Controlled Trial<sup>\*</sup>

### 3.1 Introduction

Limiting the spread of infectious disease has positive social benefits – as such, infectious disease programs often feature subsidies for prevention and treatment technologies. Financing such subsidies is obviously subject to a budget constraint, however, and it is therefore critical to ensure that subsidy dollars are spent where they have the highest return. For products that have heterogeneous returns, the introduction of a subsidy creates a tradeoff between access and targeting. That is, subsidies for the product are likely to increase demand among both appropriate users, for whom the returns are indeed high, and among inappropriate users, for whom the benefits are marginal. This is the “menu-setting problem” described by Olmstead and Zeckhauser (1999).

This tradeoff between affordability and over-consumption is magnified for products for which overuse has negative social spillovers. For example, the (ineffective but quite common) use of antibiotics to treat viral infections contributes to antibiotic resistance. Likewise, antimalarial treatment in the absence of malaria can contribute to antimalarial resistance. When people are uncertain about the cause of their ailment and the costs of under-treating can be deadly (e.g., untreated malaria is a major cause of childhood mortality in Africa), presumptive treatment is likely to be privately optimal if the treatment is subsidized and thus affordable (provided side-effects are minimal). This makes the menu-setting problem even more pressing: the trade-off is not just between affordability and cost-ineffective consumption at a single point in time, but a trade-off between affordability

---

<sup>\*</sup>This chapter is co-authored with Jessica Cohen and Pascaline Dupas.

today and effectiveness in the future.

This paper studies the menu-setting problem for the latest class of antimalarials, artemisinin combination therapies (ACTs). Artemisinin-based therapies now constitute the only effective class of antimalarials in Africa, where drug resistance has rendered all earlier generations of antimalarials (quinine, chloroquine, amodiaquine, sulfadoxine-pyrimethamine) mostly ineffective. The use of artemisinin derivatives by themselves as monotherapies is highly discouraged by the World Health Organization, however, due to concerns that malaria parasites have already started developing resistance to artemisinin. Instead, the WHO encourages the use of ACTs, which combine an artemisinin derivative with a partner drug (such as mefloquine or lumefantrine), and thereby help protect the artemisinin derivatives from resistance.<sup>1</sup> The unsubsidized price of ACTs is prohibitive for the great majority of rural African households and as a result, in 2008, 6 years after ACTs were placed on the WHO's essential drugs list, fewer than 15 percent of African children with malaria were treated with ACTs (World Health Organization 2009). In response, a call was made for a global ACT subsidy to achieve two main goals: (1) immediately *save lives*, by increasing access to ACTs, and (2) *buy time*, by crowding-out monotherapies and thereby delaying resistance (Arrow et al. 2004). The Affordable Medicines Facility for malaria (AMFm) initiative, financed by major international aid agencies, was subsequently established to roll out a 95 percent subsidy to first line buyers of ACTs throughout Africa. At the time of writing, the subsidy was being piloted in 8 countries.

The AMFm subsidy was explicitly designed to impact the price of ACTs in the private sector, as many people seeking malaria treatment do so in loosely regulated, informal private-sector drug shops where they receive no formal diagnosis. In this context, it is quite likely that a 95 percent decrease in ACT prices will be associated with increases in not only appropriate but also inappropriate ACT use. A high rate of overtreatment with ACTs is problematic for several reasons. First, it is a waste of a vast amount of subsidy money. The co-payments alone for the AMFm are estimated to cost \$216 million in the pilot phase (Global Fund 2010). Second, if the retail-sector ACT subsidy draws malaria-negative people from health clinics to the drug shop (reducing the chances they receive diagnostic confirmation), it could delay or preclude proper treatment for the true cause of illness (Reyburn et al. 2004). Finally, a high rate of overtreatment for malaria may contribute to the selection of drug resistant parasites (Perkins and Bell 2008; White 2004). This means that, although ACT subsidies would have a first order (positive) effect on resistance because of artemisinin monotherapy crowd-out, there could be a second order negative effect of accelerating resistance from overtreatment with ACTs.

We use data from a randomized controlled trial conducted with over 2,900 households in rural Kenya to study the tradeoffs between ACT affordability and overuse in the context of the AMFm subsidy. Our research design also tests an alternative to the AMFm subsidy regime that explicitly

---

<sup>1</sup>Combination therapies slow resistance because in order for a resistant parasite to arise, it must develop mutations that make it resistant to all drugs in the combinations. When the combined drugs have differing modes of action, the probability of this event occurring is substantially lower than the probability of resistance developing to any single drug alone (World Health Organization 2010a).



acknowledges the problem of overuse by providing access to a subsidized rapid diagnostic test for malaria (RDT) in tandem with subsidized ACTs.<sup>2</sup>

We show that subsidies for ACTs and RDTs can successfully broaden access to these technologies, and that including RDT subsidies under any ACT subsidy policy could be an effective way to improve the targeting of subsidized ACTs to people with confirmed malaria. We also show that this RDT subsidy could be financed by reducing the ACT subsidy somewhat, an approach that could be more cost-effective than the proposed AMFm subsidy alone. This is primarily due to two stark results from our experiment:

1. Over-diagnosis of malaria is extremely common in our study context, particularly among teenagers and adults. As a result, when ACTs are heavily subsidized, only 33 percent of ACT takers over age 13 actually have malaria. This implies that improving diagnostic access has the potential to considerably reduce over-treatment.
2. The demand for ACTs is highly price-elastic above a certain range (which includes the unsubsidized price of ACTs), but very inelastic at low prices, especially for children. Specifically, we see a modest 13 percent decline in ACT purchases at the drug shop when the subsidy level declines from 92 to 80 percent, corresponding to a 150 percent price increase. For children, who are much more likely to actually have malaria and for whom malaria is most dangerous, there is no significant price sensitivity in this range. This implies that some reduction in the ACT subsidy (compared to the current AMFm plan) is unlikely to meaningfully reduce access.

In order for the RDT subsidy to be cost-effective (relative to an ACT subsidy alone), it is critical for people to be both willing to take the test and compliant with the test result. We find that willingness to test is very high: when offered a voucher for subsidized RDTs, more than 80 percent of households who visited the drug shop chose to get the patient tested with an RDT prior to making their ACT purchase decision. This is despite the fact that only 15 percent of households had ever heard of RDTs prior to our experiment. Compliance with the test results is not as high, however. In our context, about 49 percent of patients over the age of 5 who tested negative went on to purchase the ACT. This is not surprising given that the status quo testing technology (a microscopic test offered at health centers) has a relatively high rate of false negatives when used by health workers in the field (31 percent according to a 2002 study in Kenya (Zurovac et al. 2006)), and health practitioners themselves tend to ignore test results and prescribe antimalarials to those who test negative. For example, a 2002 study in Kenya found that nearly 80 percent of patients who tested negative for malaria were prescribed antimalarials regardless (Zurovac et al. 2006). This had decreased somewhat by 2010, but remained as high as 50 percent among those above the age of 5,

---

<sup>2</sup>RDTs for malaria work similarly to rapid tests for HIV and do not require specialized equipment, such as a microscope, or electricity. A small sample of blood is collected through a finger prick and placed on a testing cassette. The blood sample is exposed to a buffer solution, and the presence of malaria antibodies can be determined within approximately 15 minutes.

despite the introduction of strict guidelines for health workers to test and adhere to test results for patients above 5 (Juma and Zurovac 2011). While RDTs have a much lower rate of false negatives than microscopy (generally under 5 percent in lab settings (World Health Organization 2010b) and around 8 percent when used by health workers (de Oliveira et al. 2009)), it might take some time for households to learn this.

Overall, in the absence of any information or marketing campaign on RDTs, our estimates suggest that moving from the AMFm subsidy level to an 80 percent ACT subsidy with RDTs could increase the share of ACT takers who are malaria positive at the drug shop by around 22 percentage points. The majority (18 percentage points) of this impact comes from selection induced by the higher ACT price. However, the total impact could be substantially increased if full adherence to RDT results were achieved. Overall, our results suggest that, in this context, taking some of the planned ACT subsidy money away from ACTs and putting it towards subsidizing and promoting RDTs could improve targeting and be particularly cost-effective among older children and adults if adherence to test results can be improved.

It is important to point out that this subsidy regime is a second-best strategy. The first-best would be to make the ACT subsidy conditional on having a positive malaria test result. This first-best is unlikely to be enforceable at a reasonable cost, however. Overuse of prescription-only drugs are common even in highly regulated health care markets such as the US and Europe, due to physician agency problems (McGuire 2000). Similar agencies issues are likely to be widespread in developing countries where monitoring of both private and public health care sectors is extremely limited (World Bank 2004).

While our results suggest that the combined ACT+RDT subsidy regime would be preferred to an ACT subsidy alone (especially in the long run), our results make it very clear that an ACT subsidy will considerably increase access among the needy. We proxy SES by whether a household's female head is illiterate (about 38 percent of our sample) and find a substantial access gap in the absence of a subsidy: literate-headed households are over three times more likely to treat an illness episode with an ACT. However, under the ACT subsidy regime literate-headed households are just 17 percent more likely to treat an illness with an ACT. The subsidy would also primarily benefit children, who are much more susceptible to malaria than adults. In fact, for children, the benefit of the RDT subsidy would not be in averting over-treatment, but rather in ensuring that in those cases where a child suspected of having malaria does not actually have malaria, the true illness is diagnosed and treated faster.

Beyond its immediate relevance to the proposed AMFm subsidy initiative, which will affect millions of households in rural Africa in both the short-run (affordability) and long-run (drug resistance), our paper contributes to the literature in three main ways. First, our paper adds to a fast-growing experimental literature on user fees for health services and health products whose appropriate use generates positive externalities. So far this literature has focused on optimal pricing

for preventative health products, such as water purification kits or bednets,<sup>3</sup> for which overuse is not a problem, and for which the objective of the social planner is to expand access while limiting *underuse* among subsidy beneficiaries. In contrast, this paper considers the price-setting problem that arises when overuse generates negative externalities (in our case, through drug resistance).

Second, we contribute to the literature on overdiagnosis and overtreatment, two major contributors to health care costs and a source of concern throughout the world (Welch, Schwartz, and Woloshin 2011). Finally, we contribute to the nascent literature on treatment-seeking behavior in resource-constrained environments, along with the earlier contributions of Leonard, Mliga, and Mariam (2002) in Tanzania, Banerjee, Deaton, and Duflo (2004) in Rajasthan (India), and of Leonard (2007, 2009) in Tanzania and Cameroun, respectively.

Our results generate a number of important questions for future research. The first obvious question is that of learning about the effectiveness of ACTs and the reliability of RDTs. Limiting overtreatment with ACTs is likely to improve inference about ACTs' effectiveness among the general population (Adhvaryu 2009). The timeline of our experiment was too short to study this question, however. In a companion paper (Cohen, Dupas, and Schaner 2011) we find that exposure to RDTs via neighbors increases demand for RDTs, although learning about the reliability of RDTs might take some time.

The second question is how to ensure optimal provider incentives. As discussed in Cohen and Dickens (2011), drug shops, which make a profit from selling ACTs whether their clients are truly malaria positive or not, might not have any incentive to sell a cheap diagnostic test that will result in fewer ACT purchases. The problem of RDT provision is an incentive problem similar to that of "informed experts" who sell both their diagnostic of a problem and the solution to the problem, such as surgeons or auto repair shops (Wolinsky 1993). One possibility for increasing RDT provision would be to decouple the supply of medication and diagnostic services. For example, RDTs could be made available at general stores, rather than drug shops, and be made simple enough to use for households to self-administer the test. Alternatively, drug shops might have an interest in building trust and ensuring that their clients' beliefs about the drug's effectiveness remain high. This dynamic incentive could be enough to ensure availability of RDTs at drug shops.

The remainder of the paper proceeds as follows: Section 3.2 describes some background facts on the malaria burden and treatment options in rural Africa, as well as the proposed AMFm subsidy. Section 3.3 develops a model of treatment-seeking behavior in this environment, and identifies the key trade-offs inherent in heavily subsidizing ACTs. Section 3.4 describes our experimental design and data. We discuss the results in Section 3.5 and performs a cost-effectiveness analysis in Section 3.6, before we conclude in Section 3.7.

---

<sup>3</sup>See Cohen and Dupas (2010), Dupas (2010), Hoffmann (2009), and Tarozzi et al. (2011) on bednets; and Ashraf, Berry, and Shapiro (2010) and Kremer et al. (2011) on water purification.

## 3.2 Background

### 3.2.1 Background on Malaria

Malaria is a disease caused not by “bad air”, as was once believed, but by a blood parasite called *Plasmodium*, which is transmitted from human to human by female anopheles mosquitoes. Malaria is estimated to cause 200 million illnesses and to kill close to one million people every year – the great majority of them in Africa, and the great majority of them under the age of five (World Health Organization 2009).

Despite major strides in malaria eradication in the early and mid-20th century, notably in the Americas (Bleakley 2010; Lucas 2010), efforts to eradicate malaria worldwide were abandoned in the 1970s. Recently, efforts to control malaria transmission have rejuvenated with the introduction of highly effective prevention tools, such as long-lasting insecticide treated bednets. Such nets have been distributed on a massive scale in the past five years, contributing to reductions in malaria incidence and deaths in some countries (Otten et al. 2009). The morbidity burden of malaria remains considerable, however, and there is no malaria vaccine on the near horizon. Given this, policy-makers and donors have recently been turning their attention to malaria treatment, an aspect of malaria control where less progress has been made.

Because immunity to malaria develops with repeated exposure, children under 5 are most vulnerable to acquiring and dying from malaria. How readily these children can access effective anti-malarials when they get infected is thus a very important determinant of overall malaria morbidity and mortality. According to the 2009 World Malaria Report, fewer than 15 percent of young children with presumed malaria in endemic countries were treated with effective antimalarial drugs. This crisis in access has been fueled by the spread of drug resistant malaria parasites. Early malaria control efforts relied heavily on chloroquine as a cheap, effective treatment. *Plasmodium falciparum*, the most common and deadly of the five strains of malaria, started becoming resistant to chloroquine in the 1960s and rendered the drug ineffective by the early 1990s, contributing to a substantial rebound in malaria mortality (Trape 2001). Subsequent innovations in antimalarial medicines have been successively less able to withstand parasite resistance (D’Alessandro and Buttiens 2001).

Currently, the only effective antimalarial against the *P. falciparum* parasite is artemisinin, a compound derived from Chinese wormwood trees that is significantly more expensive to produce than older, synthetic forms of malaria medicine. Artemisinin acts quickly to bring down the parasite load (often people feel significantly better within 24 hours) and has only mild side effects. The retail price of artemisinin-based antimalarials is roughly \$6-8 in Sub-Saharan Africa.<sup>4</sup> In most populations dealing with endemic malaria this cost of treatment is unaffordable.

---

<sup>4</sup>ACT Watch, Population Services International, Outlet Surveys (<http://www.actwatch.info>). The median price of Artemether Lumefantrine (the drug used in this study) in drug shops is \$5.26 in Uganda, \$6.03 in Benin, \$4.58 in DRC, \$5.36 in Nigeria and \$5.36 in Zambia. In most cases, other ACTs are \$1 more expensive, and all ACTs are more expensive in pharmacies than in drug shops.

As discussed in the introduction, the two major challenges in malaria control today are: (1) How can effective malaria medicines be made accessible and affordable? And (2) How can resistance to the only remaining effective treatment be forestalled?

### **3.2.2 The Affordable Medicines Facility for Malaria (AMFm)**

In response to the two challenges above, the Gates Foundation, UNICEF, the Global Fund and others have recently created a “global subsidy” policy called the Affordable Medicines Facility for Malaria (AMFm). Through a co-payment to ACT manufacturers, the program would reduce the price of ACTs by roughly 95 percent to first line buyers, such as governments, NGOs and wholesalers (Global Fund to Fight AIDS, TB and Malaria 2010). The final price to consumers in the private sector is unrestricted, but the aim is for retail sector ACTs to be cheap enough for most rural, poor populations to afford them and to be competitive with older antimalarials like chloroquine. The AMFm launched in early 2011 as a pilot in 8 countries (including Kenya). Our study was conceived and implemented in 2008/2009, when the AMFm was under consideration but had not yet started its pilot.

### **3.2.3 Health Providers and Health Treatment Seeking in Rural Kenya**

To provide context for our theoretical framework and experimental design, this subsection provides some descriptive evidence on health care choices that are available to households in rural Kenya and presents some key baseline facts about treatment-seeking behavior in our study population.

As in many developing countries, the range of health treatment providers in Kenya is vast. Public health facilities vary from small “dispensaries”, which provide very basic outpatient care for the most common types of illness, to hospitals with medical specialists and surgical capacity (Luoma 2010). As in other countries ((Leonard 2007; Das et al. 2008)), the level of practitioner training and expertise varies widely across Kenyan health facilities, even within the same tier of care. However, in rural areas, patients generally only have access to lower level facilities, as the distance and transport costs to higher level facilities in district centers and urban areas are often prohibitive. Lower level health facilities are typically staffed with nurses or medical assistants and are known to have high rates of absenteeism and stock outs of essential medicines (Kangwana et al. 2009; Chaudhury et al. 2006). Rural health facilities often do blood slide microscopy tests for malaria, though this depends on the availability of a trained lab technician and stocks of slides and reagents. Even when microscopy tests for malaria are performed, negative results are often disregarded and treatment is given anyway (Juma and Zurovac 2011; Zurovac et al. 2006).<sup>5</sup> ACTs are free in public health facilities in Kenya, but are quite often stocked out (Kangwana et al. 2009). Even with free medication, the direct and indirect costs of seeking treatment for malaria in the public sector can

---

<sup>5</sup>The reasons why negative malaria tests are so often ignored by medical practitioners in Africa is the subject of a growing body of public health research. Some explanations include historical presumptive treatment of malaria, risk aversion, lack of confidence in the test results, professional norms and patient demands (Chandler et al. 2008).

be high if fees are charged for consultation, diagnosis, etc. (as is often the case in our study area) and if it takes a long time to reach the facility and been seen by a medical professional.

Households also have the option of treating an illness with over-the-counter medication purchased at a drug shop.<sup>6</sup> The education levels and credentials of drug shop owners vary widely, but they are often asked by patients for treatment recommendations (Patouillard et al. 2010; Marsh et al. 2004). The two main benefits to treating an illness at a drug shop, rather than a public health facility, are convenience and choice. Drug shops are ubiquitous in Kenya, even in the most remote areas, whereas the average household in our sample lived more than 6.5 kilometers from the nearest health facility. These shops are also often open 7 days a week from early morning until late at night and offer a wide variety of medications to treat malaria (for example, many of the shops we visited during our pilot phase were open 12 hours a day, 6-7 days of the week). Private sector supply chains for pharmaceuticals are also more reliable than in the public sector, so drug shops are less often stocked out of medication than health facilities.<sup>7</sup> The main drawbacks to treating malaria at a drug shop rather than a public health facility is the lack of diagnostic capability, the risk of receiving lower quality or counterfeit drugs, and of course the absence of emergency medicines and equipment to treat severe malaria infections.

We conclude this section with some basic statistics on malaria treatment seeking behavior reported in our baseline survey. The results, presented in Table 2, are from self-reported behavior in response to presumed malaria episodes. (We will discuss in detail how the sample was formed and how the data was collected in Section 4). Overall presumed malaria incidence in Western Kenya is very high, with nearly 70 percent of households reporting an episode of malaria in the month before baseline. Yet rates of malaria diagnostic testing were relatively low, with 18 percent of households taking a microscopy test and 3 percent of households taking an RDT in the previous month. Most households either go to the health facility (41 percent) or to the drug shop (37 percent) to treat malaria, though a substantial minority (18 percent) does not seek care. The drug shop is the most common source of antimalarial medication, however – antimalarials are procured from a drug shop 52 percent of the time.

At baseline, households in our sample had limited access to effective antimalarials - only 21 percent of presumed malaria episodes were reported to be treated with ACTs. This rate of ACT taking is not much higher in children 13 and younger, even though young children are most at risk of severe morbidity and mortality from malaria. Roughly 35 percent of households take older, less effective drugs like amodiaquine and sulfadoxine pyrimethamine. Households spend \$1.68 per

---

<sup>6</sup>Another option is to seek care in the formal private sector (at a clinic or doctor's office, for example). The cost of this type of care is prohibitively high for our study population, however. At endline, just 4 percent of all illness episodes were treated in the formal private sector.

<sup>7</sup>Indeed it is often the case that health facilities will send patients to drug shops to purchase the medicines that they are stocked out of. We see some evidence of this in our baseline. When a presumed malaria episode was first treated at the health center, the household only reported obtaining drugs from the health center 82 percent of the time.

malaria episode (the median is substantially lower at \$0.77). This is a remarkably large sum of money given that the agricultural daily wage is around \$1.5 (Dupas and Robinson 2011). In sum, our baseline data suggests that households have frequent malaria episodes that are often treated at the drug shop with inappropriate medications, despite high out of pocket expenditures. Children are slightly more likely to go to a health center and be treated with ACTs. Literate households are more likely to go to health facilities and access ACTs than illiterate households.

### 3.3 Model

This section develops a model of treatment seeking behavior in the environment described above. The goal of the model is to highlight the trade-off inherent to subsidizing ACTs through the retail sector. The trade-off is embedded in the following two policy parameters of interest:

- The share of true malaria episodes that do not get treated with ACTs – we denote this as “ $UT$ ” for “under-treatment”.
- The share of non-malaria episodes that are treated with ACTs – we denote this as “ $OT$ ” for “over-treatment”.

The objective of the social planner is to decrease  $UT$  while limiting the increase in  $OT$ , since overtreatment has the negative social externalities we discussed earlier. In other words, the objective is to reduce the number of type II errors (false negatives) without increasing the number of type I errors (false positives) too much. The relative weight placed on these two parameters is captured by the social planner’s objective function (some  $F(UT, OT)$ ), which the planner maximizes subject to a budget constraint.

One simple objective function that prizes ACT targeting over ACT access is the fraction of ACT takers who are malaria positive, which we denote by  $T$  for “targeting”. Let  $\Pi$  represent the fraction of all illness episodes that are actually malaria. Then we can express  $T$  as follows:

$$T = \frac{(1 - UT) \Pi}{(1 - UT) \Pi + OT(1 - \Pi)}$$

In a first-best world, decreasing  $UT$  while keeping  $OT$  at a minimum could be easily achieved by simply making the ACT subsidy conditional: only those with a positive malaria test result would be allowed to buy an ACT at the subsidized price. This is the idea behind the free ACT distribution ongoing at health centers in our study context. It is clear from our baseline data that access to health centers is limited, however – hence the AMFm plan to roll out a subsidy through the retail sector. The goal of this section is to discuss how such a proposed retail sector subsidy will affect  $UT$ ,  $OT$ , and  $T$ .

### 3.3.1 Model Setup

We consider an environment where the household has three possible actions,  $a$ :

1. Buy ACTs at the drug shop:  $a = s$
2. Seek diagnosis at a formal health facility and receive ACTs if positive:  $a = h$
3. Purchase other drugs at the drug shop (e.g. antipyretics) or do nothing:  $a = n$

When a household gets an illness shock, the household observes the symptoms of the illness and subjectively assesses the probability  $\pi$  that the illness is actually malaria. The expected value of taking a particular action  $a \in \{s, h, n\}$  depends on this probability, and is denoted by  $V^a(\pi)$ . It can be decomposed into:

$$\begin{aligned} V^a(\pi) &= \pi(U_P^a(\pi) - p_P^a(\pi)) + (1 - \pi)(U_N^a(\pi) - p_N^a(\pi)) \\ &= \pi V_P^a(\pi) + (1 - \pi)V_N^a(\pi) \end{aligned}$$

where  $U_M^a(\pi)$  is the utility obtained from taking action  $a$  when the individual's true malaria status is  $M \in \{P, N\}$  (i.e., malaria positive or malaria negative) and  $p_M^a$  is the expected price paid for treatment when the individual's true malaria status is  $M$ .<sup>8</sup> Note that the utilities and prices may be a function of  $\pi$  – if, for example, the severity of symptoms is increasing as  $\pi$  increases, then individuals may expect to pay more to treat the illness, particularly when it is not actually malaria. An individual will seek ACT treatment at the drug shop if

$$V^s(\pi) = \pi(U_P^s(\pi) - p_P^s(\pi)) + (1 - \pi)(U_N^s(\pi) - p_N^s(\pi)) \geq \max\{V^h(\pi), V^n(\pi)\} \quad (3.1)$$

### 3.3.2 Impact of an ACT Subsidy at the Drug Shop

Consider a population of illness episodes. The share of episodes that are actually malaria,  $\Pi$ , is determined by the distribution  $g(\cdot)$  of actual malaria positivity in the population, where we denote true malaria positivity as  $\tilde{\pi}$ . We assume that individuals' subjective malaria assessments are accurate, in that  $\pi = \tilde{\pi}$  (an individual's self-assessed probability of having malaria is equal to the true probability). In the absence of any diagnostic testing in the retail sector, a decrease in the price of ACTs at the drug shop will decrease  $UT$  and increase  $OT$ . Generally, the relative magnitudes of the changes in  $UT$  and  $OT$  are unclear, so the net effect on  $T$  (the share of illness episodes treated with ACTs that are actually malaria) is ambiguous. However, we can refine the prediction if we make two assumptions:

---

<sup>8</sup>We assume that  $V^a : \pi \rightarrow \mathbb{R}$  is a function, not a correspondence. This is not a trivial restriction – the assumption would be violated if, for example, two illness episodes had equal malaria probability but different likelihoods of being other illnesses of differing severity, such as a cold or pneumonia.



- A1. All cases at the health center are diagnosed, and only given an ACT if the patient tests malaria positive.
- A2. The value of taking action  $a = n$  (doing nothing/taking non-ACT medication at the drug shop) becomes relatively less attractive as  $\pi$  increases:  $\frac{\partial}{\partial \pi} [V^a(\pi) - V^n(\pi)] > 0$  for  $a \in \{s, h\}$

When these two assumptions hold, we have the following result:

- **Result 1:** Suppose the subsidy strictly increases the share of illness episodetreated with ACTs at the drug shop. Then  $UT$  will decrease and  $OT$  will increase such that  $T$  decreases.

A decrease in the price of ACTs in the private sector (holding the health center price constant) will decrease both  $p_P^s(\pi)$  and  $p_N^s(\pi)$  by the amount of the price reduction (assuming the individual always buys the antimalarial upon arriving at a drug shop). This increases the left hand side of inequality 3.1 while leaving  $V^h(\pi)$  and  $V^n(\pi)$  unchanged for all values of  $\pi$ . Note that purchases of ACTs at the drug shop will *strictly* increase unless  $V^s(\pi)$  is everywhere dominated by either  $V^h(\pi)$  or  $V^n(\pi)$  after the price reduction.

Suppose ACT purchases at the drug shop increase. First consider crowd-out from the health center – by assumption A1, this crowd-out will leave  $UT$  unchanged and increase  $OT$ . This shift will clearly work to decrease  $T$ . Now consider crowd out from doing something else. By assumption A2, the curves  $V^s(\pi)$  and  $V^n(\pi)$  will only intersect once, and  $V^s(\pi)$  will cut  $V^n(\pi)$  from below. This implies that all the marginal illnesses induced to take ACTs by the subsidy will have lower malaria probabilities than those illness who would have taken ACTs in either case. This implies that  $T$  will decrease. Thus, under these conditions, an ACT subsidy that increases ACT access will always make targeting of ACTs to those with malaria worse.

However, this prediction is no longer unambiguous if there is heterogeneity in the population. Specifically, suppose there are two types of households, rich and poor. Rich households are able to afford ACTs and travel to the health center in the absence of the subsidy, whereas poor households cannot – they always either hope an illness resolves on its own or they purchase inexpensive medication at the drug shop. Figure 1 graphs the value curves for the rich (panel A) and the poor (panel B) in the case where going to the health center is preferred to buying an ACT at the drug shop for intermediate malaria probabilities. We have normalized the value functions so that  $V^n(\pi) = 0$  for all  $\pi$ .

For the rich, reducing the price of ACTs at the drug shop will lead to crowd-out from the health center among those with intermediate malaria probabilities, and, if the ACT subsidy is large enough, crowd-out from subtherapeutic options among those with a low malaria probability. By Result 1, targeting among rich households will decrease after the subsidy.

In contrast, since none of the poor take an ACT in the absence of the subsidy, introducing the subsidy crowds in a mass of high-positivity ACT takers, which will substantially decrease  $UT$  while only slightly increasing  $OT$ . This crowd-in could actually improve overall targeting if there are

enough poor households in the population. (This would also be true if we considered a less extreme case where, in the absence of the subsidy, the poor only take ACTs when they are very certain that they have malaria). This underscores that it is particularly important to pay attention to distributional impacts of the ACT subsidy. In particular, the subsidy would be especially attractive if it increased takeup among high-positivity populations who didn't have access to ACTs before (this is certainly the intent of the AMFm). On the other hand, it is possible that the subsidy would mostly go to populations who would have gotten the ACT regardless of the subsidy policy (at a health center, for example), or to very low positivity populations – in this case the policy would be mostly wasteful.

### 3.3.3 Impact of Adding an RDT Subsidy at the Drug Shop

Suppose that at a cost of  $p_R$ , an individual can take a diagnostic test for malaria at the drug shop – he or she must pay this cost with certainty. We assume that the diagnostics are perfectly accurate and that all individuals believe that this is the case – in this case, no one will ever take an antimalarial if they test negative. Then there are two primary advantages of taking a test:

1. If the test is negative, the individual avoids the need to pay for an antimalarial – this is particularly attractive when the price of the diagnostic is less than the price of an antimalarial, as it will reduce  $p_N^s(\pi)$ .
2. If the test is negative, it can help the individual to select appropriate medication earlier: this will increase  $U_N^s(\pi) - p_N^s(\pi)$ .

Both of these effects are captured by writing  $\widehat{V}_N^s(\pi) \geq V_N^s(\pi)$ , where  $\widehat{V}_N^s(\pi)$  is the value of seeking care at the drug shop conditional on not having malaria *and* seeing the negative RDT test result. This gives us our second result:

- **Result 2:** Adding an RDT subsidy onto the ACT subsidy will (possibly weakly) decrease  $UT$  and (possibly weakly) decrease  $OT$ , compared to the ACT subsidy regime only.  $T$  will therefore (possibly weakly) increase.

In other words, RDTs offer a way to increase access in the private sector (decrease  $UT$ ) while simultaneously improving targeting (increasing  $T$ ). To see this, we consider the intensive and the extensive margin effects of RDTs in turn. Let's begin with the intensive margin effect: this applies to individuals for whom  $V^s(\pi) \geq \max\{V^h(\pi), V^w(\pi)\}$ . These individuals will continue to seek care from the drug shop and will choose to use an RDT if:

$$\begin{aligned} \pi V_P^s(\pi) + (1 - \pi) \widehat{V}_N^s(\pi) - p_R &\geq \pi V_P^s(\pi) + (1 - \pi) V_N^s(\pi) \\ (1 - \pi) \left[ \widehat{V}_N^s(\pi) - V_N^s(\pi) \right] &\geq p_R \end{aligned} \tag{3.2}$$

that is, they will use the RDT if the expected gain in utility/savings on excess medicine exceeds the cost of the RDT. If individuals always comply with RDT results, this will leave  $UT$  unchanged while decreasing  $OT$ .

Now consider the extensive margin – there may be a set of individuals for whom  $V^s(\pi) < \max\{V^h(\pi), V^n(\pi)\}$  but for whom  $\widehat{V}^s(\pi) \geq \max\{V^h(\pi), V^w(\pi)\}$ . When this crowd out is from the health center, both  $UT$  and  $OT$  will remain unchanged. When this crowd out is from doing nothing/taking something else,  $UT$  will decrease while  $OT$  will remain unchanged. It is important to note, however, that if RDT compliance is imperfect, an RDT subsidy could potentially decrease  $T$  by increasing  $OT$ .

Some of our theoretical predictions (particularly with respect to  $T$ , which captures targeting) are ambiguous because the shapes of the value curves,  $V^a(\pi)$ , the distribution of malaria positivity in the population,  $g(\tilde{\pi})$ , and actual patterns of crowd out are unknown. However, we can learn about these objects if we can observe  $\tilde{\pi}$  and the impacts of ACT and RDT subsidies. Specifically, observing how crowd out in terms of treatment channel (whether an episode is treated at the drug shop, at the health center, or by doing nothing), and therapy choice (whether the episode is treated with an ACT, a substandard antimalarial, or something else) varies with  $\tilde{\pi}$  will allow us to infer where the value curves cross before and after the subsidy. We can also study treatment effects by socioeconomic status to assess the distributional impacts of the subsidy. The next section describes the experiment that we designed to assess the impact of ACT and RDT subsidies on policy parameters  $UT$ ,  $OT$ , and  $T$ .

## 3.4 Study Design and Data

### 3.4.1 Experimental Design

The experiment was conducted in the districts of Busia, Mumias and Samia in Western Kenya between May and December of 2009.<sup>9</sup> Malaria is endemic in this region with transmission occurring year-round, but with two peaks corresponding to heavy rain in May-July and October-November. This region is rural and poor, with the majority of household heads working as subsistence farmers. Daily agricultural wages are estimated at approximately \$1.50 (Dupas and Robinson 2011).

We selected four drug shops, in four rural market centers.<sup>10</sup> We then sampled all households in the catchment area (within a 4km radius) of each of these four drug shops. The total number of sampled households was 2,928. We then visited each household to administer a baseline survey to the female head of household, at the end of which two vouchers for ACTs and (when applicable)

---

<sup>9</sup>The study protocol was approved by the UCLA IRB, the KEMRI/Kenya National Ethical Review Committee, the Kenya Pharmacy and Poisons Board, and the IPA Kenya IRB.

<sup>10</sup>Participating drug shops were chosen on the basis of several criteria including distance from drug shops participating in other public health interventions, shop owner qualifications, length of time the shop had been in business and the number of daily customers.

two vouchers for RDTs were distributed.<sup>11</sup> Enumerators explained that ACTs are the most effective type of antimalarial and, if the household received an RDT voucher, what the RDT was for and how it worked. The vouchers stated the drug shop at which the products could be redeemed and did not have expiration dates so that households had no reason to redeem them in the absence of an illness episode. An image of an ACT voucher with English translation is in Appendix Figure A1. RDT vouchers were similar in appearance.

The experimental design is illustrated in Figure 2. Of the 2,928 households sampled during the census, 2,789 (95 percent) were reached and consented to the baseline survey. Households were randomly assigned to one of three groups, corresponding to the three policy regimes of interest. The “No Subsidy” group received vouchers to purchase unsubsidized ACTs at the market price of Ksh 500 (just under \$6.25). This treatment arm is meant to capture the no-subsidy status quo in Kenya, where over-the-counter ACTs are expensive and RDTs are not available outside a few health facilities.<sup>12</sup> The second group received the ACT subsidy only, meant to reflect outcomes under the planned form of the AMFm in Kenya (i.e. without RDTs). Within the “ACT subsidy only” group, households were randomly assigned to a subsidy level of 92, 88 or 80 percent. Since the AMFm is a top-of-the-supply-chain subsidy, with a floating final retail price, we chose subsidy levels that correspond to final prices likely to capture the range one would see across drug shops in Kenya.<sup>13</sup> The third group received vouchers for both subsidized ACTs and RDTs, with households also randomized into one of three RDT subsidy levels. The most expensive RDTs were subsidized by 85 percent, corresponding to a retail price of roughly \$0.20.<sup>14</sup>

Since ACTs are priced by dose, with the appropriate dose determined by age, the four ACT subsidy levels (0, 80, 88 and 92 percent) differed in the “price-per-pill” to which a household was entitled. Figure A2 in the Appendix demonstrates the pricing and dosing regimens in the study.<sup>15</sup> The randomization of households was done using a computerized random number assignment algo-

---

<sup>11</sup>In rare cases when there was no female head or she was not available, we interviewed the male head of household. The ACT used in this study was Coartem (Artemether Lumefantrine), produced by Novartis. The RDT was the ICT Malaria Pf test, produced by ICT Diagnostics. This type of test only detects the *Plasmodium falciparum* strain of malaria, which accounts for 98 percent of all malaria infections in Kenya and is by far the most deadly strain of malaria (Kenya Division of Malaria Control 2011).

<sup>12</sup>The rationale behind distributing a voucher for unsubsidized ACTs to the control group was to harmonize the level of “endorsement” of the local drug shop across groups, as well as harmonize the amount of information (about effectiveness and availability) provided about ACTs across groups.

<sup>13</sup>This price range also roughly corresponds to the price span from the cheapest to the most expensive non-ACT antimalarials available in drug shops in our area of study.

<sup>14</sup>Some households received RDTs for free, some received RDTs subsidized at 85 percent, and some were offered a refund for the 85 percent subsidized RDT if they tested positive. In practice, we find few substantive differences across these groups in RDT take-up and composition of ACT buyers, so we pool them together for simplicity. For additional detail on the separate RDT treatments, see Cohen et al. (2010).

<sup>15</sup>Ideal dosing is based on weight but manufacturers and the Kenyan Ministry of Health provide age guidelines as well, as it is not always feasible to weigh malaria patients. This study used the age guidelines from the Kenya Ministry of Health.

rithm and was stratified by drug shop, by the household's distance to the drug shop (in quartiles) and by the presence of children in the household.

As discussed in the previous section, the success of a subsidy policy depends upon the fraction of malaria positive people who do not receive an ACT under the policy (*UT*) and the fraction of malaria negative people taking ACTs under the policy (*OT*). To estimate these policy parameters, a sub-sample of households in all the treatment groups other than the no subsidy group was randomly selected to get a "surprise RDT".<sup>16</sup> If these households came to the drug shop to redeem their ACT voucher, but did not redeem an RDT voucher (either because they didn't have one or because they chose not to) they were asked whether they would be willing to take an RDT for free, after they had paid for the ACT. Over the three-month period during which we conducted this exercise, 93 percent of those offered the surprise RDT consented to be tested (or consented for their sick dependent to be tested). All RDTs given at the drug shop were performed by trained study officers posted at the shop. If the patient (the person for whom the ACT voucher was redeemed) had not come to the shop, one of the two study officers accompanied the client back home in order to perform the test on the patient. No one was selected for a surprise RDT in the no subsidy group because we anticipated very low take-up of the ACT in this group. (Indeed, only 8 households redeemed an ACT voucher in this group).

At the end of the experiment we visited households again to administer an endline survey. Only 5 percent of households surveyed at baseline were not reached at endline, and attrition was balanced across treatment arms (Figure 2).

### 3.4.2 Data

We use two types of data in the analysis that follows. The first is the administrative data based on redemptions at the drug shop, and the second is survey data from baseline and endline surveys administered at the beginning and end of the study.

The administrative data captures the details of the drug shop transaction (including medicines bought, symptoms, patient characteristics, and true malaria status in case an RDT was administered). This data was recorded by a trained enumerator posted at the drug shop during opening hours, every single day throughout the study period. This data includes information on over 1,700 drug shop visits made by study households over a four-month period.

The endline survey was administered about four months after the vouchers had been distributed. It asked households to recall all illness episodes that involved fever, chills, headache, sweats, nausea, cough, or diarrhea, that the household experienced in the previous four months. For each of these episodes, we collected information about symptoms, where treatment was sought, what type of malaria test (if any) was taken and what medications were purchased.

We use these two sources of data in combination. First, because it includes malaria status (based

---

<sup>16</sup>49 percent of households not offered an RDT, 71 percent of households offered an RDT for Ksh 15, and 100 percent of households offered a free RDT were selected for surprise testing.

on an RDT test result) for over 1,300 presumed malaria episodes, the administrative data enables us to estimate the relationship between reported symptoms and malaria positivity. The administrative data is also used to explore the impact of subsidies on the uptake of ACTs and RDTs and on the targeting of ACTs toward malaria positive people at the drug shop. In particular, we use voucher redemption data to measure the demand for ACTs and RDTs at the drug shop. Furthermore, we use the subsample of households selected for a surprise RDT to tell us what fraction of ACT buyers at the drug shop are truly malaria positive.

In order to gauge the overall impact of the ACT and RDT subsidies on uptake and targeting, however, we have to know how our sample treated illnesses that did not appear at the drug shop. The endline data allows us to estimate the impact of the subsidies on ACT and RDT uptake for all illness episodes, regardless of where they were treated, but the endline data alone is not informative about the issue of targeting. For that, we would need to know the true malaria status of each illness treated by our households throughout the study. While we cannot observe actual positivity, we can proxy for it based on reported symptoms, since we know from the administrative data what the mapping is between reported symptoms and actual malaria status. (The details of how we estimate malaria positivity are presented in Section 5.1).

One concern is that households in our ACT subsidy treatments would remember significantly more illness episodes, (if the vouchers served to jog their memory or make an illness episode more salient, for example). We do not find any significant evidence of this – Appendix Table A1 shows that there are no systematic differences in illness reporting at endline across treatment groups. After the endline survey we informed households that their vouchers had expired, and we collected unused vouchers back from households.<sup>17</sup>

### 3.4.3 Characteristics of Study Sample

In Table 1 we present basic household characteristics and test for balance across treatment groups. We interviewed the female household head roughly 90 percent of the time. These women are typically married, with five years of education and four dependents. Literacy rates are roughly 60 percent. On average, households live 1.66 kilometers from the drug shop for which vouchers were given and 6.5 kilometers from the nearest public health facility. While roughly 40 percent of households had heard of ACTs at baseline, less than 15 percent had heard of RDTs. Columns (4)-(6) present p-values on F-tests for differences in baseline characteristics across treatment groups. There are no significant differences across treatment groups, other than for the number of acres owned and the age distribution in the household. In particular, our control group has slightly older household heads, with, as a consequence, a significantly higher fraction of adults and lower fraction of infants. Since age is highly correlated with malaria incidence, a lack of balance across treatment groups in the age composition of households could confound estimates of treatment assignment on uptake and

---

<sup>17</sup>As compensation, all households were given a tin of cooking fat at endline regardless of whether or not they returned any vouchers to us.

targeting, even though the magnitudes of the age differences is not large. In all of the results that follow we therefore control for the age of the household head.

## 3.5 Results

### 3.5.1 Predicting Malaria Positivity

As highlighted by the model, a household’s optimal treatment plan for a given illness should vary with the household’s subjectively assessed probability that the illness is malaria. Furthermore, in order to estimate the impact of ACT and RDT subsidies on targeting, we need to be able to assess the malaria status of all illness episodes, not just those treated at the drug shop. Unfortunately, we can only observe actual malaria status for those illness episodes given an RDT at the drug shop.

To address this, we use data on illness-specific characteristics to impute a malaria probability to the universe of illness episodes enumerated at endline. We impute probabilities based on the following probit model, fit to all illnesses that were RDT tested at the drug shop (either due to voluntary redemption or surprise testing) over the course of the study:

$$pos_{eh} = \beta_0 + symp'_{eh}\delta + age'_{eh}\lambda + (symp \times over14)'_{eh}\gamma + \varepsilon_{eh}$$

where  $pos_{eh}$  is a dummy variable equal to 1 if episode  $e$  in household  $h$  tested RDT positive for malaria,  $symp_{eh}$  is a vector of symptom dummies including cough, chills, headache, diarrhea, runny nose, vomiting, body pain, malaise/fatigue, and poor appetite,  $age_{eh}$  is a vector including the patient’s age, age squared, and a dummy variable indicating that the patient is aged 14 or older (an “adult”). We also interact all the symptom dummies with this indicator, to allow for a different relationship between malaria positivity and symptoms among younger and older patients.<sup>18</sup>

The results of this regression are presented in Table 3. Our estimates are consistent with clinical indicators of malaria (CDC 2011) – chills and body pain are positively correlated with malaria positivity, while cough is robustly negatively correlated with malaria positivity. Table 3 also reveals that age correlates very strongly with malaria positivity. Although the interaction terms make the trend somewhat difficult to infer, children (aged 13 and under) who seek care at the drug shop are substantially more likely to actually have malaria as compared to adults (the relevant fractions testing positive are 38 percent for adults and 83 percent for children). While striking, these results are not entirely unexpected – young children are substantially more vulnerable to malaria, as they do not benefit from the acquired immunity that develops with repeated exposure to the parasite. On the other hand, our results are for a selected sample of episodes that households suspected to be malaria and therefore chose to treat at the drug shop. In this case, it is not obvious that an age

---

<sup>18</sup>We do not include the most commonly cited symptom of malaria, fever, in order to avoid endline reporting bias. In Kiswahili (the interview language for our respondents), the word for “fever” and “malaria” are the same – “*homa*”. A concern is that if the ACT subsidy increased ACT taking, respondents would be more likely to identify the illness as malaria, and therefore report *homa* as a symptom at endline.

gradient should be apparent. The presence of the age gradient suggests that households are less adept at identifying malaria in older individuals, possibly because the symptoms are less severe and overlap with other milder diseases, such as cold and flu.

One concern with our predicted positivity measure is that the probit model is fit to a selected sample of illness episodes, while we use the results to impute malaria positivity to the universe of illness episodes (including illnesses very unlikely to be malaria, such as allergies). As such, the imputed probabilities are likely biased upwards at the bottom of the distribution. However, when making ordinal, rather than cardinal comparisons with respect to predicted probability our measure will still be useful as long as the predicted probabilities are directionally correct. Similarly, when using predicted probability to assess targeting (the impact of RDTs on policy parameter  $T$ , the share of ACT takers who have malaria), we only need that predicted probability is consistent for actual probability *conditional* on taking an ACT or RDT. When this is the case, the Appendix illustrates that use of predicted probability will bias our estimates down, so our reliance on predicted probability implies that our results will be lower bounds.

### 3.5.2 Status Quo Treatment-Seeking Behavior

As highlighted by the model, the impact of the two subsidy regimes under study will depend on the relative value of the three possible malaria treatment-seeking behaviors (buying ACTs at the drug shop, going to the health center, other) across malaria risk levels. To get a sense of how these three options compare in the absence of a subsidy, Figure 3 plots the frequency of these three possible actions by predicted positivity among the control group. The figure graphs results of local linear regressions of the following form:

$$y_{eh} = g(\text{predpos}_{eh}) + \varepsilon_{eh} \quad (3.3)$$

where  $y_{eh}$  is the outcome of interest for episode  $e$  in household  $h$  and  $\text{predpos}_{eh}$  is predicted positivity. We present the results for all control group households in Panel A, and separately by SES (proxied by head literacy) in Panels B and C. To avoid overweighting households with many illness episodes and to ensure that results are consistent with the analysis that follows, we only include each household's first illness episode following the baseline survey in all the regressions. Solid gray vertical lines demarcate overall tertiles of predicted positivity, while the dashed gray vertical line demarcates the median.<sup>19</sup>

The figure highlights a sharp contrast in treatment-seeking behavior by SES. For literate households, the likelihood of taking a non- or sub-therapeutic action is clearly decreasing with malaria positivity, in favor of health center visits, while purchase of ACTs at the drug store begins increasing only in the top two tertiles of the malaria positivity distribution. We can draw a number of

---

<sup>19</sup>We calculated quantiles using all first illness episodes for both treatment groups and the control group. We do not update these quantiles when conducting subgroup analysis.



conclusions from these patterns. First, they suggest that our predicted positivity measure captures important heterogeneity in illness episodes that determine treatment seeking behavior. Second, literate households appear to make accurate assessments as to their malaria likelihood, and their treatment decisions appear consistent with the scenario for “the rich” described in the theory section and illustrated in Figure 1. For literate households, we therefore expect that the ACT subsidy will have a modest impact on ACT access for the truly sick (a modest decrease in  $UT$  – the share of true malaria episodes that remain untreated with an ACT), and a potentially large increase in  $OT$  (the share of non-malaria episodes that are treated with an ACT), since most of the crowd out will come from the health center.

The patterns for illiterate households in Panel C are notably different. The share of illness episodes treated at the health center is very low overall and declines sharply in the upper tertile of positivity. The share of episodes for which an ACT is bought at the drug shop is exceptionally low (likely due to the high retail price of ACTs) and increases only weakly with malaria positivity. This is consistent with the scenario for “the poor” discussed in the theory section and illustrated in panel B of Figure 1. The fact that low SES households overwhelmingly do not seek care at the health center, despite the fact that the health center provides (at least in principle) free ACTs, implies that the cost of visiting the health center must be quite high. Even though treatment at the health center is free, the health center is on average 6.5 kilometers away as the crow flies (compared to 1.7 km away for the drug shop). Reaching the health center might therefore require taking costly public transportation, or walking for 1 to 2 hours, something that might be difficult during an illness episode. Overall, the baseline treatment-seeking patterns we observe among illiterate households suggest that, under an ACT subsidy regime, crowd-out of the health center is likely to be minimal for this group, since they are much less likely to seek care at the health center in any case. Instead, the crowd-out is more likely to come from those that choose non- or sub-therapeutic options. Furthermore, if the ACT subsidy draws high positivity illiterate households to the drug shop (that is, crowd-in from illiterates substantially decreases  $UT$ ), this could have attractive implications in terms of both ACT access and targeting.

### 3.5.3 The Impact of Subsidies on Treatment Seeking Behavior

We now study whether two different treatment regimes – ACT subsidy only and ACT+RDT subsidy significantly impact households’ treatment seeking behavior *vis-à-vis* the status quo. To focus on the subsidy versus no subsidy comparison, we pool the three ACT subsidy treatments (92 percent, 88 percent, and 80 percent) into a single group.<sup>20</sup> Since households were only given two ACT and

---

<sup>20</sup>The mix of subsidized ACT prices in the “no RDT subsidy” and “RDT subsidy” treatment groups are slightly different. As such, any differential effects due to the different price mixes will load onto the “RDT subsidy” dummy. However, results for the RDT treatment are nearly identical if we separately dummy out ACT prices and constrain the RDT treatment effect to be constant across subsidized ACT price levels. We therefore present the pooled results for ease of interpretation.

(when relevant) RDT vouchers, we limit the roster of endline illness episodes to the first episode following the baseline survey to ensure that all households had the option of using a voucher if they so desired.<sup>21</sup> We also exclude all ACT-only households who were randomly selected for a surprise RDT test from the endline data analysis, as the surprise test could impact their choice of treatment. We do *not* exclude surprise tested households who received RDT vouchers – 80 percent of these households chose to use their RDTs when going to the drug shop anyway, and F-tests of the significance of surprise testing selection for the ACT+RDT group confirm that the surprise testing had no significant impact on behavior.<sup>22</sup>

We first examine the impact of subsidies graphically by presenting the results of treatment-group specific local linear regressions specified by equation 3.3 in Figure 4. Panel A illustrates results separately for each of the three treatment-seeking behaviors discussed earlier (buying ACTs at the drug shop, visiting the health center, or doing something else). We present evidence of the impacts for the full sample in panel A. Since literate and illiterate households exhibit very different behavior under the status quo, we present evidence of the impact of the two subsidy regimes separately by literacy status in panels B and C.

Panel A shows that both subsidy regimes led to an increase in the likelihood that households buy an ACT at the drug shop across the entire range of predicted positivity. The uniform increase masks different crowd out patterns by predicted positivity, however. In particular, crowd out of the health center is largely concentrated in the upper tertiles of predicted positivity, while crowd out of no or sub-therapeutic care occurs across all ranges of predicted positivity. (Our subsequent analysis will show that this pattern captures two different types of crowdout – doing nothing at low malaria probabilities, versus purchasing a substandard malaria treatment at higher malaria probabilities).

Panels B and C show that the effects of the subsidy regime are very different across literacy groups. All the crowd-out for illiterate households is from no or sub-therapeutic care. For literate households, there is important crowd out of the health center for the middle and upper tertiles of positivity, and crowd out of no/subtherapeutic care in the lower tertile. This is consistent with our theoretical scenario with heterogeneous impacts, as sketched in Figure 1.

In what follows, we analyze the impacts of the subsidy regimes on treatment-seeking behavior in more detail. To do so, we unpack treatment seeking into three domains: (a) provider choice (where to seek treatment) (b) use of diagnostic testing, and (c) drug choice. We then go on to study how ACT price variation *within* a range of subsidized prices impacts ACT demand and targeting.

**Provider Choice** Table 4 estimates the effects of the subsidy regimes on health provider choice. The first panel examines overall mean effects of the different subsidy regimes by presenting results

---

<sup>21</sup>Some households reported more than one member getting sick at once. In these cases, we include all concurrent first episodes, and therefore cluster the standard errors in all illness episode regressions at the household level. Results are similar, though slightly attenuated, if we also include second illness episodes following the baseline survey.

<sup>22</sup>Consequently, our results are largely unchanged, though less precisely estimated, when excluding these surprise tested households.

from the following regression:

$$y_{eh} = \delta + \alpha ACTsub_h + \beta RDTsub_h + age'_{eh}\gamma + \lambda_{strata} + \varepsilon_{eh} \quad (3.4)$$

where  $y_{eh}$  is the outcome of interest,  $ACTsub_h$  is a dummy variable equal to 1 if the household was randomly allocated to either the ACT or ACT+RDT subsidy treatment group,  $RDTsub_h$  is a dummy variable equal to 1 if a household was randomly allocated to the ACT+RDT subsidy treatment group,  $age_{eh}$  is the household head's age, and  $\lambda_{strata}$  are strata fixed effects.<sup>23</sup> Since  $RDTsub$  is only equal to 1 when  $ACTsub$  is equal to 1, the coefficient  $\alpha$  gives the difference between the ACT only treatment group and the control, while  $\beta$  gives the difference between the ACT only subsidy group and the ACT+RDT subsidy group.

The first three columns present results for all households and show that the ACT subsidy increased treatment seeking at the drug shop by 15.9 percentage points (32 percent), while decreasing treatment seeking at the health center by 7.6 percentage points (26 percent). Furthermore, the subsidy encouraged care-seeking for a substantial number of illness episodes – the fraction of households not seeking any care decreased by 9 percentage points (41 percent) in both subsidy treatment groups. These effects are significant at conventional levels (though only marginally so for the health center). While the subsidy decreased rates of not seeking care for both illiterate- and literate-headed households (though our estimates are just short of marginal significance for illiterates), only literate-headed households were crowded out of the health center, as we had hypothesized given baseline treatment-seeking behaviors. Our estimates of  $\beta$  illustrate that provider choice patterns in the ACT and ACT+RDT treatment groups were essentially identical. This suggests that RDTs had a limited extensive margin effect (they do not draw more illness episodes to the drug shop). We therefore expect that any targeting effects of RDTs will be driven primarily by changes in medication taking among households who would have come to the drug shop anyway (i.e. the intensive margin).

The second panel of Table 4 examines impacts of the subsidy treatments by tertile of predicted malaria positivity. That is, we report results from the following regression:

$$y_{eh} = \delta + \sum_{j=1}^3 (\alpha_j ACTsub_h \times tert_{jeh} + \beta_j RDTsub_h \times tert_{jeh}) + \gamma_{age} + \lambda_{strata} + \varepsilon_{eh} \quad (3.5)$$

where  $tert_{jeh}$  is a dummy variable equal to 1 if episode  $e$  in household  $h$  is in tertile  $j$  of overall predicted malaria positivity.<sup>24</sup> The first three columns in this panel illustrate that crowding into the drug shop is most substantial in the first two tertiles of predicted positivity – with crowding out

---

<sup>23</sup>As mentioned earlier, we control for household head's age because the age composition of control households is tilted more towards adults, as illustrated by Table 1. When head age is missing we recode it to the mean and separately dummy these observations out.

<sup>24</sup>Note that although the tertile dummies are generated regressors, the null hypothesis of interest specifies that their coefficients are equal to zero. In this case, traditional standard errors are consistent (Newey and McFadden 1994).

from doing nothing concentrated in the lower tertile of predicted positivity and crowding out from the health center concentrated in the middle tertile of predicted positivity (note, however, that we can only reject that crowd-out patterns are the same across tertiles for “sought no care”).

These impacts are largely concentrated among literate-headed households – this observation, coupled with the increase in ACT purchases by illiterate-headed households observed in Figure 4, suggests that illiterate-headed households may not have substantially changed where they went, but may have substantially changed what they *bought*. That said, general patterns with respect to health center and no care crowd out are similar for illiterate-headed households (indeed, the overall point estimate for “sought no care” suggests a nontrivial reduction for illiterate-headed households due to the subsidy). A notable exception is that we estimate that the ACT and ACT+RDT subsidy treatments significantly *increase* rates of seeking no care for the middle tertile of illiterate-headed households. However, this result is most likely aberration – the control group in the middle tertile cell is relatively small, and the estimates are not robust to recalculating positivity tertiles within the subset of illness episodes among illiterate-headed households.

**Diagnostic Testing** The ACT subsidies clearly draw additional households to the drug shop, sometimes at the expense of the health center. A primary concern with this particular form of crowd out is that it will reduce access to diagnostic services. Figure 5 and Table 5 conduct analyses analogous to those just discussed, but with indicators of malaria testing as outcomes of interest. Panel A of Figure 5 graphs how reported rates of RDT testing, microscopy testing, and the superset, “any malaria test” vary with treatment group and predicted malaria positivity. The ACT+RDT subsidy significantly increased rates of RDT taking, which is to be expected. Furthermore, we find no evidence of crowd out of microscopy overall, despite health center crowd out. (Although microscopy results in Panels B and C of Figure 5 suggest some impacts of the treatments on rates of microscopy testing, our analysis by positivity tertile in Table 5 finds no significant impacts.) Overall, the ACT+RDT subsidy regime nearly doubled the share of illness episodes tested for malaria, from a base of 21.6 percent in the control group up to 42.6 percent.<sup>25</sup>

These large impacts reflect a very high willingness to experiment with RDTs in our sample. As mentioned earlier, over 80 percent of the ACT+RDT treatment households who sought care at the drug shop chose to take an RDT test before deciding whether or not to purchase an ACT. One important caveat is that we cannot tell whether health center crowd out leads to reduced use of *other* types of diagnostic tests. Although an RDT test provides a very useful signal as to whether or not to take an antimalarial, households faced with a negative test result may then face a great deal of uncertainty as to what to take instead. In this situation, consultation with a trained health professional could be particularly valuable.

Table 5 also reveals some evidence that exposure to the ACT-only subsidy increased rates of

---

<sup>25</sup>This result is not driven by including the surprise tested RDT households. When excluding them, we estimate a strikingly similar increase in testing of 21.5 percentage points.

RDT test taking for illiterate households. This is likely due to the fact that, at each study drug shop, we had to make RDTs available for general purchase at cost (Ksh 50, or \$0.64) so as not to deny any households access to diagnostics if they had not received an RDT voucher but had heard about tests being performed at the drug shop and asked for one. We did not advertise for this, but nevertheless some households coming to the study drug shop to redeem ACT vouchers asked to purchase an RDT at the general price.

**Drug Choice** Finally, we study the impact of the subsidy regimes on access to ACTs and other medications. Figure 6 presents local linear regression results for whether or not an illness was: 1) treated with an ACT, 2) treated with some other ineffective antimalarial or an antipyretic (which could improve symptoms but not clear malaria infection), or 3) treated with an antibiotic. Table 6 presents the analogous regression analysis. Overall, both subsidy treatments increase access to ACTs by almost 60 percent. The results by literacy reveal desirable distributional properties of the subsidy. In the control group, literate-headed households are substantially more likely to take an ACT at all levels of malaria positivity (36.5 percent of illness episodes in control literate-headed households were treated with an ACT, as compared to 10.8 percent of episodes in control illiterate-headed households). However, the two subsidy treatments substantially decrease the access gap – illiterate-headed households in the ACT treatment groups took an ACT 38 percent of the time, whereas their literate-headed counterparts took an ACT 44.6 percent of the time.

Furthermore, the subsidy regimes crowded out the use of less effective malaria therapies and antibiotics at higher levels of predicted positivity. Inspection of the graphs and Table 6 reveal that this is concentrated among literate-headed households. In the upper tertile of predicted positivity, the subsidy treatments decreased these households' use of substandard malaria therapies by 28-29 percentage points and antibiotics by 19-20 percentage points. Yet at the same time, ACT access did not significantly increase. This could be driven by a "lowest cost first" approach to malaria treatment. Specifically, a household may first treat a suspected case of malaria with an antipyretic or low-cost antimalarial, hoping that the illness gets better. If the illness does not improve, the household may then try taking a more expensive ACT. If literate-headed households were following this approach and the ACT subsidy made ACTs the "first response" choice to suspected malaria cases, this could generate the patterns in our data.

Overall, rates of ACT taking are similar in the ACT and ACT+RDT treatment groups. The RDT treatment did not significantly change rates of ACT taking in any of the predicted positivity tertiles. However, there is some evidence that RDTs actually *increased* use of antibiotics in the lowest tertile of malaria positivity. This suggests that some individuals may have still taken an ACT just to be safe, but also took other medication in case they were indeed malaria negative.

**Summary** To summarize, our results comparing the subsidy treatments to the control imply that:

- Both subsidies increased treatment seeking at the drug shop. No care is crowded out at the

lowest malaria positivity tertile, while the health center is crowded out in the middle tertile.

- The add-on RDT subsidy significantly increases access to diagnostics for malaria, doubling the share of illness episodes that received a malaria test. The ACT-only subsidy does not significantly reduce access to microscopy.
- Both subsidy regimes significantly increase access to ACTs. The gain is particularly pronounced among illiterate-headed households, who have the lowest rates of access in the control group.

Both subsidy regimes represent a substantial improvement versus the status quo. However, these results are not informative about differences in outcomes between different subsidized price levels. This is important for two reasons: First, the final retail price of ACTs is uncertain under the AMFm. If access falls sharply as the retail price creeps up from the official target price, a larger subsidy may be warranted. Second, an important policy alternative to the AMFm would be to divert some of the ACT subsidy to subsidies for RDTs. The attractiveness of this alternative will depend both on the price elasticity of demand for ACTs across a range of subsidized prices, as well as the benefits of RDTs in terms of improving targeting and access to illness-appropriate treatment. The next subsection investigates these issues by exploiting the within-subsidy price variation in our experimental design.

### 3.5.4 Within-Subsidy Price Variation, ACT Access, and ACT Targeting

We begin by asking whether higher subsidized ACT prices significantly decrease ACT access. To do so, we make use of two different data sources. First, we use administrative data from the drug shop to determine whether or not higher voucher prices resulted in fewer ACT purchases. These results shed light on the impact of price variation on ACT demand *within the private sector*. However, overall changes in access will depend on public sector crowd out as well. Consider increasing the price of an adult ACT dose from Ksh 40 (a 92 percent subsidy) to Ksh 100 (an 80 percent subsidy). If the marginal episodes crowded out of the drug shop instead go to the health center and obtain an ACT anyway, then the net impact on access will be zero. In contrast, if the marginal episodes instead do nothing or take a less effective antimalarial, then overall access will decline. To study overall impacts on access, we exploit our endline data (again, excluding those households who received no RDT voucher and who were selected for surprise testing).

Table 7 contrasts results from our administrative drug shop data and endline data. In order to focus on within-subsidy impacts, we exclude the control group from this analysis. For the drug shop analysis, we include all households in all treatment arms, and present results of the following regression:

$$y_h = \beta_0 + \beta_1 ACT88_h + \beta_2 ACT80_h + \beta_3 ACT92 \times RDT_h + \beta_4 ACT88 \times RDT_h + \beta_5 ACT80 \times RDT_h + ag e'_{eh} \gamma + \lambda_{strata} + \varepsilon_h \quad (3.6)$$

where  $ACT88_h$ ,  $ACT80_h$ , and  $ACT92_h$  are dummy variables for the three different ACT subsidy treatments,  $RDT_h$  is a dummy variable for the RDT subsidy treatment, and  $y_h$  is the outcome of interest. The table presents results for three outcomes: whether or not a household used an ACT voucher at the drug shop (this is equal to zero for households who never redeemed any vouchers), whether the household used an ACT voucher for a patient aged 13 and below, and whether the household used an ACT voucher for a patient aged 14 and older. We only consider voucher redemption for the first visit to the drug shop, as surprise testing could have changed a household's subsequent redemption behavior. We also present results of a specification where we constrain the impact of ACT price (in USD) on outcomes to be linear:

$$y_h = \beta_0 + \beta_1 ACTprice_h + \beta_2 RDT_h + \beta_3 ACTprice \times RDT_h + age'_{eh}\gamma + \lambda_{strata} + \varepsilon_h \quad (3.7)$$

The first column of Table 7 reveals minimal impacts of higher ACT prices on ACT access at the drug shop. Increasing the ACT price by 150 percent (from Ksh 40 to Ksh 100) decreases the share of households using an ACT voucher by 5.5 percentage points (a decline of 13 percent), which implies a price elasticity of demand of just -0.084 over the subsidy range we consider. This is considerably lower than the price elasticity of the demand for malaria-preventing bednets estimated by Cohen and Dupas (2010) in the same area of Kenya. However, a comparison of columns 2 and 3 reveal strikingly different patterns by age. Specifically, households are slightly *more* likely to use an ACT voucher for a child at higher prices, while they are significantly *less* likely to use an ACT voucher for an adult (the implied price elasticity of demand for adults is -0.318). This likely reflects the fact that the price of an ACT dose declines with age. Since we only use information on the first voucher redemption, this could generate the appearance of an upward sloping demand curve for doses for young children if households are willing to treat all ages at the high subsidy level, but only young children at the lower subsidy level (in this case, the overall estimated elasticity of demand will correspond to the elasticity of demand for children). Since malaria positivity is substantially higher at younger ages, this price selection is potentially advantageous from a targeting perspective.

The first three columns do not reveal many significant impacts of RDTs on ACT access. However, this result does not have a clear interpretation. RDTs may reduce rates of ACT taking by screening out malaria negative patients. At the same time, RDTs may also have a selection effect (though our earlier results suggest this may be limited) by which more households may choose to come to the drug shop to use their vouchers. The combined impact of these two effects are of ambiguous sign.

The next three columns present comparable results from our endline data. The first column presents results using all first illness episodes, whereas the second column limits the sample to first episodes in which the patient was aged 13 or younger. The final column presents results for first episodes among individuals aged 14 and above. Although less precisely estimated, our endline data generates similar point estimates and similar patterns of demand by age. This implies that the adults screened out by the higher price at the drug shop did not obtain ACTs elsewhere. Moreover,

there is suggestive (though statistically insignificant) evidence that RDTs were useful for targeting, particularly at the highest subsidy level – they increased the share of ACT takers among children while decreasing the share among adults, who were less likely to be malaria positive.

Table 8 studies targeting more directly, again making use of both our administrative (redemption) data and endline data. All columns present regressions of the same forms specified by equations 3.6 and 3.7, but limited to ACT takers. We also omit strata and age controls so as not to absorb selection effects.<sup>26</sup> In the first three columns, we use the sample of ACT redemptions in which the patient was also selected to receive a surprise RDT test to study the relationship between ACT prices, RDTs, and actual positivity at the drug shop (that is, we study the impact of the RDT subsidy on policy parameter  $T$  at the drug shop). The results indicate that both prices and RDTs improve targeting – higher prices and RDT provision are associated with higher rates of positivity. Furthermore, RDTs appear to be particularly valuable at the highest level of ACT subsidy. These impacts are concentrated among adult patients, and are very large. Just 21 percent of adults who took ACTs in the 92 percent ACT subsidy only group were actually malaria positive. Increasing the ACT price by 50 percent (to the 88 percent subsidy) more than doubles this share, as does providing an RDT at the highest subsidy level.

The next three columns of Table 8 examine impacts on *predicted* positivity (here we include ACT takers who were and were not selected for the surprise test). If our predicted positivity measure performs well, patterns should be similar to what we observe using actual positivity data. This is broadly what we observe, though point estimates are substantially attenuated. This is expected – as illustrated in the Appendix, using predicted positivity instead of actual positivity will bias targeting estimates down. Our endline results should therefore be taken as a conservative lower bound on actual targeting impacts.

The final three columns make use of our endline sample of first illness episodes and the predicted positivity measure. Point estimates in column 7 suggest that higher prices and RDTs increase positivity among ACT takers, though estimates are not significantly different from zero. However, our results are similar to the administrative results using predicted positivity – this suggests our initial results using actual positivity in the surprise tested population may be an accurate assessment of overall targeting impacts.

However, comparing the overall endline results to separate results by age group reveals an apparent puzzle. Consider specification 2 – the RDT main effect is positive overall, but negative when conducting the analysis for subgroups based on age. This is because the overall positive effect is driven by shifting the age distribution of ACT takers towards children aged 13 and younger. Since younger children have higher predicted positivity, cutting the sample by age will wash out this composition effect.

Figure 7 provides a graphical summary of our targeting results, comparing the distribution of predicted positivity among endline ACT takers in four different policy regimes: no ACT subsidy

---

<sup>26</sup>Our three subsidy price treatment groups are balanced in terms of age of the household head.



(the control group), high (92 percent) ACT subsidy-no RDT, low (80 percent) ACT subsidy-no RDT, and low ACT subsidy-RDT. Overall, the low ACT subsidy-RDT regime appears to perform the best, though we cannot reject that the distributions are all equivalent.

Taking our point estimates at face value, we estimate that moving from the 92 percent ACT subsidy-no RDT regime to the 80 percent ACT subsidy with RDT regime would increase predicted positivity among ACT takers by 5 percentage points (off of a base of 65.8 percent) while leaving the share of illness episodes treated with an ACT virtually unchanged. This estimate relies on predicted positivity and may be a substantial underestimate, however – our estimates using actual positivity among drug shop clients imply the targeting benefit would be around 24 percentage points. How beneficial are these changes? The next section takes our estimates and puts them in sharper focus by calculating a variety of cost-effectiveness metrics for the different subsidy regimes.

## 3.6 Cost Effectiveness

### 3.6.1 Methodology

In order to assess the benefits of the different subsidy regimes, we construct estimates of the three policy parameters outlined in Section 3:

- *UT*: The share of malaria episodes that are not treated with ACTs
- *OT*: The share of non-malaria episodes that are treated with ACTs
- *T*: The share of ACT takers who actually have malaria

These parameters assess how well the subsidy regimes perform in terms of access (making sure malaria positive patients get ACTs) and targeting (minimizing the number of malaria negative patients taking ACTs). However, the subsidy optimization problem is subject to a budget constraint, so we also calculate the following measures of subsidy cost:

- The subsidy cost per malaria episode – this captures overall program costs
- The subsidy cost per ACT taken by a malaria positive individual

We make use of our endline data and predicted positivity estimates to calculate the metrics above. Specifically, among the first illness episode sample, we run regressions of the form specified by equation 3.7, where outcomes of interest include (1) the share of episodes treated with an ACT, (2) the share of episodes treated with an RDT, and (3) predicted positivity among ACT/RDT takers and non-takers. We then predict values of each outcome for each ACT-RDT subsidy regime included in our experimental design. Since the subsidy cost of ACTs, malaria positivity, and the

price elasticity of ACT demand varies with patient age, we perform these calculations separately for adults (aged 14 and older) and children (aged 13 and younger).<sup>27</sup>

In terms of costs, we assume that each RDT costs \$0.51 to subsidize. (This is equal to 85 percent of the cost we paid to obtain the RDTs – subsidizing RDTs on a very large scale could bring this cost down since we only ordered a small quantity of tests). The Global Fund currently provides two different scenarios for per-dose subsidy amounts (Roll Back Malaria 2011). For each age group we took the midpoint of the two scenarios, assumed this would be the subsidy cost in our 92 percent scenario, and scaled the other subsidy amounts accordingly. We then combined the subsidy cost and demand/targeting estimates and aggregated up to the population level using the observed age distribution of illness episodes to calculate our measures of subsidy performance.

In addition to our ordinary estimates, we calculated cost effectiveness measures for the ACT+RDT regimes in a “perfect adherence” scenario. Here, we assumed that demand for RDTs was unchanged, but that compliance with test results was perfect, in that only patients who test malaria positive take ACTs.

### 3.6.2 Results

Table 9 presents results of our cost effectiveness calculations. Panel A reveals results that are consistent with our earlier targeting and demand results. Most regimes perform comparably in terms of  $UT$ , although somewhat surprisingly the ACT 92, no RDT regime performs worst, while the ACT 92+RDT regime performs best. The driver of this result can be found in Table 7 – ACT taking is highest among adults and lowest among children in the high subsidy-no RDT regime, while highest among children and lower among adults once RDTs are provided.

The results for policy parameter  $OT$  reveal that both price and RDTs are effective tools for limiting overtreatment. Here, RDTs appear to be most valuable at higher ACT subsidy levels. This is intuitive – when diagnosis is uncertain lower ACT prices will select more “marginal” illnesses that are less likely to be malaria. When RDTs are sufficiently low cost to be attractive to these marginal cases, the test has the potential to provide a great deal of screening value. Indeed, this is apparent in Table 8 – the benefit of RDTs is largest at the highest ACT subsidy level.

We also note that comparing changes in  $UT$  and  $OT$  across regimes conforms to our theoretical predictions regarding RDTs. The introduction of RDTs reduce both  $UT$  and  $OT$ , so  $T$  consequently increases (this is Result 2). Although  $UT$  increases slightly with price across the ACT+RDT regimes (as predicted by Result 1), the measure is essentially flat in the ACT-only regimes – this is largely driven by the very low price elasticity of demand across the range of subsidies that we consider. Taken together, the results in columns 1-3 suggest that a slight increase in the subsidized price

---

<sup>27</sup>As illustrated by Figure A2, ACTs come in 4 different dose sizes, determined by age. Since we do not have sufficient sample size to estimate predicted outcomes for each age group separately, we assume that outcomes are equivalent for all young children. We then combine these estimates with the age distribution in the population and the four different dose sizes to calculate cost effectiveness metrics.

of ACTs could have beneficial targeting impacts while not harming ACT access. However, these results do *not* imply that a subsidy is unwarranted – the stark differences in access between our subsidy group and the control group presented in Table 6 make this very clear.

Should this saved subsidy money be shifted to RDTs? The RDT regimes perform better in terms of targeting, but RDTs are also costly, as reflected by the last two columns of Table 9. In particular, the 80 percent ACT subsidy with no RDTs performs just as well in terms of *UT*, *OT*, and *T* as compared to the same subsidy level plus RDTs, but costs 19 percent less per malaria positive episode treated with an ACT.

However, this does not imply that RDTs do not have the potential to be cost effective. RDT noncompliance in our population was high – while we explicitly advised that patients aged 5 and under take an ACT even when testing negative (consistent with WHO guidelines at the time of the study), 49 percent of patients over 5 still took an ACT when RDT negative. This “cautiousness” in learning from test results is not surprising given the fact that the status quo diagnostic technology is often ignored by health practitioners and has a high rate of false negatives. While RDTs have a lower rate of false negatives (5 percent), it might take some time for households to learn this. Another possible explanation for the high ACT purchase rate after a negative RDT result observed in our experiment is hoarding – households might have decided to buy the ACT dose and to keep it for later (the next malaria episode). Such hoarding could have been encouraged by the experimental design (if households were afraid the vouchers would expire or that the supply of ACTs at drug shops would dry up). Both these issues (lack of information about RDTs and hoarding) could disappear if an ACT+RDT subsidy were to be implemented as the steady state. It is therefore relevant to consider the potential cost-effectiveness of RDTs in the case where compliance with RDT results is improved. We present the “best case scenario” in Panel B of Table 9, where we assume tested individuals only take ACTs when the RDT is positive.

Calculations for *OT* suggest that RDTs have great potential to limit overtreatment and improve targeting. Furthermore, when RDT compliance is high, the additional cost of subsidizing them is lower as they reduce the number of subsidized ACT doses consumed. Our results suggest that moving from the 92 percent ACT subsidy regime to an 80 percent subsidy combined with RDTs would reduce the cost per ACT to malaria positive person by 5 percent while reducing the share of malaria negative illnesses treated with an ACT by 58 percent. Additional research is needed to understand longer run use of and adherence to RDTs – without this information, it is difficult to say how close steady state policy could come to our best case scenario.

Another important benefit of RDTs that are not captured by our calculations is that they may increase the likelihood that a non-malaria illness is treated with appropriate medication promptly. Indeed, we find that illnesses least likely to be malaria were more likely to be treated with an antibiotic in the ACT+RDT treatment group. Given that pneumonia, a bacterial illness, is the largest cause of childhood mortality, this benefit could be substantial, even if individuals who test RDT negative continue to take ACTs.

### 3.7 Conclusion

There is a large class of health issues for which both under-medication and over-medication generate negative spillovers. Under-medication is a public bad for any communicable disease, since the number of untreated individuals increases transmission rates. Over-medication is a public bad whenever the cost of treatment is subsidized. Over-medication is also a public bad when it leads to drug resistance. For this class of health issues, it is thus critical to find the right balance between, on one hand, access and affordability when the medicine is truly needed, and on the other hand, disincentive to overuse the medicine.

Malaria is by far the deadliest in this class of health issues. Malaria kills close to 1 million people each year because of lack of access to effective treatment (World Health Organization 2009). At the same time, parasite resistance to treatment has been developing faster and faster with each new generation of antimalarials. Learning how to reduce malaria mortality and morbidity through prompt access to effective treatment, while at the same time limiting resistance to the latest generation of antimalarials, the ACT, is one of the most pressing and important questions facing the global health community today.

This paper is one step forward in the direction of answering this question. We use detailed data on treatment-seeking behavior among close to 3,000 households in a malaria-endemic area of Kenya, combined with an innovative experimental design that enables us to identify essential pieces of the puzzle: the price elasticity of demand for effective medication, how demand for ACTs varies by malaria risk level, and how access to proper diagnosis affects the demand for medication and targeting. Our analysis leads to three main findings.

First, we find that the demand for ACTs is very elastic at unsubsidized prices, but inelastic over a relatively large range of subsidized prices. This suggests that subsidies for ACTs are clearly needed in order to increase rates of effective treatment among those that suffer from malaria, but these subsidies need not be as large as currently planned by the donor community. Furthermore, we find evidence that price is a useful tool for selection – slightly higher ACT prices reduce ACT taking among adults, who are substantially less likely to be malaria positive, while leaving access among young children unchanged. Second, we find that overdiagnosis of malaria is extremely common, therefore ACT subsidies alone would lead to an important increase in inappropriate use of ACTs. Third, we find that demand for Rapid Diagnostic Testing is extremely high when it is readily affordable and available. Put together, these three findings suggest that bundling the currently proposed ACT subsidy with a subsidy for Rapid Diagnostic Testing is critical in order to achieve the goal of reducing the burden of malaria not only today, but also tomorrow.

### 3.A Appendix: Predicted Positivity and Regression Bias

Here we illustrate the downward bias inherent in using predicted malaria positivity to assess targeting impacts. We are interested in differences in outcomes between two groups: ACT subsidy, no RDT ( $RDT = 0$ ) and ACT+RDT subsidy ( $RDT = 1$ ). We make the following assumptions:

1. Predicted positivity is an unbiased measure of a patient’s actual malaria probability, *conditional on seeking either an ACT or RDT test* :  $E[\hat{\pi}_{eh}] = \tilde{\pi}_{eh}$ .
2. A share  $\Omega$  of patients in the RDT treatment group choose to take an RDT.
3. 100 percent of patients testing RDT positive for malaria take an ACT (this matches our data).
4. A fraction  $\gamma \in [0, 1]$  of patients testing RDT negative for malaria also take an ACT (RDT adherence may depend on  $\tilde{\pi}$ ).

To measure targeting, we wish to estimate the difference in actual malaria positivity among ACT takers in the RDT treatment and control groups. First, consider the sample of episodes that seek malaria treatment (they take an ACT in the control, or they take an ACT or RDT in the treatment). The share of malaria positive individuals who seek care in the control group can be expressed as:

$$E[\tilde{\pi} \mid RDT = 0, care = 1] = \Pi_0$$

In the treatment group, there are two subgroups: those who choose to take an RDT ( $test = 1$ ) and those who choose to take an ACT without first redeeming their RDT voucher ( $test = 0$ ). The share of malaria positive individuals in these two groups can be expressed as:

$$\begin{aligned} E[\tilde{\pi} \mid RDT = 1, care = 1, test = 0] &= \Pi_1^0 \\ E[\tilde{\pi} \mid RDT = 1, care = 1, test = 1] &= \Pi_1^1 \end{aligned}$$

We can then express overall malaria positivity among care seekers in the RDT treatment group as:

$$E[\tilde{\pi} \mid RDT = 1, care = 1] = \Omega \Pi_1^1 + (1 - \Omega) \Pi_1^0 = \Pi_1$$

Note that if the RDT offer changes selection into treatment seeking, we may have  $\Pi_1 \neq \Pi_0$ . In particular, the RDT may select in “marginal” suspected malaria cases, in which case  $\Pi_1 < \Pi_0$ .

In the control group,  $\Pi_0$  is the share of ACT takers who test positive. In the treatment group, the RDT will improve targeting: all of the  $\Pi_1$  positive episodes will take an ACT, but only  $\Omega\gamma(1 - \Pi_1^1) + (1 - \Omega)(1 - \Pi_1^0)$  of the negative episodes they will take an ACT. So, the share of

ACT takers who are malaria positive in the treatment and control groups can be written as

$$\begin{aligned}
S_0 &= \Pi_0 \\
S_1^0 &= \Pi_1^0 \\
S_1^1 &= \frac{\Pi_1^1}{\Pi_1^1 + \gamma(1 - \Pi_1^1)} > \Pi_1^1 \\
S_1 &= \psi S_1^1 + (1 - \psi) S_1^0 = \frac{\Pi_1}{\Pi_1 + \Omega\gamma(1 - \Pi_1^1) + (1 - \Omega)(1 - \Pi_1^0)} > \Pi_1
\end{aligned}$$

where  $\psi = \frac{\Omega(\Pi_1^1 + \gamma(1 - \Pi_1^1))}{\Omega(\Pi_1^1 + \gamma(1 - \Pi_1^1)) + (1 - \Omega)}$  is the share of treatment group patients taking the ACT who were first tested with an RDT. If we could observe true positivity among all ACT takers, the regression:

$$pos = \beta_0 + \beta_1 RDT + \varepsilon$$

(where we limit the sample to those who take ACTs) would give us what we seek:

$$E[\hat{\beta}_1] = S_1 - S_0 = \beta_1$$

Which approaches  $(1 - \Pi_0)$  when  $\gamma \rightarrow 0$  and  $\Omega \rightarrow 1$ . However, if we can only measure *predicted* positivity, we will get:

$$\begin{aligned}
E[\hat{\beta}_1] &= E[\hat{\pi} | RDT = 1, ACT = 1] - E[\hat{\pi} | RDT = 0, ACT = 1] \\
&= E[\tilde{\pi} | RDT = 1, ACT = 1] - E[\tilde{\pi} | RDT = 0, ACT = 1]
\end{aligned}$$

where the second equality holds by assumption 1. Furthermore, since there is no screening in the control group:  $E[\tilde{\pi} | RDT = 0, ACT = 1] = \Pi_0$ . However, RDT-induced screening in the treatment group presents problems. First, we can decompose  $E[\tilde{\pi} | RDT = 1, ACT = 1]$  into contributions from the tested and untested groups:

$$\begin{aligned}
E[\tilde{\pi} | RDT = 1, ACT = 1] &= (1 - \psi) E[\tilde{\pi} | RDT = 1, ACT = 1, test = 0] + \\
&\quad \psi E[\tilde{\pi} | RDT = 1, ACT = 1, test = 1]
\end{aligned}$$

Again, by assumption 1, the untested population poses no problems:

$E[\tilde{\pi} | RDT = 1, ACT = 1, test = 0] = S_1^0$ . The issue lies with the tested group – expected positivity in the group that actually tests positive will always be less than one (too low), while expected positivity in the group that actually tests negative (but still takes an ACT) will always

be greater than zero (too high).

$$E[\tilde{\pi} \mid ACT = 1, test = 1] = \frac{\Pi_1^1}{\Pi_1^1 + \gamma(1 - \Pi_1^1)} E[\tilde{\pi} \mid ACT = 1, test = 1, pos = 1] + \frac{\gamma(1 - \Pi_1^1)}{\Pi_1^1 + \gamma(1 - \Pi_1^1)} E[\tilde{\pi} \mid ACT = 1, test = 1, pos = 0]$$

Now, by assumption 3

$$E[\tilde{\pi} \mid ACT = 1, test = 1, pos = 1] = E[\tilde{\pi} \mid test = 1, pos = 1] = \frac{E[\tilde{\pi}^2 \mid test = 1]}{\Pi_1^1}$$

Similarly, if we further assume that RDT noncompliance,  $\gamma$ , is unrelated to  $\pi$

$$E[\tilde{\pi} \mid ACT = 1, test = 1, pos = 0] = E[\tilde{\pi} \mid test = 1, pos = 0] = \frac{E[\tilde{\pi}(1 - \tilde{\pi}) \mid test = 1]}{1 - \Pi_1^1}$$

putting all of this together

$$E[\tilde{\pi} \mid ACT = 1, test = 1] = \hat{S}_1^1 = \frac{E[\tilde{\pi}^2 \mid ACT = 1, test = 1](1 - \gamma) + \gamma\Pi_1^1}{\Pi_1^1 + \gamma(1 - \Pi_1^1)}$$

since  $E[\tilde{\pi}^2 \mid ACT = 1, test = 1] < \Pi_1^1$ , the numerator is less than  $\Pi_1^1$  so  $\hat{S}_1^1 < S_1^1$ . Specifically:

$$\begin{aligned} E[\hat{\beta}_1] &= \psi\hat{S}_1^1 + (1 - \psi)S_1^0 - S_0 \\ &= \beta_1 + \psi(\hat{S}_1^1 - S_1^1) \end{aligned}$$

This illustrates that our targeting estimates will be biased downwards – so much so that the sign of  $\hat{\beta}_1$  could actually flip to be negative. Of course, if some of the assumptions that we made are violated in practice, the bias on our targeting estimates may be more complicated. (For example, if RDT noncompliance is positively correlated with  $\pi$ , then we could have  $\hat{S}_1^1 > S_1^1$ ). Very generally, we can decompose the bias on  $\hat{\beta}_1$  into three parts – one from estimation of predicted positivity in the control group, one from estimation of predicted positivity among RDT non-takers in the treatment group, and a final part from RDT takers in the treatment group:

$$E[\hat{\beta}_1] = \beta_1 + (S_0 - \hat{S}_0) + (1 - \psi)(\hat{S}_1^0 - S_1^0) + \psi(\hat{S}_1^1 - S_1^1)$$

### 3.B Appendix: Tables and Figures



Table 3-1: Summary Statistics

<i>Characteristics of Interviewed Household Head</i>							
Female	0.867	0.895	0.907	0.292	0.125	0.333	2789
Age (years)	41.7	38.8	38.8	0.041**	0.036**	0.981	2646
Education (years)	5.10	5.36	5.54	0.424	0.158	0.253	2774
Literate	0.575	0.621	0.621	0.258	0.236	0.973	2782
Married	0.783	0.789	0.777	0.860	0.841	0.456	2784
Number Dependents	4.12	4.07	4.13	0.822	0.979	0.586	2663
<i>Household Characteristics</i>							
Number members	5.48	5.29	5.34	0.382	0.521	0.585	2789
Fraction Adults (Ages 14+)	0.623	0.582	0.580	0.044**	0.029**	0.836	2337
Fraction Infants (Under 4)	0.113	0.139	0.141	0.033**	0.018**	0.790	2337
Acres Land	2.72	2.08	2.28	0.045**	0.175	0.087*	2250
Distance from drug shop (km)	1.68	1.66	1.67	0.873	0.966	0.809	2788
Distance from closest clinic (km)	6.57	6.55	6.60	0.919	0.891	0.635	2785
<i>Baseline Malaria Knowledge and Health Practices</i>							
Number bednets	1.77	1.77	1.78	0.994	0.929	0.875	2784
Share HH members slept under net	0.561	0.585	0.573	0.450	0.698	0.455	2661
Heard of ACTs	0.399	0.425	0.427	0.519	0.467	0.904	2771
Heard of RDTs	0.128	0.153	0.140	0.365	0.646	0.375	2786
Treats water regularly	0.408	0.390	0.416	0.648	0.841	0.190	2779
Number of presumed malaria episode last month	1.20	1.20	1.23	0.985	0.744	0.508	2789
<i>Cost Per Episode (Among Those Seeking Any care)</i>							
Total Cost (US \$)	1.63	1.54	1.68	0.694	0.825	0.405	1319

Notes: Household averages. \*\*\*, \*\*, and \* indicate significance at the 99, 95, and 90 percent levels respectively. The exchange rate at the time of the study was around 78 Ksh to US\$1.

Table 3-2: Baseline Treatment Seeking Behavior

	By Household SES				By Patient's Age		
	All	Illiterate	Literate	p-value Illit=Lit	Patient 13 or Younger	Patient 14 or Older	p-value Child=Adult
Household Level Malaria and Diagnostic Incidence							
Number of presumed malaria episodes last month	1.22	1.36	0.994	0.000***	0.617	0.568	–
At least one presumed malaria episode last month	0.685	0.739	0.600	0.000***	0.435	0.387	–
Household member took RDT test (in last month)	0.029	0.034	0.023	0.077*	–	–	–
Household member took microscopy test (in last month)	0.180	0.209	0.133	0.000***	–	–	–
Treatment Seeking for All Presumed Malaria Episodes							
Did not seek care	0.182	0.260	0.147	0.000***	0.139	0.218	0.000***
Went to health center	0.413	0.331	0.448	0.000***	0.470	0.364	0.000***
Went to drug shop	0.369	0.354	0.376	0.337	0.357	0.382	0.159
Medication for All Presumed Malaria Episodes							
No antimalarial taken	0.221	0.302	0.186	0.000***	0.184	0.252	0.000***
Took ACT	0.213	0.120	0.255	0.000***	0.240	0.193	0.002***
Took Sulfadoxine-Pyrimethamine (SP)	0.100	0.074	0.112	0.004***	0.075	0.130	0.000***
Took Amodiaquine (AQ)	0.181	0.166	0.187	0.240	0.212	0.153	0.000***
Took Other Antimalarial	0.072	0.055	0.079	0.029**	0.095	0.050	0.000***
Forgot Name of Antimalarial Taken	0.217	0.285	0.185	0.000***	0.198	0.225	0.089*
Source of Antimalarials (Among Antimalarial Takers)							
Health Center	0.444	0.413	0.454	0.130	0.475	0.416	0.005***
Drug Shop	0.523	0.540	0.518	0.437	0.498	0.552	0.011**
Another Source	0.033	0.048	0.028	0.069*	0.027	0.032	0.414
Cost Per Episode (Among Antimalarial Takers)							
Total Cost (\$US)	1.68	1.38	1.80	0.014**	1.44	1.97	0.000***

Notes: Standard errors clustered at household level for episode-level statistics. \*\*\*, \*\*, and \* indicate significance at the 99, 95, and 90 percent levels respectively.

Table 3-3: Predicting Malaria Positivity – Probit Marginal Effects

	Coefficient	Standard Error
Cough	-0.107***	(0.038)
Chills	0.096**	(0.043)
Headache	0.021	(0.048)
Diarrhea	-0.031	(0.046)
Runny Nose	-0.020	(0.066)
Vomiting	0.008	(0.033)
Body Pain	0.118	(0.085)
Malaise	-0.051	(0.087)
Poor Appetite	-0.013	(0.038)
Age 14 or Above	-0.017	(0.127)
Age	0.081***	(0.017)
Age Squared	-0.007***	(0.001)
(Age 14 or Above) x Cough	0.017	(0.062)
(Age 14 or Above) x Chills	-0.051	(0.074)
(Age 14 or Above) x Headache	0.016	(0.065)
(Age 14 or Above) x Diarrhea	0.057	(0.093)
(Age 14 or Above) x Runny Nose	-0.365	(0.240)
(Age 14 or Above) x Vomiting	0.074	(0.051)
(Age 14 or Above) x Body Pain	-0.168	(0.131)
(Age 14 or Above) x Malaise	0.055	(0.091)
(Age 14 or Above) x Poor Appetite	0.052	(0.075)
(Age 14 or Above) x Age	-0.094***	(0.018)
(Age 14 or Above) x Age Squared	0.007***	(0.001)
DV Mean / N	0.702	1386

Notes: Standard errors in parentheses. Sample includes all individuals who were tested with an RDT by the research team at the drugstore and had nonmissing symptom and age data. \*\*\*, \*\*, and \* indicate significance at the 99, 95, and 90 percent levels respectively.

Table 3-4: Impact of ACT and RDT Subsidy on Treatment Seeking by Literacy and Predicted Malaria Probability

	All			Illiterate			Literate		
	Sought Care at Drug Shop	Sought Care at Health Center	Sought No Care	Sought Care at Drug Shop	Sought Care at Health Center	Sought No Care	Sought Care at Drug Shop	Sought Care at Health Center	Sought No Care
<i>Specification 1 – Main Effect</i>									
$\alpha$ ACT Subsidy	0.159*** (0.047)	-0.076* (0.043)	-0.091*** (0.036)	0.072 (0.074)	0.002 (0.056)	-0.096 (0.061)	0.215*** (0.061)	-0.135** (0.059)	-0.085** (0.043)
$\beta$ ACT Subsidy $\times$ RDT Subsidy	0.000 (0.027)	-0.006 (0.023)	0.006 (0.018)	-0.031 (0.046)	0.034 (0.036)	0.013 (0.036)	0.018 (0.033)	-0.025 (0.030)	0.001 (0.021)
p-val for F-Test $\alpha+\beta=0$	0.000***	0.042**	0.010***	0.530	0.468	0.138	0.000***	0.003***	0.037**
<i>Specification 2 – Impact by Predicted Probability</i>									
$\alpha_1$ ACT Subsidy * Lower Tertile	0.195*** (0.071)	-0.021 (0.060)	-0.186*** (0.062)	0.054 (0.109)	0.080 (0.077)	-0.153 (0.101)	0.298*** (0.097)	-0.088 (0.088)	-0.221*** (0.078)
$\alpha_2$ ACT Subsidy * Middle Tertile	0.171** (0.083)	-0.182*** (0.075)	0.010 (0.054)	0.091 (0.165)	-0.213 (0.147)	0.124** (0.060)	0.201** (0.096)	-0.180** (0.088)	-0.024 (0.069)
$\alpha_3$ ACT Subsidy * Upper Tertile	0.072 (0.089)	-0.047 (0.088)	-0.040 (0.047)	0.023 (0.119)	0.027 (0.088)	-0.097 (0.087)	0.120 (0.129)	-0.142 (0.130)	0.021 (0.048)
$\beta_1$ (ACT $\times$ RDT) * Lower Tertile	-0.009 (0.045)	-0.012 (0.037)	0.024 (0.034)	0.056 (0.070)	-0.010 (0.056)	-0.028 (0.060)	-0.064 (0.060)	-0.012 (0.052)	0.072* (0.042)
$\beta_2$ (ACT $\times$ RDT) * Middle Tertile	-0.026 (0.046)	0.018 (0.040)	-0.005 (0.033)	-0.142* (0.084)	0.054 (0.067)	0.081 (0.063)	0.036 (0.057)	-0.004 (0.050)	-0.047 (0.039)
$\beta_3$ (ACT $\times$ RDT) * Upper Tertile	0.036 (0.043)	-0.024 (0.039)	-0.002 (0.023)	-0.016 (0.083)	0.065 (0.065)	-0.011 (0.053)	0.054 (0.052)	-0.055 (0.050)	-0.002 (0.025)
p-val: $\alpha_1 + \beta_1 = 0$	0.004***	0.542	0.005***	0.253	0.279	0.043**	0.008***	0.214	0.049**
p-val: $\alpha_2 + \beta_2 = 0$	0.056*	0.019**	0.920	0.740	0.259	0.000***	0.006***	0.021**	0.245
p-val: $\alpha_3 + \beta_3 = 0$	0.192	0.397	0.343	0.942	0.246	0.180	0.158	0.111	0.674
p-val: $\alpha_1 = \alpha_2 = \alpha_3$	0.535	0.217	0.049**	0.944	0.202	0.024**	0.507	0.755	0.030**
p-val: $\alpha_1 + \beta_1 = \alpha_2 + \beta_2 = \alpha_3 + \beta_3$	0.742	0.304	0.081*	0.620	0.275	0.000***	0.907	0.688	0.129
DV Mean (Control Group)	0.494	0.290	0.216	0.585	0.154	0.262	0.438	0.375	0.188
N	2042	2042	2042	705	705	705	1332	1332	1332

Notes: Robust standard errors in parentheses, clustered at household level. Regressions include first illness episode that occurred after baseline. A few households have more than one first illness episode if two family members were sick simultaneously. All regressions control for household head age and a full set of strata dummies. \*\*\*, \*\*, and \* indicate significance at the 99, 95, and 90 percent levels respectively.

Table 3-5: Impact of ACT and RDT Subsidy on Diagnosis Seeking, by Literacy and Predicted Malaria Probability

		All			Illiterate			Literate		
		RDT Test	Microscopy Test	Any Malaria Test	RDT Test	Microscopy Test	Any Malaria Test	RDT Test	Microscopy Test	Any Malaria Test
<i>Specification 1 - Main Effect</i>										
$\alpha$	ACT Subsidy	0.008 (0.025)	-0.019 (0.033)	-0.011 (0.038)	0.054 (0.034)	-0.005 (0.046)	0.049 (0.056)	-0.005 (0.036)	-0.036 (0.047)	-0.041 (0.053)
$\beta$	ACT Subsidy $\times$ RDT Subsidy	0.231*** (0.019)	-0.010 (0.019)	0.221*** (0.024)	0.154*** (0.029)	0.015 (0.027)	0.169*** (0.037)	0.266*** (0.024)	-0.016 (0.025)	0.250*** (0.031)
	p-val for F-Test $\alpha+\beta=0$	0.000***	0.327	0.000***	0.000***	0.820	0.000***	0.000***	0.219	0.000***
<i>Specification 2 - Impact by Predicted Probability</i>										
$\alpha_1$	ACT Subsidy * Lower Tertile	0.023 (0.030)	-0.017 (0.051)	0.006 (0.057)	0.075* (0.042)	0.037 (0.072)	0.112 (0.082)	0.009 (0.049)	-0.047 (0.075)	-0.037 (0.086)
$\alpha_2$	ACT Subsidy * Middle Tertile	0.012 (0.048)	-0.010 (0.059)	0.002 (0.070)	0.054 (0.043)	-0.082 (0.112)	-0.029 (0.113)	0.009 (0.064)	-0.012 (0.070)	-0.002 (0.087)
$\alpha_3$	ACT Subsidy * Upper Tertile	0.001 (0.058)	-0.031 (0.063)	-0.030 (0.076)	0.016 (0.084)	-0.006 (0.064)	0.010 (0.108)	-0.007 (0.079)	-0.063 (0.102)	-0.071 (0.106)
$\beta_1$	(ACT $\times$ RDT) * Lower Tertile	0.155*** (0.028)	-0.027 (0.032)	0.128*** (0.040)	0.111*** (0.043)	-0.029 (0.048)	0.082 (0.060)	0.177*** (0.038)	-0.024 (0.042)	0.153*** (0.056)
$\beta_2$	(ACT $\times$ RDT) * Middle Tertile	0.219*** (0.031)	0.002 (0.033)	0.221*** (0.041)	0.177*** (0.047)	0.018 (0.052)	0.195*** (0.065)	0.237*** (0.039)	0.006 (0.042)	0.243*** (0.050)
$\beta_3$	(ACT $\times$ RDT) * Upper Tertile	0.292*** (0.036)	-0.001 (0.029)	0.291*** (0.042)	0.216*** (0.065)	0.064* (0.035)	0.279*** (0.072)	0.319*** (0.043)	-0.029 (0.039)	0.289*** (0.052)
	p-val: $\alpha_1 + \beta_1 = 0$	0.000***	0.324	0.009***	0.000***	0.906	0.006***	0.000***	0.297	0.144
	p-val: $\alpha_2 + \beta_2 = 0$	0.000***	0.886	0.001***	0.000***	0.548	0.131	0.000***	0.930	0.003***
	p-val: $\alpha_3 + \beta_3 = 0$	0.000***	0.587	0.000***	0.005***	0.384	0.005***	0.000***	0.339	0.028**
	p-val: $\alpha_1 = \alpha_2 = \alpha_3$	0.938	0.969	0.923	0.796	0.658	0.525	0.983	0.899	0.879
	p-val: $\alpha_1 + \beta_1 = \alpha_2 + \beta_2 = \alpha_3 + \beta_3$	0.178	0.874	0.307	0.681	0.657	0.690	0.395	0.673	0.528
	DV Mean (Control Group)	0.068	0.148	0.216	0.031	0.108	0.138	0.094	0.177	0.271
	N	2042	2042	2042	705	705	705	1332	1332	1332

Notes: Robust standard errors in parentheses, clustered at household level. Regressions include first illness episode that occurred after baseline. A few households have more than one first illness episode if two family members were sick simultaneously. All regressions control for household head age and a full set of strata dummies. \*\*\*, \*\*, and \* indicate significance at the 99, 95, and 90 percent levels respectively.

Table 3-6: Impact of ACT and RDT Subsidy on Diagnosis Seeking and Medication Choice by Literacy and Predicted Malaria Probability

	All				Illiterate				Literate			
	Took Other				Took Other				Took Other			
	Anti-malarial or Took ACT	Anti-Pyretic	Anti-biotic	Took No Medicine	Anti-malarial or Took ACT	Anti-Pyretic	Anti-biotic	Took No Medicine	Anti-malarial or Took ACT	Anti-Pyretic	Anti-biotic	Took No Medicine
<i>Specification 1 - Main Effect</i>												
$\alpha$ ACT Subsidy	0.153*** (0.041)	-0.052 (0.045)	-0.071** (0.035)	-0.017 (0.022)	0.272*** (0.056)	-0.008 (0.075)	-0.062 (0.052)	-0.012 (0.032)	0.081 (0.057)	-0.092 (0.056)	-0.089* (0.047)	-0.012 (0.022)
$\beta$ ACT Subsidy $\times$ RDT Subsidy	0.006 (0.027)	-0.002 (0.027)	0.022 (0.018)	0.012 (0.011)	-0.055 (0.046)	0.009 (0.027)	0.039 (0.019)	0.008 (0.035)	0.037 (0.035)	-0.015 (0.035)	0.014 (0.023)	0.005 (0.013)
p-val for F-Test $\alpha+\beta=0$	0.000***	0.187	0.134	0.775	0.985	0.628	0.894	0.021**	0.021**	0.035**	0.091*	0.731
<i>Specification 2 - Impact by Predicted Probability</i>												
$\alpha_1$ ACT Subsidy * Lower Tertile	0.157*** (0.056)	0.054 (0.072)	-0.050 (0.051)	-0.054 (0.040)	0.222*** (0.070)	-0.004 (0.112)	-0.009 (0.075)	-0.005 (0.053)	0.122 (0.088)	0.089 (0.096)	-0.084 (0.074)	-0.086 (0.053)
$\alpha_2$ ACT Subsidy * Middle Tertile	0.123 (0.079)	-0.089 (0.083)	-0.060 (0.068)	0.020 (0.018)	0.222 (0.143)	0.042 (0.168)	-0.165 (0.145)	-0.014 (0.020)	0.071 (0.096)	-0.129 (0.094)	-0.028 (0.077)	0.039 (0.026)
$\alpha_3$ ACT Subsidy * Upper Tertile	0.160* (0.083)	-0.175** (0.076)	-0.120* (0.064)	0.001 (0.023)	0.330*** (0.111)	-0.063 (0.122)	-0.063 (0.076)	0.006 (0.042)	0.021 (0.107)	-0.282*** (0.090)	-0.193** (0.092)	0.012 (0.016)
$\beta_1$ (ACT $\times$ RDT) * Lower Tertile	-0.029 (0.043)	-0.010 (0.046)	0.059* (0.033)	0.027 (0.019)	-0.032 (0.063)	-0.008 (0.070)	0.044 (0.048)	-0.004 (0.034)	-0.034 (0.060)	-0.016 (0.064)	0.062 (0.045)	0.040* (0.023)
$\beta_2$ (ACT $\times$ RDT) * Middle Tertile	0.013 (0.045)	-0.044 (0.046)	0.012 (0.032)	0.020 (0.018)	-0.112 (0.079)	-0.095 (0.086)	0.016 (0.050)	0.052** (0.022)	0.074 (0.056)	-0.033 (0.056)	0.015 (0.040)	-0.002 (0.025)
$\beta_3$ (ACT $\times$ RDT) * Upper Tertile	0.017 (0.046)	0.039 (0.046)	0.001 (0.028)	-0.013 (0.015)	0.008 (0.088)	0.140 (0.087)	0.040 (0.044)	-0.032 (0.035)	0.018 (0.055)	-0.011 (0.055)	-0.013 (0.035)	-0.012 (0.016)
p-val: $\alpha_1 + \beta_1 = 0$	0.007***	0.484	0.837	0.484	0.000***	0.907	0.595	0.865	0.258	0.383	0.748	0.395
p-val: $\alpha_2 + \beta_2 = 0$	0.062*	0.077*	0.465	0.000***	0.404	0.730	0.293	0.127	0.097*	0.055*	0.862	0.003***
p-val: $\alpha_3 + \beta_3 = 0$	0.020**	0.046**	0.053*	0.536	0.000***	0.478	0.753	0.470	0.690	0.000***	0.019**	0.958
p-val: $\alpha_1 = \alpha_2 = \alpha_3$	0.929	0.082*	0.670	0.211	0.695	0.865	0.620	0.908	0.759	0.019**	0.396	0.096*
p-val: $\alpha_1 + \beta_1 = \alpha_2 + \beta_2 = \alpha_3 + \beta_3$	0.861	0.090*	0.241	0.038**	0.289	0.737	0.488	0.330	0.728	0.007***	0.178	0.038**
DV Mean (Control Group)	0.259	0.494	0.185	0.049	0.108	0.446	0.138	0.062	0.365	0.531	0.219	0.042
N	2042	2042	2042	2042	705	705	705	705	1332	1332	1332	1332

Notes: Robust standard errors in parentheses, clustered at household level. Regressions include first illness episode that occurred after baseline. A few households have more than one first illness episode if two family members were sick simultaneously. All regressions control for household head age and a full set of strata dummies. \*\*\*, \*\*, and \* indicate significance at the 99, 95, and 90 percent levels respectively.

Table 3-7: Price Elasticity, Within ACT Subsidy Group

	(1)	(2)	(3)	(4)	(5)	(6)
	Administrative Data			Endline Data		
	HH Used ACT Voucher	HH Used ACT Voucher for Patient Age 13 or Below	HH Used ACT Voucher for Patient Age 14 or Above	Dep. Var: Took ACT at First Illness Episode		
				All	Age 13 and Below	Age 14 and Above
<i>Specification 1: ACT Price Dummies (Omitted Category: ACT 40 Only)</i>						
ACT Subsidy = 88%	-0.027 (0.038)	0.032 (0.034)	-0.058** (0.027)	-0.042 (0.060)	0.001 (0.081)	-0.128 (0.087)
ACT Subsidy = 80%	-0.055 (0.037)	0.027 (0.034)	-0.082*** (0.026)	-0.017 (0.058)	0.021 (0.080)	-0.091 (0.083)
ACT 92% x RDT Subsidy	-0.005 (0.036)	0.023 (0.032)	-0.028 (0.027)	0.000 (0.052)	0.063 (0.069)	-0.107 (0.075)
ACT 88% x RDT Subsidy	0.005 (0.033)	0.024 (0.031)	-0.018 (0.021)	0.028 (0.046)	0.053 (0.061)	-0.008 (0.067)
ACT 80% x RDT Subsidy	-0.039 (0.032)	-0.035 (0.030)	-0.004 (0.019)	-0.007 (0.044)	0.011 (0.060)	-0.020 (0.061)
Mean DV (ACT 92%, no RDT)	0.439	0.268	0.171	0.457	0.462	0.450
N	2609	2609	2609	1880	1085	794
Total Effect (ACT 80% & RDT)	-0.094	-0.008	-0.086	-0.024	0.032	-0.112
P-val (ACT 80% & RDT)	0.004***	0.777	0.000***	0.625	0.633	0.111
<i>Specification 2: Linear ACT Price</i>						
ACT Price (US\$)	-0.069 (0.047)	0.028 (0.043)	-0.097*** (0.032)	-0.010 (0.073)	0.030 (0.101)	-0.085 (0.104)
RDT Subsidy	0.034 (0.057)	0.085 (0.052)	-0.051 (0.039)	-0.019 (0.085)	-0.071 (0.115)	0.093 (0.120)
ACT Price (US\$) x RDT Subsidy	-0.054 (0.060)	-0.091 (0.055)	0.037 (0.040)	0.022 (0.081)	0.103 (0.108)	-0.128 (0.115)
Mean DV (No RDT)	0.414	0.291	0.123	0.426	0.490	0.344
N	2609	2609	2609	1880	1085	794

Notes: Robust standard errors clustered at the household level in parentheses. All regressions control for a full set of strata dummy variables and age of the household head. \*\*\*, \*\*, and \* indicate significance at the 99, 95, and 90 percent levels respectively.

Table 3-8: Targeting Effects of ACT Price and RDTs

	ACT Takers at the Drug Shop (Administrative Records)						ACT Takers at Endline (Self Reports)		
	Dep Var: Tested Positive			Dep Var: Predicted Positivity			Dep Var: Predicted Positivity		
	All	Age<14	Age 14+	All	Age<14	Age 14+	All	Age<14	Age 14+
Specification 1: ACT Price Dummies (Ksh 40 Omitted)									
ACT Subsidy = 88%	0.187** (0.080)	0.060 (0.081)	0.256* (0.146)	0.074* (0.039)	0.007 (0.017)	0.046 (0.039)	0.057 (0.044)	-0.004 (0.018)	-0.021 (0.048)
ACT Subsidy = 80%	0.182** (0.084)	0.066 (0.083)	0.170 (0.158)	0.107*** (0.038)	0.015 (0.015)	0.086* (0.051)	0.030 (0.043)	-0.003 (0.015)	-0.011 (0.035)
ACT 92% x RDT Subsidy	0.163** (0.070)	0.053 (0.073)	0.251** (0.110)	0.064* (0.034)	0.016 (0.013)	0.026 (0.031)	0.053 (0.036)	-0.011 (0.015)	-0.004 (0.034)
ACT 88% x RDT Subsidy	0.018 (0.062)	0.015 (0.059)	-0.001 (0.143)	0.019 (0.030)	0.005 (0.014)	-0.036 (0.037)	-0.001 (0.035)	-0.011 (0.016)	0.003 (0.043)
ACT 80% x RDT Subsidy	0.061 (0.067)	0.036 (0.061)	0.142 (0.160)	-0.020 (0.030)	0.000 (0.012)	-0.080 (0.050)	0.021 (0.033)	-0.001 (0.012)	0.013 (0.027)
Specification 2: Linear ACT Price									
ACT Price (US\$)	0.214** (0.107)	0.075 (0.104)	0.227 (0.206)	0.131*** (0.048)	0.019 (0.019)	0.111* (0.066)	0.025 (0.054)	-0.004 (0.019)	-0.011 (0.044)
RDT	0.181 (0.113)	0.046 (0.111)	0.260 (0.206)	0.117** (0.053)	0.023 (0.021)	0.087 (0.060)	0.052 (0.057)	-0.021 (0.023)	-0.021 (0.054)
ACT Price (US\$) x RDT	-0.116 (0.123)	-0.013 (0.118)	-0.143 (0.251)	-0.111* (0.057)	-0.019 (0.023)	-0.136* (0.075)	-0.030 (0.061)	0.015 (0.024)	0.027 (0.054)
DV Mean (Act 40, No RDT)	0.563	0.791	0.214	0.637	0.822	0.358	0.658	0.846	0.407
N	687	499	188	806	591	215	816	569	247

Notes: Robust standard errors in parentheses, clustered at the household level when applicable (columns 6-9). \*\*\*, \*\*, and \* indicate significance at the 99, 95, and 90 percent levels respectively.



Table 3-9: Cost Effectiveness Estimates

	UT: Share Malaria+ Not Treated With ACT	OT: Share Malaria- Treated With ACT	T: Fraction of ACT Takers Who Are Malaria+	Cost/Malara Episode	Cost/ACT to Malaria Positive Episode	Share of Total Subsidy Budget Spent on Malaria+ Episodes
<i>A. Actual RDT Adherence</i>						
ACT 92%	0.550	0.414	0.676	0.667	1.48	0.605
ACT 92% - RDT	0.517	0.370	0.705	0.868	1.80	0.653
ACT 88%	0.548	0.392	0.687	0.633	1.40	0.620
ACT 88% - RDT	0.523	0.366	0.708	0.815	1.71	0.658
ACT 80%	0.543	0.349	0.711	0.570	1.25	0.653
ACT 80% - RDT	0.536	0.357	0.713	0.713	1.54	0.669
<i>B. Perfect RDT Adherence</i>						
ACT 92%	0.550	0.414	0.676	0.667	1.48	0.609
ACT 92% - RDT	0.517	0.122	0.830	0.757	1.57	0.749
ACT 88%	0.548	0.392	0.687	0.633	1.40	0.629
ACT 88% - RDT	0.523	0.140	0.820	0.721	1.51	0.744
ACT 80%	0.543	0.349	0.711	0.570	1.25	0.674

Notes: Assumes unsubsidized costs of \$1.46, \$1.20, \$0.84, and \$0.43 for ACT doses for adults, teens, children, and infants respectively. Assumes \$0.51 subsidy cost for RDTs. "Care seekers" are defined to be patients purchasing an ACT in the ACT-only regimes and patients purchasing an ACT or RDT in the ACT-RDT regimes. The perfect RDT adherence scenario assumes perfect compliance with RDT test results (for ACT-RDT regimes only) and no change in RDT takeup.

Figure 3-1: Theoretical Treatment Seeking Scenarios

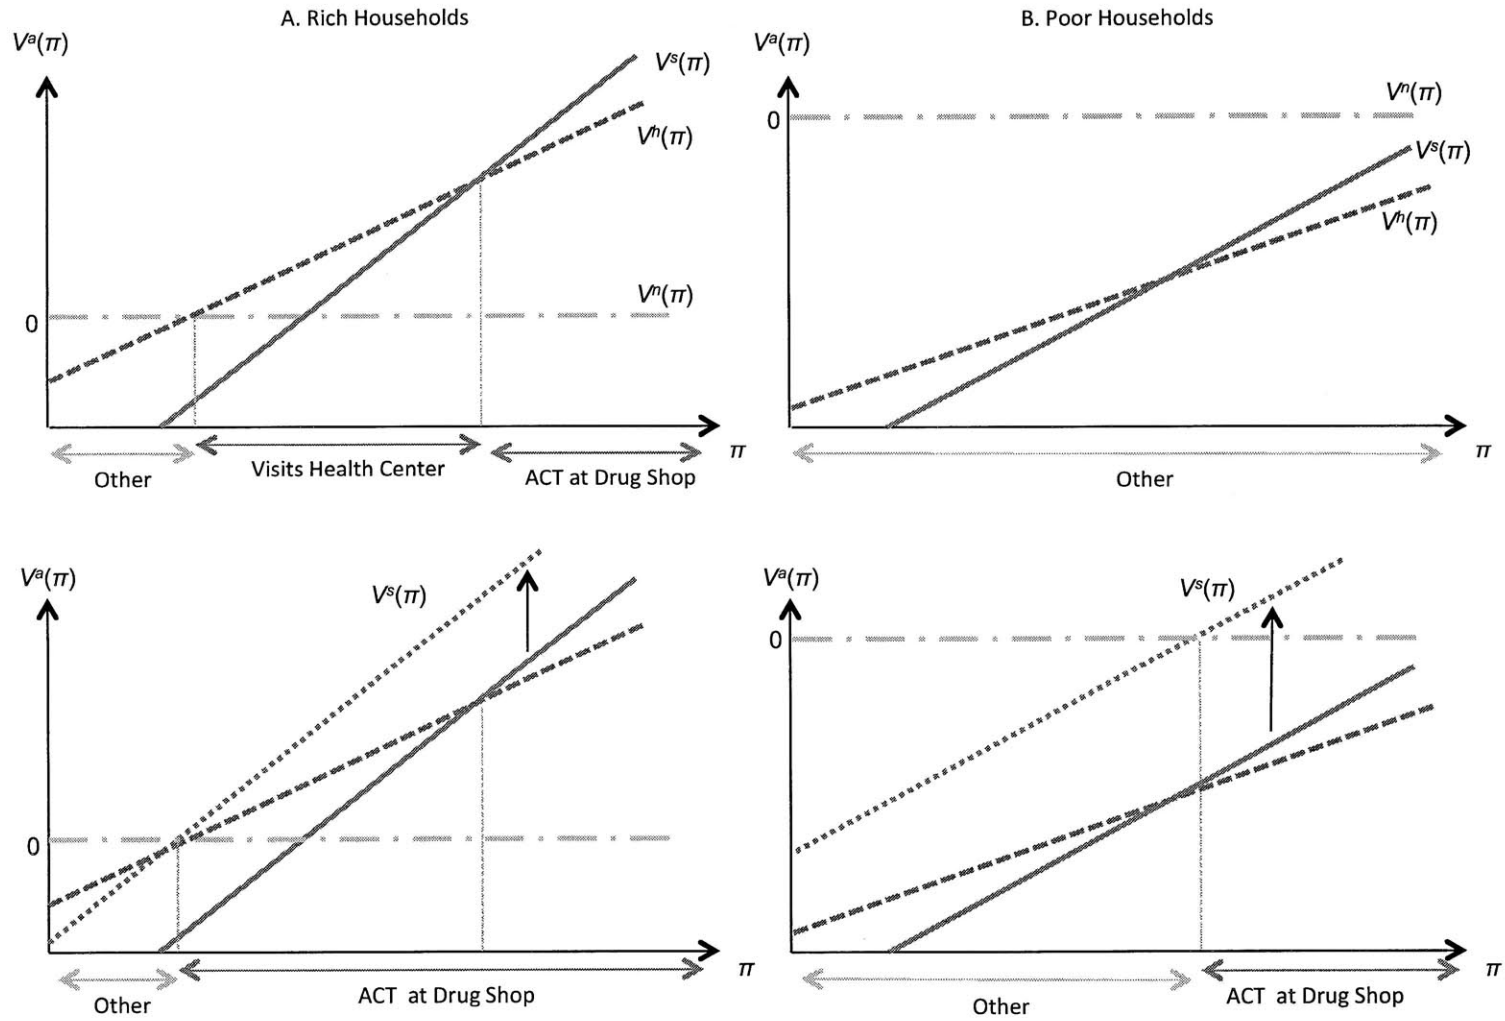
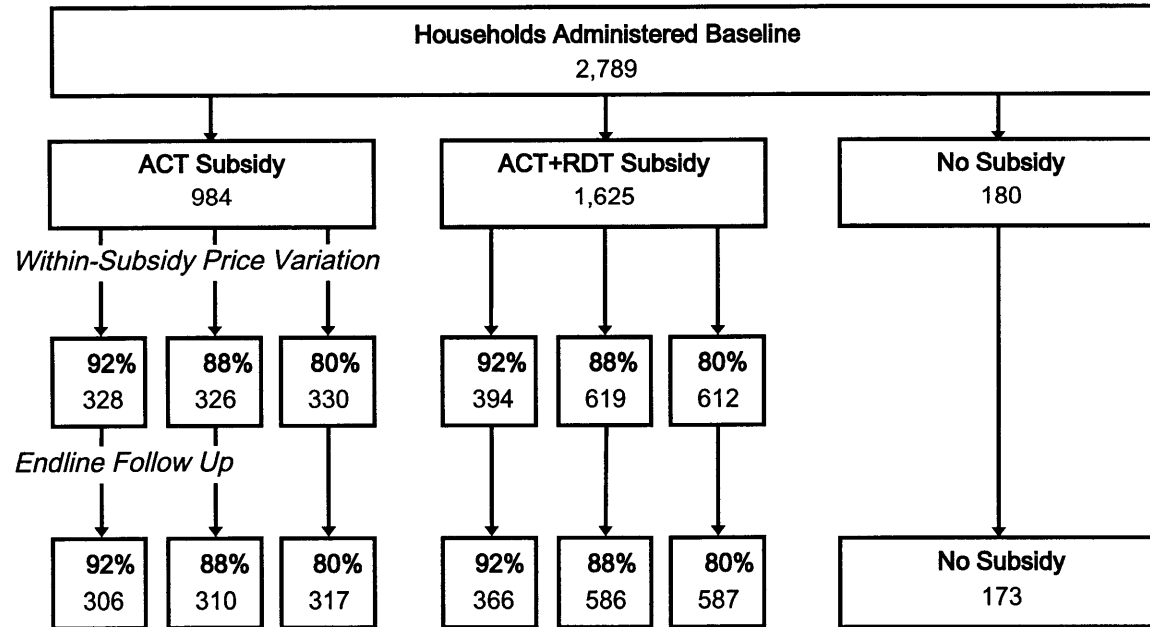


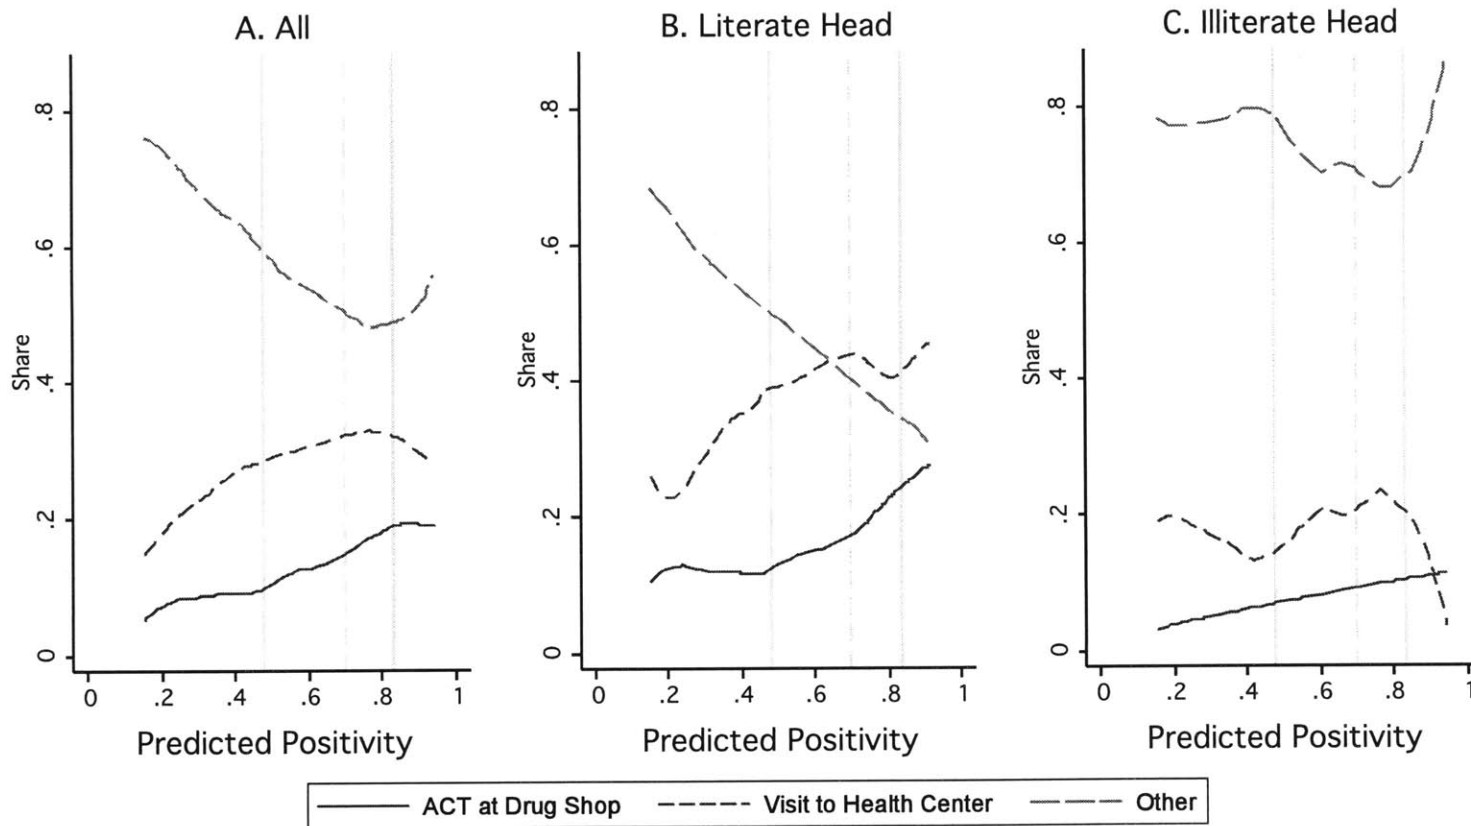
Figure 3-2: Experimental Design and Attrition: Number of Households per Study Arm

Catchment Area Census: Target 2,928 Households



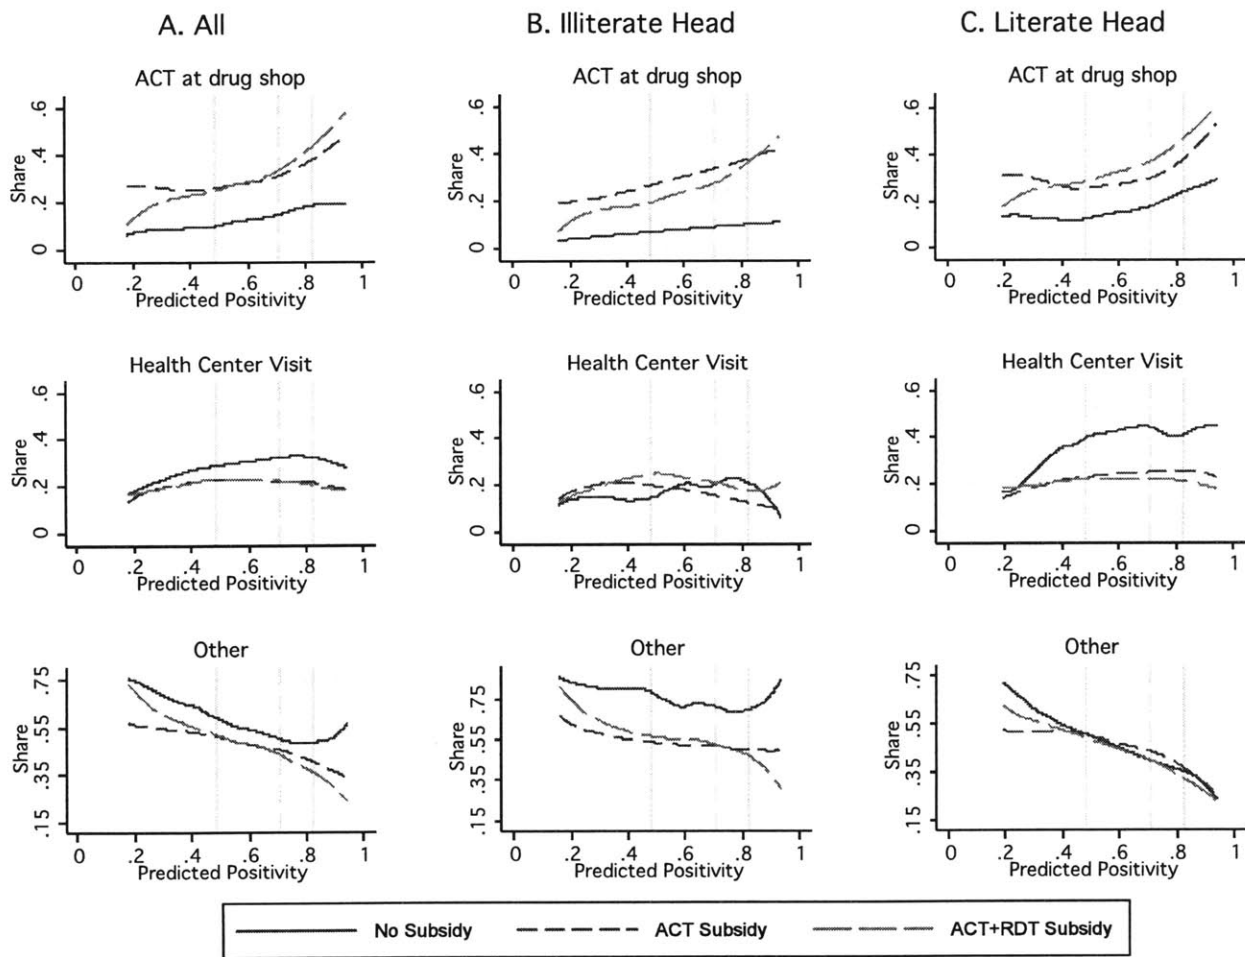
Notes: 49 percent of ACT subsidy only households and 80 percent of ACT+RDT Subsidy households were selected for surprise RDT testing at the drug shop.

Figure 3-3: Baseline Malaria Treatment Seeking Behavior by Predicted Positivity and SES



Notes: Data from "No Subsidy" group. Local linear regression lines trimmed at 2.5 percent. Tertiles demarcated by gray vertical lines. Median demarcated by dashed gray vertical line.

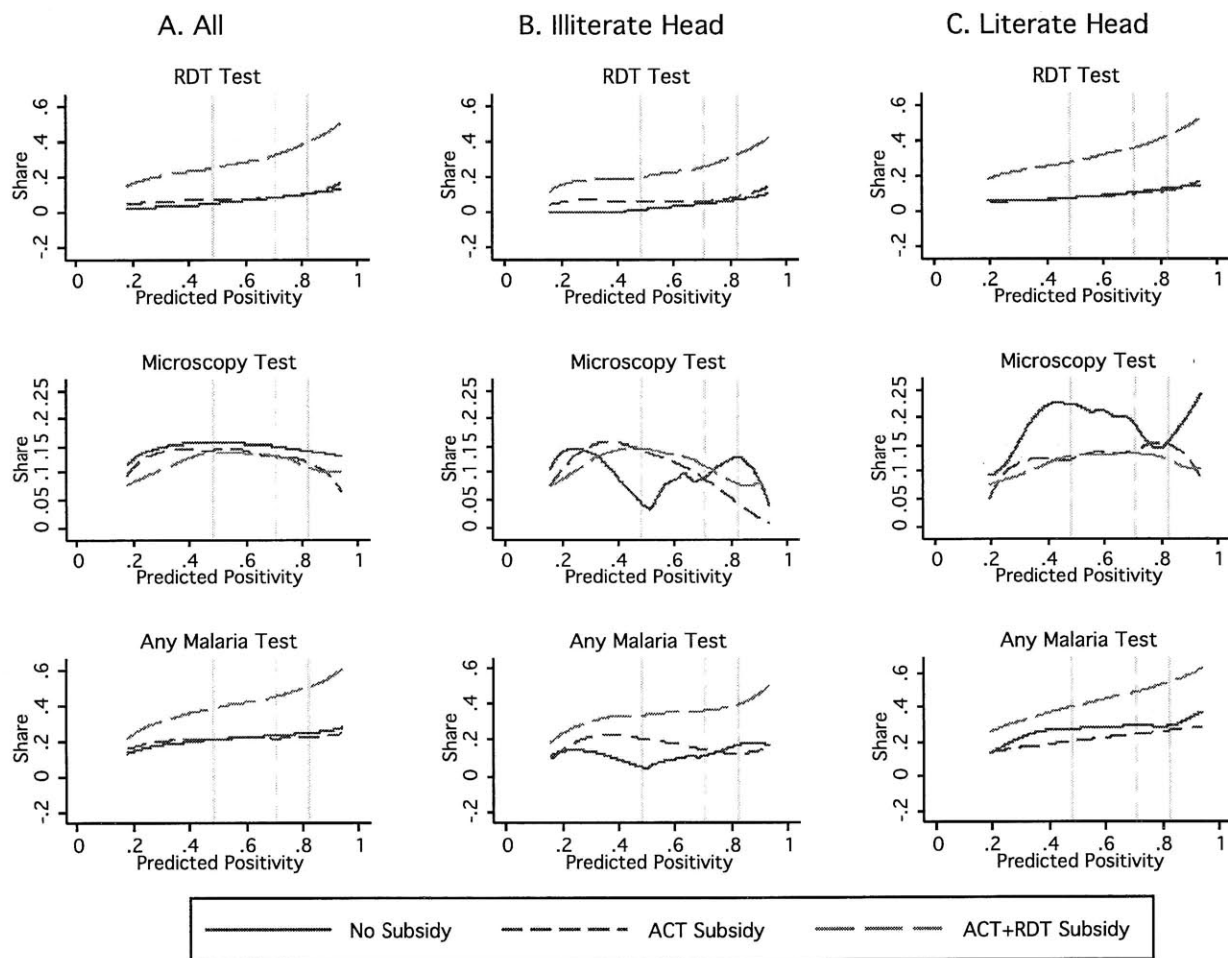
Figure 3-4: Impacts of Subsidy Regimes on Treatment Seeking Behavior by Predicted Positivity and SES



169

Notes: Local linear regression lines trimmed at 2.5 percent. Gray vertical lines demarcate tertiles. Excludes households without RDT vouchers but randomly selected for surprise RDT testing at drug shop.

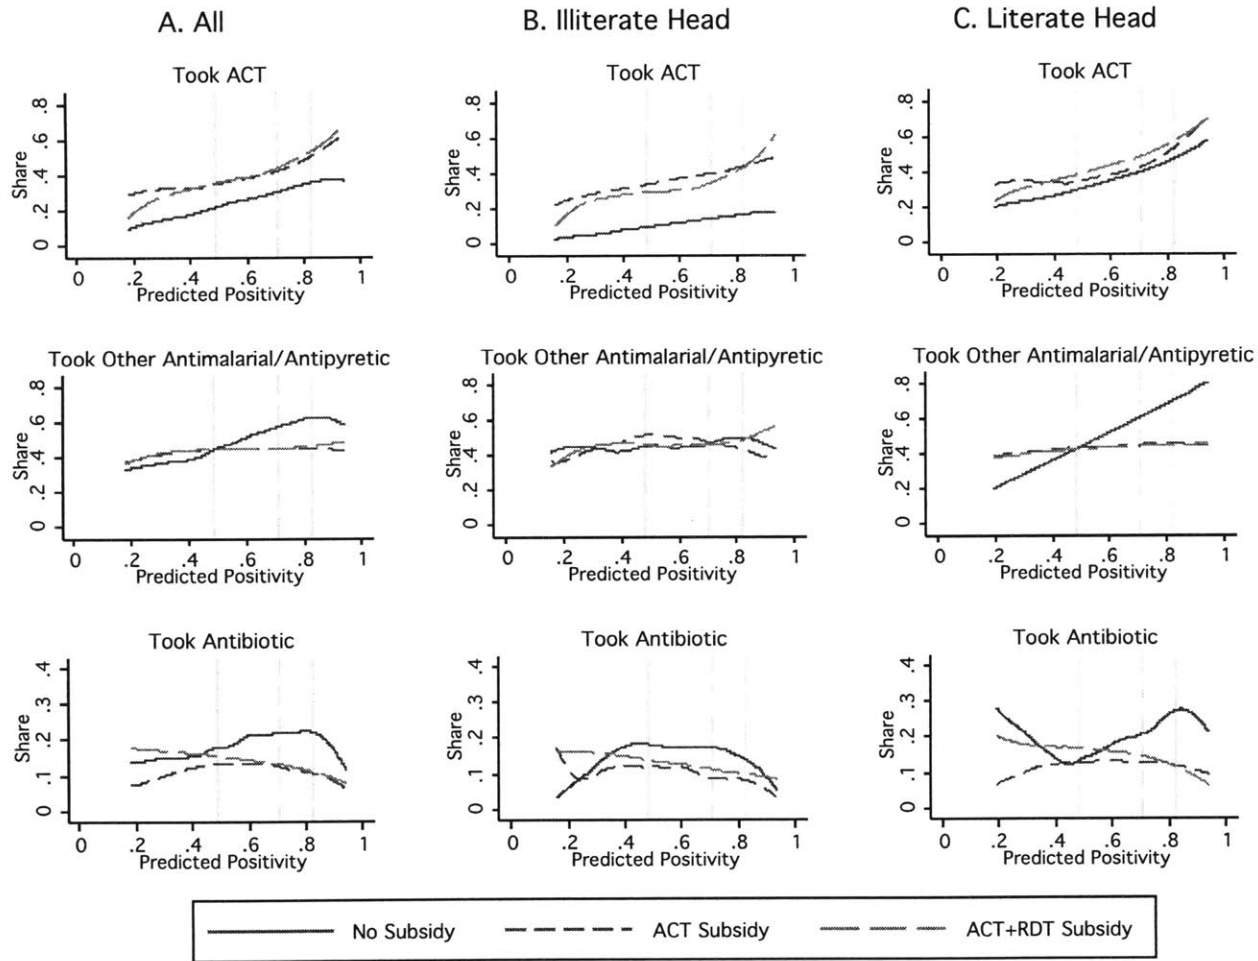
Figure 3-5: Malaria Testing by Predicted Positivity, Subsidy Treatment and Head Literacy



170

Notes: Local linear regression lines trimmed at 2.5 percent. Gray vertical lines demarcate tertiles. Excludes households without RDT vouchers but randomly selected for surprise RDT testing at drug shop.

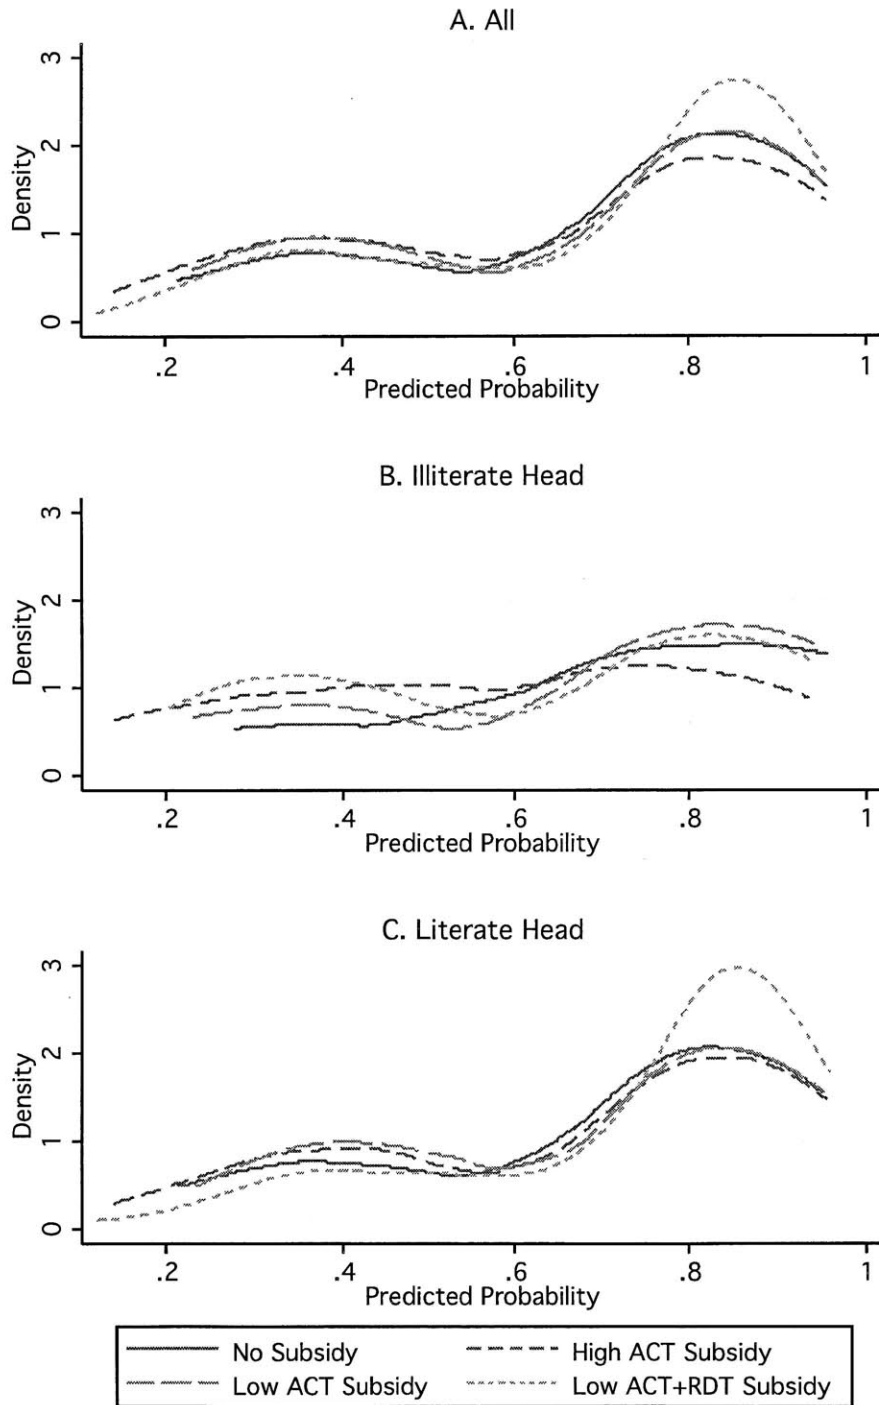
Figure 3-6: Drug Choice by Predicted Positivity, Subsidy Treatment and Head Literacy



171

Notes: Local linear regression lines trimmed at 2.5 percent. Gray vertical lines demarcate tertiles. Excludes households without RDT vouchers but randomly selected for surprise RDT testing at drug shop.

Figure 3-7: Distribution of Predicted Positivity Among ACT Takers by Subsidy Treatment



Notes: Excludes households without RDT vouchers but randomly selected for surprise RDT testing at drug shop.





Table 3-A1: Reporting Bias With Endline Illness Episodes: Comparison Across Subsidy Levels

	Reported Any Illness Episode	Number Episodes Reported	Predicted Malaria Positivity – First Episode	Days Ago – First Episode	Patient Age – First Episode
ACT 92%	0.015 (0.020)	0.024 (0.157)	0.037 (0.023)	1.73 (3.86)	-1.71 (1.65)
ACT 88%	0.002 (0.021)	-0.063 (0.155)	0.039* (0.023)	4.72 (3.75)	-2.92* (1.61)
ACT 80%	-0.020 (0.021)	-0.168 (0.155)	0.031 (0.023)	3.19 (3.78)	-1.69 (1.62)
Any RDT	0.006 (0.010)	-0.025 (0.078)	0.004 (0.010)	-1.27 (1.87)	0.906 (0.777)
Ex-Post Tested	0.001 (0.010)	0.089 (0.079)	-0.017 (0.011)	5.09*** (1.95)	0.988 (0.797)
P-value (92=88=80)	0.005***	0.101	0.765	0.388	0.221
DV Mean	0.950	3.05	0.627	64.7	19.1
N	2621	2621	2473	2438	2473

Notes: Robust standard errors (clustered at the household level when relevant) in parentheses. All regressions include full set of strata dummies and a control for household head age. \*\*\*, \*\*, and \* indicate significance at the 99, 95, and 90 percent levels respectively.

Figure 3-A1: ACT Voucher and Translation

ID: <input type="text"/>	<b>Bei nafuu kwa wagonjwa wanaougua malaria!</b>
	Peleka hii kadi:
	<b>Kwa Duka la Dawa la Funyula Chemist</b> Lililo karibu na Pop Inn Hotel
	Ili kupata dawa ya malaria ya Coartem (AL):
	<b>Ksh 100/-</b> walio na umri wa miaka 14 au zaidi <b>Ksh 75/-</b> watoto kati ya miaka 9 hadi 13 <b>Ksh 50/-</b> watoto kati ya miaka 4 hadi 8 <b>Ksh 25/-</b> watoto wa umri wa miezi 3 hadi miaka 3
	

**A special value for those sick with malaria!**

Bring this card to:

**Funyula Chemist**

Directions

**To obtain the anti-malarial Coartem at a special price:**

**Ksh 100/-** for individuals 14 and older

**Ksh 75/-** for children aged 9-13

**Ksh 50/-** for children aged 4-8

**Ksh 25/-** for children aged 3 months-3 years

***Kadi inakubalika kuanzia saa tatu – hadi saa kumi na moja (Jumatatu hadi Jumamosi)***

Coartem (artemether–lumefantrine, or AL) ni dawa mpya ya malaria iliyo na nguvu kuliko dawa zingine za malaria.

Hii kadi inakuwezesha kununua dawa ya malaria ya coartem iwapo mtu atangojeka kwa nyumba yako. Lazima mtu huyo aje na hii kadi kwenye duka la dawa ili aweze kununua dawa hii. Kumbuka: Bei ni tofauti kulingana na umri kwa sababu watoto humeza kiwango kidogo kuliko watu wazima.

Muhimu: Watoto chini ya miezi 3 na wamama wajawazito kwa miezi 3 za kwanza hawatakikani kumeza hii dawa ya Coartem

***This voucher can only be redeemed from 9AM-5PM Monday-Saturday***

Coartem (artemether–lumefantrine, or AL) is a new anti-malaria drug that is more effective than other drugs currently available to you.

This card may only be used to purchase Coartem for someone in your household. A household member must come with this card to the chemist to make the purchase. Note: dose prices vary by age because children need less medicine than adults.

It is important that infants under the age of 3 months and women in the first trimester of pregnancy do not take Coartem.

Notes: Figure shows voucher front, front translation, voucher back, and back translation respectively.

Figure 3-A2: ACT Price and Dosing Guide

		<i>Recommended Dose and Corresponding Dose Cost for:</i>			
		<b>Adult (14+)</b>	<b>Ages 9-13</b>	<b>Ages 4-8</b>	<b>Ages 3m-3y</b>
<i>Price Per Pill</i>	<i>Dose</i>	4 pills, twice a day for three days	3 pills, twice a day for three days	2 pills, twice a day for three days	1 pill, twice a day for three days
	Ksh 20.83 (Control)	Ksh 500	Ksh 375	Ksh 250	Ksh 125
	Ksh 4.16 (92% Subsidy)	Ksh 100	Ksh 75	Ksh 50	Ksh 25
	Ksh 2.50 (88% subsidy)	Ksh 60	Ksh 45	Ksh 30	Ksh 15
	Ksh 1.66 (80% subsidy)	Ksh 40	Ksh 30	Ksh 20	Ksh 10



# References

- Adhvaryu, A. R. (2009). The Effect of Introducing New Malaria Therapy on Health Care Demand in Tanzania. Mimeo.
- Anderson, S. and J.-M. Baland (2002). The Economics of Roscas and Intra-household Resource Allocation. *Quarterly Journal of Economics* 117(3), 963–995.
- Angrist, J. (2002). How do Sex Ratios Affect Marriage and Labor Markets? Evidence from America’s Second Generation. *Quarterly Journal of Economics* 117(3), 997–1038.
- Aportela, F. (1999). Effects of Financial Access on Savings by Low-Income People. Mimeo, Banco de Mexico.
- Arrow, K., C. Panosian, and H. Gelband (Eds.) (2004). *Saving Lives, Buying Time: Economics of Malaria Drugs in an Age of Resistance*. Washington, DC: Institute of Medicine: National Academies Press.
- 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. Aycinena, C. Martinez A., and D. Yang (2010). Remittances and the Problem of Control: A Field Experiment Among Migrants from El Salvador. Mimeo.
- Ashraf, N., J. Berry, and J. M. Shapiro (2010). Can Higher Prices Stimulate Product Use? Evidence from a Field Experiment in Zambia. *American Economic Review* 100(5), 2383–2413.
- Ashraf, N., D. Karlan, and W. Yin (2006). Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines. *Quarterly Journal of Economics* 121(2), 635–672.
- Ashraf, N., D. Karlan, and W. Yin (2010). Female Empowerment: Impact of a Commitment Savings Product in the Philippines. *World Development* 38(3), 333 – 344.
- Baland, J.-M., C. Guirkinge, and C. Mali (2007). Pretending to be Poor: Borrowing to Escape Forced Solidarity in Credit Cooperatives in Cameroon. Mimeo.
- Banerjee, A., A. Deaton, and E. Duflo (2004). Health Care Delivery in Rural Rajasthan. *Economic and Political Weekly* 39(9).

- Banerjee, A. V. and E. Duflo (Winter 2007). The Economic Lives of the Poor. *The Journal of Economic Perspectives* 21(27), 141–167.
- Banerjee, A. V. and S. Mullainathan (2010). The Shape of Temptation: Implications for the Economic Lives of the Poor. Mimeo.
- Bauer, M. and J. Chytlovic̆ (2009). Women, Children and Patience: Experimental Evidence from Indian Villages. IZA Discussion Paper 4241, Institute for the Study of Labor.
- Baumol, W. J. (1952). The Transactions Demand for Cash: An Inventory Theoretic Approach. *The Quarterly Journal of Economics* 66(4), pp. 545–556.
- Bill and Melinda Gates Foundation (2010, November). Melinda Gates Challenges Global Leaders: Create Savings Accounts and Bring Financial Security to the World’s Poorest. Press Release.
- Bleakley, H. (2010). Malaria Eradication in the Americas: A Retrospective Analysis of Childhood Exposure. *American Economic Journal: Applied Economics* 2(45), 1–45.
- 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. and P. A. Chiappori (1998). Efficient Intra-Household Allocations: A General Characterization and Empirical Tests. *Econometrica* 66(6), 1241–1278.
- Bruce, J. (1989). Homes Divided. *World Development* 17(7), 979 – 991.
- Bruhn, M. and I. Love (2009). The Economic Impact of Banking the Unbanked: Evidence from Mexico. World Bank Policy Research Working Paper 4981.
- Brune, L., X. Gine, J. Goldberg, and D. Yang (2010). Commitments to Save: A Field Experiment in Rural Malawi. Mimeo.
- Burgess, R. and R. Pande (2005). Do Rural Banks Matter? Evidence from the Indian Social Banking Experiment. *The American Economic Review* 95(3), pp. 780–795.
- Caplin, A. and J. Leahy (2004). The Social Discount Rate. *The Journal of Political Economy* 112(6), 1257–1268.
- CDC (2011). About Malaria. <http://www.cdc.gov/malaria/about/disease.html>. Accessed April 14, 2011.
- Central Bank of Kenya (2009, December). Statistical Bulliten. Technical report. Revised.
- Central Bank of Kenya (2010, June). Statistical Bulliten. Technical report.
- Chaia, A., A. Dalal, T. Goland, M. J. Gonzalez, J. Morduch, and R. Schiff (2009). Half the World is Unbanked. Financial Access Initiative Framing Note.

- Chandler, C., C. Jones, G. Boniface, K. Juma, H. Reyburn, and C. Whitty (2008). Guidelines and Mindlines: Why do Clinical Staff Over-Diagnose Malaria in Tanzania? A Qualitative Study. *Malaria Journal* 7(1), 53.
- Chaudhury, N., J. Hammer, M. Kremer, K. Muralidharan, and F. H. Rogers (2006). Missing in Action: Teacher and Health Worker Absence in Developing Countries. *The Journal of Economic Perspectives* 20(26), 91–116.
- 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.
- Chin, A., L. Karkoviata, and N. Wilcox (2010). Impact of Bank Account on Migrant Savings and Remittances: Evidence from a Field Experiment. Mimeo.
- Cohen, J. and W. Dickens (2011). Adoption of Over-the-Counter Diagnostics in Africa: The Role of Subsidies, Beliefs, Externalities and Competition. In *The Value of Information*. Resources for the Future.
- Cohen, J. and P. Dupas (2010, February). Free Distribution or Cost-Sharing? Evidence from a Randomized Malaria Prevention Experiment. *The Quarterly Journal of Economics* 125(1), 1–45.
- Cohen, J., P. Dupas, and S. Schaner (2010). Prices, Diagnostic Tests and the Demand for Malaria Treatment: Evidence from a Randomized Trial. Mimeo.
- Cohen, J., P. Dupas, and S. Schaner (2011). Information, Peers, and Treatment-Seeking Behavior: Experimental Evidence from Rural Kenya. Mimeo.
- 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.
- D'Alessandro, U. and H. Buttiens (2001). History and Importance of Antimalarial Drug Resistance. *Tropical Medicine & International Health* 6(11), 845–848.
- Das, J., J. Hammer, and K. Leonard (2008). The Quality of Medical Advice in Low-Income Countries. *The Journal of Economic Perspectives* 22(2), pp. 93–114.
- de Laet, 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.
- de Oliveira, A. M., J. Skarbinski, P. O. Ouma, S. Kariuki, J. W. Barnwell, K. Otieno, P. Onyona, L. M. Causer, K. F. Laserson, W. S. Akhwale, L. Slutsker, and M. Hamel (2009). Performance of malaria rapid diagnostic tests as part of routine malaria case management in Kenya. *Am J Trop Med Hyg* 80(3), 470–474.

- 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. (2010, August). Short-Run Subsidies and Long-Run Adoption of New Health Products: Evidence from a Field Experiment. Working Paper 16298, National Bureau of Economic Research.
- Dupas, P. and J. Robinson (2011). Savings Constraints and Microenterprise Development: Evidence from a Field Experiment in Kenya. Mimeo, UCLA.
- Feldstein, M. S. (1964). The Social Time Preference Discount Rate in Cost Benefit Analysis. *The Economic Journal* 74(294), 360–379.
- Fudenberg, D. and D. K. Levine (2006). A Dual-Self Model of Impulse Control. *The American Economic Review* 96(5), 1449–1476.
- Global Fund to Fight AIDS, TB and Malaria (2010). AMFm Frequently Asked Questions. Technical report.
- Gollier, C. and R. Zeckhauser (2005). Aggregation of Heterogeneous Time Preferences. *Journal of Political Economy* 113(4), 878–896.
- Gul, F. and W. Pesendorfer (2004). Self-Control and the Theory of Consumption. *Econometrica* 72(1), 119–158.
- Harris, C. and D. Laibson (2001). Dynamic Choices of Hyperbolic Consumers. *Econometrica* 69(4), 935–957.
- 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.
- Hertzberg, A. (2010). Exponential Individuals, Hyperbolic Households. Mimeo.
- Hoffmann, V. (2009, May). Intrahousehold Allocation of Free and Purchased Mosquito Nets. *American Economic Review* 99(2), 236–41.
- Jack, W. and T. Suri (2011). Mobile Money: The Economics of M-PESA. Mimeo.
- Jackson, M. O. and L. Yariv (2010). Collective Dynamic Choice: The Necessity of Time Inconsistency. Mimeo.
- Jakiela, P. and O. Ozier (2011). Does Africa Need a Rotten Kin Theorem? Experimental Evidence from Village Economics. Mimeo.
- Juma, E. and D. Zurovac (2011). Changes in Health Workers' Malaria Diagnosis and Treatment Practices in Kenya. *Malaria Journal* 10(1), 1.



- Kaboski, J. P. and R. M. Townsend (2005). Policies and Impact: An Analysis of Village-Level Microfinance Institutions. *Journal of the European Economic Association* 3(1), pp. 1–50.
- Kangwana, B. B., J. Njogu, B. Wasunna, S. V. Kedenge, D. N. Memusi, C. A. Goodman, D. Zurovac, and R. W. Snow (2009). Short Report: Malaria Drug Shortages in Kenya: A Major Failure to Provide Access to Effective Treatment. *American Journal of Tropical Medicine and Hygiene* 80(5), 737–738.
- Karlan, D. and J. Morduch (2010). *The Handbook of Development Economics*, Chapter 5, Access to Finance. North-Holland.
- Kendall, J., N. Mylenko, and A. Ponce (2010). Measuring Financial Access around the World. World Bank Policy Research Working Paper 5253.
- Kenya Division of Malaria Control (2011). Kenya Malaria Fact Sheet. <http://www.nmcp.or.ke/section.asp?ID=4>. Accessed May 4, 2011.
- Kling, J. R., J. B. Liebman, and L. F. Katz (2007). Experimental Analysis of Neighborhood Effects. *Econometrica* 75(1), pp. 83–119.
- Kremer, M., J. Leino, E. Miguel, and A. Zwane (2011). Spring Cleaning: Rural Water Impacts, Valuation, and Property Rights Institutions. *The Quarterly Journal of Economics* 126(1).
- Lafortune, J. (2010). Making Yourself Attractive: Pre-Marital Investments and the Returns to Education in the Marriage Market. Mimeo.
- Laibson, D. (1997). Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics* 112(2), 443–477.
- Leonard, K. L. (2007). Improving Health Outcomes by Choosing Better Doctors: Evidence of Social Learning about Doctor Quality from Rural Tanzania. Mimeo.
- Leonard, K. L. (2009). The Cost of Imperfect Agency in Health Care: Evidence from Rural Cameroun. *Journal of Development Economics* 88(2), 282 – 291.
- Leonard, K. L., G. R. Mliga, and D. H. Mariam (2002). Bypassing Health Centers in Tanzania: Revealed Preferences for Observable and Unobservable Quality. *Journal of African Economies* 11(4). Mimeo.
- Ligon, E. (2002). Dynamic Bargaining in Households (With Application to Bangladesh). Mimeo.
- Lucas, A. M. (2010). Malaria Eradication and Educational Attainment: Evidence from Paraguay and Sri Lanka. *American Economic Journal: Applied Economics* 2(2), 46–71.
- 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.
- Luoma, M. e. a. (2010). Kenya Health System Assessment 2010. Technical report, Abt Associates Inc., Bethesda, MD: Health Systems 20/20 Project.

- Mani, A. (2010). Mine, Yours, or Ours? The Efficiency of Household Investment Decisions: An Experimental Approach. Mimeo.
- Marglin, S. A. (1963). The Social Rate of Discount and The Optimal Rate of Investment. *The Quarterly Journal of Economics* 77(1), 95-111.
- Marsh, V. M., W. M. Mutemi, A. Willetts, K. Bayah, S. Were, A. Ross, and K. Marsh (2004). Improving Malaria Home Treatment by Training Drug Retailers in Rural Kenya. *Tropical Medicine & International Health* 9(4), 451-460.
- Mazzocco, M. (2007). Household Intertemporal Behaviour: A Collective Characterization and a Test of Commitment. *Review of Economic Studies* 74(3), 857 – 895.
- McGuire, T. G. (2000). Chapter 9: Physician Agency. Volume 1, Part 1 of *Handbook of Health Economics*, pp. 461 – 536. Elsevier.
- 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.
- O'Donoghue, T. and M. Rabin (1999). Doing It Now or Later. *The American Economic Review* 89(1), 103-124.
- Olmstead, T. and R. Zeckhauser (1999). The Menu-Setting Problem and Subsidized Prices: Drug Formulary Illustration. *Journal of Health Economics* 18(5), 523 – 550.
- Otten, M., M. Aregawi, W. Were, C. Karema, A. Medin, W. Bekele, D. Jima, K. Gausi, R. Komatsu, E. Korenromp, D. Low-Beer, and M. Grabowsky (2009). Initial Evidence of Reduction of Malaria Cases and Deaths in Rwanda and Ethiopia Due to Rapid Scale-up of Malaria Prevention and Treatment. *Malaria Journal* 8(1), 14.
- Patouillard, E., K. Hanson, and C. Goodman (2010). Retail Sector Distribution Chains for Malaria Treatment in the Developing World: A Review of the Literature. *Malaria Journal* 9(1), 50.
- Perkins, M. and D. Bell (2008). Working Without a Blindfold: The Critical Role of Diagnostics in Malaria Control. *Malaria Journal* 7(Suppl 1), S5.
- Platteau, J.-P. (2000). *Institutions, Social Norms, and Economic Development*. Amsterdam: Harwood Academic.
- Prina, S. (2011). Do Simple Savings Accounts Help the Poor to Save? Evidence from a Field Experiment in Nepal. Mimeo, Case Western Reserve University.
- Reyburn, H., R. Mbatia, C. Drakeley, I. Carneiro, E. Mwakasungula, O. Mwerinde, K. Saganda, J. Shao, A. Kitua, R. Olomi, B. Greenwood, and C. Whitty (2004). Overdiagnosis of Malaria

- in Patients with Severe Febrile Illness in Tanzania: A Prospective Study. *British Medical Journal* 329(7476), 1212.
- Robinson, J. (2008). Limited Insurance Within the Household: Evidence from a Field Experiment in Kenya. MPRA Paper 8314.
- Roll Back Malaria (2011). Fact Sheet: Updated ACT Prices Under the Affordable Medicines Facility-malaria, 1 March 2011.
- Rutherford, S. (1999). Savings and the Poor: The Methods, Use and Impact of savings by the Poor of East Africa. Technical report, MicroSave-Africa, Kampala.
- Rutherford, S. (2000). *The Poor and Their Money*. Oxford University Press.
- Schaner, S. G. (2010). One Fee Does Not Fit All: The Impact of ATM Card Provision on Formal Savings Account Use in Kenya. (Mimeo).
- Schaner, S. G. (2011). Intrahousehold Preference Heterogeneity, Commitment, and Strategic Savings: Theory and Evidence from Kenya. Mimeo, MIT.
- Shapiro, J. (2010). Discounting for You, Me, and We: Time Preference in Groups and Pairs. Mimeo.
- Steel, W. F. and E. Aryeetey (1994). Informal Savings Collectors in Ghana: Can they Intermediate? *Finance and Development March*, 36–37.
- 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.
- Tarozzi, A., A. Mahajan, B. Blackburn, D. Kopf, L. Krishnan, and J. Yoong (2011). Microloans, Insecticide-Treated Bednets and Malaria: Evidence from a randomized controlled trial in Orissa (India). Mimeo.
- Thaler, R. H. (1990). Anomalies: Saving, Fungibility, and Mental Accounts. *The Journal of Economic Perspectives* 4(1), 193–205.
- Thomas, D. (1994). Like Father, like Son; Like Mother, like Daughter: Parental Resources and Child Height. *The Journal of Human Resources* 29(4), 950–988.
- Tobin, J. (1956). The Interest-Elasticity of Transactions Demand For Cash. *The Review of Economics and Statistics* 38(3), pp. 241–247.
- Trape, J. (2001). The Public Health Impact of Chloroquine Resistance in Africa. *Am J Trop Med Hyg* 64(1), 12–17.
- Udry, C. (1996). Gender, Agricultural Production, and the Theory of the Household. *The Journal of Political Economy* 104(5), 1010–1046.
- Weitzman, M. L. (2001, March). Gamma Discounting. *American Economic Review* 91(1), 260–271.

- Welch, H. G., L. M. Schwartz, and S. Woloshin (2011). *Overdiagnosed: Making People Sick in the Pursuit of Health*. Beacon Press.
- White, N. (2004). Antimalarial drug resistance. *J Clin Invest.* 113(8), 1084–1092.
- Wolinsky, A. (1993, Autumn). Competition in a Market for Informed Experts' Services. *RAND Journal of Economics* 24(3), 380–398.
- World Bank (2004). World Development Report 2004: Making Services Work for Poor People. Technical report.
- World Health Organization (2009). World Malaria Report. Technical report, World Health Organization.
- World Health Organization (2010a). Guidelines for the Treatment of Malaria. Technical report.
- World Health Organization (2010b). Malaria Rapid Diagnostic Test Performance. Technical report.
- Zuber, S. (2010). The aggregation of preferences: can we ignore the past? *Theory and Decision*, 1–18.
- Zurovac, D., B. Midia, S. A. Ochola, M. English, and R. W. Snow (2006). Microscopy and Outpatient Malaria Case Management among Older Children and Adults in Kenya. *Tropical Medicine & International Health* 11(4), 432–440.