Credit Constraints, Discounting and Investment in Health: Evidence from Micropayments for Clean Water in Dhaka*

Raymond Guiteras
University of Maryland

David I. Levine
U.C. Berkeley Haas

Thomas Polley
Duke University

Brian Quistorff
University of Maryland

October 24, 2015

Abstract

Low rates of adoption of and low willingness to pay for preventative health technologies pose an ongoing puzzle in development economics. In the case of water-borne disease, the burden is high both in terms of poor health and cost of treatment. Inexpensive preventative technologies are available, but willingness to pay (WTP) for products such as chlorine treatment or ceramic filters has been observed to be low in a number of contexts. In this paper, we investigate whether time payments (micro-loans or dedicated micro-savings) can increase WTP for a high-quality ceramic water filter among 400 households in slums of Dhaka, Bangladesh, where water quality is poor and the burden of water-borne disease high. We use a modified Becker-Degroot-Marschak mechanism to elicit WTP for the filter under a variety of payment plans. Crucially, we obtain valuations from each household across all payment plans, which (a) increases power and (b) allows us to investigate the mechanisms behind differences in WTP across plans. We find that time payments significantly increase WTP: compared to a lump-sum up-front purchase, median WTP increases 83% with a six-month loan and 115% with a 12-month loan. Similarly, coverage can be greatly increased: at an unsubsidized price (50% subsidy) coverage is 12% (27%) under a lump-sum but as high as 45% (71%) given time payments. We use our rich within-household WTP data, the design of the payment plans, and a simple structural model of time preference and credit constraints to investigate the mechanisms. We find that households are not impatient with respect to health goods and that therefore time-preferences do not contribute to low baseline WTP. We find strong evidence for the presence of credit constraints.

Keywords: Microfinance, Savings, Health
JEL Codes: C93, D14, D91, G21, I15, O16, Q53, Q56

---

*Contact: guiteras@econ.umd.edu, levine@haas.berkeley.edu, thomas.polley@duke.edu, quistorff@econ.umd.edu. We thank the Bill and Melinda Gates Foundation and 3ieimpact for financial support. Ariel BenYishay, Jessica Goldberg, Glenn Harrison, John Rust, Steve Stern, Ken Train, Song Yao and participants in the Columbia University Sustainable Development Seminar, Georgetown University gui2de Development Economics Seminar, WADES, UC Berkeley Development Lunch, University of Maryland

1 Introduction

Low rates of adoption of and low willingness to pay for preventative health technologies pose an ongoing puzzle in development economics (Dupas, 2011; Abdul Latif Jameel Poverty Action Lab, 2011). In the case of water-borne disease, the burden is high both in terms of poor health and cost of treatment, and inexpensive preventative technologies are available, but willingness to pay for products such as chlorine treatment or ceramic filters has been observed to be low in a number of contexts (Ahuja et al., 2010; Ashraf et al., 2010; Luoto et al., 2011; Berry et al., 2015; Guiteras et al., 2015).

Many explanations for this puzzle have been proposed. We focus on one common characteristic of many health technologies: a relatively large up-front investment is required, while the benefits accrue over time. This is problematic for a number of interdependent reasons. First, households may find it difficult to borrow, especially for non-business purposes. Second, poor households may have high discount rates or be close to subsistence levels of consumption and therefore be unwilling to sacrifice a large amount of current consumption. Third, households may exhibit time-inconsistency in the form of present bias or hyperbolic discounting (Ashraf et al., 2006). Fourth, households may be unwilling to sink a large sum into a new technology when they are unsure of its benefits. These barriers suggest a number of interventions to increase adoption and improve welfare. Consumers who face liquidity constraints or exhibit present bias may find it difficult to fund purchases even if they are willing to pay substantial amounts over time (Holla and Kremer, 2009). As a result, time payments, either micro-loans or layaways (dedicated savings), may increase adoption and improve welfare (Tarozzi and Mahajan, 2011; Dupas and Robinson, 2013). When consumers have uncertain valuation of a new product, a free trial or money-back guarantee can allow learning at low risk (Levine and Cotterman, 2012).

Applied Micro Lunch, University of New South Wales and Yale / Citi China India Insights Conference 2013 provided helpful comments. We are grateful to Raihan Hossain and Ashraful Islam for research assistance and project management, and icddr,b for collaboration in field implementation. This work was motivated by a previous collaboration with Kaniz Jannat, Stephen P. Luby and Leanne Unicomb. All errors are our own.

In this paper, we examine how time payment plans (either micro-loans or layaways) and interventions to decrease the risk incurred while learning (free trial, money-back guarantee) affect willingness to pay (WTP) and attempt to understand the mechanisms at work. Both of these are empirically challenging. First, individuals with greater access to finance may have a greater taste for health relative to consumption or more resources overall. Second, even if access to finance were randomly assigned, there are many variations possible and we would typically only observe one choice per individual, so it would require an enormous sample size to determine which policies are most attractive. Third, many of the underlying reasons for increased willingness to pay (liquidity constraints, high discount rates, present bias / hyperbolic discounting, value of low-risk learning) have similar empirical implications.

To address these questions, we measure WTP for a high-quality ceramic water filter in 400 households in slums of Dhaka, Bangladesh, where water quality is poor and the burden of water-borne disease high. We use a modified Becker-Degroot-Marschak (BDM) mechanism to elicit WTP under a variety of time payment plans, including a lump-sum paid immediately, micro-loans and dedicated micro-savings plans of varying duration. Crucially, we obtain valuations from each household across all payment plans, which (a) vastly increases power and (b) allows us to investigate the mechanisms behind differences in WTP across plans.

We find that the availability of time payments dramatically increases willingness to pay. While the retail price is BDT 2100 (USD 28)[1], median WTP under a lump-sum, up-front payment is BDT 755 (USD 10.07), but increases to BDT 1260 (USD 16.80) with a simple 6-month loan and BDT 1530 (USD 20.40) for a 12-month loan. To separate time preference from liquidity constraints, we elicited WTP from subjects given layaway (dedicated micro-savings) plans with the same payment schedule as the loans. The intuition for this approach is that, while layaway plans should be less appealing than loans to all consumers, patient consumers who are liquidity constrained will find the layaway relatively more appealing than will impatient consumers. To our surprise, we found that for almost all households,

[1]The exchange rate during the study was roughly BDT 75 = USD 1.

WTP with a loan is virtually identical to WTP with a layaway plan with the same payment schedule – that is, they are willing to pay the exact same amount over 6 months to receive the filter in 6 months as they are to receive the filter today. In a standard model where all forms of consumption are discounted at the same rate, this suggests that liquidity constraints are more important than time preference in explaining the large increase in WTP from time payments. Alternatively, households could discount future general consumption (i.e. money) heavily, but do not discount the use of the filter at all. To investigate the mechanisms at work, we estimate a simple structural model of liquidity constraints and time preference, and find strong evidence for the existence of credit constraints.

This paper proceeds as follows. In Section 2, we provide a brief literature review and conceptual framework. In Section 3, we describe the experimental design. In Section 4, we discuss the reduced-form evidence provided by our data. In Section 4, we propose and estimate a simple structural model of time preferences and credit constraints. Finally, in Section 6 we conclude.

2 Literature

Under-investment in welfare-enhancing or profitable technologies is thought to be a commonplace problem in developing countries. There are a variety of products, ranging from modern fertilizer to efficient cookstoves, that many poor people do not purchase, in spite of what would appear to be large benefits. While there are many potential explanations for this seeming underinvestment, in this section we focus on research related to time preference, liquidity constraints and consumers' lack of information on the effectiveness of the new product.

Tarozzi and Mahajan (2011) (TM) and Dupas and Robinson (2013) (DR) both examine the relationship between non-standard time preferences and health investments, TM studying loans for bednet purchases in Orissa, India, and DR studying commitment savings for subject-chosen health products in Kenya. We highlight two differences between our study and these. First, we directly compare behavior under savings and borrowing. This is useful for policymakers as well as for understanding behavioral mechanisms. Second, we measure effects on WTP rather than share purchasing at a single price (TM) or total health investment or savings accumulated (DR), so our results are informative for pricing policy.

While the relationship between liquidity constraints and consumption has a long history (Deaton, 1991), recent research in developing countries has focused on the production side. In addition to the large literature on microfinance, Banerjee and Duflo (2005) review estimates of returns to capital in small-scale productive activities in developing countries and finds them varying drastically within country. de Mel et al. (2008) find the average real return to capital (distributed in a randomized experiment in Sri Lanka) is substantially higher than market interest rates (at least for male entrepreneurs). They interpret their results as largely consistent with liquidity constraints. Banerjee and Duflo (2014) study a change in the rules defining what firms are eligible for earmarked credit from Indian banks. They estimate very high rates of return for firms that gained easier access to credit due to the change in rules, suggesting that liquidity constraints are binding even for relatively large, formal enterprises. However, several recent studies have found that microfinance lending programs have had at most small effects on most economic outcomes of interest (Banerjee et al., 2015).

Both consumers and producers are likely to be uncertain about the returns to a new technology, and experimentation can be risky (Foster and Rosenzweig, 1995). Recent empirical research on the relationship between experimentation and adoption has been mixed. Dupas (2014) finds that short-run subsidies increase long-run adoption of insecticide-treated bednets in Kenya. Levine and Cotterman (2012) found that adding a free trial, time payments, and the right to return increased uptake of an efficient charcoal stove from 5 percent to 45 percent. That study showed that either the free trial or time payments increased uptake by about half the total effect, but did not identify what barriers the sales offers overcame. However, experimentation can also lead to decreased adoption if consumers find the product inconvenient or unpleasant to use (Mobarak et al., 2012; Luoto et al., 2012).

There is substantial evidence that many people have present bias, meaning that their subjective discount rate for short-term decisions today is higher than their subjective discount rate for short-term decisions in the future. The most common formulation within economics is a model that assumes there is an exponential discount rate δ for most decisions, but an additional present bias discount rate $\beta < 1$ for all future periods (Laibson, 1997; O'Donoghue and Rabin, 1999).

3 Experimental Design and Data Collection

3.1 Context and Object of Sale

The target population consists of poor households with young children in slums of Dhaka, Bangladesh. This population is of particular interest because of the low-quality piped water in these neighborhoods and high burden of water-borne disease, both generally and among young children.

The core intervention is the offer for sale of a long-lasting ceramic water filter with a retail price of approximately BDT 2100. We are interested in the demand for water filters because in previous research in this population found a strong distaste for chlorine-based treatment: WTP is low, and use is low even when provided free (Guiteras et al., 2015). The ceramic filter was popular in consumer testing in a similar population elsewhere in Dhaka, although few households purchased the filter at the break-even price.

We begin with a simple household survey to collect basic data on demographics, socioeconomic status, risk preferences and recent episodes of water-borne disease. We then conduct a marketing meeting to promote the filter to the subject households and explain the dangers of local water. The promotional message draws on our previous work in Dhaka with similar compounds, and combines both a positive health message as well as a message emphasizing disgust at ingesting fecal matter in unfiltered water. We inform the subject of the possible payment plans that might be offered in the sales visit and instruct her to think how much she (and possible the household) would pay for each option. We also explain the modified Becker-DeGroot-Marschak (BDM) mechanism (Becker et al., 1964), described below, that we use to elicit WTP. To increase understanding we practice BDM using real goods and money.

Two weeks later, we return for a sales visit, in which we use BDM to obtain the households' WTP under several different payment plans, listed in Table 1 and described at greater length below. There are two basic time-payment types, loans and dedicated savings / layaway plans. The plans also differ in duration and whether the first payment is made immediately or with a one-month delay. The subject will randomly receive an offer for which she has already stated whether she would accept or reject. If she purchases a filter under non-delay plans, the filter will be delivered by the end of the next day and payments begin. Thereafter payments are collected monthly and the collections officer records at each visit if the filter has been used recently.

[Table 1 about here.]

3.2 Willingness-to-pay data and the Becker-DeGroot-Marschak mechanism

To obtain precise data on WTP, for each offer type, we conduct a series of Becker-DeGroot-Marschak mechanism (BDM) procedures, one for each offer type. In the standard implementation of BDM for a single good, the subject states her maximum WTP ("bid"). If there was only one offer type the bid is then compared against a random price ("offer"). If her bid is less than the offer price, she does not purchase the filter. If her bid is greater than or equal to the offer price, she purchases the filter at the offer price. Under fairly weak assumptions, her best strategy is to bid her maximum WTP truthfully. To obtain a subject's WTP for a number of different offer types, we adapt BDM into a two-stage procedure. In the first stage, we obtain the subject's bid for each of the 8 offer types shown in Offer types, giving us a vector of WTP amounts for each household $\text{WTP}_h = (\text{WTP}_{h,1}, \ldots, \text{WTP}_{h,8})$. Then, in the second stage, we randomize one offer type for which the BDM draw is actually taken. That is, a random offer type $t \in \{1, \ldots, 8\}$ is chosen, we draw a random offer price $p_{h,t}$, and proceed as in a single-item BDM: if $p_{h,t} \leq \text{WTP}_{h,t}$, the household receives the filter and pays $p_{h,t}$; if $p_{h,t} > \text{WTP}_{h,t}$, the household cannot buy the filter. One disadvantage of our implementation was that, after extensive piloting, we found that it was necessary to provide participants with the minimum and maximum possible lottery prices, and to cap this range at the approximate break-even retail price of BDT 2100. This was necessary to improve participant understanding and to maintain a sense of fairness. However, it does mean that our WTP measure is censored, in that if a household has a very high WTP, we will observe only the top-coded value of BDT 2100. Because of this censoring, we will focus on quantile (median) estimates for demand data.

3.2.1 Offer Types

Table 1 lists the main offer types. The simplest offer is a lump sum paid on delivery (either the same day or the next day). Next, we offer loans which begin immediately and involve 3, 7 and 12 monthly payments.[2] A parallel set of plans (3 and 7 payments) are for layaway, in which households make regular payments into a dedicated lockbox, according to the payment schedule, until they have accumulated the offer amount. These plans are soft commitments: even though the lockbox key is held by ICDDR,B, the lockbox itself remains with the household and the savings will not be confiscated if the household "defaults" by not following through on its commitment. At the time the household is scheduled to make a deposit, field staff visit to confirm that the deposit has been made. Households also have the

[2] The prompts for BDM bids are framed in terms of the monthly payment rather than the total (e.g. "three monthly payments of BDT 400," rather than "BDT 1,200 over three months." However, we also provide subjects with the total amount implied by their monthly payments if they ask, as most pilot subjects have done. The BDM draw, which determines the allocation and total price paid, is in terms of the total amount, which is then converted back into monthly payments for the relevant payment plan. We conduct the BDM draw in terms of the total amount for operational simplicity – otherwise, surveyors would have to carry separate price envelopes for each offer.

option to "deposit" their money with the field staff in exchange for a receipt.[3] The plans were presented to the households in one of four possible orders differing in (a) whether loan or layaway plans came first, and (b) the different lengths were arranged from shortest to longest or the reverse.

3.2.2 Randomized Treatments

Two randomized treatments were given to our sample. They were each given to half of the population, orthogonally, so that a quarter received both treatments, a quarter received neither, and a quarter each received only one of the treatments.

Randomized Treatment 1: Free trial. The first treatment is a two-week free trial, giving households an opportunity to learn to use the filter and to confirm whether ease of use, taste, and other characteristics are acceptable. These households received the filter at the marketing meeting and have if for use until the sales meeting. For risk-averse consumers one would expect the free trial to increase WTP (Levine and Cotterman, 2012), although there are counterexamples (Mobarak et al., 2012; Luoto et al., 2012).

Randomized Treatment 2: Money-back guarantee or rent-to-own. One potential barrier to adoption is that households may incur income, health or consumption shocks that ex-post mean that money spent on a filter would have been better spent on something else. To test whether this is an important determinant of WTP, we randomize whether the loan offer gives the household the option to return the product for a partial or full refund up to a year from the sales date. With no refund, a time payment plan is similar to a "rent-to-own" scheme, in which the subject risks losing only accumulated payments, rather than the full lump sum.

[3] An alternative approach to identifying liquidity constraints is to give the subject the good in question and perform a reverse BDM in which the subject reveals the minimum amount she is willing to accept (WTA) in exchange for the good. The idea is to remove the liquidity constraint so that any variation in minimum WTA across payment plans could be attributed to time preference. This was not successful in piloting, for two main reasons. First, a large majority of pilot subjects stated that they would not accept any amount in exchange for the filter. We interpret this as some combination of a please-the-implementer effect and an endowment effect, with the former being more likely given that subjects even refused amounts higher than the going retail price. Second, reverse time payment plans were not perceived as credible by the subjects – many were skeptical that we would return multiple times over several months to give them money.

With a full or partial refund, the time payment plan comes to resemble the layaway plan, but with the household receiving the flow of benefits from the product while payments are being made.

3.3 Data Collection and Summary Statistics

We conducted a baseline household survey at the time of the marketing meeting to collect basic data on demographics, socioeconomic status, risk preferences and recent episodes of water-borne disease. A final end line survey is conducted 6 months later. Main survey collections occurred from September of 2012 to June of 2013.

Table 2 shows the means of demographic characteristics in our sample. It also reports the difference across the free trial treatments (mean of those with free trial minus those without) and across the guarantee arms. There do not appear to be large differences across these treatments. Our measure of income is fairly noisy because of several outliers.

[Table 2 about here.]

4 Reduced-Form Evidence

4.1 Time Payments

The most salient result from the study is that time payments dramatically increase WTP. Figure 1 compares the share of households willing to purchase the filter given a lump-sum offer with the share using the household's maximum bid across offers. All figures deflate cash flows by the average local business rate (13 percent). Time payments increase demand by 30 percentage points or more at all prices above BDT 700. Median WTP increases from BDT 755 to 1530 for a 12-month loan. Figure 2 examines differences in individual household WTP. Among households that are not censored (i.e. (i) do not have all bids at the top bid amount, and (ii) express some positive WTP for any offer), WTP increases for most households, with a median increase of BDT 600 (min. 0, IQR 200-1100, max. 1900).

[Figure 1 about here.]

[Figure 2 about here.]

Even a short-term (3-month) loan significantly increases demand, which continues as the term of the loan lengthens. This can be seen in Figure 3, which plots the share of subjects willing to purchase given each loan offer.

[Figure 3 about here.]

4.2 Discounting and Layaway Plans

Surprisingly, WTP given time payment layaway plans are almost identical to loans. Figure 4 shows that the demand curves lie almost on top of each other, and Figure 5 shows that nearly all households have identical WTP for loans and layaway plans of the same duration. This suggests that households do not discount health benefits in the same way as utility from general consumption.[4] Anchoring effects or later inattention seem unlikely as the (randomized) order of the offers did not significantly affect the WTP prices.

[Figure 4 about here.]

[Figure 5 about here.]

4.3 Randomized Treatments

The results from our randomized treatments are somewhat less striking. In neither case (free trial, Figure 6; guarantee, Figure 7) do we see strong evidence for an increase in demand.

[Figure 6 about here.]

[Figure 7 about here.]

[4] It is possible that the layaway plan provides unique benefits, such as access to a lockbox to help save,

4.4 Ability to Save

A simple measure of the ability of households to save is to compare the up-front lump-sum payment with the one-month delayed plan. We find that on average the delayed plan allows households to increase the WTP by 20.6%. This increase by allowing an extra month is roughly consistent in terms of magnitudes with the 45.5% from having two extra months in the 3-month plan. While a large fraction appear to save, we find that 53.7% of the sample does not have a higher WTP for the delayed plan than lump sum implying that the majority have difficulty saving.

4.5 Ability to Borrow

Following Attanasio et al. (2008) and Karlan and Zinman (2008), we can test for the presence of credit constraints by measuring whether household WTP (in net present value terms) increases with the length of the loan. If a household does not face credit constraints, i.e. can borrow or save at a prevailing market interest rate without restriction, then the household's WTP would be equal (in net present value) across all payment plans, regardless of the household's time preferences. Therefore, if the household's WTP increases with the length of the loan maturity, this is evidence of credit constraints.[5] We test for this effect in two ways, first with a constant interest rate of 26 percent (twice the average business business loan rate), and second using individual-specific interest rates derived from the 3-installment plan (compared to lump-sum amount). The regression results in Table 3 show that the maturity length has a consistently significant effect on the present value of the loan size implying the presence of liquidity constraints.

[Table 3 about here.]

that offset any utility loss to delaying consumption of the filter. However, it is unlikely that these effects would almost exactly cancel out for almost all households.

[5]Friction costs associated with meetings would bias our results in the opposite direction.

Additional reduced-form evidence for credit constraints can be found by looking at the 1-month delay and 3-month plan. With perfect liquidity, we would observe that households' WTP are roughly equal in net present value terms:

$$\frac{p_{1dly}}{R} \approx p_{3m} + \frac{p_{3m}}{R} + \frac{p_{3m}}{R^2},$$

where R is the repayment rate (one plus the interest rate). Then

$$p_{1dly} = p_{3m}\left(R + 1 + R^{-1}\right).$$

With monthly interest rates small (i.e. $R \approx 1$), then the nominal totals of both (p_{1dly} and $3p_{3m}$) should be approximately the same. In reality, the 3-month loan value is 28.9% higher, which is economically significant.

5 Estimating Preferences and Constraints

Our reduced-form empirical analysis provides strong evidence that micro-loans and micro-savings significantly increase WTP. To assess the relative importance of financial constraints or time preference in explaining this fact, we turn to a simple structural model incorporating time preferences and credit constraints. We designed our experiment to provide clean identification of preferences versus constraints by using within-household differences in WTP between micro-loans and micro-savings. However, the surprising fact that subjects' WTP for micro-loans and micro-savings plans were almost identical means that this strategy is no longer viable. Instead, to estimate the structural model, we exploit differences in households' WTP across loan offers with different loan duration and different timing of payments.

5.1 Utility

We assume that the household maximizes utility over a finite horizon. We assume that a family intends to pay to continue using the filter after its lifetime, so all financial transactions with regards to the filter should be taken care of during its lifetime. As a single filter-element usually lasts a family 12 months, we choose this as the planning horizon which we divide into monthly periods.[6] Given the data on lay-away plans, we assume there is no discounting of health-related utility, but there is discounting over utility from general consumption. We assume that this discounting is exponential.

The household's WTP for a particular plan is the highest total price such that the household is indifferent between purchasing at that price and not purchasing at all. Equivalently, it is the price such that the household's stream of lowered utilities from non-health activities is equal to the utility gain from having the filter. Let p be the monthly price associated with the plan total. If the household makes final payments of $\{\bar{p}_t\}_{t=0}^{11}$ over the course of 12 months (we allow for borrowing from other sources to smooth out-of-pocket payments so that $\bar{p}$ may be payments to us or to other sources) then

$$\sum_{i=0}^{11} \frac{1}{(1+\delta)^t}[u(y) - u(y - \bar{p}_t)] = B \tag{1}$$

where u is the utility function over non-health activities, y is monthly income (assumed for simplicity to be constant), and B is the present-value of the filter. Taking a second-order approximation of the difference yields

$$\sum_{i=0}^{11} \frac{1}{(1+\delta)^t}\left[\bar{p}_t + \frac{1}{2}\eta\bar{p}_t^2\right] = w \tag{2}$$

where $\eta = -u''(y)/u'(y)$ measures utility curvature (the coefficient of absolute risk aversion) and $w = B/u'(y)$ is the value of the filter normalized by the marginal utility of income.

[6]The 12 months is from the sales meeting. Since we assume no discounting of health benefits, whether the family receives the filter initially or later in a lay-away plan is unimportant in the model.

5.2 Credit environment

Given a monthly repayment rate R_1 (one plus the interest rate) an actuarially fair loan contract will set the present-value of the monthly payments (p) to be equal to the amount borrowed (b).

$$\sum_t \frac{1}{R^t} p = b$$

Rather than build credit constraints from microfoundations, we take a reduced-form approach and model credit constraints as a nonlinear cost-of-borrowing function, which for simplicity we approximate as a quadratic.

$$\sum_t \frac{1}{R^t} p = \tilde{R}_0 + b + \tilde{R}_2 b^2,$$

where $\tilde{R}_0 + \tilde{R}_2 b^2$ is the penalty for being constrained. Putting all the repayment terms in similar units, define $q(b)$ as the repayment amount of a one period loan and define $R_0 = \tilde{R}_0 \cdot R$, $R_1 = R$, and $R_2 = \tilde{R}_2 \cdot R$:

$$q(b) = R_0 + R_1 b + R_2 b^2 \tag{3}$$

This extends the standard transaction-cost model of loans (e.g., Helms and Reille, 2004) by adding a quadratic term. Adding the quadratic term is attractive for several reasons: (a) the observed repayment rate $q(b)/b$ is not necessarily declining in b; (b) it is better able than the simple linear model to approximate situations where there are fixed limits on the amounts a household can borrow (where the borrowing costs function would be vertical). Possible micro-foundations for a quadratic-shaped borrowing cost function include (a) it incorporates the idea that with multiple sources of limited funds a household will choose the cheaper options first, and (b) if lender's expect that larger loans are less likely to be paid back they will charge higher effective interest rates for larger loans. We assume that if the borrower wants to repay a loan at a different date (or via installments), then they just need to repay the lender the same net-present value as the $q(b)$ function.

We assume that each household evaluates whether to borrow through outside lenders in order to smooth consumption. We assume that if they borrow from outside lenders then they borrow from outside a fixed monthly amount while in our plan, and that after our time-payments plan is over they repay another fixed monthly amount back. As most money lending is this environment is short-term we set the repayment length to outside moneylenders at 3 months. For example, when considering how much they would pay us monthly (p) for a three-month plan, they think about borrowing possible monthly amounts b (so their consumption only drops by $\bar{p}_{t<3} = p - b$) for three months and then paying back a monthly amount $\bar{p}_{t\geq 3} = e$ for three months (where e is determined by $q(b)$). Outside credit would be most attractive for shorter plans as there is more opportunity for smoothing.

5.3 Identification

To show how the profile of WTP identifies the parameters, we show in Figure 8 the effects of varying each parameter individually. R_0 has a discontinuous effect on the profile. As it increases it make borrowing not worthwhile at earlier and earlier durations. Higher values of R_1 depress the WTP profile, though slightly more at shorter durations. Higher values of R_2 dampen WTP for shorter bids (as households are likely to borrow more per month) and make WTP profile more concave. Higher values of δ have minimal effects at short durations (even if they borrow outside that if finished up soon) and has a maximal effect at 7 months as this is where they are paying for quite a while and they still may borrow. Utility curvature (η) and the monetized filter value (w) are not well identified separately in the current model; they both raise the WTP profile quite evenly. Only one of these is estimated at the individual level.

[Figure 8 about here.]

5.4 Estimation

The parameters of our model are $\{B, \delta, \eta, R_0, R_1, R_2\}$, where $\{B, \delta, \eta\}$ describe a household's preferences (valuation of the filter, discounting and utility curvature, respectively) and $\{R_0, R_1, R_2\}$ describe the household's credit environment. In principle, all of these parameters vary at the household level, but we do not have rich enough data to robustly estimate all parameters at the household level given the number of parameters is close to the number of individual WTP observations. We divide the parameters into population-level and individual-level groups. For individual parameters we choose $\omega_i = \{B_i, R_{2i}, \delta_i\}$, as we believe these are likely to vary the most in the population and allow us to be able to distinguish basic preferences from a measure of credit-constraints. Population parameters are then $\alpha = \{R_0, R_1, \eta\}$. For each household i, our data consist of the household's bids on each of the M offers, $p_i = (p_{i,1}, \ldots, p_{i,M})$. We estimate the model in an iterative two-step process. To build intuition, we first describe this process as if p_i were not censored (recall that bids were top-coded at the approximate break-even price of BDT 2100), and then describe our modification to account for censoring.

First, note that equation 2 provides a mapping from parameters (ω_i, α) to predicted WTP $\hat{p}_{i,m} = (\hat{p}_{i,1}, \ldots, \hat{p}_{i,M})$. For each plan this is maximum of the predicted WTP assuming no other borrowing and of the WTP assuming outside borrowing. For the no borrowing condition we solve Equation 2 for when monthly forgone consumption is the monthly filter price and there are no effects after the plan is finished. For the borrowing condition, we use the aggregate repayment equation plus the first order condition from Equation 2 to determine the maximum WTP.

The intuition for our estimation strategy is to choose the value of the parameters that minimize the difference between actual bids and predicted bids. We start with initial guesses for the population-level parameters $\alpha^{(0)}$. We then iterate the following procedure:

1. Given current population level estimates $\alpha^{(j)}$, we choose household-specific parameters

$\omega_i^{(j+1)}$ to maximize:

$$\Gamma_i\left(\omega_i|\alpha^{(j)}\right) = \sum_m \Psi_{im}(\omega_i, \alpha^{(j)}, p_{im})\ \forall i \tag{4}$$

$$\Psi_{im}(\omega_i, \alpha^{(j)}, p_{im}) = \ln \frac{1}{\sigma}\phi\left(\frac{p_{im} - \hat{p}_m(\omega_i, \alpha)}{\sigma}\right) \tag{5}$$

where Ψ_{im} is the log-likelihood of the parameters yielding the stated WTP for person i and plan m, the summation is over plans m, p_{im} is the observed WTP for individual i for plan m, and $\hat{p}_m(\omega_i, \alpha)$ is the predicted WTP for plan m given the parameters (i.e. the maximum WTP implied by Equation 2 given parameters $\{\omega_i, \alpha^{(j)}\}$).

2. Given current individual-specific parameters for the sample $\omega^{(j+1)} = \left\{\omega_i^{(j+1)}\right\}_{i=1}^N$, we choose population level parameters $\alpha^{(j+1)}$ to maximize the sample log-likelihood:

$$\begin{aligned}\Gamma\left(\alpha^{(j+1)}|\omega^{(j+1)}\right) &= \sum_i \Gamma_i\left(\omega_i|\alpha^{(j)}\right) \\ &= \sum_i \sum_m \Psi_{im}(\omega^{(j+1)}, \alpha^{(j+1)}, p_{im})\end{aligned}$$

where the outer summation is over subjects i.

3. We repeat steps 1-2 until convergence.

In each step we estimate the parameters of interest via maximum likelihood. With current candidate parameters and the parameters taken as given in each round, we predict the WTP for each individual. The WTP is the highest price that allows the family through some amount of borrowing to be indifferent between purchasing the filter at the price and having no filter. We solve then for the amount of borrowing from outside money-lenders that maximizes the WTP while keeping the family indifferent. We can then determine the error between predicted and observed WTPs which we assume is normally distributed. We weight deviations between predicted and observed bids equally for each offer.

As observed WTPs are censored from above, we adjust the likelihood function in a

Tobit-style fashion, replacing $\Psi_{im}(\omega_i, \alpha^{(j)}, p_{im})$ from Equation 5 with

$$\Psi_{im}(\omega_i, \alpha^{(j)}, p_{im}) = 1\{p_{im} < p_{top}\} \cdot \ln\left[\frac{1}{\sigma}\phi\left(\frac{p_{im} - \hat{p}_m(\omega_i, \alpha)}{\sigma}\right)\right] + 1\{p_{im} = p_{top}\} \cdot \ln\left[1 - \Phi\left(\frac{p_{im} - \hat{p}_m(\omega_i, \alpha)}{\sigma}\right)\right], \tag{6}$$

where p_{top} is the top-coded amount and $1\{\cdot\}$ is the indicator function.

During estimation, parameters are constrained so that they are always positive.

5.5 Structural Results

The values of the structural parameters are reported in Table 4, and distributions for individual-level parameters are shown in Figures 9 - 11. All the parameters point estimates are significant at the 1% level except for R_2

[Table 4 about here.]

For the credit environment, all parameters are estimated significantly. R_1 is equivalent to a 14.9% annual interest rate. The median R_2 is close zero. Though the point estimate is not significant, the share (about one half) of those with non-zero values is significant. For a loan the size of a full priced filter (BDT 2100), the penalty from R_2 is equal to BDT 263, roughly twice the fixed cost of borrowing (R_1). These parameters imply that borrowing for both small amounts and large amounts are economically significant.

We estimate quite a high discount rate for non-health utility (18.0%). Intuitively, this is driven by the fact that total WTP for 12-months is significantly higher than for 7-month. Credit constraints are mostly relaxed by 7-months (given plausible ranges of utility curvature) so this large increase determines the discount rate. Our estimate for the monetized utility value of the filter is close to the full market price. Finally, the estimated utility curvature is estimated at .01.

Using these estimated parameters, we look at the effect of counterfactuals for eliminating credit constraints. We find that by eliminating credit constraints (setting $R_0 = R_2 = 0$) that the median WTP for a 3-month payment increases by 347, which is equivalent to an increase in the value of the filter by 53%.

6 Conclusion

In this paper we show results of a detailed study of the willingness-to-pay (WTP) for household water filters under a variety of payment plans. Demand for the filter is increased dramatically by having a longer payment plans. We find reduced form evidence for liquidity constraints, both for saving and borrowing. Surprisingly, we find that households are indifferent between receiving the filter now or at the end of the payment plan (up to 12 months later) implying discounting for health related goods that is close to zero. Using a structural model we find evidence of high discounting (for non-health utility) and credit constraints. Using the estimated results, we find a counterfactual setting that removing credit constraints is equivalent to increasing the household's value of the filter by 53%.

References

Abdul Latif Jameel Poverty Action Lab. The price is wrong. Cambridge, MA, April 2011. URL `http://www.povertyactionlab.org/publication/the-price-is-wrong`.

Amrita Ahuja, Michael Kremer, and Alix Peterson Zwane. Providing safe water: Evidence from randomized evaluations. *Annual Review of Resource Economics*, 2(1):237–256, October 2010. doi:10.1146/annurev.resource.012809.103919.

Nava Ashraf, Dean Karlan, and Wesley Yin. Tying odysseus to the mast: Evidence from a commitment savings product in the Philippines. *The Quarterly Journal of Economics*, 121 (2):635–672, 2006. doi:10.1162/qjec.2006.121.2.635.

Nava Ashraf, James Berry, and Jesse M. Shapiro. Can higher prices stimulate product use? Evidence from a field experiment in Zambia. *American Economic Review*, 100(5): 2383–2413, December 2010. doi:10.1257/aer.100.5.2383.

Orazio P. Attanasio, Pinelopi Koujianou Goldberg, and Ekaterini Kyriazidou. Credit constraints in the market for consumer durables: Evidence from micro data on car loans. *International Economic Review*, 49(2):401–436, may 2008. doi:10.1111/j.1468-2354.2008.00485.x.

A. V. Banerjee and E. Duflo. Do firms want to borrow more? Testing credit constraints using a directed lending program. *The Review of Economic Studies*, 81(2):572–607, feb 2014. doi:10.1093/restud/rdt046.

Abhijit Banerjee, Dean Karlan, and Jonathan Zinman. Six randomized evaluations of microcredit: Introduction and further steps. *American Economic Journal: Applied Economics*, 7 (1):1–21, jan 2015. doi:10.1257/app.20140287.

Abhijit V. Banerjee and Esther Duflo. Chapter 7 growth theory through the lens of development economics. *Handbook of Economic Growth*, pages 473–552, 2005. doi:10.1016/s1574-0684(05)01007-5.

Gordon M. Becker, Morris H. Degroot, and Jacob Marschak. Measuring utility by a single-response sequential method. *Behavioral Science*, 9(3):226–232, 1964. doi:10.1002/bs.3830090304.

James Berry, Gregory M. Fischer, and Raymond P. Guiteras. Eliciting and utilizing willingness to pay: Evidence from field trials in northern Ghana. July 2015. URL `http://ibread.org/bread/working/459`.

Suresh de Mel, David McKenzie, and Christopher Woodruff. Returns to capital in microenterprises: Evidence from a field experiment. *Quarterly Journal of Economics*, 123(4): 1329–1372, Nov 2008. doi:10.1162/qjec.2008.123.4.1329.

Angus Deaton. Saving and liquidity constraints. *Econometrica*, 59(5):1221, Sep 1991. doi:10.2307/2938366.

Pascaline Dupas. Health behavior in developing countries. *Annual Review of Economics*, 3 (1):425–449, 2011. doi:10.1146/annurev-economics-111809-125029.

Pascaline Dupas. Short-run subsidies and long-run adoption of new health products: Evidence from a field experiment. *Econometrica*, 82(1):197–228, 01 2014. doi:10.3982/ECTA9508.

Pascaline Dupas and Jonathan Robinson. Why don't the poor save more? Evidence from health savings experiments. *American Economic Review*, 103(4):1138–1171, Jun 2013. doi:10.1257/aer.103.4.1138.

Andrew D. Foster and Mark R. Rosenzweig. Learning by doing and learning from others: Human capital and technical change in agriculture. *Journal of Political Economy*, 103(6): 1176–1209, 1995. doi:10.1086/601447.

Raymond P. Guiteras, David I. Levine, Stephen P. Luby, Thomas H. Polley, Kaniz Khatun-e Jannat, and Leanne Unicomb. Disgust, shame and soapy water: Tests of novel interventions to promote safe water and hygiene. August 2015. URL `http://papers.ccpr.ucla.edu/abstract.php?preprint=1092`.

Brigit Helms and Xavier Reille. Interest rate ceilings and microfinance: The story so far. Technical report, CGAP, 2004. URL `http://www.cgap.org/publications/interest-rate-ceilings-and-microfinance-story-so-far`.

Alaka Holla and Michael Kremer. Pricing and access: Lessons from randomized evaluations in education and health. Working Paper 158, Center for Global Development, 2009. URL `http://www.cgdev.org/content/publications/detail/1420826/`.

Dean S Karlan and Jonathan Zinman. Credit elasticities in less-developed economies: Implications for microfinance. *American Economic Review*, 98(3):1040–1068, may 2008. doi:10.1257/aer.98.3.1040.

David Laibson. Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2):443–477, 1997. doi:10.1162/003355397555253.

David I. Levine and Carolyn Cotterman. What impedes efficient adoption of products? Evidence from randomized variation in sales offers for improved cookstoves in Uganda. Working paper series, UC Berkeley: Institute for Research on Labor and Employment, March 27 2012. URL `https://ideas.repec.org/p/cdl/indrel/qt7qk8m53w.html`.

Jill Luoto, Nusrat Najnin, Minhaj Mahmud, Jeff Albert, M. Sirajul Islam, Stephen Luby, Leanne Unicomb, and David I. Levine. What point-of-use water treatment products do consumers use? Evidence from a randomized controlled trial among the urban poor in Bangladesh. *PLoS ONE*, 6(10), 10 2011. doi:10.1371/journal.pone.0026132.

Jill Luoto, Minhaj Mahmud, Jeff Albert, Stephen Luby, Nusrat Najnin, Leanne Unicomb, and David I. Levine. Learning to dislike safe water products: Results from a randomized controlled trial of the effects of direct and peer experience on willingness to pay. *Environmental Science & Technology*, 46(11):6244–6251, 2012. doi:10.1021/es2027967.

Ahmed Mushfiq Mobarak, Puneet Dwivedi, Robert Bailis, Lynn Hildemann, and Grant Miller. Low demand for nontraditional cookstove technologies. *Proceedings of the National Academy of Sciences*, 109(27):10815–10820, 2012. doi:10.1073/pnas.1115571109.

Ted O'Donoghue and Matthew Rabin. Doing it now or later. *American Economic Review*, 89 (1):103–124, Mar 1999. doi:10.1257/aer.89.1.103.

A. Tarozzi and A. Mahajan. Time inconsistency, expectations and technology adoption: The case of insecticide treated nets. Economic Research Initiatives at Duke (ERID) Working Paper 105, Duke, 2011. URL `https://ideas.repec.org/p/duk/dukeec/11-14.html`.

Figure 1: Demand: Time Payments vs. Lump Sum

(a) Levels

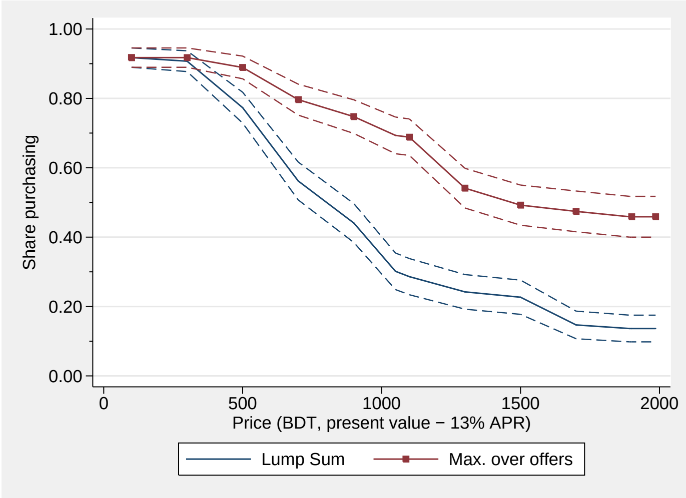

(b) Difference

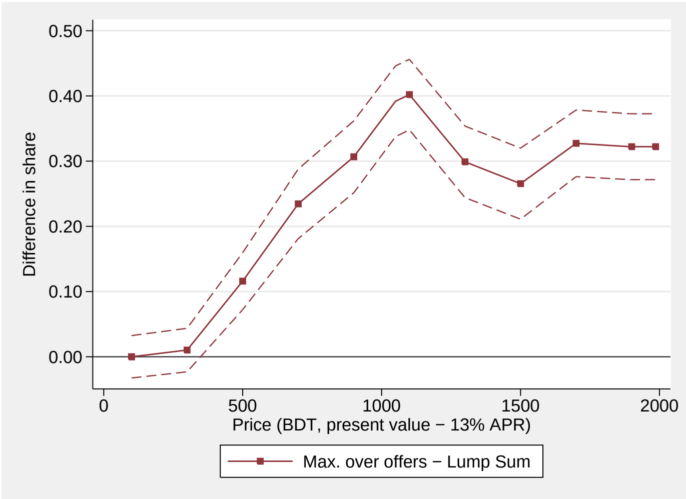

Notes: The top figure plots BDM demand curves, with 90% confidence bands, using households' maximum WTP across all offers (square markers) and households' maximum WTP for an immediate lump sum (no markers). The bottom figure plots the estimated differences (max. across all offers relative to lump sum). Pointwise inference from logit regressions (at prices BDT 100, 300, 500, . . . , max). Standard errors clustered at the compound level. 388 observations.

Figure 2: Distribution of household difference in WTP
Time Payments vs. Lump Sum

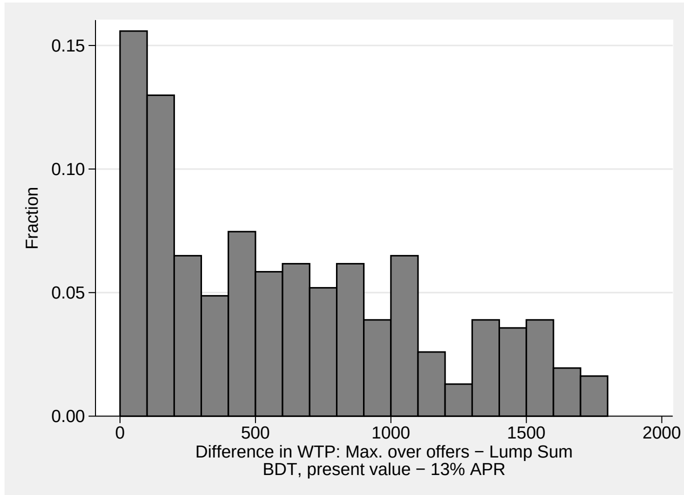

Notes: This figure plots the distribution of difference in household willingness to pay (WTP) under time payments (i.e. the maximum nominal amount across all loan and layaway offers) relative to an up-front lump-sum payment. We exclude 48 households that were top-coded, i.e. both their lump-sum and maximum time payment WTP were at the upper bound price, and the 32 households with zero WTP under all offers (including attriters and refusals), leaving 308 observations.

Figure 3: Demand Across Loan Offers

(a) Levels

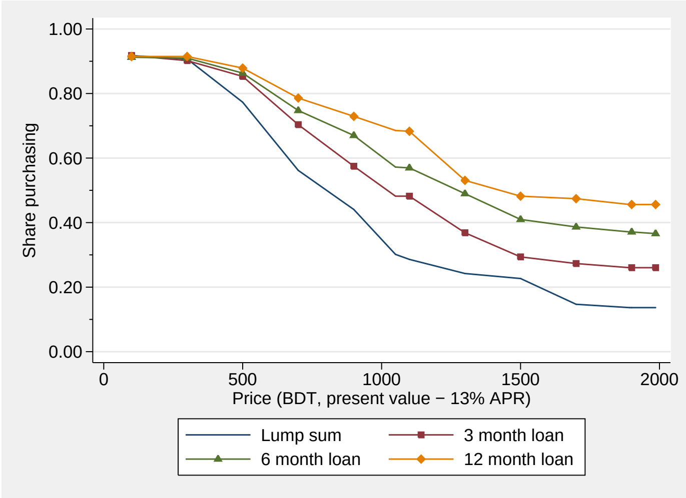

(b) Difference

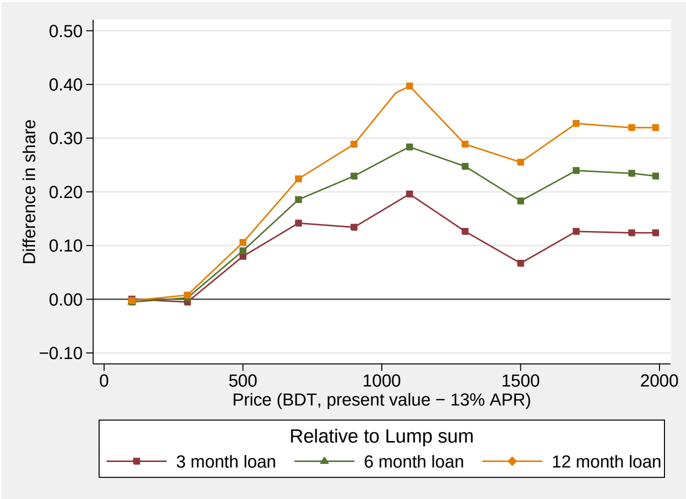

Notes: The top figure compares BDM demand curves across, with 90% confidence bands, loan offers: lump-sum (no markers), 3-month (square markers), 6-month (triangles) and 12-month (diamonds). The bottom figure plots the estimated differences for the three loan plans relative to lump-sum. Pointwise inference from logit regressions (at prices BDT 100, 300, 500, . . . , max). Standard errors clustered at the compound level. 388 observations.

Figure 4: Demand: Loans vs. Layaways

(a) 3 months

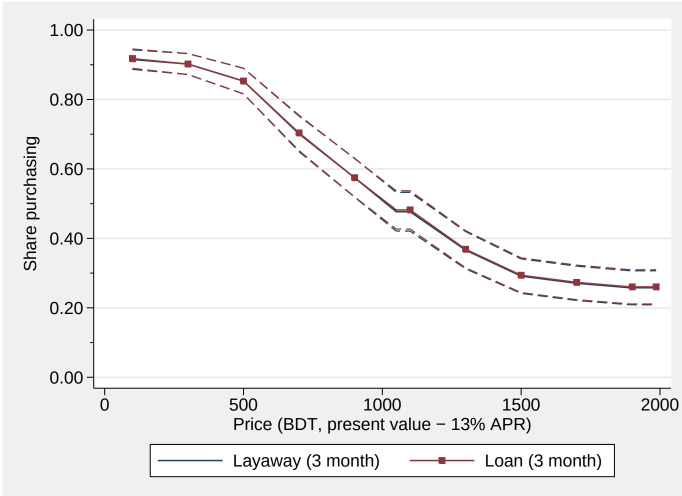

Notes: The figure compares BDM demand curves, with 90% confidence bands, for 3-month loans (square markers) and 3-month layaway plans (no markers). Pointwise inference from logit regressions (at prices BDT 100, 300, 500, . . . , max). Standard errors clustered at the compound level. 388 observations.

(b) 7 months

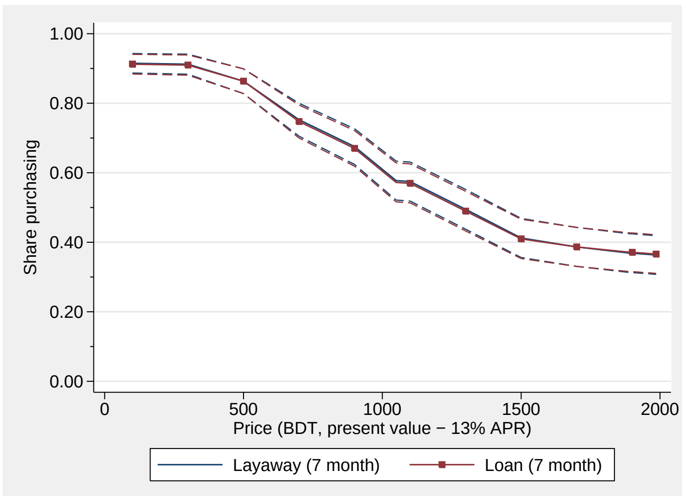

Notes: The figure compares BDM demand curves, with 90% confidence bands, for 7-month loans (square markers) and 7-month layaway plans (no markers). Pointwise inference from logit regressions (at prices BDT 100, 300, 500, . . . , max). Standard errors clustered at the compound level. 388 observations.

Figure 5: Difference in household WTP: Loans vs. Layaways

(a) 3 months

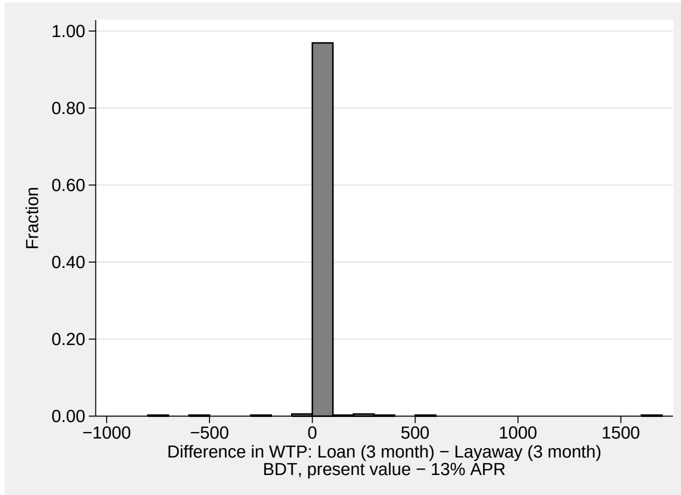

Notes: This figure plots the distribution of difference in household willingness to pay (WTP) for 3-month loans relative to 3-month layaway plans. We exclude 0 households that were top-coded, i.e. both their lump-sum and maximum time payment WTP were at the upper bound price, and 32 households with zero WTP for both offers (including attriters and refusals), leaving 356 observations.

(b) 6 months

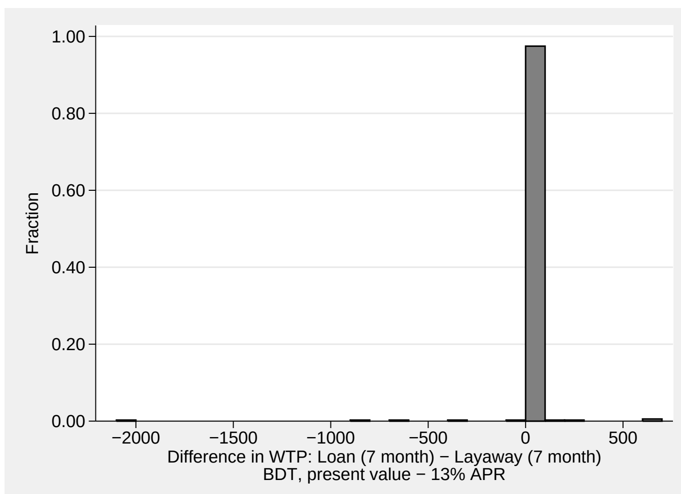

Notes: This figure plots the distribution of difference in household willingness to pay (WTP) for 7-month loans relative to 7-month layaway plans. We exclude 0 households that were top-coded, i.e. both their lump-sum and maximum time payment WTP were at the upper bound price, and 33 households with zero WTP for both offers (including attriters and refusals), leaving 355 observations.

Figure 6: Effect of Free Trial Treatment on Demand

(a) Lump-sum

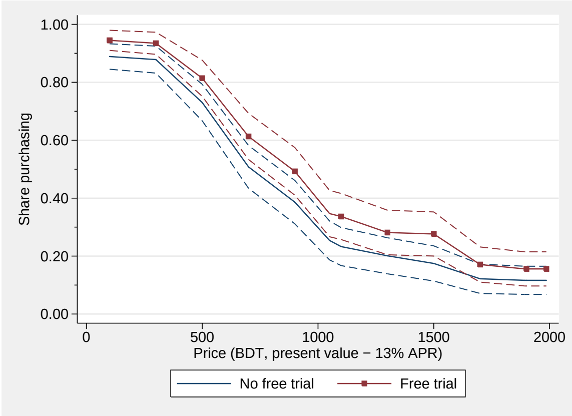

Notes: The figure compares BDM demand curves, with 90% confidence bands, between free trial and no free trial households, given a offer. Standard errors clustered at the compound level. 189 observations.

(b) 6-month loan

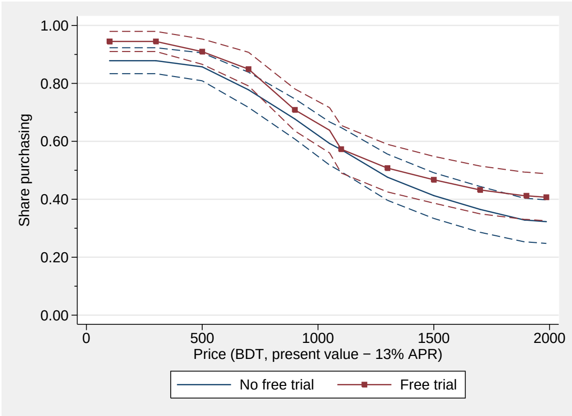

Notes: The figure compares BDM demand curves, with 90% confidence bands, between free trial and no free trial households, given a offer. Standard errors clustered at the compound level. 189 observations.

(c) Max. WTP across all offers

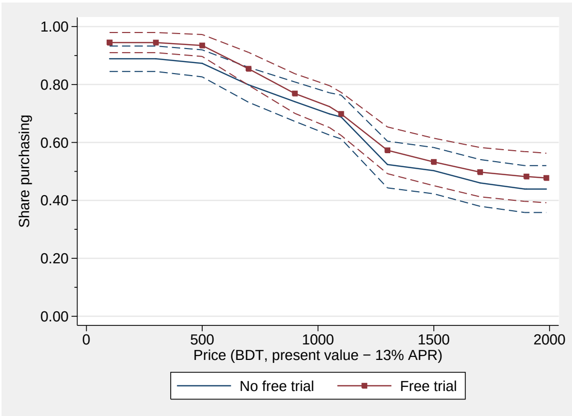

Notes: The figure compares BDM demand curves, with 90% confidence bands, between free trial and no free trial households, for the household's maximum (nominal) WTP across all offers. Standard errors clustered at the compound level. 189 observations.

Figure 7: Effect of Money-Back Guarantee on Demand

(a) Lump-sum

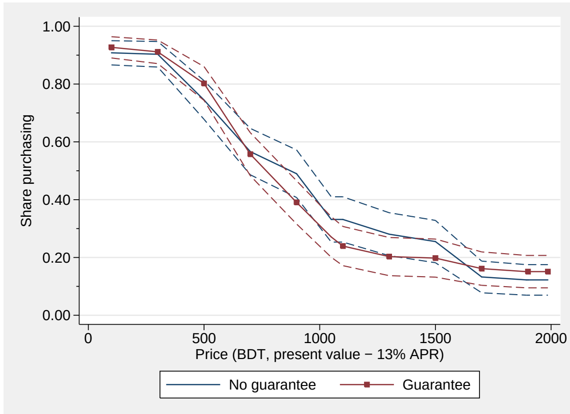

Notes: The figure compares BDM demand curves, with 90% confidence bands, between guarantee and no guarantee households, given a offer. Standard errors clustered at the compound level. 192 observations.

(b) 6-month loan

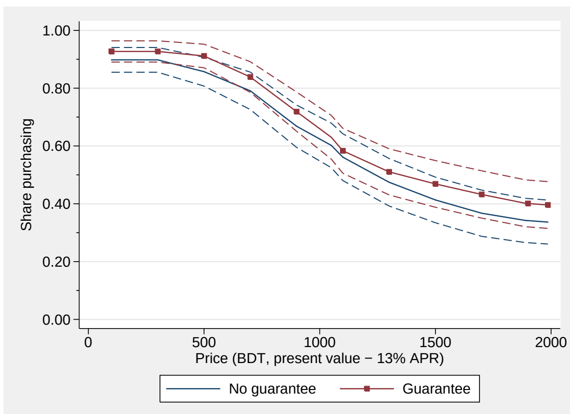

Notes: The figure compares BDM demand curves, with 90% confidence bands, between guarantee and no guarantee households, given a offer. Standard errors clustered at the compound level. 192 observations.

(c) Max. WTP across all offers

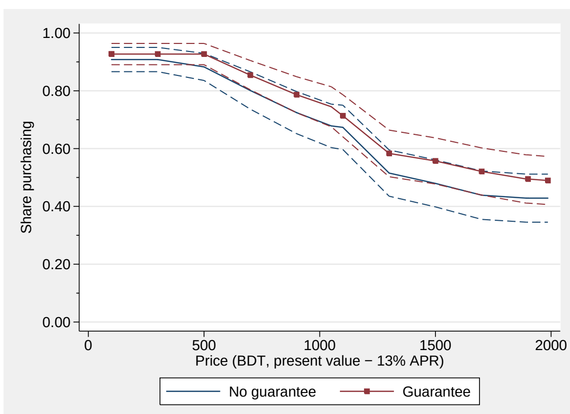

Notes: The figure compares BDM demand curves, with 90% confidence bands, between guarantee and no guarantee households, for the household's maximum (nominal) WTP across all offers. Standard errors clustered at the compound level. 192 observations.

Figure 8: Parameter effects on WTP profile

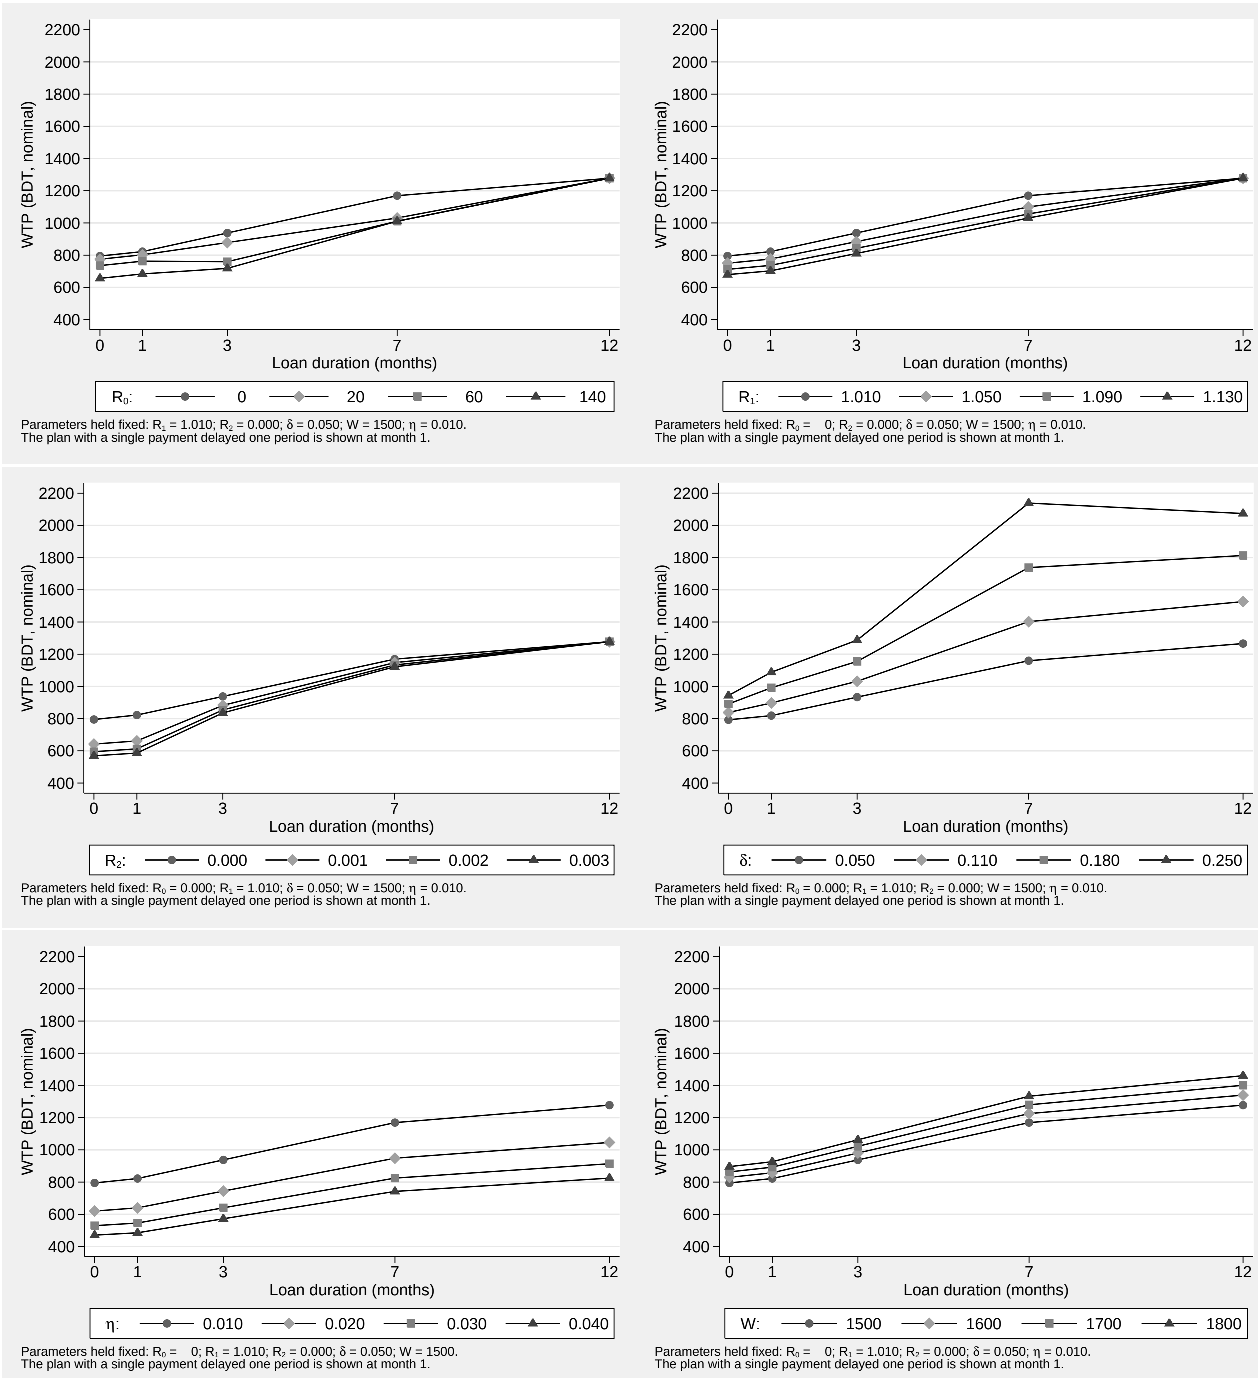

Table 1: Offer types

Offer type	Time of payment(s) (months)	Filter received (month)
Lump sum	0	0
3-month loan	0, 1, 2	0
3-month layaway	0, 1, 2	2
7-month loan	0, 1, 2, 3, 4, 5, 6	0
7-month layaway	0, 1, 2, 3, 4, 5, 6	6
12-month loan	0, 1, 2, . . . , 11	0
1-month delay	1	0
"75%, X, X"	0, 1, 2	2

In the "75%, X, X" offer, we fix the household's first payment at 75% of the maximum payment agreed to for a three-month loan, and the household then bids on the amount of the last 2 payments (X). The purpose is to provide variation between current and future payments to help identify present bias. The immediate (month=0) payment was due that day or next upon delivery.

Table 2: Randomization Check

	(1) Mean	(2) Free Trial(1-0)	(3) Guarantee(1-0)
Female	0.877	0.0317	0.0129
Age	31.67	-0.222	1.266
HH Education (years)	5.527	0.713*	0.0632
Married	0.909	0.0137	0.0137
HH size	4.188	0.0309	0.152
Rent > USD 27/month	0.779	-0.0539	0.0274
Has gas-line	0.992	0.000475	0.0157**
Water in compound	0.990	-0.0113	0.0118
HH Income (USD, monthly)	368.8	61.06	10.63

471 observations. Differences in means between those with and without the free trial (guarantee) are shown in column 2 (3).

* p<.1, ** p<.05, *** p<.01

Table 3: Maturity Effect

	(1) Loan PV (H)	(2) Loan PV (Ind)
Maturity	25.71***	22.62***
	(1.725)	(2.004)
constant	850.1***	783.8***
	(10.14)	(10.98)
N	1262	779
HH FEs	yes	yes

Standard errors in parentheses

The first model calculates the loan present value using a constant annual 26%. The second model calculates loan present values using individual rates derived from the 3-installment plan (compared to lump-sum). Top-coded values are excluded.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Estimated Structural Parameters

	(1) Estimates/std_errs
R0	116.9***
	(11.50)
R1	1.012***
	(0.0117)
R2 (median)	0.0000597
	(0.000139)
R2 (share positive)	0.509***
	(0.0288)
Monthly discount rate (median)	0.180***
	(0.0301)
Monetized filter value, taka (median)	1768.7***
	(174.7)
Utility curvature	0.0102***
	(0.000194)
Observations	291

Estimate of structural parameters from WTP data.
Bootstrap p-values (399 reps). * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

A Extra Figures

[Figure 9 about here.]

[Figure 10 about here.]

[Figure 11 about here.]

B Appendix - Secure Field Randomizations

Study participants were present during the randomized treatments such as selecting the "live" offer, selecting the BDM price for that offer, and the randomizations in the preference games. In order to minimize experimenter effects[7] we created a method that allows both the participant and the non-present investigator to ensure that the randomization was fair. This is an improvement on existing methods such as rolling dice (where it is difficult for the investigator to ensure fairness) or providing enumerators with sealed envelopes (where it is difficult for the participant to ensure fairness).

Procedure

The easiest way for two, distant, untrusting parties such as the participant and the investigator, to generate a random number $i \in \{0, ..., M-1\}$ is to have them secretly choose their own numbers (i_P and i_I) in this range, simultaneously exchange them, and have $i \equiv (i_P + i_I) \mod M$. In field experiments, it is costly to simultaneously exchange information between the participant and the investigator, so a potentially untrusted third-party is used, the enumerator[8]. This necessitates that the participant and the investigator use a "commitment scheme" so that the participant commits his i_P before knowing i_I and the investigator commits to i_I before knowing i_P.

The investigator uses an opaque, signed and sealed envelope (with the participant ID on it) with his i_I inside as his commitment. The participant commits i_P by the why the envelope is opened. For example, for $M = 6$ one could have an envelope like that in Figure 12. The envelope should only be opened across on the dashed lines and one can either cut above or below each number. We commit i_P by cutting below this number but above all the others (connecting with vertical cuts as necessary). This can be done with a small utility knife and a clipboard. [9]

[7] If participants suspect false randomization they may change their actions prior to randomization because (a) they attempt to change the outcome, (b) differences in reaction to risk and ambiguity (what hopefully was interpreted as risk might now be ambiguous), (c) general distrust of the experiment.

[8] In our experience unwelcome actions by the enumerator are usually low-effort such as doing things out of order or faking the outcome of a rolled die. We note here that our procedure does not prevent high-effort attacks such as presenting fake data to either the participant or the investigator. We thought this would be unlikely, so did not do the large amount of work required to solve this problem.

[9] An alternative is to (1) commit i_P by writing it on the envelope, (2) take a picture of that side of the envelope, (3) reveal i_I by opening the envelope such that the opening is entirely on the same side of the

[Figure 12 about here.]

Depending on what the randomization is used for, one might also want to use the commitment mechanism to ensure that additional information is recorded before the envelope is revealed. For example, in a BDM procedure one wants to record the bid before revealing the unknown price. Multiple sets of numbers can be printed on the tab-line (and cut around). Additionally, if one wants to ensure that certain procedures are done before others, one can nest the envelopes.

Implementation

General Implementation details:

- It may often be cumbersome for the enumerator to compute $(i_P + i_I) \mod M$, which it was in our case. We found it easier to have the investigator put in the envelope a permuted list of outcome identifiers and have the participant's number end up selecting the i_Pth item.
- To create a non-uniform probability distribution over N outcome, let $M > N$ and have repeated elements. The investigator can choose how much information about this distribution to divulge before the envelope is opened. For the BDM prices we divulged the support.
- This whole process was explained to the participant before the participant makes decisions regarding something to be randomized over. We found it helpful to show the participant an example list like the one in the envelope.

envelope (and away from the edges) and (4) return the photo and envelope to Bob. This has the disadvantage of requiring a camera and sufficient lighting.

Figure 9: Density of estimated individual monthly discount rate

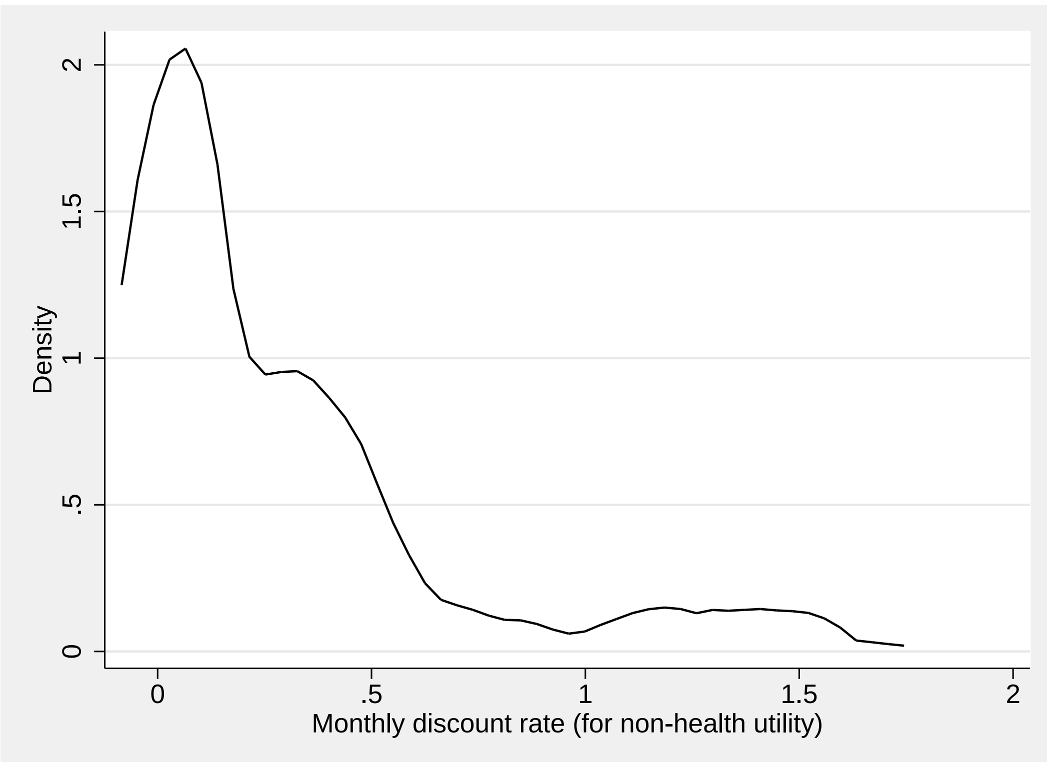

Figure 10: Density of estimated individual filter values

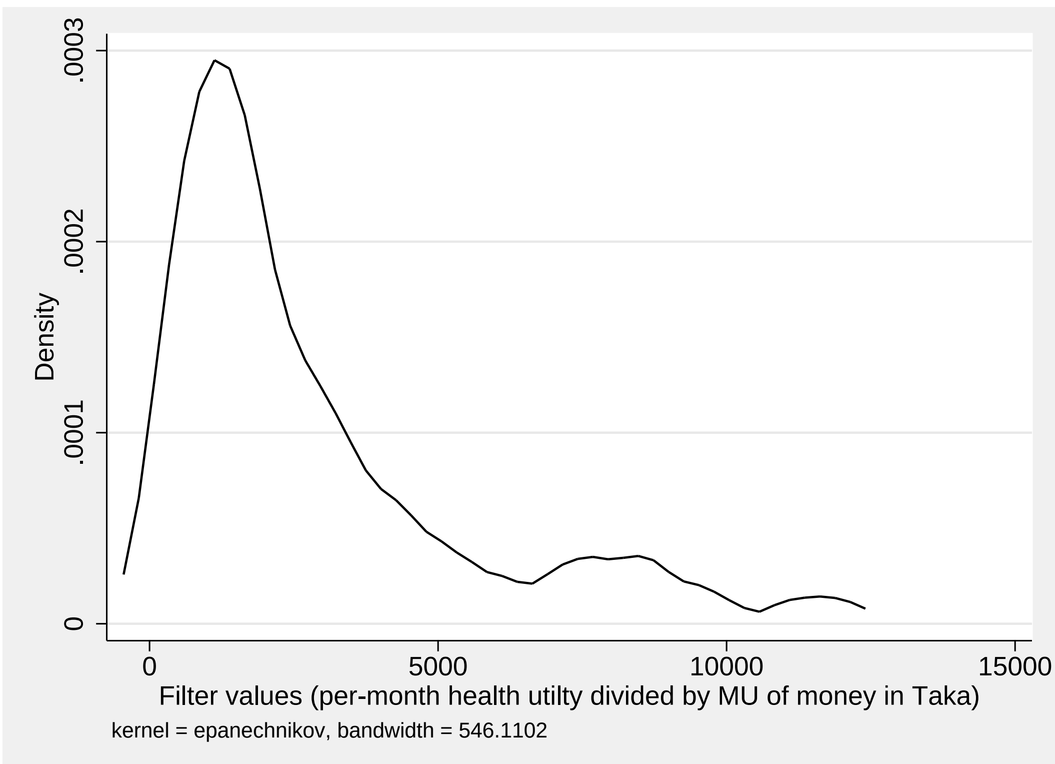

Figure 11: Density of estimated individual absolute risk aversion

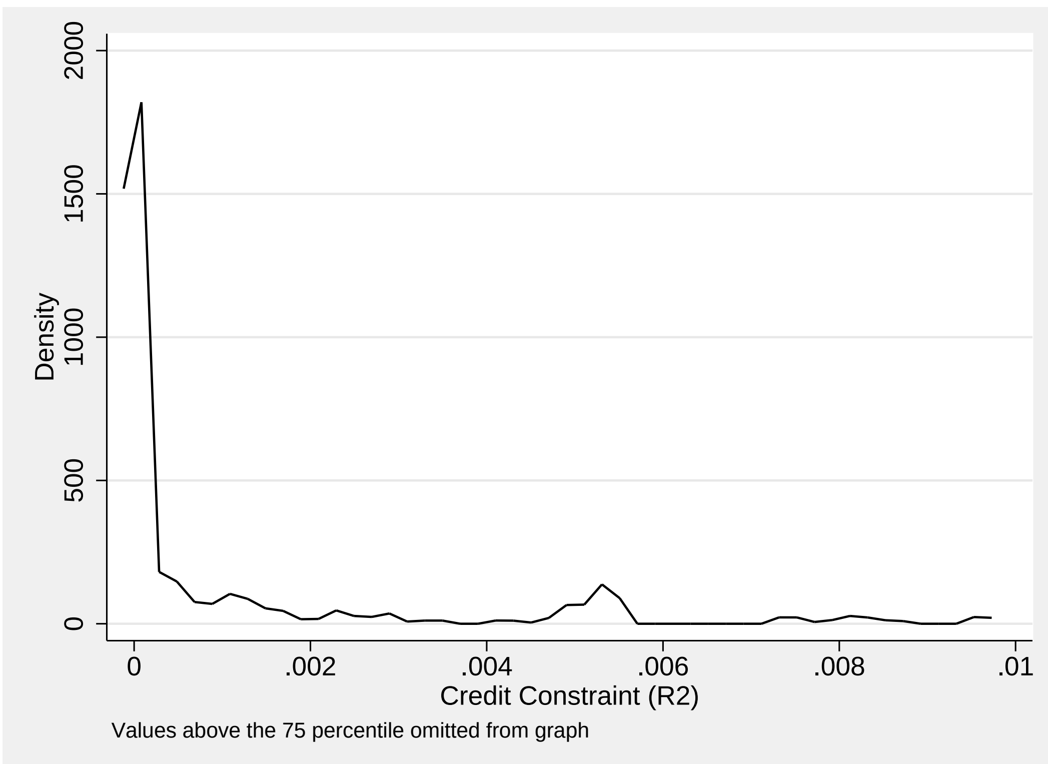

Figure 12: Envelope diagram