

# Contracting in the Presence of Poor Storage, Transaction Costs, and Liquidity Constraints

Megan Lang \*

JOB MARKET PAPER

([latest version](#))

**Preliminary; please do not cite without author's permission**

September 1, 2020

## Abstract

Consumers face a double bind when buying a good that does not store well and involves transaction costs. Buying in bulk minimizes transaction costs but creates waste. Eliminating waste by making small purchases raises costs, but may be the only option available to liquidity constrained consumers. I explore consumer responses to this problem using pay as you go (PAYGo) solar access time in Rwanda, a strictly non-storable good. I randomly offer 2,000 current solar customers a line of credit for PAYGo solar access time. In practice, the line of credit both reduces liquidity constraints and lowers transaction costs. Responses to the line of credit are strikingly heterogeneous: changes in demand range from -6.4% to 88%. Given that liquidity constraints and transaction costs render willingness to pay an inaccurate measure of consumer welfare, I estimate consumer surplus from electricity under the less distorted conditions created by my experiment. My estimates suggest that the cost-benefit proposition for universal electrification is more attractive than previously believed, but that marginal households' willingness to pay for electricity is still well below current cost-covering levels.

---

\*University of California, Berkeley; [megan.lang@berkeley.edu](mailto:megan.lang@berkeley.edu). I thank the International Growth Centre for generous financial support. Thanks to Marco Gonzalez-Navarro, Meredith Fowlie, Alain de Janvry, Ethan Ligon, Jeremy Magruder, Edward Miguel, 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 rapidly depreciate through storage, theory predicts that consumers will buy small quantities frequently to avoid waste. Transaction costs for goods with poor storage create a trade-off. Consumers can buy in bulk to avoid repeatedly incurring transaction costs, but doing so generates waste. Buying small quantities eliminates waste but imposes a heavy burden of transaction costs. Regardless of the constrained optimum, liquidity constraints push consumers towards frequent, small purchases, effectively increasing the price they face. Rapid depreciation combined with transaction costs and liquidity constraints limit consumer surplus and impact consumers' willingness to pay, biasing revealed preference measures of consumer welfare.

I implement a randomized control trial (RCT) with consumers of pay as you go (PAYGo) solar in Rwanda to study consumer behavior in the presence of rapid depreciation, transaction costs, and liquidity constraints. Consumers with PAYGo solar contracts make a small down payment to have a solar home system installed, then “pay as they go” to use it by purchasing system access time with mobile money. I document that consumers in Rwanda face significant transaction costs when buying access time. Once purchased, access time runs down continuously: a consumer cannot buy a day of access time and save it for later, they have to start using it immediately. While PAYGo solar provides a useful illustration, consumers may 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 significant transaction costs for rural consumers and can store poorly in practice if consumers are inattentive or lack tools to track consumption. As a strictly non-storable good, PAYGo access time starkly illustrates the trade-offs consumers make between waste and transaction costs.

I consider how consumers trade off waste and transaction costs by incorporating a non-storable good with transaction costs into Deaton's (1991) precautionary savings model. The model generates three key predictions. First, providing credit and lowering transaction costs unambiguously increases demand for the non-storable good among liquidity constrained consumers. Guaranteeing access to credit increases demand among liquidity constrained consumers regardless of whether they actually borrow by reducing the precautionary savings motive: consumers who know they have access to credit do not need to save as much to smooth consumption over time. Second, when transaction costs are high enough, consumers who are not severely liquidity constrained choose to buy in bulk even though doing so generates some waste. <sup>1</sup> It follows that reducing transaction costs may lead to lower demand

---

<sup>1</sup>I assume that consumers do not have a preference for waste minimization and that there are no variable

for the non-storable good among select consumers by allowing them to eliminate the wasteful part of their consumption. Third, reducing transaction costs will lead to greater reductions in demand among consumers who have a higher variance in the utility they realize from the non-storable good.

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 relaxes short-term liquidity constraints, but also directly lowers transaction costs since consumers call the solar company to use it rather than visiting a mobile money agent.<sup>2</sup> I stratify my sample based on pre-experimental demand, which I show to be a proxy for the severity of liquidity constraints facing consumers.

Taken together, my results paint a picture of consumers paying transaction costs for a non-storable good who optimize their consumption differently based on the severity of liquidity constraints they face. Consumer responses to the line of credit vary strikingly across the distribution of pre-experimental demand, with average treatment effects ranging from -6.4% to 88%. Average treatment effects vary in a manner consistent with the three key predictions in my model. First, the largest increases come from consumers with low demand prior to the experiment, who I show are most likely to be liquidity constrained. Average treatment effects cannot be driven solely by consumers who use the line of credit. Guaranteed access to credit, irrespective of actual borrowing, changes behavior for consumers who are most likely to be liquidity constrained. Second, consumers who previously bought in bulk reduce the quantity of access time they buy. Third, treatment effects are decreasing in the pre-experimental variance in electricity use on purchased days, which I use to proxy for the variance in utility realized from electricity access. My results suggest that the line of credit relaxes liquidity constraints for low-demand consumers while allowing high-demand consumers to buy access time to more precisely meet their needs, reducing waste by lowering transaction costs.

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 provide further descriptive evidence showing that use of the line of credit is positively associated with distance to the nearest mobile money agent, indicating that consumers facing higher transaction costs use the line of credit more. I go on to consider an alternative underlying mechanism that may generate the same results. If pre-experimental demand is positively correlated with present focus, then the negative treatment effects I estimate could be the result of high-demand consumers

---

transport costs associated with buying in bulk.

<sup>2</sup>Note that the line of credit may indirectly lower transaction costs for liquidity constrained consumers by allowing them to buy in bulk.

borrowing then procrastinating on repayment. In other words, the line of credit could be causing some consumers to borrow, delay repayment due to present focus, and ultimately buy less electricity than they would under a strictly prepaid regime.

I use results from a separate experiment to test whether demand is positively correlated with present focus. I randomly offer consumers a bulk discount and a monthly reward for solar purchases. Although the monthly reward and the bulk discount offer equivalent average price reductions, the bulk discount requires consumers to incur large costs in the present for benefits far into the future relative to the monthly reward. I cannot reject that the increase in demand is the same for the bulk discount and the monthly reward, suggesting that pre-experimental demand is not closely correlated with present focus. I additionally 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, whereas naive, present focused consumers should underestimate borrowing. Both pieces of evidence indicate that present focus is likely not highly correlated with pre-experimental demand, and therefore not the mechanism driving my empirical results.

Consumer responses to the line of credit point to transaction costs and liquidity constraints limiting consumer welfare from solar and shaping demand for solar. Methodologically, revealed preference estimates of consumer welfare from solar will not be accurate unless they account for these market frictions. I estimate a conservative lower bound on consumer surplus from electrification using random variation in the fee charged on the line of credit. My estimate suggests that the cost-benefit proposition for electrification is more attractive than previously believed, although marginal households' willingness to pay for electricity is still well below cost-covering levels.

My work speaks to three segments of the development economics 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 not perfectly storable. To date, the literature on transaction costs faced by consumers in low income countries has focused on financial services. Collins et al. (2009), Beck et al. (2007), Beck et al. (2008), Dupas et al. (2018), and Ashraf, Karlan, and Yin (2006) document the high transaction costs associated with providing financial services to the poor. Jack and Suri (2014), Aycinena, Martinez, and Yang (2010), and Suri, Jack, and Stoker (2012) further show that when mobile money reduces the transaction costs associated with sending and receiving remittances, consumers are better able to cope with negative shocks.

My work advances our understanding of transaction costs in a different but increasingly common setting: prepaid contracts for imperfectly storable goods. Unlike financial services,

where reducing transaction costs leads to increased use, I show that reducing transaction costs for imperfectly storable goods can lead to reduced demand among consumers who are not severely liquidity constrained. If firms selling such goods have market power, it may not be profit maximizing for them to invest in reducing transaction costs. Better understanding consumer responses to transaction costs in low income countries clarifies which market frictions private investment will alleviate over time and which will require public investment.

In my experiment, consumers respond to small amounts of easily accessible credit designed to overcome short-term liquidity constraints, contributing 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 role 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 some of the first empirical evidence, 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) and Jack and Smith (2020) in studying contracts for electricity with poor 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 prepaid contracts are efficacious from the firm's perspective relative to more flexible arrangements. However, firm profits from prepaid contracts in my setting are partially a function of market distortions, 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 particularly 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.<sup>3</sup>

Beyond the policy relevance of PAYGo solar, I measure demand for electricity under experimental conditions that deliberately reduce key frictions consumers encounter when paying for electricity, adding a new estimate of consumer welfare from electrification to a rich literature with varied findings (Khandker, Barnes, and Samad (2009), Bensch, Kluge, 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 relevant friction in my setting. 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 WTP for the first units of electricity, a critical margin for electrification policy. My conservative lower bound on consumer surplus is 44% greater than the most similar estimate in the literature, suggesting that the cost-benefit proposition for electrification is more appealing than previously believed. Although my estimates are a lower bound, my results indicate that currently non-electrified households likely 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 background and context for PAYGo solar in Rwanda. Section 3 details the experimental design and describes my sample of consumers. 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.

---

<sup>3</sup>See [https://www.gogla.org/sites/default/files/resource\\_docs/global\\_off-grid\\_solar\\_market\\_report\\_h2\\_2018\\_opt.pdf](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.

## 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 home<sup>4</sup>.

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 solar system access time using mobile money. As soon as a consumer has purchased system access time, they have unlimited access to their SHS for the duration of the purchased time.<sup>5</sup> 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 non-storability of access time as well as 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. Continuous rundown renders access time a strictly non-storable good.

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

---

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

<sup>5</sup>In 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.

In theory, consumers can reduce the transaction costs associated with paying for solar by depositing money in their mobile money wallet when they happen to be near an agent and later using those funds to buy solar. In a phone survey conducted among 1,732 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. 66% visited a mobile money agent all five times. As figure 2 shows, 78% of consumers visited a mobile money agent to pay for solar for at least four of their last five purchases. Figure 3 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. These patterns are likely the result of low mobile money penetration. A 2018 report by the World Bank found that only 31.1% of adults in Rwanda had mobile money accounts (WBG (2018)). When consumers cannot use mobile money for most transactions, it becomes less liquid than cash due to withdrawal fees. 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. Figure 4 shows the distribution of days purchased in a single transaction in the 90 days prior to the experiment. The median transaction is 6.25 days. The prevalence of small transactions suggests that many consumers are either liquidity constrained or prefer to buy small quantities to limit wasted access time.

Taking all features of the setting together, the non-storability 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 simultaneously coping with 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: an in-kind line of credit for PAYGo system access time. The in-kind 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 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. It reduces transaction costs 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.<sup>6</sup> Half of consumers pay a 10% flat fee on borrowed days and half pay a 2% fee.<sup>7</sup> 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 regions 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 those regions 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 roughly 1,200 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, I construct a wealth index to compare the population in my sample to the distribution of rural households in Rwanda.<sup>8</sup> Figure 5 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. Households 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.

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

---

<sup>6</sup>Consumers in both groups have the option to borrow less than the maximum amount available.

<sup>7</sup>10% 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.

<sup>8</sup>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.

distribution, I create four stratification bins: 0-30%, 30%-65%, 65%-80%, and 80%-100%. Figure 6 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 a rough proxy for the severity of liquidity constraints consumers face. Table 1 shows that consumers with the highest pre-experimental demand make significantly larger purchases than consumers with the lowest pre-experimental demand. Table 2 shows differences in self-reported borrowing to pay for solar when the line of credit is not available.<sup>9</sup> 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 3 shows differences between stratification bins along various dimensions of economic well-being. Consumers with low pre-experimental demand have lower wealth scores than those with high pre-experimental demand, but the pattern is not monotonic across the distribution of pre-experimental demand. Columns two and three in table 3 suggest an explanation: consumers with the lowest pre-experimental demand are also more likely to have grid connections, indicating that low demand may be a product of substitution for a small subset of consumers. Taken together, the evidence from the phone survey indicates that pre-experimental demand loosely corresponds to the severity of liquidity constraints consumers face, but is an imperfect proxy.

### 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. 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.<sup>10</sup> Consumers could request to 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

---

<sup>9</sup>Self-reports come from the 2019 phone survey of 1,200 solar customers.

<sup>10</sup>64% of consumers in the 0%-30% stratification bin were reached on the phone during at least one of the rounds of marketing calls. 86%, 90%, and 96% of consumers were reached in the 30%-65%, 65%-80%, and 80%-100% bins.

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 on payments for all consumers in my sample to estimate changes in the monthly utilization rate.

I augment my primary data sources 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.<sup>11</sup>

I use administrative data on the pre-experimental period to check for balance. Table 4 shows that randomization yields balanced treatment and control groups on a range of observable characteristics after controlling for each consumer’s stratification bin. Having verified that the randomization worked, I proceed to build a model that generates predictions about the expected effects of offering the line of credit.

## 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 as  $q \in \mathbb{Z}^+$ . I denote the stock of days of solar access as  $e$ . The consumer’s stock of electricity access  $e$  evolves according to a simple law of motion

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

Consumers get stochastic utility  $\alpha$  if they have at least one day of electricity access ( $e + q \geq 1$ ) in a given day. I assume  $\alpha$  is drawn from some distribution with cdf  $F(\cdot)$  with support over  $[\alpha_m, \alpha_M]$ , where  $\alpha_m \geq 0$  and  $\alpha_M$  is some finite constant. At the start of each time period, consumers learn their draw of  $\alpha$  and receive a stochastic endowment of income

---

<sup>11</sup>Note 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 survey only included treated consumers.

$y \in (0, y_M]$ , 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  $\tau$  each time they buy solar. I make the simplifying assumption that  $\tau$  is constant across time and consumers and represents the costs of reaching the nearest mobile money agent. Any resources 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 stock of resources 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 satisfy

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

#### 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 value function and Bellman equation are

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

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  $\alpha$ . Given that I hold  $p$  and  $\tau$  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  $\alpha'$  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  $\mu_q$  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 the first order condition

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

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

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

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

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

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

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

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  $\alpha'$  and  $y'$ . Expectations over  $\alpha'$  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  $\alpha$  where equation (11) holds with equality  $\alpha^*$ . 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  $\alpha'$ .

A consumer choosing between  $q = 1, 2, \dots$  will make 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 if

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

Intuitively, since the consumer gains  $\alpha$  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 consumers 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  $\alpha$  is drawn from a uniform distribution on  $[\alpha_m, \alpha_M]$ . I make two simplifying assumptions. First, I assume that  $y$  follows a two-state Markov chain  $y[z]$  with state  $z \in \{0, 1\}$  and transition matrix  $P$ . Intuitively, this means that the consumer receives either a high or a low draw of income on each day. Second, I limit choices of  $q$  to  $\{0, 1, 2\}$ , which allows me to reduce the state space to  $s, z$ , and  $\alpha$  because the consumer only chooses  $q > 0$  when  $e = 0$ . With these assumptions in hand, I can generate predictions about the expected outcomes from my experiment.

### 4.3 Prediction 1: The line of credit increases demand among liquidity constrained consumers.

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  $\frac{d\alpha^*}{dl}$ .

Figures 7 and 8 illustrate how lowering transaction costs and relaxing liquidity constraints change  $\alpha^*$ . In both figures, the point where  $V_0$  and  $V_1$  intersect represents  $\alpha^*$ , the lowest draw of  $\alpha$  for which a consumer will opt to buy electricity rather than forgoing it. Figure 7 shows that lowering transaction costs reduces  $\alpha^*$ , increasing the probability that the consumer chooses to buy one day of solar. Similarly, figure 8 shows that increasing the borrowing limit  $l$  reduces  $\alpha^*$ . For consumers who lack the liquidity to buy in bulk, the line of credit unambiguously increases demand for solar. As expected, the change in  $\alpha^*$  is larger when realizations of  $s$  and  $y$  are lower.

Figures 7 and 8 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 in both figures. When transaction costs are lower, consumers can save less today if they want to buy solar tomorrow. Similarly, when consumers have guaranteed access to credit they don't need to save as much today to ensure that they can buy solar tomorrow. 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.

#### 4.4 Prediction 2: 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  $\alpha$ . I instead consider the range of  $s$  over which the consumer prefers  $q = 1$  to  $q = 2$  for a given realization of  $y$ .

In figures 9 and 10, the green circles indicate the level of  $s$  above which  $V_2$  exceeds  $V_1$  for low and high values of  $\tau$  and  $l$ . Figure 10 shows that lowering transaction costs reduces the range of assets where the consumer prefers  $q = 2$ . Figure 9 shows that relaxing liquidity constraints can either widen or narrow the range of assets where the consumer prefers  $q = 2$ , depending on the relative size of  $p$  and  $\tau$ .

For consumers with sufficient liquidity to buy in bulk prior to the experiment, offering the line of credit may operate solely 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  $\alpha$ .

#### 4.5 Prediction 3: Negative treatment effects will be larger for consumers with a higher variance in $\alpha$

Figure 10 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  $\alpha$ . Intuitively, these are the consumers who will benefit the most from better targeting their consumption. Lowering transaction costs still leads to a potential reduction in demand for consumers with a low variance in  $\alpha$ , but the change is much smaller. The final prediction of the model is that treatment effects from the line of credit should be decreasing in the variance in  $\alpha$ .

The final prediction offers a way to clarify the ambiguous predictions about the change in the probability that consumers buy in bulk. Even though the line of credit can either increase or decrease the likelihood of buying in bulk, 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.

The model generates three predictions I can empirically test about offering the line of credit. First, it will increase demand among consumers who lack sufficient liquidity to buy in bulk. The increase in demand among consumers not buying in bulk should occur regardless of borrowing status: providing guaranteed access to credit can change consumer behavior even if consumers do not use it during the period of the experiment. Second, the line of



credit may lower demand among consumers who previously bought in bulk. Third, 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  $\alpha$ . 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.

## 5 Results

I measure consumer responses to the line of credit by estimating heterogeneous average treatment effects on monthly utilization rates. Let  $i$  index consumers,  $j$  index stratification bins, and  $t$  index months of the experiment. Using a 90-day pre-period to increase precision, I estimate

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

Figure 11 shows that average treatment effects from offering the line of credit follow the first two 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%. It appears that the line of credit increases demand for consumers who are most likely to be liquidity constrained while allowing consumers who are less liquidity constrained to better target their consumption.

A key feature of the model is that credit availability reduces the precautionary savings motive for consumers who periodically face a binding liquidity constraint. Even in periods where the liquidity constraint does not bind, consumers do not need to save as much to smooth future 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 need to 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

borrowers as

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

Figure 12 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 11 is driven entirely by borrowers, it would imply an impossibly large LATE of 244pp.<sup>12</sup> 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 liquidity constrained consumers engaging in precautionary saving.

The final prediction in my model is that treatment effects are decreasing in the variance in  $\alpha$ , 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  $\alpha$  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 non-liquidity constrained consumers, the effect of offering the line of credit should be primarily driven by liquidity constrained consumers who unambiguously increase their demand. 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 13 shows that consumers with the lowest variance in use are the only group with

---

<sup>12</sup>Implied 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.

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 5 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 theoretical predictions of my theoretical framework and rule out a number of alternative mechanisms. However, one alternative explanation for my estimated treatment effects could be that pre-experimental demand is correlated with present focus. Under pure prepayment, present focused consumers pay prior to enjoying solar access. With the line of credit, present focused consumers prefer to borrow access time but then procrastinate on repayment, potentially leading to a reduction in demand. I provide additional evidence to evaluate the importance of present focus in my setting.

## **5.1 Evidence on the Importance of Present Focus**

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 three pieces of evi-

dence. 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. Finally, I provide survey evidence from a subset of consumers in the experiment that elicited their expectations about using the line of credit. All three 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, suggesting that a positive correlation between present focus and pre-experimental demand are not driving my results. However, heterogeneous responses based on variance are a poor validation of the model if variance is strongly correlated with pre-experimental demand. Figure 14 shows that pre-experimental demand is not closely correlated with the standard deviation in wH used on days when a consumer has solar access. I interpret heterogeneous responses based on variance in use as additional evidence for transaction costs as a critical driver of consumer behavior.

Further evidence on the importance of present focus comes from a different 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 6 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 reward and the monthly reward, leading to equivalent reductions in average price between the two incentives. The bulk incentive should perform worse than the monthly reward if consumers are present focused: the bulk incentive requires that consumers forgo more consumption today to gain consumption farther into the future than the monthly reward. I stratify the sample by pre-experimental demand over the entire duration of their 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 15 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 one final piece of evidence on present focus among consumers in my sample. After an initial round of marketing, representatives from the solar company called consumers again to remind them about the line of credit. During this 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, they will underestimate their use of the line of credit.

Across all consumers surveyed, only 1% underestimate their use of the line of credit. Table 7 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 7 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 and results of the RCT on incentives for solar payments, they help rule out present focus as the primary underlying mechanism driving consumer responses to the line of credit. Beyond ruling out present focus, consumer predictions 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 non-storable 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 market frictions common in low-income countries. Although the complete non-storability of PAYGo solar access time is an extreme case, even goods that store poorly like perishable food or prepaid services whenever consumers are inattentive or cannot easily track consumption will force consumers to make similar trade-offs. My results point to a range of concerns for policymakers seeking to protect consumers and create more 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 the treatment effects on the utilization rate to be representative of the distribution of consumers in figure 6, I find that offering the line of credit does not significantly increase revenue collection for the firm. Figure 16 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. The results point to a new source of profitability for firms using prepaid goods and services to contract with poor households: overconsumption among a subset of households. The experiment shows that transaction costs cause consumers with sufficient liquidity to buy in bulk, purchasing more than they would in the absence of transaction costs.

One argument in favor of the status quo is that overconsumption among a subset of consumers allows for some degree of redistribution within the firm. For instance, high-consumption customers may allow the solar firm to enact less stringent repossession rules, benefiting consumers with fewer resources. To the extent that overconsumption is correlated with income or wealth, it may function as an informal progressive tax. The problem is that consumers who cannot buy in bulk still incur transaction costs each time they pay for solar. Even if they face less risk of repossession, they enjoy fewer benefits from their solar home systems than they would in the absence of market frictions.

An alternative argument is that competition will eventually reduce or eliminate the distortions observed in my experiment. Firms will compete on the convenience of their service or the transaction costs associated with using their service. While there is a degree of competition in the PAYGo solar industry, on-grid utilities are natural monopolies which may benefit from market frictions in a manner similar to the solar firm in my experiment. Even among PAYGo solar firms, contracts with consumers tend to be long-term, potentially muting incentives to reduce transaction costs. On the whole, it is unlikely that competition will work to swiftly eliminate transaction costs in prepaid contracts.

Given that market frictions distort consumer demand for prepaid services and that firms

profit from those distortions, policymakers can act to protect consumers. In the short term, it will be important for policymakers to evaluate how well consumers can effectively track and target their consumption of prepaid goods and services, even when such products are in theory storable. Prepayment is an important tool to facilitate contracting with poor consumers, but ensuring that contracts work for both firms and consumers will require careful study and adaptation. Longer term, governments may need to lead the way on investments in infrastructure to reduce transaction costs if firms lack strong enough incentives to pursue private investments.

Methodologically, the distortions created by transaction costs and liquidity constraints render revealed preference measures of welfare inaccurate. In the next section, I re-estimate consumer welfare from electrification using the less distorted demand observed in my experiment.

## 6 Welfare

Demand observed in the presence of market frictions does not provide an accurate measure of consumers' willingness to pay. The results from the experiment show that consumers significantly alter their demand for solar when I diminish key market frictions. 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 prepay 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 days 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  $\frac{dQ(F)}{dP(F)} = \frac{dQ(F)/dF}{dP(F)/dF}$ . 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  $\frac{dQ_p(F)}{dF}$  and  $\frac{dQ_b(F)}{dF}$  using simple regressions of the form

$$Q_i = \alpha + \beta F e e_i + \delta X_i + \epsilon_i,$$

where  $X_i$  controls for the consumer's daily rate and pre-experimental demand to increase precision. Adding estimates for  $\frac{dQ_p}{dF}$  and  $\frac{dQ_b}{dF}$  yields  $\frac{dQ(F)}{dF}$ . I take means of  $Q_p(F)$ ,  $Q_b(F)$ , and  $F$  in combination with the estimated coefficients to estimate  $\frac{dP(F)}{dF}$  and, in turn,  $\frac{dQ(F)}{dP(F)}$ . I bootstrap all standard errors and confidence intervals.

Figure 17 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 18.

Table 8 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 \$108 per household per year. As I move along the demand curve to increasingly marginal consumers, the lower bound on surplus drops under \$14 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 three years, which includes the solar home system, basic appliances, maintenance, 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 \$77 per year. Column (2) in table 8 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 20% beyond the range of prices with empirical support. Altogether, 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 extend the demand curve far beyond the range of prices with empirical support. 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 8 additionally facilitate comparisons to other recent measures of consumer surplus from electrification in the literature. In table 9, I take the weighted average of my estimated lower bound on consumer surplus and compare it to three recent measures in the literature: Grimm et al (2020), Lee et al (2020), and Burgess et al (2020). My lower bound on consumer surplus is greater than estimates of total consumer surplus in all three papers, with the exception of the upper range of estimates in Lee et al (2020). Were I to extend the range of prices included in my demand curve, I would 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. 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 relatively high willingness to pay for solar. 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 an imperfectly storable good that involves 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 are least likely to be liquidity constrained 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, pointing to a role for government to protect consumers.

Consumer responses to the line of credit suggest that market frictions significantly distort demand for electricity in low-income countries. It follows that revealed preference measures of welfare from electrification that cannot account for market frictions provide inaccurate estimates. Using the less distorted 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 substantially below cost-covering levels. Given that the average consumer in my

sample has a substantially higher wealth index than the average rural Rwandan household, universal electrification will likely require some form of 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 two 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.

## References

2018. “The Little Data Book on Financial Inclusion 2018.” Tech. rep., International Bank for Reconstruction and Development.
- 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.
- 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.
- Jack, Kelsey and Grant Smith. 2015. “Pay as you go: Pre-paid metering and electricity expenditures in South Africa.” *American Economic Review Papers and Proceedings* 105 (5):237–241.
- . 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.
- 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.
- 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

	<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.

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

	<i>Dependent variable:</i>		
	Borrowed for Solar	Amount Borrowed	Cannot Borrow
	(1)	(2)	(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 <sup>2</sup>	0.002	0.004	0.011
Adjusted R <sup>2</sup>	-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 2: Differences in Borrowing Behavior 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 <sup>2</sup>	0.008	0.004	0.004
Adjusted R <sup>2</sup>	0.006	0.002	0.002

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

Notes: I use the following variables, common to both surveys, 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.

Table 3: Differences in Wealth and Energy Spending between Stratification Bins

	<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 <sup>2</sup>	0.954	0.017	0.045	0.003	0.005	0.042
Adjusted R <sup>2</sup>	0.954	0.017	0.044	0.003	0.005	0.042

*Note:*

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

Table 4: Balance Table

	<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 <sup>2</sup>	0.016
Adjusted R <sup>2</sup>	0.011

*Notes:*

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

Table 5: Heterogeneity in Take-Up by Distance to Mobile Money Agent

	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 6: Cross-Randomized Experimental Design for Solar Incentives

<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 <sup>2</sup>	0.013	0.038
Adjusted R <sup>2</sup>	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 7: Differences in Consumer Expectations and Realizations of Credit Use

Pre-Experimental Utilization Rate	Estimated Slope	Q <sub>1</sub> (weighted)	Max % Δ P	CS Lower Bound (hh/year), current prices (1)	CS Lower Bound (hh/year), fully subsidized (2)
0%-30%	-0.267	15,487	2.5%	\$13.84	\$20.67
30% - 65%	-0.019	56,473	5.1%	\$48.38	\$71.97
65% - 80%	-0.027	54,693	6.2%	\$67.06	\$99.58
80% - 100%	-0.005	709,445	4.5%	\$108.28	\$161.22
Weighted Mean			4.5%	\$86.68	\$129.04

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 8: Consumer Surplus from Solar

Source	Location	Electricity Type	Price	CS (hh/year)
This paper	Rural Rwanda	PAYGo Solar	Zero	\$129.04
This paper	Rural Rwanda	PAYGo Solar	Marginal Cost	\$86.68
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.

Table 9: Consumer Surplus from Solar

## 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 4: Distribution of Payment Sizes 90 Days Prior to the Experiment





Figure 5: Distribution of Wealth Among Rural Households in Rwanda



*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 3: Self-Reported Knowledge of Using Mobile Money to Buy Solar




Figure 6: Distribution of Demand Prior to the Experiment




Figure 7: Change in  $\alpha^*$  from reducing transaction costs.



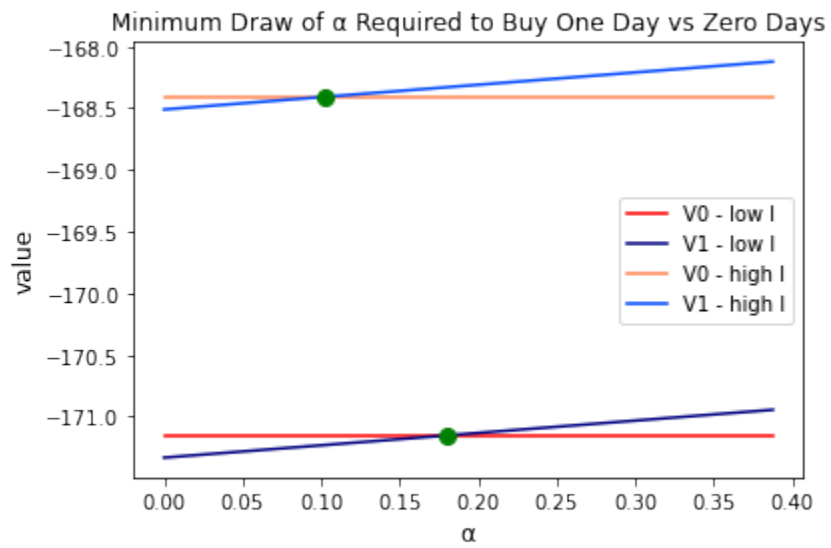


Figure 8: Change in  $\alpha^*$  from relaxing liquidity constraints.

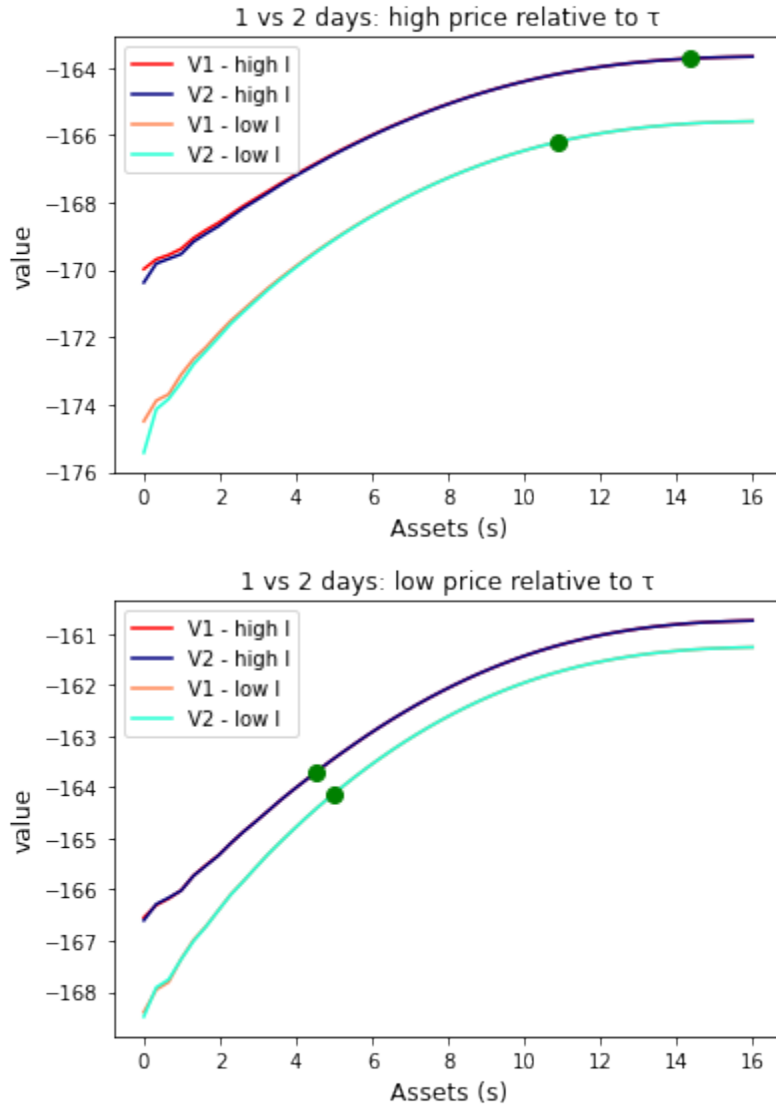


Figure 9: Change in the single vs bulk decision from relaxing liquidity constraints.






Figure 10: Size of the treatment effect for consumers with a high vs low variance in  $\alpha$



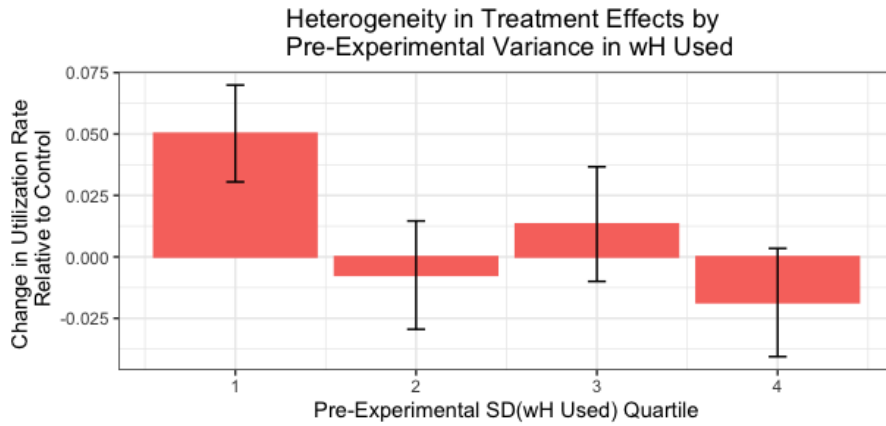
*Note:* 95% confidence intervals calculated using standard errors clustered at the level of the individual consumer. Estimates for the 30%-65% and 80%-100% bins are not significant after pre-registered multiple inference corrections.

Figure 11: 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 12: 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 first quartile remain statistically significant after applying pre-registered multiple inference corrections.

Figure 13: Heterogeneous Impacts by Pre-Experimental Variance in Use


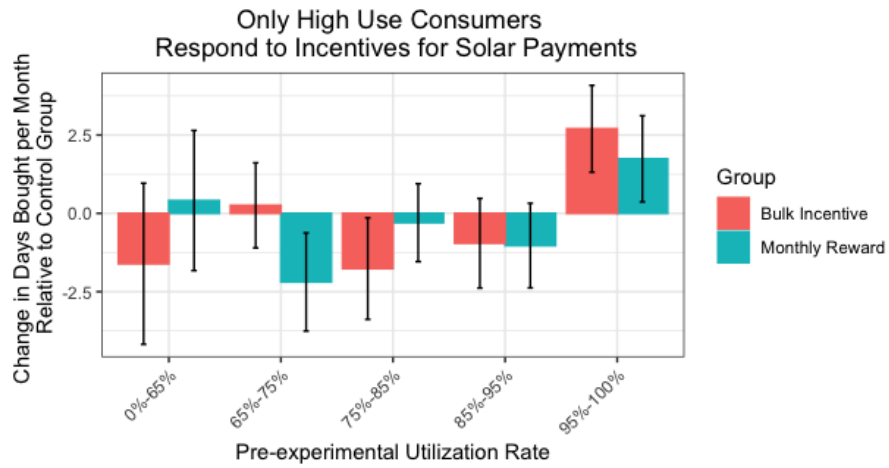


Figure 14: No Correlation Between Pre-experimental Demand and Variance in Use



*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 15: Heterogeneous Impacts of Incentives for Solar Payments



*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 conventional levels.

Figure 16: Heterogeneous Impacts on Default by Pre-Experimental Demand




Figure 17: Estimated Slopes of the demand curve for solar

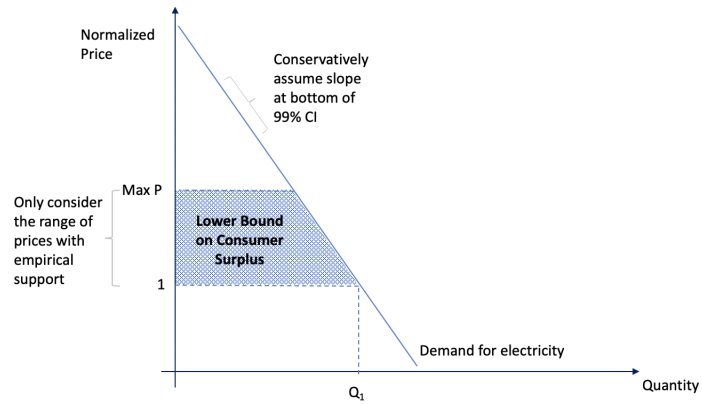


Figure 18: Estimated Slopes of the demand curve for solar