

Consuming Perishable Goods in the Presence of Transaction Costs and Liquidity Constraints

Megan Lang *

JOB MARKET PAPER
([latest version](#))

October 21, 2020

Abstract

Consumers face a trade-off when buying a perishable good with transaction costs. Buying in bulk minimizes transaction costs but creates waste. Eliminating waste by making small purchases raises costs, a problem compounded by liquidity constraints. I explore this trade-off using prepaid access time for solar electricity in rural Rwanda, a strictly non-storable good with transaction costs. I randomly offer 2,000 current solar customers a line of credit for solar access time, which alleviates liquidity constraints and lowers transaction costs. Consumers who previously bought in bulk respond by eliminating wasteful consumption, reducing demand by up to 6.4%. Those who are the most likely to be liquidity constrained increase demand by 88%. My results illustrate that transaction costs for perishable goods distort willingness to pay in opposite directions for different subsets of consumers. I show that reducing this distortion leads to a substantially higher estimate of consumer surplus from electricity than others in the literature have found. However, marginal households' willingness to pay still falls below current cost-covering levels.

*University of California, Berkeley; megan.lang@berkeley.edu. I thank the International Growth Centre for generous financial support. Thanks to Marco Gonzalez-Navarro, Meredith Fowlie, Alain de Janvry, Supreet Kaur, Ethan Ligon, Jeremy Magruder, Edward Miguel, Aprajit Mahajan, Elisabeth Sadoulet, Catherine Wolfram, and seminar participants for helpful comments. This study is pre-registered with the AEA registry (AEARCTR-0004887). It has approval from the Rwandan National Ethics Committee (No.0018/RNEC/2019) and IRB approval in the U.S.(U.C. Berkeley CPHS 2019-03-11994).

1 Introduction

When goods are perishable, buying in small quantities ensures that nothing goes to waste. Transaction costs for perishable goods create a trade-off. Buying in bulk incurs few transaction costs but generates waste when the perishable good expires. Buying small amounts minimizes waste but imposes a heavy burden of transaction costs. Liquidity constraints compound the problem: consumers with little liquidity get pushed toward frequent, small purchases and the associated high burden of transaction costs even when buying in bulk would be otherwise optimal. In both cases, observed willingness to pay yields an inaccurate measure of consumer surplus: consumers buying in bulk would be willing to pay a higher price for a lower quantity in the absence of transaction costs, and consumers buying in small increments would be willing to buy more at the going price.

I study the relationship between waste, liquidity, and transaction costs using a randomized control trial (RCT) with pay as you go (PAYGo) solar in rural Rwanda. PAYGo solar is an ideal good to study because it is strictly non-storable and involves substantial transaction costs for rural consumers. Consumers of PAYGo solar make a small down payment to have a solar home system installed. Then they “pay as they go” to use the electricity it generates by purchasing access time with mobile money. Consumers face transaction costs in the form of time: the average consumer in my sample travels fifty minutes, one way, to reach the nearest mobile money agent to buy access time. Once bought, access time runs down continuously. A consumer cannot save it for later once they buy it. PAYGo solar starkly illustrates the trade-off between waste and transaction costs, but rural consumers face similar trade-offs whenever they buy perishable food or prepay for stocks of goods like airtime, metered electricity, or mobile money. Such goods often involve transaction costs and can store poorly if inattentive consumers run through stocks faster than intended.

I randomly offer a short-term line of credit for PAYGo access time to 2,000 current solar customers in rural Rwanda. The line of credit may relax short-term liquidity constraints or provide credit at a lower price. It also directly lowers transaction costs for consumers who use it, since consumers call the solar company to borrow rather than traveling to a mobile money agent.¹ I stratify my sample based on pre-experimental demand, which is correlated with consumers’ propensity to buy in bulk and their self-reported access to credit. Stratification ensures that I am well powered to estimate heterogeneous treatment effects.

I generate predictions about the outcomes of the experiment by incorporating PAYGo access time with transaction costs into Deaton’s (1991) buffer stock model. I assume that the

¹Note that the line of credit may indirectly lower transaction costs for liquidity constrained consumers by allowing them to buy in bulk.

utility consumers gain from solar access is stochastic. This assumption is what leads to waste in the model: when consumers buy in bulk, they may end up having access to electricity on a day when they gain little to no utility from it. The model makes two sets of contrasting predictions about how consumers with limited versus ample access to liquidity will react to the offer of credit. First, lowering transaction costs increases demand for access time for consumers with limited liquidity who are not buying in bulk because it effectively lowers the price. Second, guaranteed access to credit increases demand for these consumers regardless of whether they actually borrow by reducing the precautionary savings motive. Consumers with guaranteed access to credit do not need to save as much to smooth consumption. By contrast, consumers with ample liquidity will buy in bulk when transaction costs are high, even though it generates waste.² It follows that reducing transaction costs using the line of credit may reduce demand among these consumers by allowing them to eliminate waste. Importantly, these reductions in demand will be the largest among consumers with the highest variance in the utility they gain from electricity: consumers with a high variance in utility benefit the most from more precisely targeting their consumption.

My estimated treatment effects are consistent with the four predictions from my model. First, consumers with low demand prior to the experiment, who buy in the smallest increments and report having little access to credit, dramatically increase demand by 88%. Second, I perform a bounding exercise to show that this 88% increase cannot be driven solely by consumers who actually use the line of credit, indicating that the effect is not only the result of lowering transaction costs. Guaranteed access to credit changes behavior among consumers who self-report having the least access to credit, consistent with the reduced precautionary savings motive in my model. Third, consumers who buy in bulk prior to the experiment reduce the quantity of access time they buy by up to 6.4%: lowering transaction costs appears to allow consumers to better target their consumption. Finally, treatment effects are decreasing in the pre-experimental variance in electricity used on purchased days, which is my proxy for the variance in the utility gained from electricity access. Consumers with the most variability in electricity use reduce their demand the most when I lower transaction costs. Put otherwise, the largest reductions in demand occur among consumers who can benefit the most from better targeting their electricity purchases.

Although the empirical results are consistent with my theoretical framework, I do not directly observe the liquidity constraints and transaction costs facing individual consumers throughout the course of my experiment. I consider an alternative underlying mechanism that may generate the same results. If pre-experimental demand is positively correlated

²I assume that consumers do not have a preference for waste minimization and that there are no variable transport costs associated with buying in bulk.

with present focus, also sometimes called present bias, then the negative treatment effects I estimate could be the result of high-demand consumers borrowing then procrastinating on repayment. In other words, the line of credit could be causing some consumers to borrow, delay repayment, and ultimately buy less electricity than they would under a strictly prepaid regime. Present focus is an important alternative to consider because if it is driving my results, offering the line of credit may not be welfare enhancing for consumers.

I present three pieces of evidence that rule out present focus. First, I show results from a separate experiment where I randomly offer consumers a bulk discount and a monthly reward for solar purchases. Both incentives offer equivalent average price reductions, but the bulk discount requires consumers to incur large costs in the present for benefits far into the future relative to the monthly reward. It follows that present focused consumers should respond more to the monthly reward. I cannot reject that the increase in demand is the same for both incentives, suggesting that consumers are not present focused. Second, I show that consumers across the distribution of pre-experimental demand are equally likely to choose a voluntarily lower borrowing limit when offered the line of credit. Equal take-up of a commitment device suggests that there is no correlation between pre-experimental demand and the proportion of sophisticated present focused consumers. Third, I present survey results showing that the vast majority of consumers across the distribution of pre-experimental demand overestimate their use of the line of credit. Naive, present focused consumers should underestimate borrowing. All three results indicate that present focus is not highly correlated with pre-experimental demand, and thus does not drive my results.

Consumer responses to the line of credit suggest that transaction costs and liquidity constraints limit consumer welfare from solar. However, offering the line of credit is not profitable for the solar firm. The reductions in demand that result from offering credit occur among the firm's most profitable customers. These consumers represent a much larger proportion of the firm's customer base than those who increase demand in response to the line of credit, causing the negative effects to swamp the positive effects when I re-weight my estimates to be representative of the customer population. I similarly find no significant net effect on repossessions, one of the firm's largest costs. My results highlight a potential mismatch between firm and consumer incentives to lower transaction costs for perishable goods, pointing to a possible role for policymakers to intervene.

Methodologically, revealed preference estimates of consumer welfare from solar are inaccurate unless they account for the distortions caused by liquidity constraints and transaction costs. I estimate a conservative lower bound on consumer surplus from electricity using random variation in the fee charged on the line of credit. My estimate suggests that consumer surplus from electrification is higher than previously believed: my lower bound is equal to

the most contextually similar estimate for total consumer surplus in the literature. Even so, marginal households' willingness to pay for electricity falls well below cost-covering levels.

My paper speaks to three strands of literature: impacts of transaction costs, the role of credit for poor households, and strategies for and outcomes of rural electrification. I contribute to the literature on transaction costs by providing empirical evidence on the impacts of transaction costs when goods are perishable. To date, the literature on transaction costs faced by consumers in low income countries has focused on financial services (Jack and Suri (2014), Suri, Jack, and Stoker (2012), Aycinena, Martinez, and Yang (2010), Collins et al. (2009), Beck et al. (2007), Beck et al. (2008), Dupas et al. (2018), and Ashraf, Karlan, and Yin (2006)). I show that transaction costs specifically for perishable goods can lead to some consumers buying inefficiently low quantities and others buying inefficiently high quantities. In both cases, transaction costs act as a tax on consumers. Transaction costs will persist in competitive markets when it is not profitable for firms to lower them, which is likely to happen in rural areas and when firms are serving low-income consumers. Better understanding consumer and firm responses to transaction costs clarifies which inefficiencies and inequities private investment will alleviate over time and which will require public investment.

My work contributes to two sub-literatures on credit in low income countries. The first is a nascent literature on the impacts of digital credit (see Francis, Blumenstock, and Robinson (2017) for an overview). I provide early causal evidence that small amounts of easily accessible credit can facilitate short-term consumption smoothing. I also contribute to a small literature that empirically examines the impact of guaranteed credit access on household behavior in low-income countries. Deaton (1991) establishes that credit access reduces precautionary savings motives in theory, but few consumers in low-income countries enjoy guaranteed access to credit. Lane (2020) provides empirical evidence over long time horizons, showing that guaranteeing credit in the event of a negative weather shock significantly increases upfront investment among farmers in Bangladesh. My work demonstrates the potential for small amounts of guaranteed, formal credit to significantly improve consumption smoothing for households over short time horizons.

My paper makes two contributions to the growing literature on electrification in low income countries. I join Jack and Smith (2015, 2020) in studying contracts for electricity with low-income households. I find that offering a line of credit significantly alters demand, but that it is not profitable for the solar firm to offer the more flexible contract. Like Jack and Smith (2020), my work shows that strictly prepaid contracts are efficacious from the firm's perspective relative to more flexible arrangements. However, firm profits from prepaid contracts in my setting are a function of market frictions, whereas increased profits in Jack and Smith (2020) primarily reflect reduced enforcement costs. The differences between the

two studies underline the importance of local market conditions when designing contracts with low-income consumers, as frictions like transaction costs are much more important in my rural setting than in Jack and Smith’s urban setting.

My second contribution to the literature on rural electrification is a novel estimate of consumer surplus from electricity. Estimating consumer surplus from PAYGo solar is important in its own right because PAYGo solar has the potential to become a key stepping stone in the global push to achieve universal access to electricity. In areas where expanding the grid is infeasible or households cannot afford grid connections, solar home systems provide reliable access to basic electricity. PAYGo solar is well-suited to low-income populations because it lowers barriers to adoption by allowing consumers to pay off costly solar home systems over time rather than making a single large purchase (Zollman et al. (2017)). In 2018 alone, PAYGo solar companies sold nearly 1 million solar home systems.³

Beyond the policy relevance of PAYGo solar, I measure demand for electricity under experimental conditions that deliberately reduce transaction costs and liquidity constraints. My work adds to a rich literature on the impacts of rural electrification with varied findings (Khandker, Barnes, and Samad (2009), Bensch, Kluve, and Peters (2011), Dinkleman (2011), Lipscomb, Mobarak, and Barham (2013), Khandker et al. (2014), Burlig and Preonas (2016), Chaplin et al. (2017), Lenz, Munyehirwe, and Sievert (2017), and van de Walle et al. (2017)). My estimate of consumer surplus directly builds upon the work of Lee, Miguel, and Wolfram (2020), Grimm et al. (2020), and Burgess et al. (2020) who provide estimates of consumer surplus from electricity in Kenya, Rwanda, and Bihar, India. Unlike other estimates in the literature, I measure demand on the use rather than the adoption margin. Focusing on the intensive margin rules out imperfect information as a key mechanism. PAYGo systems are also highly reliable, allowing me to measure demand absent concerns about supply-side reliability. Given that PAYGo systems generate small quantities of electricity relative to grid connections, my estimates focus on willingness to pay (WTP) for the first units of electricity, a critical margin for electrification policy. My conservative lower bound on consumer surplus is equal to the most similar estimate of total consumer surplus in the literature (Grimm et al. (2020)), suggesting that the benefits of electrification are higher than previously believed. However, my results indicate that currently non-electrified households will not be able to pay cost-covering levels for solar, highlighting the continued need for public investment to achieve universal electrification.

The rest of my paper is organized as follows. Section 2 describes the context. Section 3

³See https://www.gogla.org/sites/default/files/resource_docs/global_off-grid_solar_market_report_h2_2018_opt.pdf, which documents a 30% growth rate in PAYGo systems sold in the second half of 2018.

details the experimental design and describes my sample. Section 4 provides a theoretical framework to derive predictions about the impact of offering a line of credit for solar access. Section 5 presents reduced form results. Section 6 presents my estimated lower bound on consumer surplus from electrification, and section 7 concludes.

2 Background

In PAYGo solar contracts, consumers choose to adopt a solar home system that is typically bundled with high-efficiency appliances such as light bulbs, rechargeable radios, portable torches, phone chargers, or televisions. The more appliances the consumer opts to include in their bundle, the higher the price of the bundle. Once a consumer has selected their bundle, they make a down payment and have the solar panels, a battery for storing electricity, and all appliances installed in their home⁴.

After the solar home system (SHS) has been installed, consumers “pay as they go.” The solar company sets a daily rate for solar access time based on the number of appliances included in the SHS. Consumers prepay for SHS access time using mobile money. As soon as a consumer has purchased access time, they have unlimited access to their SHS for the duration of the purchased period.⁵ When access time runs out, the solar company remotely locks the consumer out of their SHS, preventing them from using it until they prepay for additional time. If the consumer does not purchase access for over thirty consecutive days, the solar company may repossess the SHS. Remote lockout and a credible threat of repossession render PAYGo solar contracts highly enforceable even in settings with weak institutions.

PAYGo contracts are designed to provide low-income households a degree of flexibility in paying for a solar home system, but such flexibility is limited by the perishable nature of access time combined with transaction costs. System access time runs down continuously regardless of how much a consumer actually uses their solar home system. Consumers cannot choose to delay the start of their purchased time, and they cannot choose to voluntarily shut down their solar home system and save some of their access time for later. For instance, if a consumer pays for three days of solar and then gets called away from their home for a day, they cannot recoup that day to use at a later time. Even though the SHS includes a battery that stores power generated by the solar panels, when the consumer is locked out of their SHS they cannot access the electricity stored in the battery. Continuous rundown combined with remote lockout from the entire SHS render access time strictly non-storable.

Traveling to a mobile money agent to purchase solar access time represents a transaction

⁴In my setting, the down payment amounts to 3-5% of the total value of the PAYGo contract.

⁵In my context, the battery that stores electricity generated by the solar panels is large enough that consumers are rarely constrained by the capacity of the system.

cost for the consumer. In phone surveys with two separate samples of solar customers in Rwanda, I asked how long it takes to reach the nearest mobile money agent. Figure 1 shows the distribution of travel time to reach the nearest mobile money agent, combining both survey samples. The average time is 50 minutes and the median time is 30 minutes one-way, although true transaction costs for purchasing solar likely vary depending on the timing of other tasks that might bring consumers close to a mobile money agent.

In theory, consumers can reduce the transaction costs associated with paying for solar by depositing money in their mobile money wallet when they are near an agent and later using those funds to buy solar. In a phone survey conducted among 1,229 solar customers in 2019, I asked consumers how many times they visited a mobile money agent to pay for solar out of their last five purchases. Figure 2 shows that 66% visited a mobile money agent all five times and 78% visited four of the last five times.⁶ This pattern is likely the result of limited mobile money take-up. A 2018 report by the World Bank found that only 31.1% of adults in Rwanda had mobile money accounts (WBG (2018)). While transactions conducted with mobile money are free, consumers have to pay withdrawal fees to convert mobile money into cash. With low take-up, consumers cannot use mobile money for most transactions and so withdrawal fees render it less liquid than cash. Even though consumers could use mobile money wallets to lower transaction costs, the survey evidence demonstrates that, in practice, consumers frequently incur transaction costs when paying for solar.

Transaction costs will be particularly burdensome for consumers without sufficient liquidity to buy in bulk. The median transaction size is 6.25 days of solar access.⁷ The prevalence of small transactions suggests that many consumers are either liquidity constrained or prefer to buy in small increments to limit wasted access time.

Taking all features of the setting together, the perishable nature of solar system access time combined with transaction costs creates stark trade-offs for consumers. They need to align cash flows with their demand for solar while minimizing transaction costs.

3 Experimental Design

I partner with a solar company in Rwanda to offer existing PAYGo solar customers a product designed to alleviate liquidity constraints and reduce transaction costs: a line of credit specifically for PAYGo system access time. The line of credit allows consumers to call the solar company and request to use up to one or two weeks of system access time before paying

⁶Appendix figure A1 indicates that trips to the mobile money agent are not driven by lack of knowledge about mobile money: nearly 80% of consumers report that they know how to use mobile money to buy solar if they have enough in their mobile money wallets.

⁷Figure A2 shows the full distribution of days purchased in a single transaction in the 90 days prior to the experiment.

for it. When a consumer makes a PAYGo payment after borrowing, the funds first go to repaying the time they borrowed plus a flat fee. Any funds that are left after repaying the line of credit go to pre-paying for additional system access time. In this way, consumers cannot default on the line of credit without defaulting on their entire PAYGo contract. Consumers can use the line of credit as many times as they like over the course of the experiment.

The line of credit simultaneously addresses liquidity constraints and transaction costs. It alleviates liquidity constraints by allowing consumers to purchase solar access time when they do not have cash on hand. Even if consumers are not strictly liquidity constrained, it may provide a less expensive source of credit than would otherwise be available. As I will demonstrate in the model, having guaranteed access to credit can allow consumers to increase demand even if they do not borrow by providing a tool to help them smooth consumption. The line of credit also reduces transaction costs for consumers who actually use it because consumers use the line of credit by calling the solar company rather than traveling to a mobile money agent. The line of credit enables consumers to decouple cash flows with their demand for electricity and time trips to the mobile money agent to better suit their convenience.

I cross-randomize the terms of the line of credit along three dimensions. Half of the consumers in the treatment group can only borrow up to seven days of solar access time, while the other half can borrow up to fourteen days. All consumers have the option to choose a voluntarily lower borrowing limit than the one originally offered, which allows consumers to commit to borrowing smaller amounts. Half of consumers pay a 10% flat fee on borrowed days and half pay a 2% fee.⁸ Finally, half of consumers lose access to the line of credit if they do not repay their borrowed days plus the fee within one week of their borrowed time running out. The other half do not face any such time limit, but like all PAYGo customers they get remotely locked out of their system when they run out of access time.

In total, the solar company offered the line of credit to 2,000 randomly selected existing solar customers in the Northern and Southern provinces of Rwanda who had signed PAYGo contracts at least 90 days prior to the start of the experiment. The control group consists of all other existing customers in the Northern and Southern provinces who had signed contracts at least 90 days prior to the experiment: 9,360 consumers.

Consumers in my sample are self-selected, as they have all opted to sign PAYGo solar contracts. I combine responses to a phone survey conducted with 1,229 solar customers in 2019 with the latest Integrated Household Living Survey, a nationally representative survey of Rwandan households last conducted in 2016-2017. Using questions common to both surveys,

⁸10% is comparable to rates charged by telecommunications companies in Rwanda when consumers borrow airtime, which is the most similar product I identified in rural markets.

I construct a wealth index to compare the population in my sample to the distribution of rural households in Rwanda.⁹ I find that consumers in my sample are wealthier than the average rural Rwandan household, a feature I return to in my discussion of the welfare impacts of rural electrification.¹⁰

I stratify my treated sample based on the 90-day utilization rate (UR) prior to the start of the experiment. The utilization rate is the proportion of days a consumer has purchased system access. To understand consumer responses to the line of credit across the 90-day UR distribution, I create four stratification bins: 0-30%, 30%-65%, 65%-80%, and 80%-100%. Figure 3 shows the distribution of UR in the 90 days prior to the start of the experiment, along with lines denoting the stratification bins.

Descriptive information about differences between consumers in each utilization bin shows that pre-experimental demand is correlated with access to liquidity and consumers' propensity to buy in bulk. Table 2 shows differences in self-reported borrowing to pay for solar when the line of credit is not available.¹¹ Consumers with the lowest pre-experimental demand are significantly less likely to have borrowed for solar than consumers in other utilization bins, and those who do borrow appear to borrow less. Of the consumers who have not borrowed to pay for solar, those with the lowest pre-experimental demand are more likely to report that they did not borrow because they were unable to find credit. Table 1 shows that consumers with the highest pre-experimental demand are more likely to buy in bulk: average purchase sizes for consumers with the highest pre-experimental demand are more than double those for consumers with the lowest pre-experimental demand.¹²

Correlations between pre-experimental demand and liquidity constraints as well as the propensity to buy in bulk illustrate why it is important to stratify based on pre-experimental demand. Consumers who lack easy access to liquidity versus those who have sufficient liquidity to buy in bulk are likely to respond to the line of credit differently. Stratifying on the basis of pre-experimental demand provides me with the statistical power necessary to estimate differential average treatment effects across the distribution of consumers.

⁹I use the following variables to construct the wealth index: ubudehe category (a government-assigned category designed to summarize the socio-economic status of a household), roof material, wall material, floor material, primary source of electricity (if any), primary source of light, whether or not the household is connected to the national grid, and weekly energy expenditures.

¹⁰Figure A3 shows the nationally representative distribution of wealth scores among rural consumers, with the red line representing the mean among all rural Rwandan households and the blue line representing the mean in my sample.

¹¹Self-reports come from the 2019 phone survey of 1,229 solar customers.

¹²Table A2 shows differences between stratification bins along other dimensions of economic well-being that are less closely connected to the decision to buy in bulk versus buy small quantities frequently.

3.1 Timeline and Data

The solar company marketed the line of credit starting on October 14, 2019. Marketing involved calling each customer in the treatment group to explain the terms of the line of credit and how to access it. During the initial call, all consumers were also offered the option to choose a voluntarily lower credit limit, a form of commitment device for consumers concerned about borrowing too much. All consumers also received a SMS message containing details of the line of credit. After completing the initial round of marketing calls, the solar company attempted to call every treated consumer again to complete a short survey and to further educate customers about the line of credit starting in late November.¹³ Consumers could use the line of credit through February 14, 2020. All treated consumers received a SMS message on the last day of the experiment informing them that the program had ended.

My primary source of data is the administrative records of the solar company. The dataset of loan requests and repayments allows me to estimate differences in use of the line of credit as well as differences in price sensitivity between treatment groups and stratification bins. I use administrative data to conduct initial balance checks and to estimate changes in the monthly utilization rate.¹⁴

I augment these data with two other datasets that enable me to test for heterogeneity along relevant dimensions other than pre-experimental demand. Data generated by the solar home systems provide a measure of the amount of electricity actually consumed on each day a solar home system is switched on. I use daily totals of watt hours consumed when systems are switched on to check for heterogeneous treatment effects based on the pre-experimental variance in watt hours used. I also use information from a short phone survey conducted by the solar company midway through the experiment to check for differences in usage rates based on the distance from the nearest mobile money agent and to assess the accuracy of consumer expectations about borrowing.¹⁵

Next, I turn to a theoretical framework to generate predictions about the outcomes of my experiment. The model focuses on two key elements of my experimental design. First, my sample stratification is based on observed heterogeneity in consumer decisions around PAYGo solar prior to the experiment. This suggests heterogeneity in access to liquidity. Second, the line of credit induces exogenous changes in the availability of liquidity and the

¹³64% of consumers in the 0%-30% stratification bin were reached on the phone during at least one of the rounds of calls. 86%, 90%, and 96% of consumers were reached in the 30%-65%, 65%-80%, and 80%-100% bins.

¹⁴Table A1 shows that randomization yields balanced treatment and control groups on a range of observable characteristics after controlling for each consumer's stratification bin.

¹⁵Note that the phone survey conducted by the solar company mid-experiment is distinct from the 2019 phone survey of solar customers I use to generate descriptive statistics on consumers in different stratification bins. Importantly, the mid-experiment survey only included treated consumers.

size of transaction costs. Therefore, my theoretical framework will consider how heterogeneously liquidity constrained consumers will respond to exogenous changes in liquidity and transaction costs.

4 Theoretical Framework

A representative consumer gets utility from consuming some composite consumption good c and from using the appliances that can be powered by their solar home system. The consumer chooses how many days of solar access to buy each day. I denote the quantity of days bought on a given day as $q \in \mathbb{Z}^+$ and the overall stock of days of solar access as $e \in \mathbb{Z}^+$. The consumer's stock of electricity access evolves according to a simple law of motion that describes the continuous rundown of system access time

$$e' = e + q - (\mathbf{1}(e + q \geq 1)). \quad (1)$$

Consumers get stochastic utility α if they have at least one day of electricity access ($e + q \geq 1$) in a given day. I assume α is drawn from some distribution with cdf $F(\cdot)$ with support over $[\alpha_m, \alpha_M]$, where $\alpha_m \geq 0$ and α_M is some finite constant. Combined with the continuous rundown described in equation (1), stochastic utility generates waste: if a consumer pays for solar access into the future, they may end up having access on a day with a low realization of α when they would not have opted to purchase access otherwise. At the start of each time period, consumers learn their draw of α and receive a stochastic endowment of income $y \in (0, y_M]$, where y_M is some finite constant.

Consumers choose how many days of solar access to purchase and how much of the composite consumption good to consume, which leaves some level of savings to be carried forward to the next day. As equation (1) indicates, consumers can purchase multiple days of solar at once but they cannot store it in the sense that once they have purchased solar access, the stock decreases every day until it is zero or until the consumer buys additional access time.

The consumer's preferences are represented by

$$U = \mathbb{E} \left[\sum_{t=0}^{\infty} \beta^t (\alpha_t \mathbf{1}(e_t + q_t \geq 1) + u(c_t)) \right]. \quad (2)$$

$\beta \in (0, 1)$ is the discount rate. I assume $u'(\cdot) > 0$, $u''(\cdot) < 0$, $u'(0) = \infty$, and $u'(\infty) = 0$.

The relative price of solar is p . Consumers incur a transaction cost τ each time they buy solar. I make the simplifying assumption that τ is constant across time and consumers

and that it represents the costs of reaching the nearest mobile money agent. I denote the consumer's non-solar asset stock as s . Any assets that the consumer saves in the current period earn a return $1 + r$, or if consumers borrow they have to repay $1 + r$ multiplied by the borrowed amount in the following period. The consumer's asset stock is governed by the law of motion

$$s' = (1 + r)(s - c - qp - \mathbf{1}(q \geq 1)\tau) + y'. \quad (3)$$

Finally, the consumer faces a borrowing constraint l . In each period, the consumer needs to choose c , s' , and q to satisfy¹⁶

$$-l \leq s'. \quad (4)$$

4.1 Characterizing Optimal Choices

Each day, the consumer chooses c and q to maximize (2) subject to (1), (3), and (4). For a given q , the Bellman equation is

$$\begin{aligned} V_q(s, e, y, \alpha) = \max_{s'} & \left[\alpha \mathbf{1}(e + q \geq 1) + u(c(q)) + \beta \mathbb{E}[V(s', e', y', \alpha')] \right] \\ \text{s.t. } & s' = ((1 + r)s - c(q) - qp - \mathbf{1}(q \geq 1)\tau) + y', \\ & e' = e + q - \mathbf{1}(e + q \geq 1), \text{ and } -l \leq s'. \end{aligned} \quad (5)$$

The consumer chooses the value of q that maximizes current and future utility, meaning that they choose q to satisfy

$$V(s, e, y, \alpha) = \max_q \{V_q(s, e, y, \alpha)\}. \quad (6)$$

When $e > 0$, the consumer enjoys access to solar regardless of the realization of α . Given that I hold p and τ constant over time, choosing $q > 0$ when $e > 0$ weakly reduces utility today. While choosing $q > 0$ could raise expected utility tomorrow, the consumer can costlessly wait for α' and y' to be realized and then make the optimal decision. It follows that I only need to consider the consumer's choice of q when $e = 0$.

Let μ_q be the Lagrange multiplier on the liquidity constraint (4) for a given choice of q . Let $\mathbb{E}[\bar{V}_q(s', e', y', \alpha')] = \mathbb{E}[V(s', e', y', \alpha')|q]$ be the maximal expected V for a given choice of q . The interior solutions to the sub-problems defined by equation (5) are characterized by

¹⁶The model generates similar predictions if instead of assuming consumers are liquidity constrained, I assume that they can borrow at a high price which I lower by offering the line of credit.

the first order condition

$$\beta \mathbb{E} \left[\frac{\partial \bar{V}_q(s', e', y', \alpha')}{\partial s'} \middle| s, e, y, \alpha \right] - \mu_q = \frac{du(c(q))}{dc(q)}. \quad (7)$$

The envelope condition allows me to write $\mathbb{E} \left[\frac{\partial \bar{V}_q(s', e', y', \alpha')}{\partial s'} \middle| s, e, y, \alpha \right]$ as

$$\mathbb{E} \left[\frac{\partial \bar{V}_q(s', e', y', \alpha')}{\partial s'} \middle| s, e, y, \alpha \right] = (1 + r) \mathbb{E} \left[\frac{du(c'(q'))}{dc'(q')} \middle| s, e, y, \alpha \right]. \quad (8)$$

I then substitute (8) into (7) to obtain the Euler equation:

$$\beta(1 + r) \mathbb{E} \left[\frac{du(c'(q'))}{dc'(q')} \middle| s, e, y, \alpha \right] - \mu_q = \frac{du(c(q))}{dc(q)}. \quad (9)$$

To simplify notation, let $\frac{du(c(q))}{dc(q)} = u'(c(q))$, similarly let $\frac{du(c'(q'))}{dc'(q')} = u'(c'(q'))$. Then I can re-write the consumer's optimal choice of c as

$$u'(c(q)) = \max \left\{ \beta(1 + r) \mathbb{E}[u'(c'(q')) \middle| s, e, y, \alpha], u'(s - qp - \mathbf{1}(q \geq 1)\tau + l) \right\}. \quad (10)$$

Equation (10) characterizes the consumer's optimal choice of c for a given q . It is straightforward to show that under certain conditions, there will exist a policy function $\sigma_q(s, e, y, \alpha)$ that defines optimal consumption for a given realization of the state. Importantly, expectations are taken over both α' and y' . Expectations over α' speak to the need for an additional policy function that governs how the probabilities of choosing different levels of q in the future change based on choices of c and q today.

The consumer chooses $q = 1$ rather than $q = 0$ if and only if

$$\alpha \geq u(\sigma_0(s, e, y, \alpha)) - u(\sigma_1(s, e, y, \alpha)) + \beta \mathbb{E}V((1 + r)(s - \sigma_0(s, e, y, \alpha)), e', y', \alpha') - \beta \mathbb{E}V((1 + r)(s - \sigma_1(s, e, y, \alpha) - p - \tau), e', y', \alpha'). \quad (11)$$

Call the threshold level of α where equation (11) holds with equality α^* . For a given realization of s, e , and y , the consumer prefers buying zero days to one day of solar with probability $F(\alpha^*(s, e, y))$ and prefers buying one day to zero days with probability $1 - F(\alpha^*(s, e, y))$. The function $\alpha^*(s, e, y)$ allows me to take expectations over α' .

A consumer choosing between $q = 1, 2, \dots$ will only condition their choice on s and y . To see why, note that a consumer choosing between $q = i$ and $q = j$ with $i, j \geq 1$ will prefer

i to j if and only if

$$\alpha + u(\sigma_i(s, e, y, \alpha)) + \beta \mathbb{E}V(s'(i), e', y', \alpha') \geq \alpha + u(\sigma_j(s, e, y, \alpha)) + \beta \mathbb{E}V(s'(j), e', y', \alpha'),$$

or, rearranging,

$$u(\sigma_i(s, e, y, \alpha)) - u(\sigma_j(s, e, y, \alpha)) \geq \beta \mathbb{E}V(s'(j), e', y', \alpha') - \beta \mathbb{E}V(s'(i), e', y', \alpha').$$

Intuitively, since the consumer gains α regardless of the choice of $q \geq 1$, the decision depends only on the other state variables. Since I've already shown that the consumer only chooses $q > 0$ when $e = 0$, it follows that the consumer's choice only depends on s and y .

For a given realization of (s, y) , the consumer knows whether i is preferred to j . I assume that if the consumer prefers $q = 0$ to $q = 1$, they will also prefer $q = 0$ to $q > 1$. With this simplifying assumption in place, I can write the consumer's expectations as

$$\mathbb{E}V(s', e', y', \alpha') = \begin{cases} \mathbb{E}_y V_0(s', e', y', \alpha') & \text{if } e' > 0 \\ F(\alpha^*(s', e', y')) \mathbb{E}_y V_0(s', e', y', \alpha') + \\ (1 - F(\alpha^*(s', e', y'))) \mathbb{E}_y \max\{V_1(s', e', y', \alpha'), V_2(s', e', y', \alpha'), \dots\} & \text{if } e' = 0. \end{cases} \quad (12)$$

4.2 Solving the model

Taken together, the consumer's decisions can be fully characterized using the set of optimal consumption functions $\sigma_0(s, e, y, \alpha)$, $\sigma_1(s, e, y, \alpha), \dots$, the set of value functions $V_0(s, e, y, \alpha)$, $V_1(s, e, y, \alpha), \dots$, and the function $\alpha^*(s, e, y)$.

I use numerical maximization to obtain the functions that characterize the solution to the consumer's problem. I assume the constant relative risk aversion utility function so that

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma},$$

with $\gamma > 1$. I assume that α is drawn from a uniform distribution on $[\alpha_m, \alpha_M]$, although later I will allow α_m and α_M to be different for different consumers. I make two additional assumptions for computational simplicity. First, I assume that y follows a three-state Markov chain $y[z]$ with state $z \in \{0, 1, 2\}$ and transition matrix P . Intuitively, this means that the consumer receives either a high, medium, or a low draw of income on each day with some probability. Second, I limit choices of q to 0, 1, or 2, which allows me to reduce the state space to s, z , and α . With these assumptions in hand, I can generate predictions about the

expected outcomes from my experiment.

Prediction 1: lowering transaction costs increases demand among consumers not buying in bulk. I assume that liquidity constrained consumers cannot buy in bulk, making the relevant choice the one between $q = 0$ and $q = 1$. Consumers choose $q = 0$ rather than $q = 1$ with probability $F(\alpha^*(s, z))$, so the relevant comparative statics for my experiment are $\frac{d\alpha^*}{d\tau}$ and $\frac{d\alpha^*}{dl}$.

Figure 4 illustrates how lowering transaction costs changes α^* . In the figure, the point where V_0 and V_1 intersect represents α^* , the lowest draw of α for which a consumer will opt to buy electricity rather than forgoing it. Figure 4 shows that lowering transaction costs reduces α^* , increasing the probability that the consumer chooses to buy one day of solar. As expected, the change in α^* is larger when realizations of s and y are lower.

Prediction 2: relaxing liquidity constraints increases demand among consumers not buying in bulk, even if they do not use the line of credit. Figure 5 tells a similar story: increasing the borrowing limit l reduces α^* . For consumers who lack the liquidity to buy in bulk, the line of credit unambiguously increases demand for solar. Importantly, figure 5 illustrate another reason that the line of credit increases demand among liquidity constrained consumers: it reduces the precautionary savings motive. V_0 and V_1 are higher under the experimental condition. When consumers have guaranteed access to credit they don't need to save as much today to ensure that they can buy solar tomorrow.

Lowering transaction costs also reduces the precautionary savings motive: when consumers know they face lower transaction costs tomorrow, they do not need to save as much today. However, recall that consumers in my experiment only experience a reduction in transaction costs when they actually borrow. It follows that the reduction in the precautionary savings motive from relaxing liquidity constraints is particularly important because it implies that offering the line of credit can change consumer behavior even if consumers do not use the line of credit.

Prediction 3: The line of credit may reduce demand among consumers who buy in bulk. I turn now to the choice for consumers who have sufficient liquidity to buy in bulk. Recall that the choice between $q = 1$ and $q = 2$ does not depend on α . I instead consider the range of assets s over which the consumer prefers $q = 1$ to $q = 2$ for a given realization of y .

In figures 6 and 7, the green points indicate the level of s above which V_2 exceeds V_1 for low and high values of τ and l . The range to the right of the green points is the range of asset realizations over which the consumer prefers to buy in bulk. If the green point moves to the left as a result of my experiment, it follows that the experiment increases the probability that a consumer will prefer to buy in bulk. Figure 7 shows that lowering transaction costs

reduces the range of assets where the consumer prefers to buy in bulk ($q = 2$). Figure 6 shows that relaxing liquidity constraints can either widen or narrow the range of assets where the consumer prefers to buy in bulk depending on the relative size of p and τ .

For consumers with sufficient liquidity to buy in bulk prior to the experiment, offering the line of credit may operate primarily as a reduction in transaction costs. Given that lowering transaction costs makes buying in bulk less appealing, the line of credit could lower overall demand among consumers previously buying in bulk. Such consumers may stop buying in bulk and instead target their purchases to days when they receive high realizations of α .

Prediction 4: Negative treatment effects will be larger for consumers with a higher variance in α . Figure 7 illustrates the final prediction from my model. Comparing the top and bottom figures, I show that the potential reduction in demand as a result of lowering transaction costs is larger for consumers with a higher variance in α . Intuitively, these are the consumers who will benefit the most from better targeting their consumption. In fact, with the realization of income I selected for the sake of illustration in figure 7, lowering transaction costs leads to virtually no change in the probability of buying in bulk for consumers with a low variance in α . The final prediction of the model is that treatment effects from the line of credit should be decreasing in the variance in α .

The final prediction offers a way to clarify the ambiguous predictions about the change in the probability that consumers buy in bulk. Alleviating liquidity constraints can either increase or decrease the likelihood of buying in bulk, but lowering transaction costs should lower bulk demand. However, if consumers respond to the line of credit by reducing demand then those reductions should be greater for consumers with a higher variance in the utility realized from solar access if they are being driven by the reduction in transaction costs.

The model generates four empirically testable predictions about offering the line of credit. First, lowering transaction costs will increase demand among consumers who lack sufficient liquidity to buy in bulk. Second, relaxing liquidity constraints will increase demand among consumers not buying in bulk regardless of their borrowing status: providing guaranteed access to credit can change consumer behavior even if consumers do not use it. Third, the line of credit may lower demand among consumers who previously bought in bulk. Fourth, if certain consumers do reduce demand in response to being offered the line of credit, then treatment effects from the line of credit will be decreasing in the variance of α if the reduction in demand is the result of lowering transaction costs. Consumers with the most variance in their utility from electricity have the strongest incentive to stop buying in bulk and target their consumption when transaction costs are lower. Next, I present empirical results to evaluate how well the model describes consumer behavior, and to consider potential alternatives explanations.

5 Results

I measure consumer responses to the line of credit by estimating heterogeneous average treatment effects on monthly utilization rates, or the proportion of days each month that a consumer has access to their solar home system. Let i index consumers, j index stratification bins, and t index months of the experiment. Tmt_{it} is a dummy variable equal to one if consumer i had access to the line of credit in month t . Bin_{ij} is a dummy variable equal to one if consumer i is in stratification bin j based on their pre-experimental utilization rate. γ_i and γ_t are consumer and month fixed effects. Using a three month pre-period to increase precision, I estimate

$$UtilizationRate_{it} = \alpha + \sum_{j=1}^4 \beta_j (Tmt_{it} \times Bin_{ij}) + \gamma_i + \gamma_t + \epsilon_{it}.$$

Figure 8 shows that average treatment effects from offering the line of credit follow the first and third theoretical predictions in my model. Consumers with the lowest pre-experimental demand increase their monthly utilization rates by 11pp as a result of being offered the line of credit, an increase of 88% over the control group. Consumers in the second-lowest stratification bin significantly increase utilization rates as a result of being offered the line of credit, although at a more modest magnitude of 5.3%. By contrast, consumers in the second highest stratification bin reduce their utilization rate by 6.4%, while consumers with the highest pre-experimental demand reduce utilization by 1.6%.¹⁷ It appears that the line of credit increases demand for consumers who are most likely to be liquidity constrained while allowing consumers who previously bought in bulk to better target their consumption.

The second prediction of the model is that credit availability reduces the precautionary savings motive for liquidity constrained consumers. Even in periods where the liquidity constraint does not bind, consumers do not need to save as much to smooth consumption when they have guaranteed access to credit. Reducing the precautionary savings motive may increase demand for solar even for consumers who do not ultimately use the line of credit.

I want to test whether the line of credit alters consumer behavior among non-borrowers; however, I cannot estimate separate effects for borrowers and non-borrowers because I do not know which consumers in the control group would have borrowed. Instead, I consider the hypothesis that borrowers drive all estimated treatment effects. If so, perfect compliance with my randomization allows me to calculate local average treatment effects (LATEs) for

¹⁷The reduction in demand between consumers in the 65%-80% and 80%-100% stratification bins is not statistically significantly different.

borrowers as

$$LATE = \frac{\Delta UR}{ProportionBorrowers}.$$

Figure 9 shows the proportion of consumers in each stratification bin that use the line of credit over the course of the experiment. Only 4.5% of consumers with the lowest pre-experimental demand use the line of credit. If the average treatment effect for low-demand consumers in figure 8 is driven entirely by borrowers, it would imply an impossibly large LATE of 244pp.¹⁸ It follows that the average treatment effect must be driven in part by consumers who do not borrow, at least among those with the lowest pre-experimental demand. Increased demand as a result of guaranteed access to credit, irrespective of credit use, is consistent with consumers engaging in precautionary saving as described by my model.

The final prediction in my model is that treatment effects are decreasing in the variance in α , the utility consumers realize from electricity access on a given day. Although I cannot observe utility from solar access, I do observe the number of watt hours (wH) used on each day a consumer's solar home system is switched on. I use wH consumed as a proxy for the utility gained from solar access. For each consumer, I calculate the standard deviation of wH used on days when they system is switched on in the 90 days prior to the experiment. I divide the distribution of standard deviation in use into quartiles. Letting k index quartiles of standard deviations in use, I estimate heterogeneous treatment effects using the specification

$$UR_{it} = \alpha + \sum_{k=1}^4 \beta_k (Tmt_{it} \times SDUseQuartile_{ik}) + \gamma_i + \gamma_t + \epsilon_{it}.$$

Note that the variability in α only matters for consumers who previously bought in bulk, as variance has no bearing on the decision to buy one day or forgo access. Low variance consumers are more likely to continue buying in bulk even after being offered the line of credit since precisely targeting consumption matters less. Pooling liquidity constrained and consumers with those previously buying in bulk, the effect of offering the line of credit should be primarily driven by the liquidity constrained consumers who unambiguously increase their demand in the low variance quartile. Consumers with a high variance in their utility from solar access will stop buying in bulk because the line of credit allows them to better target their consumption. The total effect for consumers with a high variance in utility from solar will be a weighted average of increased demand from liquidity constrained consumers and reduced demand from consumers buying in bulk.

Figure 10 shows that consumers with the lowest variance in use are the only group with

¹⁸Implied LATEs for the other stratification bins are less extreme: 11pp, -14pp, and -6pp. Even if I only take the proportion of borrowers out of the 64% of consumers who actually answered a marketing call in the 0%-30% stratification bin, the implied LATE is 157pp.

a significantly positive treatment effect. On average, consumers with the lowest variance in use increase their monthly utilization rates by 5pp in response to being offered the line of credit. Consumers in the second and third quartiles do not exhibit any significant treatment effect from being offered the line of credit. Consumers with the highest variance in use reduce utilization by around 2pp in response to being offered the line of credit, although the effect is only significant at the 10% level. Consumers with the least variance in use have a significantly higher treatment effect than those with the most variance in use. The estimated results are consistent with high demand, high variability consumers better targeting their solar consumption when provided with a way to reduce the burden of transaction costs, providing further evidence in support of my theoretical framework.¹⁹

While the pattern of average treatment effects is consistent with the mechanisms proposed in my model, I cannot directly observe transaction costs or liquidity constraints for all of the consumers in my sample throughout the course of the experiment. I provide descriptive evidence to further explore the role of transaction costs. Unlike the reduction in the precautionary savings motive, which occurs regardless of use, the line of credit only reduces transaction costs if consumers use it. I expect consumers who face the highest transaction costs to use the line of credit the most. While I do not have data on transaction costs for all consumers in my sample, I know the time to reach the nearest mobile money agent for most consumers in the treatment group. I use this data to run a descriptive regression to see whether consumers who face higher transaction costs are more likely to use the line of credit. Table 3 shows that living one hour further from the nearest mobile money agent is associated with a 5.8pp, or nearly 50%, increase in the likelihood that a consumer uses the line of credit, providing support for the importance of transaction costs as a key mechanism underlying my estimated treatment effects.

Taken together, my empirical results closely match the predictions of my theoretical framework and rule out a number of alternative mechanisms. In the next section, I consider an important alternative explanation for my estimated treatment effects: pre-experimental demand is correlated with present focus, or present bias (Laibson (1997), O’Donoghue and Rabin (1999)). Of the range of possible alternative mechanisms, present focus merits special consideration because it changes the welfare implications of my results. If consumers with the highest pre-experimental demand are reducing solar purchases as the result of behavioral

¹⁹The pre-experimental variance in watt hours used is an imperfect proxy for α , the utility gained from solar access, because I only observe it on days when consumers have access to their systems. I provide a further robustness check for this effect by re-estimating it only for consumers in the top stratification bin, who have the highest pre-experimental demand. These consumers have access to their systems almost all of the time in the pre-experimental period. Figure A4 shows that the same pattern holds among these consumers: those with the lowest pre-experimental variance in use exhibit significantly higher treatment effects than those with the highest pre-experimental variance in use.

biases, it is no longer clear that offering the line of credit is welfare-enhancing for consumers. Methodologically, standard revealed preference measures of consumer welfare are inaccurate when consumers are present focused (Bernheim and Rangel (2009), Allcott and Taubinsky (2015)). For both reasons, it is important to establish that my results are not being driven by present focus before I proceed to welfare estimation.

5.1 Evidence on the Importance of Present Focus

Under pure prepayment, consumers pay prior to enjoying solar access. With the line of credit, present focused consumers will prefer to borrow access time to delay the costs of electricity access while still enjoying the benefits. They are then likely to procrastinate on repayment, potentially leading to a reduction in demand given that consumers cannot buy more access time until they pay back the line of credit. If the degree of present focus is positively correlated with pre-experimental demand, then the negative treatment effects I estimate for consumers with the highest pre-experimental demand could be the result of present focus rather than consumers better targeting their solar purchases.

I lack direct measures of present focus among consumers in my sample. Instead, I evaluate whether present focus is driving my experimental results by providing four pieces of evidence. First, I show that the two dimensions of heterogeneity I examined in the previous section, both of which yielded results consistent with my model, are uncorrelated. Second, I summarize results from a separate experiment where customers of the same solar company were offered time-varying incentives for solar payments. Third, I show how many consumers in each stratification bin opted for a voluntarily lower borrowing limit, the commitment device offered at the start of the experiment. Finally, I evaluate the accuracy of consumer expectations about borrowing. All four pieces of evidence suggest that present focus is not driving my results.

Both heterogeneity in pre-experimental demand and pre-experimental variance in use generate results that are consistent with my theoretical framework. If present focus is driving results along both dimensions of heterogeneity, then there should be a positive correlation between pre-experimental demand and pre-experimental variance in electricity use. Figure 11 shows that pre-experimental demand is not closely correlated with the standard deviation in wH used on days when a consumer has solar access.

Further evidence on the importance of present focus comes from a separate randomized control trial with a different sample of customers from the same solar company. I randomly offered 1,600 consumers incentives to buy more days of solar access. Half of the treatment group received x access days for free if they bought y days in bulk, and half received x free access days if they bought y days over the course of a calendar month.

Table 4 outlines the full cross-randomization. I hold constant the total number of days required to qualify for the incentives and the number of free days consumers could earn between the bulk discount and the monthly reward, leading to equivalent reductions in average price between the two incentives. Consumers should respond more to the monthly reward if they are present focused: the bulk incentive requires that consumers forgo more consumption today to gain consumption farther into the future relative to the monthly reward. I stratify the sample by pre-experimental demand over the entire duration of consumers' tenure with the firm. While slightly different than the stratification I use to test the line of credit, it still allows me to measure heterogeneous effects across the distribution of demand.

Figure 12 shows the effect of the incentives for solar payments on the number of days bought per month across the distribution of pre-experimental demand. Treatment effects are null or significantly negative for all consumers except those with the highest pre-experimental demand. Consumers with the highest pre-experimental demand respond to the incentives by increasing purchases by 2-2.5 days per month, a 6-8% increase relative to the control group. Critically, I cannot reject that bulk incentives and the monthly reward have the same effect for consumers with the highest pre-experimental demand. Consumer responses to the incentives indicate that consumers are likely not present focused in a manner that is strongly correlated with pre-experimental demand.

I provide two final pieces of evidence on present focus among consumers in my sample. First, during the initial round of marketing calls for the line of credit, all consumers in the treatment group were given the option to select a voluntarily lower borrowing limit. Selecting a voluntarily lower borrowing limit is a commitment device. I expect that consumers with standard preferences will not choose to limit their choice set, and that consumers who are naive about their own present focus will not choose to limit their choice set. It follows that only consumers who are (partially) sophisticated and present focused will use the commitment device. Figure 13 shows that there are no statistically significant differences between stratification bins in the proportion of consumers opting to use the commitment device. If present focus is positively correlated with pre-experimental demand, it must be driven by a greater proportion of naive present-focused consumers with high pre-experimental demand.

After the initial round of marketing, representatives from the solar company called consumers again to remind them about the line of credit. During the second round of calls, they asked a random subset of 596 treated consumers how many times they expected to use the line of credit over the next month, the likelihood that the consumer would use the line of credit one more time than they expected, and the likelihood that the consumer would use the line of credit one fewer time than expected. I compute the actual number of times each consumer used the line of credit over the month following the phone call and calculate the

difference between the consumer’s expectation and their actual use of the line of credit. If consumers are present focused and (partially) naive, I expect them to underestimate their use of the line of credit.

Across all consumers surveyed, only 1% underestimate their use of the line of credit. Table 5 shows that consumers with the lowest pre-experimental demand overestimate their use the most, but all groups of consumers expect to use the line of credit at least one more time per month than they actually do. Column (2) of table 5 shows variation between groups in consumers’ confidence in their prediction. 86.5% of consumers with the lowest pre-experimental demand believe that there is a less than 10% chance that they will use the line of credit one fewer time than they predict. Consumers in other stratification bins are less certain, with 60.4%-71.7% believing that there is a less than 10% chance that they will use the line of credit one fewer time than predicted.

Consumer predictions are not incentivized, but combined with multiple dimensions of heterogeneity, results of the RCT on incentives for solar payments, and evidence that there are no differences between stratification bins in the proportion of sophisticated present focused consumers, they help rule out present focus as the primary mechanism driving consumer responses to the line of credit. Beyond ruling out present focus, consumer predictions about borrowing provide additional support for the role guaranteed credit plays in reducing the precautionary savings motive. Survey responses show that consumers with the lowest pre-experimental demand expect to be liquidity constrained significantly more than consumers in other stratification bins even though, in practice, a much smaller proportion actually use the line of credit. Consistent with the model, consumers with the highest expectations of future liquidity constraints increase demand the most in response to being offered guaranteed access to credit.

5.2 Discussion

Taken together, the evidence supports a model where consumers buying a perishable good with transaction costs respond to the line of credit differentially depending on the severity of liquidity constraints they face. Transaction costs and liquidity constraints are common market frictions in low-income countries. Although the strict non-storability of PAYGo solar access time is an extreme case, and perishable good with transaction costs will force consumers to make similar trade-offs. My results point to a range of concerns for policymakers seeking to promote more efficient and equitable markets.

Typically, market frictions negatively impact both firms and consumers. The experimental results show that the solar firm collects less revenue from certain groups of consumers as a result of market frictions while receiving more from others. When I re-weight treatment

effects on the utilization rate to be representative of the distribution of consumers in figure 3, I find that offering the line of credit does not significantly increase revenue collection for the firm. Figure A5 shows that the line of credit creates similarly heterogeneous impacts on repossession. Consumers with the lowest pre-experimental demand are less likely to default but offering the line of credit to consumers with the highest pre-experimental demand increases the risk of repossession, although estimates are not statistically significant. In a partial equilibrium sense, offering the line of credit is not profitable for the firm.

In general equilibrium, the impact of transaction costs will vary based on market structures. A monopolist could use a two-part tariff or a menu of two-part tariffs to capture the surplus resulting from a reduction in transaction costs. For instance, a monopolist PAYGo provider could charge a higher down payment but then give consumers the enhanced flexibility of a product like the line of credit. Consumers who value flexibility could select into the high down payment option, while those who do not could select into the traditional PAYGo contract. Such a two part tariff would allow the firm to capture some of the surplus it is currently losing from consumers who use the line of credit to better target their purchases; however, liquidity constrained consumers will likely struggle to make higher down payments even when they value flexibility. In practice, a two part tariff may lead to smaller increases in demand among low-demand consumers than I find in my experiment, limiting the profitability from offering the line of credit but still leading to an outcome where consumer and producer surplus is higher than it is currently.

Firms in a competitive market will respond to transaction costs for perishable goods by competing away transaction costs up to the point where doing so is no longer profitable. Such competition could take a variety of forms in my setting. For instance, Hayes (1987) shows that firms in a competitive market may offer two-part tariffs when consumers face uncertain utility from a good, similar to how I have modeled the consumer's decision. In a competitive market, the two-part tariff entails some lump-sum fee and then a per unit price that falls below the marginal cost. This again raises the possibility that firms could offer a menu of contracts for consumers to select into based on the value they place on flexibility, although in the competitive case firms do not reap positive profits from the two part tariff. Alternatively, competition could directly lead to firms offering a product like the line of credit.

In my setting, the PAYGo solar market is relatively competitive: there are multiple providers offering similar products and contracts. Given that it is not profitable for the firm I work with to offer the line of credit, it may be that PAYGo solar firms have reached the point where competing on convenience is no longer profitable. Alternatively, we may see more firms moving toward products like the line of credit in the coming years as this

relatively new market equilibrates.

My results highlight a more general problem: transaction costs associated with buying a perishable good act as a tax on consumers, either because they pay high transaction costs or because they buy in bulk and generate waste. It follows that transaction costs for perishable goods generate deadweight loss relative to a world free of transaction costs. Transaction costs for perishable goods will persist in general equilibrium when the costs associated with reducing them outweigh the benefits. To be precise, the increase in demand among consumers previously buying in small increments plus any consumer surplus the firm can capture from eliminating waste for consumers buying in bulk must outweigh the costs associated with lowering transaction costs for the firm. In my setting, the costs to the firm primarily consist of the administrative costs associated with offering a broader menu of contracts or implementing the line of credit, but in many other cases the costs could be substantially larger. Consider consumers who live far from the nearest grocery store. Lowering the transaction costs associated with buying groceries for such consumers would necessitate that firms make investments like building new stores or starting delivery services, both of which would require a high willingness to pay for convenience among consumers to earn a positive rate of return.

The lost surplus associated with transaction costs for perishable goods points to a role for government to improve consumer welfare. If consumers are not liquidity constrained, then a policymaker could reduce the loss associated with transaction costs for a perishable good by subsidizing storage technology or, in the case of PAYGo solar, regulating contract types so that goods are not artificially perishable. However, low-income consumers are often liquidity constrained. Providing storage when consumers are liquidity constrained does not eliminate the inefficient trade-off between liquidity and transaction costs. Similarly, providing access to credit without improving storage does not resolve the tension between waste and transaction costs. The most direct policies to minimize the loss generated by transaction costs for perishable goods are those designed to directly reduce transaction costs. For instance, public investment in infrastructure or incentives for firms to reduce transaction costs.

Beyond the inefficiencies associated with transaction costs for perishable goods, there are critical considerations of equity. Firms selling perishable goods will not find it profitable to lower transaction costs when the costs of doing so outweigh the benefits. The costs of lowering transaction costs are likely to be high when transaction costs are high, such as in rural areas. The benefits are likely to be limited when the consumers who benefit are low-income or where population density is low and the resulting potential increase in demand is small. It follows that policies designed to lower transaction costs for perishable goods can lead to both more efficient and more equitable outcomes.

Methodologically, transaction costs and liquidity constraints distort observed willingness to pay, making revealed preferences measures of welfare inaccurate. In the next section, I re-estimate consumer welfare from electrification using the demand observed in my experiment where I reduce the impacts of liquidity constraints and transaction costs.

6 Welfare

Demand observed in the presence of market frictions does not provide an accurate measure of consumers' willingness to pay. The results from my experiment show that consumers significantly alter their demand for solar when I relax liquidity constraints and lower transaction costs. In this section, I use observed demand during the experiment to estimate a less distorted lower bound on consumer surplus from electricity.

I do not randomly vary the price of solar during the experiment, but I do randomly assign the fee charged on the line of credit. The fee affects the quantity of days consumers borrow and the quantity of days prepaid for over the course of the experiment. Those quantities combine with the randomly assigned fee to determine the average price a consumer pays for solar over the course of the experiment. I explicitly model the link between the exogenous fee, F , quantities demanded, and the average price paid by consumers to estimate a demand curve for solar under conditions of reduced market frictions.

Let Q_p be the number of days a consumer prepays for over the course of the experiment and Q_b be the number of days a consumer borrows. If $Q(F)$ is the total quantity of solar access demanded over the course of the experiment, then

$$Q(F) = Q_p(F) + Q_b(F)$$

and

$$P(F) = \frac{Q_p(F) + Q_b(F)(1 + F)}{Q_p(F) + Q_b(F)}.$$

The slope of the demand curve is $\frac{dP(F)}{dQ(F)} = \frac{dP(F)/dF}{dQ(F)/dF}$. Differentiating Q and P with respect to F , I get the following expressions.

$$\frac{dQ(F)}{dF} = \frac{dQ_p(F)}{dF} + \frac{dQ_b(F)}{dF}. \quad (13)$$

$$\frac{dP(F)}{dF} = \frac{\frac{dQ_p}{dF} + \frac{dQ_b}{dF}(1 + F) + Q_b(F)}{Q_p(F) + Q_b(F)} - \frac{Q_p(F) + Q_b(F)(1 + F)}{\left(\frac{dQ_p(F)}{dF} + \frac{dQ_b(F)}{dF}\right)^2}. \quad (14)$$

I can estimate $\frac{dQ_p(F)}{dF}$, $\frac{dQ_b(F)}{dF}$, and $\frac{dP}{dF}$ using the simple regressions

$$Q_{pi} = \alpha + \beta Fee_i + \delta X_i + \epsilon_i,$$

$$Q_{bi} = \alpha + \beta Fee_i + \delta X_i + \epsilon_i,$$

$$P_i = \alpha + \beta Fee_i + \delta X_i + \epsilon_i,$$

where X_i controls for the consumer's daily rate and pre-experimental demand to increase precision. The daily rate is the price for a day of solar, which varies depending on the number of appliances the consumer chose to include with their solar home system when making their initial adoption decision. Summing estimates for $\frac{dQ_p}{dF}$ and $\frac{dQ_b}{dF}$ yields $\frac{dQ(F)}{dF}$, which I combine with my estimate of $\frac{dP}{dF}$ to obtain an estimate of $\frac{dP(F)}{dQ(F)}$. I bootstrap all standard errors and confidence intervals.

Figure 14 shows the estimated slopes for each stratification bin. I cannot reject that demand is perfectly inelastic across the distribution of pre-experimental demand. To be conservative, I take the bottom of the 99% confidence interval around my estimates of the slope for each group in my estimation of consumer surplus.

I use the estimated demand curves along with observed demand to calculate a conservative lower bound on consumer surplus. I anchor the estimated demand curves at the total quantity demanded when the effective price is 1 for each stratification bin. To convert into monetary terms, I use the median daily rate, RWF 190 when the effective price is 1. I form a lower bound by only considering the bottom of the resulting Marshallian welfare triangle where I have empirical support for the price variation, as illustrated in figure 15.

Table 6 provides two lower bounds on consumer surplus. Column (1) shows a lower bound for consumer surplus at current prices, or a bound on the consumer surplus that solar customers would enjoy if they had access to the line of credit and paid current daily rates for solar. Inframarginal consumers enjoy the most consumer surplus, at least \$55 per household per year. As I move along the demand curve to increasingly marginal consumers, the lower bound on surplus drops under \$7 per year.

Column (2) shows a lower bound on consumer surplus if solar access time were fully subsidized so that consumers paid a price of zero. The lower bound in column (2) allows me to assess which groups of consumers have a willingness to pay exceeding the cost of the solar home system. The total value of the median PAYGo contract in my setting is \$230 paid over approximately 3.75 years, which includes the solar home system, basic appliances, maintenance, financing costs, and the labor associated with administering the PAYGo system. For the firm to break even, consumers need to have a willingness to pay of at least \$61 per year.

Column (2) in table 6 shows that 76% of current consumers have a lower bound on consumer surplus that is high enough for the firm to break even. Consumers with pre-experimental utilization of 30%-65% have a sufficiently high willingness to pay if I extend the demand curve an additional 55% beyond the range of prices with empirical support. Altogether, it seems plausible that around 88.5% of current consumers likely have a willingness to pay for solar that is high enough for the firm to break even when I examine the less distorted demand curves resulting from my experiment.

Importantly, consumers with the lowest pre-experimental demand do not have a high enough willingness to pay for solar even if I extrapolate beyond the empirical price support up to the intercept on the vertical axis. These marginal consumers point toward the challenge of electrifying the millions of rural households who have not selected into a PAYGo solar contract. They are also the consumers with the largest treatment effect from being offered the line of credit, indicating that both contract structures and prices have a role to play in making electricity accessible to such households.

The lower bounds in table 6 additionally facilitate comparisons to other recent measures of consumer surplus from electrification in the literature. In table 7, I take the weighted average of my estimated lower bound on consumer surplus and compare it to three recent estimates in the literature: Grimm et al (2020), Lee et al (2020), and Burgess et al (2020). My lower bound on consumer surplus is equal to or larger than estimates of total consumer surplus in all three papers, with the exception of the upper range of estimates in Lee et al (2020). If I extend the range of prices included in my demand curve or use my point estimates of the slope of the demand curve rather than the bottom of the 99% confidence interval, I obtain estimates for consumer surplus that are substantially higher than other estimates in the literature. My results suggest that consumer surplus from electricity is likely higher than previously believed because market frictions are distorting demand.

Multiple factors beyond reduced market frictions could contribute to the difference between my estimated lower bound on consumer surplus from electrification and other estimates in the literature. Grimm et al (2020) and Lee et al (2020) both derive their estimates from willingness to pay on the extensive margin. If consumers have imperfect information about the benefits of electrification, demand on the extensive margin will be lower than the demand I observe on the intensive margin. Unreliable supply on the grid could dampen demand for the Kenyan households in Lee et al (2020), and to a lesser extent Burgess et al (2020) relative to the solar home systems in my setting. However, differences in supply side reliability and extensive versus intensive margin demand both point toward my estimates providing less distorted estimates of consumer surplus.

The primary concern with my estimated lower bound is that my sample is positively

selected on willingness to pay for electricity: not every rural household in Rwanda chooses to sign a PAYGo solar contract. I attempt to mimic my positively selected sample in my comparison to Grimm et al (2020) by only considering the subset of consumers with a willingness to pay for solar that meets or exceeds market prices. In comparing to Lee et al (2020), I use the consumer surplus estimates that most closely reflect the types of appliances that can be powered by a solar home system. Unfortunately, I cannot mimic the positive selection in my comparisons to Burgess et al (2020), which likely explains at least part of the difference between my estimated lower bound on consumer surplus and their estimate for consumer surplus across the entire population of Bihar. The differences between my lower bound on consumer surplus and others in the literature may be overstated to the extent that I cannot accurately imitate the positive self-selection in my sample.

My results suggest that consumer surplus from electrification may be higher than previously believed for the subset of rural consumers with the highest value for electricity. Higher consumer surplus translates into a more attractive cost-benefit proposition for electrification. My welfare estimates cannot directly speak to potential consumer surplus from non-electrified households, but evidence from the marginal consumers in my sample suggests that they will likely require significant assistance to adopt and pay for electricity.

7 Conclusion

I highlight the unique problem consumers face when buying a perishable good with transaction costs and demonstrate the importance of liquidity constraints in shaping consumer responses to the problem. As my theoretical framework predicts, consumers in Rwanda respond to a line of credit for PAYGo solar access in a manner consistent with high transaction costs and heterogeneous liquidity constraints. Consumers who are most likely to be liquidity constrained increase demand in response to being offered the line of credit while consumers who previously bought in bulk significantly reduce demand.

Offering the line of credit is not profitable for the solar firm even though it enables consumers to better optimize their consumption of electricity. When transaction costs persist in competitive markets, they act as a tax on consumers and generate deadweight loss. Transaction costs persist when the costs of reducing them outweigh the benefits, which is likely to be true in rural areas or when firms are serving low-income consumers. It follows that policymakers can act to enhance efficiency and equity through policies designed to lower transaction costs.

Consumer responses to the line of credit show that market frictions significantly shape demand for electricity in low-income settings. Revealed preference measures of welfare from

electrification that cannot account for market frictions provide inaccurate estimates. Using the demand observed in my experiment, I find that consumer surplus from electrification is substantially higher than comparable estimates in the literature. A wide range of consumers in my sample have a willingness to pay for electricity that exceeds the cost of the PAYGo contract; however, demand among marginal consumers in my sample falls short of cost-covering levels. Given that the average consumer in my sample is wealthier than the average rural Rwandan household, universal electrification will likely require fiscal support such as subsidies. My work demonstrates that subsidies that build in flexible payment options will allow consumers to pay for more electricity and to better target their consumption, increasing the benefits of electrification while lowering the overall cost of subsidies.

Prepaid contracts with low-income households represent an attempt to provide services profitably in a challenging market environment. Prepayment allows for low cost contract enforcement in settings where institutions may be weak and the cost of enforcing contracts over small amounts of money are high. It also provides consumers with a degree of flexibility, allowing for non-penalized missed payments or demand reductions, up to a point. Despite these features, common market frictions like liquidity constraints and transaction costs force consumers to make costly trade-offs that shape their demand for prepaid goods and services. My work points to the continued need for innovation in contracts and products for low-income consumers, particularly in addressing liquidity constraints and transaction costs for rural consumers.

I offer three directions for future work. To achieve universal electrification, we need to understand more about demand for electricity along the intensive margin among marginal consumers. Even if grid connections or down payments for PAYGo contracts are subsidized for marginal consumers, they will not reap the full benefits of electrification if prices are too high on the intensive margin. My results suggest that willingness to pay will be low among such consumers, but that they could benefit significantly from having access to the type of short-term credit offered in my experiment. Better understanding intensive margin demand will facilitate better planning for universal electrification.

My work provides empirical support for the precautionary savings model in Deaton (1991), and shows that consumers engage in precautionary savings over short time horizons. Offering guaranteed access to small amounts of credit for such short-term consumption smoothing has traditionally been prohibitively costly, but digital credit has brought such services within reach. Better understanding under what conditions firms can provide guaranteed access to small amounts of credit for a broad range of consumers has the potential to substantially reduce critical market frictions and improve consumption smoothing.

Finally, my work highlights the inefficiencies created by transaction costs for perishable

goods. There is a large range of policies that could work to lower transaction costs: public investment in transportation and market infrastructure, government incentives for firms to locate in under-served areas, and subsidized transportation to name only a few. It will be important to understand when initial public investments can spur private investment and competition that leads to meaningful reductions in transaction costs for consumers. Similarly, combining theory and empirics to understand how to target public investments for the benefit of low-income consumers will provide critical insights to policymakers with limited funds. Alternatively, future work could focus on market conditions that will make investments in reducing transaction costs lucrative for firms. A number of PAYGo solar firms are now offering a wide range of products beyond solar home systems such as loans for school fees or additional services and appliances. If consumers have high enough demand for these expanded offerings, firm and consumer incentives around transaction costs may become better aligned. Designing effective policies to reduce transaction costs has the potential to foster more inclusive and equitable economic growth.

References

2018. “The Little Data Book on Financial Inclusion 2018.” Tech. rep., International Bank for Reconstruction and Development.
- Allcott, Hunt and Dmitry Taubinsky. 2015. “Evaluating behaviorally motivated policy: experimental evidence from the lightbulb market.” *American Economic Review* 105:2501–38.
- Ashraf, Nava, Dean Karlan, and Wesley Yin. 2006. “Deposit Collectors.” *The B.E. Journal of Economic Analysis and Policy* 5 (2):1–24.
- Aycinena, Diego, Claudia Martinez, and Dean Yang. 2010. “The Impact of Transaction Fees on Migrant Remittances: Evidence from a Field Experiment Among Migrants from El Salvador.”
- Beck, Thornsten, Demirguc-Kunt Asli, Maria Soledad, and Peria Martinez. 2008. “Banking Services for Everyone? Barriers to Bank Access and Use around the World.” *World Bank Research Observer* 22 (3):397–430.
- Beck, Thornsten, Demirguc-Kunt Asli, Maria Soledad, Peria Martinez, and Maria Soledad. 2007. “Reaching Out: Access to and use of banking services across countries.” *Journal of Financial Economics* 85 (1):234–266.
- Bensch, Gunther, Jochen Kluge, and Jorg Peters. 2011. “Impacts of Rural Electrification in Rwanda.” *Journal of Development Effectiveness* 3 (4):567–588.
- Bernheim, Douglas and Antonio Rangel. 2009. “Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics.” *Quarterly Journal of Economics* 124:51–104.
- Burgess, Robin, Michael Greenstone, Nicholas Ryan, and Anant Sudarshan. 2020. “Demand for Electricity on the Global Electrification Frontier.”
- Burlig, Fiona and Louis Preonas. 2016. “Out of the Darkness and into the Light? Development Effects of Rural Electrification.” EI @ Haas WP 268.
- Chaplin, Duncan, Arif Mamun, Ali Protik, John Schurrer, Divya Vohra, Kristine Bos, Hannah Burak, Laura Meyer, Anca Dumitrescu, Christopher Ksoll, and Thomas Cook. 2017. “Grid Electricity Expansion in Tanzania by MCC: Findings from a Rigorous Impact Evaluation.” Tech. rep., Mathematica Policy Research.

- Collins, Daryl, Jonathan Morduch, Stuart Rutherford, and Orlanda Ruthven. 2009. *Portfolios of the Poor: How the World's Poor Live on \$2 a Day*. Princeton University Press.
- Deaton, Angus. 1991. "Savings and Liquidity Constraints." *Econometrica* 59 (5):1221–1248.
- Dinkleman, Taryn. 2011. "The Effects of Rural Electrification on Employment: New Evidence from South Africa." *American Economic Review* 101 (7):3078–3108.
- Dupas, Pascaline, Dean Karlan, Jonathan Robinson, and Diego Ubfal. 2018. "Banking the Unbanked? Evidence from Three Countries." *American Economic Journal: Applied Economics* 10 (2):257–297.
- Francis, Eilin, Joshua Blumenstock, and Jonathan Robinson. 2017. "Digital Credit: A Snapshot of the Current Landscape and Open Research Questions."
- Grimm, Michael, Luciane Lenz, Jorg Peters, and Maximiliane Sievert. 2020. "Demand for Off-Grid Solar Electricity: Experimental Evidence from Rwanda." *Journal of the Association of Environmental and Resource Economists* 7:417–454.
- Hayes, Beth. 1987. "Competition and Two-Part Tariffs." *The Journal Of Business* 60:41–54.
- Jack, Kelsey and Grant Smith. 2020. "Charging ahead: Prepaid electricity metering in South Africa." *American Economic Journal: Applied Economics* 12 (2):134–168.
- Jack, William and Tavneet Suri. 2014. "Risk Sharing and Transactions Costs: Evidence from Kenya's Mobile Money Revolution." *American Economic Review* 104 (1):183–223.
- Khandker, Shahidur, Douglas Barnes, and Hussain Samad. 2009. "Welfare Impacts of Rural Electrification: A Case Study from Bangladesh." WPS 4859.
- Khandker, Shahidur, Hussain Samad, Rubaba Ali, and Douglas Barnes. 2014. "Who Benefits Most from Rural Electrification? Evidence in India." *The Energy Journal* 35 (2):76–96.
- Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112:443–478.
- Lane, Gregory. 2020. "Credit Lines as Insurance: Evidence from Bangladesh."
- Lee, Kenneth, Edward Miguel, and Catherine Wolfram. 2020. "Experimental Evidence on the Economics of Rural Electrification." *Journal of Political Economy* 128 (4):1523–1565.

- Lenz, Luciane, Anciet Munyehirwe, and Maximiliane Sievert. 2017. "Does Large-Scale Infrastructure Investment Alleviate Poverty? Impacts of Rwanda's Electricity Access Roll-Out Program." *World Development* 89:88–110.
- Lipscomb, Molly, Mushfiq Mobarak, and Tania Barham. 2013. "Development Effect of Electrification: Evidence from the Topographic Placement of Hydropower in Brazil." *American Economic Journal: Applied Economics* 5 (2):200–231.
- O'Donoghue, Ted and Matthew Rabin. 1999. "Doing it Now or Later." *American Economic Review* 89:103–124.
- Suri, Tavneet, William Jack, and Thomas Stoker. 2012. "Documenting the birth of a financial economy." *Proceedings of the National Academy of Sciences* 109 (26):10257–10262.
- van de Walle, Dominique, Martin Ravallion, Vibhuti Mendiratta, and Gayatri Koolwal. 2017. "Long-Term Gains from Rural Electrification in Rural India." *World Bank Economic Review* 31 (2):385–411.
- Zollman, Julie, Daniel Waldron, Alexander Sotiriou, and Anne Gachoka. 2017. "Escaping Darkness: Understanding Consumer Value in PAYGo Solar." Tech. rep., CGAP Access to Finance Forum.

8 Tables

Table 1: Average Pre-Experimental Payment Sizes by Stratification Bin

	<i>Dependent variable:</i>
	Mean Purchase Size (Days)
0%-30%	-8.304*** (0.538)
30%-65%	-2.970*** (0.296)
65%-80%	-3.535*** (0.353)
Daily Rate	-0.011*** (0.001)
Intercept (80%-100% Mean Payment Size)	15.425*** (0.259)
Observations	1,342

Note:

*p<0.1; **p<0.05; ***p<0.01

Notes: The mean purchase size is the average number of days a consumer bought in a single transaction in the 90 days prior to the start of the experiment. The daily rate is the price the consumer pays for a day of access time, which consumers select into based on the number of appliances they choose to include with their solar home system.

Table 2: Differences in Borrowing Behavior between Stratification Bins

	<i>Dependent variable:</i>		
	Borrowed for Solar (1)	Amount Borrowed (2)	Cannot Borrow (3)
0%-30%	-0.073* (0.044)	-1,458.365 (2,954.986)	0.051*** (0.015)
30%-65%	-0.006 (0.038)	2,367.389 (2,558.273)	0.008 (0.013)
65%-80%	-0.016 (0.046)	5,832.380* (3,078.833)	0.030* (0.016)
Constant	0.554*** (0.019)	4,440.067*** (1,268.061)	0.019*** (0.006)
Observations	1,229	1,229	1,229
R ²	0.002	0.004	0.011
Adjusted R ²	-0.0001	0.001	0.009

*p<0.1; **p<0.05; ***p<0.01

Notes: Data come from a phone survey conducted in March, 2019. Consumers self-report whether they have ever borrowed to pay for solar, the amount that they have borrowed, and, if they did not borrow, the reason why. Columns (1) and (3) are simple linear probability models and column (2) is an OLS regression with the amount borrowed (in RWF) on the left hand side. I determine the stratification bin for each consumer in the phone survey by computing their utilization rate over the 90 days prior to the experiment.

Table 3: Heterogeneity in Take-Up by Distance to Mobile Money Agent

	<i>Dependent variable:</i>
	Take-up Rate
Hours to Reach MM Agent	0.058** (0.024)
90-Day Pre-experimental Utilization Rate	0.135*** (0.034)
Daily Rate (RWF)	0.0003** (0.0001)
Hi Fee	-0.012 (0.024)
Hi Borrowing Limit	-0.007 (0.024)
Repayment Time Limit	-0.021 (0.024)
Intercept	0.120*** (0.045)
Observations	1,342
R ²	0.016
Adjusted R ²	0.011

Notes:

*p<0.1; **p<0.05; ***p<0.01
Standard errors are White robust.

Table 4: Cross-Randomized Experimental Design for Solar Incentives

	Bulk Incentive		Monthly Reward	
	4 Week Minimum	5 Week Minimum	4 Week Minimum	5 Week Minimum
Low Reward	1 free day	3 free days	1 free day	3 free days
High Reward	2 free days	4 free days	2 free days	4 free days

Notes: Each cell contains 200 current solar customers, stratified by pre-experimental utilization rates. The table shows the minimum qualifying threshold for consumers in each group to receive any free days of solar, but consumers were offered a schedule of increasing rewards for increasingly large purchases. In practice, the number of consumers who qualify for rewards above the minimum is trivial.

Table 5: Differences in Consumer Expectations and Realizations of Credit Use

	<i>Dependent variable:</i>	
	Expected Less Actual Use	High Confidence
	(1)	(2)
30% - 65%	-0.132* (0.051)	-0.177*** (0.057)
65%-80%	-0.208*** (0.051)	-0.148** (0.058)
80%-100%	-0.100 (0.052)	-0.261*** (0.057)
Intercept	1.188*** (0.034)	0.865*** (0.044)
Observations	596	547
R ²	0.013	0.038
Adjusted R ²	0.008	0.033

*p<0.1; **p<0.05; ***p<0.01

Notes: Expected less actual use is the difference between the number of times a consumer expected to use the line of credit over the month and the actual number of times the consumer used the line of credit over the course of the month. High confidence is a dummy variable equal to one if the consumer stated that there was less than a 10% chance that they would use the line of credit one fewer time than they predicted.

Table 6: Consumer Surplus from Solar

Pre-Experimental Utilization Rate	Estimated Slope (99% CI Lower Bound)	Q ₁ (weighted)	Max % Δ P	CS Lower Bound (hh/year), current prices (1)	CS Lower Bound (hh/year), fully subsidized (2)
0%-30%	-0.038	15,487	2.5%	\$6.99	\$13.84
30% - 65%	-0.073	56,473	5.1%	\$24.76	\$48.35
65% - 80%	-0.107	54,693	6.2%	\$34.51	\$67.03
80% - 100%	-0.110	709,445	4.5%	\$55.33	\$108.27
Weighted Mean			4.5%	\$44.31	\$86.67

Notes: I calculate welfare only for the population of current solar customers, consisting of 50,000 households. The weights are 11.46% for the 0%-30% stratification bin, 12.13% for the 30%-65% bin, 8.52% for the 65%-80% bin, and 67.89% for the 80%-100% bin. I assume 1 USD = 900 RWF.

Table 7: Consumer Surplus from Solar

Source	Location	Electricity Type	Price	CS (hh/year)
This paper	Rural Rwanda	PAYGo Solar	Zero	\$86.67
This paper	Rural Rwanda	PAYGo Solar	Marginal Cost	\$44.31
Grimm et al (2020)	Rural Rwanda	Solar home system	Zero	\$89.50
Lee et al (2020)	Western Kenya	Grid	Marginal Cost	\$23.40 - \$331
Burgess et al (2020)	Bihar, India	All sources	Marginal Cost	\$6.99

Notes: Grimm et al (2020) estimate demand for solar home systems when consumers can, at most, spread payments out over 5 months. I only consider households in their sample with a willingness to pay over \$120 overall in order to mimic the selection of consumers into my sample. I assume a discount rate of 15% and assume that the solar home systems will function well for only three years to get my hh/year estimate of \$89.50. Lee et al (2020) present a range of estimates depending on demand elasticities, and assume a 15% discount rate and a 30 year asset life for a grid connection. I only compare my estimates to their estimates for consumers with relatively low electricity consumption (table 4 columns 1 and 2), as these are most comparable to the rural consumers in my setting. Burgess et al (2020) provide estimates of CS for all consumers, including those who have not adopted electricity. Lacking the full demand curve in their setting, I cannot mimic the sample selection present in my experiment, so part of the difference in estimates is likely attributable to my positively self-selected sample.

9 Figures

Figure 1: Consumer Travel Times to Nearest Mobile Money Agent

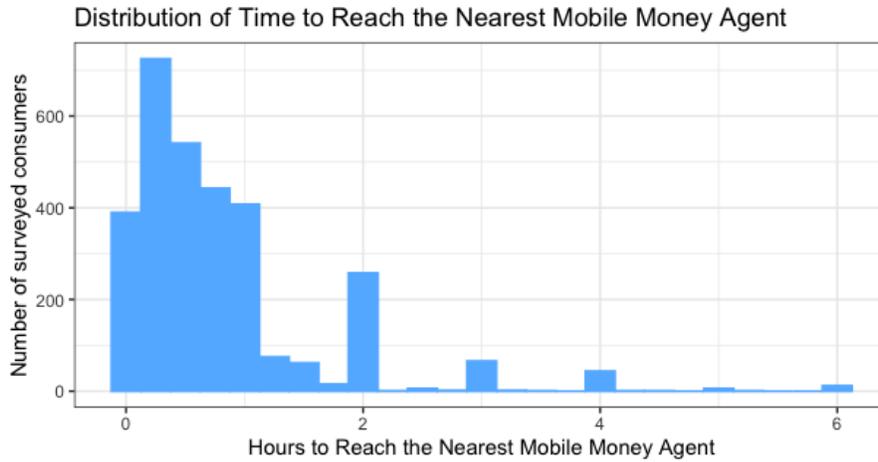


Figure 2: Self-Reported Use of Mobile Money Agents to Buy Solar

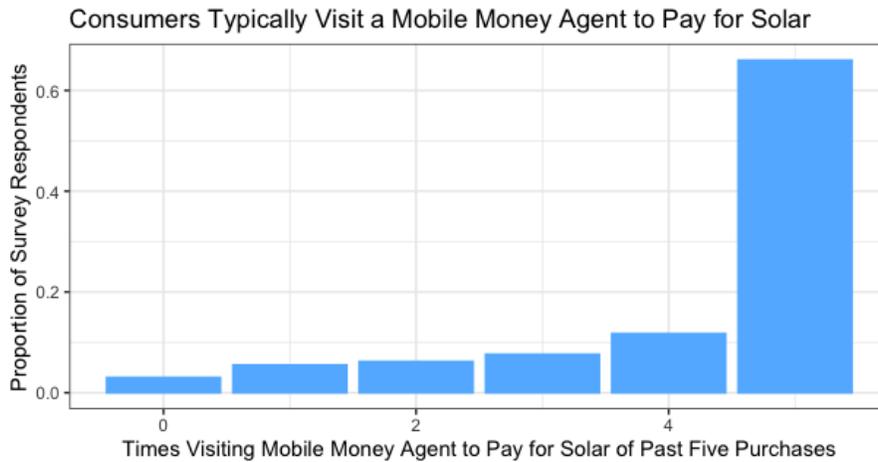


Figure 3: Distribution of Demand Prior to the Experiment

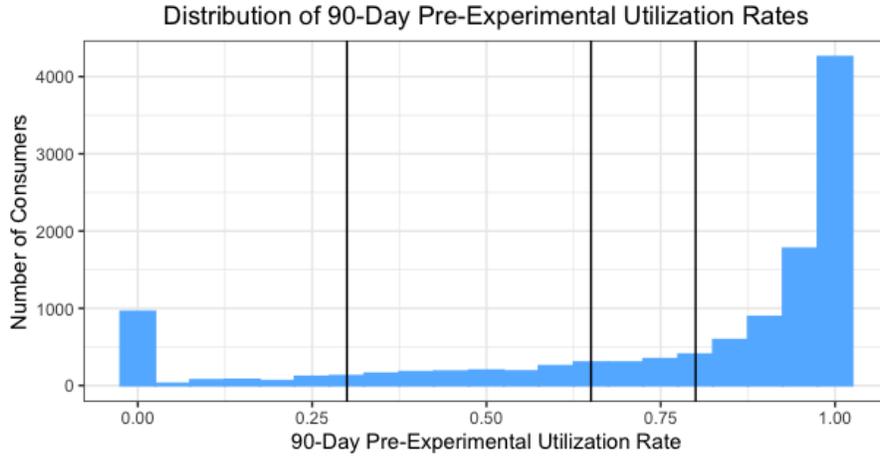
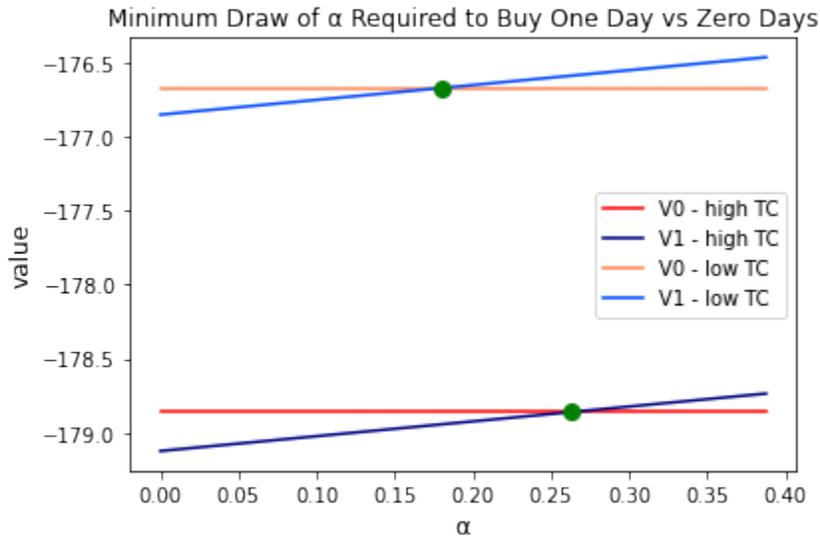
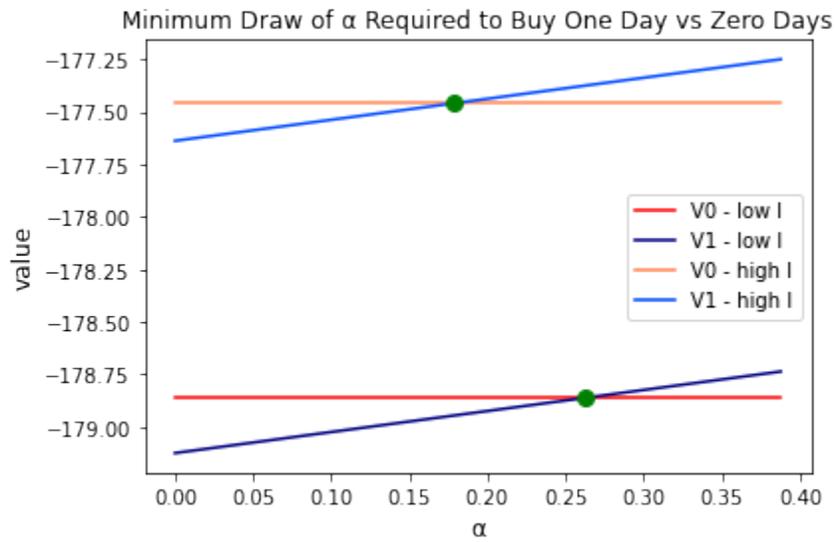


Figure 4: Change in α^* from reducing transaction costs.



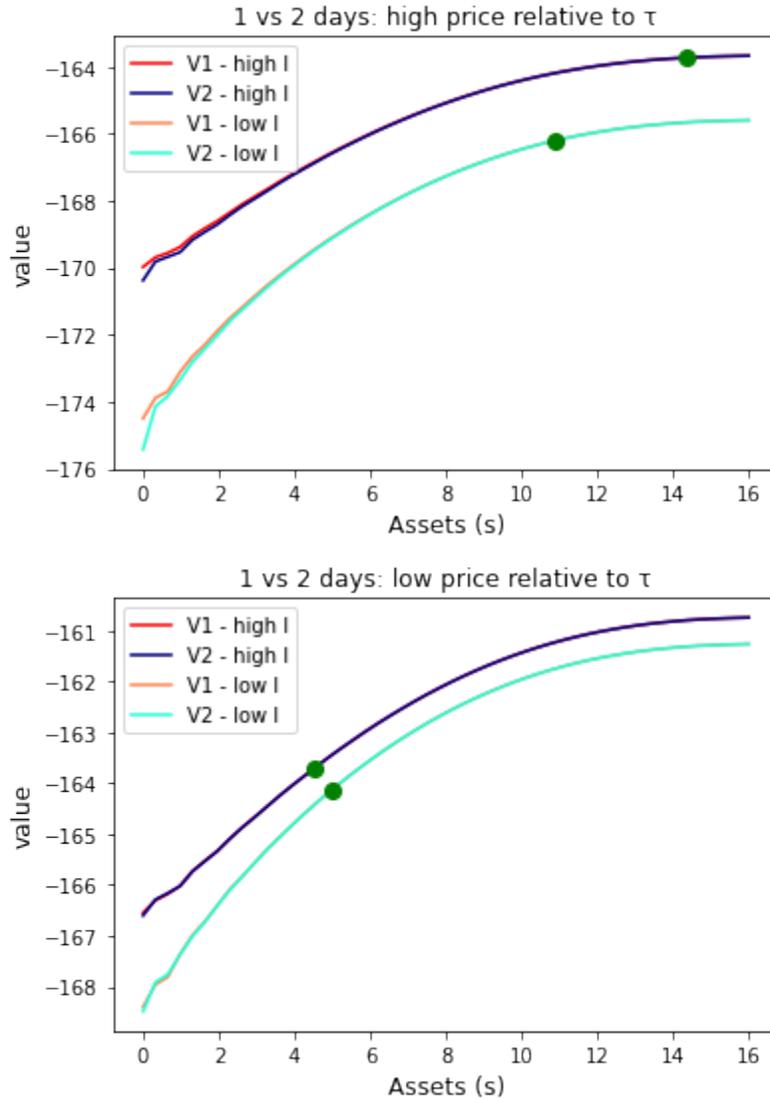
Note: The vertical axis is the level of the value function, with V_0 corresponding to the value function associated with forgoing access to electricity in the current period and V_1 corresponding to the value function associated with buying one day of access. The horizontal axis is the realization of α in the current period. Value functions are plotted for a given realization of assets and income, although the figure remains qualitatively similar for alternative realizations of assets and income. α^* is the minimum realization of α for which the consumer prefers to buy a day of solar access rather than forgoing it, or the point where V_0 and V_1 intersect. Low TC shows the value functions during the experiment where I lower transaction costs.

Figure 5: Change in α^* from relaxing liquidity constraints.



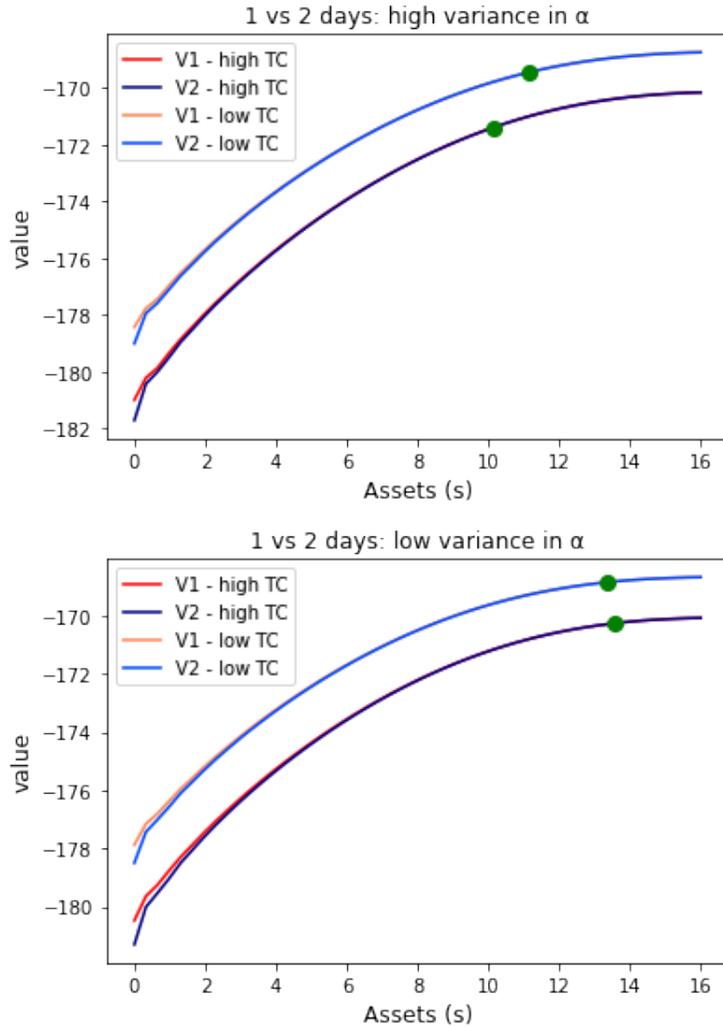
Note: The vertical axis is the level of the value function, with V0 corresponding to the value function associated with forgoing access to electricity in the current period and V1 corresponding to the value function associated with buying one day of access. The horizontal axis is the realization of α in the current period. Value functions are plotted for a given realization of assets and income, although the figure remains qualitatively similar for alternative realizations of assets and income. α^* is the minimum realization of α for which the consumer prefers to buy a day of solar access rather than forgoing it, or the point where V_0 and V_1 intersect. High I shows the value functions during the experiment where I relax liquidity constraints.

Figure 6: Change in the single vs bulk decision from relaxing liquidity constraints.



Note: The vertical axis is the level of the value function. The horizontal axis is the realization of assets s in the current period. The green points show the minimum realization of assets for which the consumer prefers to buy in bulk ($q = 2$) to buying a single day. It follows that when the point of intersection moves to the right, the consumer has a lower probability of buying in bulk. V1 corresponds to the value function when the consumer buys one day of access and V2 corresponds to the value function when the consumer buys in bulk. High I shows the value functions under the experimental condition where I relax liquidity constraints.

Figure 7: Size of the treatment effect for consumers with a high vs low variance in α



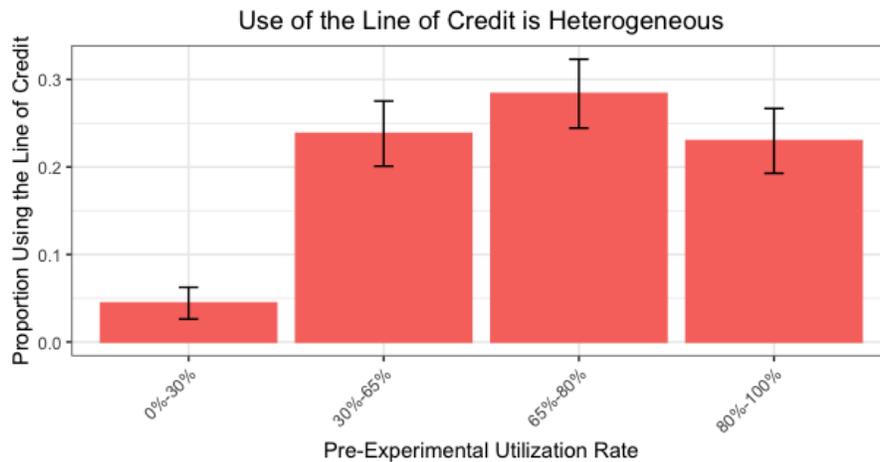
Note: The vertical axis is the level of the value function. The horizontal axis is the realization of assets s in the current period. The green points show the minimum realization of assets for which the consumer prefers to buy in bulk ($q = 2$) to buying a single day. It follows that when the point of intersection moves to the right, the consumer has a lower probability of buying in bulk. V1 corresponds to the value function when the consumer buys one day of access and V2 corresponds to the value function when the consumer buys in bulk. Low TC shows the value functions under the experimental condition where I lower transaction costs.

Figure 8: Heterogeneous Average Treatment Effects on Utilization



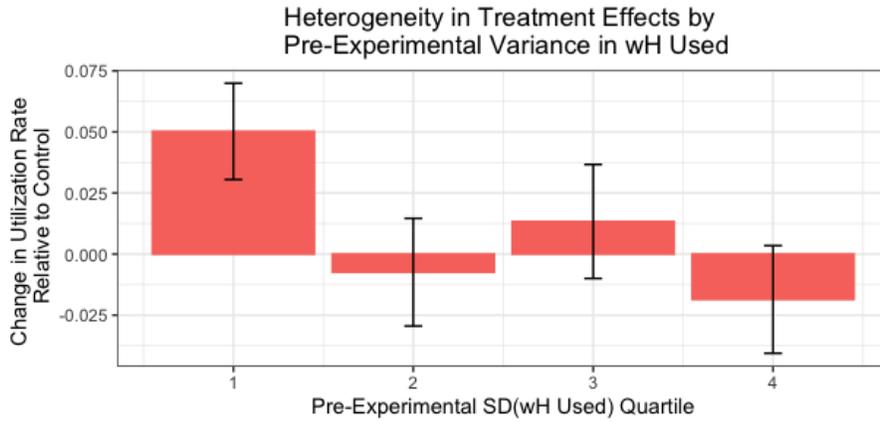
Note: 95% confidence intervals calculated using standard errors clustered at the level of the individual consumer. Estimates for the 30%-65% and 80%-100% bins are not significant at the 5% level after pre-registered multiple inference corrections.

Figure 9: Heterogeneous Average Treatment Effects on Utilization



Note: 95% confidence intervals calculated using White robust standard errors. All estimates remain statistically significant after applying pre-registered multiple inference corrections.

Figure 10: Heterogeneous Impacts by Pre-Experimental Variance in Use



Note: 95% confidence intervals calculated using standard errors clustered at the level of the individual consumer. Estimates for the first quartile remain statistically significant after applying pre-registered multiple inference corrections.

Figure 11: No Correlation Between Pre-experimental Demand and Variance in Use

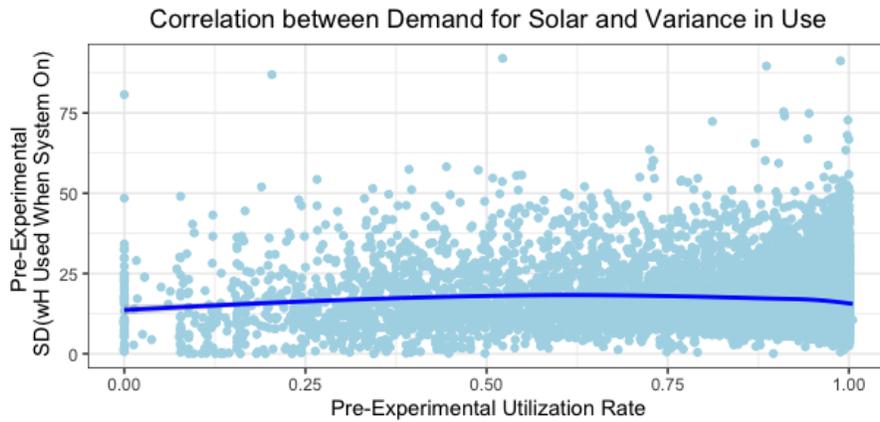
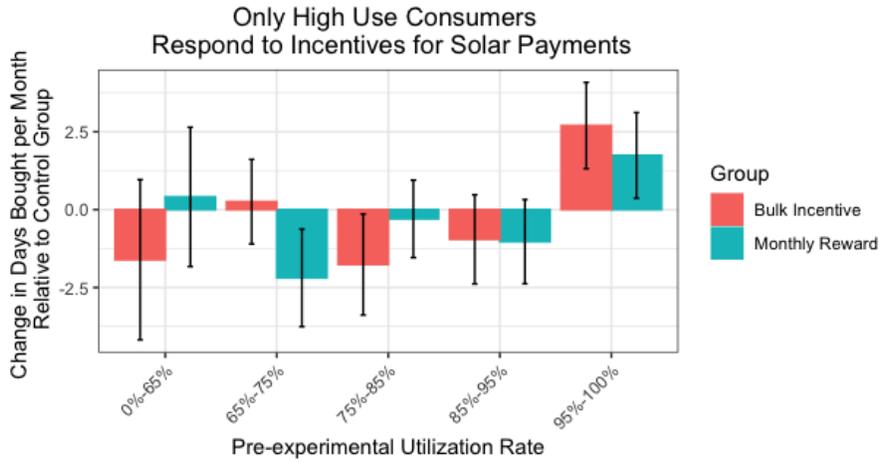


Figure 12: Heterogeneous Impacts of Incentives for Solar Payments



Note: I pool across minimum qualifying purchase sizes and reward sizes to increase power. 95% confidence intervals calculated with standard errors clustered at the level of the individual consumer.

Figure 13: No Significant Differences in Use of the Commitment Device

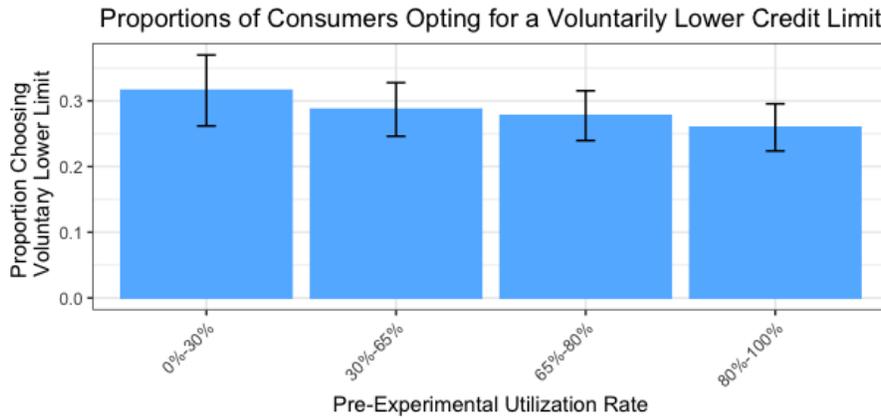


Figure 14: Estimated Slopes of the demand curve for solar

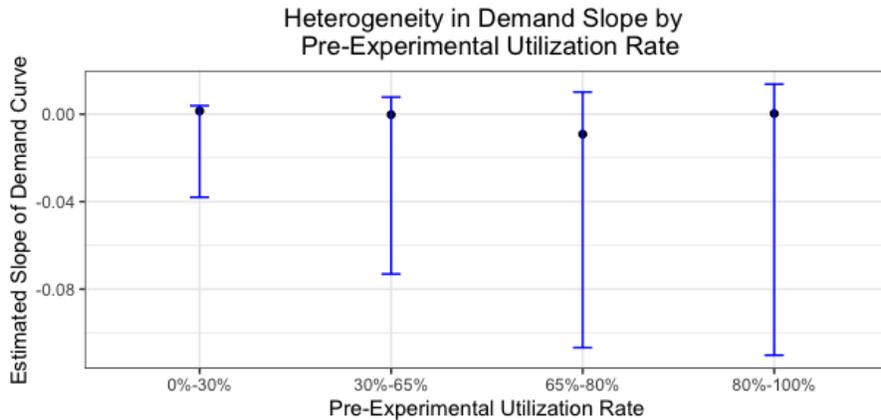
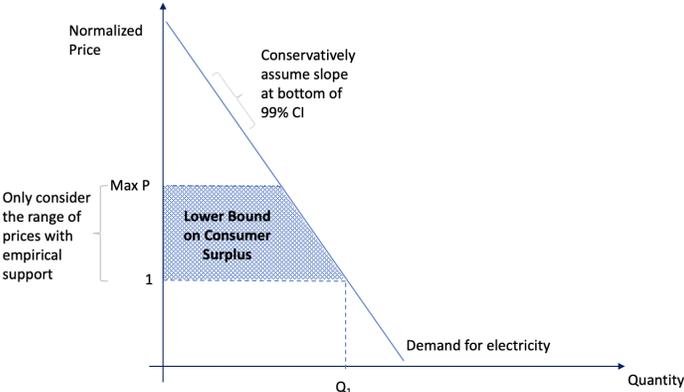
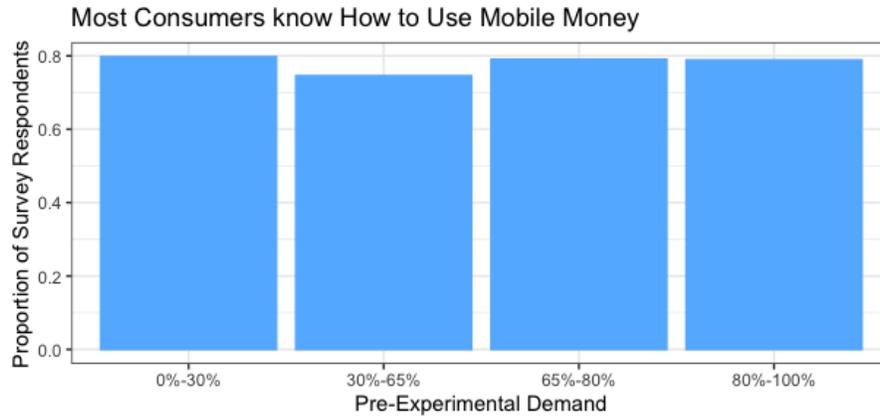


Figure 15: Estimated Slopes of the demand curve for solar



10 Appendix

Figure A1: Self-Reported Knowledge of Using Mobile Money to Buy Solar



Note: The figure shows the proportion of consumers in each stratification bin who answer "yes" to the question, "If you had mobile money already in your account and you wanted to use it to pay for solar, do you know how you would do that?"

Figure A2: Distribution of Payment Sizes 90 Days Prior to the Experiment

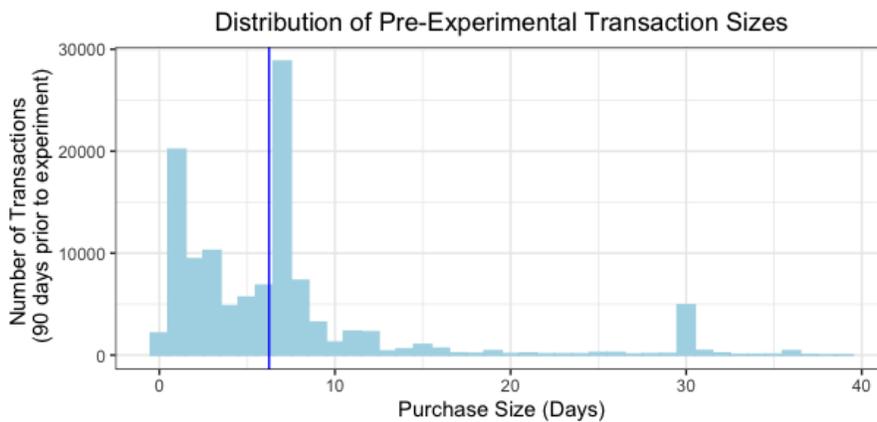
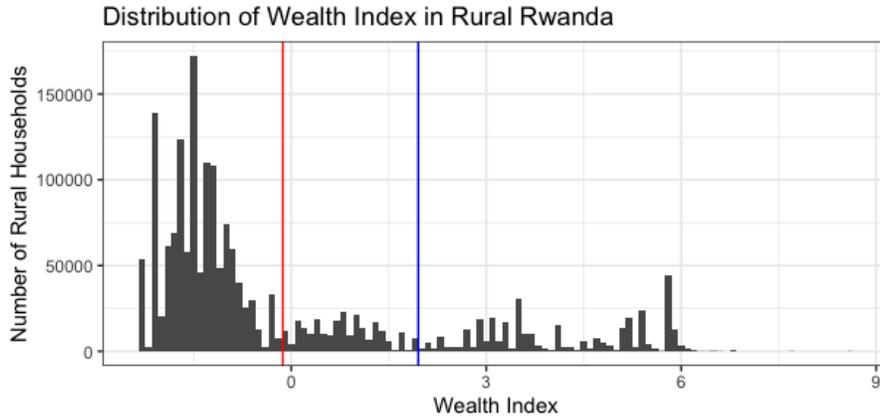
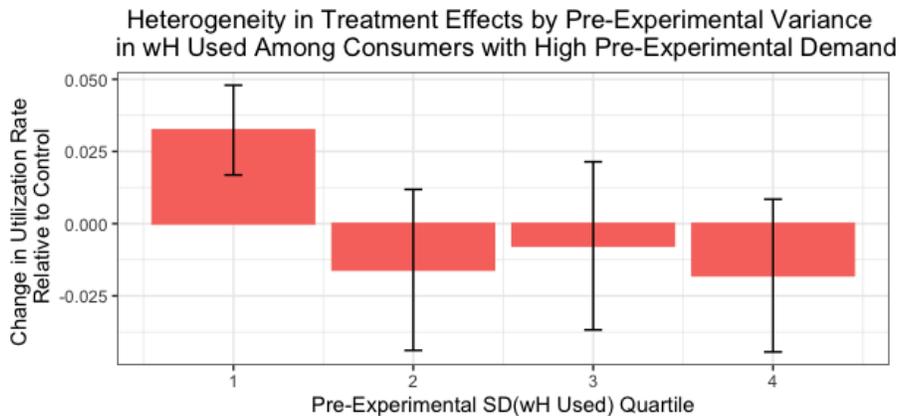


Figure A3: Distribution of Wealth Among Rural Households in Rwanda



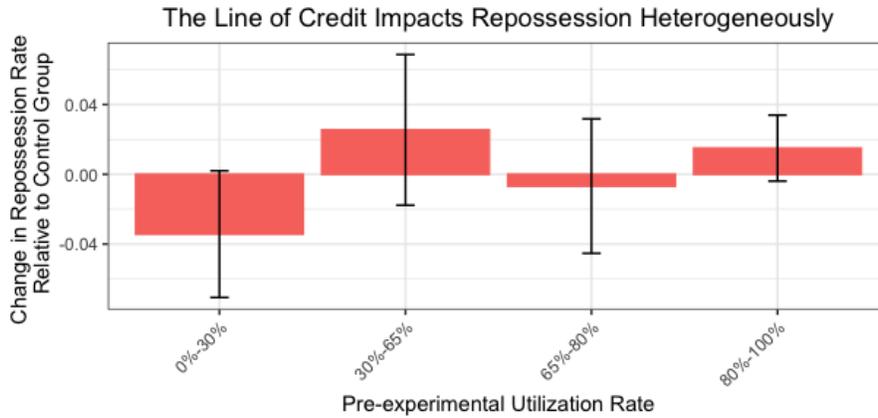
Note: The figure shows the distribution of wealth indices for a nationally representative sample of rural households in Rwanda. The red line indicates the mean level of wealth for the nationally representative sample and the blue line indicates the mean level of wealth for the consumers in my experimental sample.

Figure A4: Heterogeneous Treatment Effects: High Pre-Experimental Demand



Note: The figure shows heterogeneous treatment effects based on pre-experimental variance in wH used only for consumers with the highest pre-experimental demand. The general that I find in the entire sample holds here, among consumers who have access to their solar home systems on most days in the pre-experimental period. I can still reject that treatment effects are the same between quartiles one and four.

Figure A5: Heterogeneous Impacts on Default by Pre-Experimental Demand



Note: I measure the repossession rate as the proportion of consumers eligible for repossession one month after the experiment ends. 95% confidence intervals calculated using White robust standard errors. Estimates are not statistically significant at the 5% level.

Table A1: Balance Table

	<i>Dependent variable:</i>					
	90-day Pre UR	Daily Rate	Mean wH	SD wH	Tenure	Pmt Size
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.0002 (0.002)	1.798 (2.292)	0.407 (0.665)	0.491 (0.417)	-2.194 (5.371)	0.284 (0.368)
Bin 1	-0.901*** (0.002)	30.931*** (2.627)	-18.581*** (1.128)	1.149 (0.718)	38.890*** (6.055)	-8.767*** (0.417)
Bin 2	-0.466*** (0.002)	18.141*** (2.546)	-9.840*** (0.686)	1.556*** (0.431)	29.601*** (5.943)	-3.225*** (0.407)
Bin 3	-0.230*** (0.002)	13.721*** (3.008)	-6.631*** (0.812)	1.747*** (0.509)	12.471* (7.058)	-3.732*** (0.485)
Constant	0.957*** (0.001)	200.303*** (0.962)	53.705*** (0.254)	16.436*** (0.159)	451.822*** (2.256)	13.157*** (0.154)
Observations	11,605	11,201	10,448	10,427	11,695	11,605
R ²	0.954	0.017	0.045	0.003	0.005	0.042
Adjusted R ²	0.954	0.017	0.044	0.003	0.005	0.042

Note:

*p<0.1; **p<0.05; ***p<0.01

Table A2: Differences in Wealth and Energy Spending between Stratification Bins

	<i>Dependent variable:</i>		
	Wealth Index (1)	Connected to Grid (2)	Non-electricity Weekly Energy Expenditures (3)
0%-30%	-0.067 (0.113)	0.034** (0.015)	242.376** (118.010)
30%-65%	-0.304*** (0.097)	0.003 (0.013)	41.471 (102.167)
65%-80%	-0.113 (0.117)	0.012 (0.015)	-68.493 (122.956)
Constant	2.008*** (0.048)	0.023*** (0.006)	314.143*** (50.641)
Observations	1,208	1,229	1,229
R ²	0.008	0.004	0.004
Adjusted R ²	0.006	0.002	0.002

*p<0.1; **p<0.05; ***p<0.01

Notes: I use the following variables to construct the wealth index: ubudehe category, roof material, wall material, floor material, primary source of electricity (if any), primary source of light, whether or not the household is connected to the national grid, and weekly energy expenditures. Column (2) is a linear probability model with a dummy variable for grid access on the left hand side. Consumers with low pre-experimental demand tend to have lower wealth scores than those with higher pre-experimental demand, but the pattern is not monotonic across the distribution of pre-experimental demand. Columns (2) and (3) suggest an explanation: consumers with the lowest pre-experimental demand are also slightly more likely to have a grid connection. For a small subset of consumers, low demand could be the result of substitution between electricity sources.