Title
Essays on Demand for Solar Electricity among Rural Consumers in Sub-Saharan Africa

Author
Lang, Megan Elizabeth

Publication Date
2021

Peer reviewed|Thesis/dissertation

Essays on Demand for Solar Electricity among Rural Consumers in Sub-Saharan Africa

by

Megan E Lang

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Agricultural and Resource Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Jeremy Magruder, Chair
Professor Ethan Ligon
Professor Catherine Wolfram

Spring 2021

Essays on Demand for Solar Electricity among Rural Consumers in Sub-Saharan Africa

Abstract

Essays on Demand for Solar Electricity among Rural Consumers in Sub-Saharan Africa

by

Megan E Lang

Doctor of Philosophy in Agricultural and Resource Economics

University of California, Berkeley

Professor Jeremy Magruder, Chair

Electricity is widely considered to be critical for economic growth, but recent literature provides mixed evidence on the welfare implications of rural electrification in developing countries. Understanding the welfare impacts of electrification requires detailed knowledge about the determinants of consumer demand for electricity in rural, low-income settings. My dissertation uses a combination of field experiments and natural experiments to contribute to current knowledge of consumer demand for electricity and explore the associated implications for consumer welfare.

I employ a unique product to address current limitations in the literature: pay as you go (PAYGo) solar. PAYGo solar is explicitly designed to lower barriers to electrification for rural, low-income households. It does so my requiring down payments which are small relative to the cost of purchasing a solar home system outright or being connected to the grid, and by allowing for flexibility in ongoing payments to use the system. In each chapter of my dissertation, I partner with a PAYGo solar company operating throughout sub-Saharan African to study a distinct aspect of consumer demand for solar electricity.

The first chapter of my dissertation studies consumer responses to randomly offered bulk discounts and monthly rewards for electricity purchases in Rwanda and Kenya. Both types of incentives potentially alter the average and marginal price that consumers pay for solar electricity. I find that most consumers do not respond to either type of incentive. While this could suggest that demand for electricity is inelastic on the intensive margin, I provide suggestive evidence that uncertainty over the future marginal utility of solar as well as liquidity constraints may also play important roles in explaining consumer non-responsiveness to price.

In the second chapter of my dissertation, I use a second field experiment with solar customers in Rwanda to formally explore the role of liquidity constraints and transaction costs in shaping demand for solar electricity. The key observation underlying this chapter is that PAYGo solar contracts render electricity access a strictly perishable good, or a good that cannot be stored. I show that transaction costs for perishable goods create welfare loss.

The loss comes from the trade-off between transaction costs and the waste that occurs when perishable goods expire, a trade-off that is compounded by liquidity constraints.

I explore the trade-off between waste and transaction costs through a field experiment where I randomly offer 2,000 current solar customers a line of credit for solar access time. The line of credit both alleviates liquidity constraints and lowers transaction costs. Consumers who previously bought in bulk respond by eliminating wasteful consumption, reducing demand by up to 6.4%. Those who are the most likely to be liquidity constrained increase demand by 88%. My results illustrate that transaction costs for perishable goods distort willingness to pay in opposite directions for different subsets of consumers. I estimate consumer surplus from electricity under the less distorted conditions of my experiment. My estimates indicate a stronger cost-benefit proposition for universal electrification than others in the literature, but indicate that marginal households' willingness to pay still falls below current cost-covering levels.

In the final chapter, I turn to the aspect of demand for electricity that I could not address with either of the field experiments in the first two chapters: the initial decision to adopt electricity. I specifically ask how much electrification choices depend on the expected price of electricity on the intensive margin. I use the phased rollout of a subsidy for PAYGo solar electricity in Togo to measure impacts on electrification. The subsidy reduces the intensive margin price of electricity by 18–42%, with the largest discounts going to consumers choosing the smallest systems. The subsidy increases the adoption of small systems by 85%–135% under different specifications, implying relatively high elasticities of 2.02–3.25. Increases in adoption of medium and large systems range from 23% to 63%, implying elasticities between 1.33 and 3.54. I compare welfare estimates from the subsidy in Togo to my work on the intensive margin of PAYGo in Rwanda in chapter two. Discrepancies between the two studies point to limitations of revealed preference measures for durable goods: both intensive and extensive margin demand may be shaped by a variety of non-price factors that fail to accurately measure welfare. I conclude by offering directions for future research on the economics of rural electrification.

To all the runners who taught me the skills I would need to succeed in research: to work hard, find great teammates, pace myself, and to have the confidence to take risks when the time is ripe.

Contents

List of Figures

List of Tables

Acknowledgments

I could not have completed my PhD without the guidance and support of my mentors, peers, family, and friends.

First, I thank my dissertation committee. I thank Jeremy Magruder for always encouraging my independence as a researcher and providing the support I needed to pursue my own ideas from the very start of my graduate career. His enthusiasm inspired me to push through challenges and his equanimity taught me how to navigate the ups and downs of research and fieldwork. Ethan Ligon's seemingly infinite patience to work through theory has made me a better economist, empowering me to tackle problems I once viewed as too difficult. Catherine Wolfram's support gave me permission to be an energy and a development economist even though I came to energy late in the game. In doing so, she allowed me to build a research agenda that I am passionate about pursuing for years to come.

Few PhD students succeed in ARE without the advice and guidance of Betty Sadoulet. I thank Betty for her tireless efforts to keep me and my cohort on track, and for not giving up on our cohort when half of us failed our first microeconomic theory exam.

None of the work in my dissertation would have been possible without my collaborators at BBOXX. I am deeply grateful to Andrew Kent, Anthony Mabonga, Philip Sendi, and Gilbert Ishimwe for their willingness to find common ground and pursue rigorous answers to important questions about rural electrification. Thanks to Sugi Aime Colombe, Moises Dushime, Innocent Mutaganda, and the teams at the BBOXX HQ, BBOXX Rwanda and BBOXX Kenya offices for their collaboration. I would also like to thank the International Growth Centre and the Center for Effective Global Action for generous research funding.

I have been incredibly fortunate to travel through my PhD with a cohort of people who have come to be some of my closest friends. To all of you: thank you for the hours of studying, listening, puzzling, struggling, laughing, and playing that we have done over the past six years. Our tight knit community has enriched my life in so many ways.

Outside my cohort, I thank Carly Trachtman, Susanna Berkouwer, Erin Kelley, and Julia Seither for being consistent sounding boards, cheerleaders, and friends. Thanks to Tamma Carleton for countless miles run discussing everything related to economics and life and for telling me to apply to Berkeley ARE for grad school.

To my sister Jackie: thank you for showing me the path to a PhD, helping me believe I could do real analysis, and still letting me ask you math questions. Finally, thank you to my parents Linda and Marc, who cultivated my love of learning and believed in me as an economist since high school IB economics. I hope that one day I can show you the places my research has led me to love.

Chapter 1

Demand for Electricity on the Intensive Margin in Rwanda and Kenya

1.1 Introduction

Although electricity is widely considered to be critical for economic growth, recent literature provides mixed evidence on the welfare impacts of electrification in low income countries. Studies including Dinkleman 2011, Lipscomb, Mobarak, and Barham 2013, and Khandker, Samad, et al. 2014 find substantial benefits from electrification, but Grimm, Lenz, et al. 2020, Lee, Miguel, and Wolfram 2020, and others find that electrification at best has no impact on welfare and, at worst, is actually welfare reducing given its costs. Low or no welfare benefits from electrification are consistent with highly elastic demand for electricity, which is what Lee, Miguel, and Wolfram 2020 find in Kenya and what Grimm, Lenz, et al. 2020 find in Rwanda. However, limited welfare benefits from electrification are also consistent with credit constrained consumers, poor reliability in the supply of electricity, and spillovers due to theft of electricity. Such factors confound straightforward estimation of the impact of electrification on welfare.

In contrast, the literature on consumer demand for electricity in developed countries typically finds that demand for electricity is price inelastic (see Deryugina, MacKay, and Reif 2019, for example). If the same elasticities were to translate to low income settings, it would suggest that welfare gains from electrification could be large. While Ito 2014 argues that insensitivity to price in the U.S. is at least partially due to consumer inattention, it is intuitively appealing to consider electricity a necessity good for consumers in high income countries.

I study consumer demand for electricity in a developing country setting using a product that enables me to avoid many confounding factors present in other studies: pay as you go (PAYGo) solar. With PAYGo solar, consumers make a small down payment on a solar home system, then "pay as they go" by purchasing system access time. The PAYGo contract ties repayment for the solar home system to use of the solar home system, leveraging remote lockout technology to enforce a contract between the consumer and the solar firm.

Solar home systems are highly reliable, so supply side reliability is not a concern in my setting. Since PAYGo contracts are designed to allow consumers to pay off the solar home system over time, credit constraints are substantially lower on the adoption margin in my setting than in other settings where households must independently finance a solar home system or a connection to the grid. Furthermore, electrical appliances are bundled with the solar home system, so consumers face lower credit constraints both in terms of adopting electricity access and financing the capital purchases necessary to benefit from electricity. Remote lockout as well as the credible threat of system repossession enable the solar company to enforce payments, largely preventing theft. Finally, consumers in my setting spend 10-15% of their household budget on electricity (Zollman et al. 2017), and they must pre-pay for electricity. Both of these factors make consumer inattention unlikely in my setting.

I leverage this setting to explore factors that shape consumer demand for electricity along the intensive margin among poor, rural consumers in Kenya and Rwanda. This paper addresses two primary questions. First, how do PAYGo solar consumers respond to incentives that alter the average and marginal price that they pay for electricity, and what is the price elasticity implied by consumer responses to these incentives? Second, to what extent are market frictions and the structure of the PAYGo contract potentially muting consumer responsiveness to the incentives? I also provide limited evidence on behavioral biases that could be driving my experimental results. Finally, I discuss implications for the provision of PAYGo solar specifically and electricity more generally, then conclude with future directions for research.

1.2 Literature Review

Much of the literature on electrification in developing countries to date has focused on the extensive margin. These papers fall into two methodological classes, which are not necessarily mutually exclusive. The first class estimates welfare changes due to electrification using impact evaluation-style designs. The second evaluates welfare changes from electrification by studying demand curves.

Impact evaluation-style designs evaluate potential welfare benefits of electrification by estimating the reduced form impact of electrification on a range of outcomes that proxy for welfare. For instance, Burlig and Preonas 2016 estimate reduced form medium-run impacts of electrification using a regression discontinuity design and find that electrification has no effect on income and other measures of economic development, suggesting that welfare impacts from electrification may be low.

Evidence from impact evaluation-style studies is somewhat mixed. Grimm, Lenz, et al. 2020 and Lenz, Anciet Munyehirwe, and Sievert 2017 both find some positive impacts in Rwanda in the form of reduced energy expenditures, health improvements, and reduced use of dry cell batteries, but Bensch, Kluve, and Peters 2011 find no positive impacts on income or energy expenditures. In Tanzania, Chaplin et al. 2017 find mixed impacts, with electrification increasing income generating activities and per capita consumption but having negative impacts on health. Conversely, Lipscomb, Mobarak, and Barham 2013 and Dinkleman 2011 both find that electrification leads to improvements in economic development in Brazil and South Africa, respectively. Walle et al. 2017 and Khandker, Samad, et al. 2014 both find positive impacts of electrification on consumption and income in India, in contrast to what Burlig and Preonas 2016 estimate. Khandker, Barnes, and Samad 2009 also document increases in income as a result of electrification in Bangladesh.

The primary limitation of impact evaluation-style measurements is that impacts of electrification may be difficult to capture using proxies like income or consumption, particularly in the short run. Qualitative evidence from Zollman et al. 2017 suggests that consumer utility increases substantially as a result of purchasing a solar home system, but that the changes are due to difficult to measure improvements in quality of life. For example, solar home systems provide outdoor lighting that enhances security. Reduced crime as a result of electrification is likely difficult to capture using proxies for welfare, particularly in the short term, but the benefits for consumers in terms of increased safety and reduced stress could be large as they accumulate over time.

Demand curve-based estimates of welfare avoid some of the problems with using proxies for welfare. By directly estimating willingness to pay for electricity, they remain agnostic about the source of welfare and instead rely on revealed preferences.

Evidence from demand curve-based estimates of welfare suggests that welfare benefits from electrification are low. Lee, Miguel, and Wolfram 2020 find that welfare benefits of electrification are actually negative in Kenya due to highly elastic demand and high connection costs. Similarly, Grimm, Lenz, et al. 2020 find that willingness to pay for solar in Rwanda is substantially below cost-covering levels. In Tanzania, Chaplin et al. 2017 find that consumers respond to reductions in the cost of connecting, though not enough to increase adoption to desired levels.

Examining the demand curve on the adoption margin may not accurately measure welfare benefits of electrification for two reasons. First, if consumers have incomplete or inaccurate information about the benefits of adopting electricity, then their adoption decision will not reflect their true willingness to pay for electricity given complete information. In settings where electrification rates are low, it is plausible that consumers do not have complete information when making their adoption decision. Second, as Lee, Miguel, and Wolfram 2020 point out, consumers may be credit constrained. The demand curve for credit constrained consumers may appear more elastic than the demand curve when credit is available.

In addition to these concerns, both methodologies for estimating welfare impacts are often confounded by poor supply-side reliability. If electricity is not reliable, benefits of electrification will be low and willingness to pay for electricity will reflect the uncertainty in supply.

I estimate welfare from electricity using price changes on the intensive margin, thereby contributing to the literature by focusing on a post-adoption margin. The drawback of studying the intensive margin is that my results are estimated on a self-selected sample of consumers who have already adopted solar, limiting generalizability. However, the structure of the contracts with the solar company I partner with are such that the solar company implicitly extends credit to consumers, thereby at least partially alleviating credit constraints on the adoption margin. By focusing on the intensive margin, I can rule out inaccurate and incomplete information about the benefits of electricity, which may shape demand on the extensive margin. Finally, reliability of solar home systems in my setting is high, so I can rule out poor reliability as a potential confound.

There has been relatively little work on the intensive margin of electricity consumption in low income countries. Notably, K. Jack and Smith 2015 show that consumers in lower quartiles of the property value distribution in South Africa purchase more electricity when tariffs reset to low rates, suggesting that consumers may be sensitive to price on the intensive margin. Furthermore, K. Jack and Smith 2020 find that pre-paying for electricity reduces electricity consumption by 14%. They attribute the reduction in part to marginal price sensitivity, providing additional evidence that the intensive margin of electricity is price elastic for consumers in low income countries.

1.3 Background: Pay as You Go Solar

Consumers who choose to adopt a solar home system from the solar company I partner with purchase both the system and all of the appliances to be used with the system, with the exception of mobile phones.[1] Consumers make a small down payment to have the system installed, and the total value of the system and their appliances plus a fee is then converted into a daily price for solar access, to be paid off over a period of roughly three years.

The smallest solar home systems include solar panels, a battery to store power generated during daylight hours, three light bulbs, and the consumer's choice of an additional light bulb, a rechargeable portable light, or a rechargeable radio. Consumers can choose to have additional appliances for a higher price, including a television or a smart phone. Therefore, the solar home systems I study are primarily used for the purpose of household consumption rather than income generation.

Once a household has a solar home system installed, the consumer prepays for electricity access. The consumer adds money to their solar account using mobile money. The deposit purchases a certain number of days of solar access, depending on the daily price the consumer opted into based on the number of appliances they selected.

Importantly, system access time runs down continuously, regardless of system use. Therefore, a consumer who purchases three days of solar access on Wednesday will be able to use their system on Wednesday, Thursday, and Friday. They cannot opt to turn their system off on Friday and save the third day for later use. Furthermore, consumers cannot transfer

[1] Mobile phones can be charged using the solar home system, but the phone itself is not included in the package the consumer purchases from the solar company.

money out of their solar account (no refunds), and they cannot transfer solar credit to other solar consumers (e.g., there is no secondary market for solar).

If a consumer runs out of solar credit, their system is remotely switched off until they make another payment. After prolonged periods of non-payment, the solar company repossesses the solar home system, including all appliances, to be refurbished and redeployed to a new customer.

On average in both Rwanda and Kenya, consumers use fifty watt hours (wH) on days when they have purchased access to their solar home systems.[2] In Rwanda, the median purchase size among the sample of consumers in my experiment is seven days of solar while in Kenya the median purchase size is five days of solar.

1.4 Experimental Design

The solar company offered two types of incentives to randomly selected customers for 7–8 months in 2018–2019. The first type of incentive was a bulk discount: buy x weeks at once to receive y days free. The second type of incentive was a monthly bonus: buy x weeks over the course of a month to receive y days free.

The primary goal of both types of incentives was to alter the average and marginal prices consumers faced for solar. In addition, the bulk discount required consumers to forgo more liquidity than the monthly reward, as consumers could make multiple small purchases to work up to the monthly reward but had to make a large purchase to earn the bulk discount. Therefore, comparing responses to the two types of incentives provides a way to understand the importance of liquidity in driving consumer behavior. I also induced second-order variation in the price of solar and the liquidity required to qualify for incentives by cross-randomizing x and y within each incentive type. x was either four or five weeks. Within the four week group, consumers either received one or two bonus days when they qualified, and within the five week group consumers received either 3 or 4 days. Consumers were offered a schedule of incentives, with incentives increasing in the number of weeks purchased. Figure 1.1 shows the full schedule of incentives for each of the cross-randomized groups. In practice, few consumers qualify for anything above the minimum number of bonus days in a given incentive schedule.

One concern about the treatments is that the structure of the incentives may simply provide consumers with information about how to remain in good standing with the solar company. For consumers experiencing uncertainty regarding the risk of repossession, the incentives may change behavior not because they lower the effective price of solar but because they provide additional information about how to be a "good customer". To address such concerns, the solar company also agreed to implement a simple information treatment. In the information treatment, consumers received a phone call from the solar company informing

[2]Appliances used with the solar home system are designed to be highly efficient, so consumer benefit from fifty wH more than would be expected were they using equivalent appliances that had not been optimized for use with the solar home system.

	Low Bonus			High Bonus		
	Weeks to buy	Bonus days	% discount	Weeks to buy	Bonus days	% discount
Low Threshold	4	1	3.57%	4	2	7.14%
	5	3	8.57%	5	4	11.43%
	6	5	11.90%	6	6	14.29%
	8	7	12.50%	8	10	17.86%
	12	12	14.29%	12	20	23.81%
High Threshold	Weeks to buy	Bonus days	% discount	Weeks to buy	Bonus days	% discount
	5	3	8.57%	5	4	11.43%
	6	5	11.90%	6	6	14.29%
	8	7	12.50%	8	10	17.86%
	12	12	14.29%	12	20	23.81%

Figure 1.1: Full Cross-Randomized Schedule of Incentives

them how many days a month they needed to purchase to be considered a “good customer”. I compare responses to the information treatment to responses to both types of incentives to determine whether it appears that information is driving any estimated treatment effects.

In total, 800 consumers were offered the bulk discount and 800 consumers were offered the monthly reward in both Rwanda and Kenya. Thus, each cross-randomization cell contained 200 consumers. 400 consumers received the information treatment. Currently, I use a control group of 5200 consumers in Rwanda and 3800 consumers in Kenya, though these can be expanded further to increase power. I stratify the sample based on a consumers’ utilization rate at the start of the experiment. The utilization rate is the number of days a consumer has had access to their solar home system out of the total number of days they could have had access to their solar home system. I stratify based on five bins: 65% and below, 65%-75%, 75%-85%, 85%-95%, and above 95%.

A representative from the solar company’s call center called each consumer who had been selected for one of the incentives or the information treatment in the two weeks prior to the start of the experiment. Once the experiment started, consumers in the incentive treatment groups received weekly SMS messages with detailed reminders about the full schedule of incentives.

The experiment ran from July, 2018 through February, 2019 in Rwanda and from August, 2018 through February, 2019 in Kenya. In both locations, the experimental period covered the end of a harvest season, a full lean season, and a planting season.

1.5 Results

I present all results separately for Kenya and Rwanda due to differences in mobile money penetration, which is substantially higher in Kenya than it is in Rwanda.[3] High mobile money penetration renders mobile money substantially more liquid, as it can be used in nearly the same manner as cash. By contrast, when mobile money penetration is low, mobile money can only be used for a limited range of transactions. Given that there are fees to withdraw money from a mobile money account, low penetration has direct implications for how much money consumers choose to hold in their mobile money wallet. This in turn shapes the transaction costs consumers face when they pay for solar, thereby potentially altering consumer responses to the incentives being offered in the experiment.

My basic specification to estimate reduced form treatment effects takes the following form:

$$MonthlyPurchases_{it} = \alpha + \beta_1 INFO_{it} + \beta_2 BULK_{it} + \beta_3 MR_{it} + \gamma_i + \gamma_t + \epsilon_{it}. \tag{1.1}$$

$MonthlyPurchases_{it}$ is the number of days purchased by consumer i in month t. $INFO_{it}$, $BULK_{it}$, and MR_{it} are dummy variables for whether consumer i was in the information treatment, the bulk discount, or the monthly reward treatment groups, respectively, in month t. γ_i is a consumer fixed effect and γ_t is a month fixed effect. Throughout, I use a four month pre-period for Rwanda and a five month pre-period for Kenya. I cluster all standard errors at the unit of randomization, the individual consumer.

Table 1.1 shows treatment effects pooled across all cross-randomized sub-treatments as well as aggregated across all layers of stratification. In Rwanda, I can say with 95% confidence that bulk discounts increased monthly purchases by more than 4.12% and that monthly rewards increased monthly purchases by more than 4.56%. In Kenya, I can say with 95% confidence that bulk discounts increased monthly purchases by more than 1.30% and that monthly rewards increased purchases by more than 2.34%. I can also reject that either incentive had an effect that was statistically different from the information treatment in both countries at the 5% significance level.

One concern with pooling all results is that consumers in different parts of the utilization distribution may respond differently to the incentives. For instance, a consumer who typically purchases only fifteen days in a month will have to increase monthly purchases by substantially more than a consumer who usually purchases twenty-five days a month to qualify for an incentive. My sample stratification allows me to test for this type of heterogeneity.

Figures 1.2 and 1.3 show heterogeneous treatment effects in Rwanda and Kenya, respectively. The black dots in each figure represent the mean monthly purchases of the control group for a given utilization bin over the course of the experiment.

[3] See *Rwanda - Data at a Glance* n.d. and *Kenya - Data at a Glance* n.d. Financial inclusion insights estimates that only 23% of adults in Rwanda have a mobile money account, while 72% of adults in Kenya have a mobile money account. Furthermore, they find that mobile money users in Kenya are highly likely to use an advanced function, whereas the most common use of mobile money in Rwanda is paying salaries.

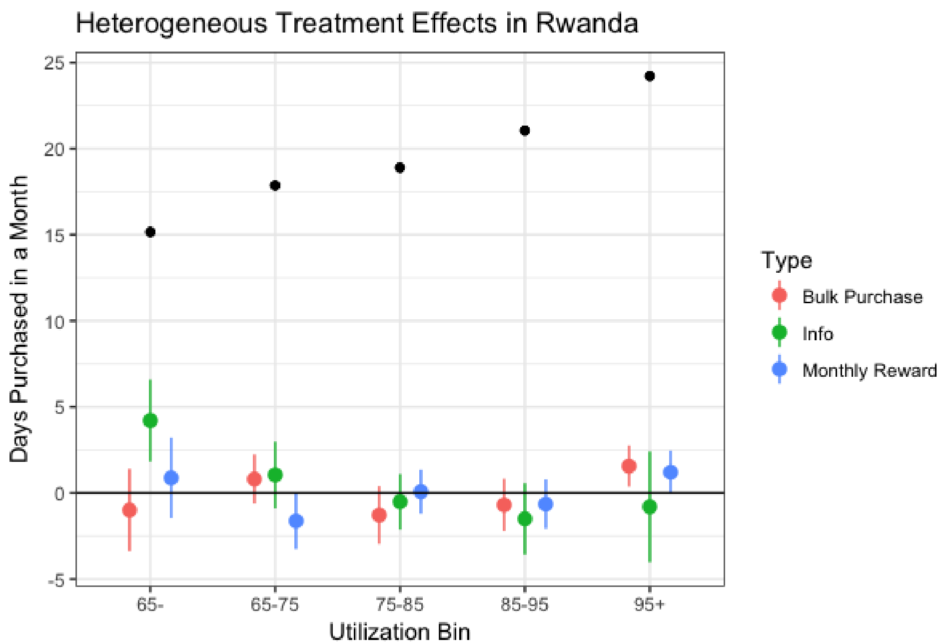

Figure 1.2: Heterogeneous Treatment Effects in Rwanda

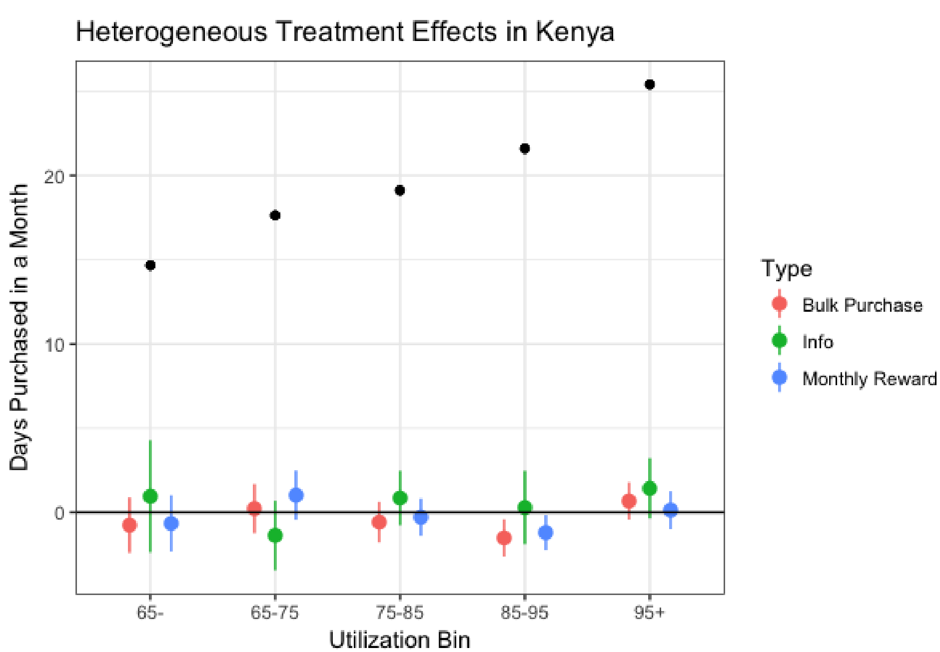

Figure 1.3: Heterogeneous Treatment Effects in Kenya

Table 1.1: Pooled Treatment Effects

	Dependent variable:	
	Treatment Effect (Days Bought in a Month)	
	Rwanda	Kenya
	(1)	(2)
Information	0.350	0.502
	(0.560)	(0.527)
Bulk Purchase	-0.095	-0.398
	(0.404)	(0.329)
Monthly Reward	-0.003	-0.196
	(0.395)	(0.326)
Consumer Fixed Effects	Yes	Yes
Month Fixed Effects	Yes	Yes
Control Mean	16.9	18.9
Observations	81,489	54,848
Adjusted R^2	0.134	0.213
Note:	*p<0.1; **p<0.05; ***p<0.01	
	Standard errors clustered at the level of the consumer.	

Figure 1.2 shows that consumers in the top utilization bin are responding to both types of incentives in Rwanda, although I cannot reject that the incentives have a different effect than the information treatment due to a noisy estimate on the information treatment. Both treatment effects represent approximately a seven percent increase over control group purchases. Throughout the rest of the utilization distribution, incentives did not have any significant effects. Notably, the information treatment increased purchases of consumers in the lowest utilization bin by over twenty-five percent - a massive effect for a one-time treatment. Unfortunately, I'm unable to shed further light on the effect of the information treatment among low-utilization consumers given limited statistical power.

Figure 1.3 shows analogous effects for Kenya. In Kenya, neither incentive nor the information treatment had any significant, positive effect on monthly purchases. In fact, both types of incentives significantly reduced purchases by consumers in the 85%-95% utilization bin by around six to seven percent, although again I cannot rule out that the effects are the same as the information treatment.

Taken together, my reduced form treatment effects show that most consumers in both Rwanda and Kenya did not respond to the price incentives offered during the experiment. In the remainder of the paper, I consider a range of possible explanations for these null effects. First, I provide two different estimates of the price elasticities implied by my experimental results. Then, I consider non-price factors that are also consistent with the estimated treatment effects.

1.6 Implications for Elasticity

One explanation for my estimated treatment effects is that the consumers in my sample are not sensitive to the price of electricity on the intensive margin. This would be consistent with estimates of the price elasticity of electricity in the U.S., but would differ from what K. Jack and Smith 2020 find in South Africa. However, since the experiment does not induce straightforward variation in prices, it is not obvious what elasticities are actually implied by my reduced form results. In this section, I build a simple model that reflects that structure of the incentives and the PAYGo contract, then use the exogenous variation induced by the experiment to estimate the parameters of the model. I supplement this with more informal estimates of the price elasticity of solar.

A Model of Demand for Solar

The treatments induce exogenous variation in the number of bonus days offered to consumers at various purchasing thresholds. I am interested in estimating the price elasticity of solar. Therefore, I need to develop a system of equations to link exogenous variation in bonus days offered to the price elasticity of solar.

To fix ideas, let D be the number of days of solar that a consumer actually purchases in a month, while B is the number of bonus days that the consumer earns. The potential size of B, should the consumer qualify for bonus days, is determined exogenously by the experimental design. Since I expect the size of the bonus to affect consumer behavior, it follows that D is a function of B which I will denote $D(B)$.

The choice of whether or not to qualify for bonus days is endogenous, creating a typical selection into treatment scenario. I can address this using two stage least squares. My first stage is

$$B_{it} = \alpha + \delta Treatment_{it} + \gamma_i + \gamma_t + \mu_{it}. \tag{1.2}$$

That is, I first estimate the number of bonus days actually earned by consumer i in month t as a function of the treatment group that the consumer was exogenously assigned to in that month. I use $\hat{\delta}$ to predict $\hat{B}$, then I estimate

$$D_{it} = \rho + \beta\hat{B} + \gamma_i + \gamma_t + \epsilon_{it}. \tag{1.3}$$

$\hat{\beta}$ is the change in the number of days purchased resulting from increasing the number of bonus days earned by one. Put otherwise,

$$\beta = \frac{dD(B)}{dB}.$$

To estimate the price elasticity of solar, I need to transform $\frac{dD(B)}{dB}$ into $\frac{dQ(B)}{dP(B)}$. To see why Q and P are both functions of B, consider that

$$Q(B) = D(B) + B \tag{1.4}$$

is the total days of solar that a consumer receives in a month. P is the average price that the consumer pays for Q days of solar in a month. I normalize the daily rate a consumer pays to one, then write the effective price as

$$P(B) = \frac{D(B)}{D(B) + B}. \tag{1.5}$$

Intuitively, a consumer pays for $D(B)$ days and receives $D(B)$ days plus any bonus days that the purchase earns. Therefore, the effective price is the price paid by the consumer divided by the total quantity received by the consumer.

From here, it is straightforward to see that

$$\frac{dQ(B)}{dP(B)} = \frac{\frac{dQ(B)}{dB}}{\frac{dP(B)}{dB}}.$$

I take the derivative of equations (1.4) and (1.5) to get

$$\frac{dQ(B)}{dB} = \frac{dD(B)}{dB} + 1 \tag{1.6}$$

and

$$\frac{dP(B)}{dB} = \frac{\frac{dD(B)}{dB}B - D(B)}{(D(B) + B)^2}. \tag{1.7}$$

I have an estimate of $\frac{dD(B)}{dB}$ from equations (1.2) and (1.3). Therefore, I simply plug in my estimate to get a point estimate of the price elasticity of solar at various values of B, and corresponding values of $D(B)$.

Elasticity Results

Table 1.2 shows the 2SLS and OLS estimates of $\frac{dD(B)}{dB}$ for Rwanda and Kenya. The first point of note is that the first stage F-statistic is borderline in Rwanda, and weak in Kenya. Therefore, I will proceed to transform these point estimates into elasticities for Rwanda, but

Table 1.2: dD/dB using 2SLS and OLS in Rwanda and Kenya

	Dependent variable:			
	2SLS (RW)	OLS (RW)	2SLS (KE)	OLS (KE)
	(1)	(2)	(3)	(4)
$\frac{dD(B)}{dB}$	4.220**	7.067***	7.442***	6.771***
	(1.657)	(0.525)	(1.889)	(0.916)
Consumer Fixed Effects	Y	Y	Y	Y
Month Fixed Effects	Y	Y	Y	Y
First Stage F-statistic	9.1646	NA	5.9949	NA
Control Mean	17.13	17.13	18.91	18.91
Observations	82,936	82,936	54,514	54,514

Note: *p<0.1; **p<0.05; ***p<0.01
Standard errors clustered at the level of the consumer.

not for Kenya since so few customers earned bonus days there that I do not have a credible first stage.

In Rwanda, I get an estimate of 4.220 for $\frac{dD(B)}{dB}$, with a 95% confidence interval of [0.97,7.47]. When I plug this into my model of consumer demand, it translates into an elasticity of -5.22 to -10.36 for values of B ranging from zero to four.[4] These are highly elastic values for the price elasticity of solar, which initially appears to contradict the null reduced form treatment effects. However, using 2SLS necessarily estimates the second stage off those consumers who actually respond by earning bonus days in the first stage. Therefore, I interpret my results as an estimate of the elasticity for those consumers who responded to the incentive, not as the average elasticity among all consumers in the experiment.

To provide a contrast with the estimated elasticities, I also do back of the envelope calculation that uses the average monthly discount among all cross-randomized treatment groups, then calculates the maximum elasticity using the pooled treatment effects in table 1.1. This calculation implies that the price elasticity of demand in Rwanda is at least between -0.59 and -0.54 in Rwanda. In Kenya, the same calculation yields a price elasticity between -0.3 and -0.17.[5]

The difference between the elasticity I compute using my model and 2SLS and what I find

[4]Bootstrapped standard errors are forthcoming.

[5]I calculate these rough elasticities by dividing the maximum average treatment effects (the highest treatment effects included in the 95% confidence interval) based on Table 1.1 for the bulk discount and the reliability incentive by the average price discount induced by the incentives. The average price discount for qualifying for the minimum incentive in each quarter of figure 1.1 is 7.68%.

using back of the envelope calculations with reduced form treatment effects suggests that there may be substantial heterogeneity in consumers' price elasticity of solar. I examine this possibility by repeating the back of the envelope calculations with the heterogeneous treatment effects that I estimate. In Rwanda, these range from elasticities of zero to elasticities of -2.76. In Kenya, they range from zero to -1.84. These elasticities cover a wide range of consumer responses, and further suggest that there may be a large degree of heterogeneity in the price elasticity of solar among different consumers.

1.7 Evidence on Consumer Demand and Contract/Market Frictions

While the estimated treatment effects could be due to inelastic demand for electricity among a large subset of consumers, there are features of the incentives, the PAYGo contracts, and the market environment that could also contribute to consumer non-responsiveness. Below, I describe these non-price features of PAYGo contracts, consider how they may affect consumer responses to the offered incentives, and provide suggestive evidence on their importance in explaining the estimated treatment effects.

Intertemporal Substitution

PAYGo contracts may enable consumers who cannot afford solar access all of the time to substitute electricity consumption intertemporally. For instance, a consumer may choose to allow their account balance to run down to zero and leave their system off for one or two days to save money. This could be a particularly effective strategy for consumers who have a number of rechargeable appliances, as they could charge all appliances the day before their system is shut off, use the appliances until the charge runs out, and only then purchase additional system access time.

Recall that one feature of the PAYGo contract is continuous rundown of system access time. If a consumer purchases ten days of access, they must use that access time for the next ten days. Therefore, if a consumer engages in intertemporal substitution, purchasing additional days of access time to qualify for an incentive may not be an attractive option, as the discounts induced by the incentives could be less than the money saved by periodically forgoing access altogether.

To assess the importance of intertemporal substitution in explaining my estimated treatment effects, I leverage system use data for the universe of solar consumers who are customers of my partner company in Rwanda and Kenya. In the use data, I observe the number of watt hours (wH) actually generated and used by each consumer's system every day. I combine this with payment data and customer covariates to better understand the prevalence of intertemporal substitution using the following specification:

$$wH_{it} = \alpha + \beta_1 Purchase_{it} + \beta_2 SwitchOn_{it} + \beta_3(Purchase_{it} \times Recharge_i) + \beta_4(SwitchOn_{it} \times Recharge_i) + \gamma_i + \gamma_t + \epsilon_{it}. \quad (1.8)$$

Here wH_{it} is the number of watt hours consumed by consumer i on day t. $Purchase_{it}$ is a dummy variable equal to one when consumer i has purchased additional system access time on day t *when their system was already switched on*. For clarity, I will refer to such purchases as continuation purchases. $SwitchOn_{it}$ is a dummy variable equal to one when consumer i has purchased additional system access time on day t that results in their system being switched on after being switched off. $Recharge_i$ is a dummy equal to one if consumer i owns rechargeable appliances. γ_i and γ_t are consumer and day fixed effects, respectively.

If consumers are engaging in intertemporal substitution, I expect that $\beta_2 > 0$: consumers increase electricity consumption more when their systems get switched back on than on days when they make a purchase that simply maintains system access. Furthermore, I expect that $\beta_4 > 0$ and $\beta_4 > \beta_2$: consumers with rechargeable appliances increase electricity consumption more when their systems get switched back on than consumers who do not have rechargeable appliances.

It is also possible that intertemporal substitution is reflected not as a spike in use on the day following a system being switched back on, but on the day prior to a system being switched off. Therefore, I estimate an analogous specification for days prior to being switched off compared to all other days that a consumer's system is turned on:

$$wH_{it} = \alpha + \mu_1 PriorOff_{it} + \mu_2 PriorOff_{it} \times Recharge_i + \gamma_i + \gamma_t + \epsilon_{it}. \quad (1.9)$$

$PriorOff_{it}$ is a dummy equal to one if consumer i's system was switched off on day $t+1$. The intuition and hypotheses for this specification are similar to those for equation (1.8). If consumers are engaging in intertemporal substitution, I expect that $\mu_1 > 0$ and $\mu_2 > 0$.

Figures 1.4 and 1.5 show the estimated differences in electricity consumption for equations (1.8) and (1.9) in Rwanda and Kenya. In Rwanda, I find that consumption is significantly higher on days when systems are switched back on relative to all other days when consumers have system access but did not make a purchase. However, I cannot reject that the difference is the same as the difference on days when a consumer makes a continuation purchase. I do find that consumers with rechargeable appliances consume significantly more electricity on the day prior to being switched off relative to all non-purchase days of system access, while the difference is negative for consumers without rechargeable appliances. While the signs of these estimates suggest that consumers are engaging in intertemporal substitution, the magnitudes are small: in Rwanda, they comprise only around 1.6% of mean electricity consumption for consumers with rechargeable appliances.

In Kenya, I find that consumers actually consume less on days when they are switched back on relative to days when they make continuation purchases. As in Rwanda, I also observe that consumers with rechargeable appliances consume significantly more electricity

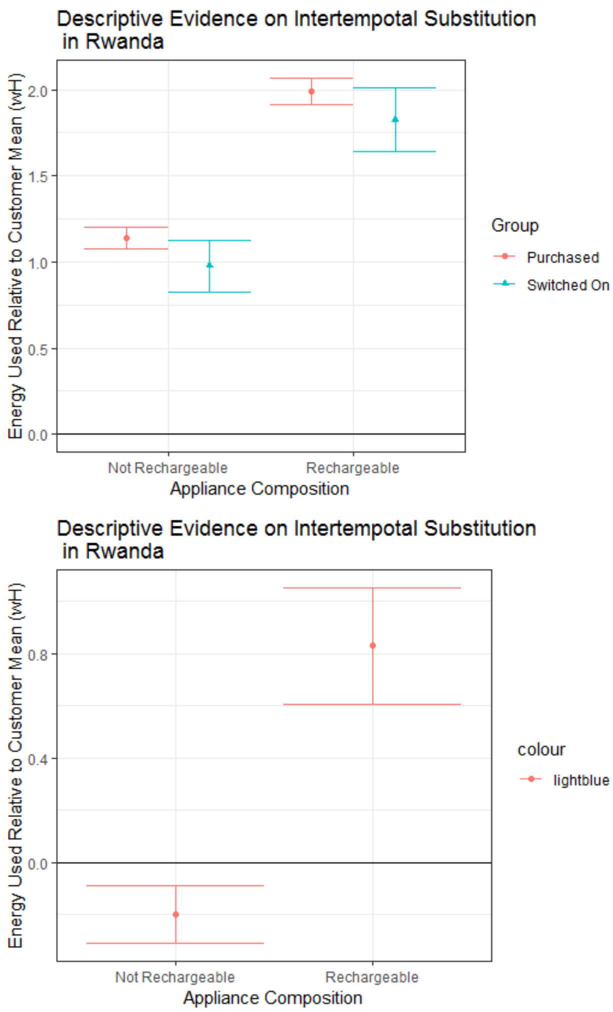

Figure 1.4: Evidence on Intertemporal Substitution in Rwanda

on the day prior to being shut off relative to all other days that they have access to their solar home system; however, the effect is again small at only around 1% of mean consumption. Taken together this evidence suggests that intertemporal substitution is likely not driving consumer non-responsiveness to incentives.

Uncertainty Regarding the Future Marginal Utility of Electricity

Another consequence of continuous rundown is that consumers may be unwilling to purchase system access time very far in advance if they experience uncertainty about their future marginal utility of solar. For instance, a consumer who may be away from home will want to wait to purchase system access time until she knows whether she will actually be at home

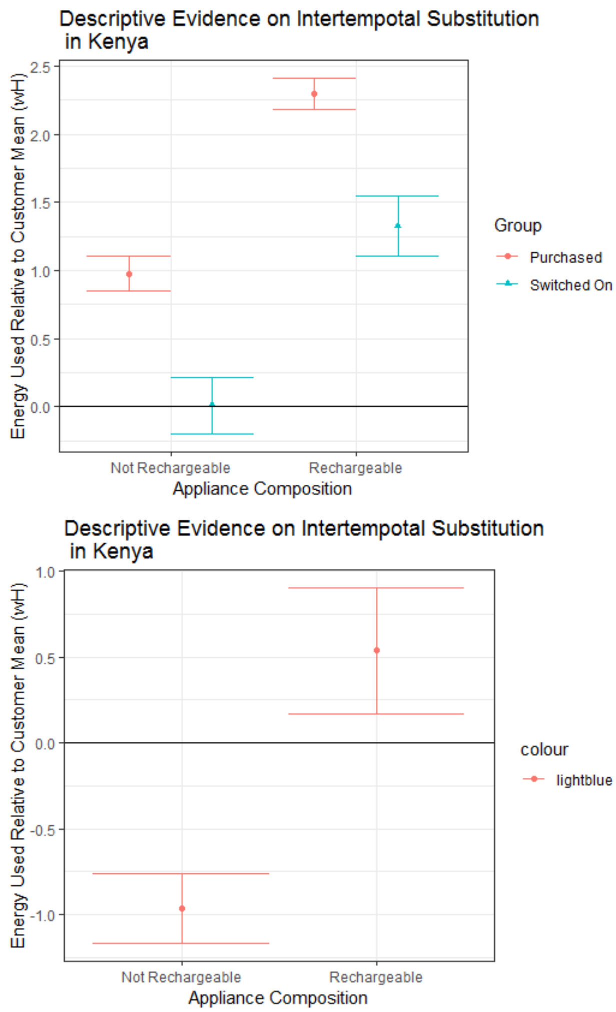

Figure 1.5: Evidence on Intertemporal Substitution in Kenya

to use it. More generally, there may be days when it is easier for consumers to forgo system access. This uncertainty about the future marginal utility of electricity will make incentives unappealing, as they require purchasing more system access time in advance than consumers otherwise would.

Figure 1.6 shows a time series of system use for two sample consumers in Rwanda during October, 2018, with randomly selected consumers represented in the background to provide a sense of the cross-sectional variation in use. As figure 1.6 shows, consumers with similar mean usage may have substantially different variance in their use. This is the variation that I leverage to better understand the importance of uncertainty over the future marginal utility of consumption.

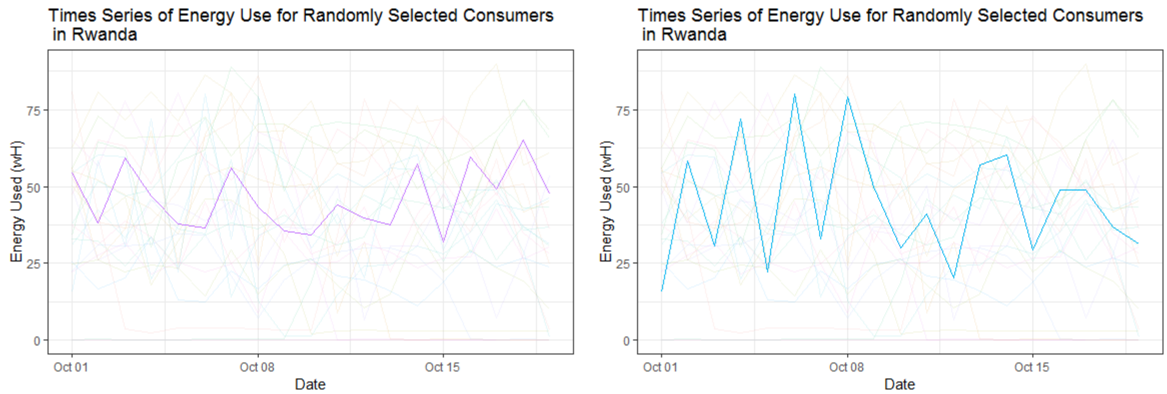

Figure 1.6: Low Versus High Variance Consumer - Rwanda

I estimate a linear probability model on consumers who were offered a monthly reward:

$$Qualify_{it} = \alpha + \beta_1 StDev(Use_i) + \beta_2 Rate_i + \beta_3 Treatment_{it} + \beta_4 InframarPur_{it} + \sum_{j=1}^{5} \eta_j (StDev(Use_i) \times Bin_j) + \epsilon_{it}. \quad (1.10)$$

$Qualify_{it}$ is a dummy equal to one if consumer i qualifies for an incentive in month t. $StDev(Use_i)$ is the standard deviation in use for consumer i, on days when consumer i has access to their solar home system. $Rate_i$ is the daily rate paid by consumer i, and simply controls for the amount of cash required for consumer i to purchase an additional day of solar access. $Treatment_{it}$ is the cross-randomized treatment group that consumer i is in during month t, which controls for both the qualifying threshold and the size of the incentive. $InframarPur_{it}$ is the total number of days purchased by consumer i in month t prior to their final, or marginal, purchase for the month. Finally, Bin_j is a dummy indicating whether consumer i is in a given bin of inframarginal purchases.

The intuition behind equation (1.10) is that I compare consumers in the same treatment group facing the same marginal purchase at the end of a month, then test to see whether consumers with a higher variance in their electricity consumption are less likely to make a large enough purchase to qualify for a monthly reward than consumers with a lower variance in their electricity consumption.

Figures 1.7 and 1.8 show the estimated results for five marginal purchase bins. In both Rwanda and Kenya, consumers who have purchased twenty days or less prior to their final purchase of the month are less likely to qualify for an incentive if they have a higher variance in their use. In Rwanda, a one standard deviation increase in a consumer's standard deviation in electricity use is associated with a 17.4% reduction in the probability of qualifying for an incentive, and in Kenya is associated with a 14.75% reduction in the probability of qualifying for consumers in the fifteen to twenty day bin.

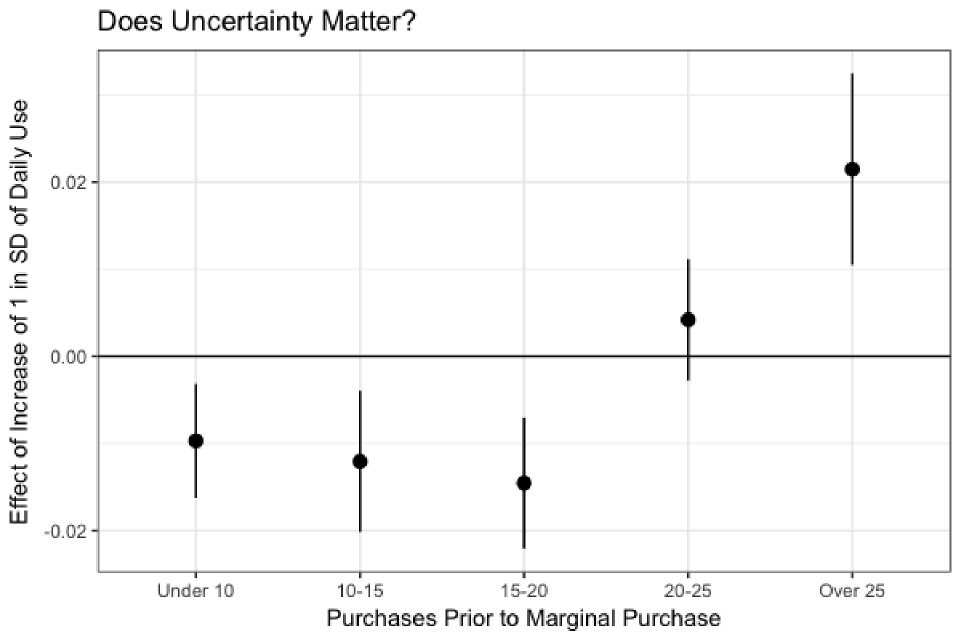

Figure 1.7: Evidence on Uncertainty in Rwanda

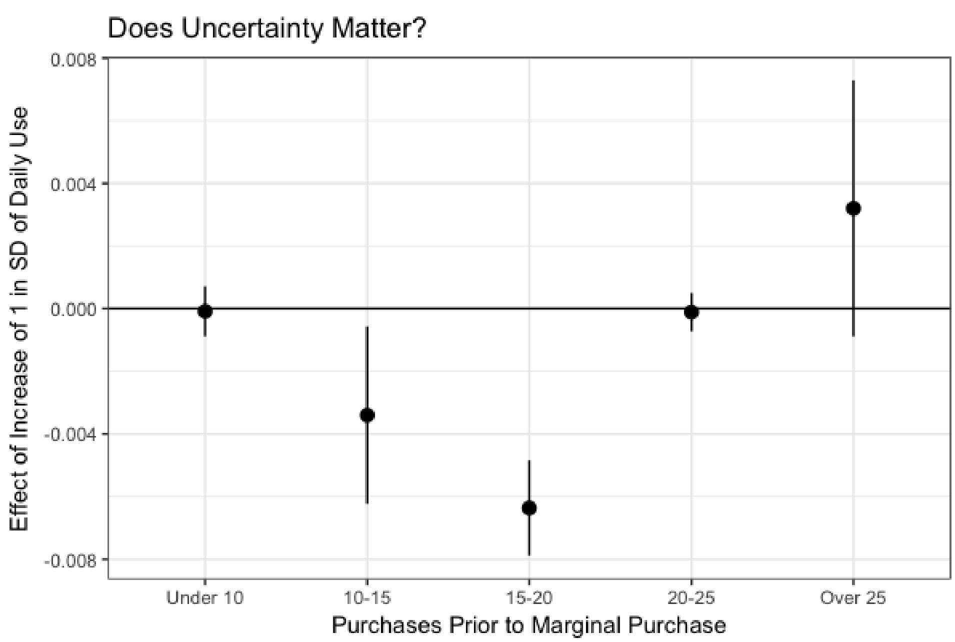

Figure 1.8: Evidence on Uncertainty in Kenya

In both countries, consumers with more than twenty-five days in inframarginal purchases are more likely to qualify for the incentives if they have a higher variance. In Rwanda, a one standard deviation increase in the standard deviation for consumers in this bin is associated with a 24.37% increase in the probability of qualifying for an incentive, and in Kenya is associated with a 8.26% increase, although in Kenya the estimate is not significantly different from zero.

Although variance in electricity consumption is not necessarily equivalent to uncertainty over the future marginal utility of solar, this evidence suggests that uncertainty may play a substantial role in explaining consumer non-responsiveness at the lower end of the utilization distribution. Interestingly, consumers with higher inframarginal purchases seem to be more likely to qualify for incentives when they face more uncertainty over their use, potentially reflecting a higher overall marginal utility of solar for these consumers.

Liquidity Constraints and Transaction Costs

Another possible explanation for consumer non-responsiveness is that the liquidity consumers have to forgo to qualify for an incentive is too high, even for the monthly reward. While I cannot provide any direct evidence on this mechanism, survey evidence suggests that consumers in Rwanda value liquidity highly.

In a survey of solar customers in Rwanda who had taken part in the experiment, I asked a series of three questions aimed at getting a sense of the transaction costs consumers incur to purchase solar: how long it takes to reach the nearest mobile money agent, how many times the consumer visited a mobile money agent to pay for solar of the last five times they paid for solar, and whether the consumer would know how to pay for solar if they did not visit a mobile money agent.

Figure 1.9 shows the distribution of time to the nearest mobile money agent. While many consumers can reach a mobile money agent in under an hour, a substantial share face travel times of one or more hours to reach the nearest agent. Figure 1.10 shows that a large majority of consumers have visited a mobile money agent to pay every single time they have purchased solar, of their last five purchases. Notably, this is true both for consumers who reported that they would know how to make a solar payment on their own if they had sufficient funds in their mobile money wallet and those that reported that they did not know how to make mobile money payments on their own.[6] This suggests that most consumers are not visiting mobile money agents to pay for solar due to poor understanding of mobile money. Finally, figure 1.11 shows that there is very little correlation between the distance to the nearest mobile money agent and the frequency with which a consumer visits a mobile money agent to pay for solar rather than paying on their own.

To summarize, this survey evidence suggests that consumers may face high transaction costs to pay for solar in the form of travel time to the nearest mobile money agent. Fur-

[6] Consumers who report that they do not know how to make mobile money payments on their own but who did not visit a mobile money agent each of the last five visits likely received assistance from a neighbor or family member.

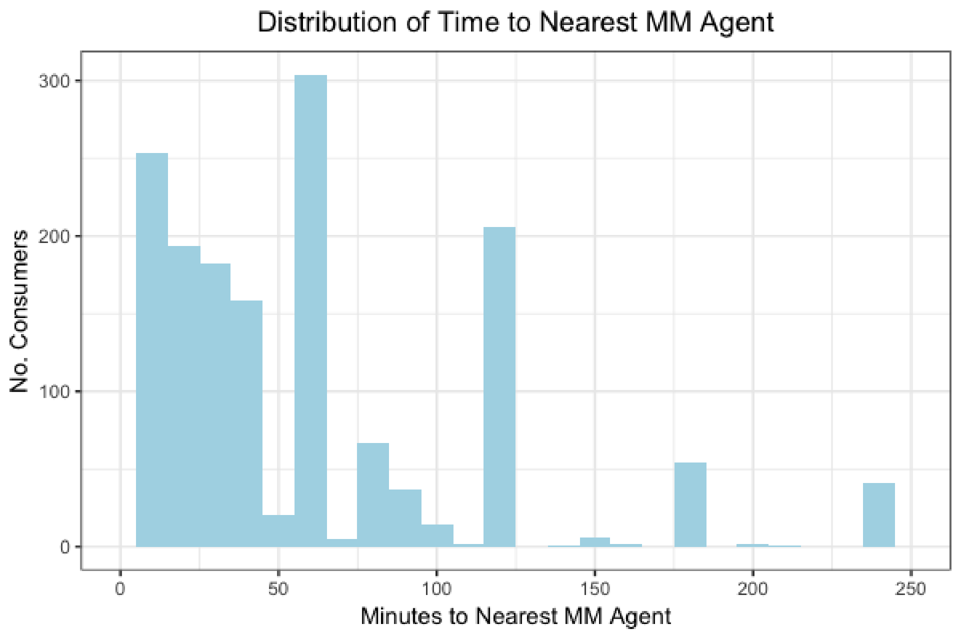

Figure 1.9: Distribution: Time to Reach Nearest Mobile Money Agent in Rwanda

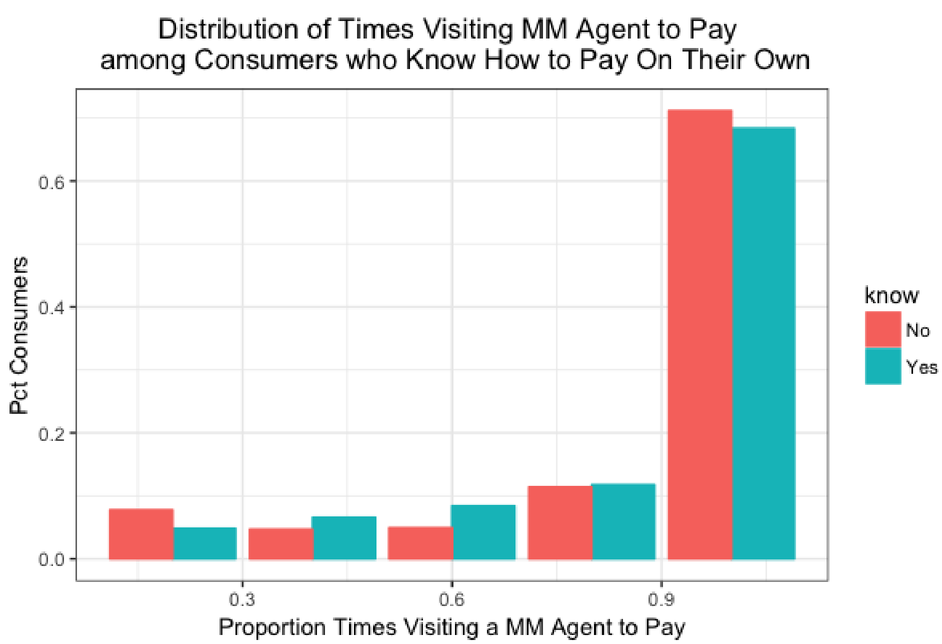

Figure 1.10: Percentage of Last Five Solar Purchases Made with MM Agent, Rwanda

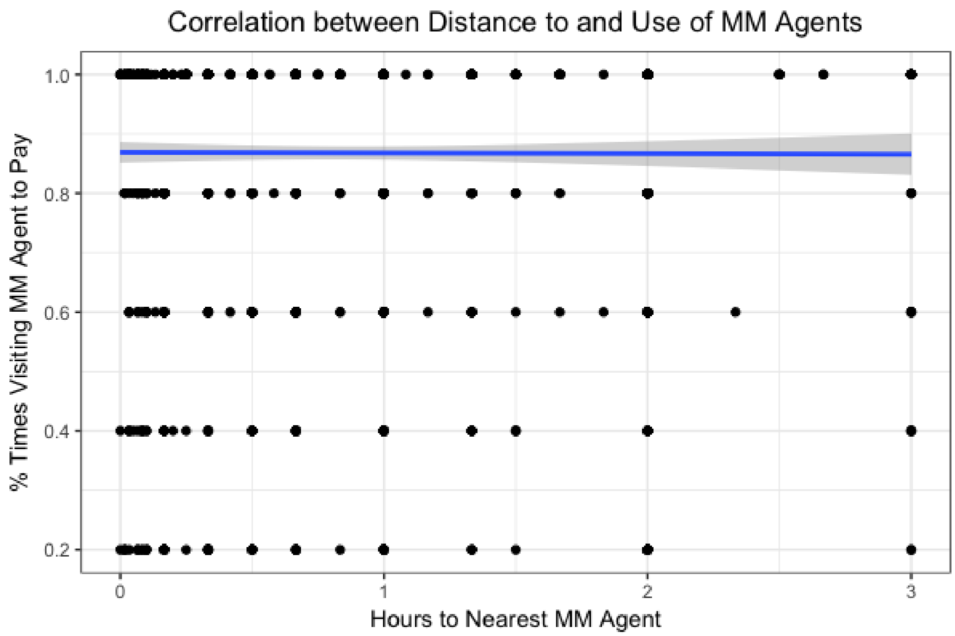

Figure 1.11: Correlation between Distance to and Use of MM Agent to Pay for Solar, Rwanda

thermore, it appears that consumers are choosing to repeatedly incur those high transaction costs rather than making larger solar purchases or simply holding mobile money in their wallets. This is true regardless of the size of the transaction cost (as shown in figure 1.11) and regardless of the consumer's proficiency with mobile money (as shown in figure 1.10). The median purchase size for consumers in my survey sample is seven days of solar. Therefore, consumers may be incurring high transaction costs on a weekly basis rather than forgoing liquidity. I take the survey responses as suggestive evidence that consumers may be liquidity constrained, potentially explaining consumer non-responsiveness to incentives.

Importantly, I only have survey evidence from Rwanda. As stated earlier, mobile money penetration is substantially higher in Kenya. High mobile money penetration effectively makes mobile money more liquid, reducing the risk of converting cash to mobile money and thereby reducing the transaction costs required to make a solar purchase. While I cannot speak directly to the importance of liquidity constraints in Kenya, these differences in the institutional setting suggest that liquidity constraints may play a smaller role there relative to Rwanda.

1.8 Evidence on Behavioral Responses to Incentives

While market frictions may partially explain consumer non-responsiveness, it is also possible that behavioral biases impeded changes to consumer behavior. I consider two behavioral

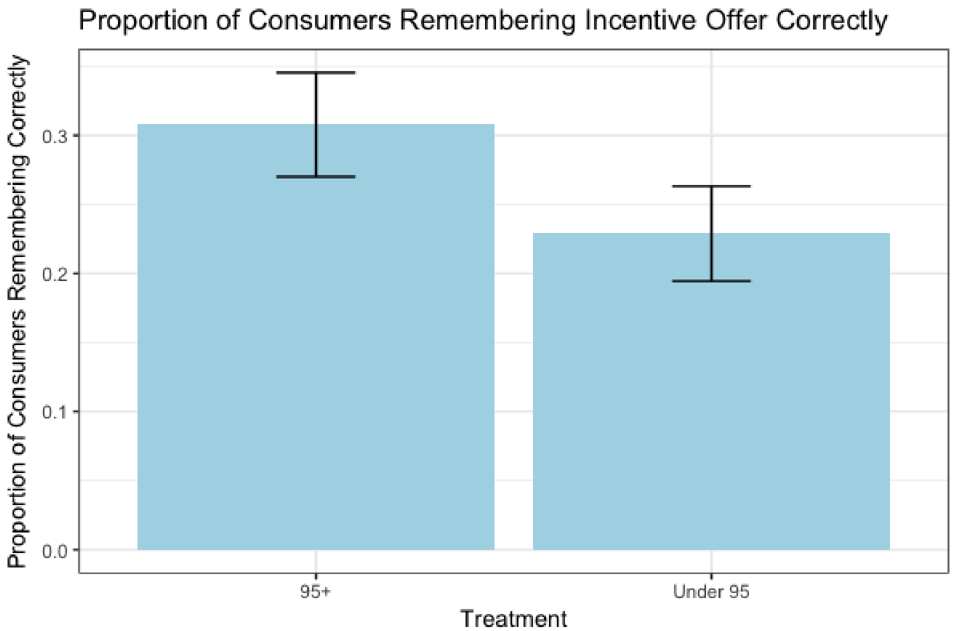

Figure 1.12: Evidence on Salience of Incentives, Rwanda

biases. First, an initial phone call and weekly SMS reminders about the incentives may not have been sufficient to make the incentives salient for consumers. Second, estimated null results for the monthly reward could be explained by high cognitive costs. To qualify for the monthly reward, consumers needed to track their purchases over the course of a month. If consumers struggled to do so, or were simply guessing about how many days they needed to purchase at the end of a month to qualify, responses may have been muted.

Salience

In Rwanda, I asked consumers whether they recalled being offered a promotion by the solar company at any point between May 2018 and the survey date, which was just after the incentives ended at the end of February. Figure 1.12 plots the proportion of treated consumers who recall being offered a promotion among consumers in the 95% and above utilization bin, where I estimate positive and significant treatment effects, and in all other utilization bins where I estimate null treatment effects. Salience is low among all sampled consumers, though it is significantly lower among consumers in utilization bins below 95% than it is for consumers in the 95% and above utilization bin.

Higher recall among consumers in the 95% and above utilization bin suggests that salience may be closely tied to actual use of the incentives. To further investigate this possibility, I plot the proportion of consumers who recall being offered a promotion as a function of the last month in which they actually qualified for an incentive. In figure 1.13, I see that

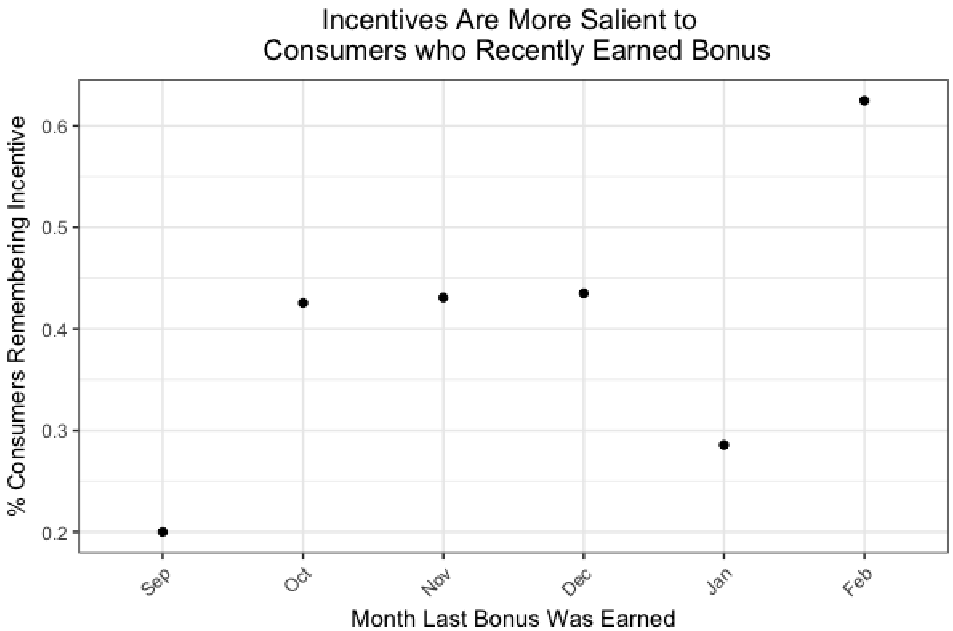

Figure 1.13: Further Evidence on Salience of Incentives, Rwanda

consumers who qualified for an incentive in February, the final month of the experiment, are substantially more likely to recall being offered a promotion relative to consumers who last qualified for an incentive earlier in the experiment. This provides further evidence that use of incentives may be closely tied to their salience, as consumers who earned bonus days earlier in the experiment forgot about the incentive by the time of the phone survey.

In summary, it does appear that salience of the incentives was low in spite of weekly SMS reminders. However, it is not clear whether low salience is the cause of estimated null treatment effects or whether consumers who found it difficult to qualify for incentives simply ignored reminders about them, engaging in some form of rational inattention. Survey evidence suggests that consumers who were more likely to use the incentives were also the most likely to remember them, but again the direction of causality between recall and use is unclear.

Cognitive Costs

To evaluate the importance of cognitive costs, I plot the proportion of consumers qualifying for an incentive as a function of the effective discount they would have received had they made a large enough purchase in their final transaction of the month to qualify for an incentive. While consumers with low inframarginal purchases are unlikely to make a large enough purchase to qualify, I would expect that close to 100% of consumers facing 50% discounts or above should qualify. However, if consumers struggle to correctly track the

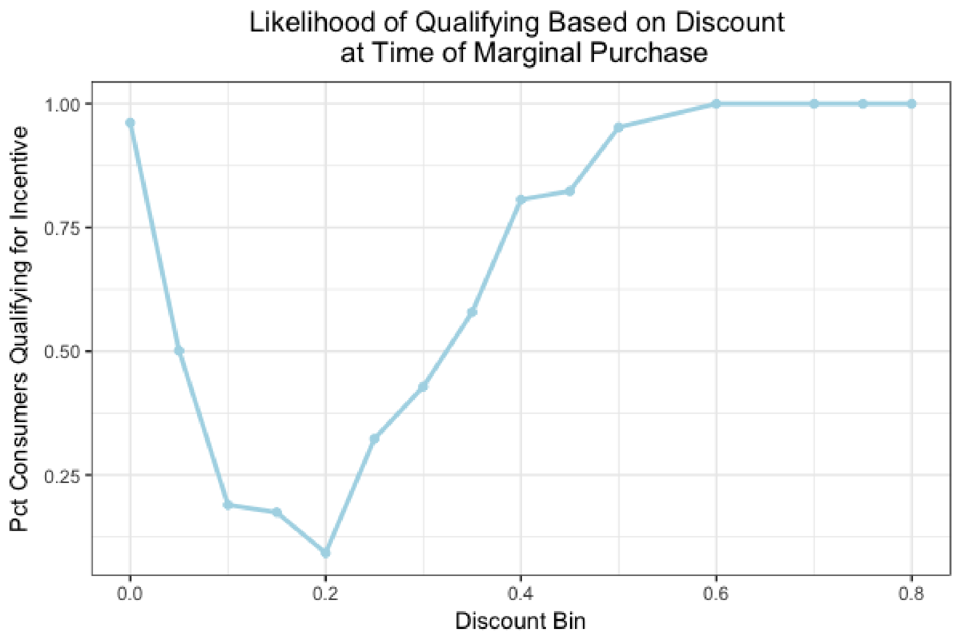

Figure 1.14: Evidence on Cognitive Costs, Rwanda

number of days they have purchased in a month, they may mistakenly purchase too little to qualify.

Figures 1.14 and 1.15 show the relationship between the effective discount and the proportion of consumers qualifying for the monthly reward in Rwanda and Kenya, respectively. The plots are fairly similar, and both show that around 90-95% of consumers facing a discount of 50% or more at the time of their marginal purchase do qualify for the incentive. Thus, it seems unlikely that cognitive costs explain the estimated null effects for the monthly reward, though the evidence on this channel is weak at best.

1.9 Discussion

My results suggest that consumers may be heterogeneous in their price elasticity of solar, and that non-price factors may play a critical role in shaping consumer demand for electricity in low-income countries. How electricity is provided and the market context in which consumers operate have significant implications for demand. The growing popularity of PAYGo contracts and pre-paid electricity meters on the grid is largely a response to market frictions and contractual challenges. Pre-paid electricity reduces enforcement costs in settings where institutions are weak, makes price more transparent and salient to low-income consumers, and prevents consumers from over-consuming electricity and failing to pay their bills (K. Jack and Smith 2020).

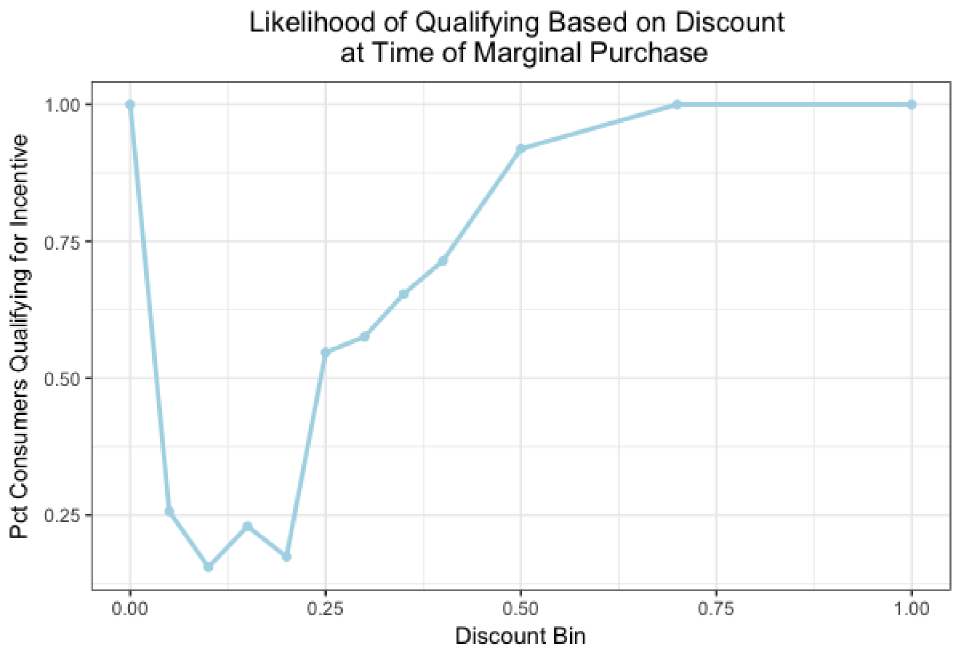

Figure 1.15: Evidence on Cognitive Costs, Kenya

Although PAYGo contracts overcome many of the challenges of electricity provision in low income settings, my results show that they may still be sub-optimal depending on the institutional context. High transaction costs coupled with liquidity constraints may dampen demand in settings where mobile money penetration is not sufficiently high. In addition to liquidity constraints, consumers may incur high transaction costs due to uncertainty over the future marginal utility of solar, or uncertainty over future cash flows. Importantly, the heterogeneity I observe in the price elasticity of solar may simply be another manifestation of liquidity constraints or uncertainty over the future marginal utility of solar.

It is clear that these market frictions and contractual features are sub-optimal for consumers, but it is unclear whether they are also sub-optimal for firms. For instance, firms could offer consumers the option to save days of solar rather than mandating continuous rundown of system access time. This would give consumers complete control over the timing of their electricity consumption, and could result in a reduction in overall demand. However, it also reduces the risk of waste for consumers, which could increase overall demand. Similarly, firms could allow consumers to overdraft their solar accounts, borrowing system access time from the firm and repaying it later. This increases the risk of overconsumption if consumers are present biased, but it could increase demand without increasing the default rate by effectively lowering transaction costs for consumers.

My results also speak to potential interlinkages between the policy objectives of electrification and financial inclusion. In a setting where mobile money is nearly as liquid as cash, credit constraints and uncertainty over the future marginal utility of solar are the only non-price factors that may explain my estimated treatment effects. Therefore, increasing mobile money penetration could be critical to increasing consumer demand for electricity as well as driving financial inclusion.

1.10 Conclusion

I study incentives that change the average and marginal price of electricity for solar consumers in Rwanda and Kenya to better understand the factors shaping consumer demand for electricity in low income settings. I find that most solar consumers in Rwanda and Kenya do not respond to bulk discounts or monthly rewards that effectively alter the average and marginal price of electricity, but that there is substantial heterogeneity in consumer responses. However, I find that uncertainty over the future marginal utility of solar as well as liquidity constraints could at least partially explain the estimated null treatment effects from incentives.

I identify two key directions for future research. First, my estimates of elasticities are limited. My I cannot estimate an elasticity using 2SLS for Kenya due to a weak first stage, and my first stage in Rwanda is borderline at best. The back of the envelope calculations I perform come from price changes that were associated with incentives rather than changes in the price of each day of solar. Given the limitations of estimating welfare based on demand on the extensive margin, we require more credible estimates of the price elasticity of demand for electricity in these settings on the intensive margin. Government subsidies for electricity use or simple price experiments would yield the required variation to obtain more accurate estimates of the price elasticity of demand.

Second, I can present at best suggestive evidence that liquidity constraints, transaction costs, and uncertainty over the future marginal utility of solar are dampening consumer demand for electricity. Ideally, variation in liquidity constraints or the removal of continuous rundown would enable a better understanding of the relative importance of these non-price factors in different contexts.

Chapter 2

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

2.1 Introduction

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

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

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

I generate predictions about the outcomes of the experiment by incorporating PAYGo access time with transaction costs into Deaton's (1991) buffer stock model. I assume that the utility consumers gain from solar access is stochastic. This assumption is what leads to waste in the model: when consumers buy in bulk, they may end up having access to electricity on a day when they gain little to no utility from it. The model makes two sets of contrasting predictions about how consumers with limited versus ample access to liquidity will react to the offer of credit. First, lowering transaction costs increases demand for access time for consumers with limited liquidity who are not buying in bulk because it effectively lowers the price. Second, guaranteed access to credit increases demand for these consumers regardless of whether they actually borrow by reducing the precautionary savings motive. Consumers with guaranteed access to credit do not need to save as much to smooth consumption. By contrast, consumers with ample liquidity will buy in bulk when transaction costs are high, even though it generates waste.[2] It follows that reducing transaction costs using the line of credit may reduce demand among these consumers by allowing them to eliminate waste. Importantly, these reductions in demand will be the largest among consumers with the highest variance in the utility they gain from electricity: consumers with a high variance in utility benefit the most from more precisely targeting their consumption.

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

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

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

action costs. Put otherwise, the largest reductions in demand occur among consumers who can benefit the most from better targeting their electricity purchases.

Although the empirical results are consistent with my theoretical framework, I do not directly observe the liquidity constraints and transaction costs facing individual consumers throughout the course of my experiment. I consider an alternative underlying mechanism that may generate the same results. If pre-experimental demand is positively correlated with present focus, also sometimes called present bias, then the negative treatment effects I estimate could be the result of high-demand consumers borrowing then procrastinating on repayment. In other words, the line of credit could be causing some consumers to borrow, delay repayment, and ultimately buy less electricity than they would under a strictly prepaid regime. Present focus is an important alternative to consider because if it is driving my results, offering the line of credit may not be welfare enhancing for consumers.

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

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

Methodologically, revealed preference estimates of consumer welfare from solar are inaccurate unless they account for the distortions caused by liquidity constraints and transaction costs. I estimate a conservative lower bound on consumer surplus from electricity using random variation in the fee charged on the line of credit. My estimate suggests that consumer surplus from electrification is higher than previously believed: my lower bound is equal to the most contextually similar estimate for total consumer surplus in the literature. Even so, marginal households' willingness to pay for electricity falls well below cost-covering levels.

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

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

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

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

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

The rest of my paper is organized as follows. Section 2 describes the context. Section 3 details the experimental design and describes my sample. Section 4 provides a theoretical framework to derive predictions about the impact of offering a line of credit for solar access. Section 5 presents reduced form results. Section 6 presents my estimated lower bound on consumer surplus from electrification, and section 7 concludes.

2.2 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

[3] See `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.

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[4].

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

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

Traveling to a mobile money agent to purchase solar access time represents a transaction 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 2.1 shows the distribution of travel time to reach the nearest mobile money agent, combining both survey samples. The average time is 50 minutes and the median time is 30 minutes one-way, although true transaction costs for purchasing solar likely vary depending on the timing of other tasks that might bring consumers close to a mobile money agent.

In theory, consumers can reduce the transaction costs associated with paying for solar by depositing money in their mobile money wallet when they are near an agent and later using those funds to buy solar. In a phone survey conducted among 1,229 solar customers in 2019, I asked consumers how many times they visited a mobile money agent to pay for solar out of their last five purchases. Figure 2.2 shows that 66% visited a mobile money agent all five times and 78% visited four of the last five times.[6] This pattern is likely the result of limited

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

[5]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.

[6]Appendix figure 2.16 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.

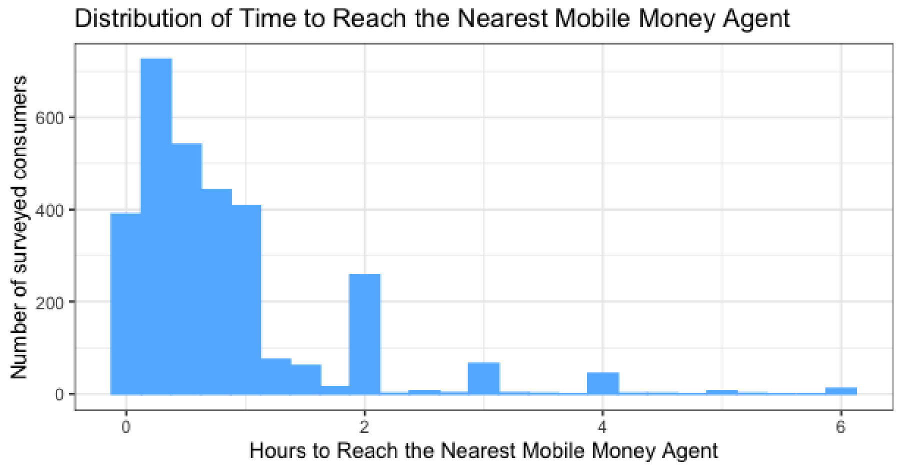

Figure 2.1: Consumer Travel Times to Nearest Mobile Money Agent

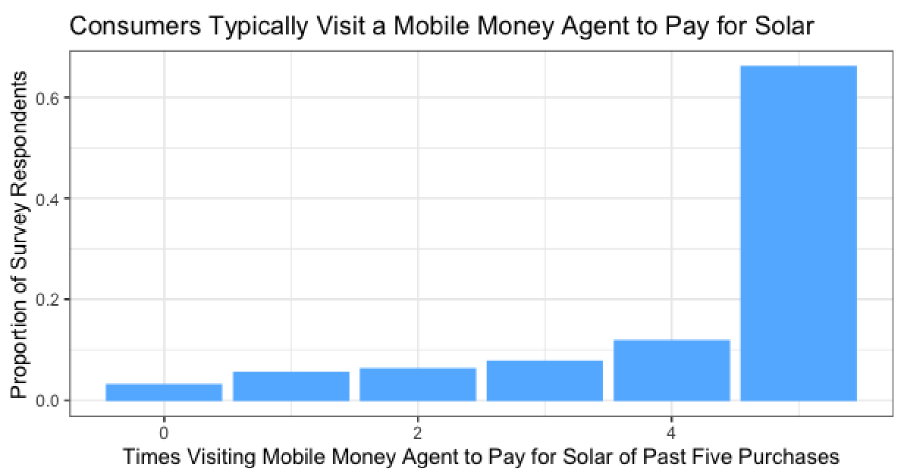

Figure 2.2: Self-Reported Use of Mobile Money Agents to Buy Solar

mobile money take-up. A 2018 report by the World Bank found that only 31.1% of adults in Rwanda had mobile money accounts (*The Little Data Book on Financial Inclusion 2018* 2018). While transactions conducted with mobile money are free, consumers have to pay withdrawal fees to convert mobile money into cash. With low take-up, consumers cannot use mobile money for most transactions and so withdrawal fees render it less liquid than cash. Even though consumers could use mobile money wallets to lower transaction costs, the survey evidence demonstrates that, in practice, consumers frequently incur transaction costs when paying for solar.

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

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

2.3 Experimental Design

I partner with a solar company in Rwanda to offer existing PAYGo solar customers a product designed to alleviate liquidity constraints and reduce transaction costs: a line of credit specifically for PAYGo system access time. The line of credit allows consumers to call the solar company and request to use up to one or two weeks of system access time before paying for it. When a consumer makes a PAYGo payment after borrowing, the funds first go to repaying the time they borrowed plus a flat fee. Any funds that are left after repaying the line of credit go to pre-paying for additional system access time. In this way, consumers cannot default on the line of credit without defaulting on their entire PAYGo contract. Consumers can use the line of credit as many times as they like over the course of the experiment.

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

I cross-randomize the terms of the line of credit along three dimensions. Half of the consumers in the treatment group can only borrow up to seven days of solar access time, while the other half can borrow up to fourteen days. All consumers have the option to choose a voluntarily lower borrowing limit than the one originally offered, which allows consumers to commit to borrowing smaller amounts. Half of consumers pay a 10% flat fee on borrowed days and half pay a 2% fee.[8] Finally, half of consumers lose access to the line of credit if they do not repay their borrowed days plus the fee within one week of their borrowed time

[7] Appendix figure 2.17 shows the full distribution of days purchased in a single transaction in the 90 days prior to the experiment.

[8] 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.

running out. The other half do not face any such time limit, but like all PAYGo customers they get remotely locked out of their system when they run out of access time.

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

Consumers in my sample are self-selected, as they have all opted to sign PAYGo solar contracts. I combine responses to a phone survey conducted with 1,229 solar customers in 2019 with the latest Integrated Household Living Survey, a nationally representative survey of Rwandan households last conducted in 2016-2017. Using questions common to both surveys, I construct a wealth index to compare the population in my sample to the distribution of rural households in Rwanda.[9] I find that consumers in my sample are wealthier than the average rural Rwandan household, a feature I return to in my discussion of the welfare impacts of rural electrification.[10]

I stratify my treated sample based on the 90-day utilization rate (UR) prior to the start of the experiment. The utilization rate is the proportion of days a consumer has purchased system access. To understand consumer responses to the line of credit across the 90-day UR distribution, I create four stratification bins: 0-30%, 30%-65%, 65%-80%, and 80%-100%. Figure 2.3 shows the distribution of UR in the 90 days prior to the start of the experiment, along with lines denoting the stratification bins.

Descriptive information about differences between consumers in each utilization bin shows that pre-experimental demand is correlated with access to liquidity and consumers' propensity to buy in bulk. Table 2.1 shows differences in self-reported borrowing to pay for solar when the line of credit is not available.[11] Consumers with the lowest pre-experimental demand are significantly less likely to have borrowed for solar than consumers in other utilization bins, and those who do borrow appear to borrow less. Of the consumers who have not borrowed to pay for solar, those with the lowest pre-experimental demand are more likely to report that they did not borrow because they were unable to find credit. Table 2.2 shows that consumers with the highest pre-experimental demand are more likely to buy in bulk: average purchase sizes for consumers with the highest pre-experimental demand are more than double those for consumers with the lowest pre-experimental demand.[12]

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

[10] Appendix figure 2.18 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.

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

[12] Appendix table 2.9 shows differences between stratification bins along other dimensions of economic well-being that are less closely connected to the decision to buy in bulk versus buy small quantities frequently.

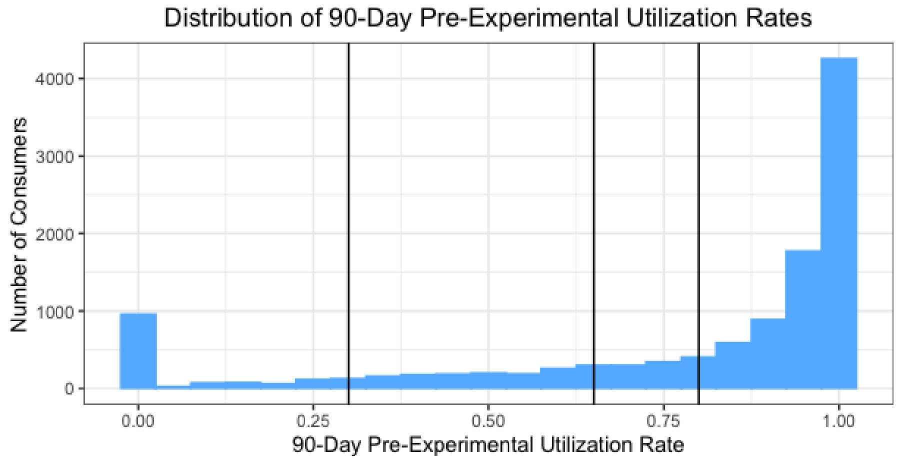

Figure 2.3: Distribution of Demand Prior to the Experiment

Table 2.1: Differences in Borrowing Behavior between Stratification Bins

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

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

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

Table 2.2: Average Pre-Experimental Payment Sizes by Stratification Bin

	Dependent variable:
	Mean Purchase Size (Days)
0%-30%	-8.304^{***}
	(0.538)
30%-65%	-2.970^{***}
	(0.296)
65%-80%	-3.535^{***}
	(0.353)
Daily Rate	-0.011^{***}
	(0.001)
Intercept (80%-100% Mean Payment Size)	15.425^{***}
	(0.259)
Observations	1,342
Note:	*p<0.1; **p<0.05; ***p<0.01

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

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

Timeline and Data

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

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

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

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

2.4 Theoretical Framework

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

[13] 64% of consumers in the 0%-30% stratification bin were reached on the phone during at least one of the rounds of calls. 86%, 90%, and 96% of consumers were reached in the 30%-65%, 65%-80%, and 80%-100% bins.

[14] Appendix table 2.8 shows that randomization yields balanced treatment and control groups on a range of observable characteristics after controlling for each consumer's stratification bin.

[15] 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 mid-experiment survey only included treated consumers.

describes the continuous rundown of system access time

$$e' = e + q - (\mathbf{1}(e + q \geq 1)). \tag{2.1}$$

Consumers get stochastic utility α if they have at least one day of electricity access ($e + q \geq 1$) in a given day. I assume α is drawn from some distribution with cdf $F(\cdot)$ with support over $[\alpha_m, \alpha_M]$, where $\alpha_m \geq 0$ and α_M is some finite constant. Combined with the continuous rundown described in equation (2.1), stochastic utility generates waste: if a consumer pays for solar access into the future, they may end up having access on a day with a low realization of α when they would not have opted to purchase access otherwise. At the start of each time period, consumers learn their draw of α and receive a stochastic endowment of income $y \in (0, y_M]$, where y_M is some finite constant.

Consumers choose how many days of solar access to purchase and how much of the composite consumption good to consume, which leaves some level of savings to be carried forward to the next day. As equation (2.1) 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

$$U = \mathbb{E}\Big[\sum_{t=0}^{\infty} \beta^t \big(\alpha_t \mathbf{1}(e_t + q_t \geq 1) + u(c_t)\big)\Big]. \tag{2.2}$$

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

The relative price of solar is p. Consumers incur a transaction cost τ each time they buy solar. I make the simplifying assumption that τ is constant across time and consumers and that it represents the costs of reaching the nearest mobile money agent. I denote the consumer's non-solar asset stock as s. Any assets that the consumer saves in the current period earn a return $1 + r$, or if consumers borrow they have to repay $1 + r$ multiplied by the borrowed amount in the following period. The consumer's asset stock is governed by the law of motion

$$s' = (1 + r)(s - c - qp - \mathbf{1}(q \geq 1)\tau) + y'. \tag{2.3}$$

Finally, the consumer faces a borrowing constraint l. In each period, the consumer needs to choose c, s', and q to satisfy[16]

$$-l \leq s'. \tag{2.4}$$

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

Characterizing Optimal Choices

Each day, the consumer chooses c and q to maximize (2.2) subject to (2.1), (2.3), and (2.4). For a given q, the Bellman equation is

$$V_q(s,e,y,\alpha) = \max_{s'} \Big[\alpha\mathbf{1}(e+q \geq 1) + u(c(q)) + \beta\mathbb{E}[V(s',e',y',\alpha')]\Big]$$
$$\text{s.t. } s' = ((1+r)s - c(q) - qp - \mathbf{1}(q \geq 1)\tau) + y',$$
$$e' = e + q - \mathbf{1}(e+q \geq 1), \text{ and } -l \leq s'. \quad (2.5)$$

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

$$V(s,e,y,\alpha) = \max_q\{V_q(s,e,y,\alpha)\}. \quad (2.6)$$

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

Let μ_q be the Lagrange multiplier on the liquidity constraint (2.4) for a given choice of q. Let $\mathbb{E}[\overline{V_q}(s',e',y',\alpha')] = \mathbb{E}[V(s',e',y',\alpha')|q]$ be the maximal expected V for a given choice of q. The interior solutions to the sub-problems defined by equation (2.5) are characterized by the first order condition

$$\beta\mathbb{E}\Big[\frac{\partial\overline{V_q}(s',e',y',\alpha')}{\partial s'}|s,e,y,\alpha\Big] - \mu_q = \frac{du(c(q))}{dc(q)}. \quad (2.7)$$

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

$$\mathbb{E}\Big[\frac{\partial\overline{V_q}(s',e',y',\alpha')}{\partial s'}|s,e,y,\alpha\Big] = (1+r)\mathbb{E}\Big[\frac{du(c'(q'))}{dc'(q')}|s,e,y,\alpha\Big]. \quad (2.8)$$

I then substitute (2.8) into (2.7) to obtain the Euler equation:

$$\beta(1+r)\mathbb{E}\Big[\frac{du(c'(q'))}{dc'(q')}|s,e,y,\alpha\Big] - \mu_q = \frac{du(c(q))}{dc(q)}. \quad (2.9)$$

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

$$u'(c(q)) = \max\Big\{\beta(1+r)\mathbb{E}[u'(c'(q'))|s,e,y,\alpha], u'(s - qp - \mathbf{1}(q \geq 1)\tau + l)\Big\}. \quad (2.10)$$

Equation (2.10) characterizes the consumer's optimal choice of c for a given q. It is straightforward to show that under certain conditions, there will exist a policy function $\sigma_q(s, e, y, \alpha)$ that defines optimal consumption for a given realization of the state. Importantly, expectations are taken over both α' and y'. Expectations over α' speak to the need for an additional policy function that governs how the probabilities of choosing different levels of q in the future change based on choices of c and q today.

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

$$\begin{aligned}\alpha \geq u(\sigma_0(s, e, y, \alpha)) - u(\sigma_1(s, e, y, \alpha)) + \beta \mathbb{E} V((1 + r)(s - \sigma_0(s, e, y, \alpha)), e', y', \alpha') - \\ \beta \mathbb{E} V((1 + r)(s - \sigma_1(s, e, y, \alpha) - p - \tau), e', y', \alpha'). \quad (2.11)\end{aligned}$$

Call the threshold level of α where equation (2.11) holds with equality α^*. For a given realization of s, e, and y, the consumer prefers buying zero days to one day of solar with probability $F(\alpha^*(s, e, y))$ and prefers buying one day to zero days with probability $1 - F(\alpha^*(s, e, y))$. The function $\alpha^*(s, e, y)$ allows me to take expectations over α'.

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

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

or, rearranging,

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

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

For a given realization of (s, y), the 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

$$\mathbb{E} V(s', e', y', \alpha') = \begin{cases} \mathbb{E}_y V_0(s', e', y'\alpha') \text{ if } e' > 0 \\ F(\alpha^*(s', e', y')) \mathbb{E}_y V_0(s', e', y', \alpha') + \\ (1 - F(\alpha^*(s', e', y'))) \mathbb{E}_y \max\{V_1(s', e', y', \alpha'), V_2(s', e', y', \alpha'), \ldots\} \text{ if } e' = 0. \end{cases} \quad (2.12)$$

Solving the model

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

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

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

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

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

Figure 2.4 illustrates how lowering transaction costs changes α^*. In the figure, the point where V_0 and V_1 intersect represents α^*, the lowest draw of α for which a consumer will opt to buy electricity rather than forgoing it. Figure 2.4 shows that lowering transaction costs reduces α^*, increasing the probability that the consumer chooses to buy one day of solar. As expected, the change in α^* is larger when realizations of s and y are lower.

Prediction 2: relaxing liquidity constraints increases demand among consumers not buying in bulk, even if they do not use the line of credit. Figure 2.5 tells a similar story: increasing the borrowing limit l reduces α^*. For consumers who lack the liquidity to buy in bulk, the line of credit unambiguously increases demand for solar. Importantly, figure 2.5 illustrate another reason that the line of credit increases demand among liquidity constrained consumers: it reduces the precautionary savings motive. V_0 and V_1 are higher under the experimental condition. When consumers have guaranteed access to credit they don't need to save as much today to ensure that they can buy solar tomorrow.

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

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

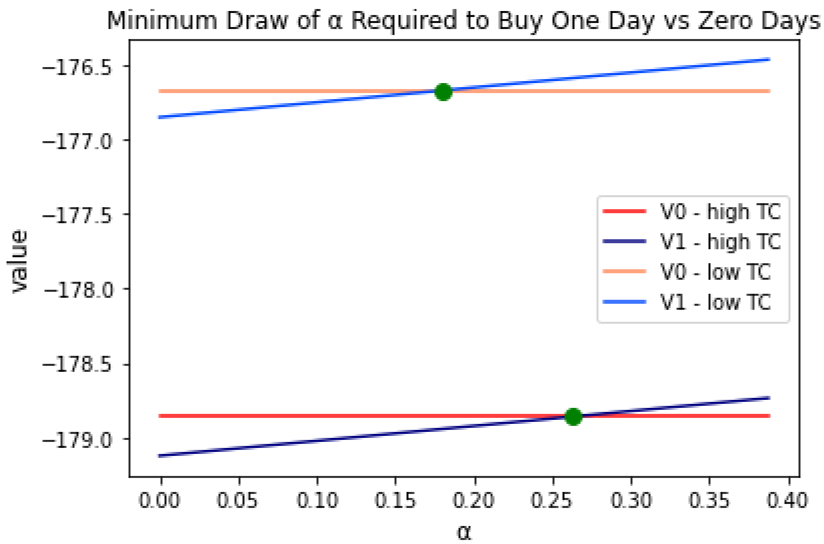

Note: The vertical axis is the level of the value function, with V0 corresponding to the value function associated with forgoing access to electricity in the current period and V1 corresponding the the value function associated with buying one day of access. The horizontal axis is the realization of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

Figure 2.4: Change in α^* from reducing transaction costs.

In figures 2.6 and 2.7, the green points indicate the level of s above which V_2 exceeds V_1 for low and high values of τ and l. The range to the right of the green points is the range of asset realizations over which the consumer prefers to buy in bulk. If the green point moves to the left as a result of my experiment, it follows that the experiment increases the probability that a consumer will prefer to buy in bulk. Figure 2.7 shows that lowering transaction costs reduces the range of assets where the consumer prefers to buy in bulk ($q = 2$). Figure 2.6 shows that relaxing liquidity constraints can either widen or narrow the range of assets where the consumer prefers to buy in bulk depending on the relative size of p and τ.

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

Prediction 4: Negative treatment effects will be larger for consumers with a higher variance in α. Figure 2.7 illustrates the final prediction from my model. Comparing the top and bottom figures, I show that the potential reduction in demand as a result of lowering

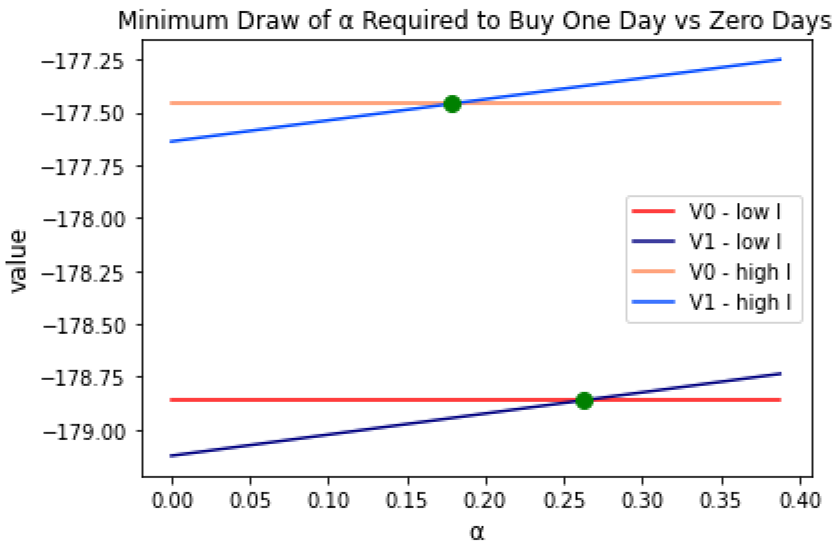

Note: The vertical axis is the level of the value function, with V0 corresponding to the value function associated with forgoing access to electricity in the current period and V1 corresponding the the value function associated with buying one day of access. The horizontal axis is the realization of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

Figure 2.5: Change in α^* from relaxing liquidity constraints.

transaction costs is larger for consumers with a higher variance in α. Intuitively, these are the consumers who will benefit the most from better targeting their consumption. In fact, with the realization of income I selected for the sake of illustration in figure 2.7, lowering transaction costs leads to virtually no change in the probability of buying in bulk for consumers with a low variance in α. The final prediction of the model is that treatment effects from the line of credit should be decreasing in the variance in α.

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

The model generates four empirically testable predictions about offering the line of credit. First, lowering transaction costs will increase demand among consumers who lack sufficient liquidity to buy in bulk. Second, relaxing liquidity constraints will increase demand among consumers not buying in bulk regardless of their borrowing status: providing guaranteed

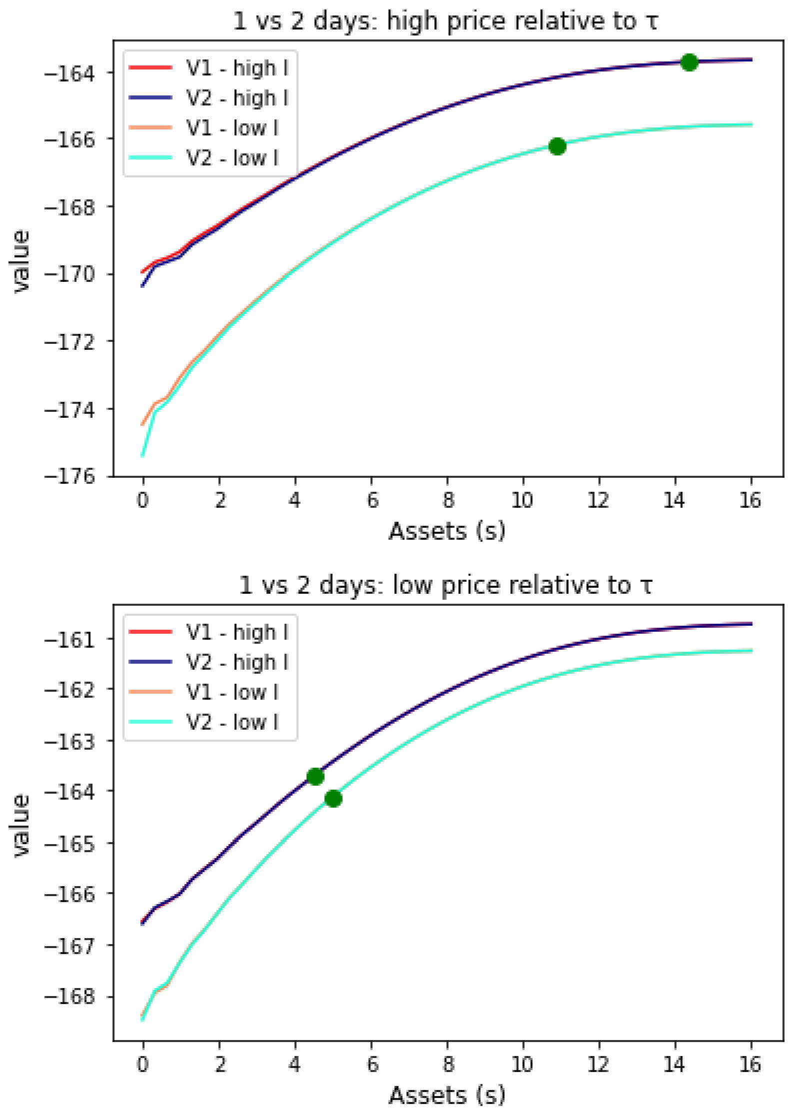

Note: The vertical axis is the level of the value function. The horizontal axis is the realization of assets s in the current period. The green points show the minimum realization of assets for which the consumer prefers to buy in bulk ($q = 2$) to buying a single day. It follows that when the point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

Figure 2.6: Change in the single vs bulk decision from relaxing liquidity constraints.

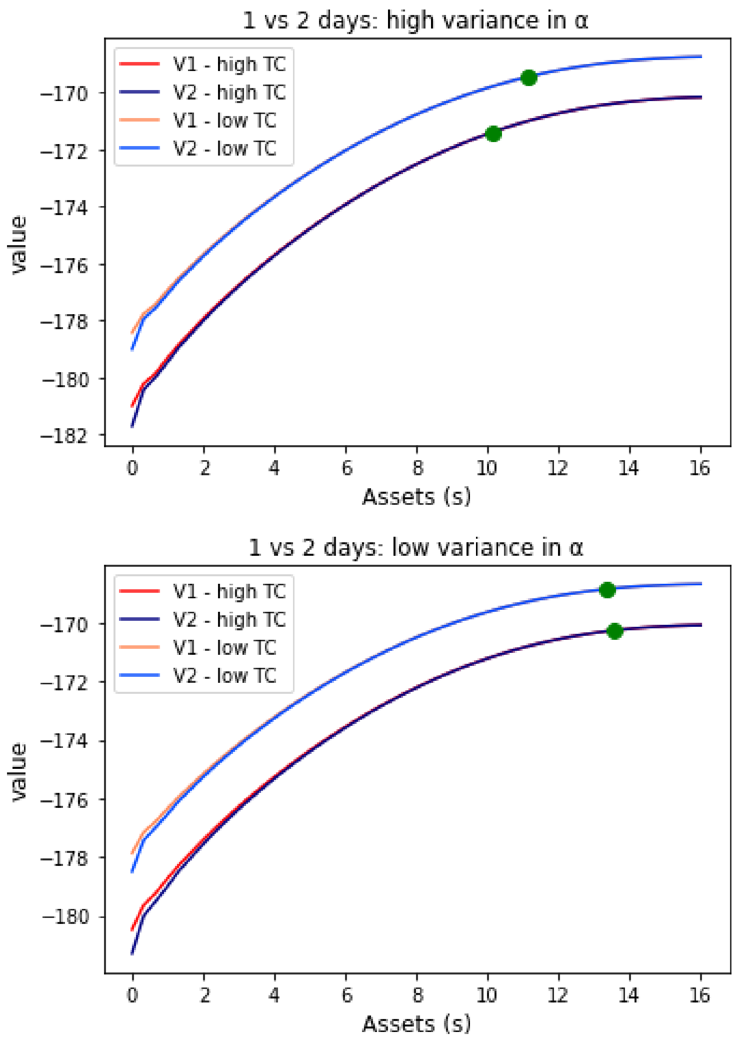

Note: The vertical axis is the level of the value function. The horizontal axis is the realization of assets s in the current period. The green points show the minimum realization of assets for which the consumer prefers to buy in bulk ($q = 2$) to buying a single day. It follows that when the point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

Figure 2.7: Size of the treatment effect for consumers with a high vs low variance in α

access to credit can change consumer behavior even if consumers do not use it. Third, the line of credit may lower demand among consumers who previously bought in bulk. Fourth, if certain consumers do reduce demand in response to being offered the line of credit, then treatment effects from the line of credit will be decreasing in the variance of α if the reduction in demand is the result of lowering transaction costs. Consumers with the most variance in their utility from electricity have the strongest incentive to stop buying in bulk and target their consumption when transaction costs are lower. Next, I present empirical results to evaluate how well the model describes consumer behavior, and to consider potential alternatives explanations.

2.5 Results

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

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

Figure 2.8 shows that average treatment effects from offering the line of credit follow the first and third theoretical predictions in my model. Consumers with the lowest pre-experimental demand increase their monthly utilization rates by 11pp as a result of being offered the line of credit, an increase of 88% over the control group. Consumers in the second-lowest stratification bin significantly increase utilization rates as a result of being offered the line of credit, although at a more modest magnitude of 5.3%. By contrast, consumers in the second highest stratification bin reduce their utilization rate by 6.4%, while consumers with the highest pre-experimental demand reduce utilization by 1.6%.[17] It appears that the line of credit increases demand for consumers who are most likely to be liquidity constrained while allowing consumers who previously bought in bulk to better target their consumption.

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

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

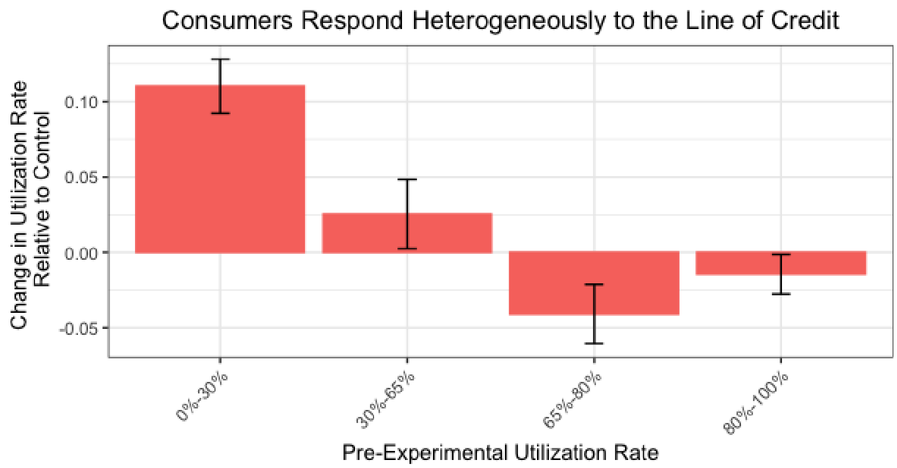

Note: 95% confidence intervals calculated using standard errors clustered at the level of the individual consumer. Estimates for the 30%-65% and 80%-100% bins are not significant at the 5% level after pre-registered multiple inference corrections.

Figure 2.8: Heterogeneous Average Treatment Effects on Utilization

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

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

Figure 2.9 shows the proportion of consumers in each stratification bin that use the line of credit over the course of the experiment. Only 4.5% of consumers with the lowest pre-experimental demand use the line of credit. If the average treatment effect for low-demand consumers in figure 2.8 is driven entirely by borrowers, it would imply an impossibly large LATE of 244pp.[18] It follows that the average treatment effect must be driven in part by consumers who do not borrow, at least among those with the lowest pre-experimental demand. Increased demand as a result of guaranteed access to credit, irrespective of credit use, is consistent with consumers engaging in precautionary saving as described by my model.

The final prediction in my model is that treatment effects are decreasing in the variance in α, the utility consumers realize from electricity access on a given day. Although I cannot observe utility from solar access, I do observe the number of watt hours (wH) used on each day a consumer's solar home system is switched on. I use wH consumed as a proxy for the

[18]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.

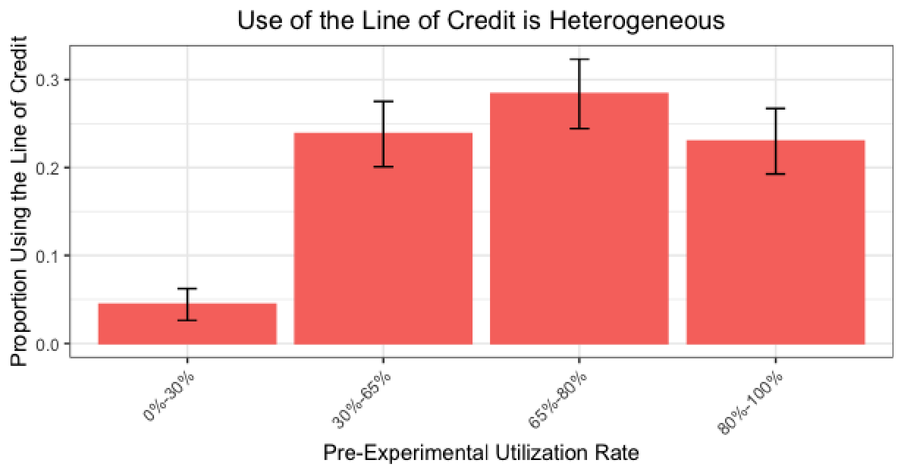

Note: 95% confidence intervals calculated using White robust standard errors. All estimates remain statistically significant after applying pre-registered multiple inference corrections.

Figure 2.9: Heterogeneous Average Treatment Effects on Utilization

utility gained from solar access. For each consumer, I calculate the standard deviation of wH used on days when they system is switched on in the 90 days prior to the experiment. I divide the distribution of standard deviation in use into quartiles. Letting k index quartiles of standard deviations in use, I estimate heterogeneous treatment effects using the specification

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

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

Figure 2.10 shows that consumers with the lowest variance in use are the only group with 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

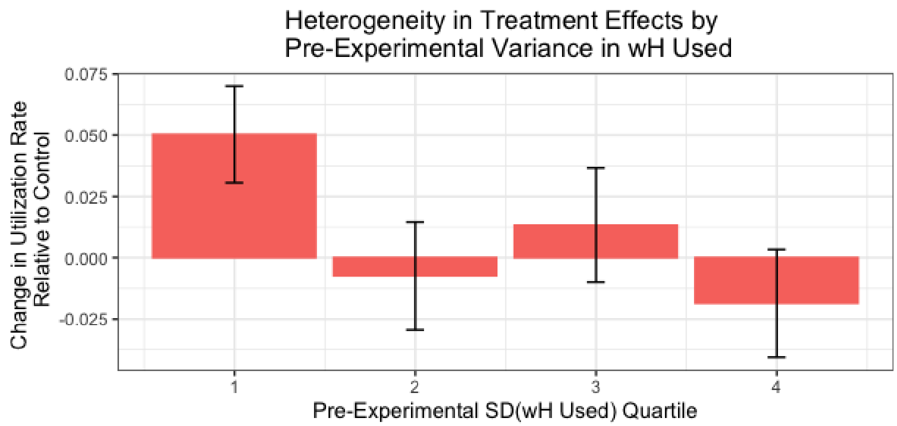

Note: 95% confidence intervals calculated using standard errors clustered at the level of the individual consumer. Estimates for the first quartile remain statistically significant after applying pre-registered multiple inference corrections.

Figure 2.10: Heterogeneous Impacts by Pre-Experimental Variance in Use

effect from being offered the line of credit. Consumers with the highest variance in use reduce utilization by around 2pp in response to being offered the line of credit, although the effect is only significant at the 10% level. Consumers with the least variance in use have a significantly higher treatment effect than those with the most variance in use. The estimated results are consistent with high demand, high variability consumers better targeting their solar consumption when provided with a way to reduce the burden of transaction costs, providing further evidence in support of my theoretical framework.[19]

While the pattern of average treatment effects is consistent with the mechanisms proposed in my model, I cannot directly observe transaction costs or liquidity constraints for all of the consumers in my sample throughout the course of the experiment. I provide descriptive evidence to further explore the role of transaction costs. Unlike the reduction in the precautionary savings motive, which occurs regardless of use, the line of credit only reduces transaction costs if consumers use it. I expect consumers who face the highest transaction costs to use the line of credit the most. While I do not have data on transaction costs for all consumers in my sample, I know the time to reach the nearest mobile money agent for most consumers in the treatment group. I use this data to run a descriptive regression to see whether consumers who face higher transaction costs are more likely to use the line of

[19]The pre-experimental variance in watt hours used is an imperfect proxy for α, the utility gained from solar access, because I only observe it on days when consumers have access to their systems. I provide a further robustness check for this effect by re-estimating it only for consumers in the top stratification bin, who have the highest pre-experimental demand. These consumers have access to their systems almost all of the time in the pre-experimental period. Appendix figure 2.19 shows that the same pattern holds among these consumers: those with the lowest pre-experimental variance in use exhibit significantly higher treatment effects than those with the highest pre-experimental variance in use.

credit. Table 2.3 shows that living one hour further from the nearest mobile money agent is associated with a 5.8pp, or nearly 50%, increase in the likelihood that a consumer uses the line of credit, providing support for the importance of transaction costs as a key mechanism underlying my estimated treatment effects.

Table 2.3: Heterogeneity in Take-Up by Distance to Mobile Money Agent

	Dependent variable:
	Take-up Rate
Hours to Reach MM Agent	0.058**
	(0.024)
90-Day Pre-experimental Utilization Rate	0.135***
	(0.034)
Daily Rate (RWF)	0.0003**
	(0.0001)
Hi Fee	−0.012
	(0.024)
Hi Borrowing Limit	−0.007
	(0.024)
Repayment Time Limit	−0.021
	(0.024)
Intercept	0.120***
	(0.045)
Observations	1,342
R^2	0.016
Adjusted R^2	0.011
Notes:	*p<0.1; **p<0.05; ***p<0.01
	Standard errors are White robust.

Taken together, my empirical results closely match the predictions of my theoretical framework and rule out a number of alternative mechanisms. In the next section, I consider an important alternative explanation for my estimated treatment effects: pre-experimental demand is correlated with present focus, or present bias (Laibson 1997, O'Donoghue and Rabin 1999). Of the range of possible alternative mechanisms, present focus merits special consideration because it changes the welfare implications of my results. If consumers with the highest pre-experimental demand are reducing solar purchases as the result of behavioral biases, it is no longer clear that offering the line of credit is welfare-enhancing for consumers. Methodologically, standard revealed preference measures of consumer welfare are inaccurate when consumers are present focused (Bernheim and Rangel 2009, Allcott and Taubinsky

2015). For both reasons, it is important to establish that my results are not being driven by present focus before I proceed to welfare estimation.

Evidence on the Importance of Present Focus

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

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

Both heterogeneity in pre-experimental demand and pre-experimental variance in use generate results that are consistent with my theoretical framework. If present focus is driving results along both dimensions of heterogeneity, then there should be a positive correlation between pre-experimental demand and pre-experimental variance in electricity use. Figure 2.11 shows that pre-experimental demand is not closely correlated with the standard deviation in wH used on days when a consumer has solar access.

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

Table 2.4 outlines the full cross-randomization. I hold constant the total number of days required to qualify for the incentives and the number of free days consumers could earn between the bulk discount and the monthly reward, leading to equivalent reductions in average price between the two incentives. Consumers should respond more to the monthly reward if they are present focused: the bulk incentive requires that consumers forgo more consumption today to gain consumption farther into the future relative to the monthly reward. I stratify the sample by pre-experimental demand over the entire duration of consumers' tenure with

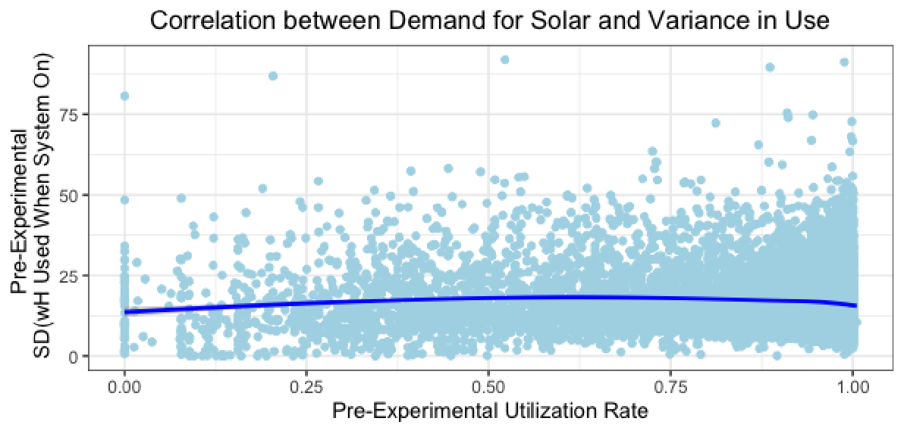

Figure 2.11: No Correlation Between Pre-experimental Demand and Variance in Use

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.

Table 2.4: Cross-Randomized Experimental Design for Solar Incentives

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

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

Figure 2.12 shows the effect of the incentives for solar payments on the number of days bought per month across the distribution of pre-experimental demand. Treatment effects are null or significantly negative for all consumers except those with the highest pre-experimental demand. Consumers with the highest pre-experimental demand respond to the incentives by increasing purchases by 2-2.5 days per month, a 6-8% increase relative to the control group. Critically, I cannot reject that bulk incentives and the monthly reward have the same effect for consumers with the highest pre-experimental demand. Consumer responses to the incentives indicate that consumers are likely not present focused in a manner that is strongly correlated with pre-experimental demand.

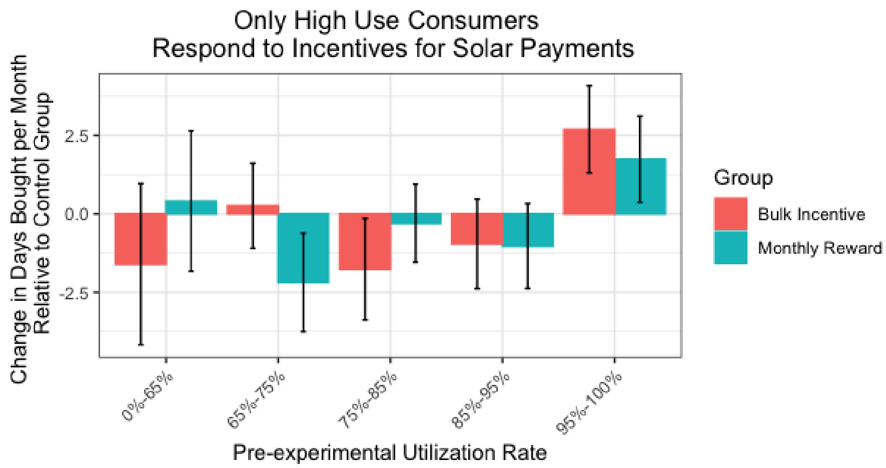

Note: I pool across minimum qualifying purchase sizes and reward sizes to increase power. 95% confidence intervals calculated with standard errors clustered at the level of the individual consumer.

Figure 2.12: Heterogeneous Impacts of Incentives for Solar Payments

I provide two final pieces of evidence on present focus among consumers in my sample. First, during the initial round of marketing calls for the line of credit, all consumers in the treatment group were given the option to select a voluntarily lower borrowing limit. Selecting a voluntarily lower borrowing limit is a commitment device. I expect that consumers with standard preferences will not choose to limit their choice set, and that consumers who are naive about their own present focus will not choose to limit their choice set. It follows that only consumers who are (partially) sophisticated and present focused will use the commitment device. Figure 2.13 shows that there are no statistically significant differences between stratification bins in the proportion of consumers opting to use the commitment device. If present focus is positively correlated with pre-experimental demand, it must be driven by a greater proportion of naive present-focused consumers with high pre-experimental demand.

After the initial round of marketing, representatives from the solar company called consumers again to remind them about the line of credit. During the second round of calls, they asked a random subset of 596 treated consumers how many times they expected to use the line of credit over the next month, the likelihood that the consumer would use the line of credit one more time than they expected, and the likelihood that the consumer would use the line of credit one fewer time than expected. I compute the actual number of times each consumer used the line of credit over the month following the phone call and calculate the difference between the consumer's expectation and their actual use of the line of credit. If consumers are present focused and (partially) naive, I expect them to underestimate their use of the line of credit.

Across all consumers surveyed, only 1% underestimate their use of the line of credit. Table 2.5 shows that consumers with the lowest pre-experimental demand overestimate their

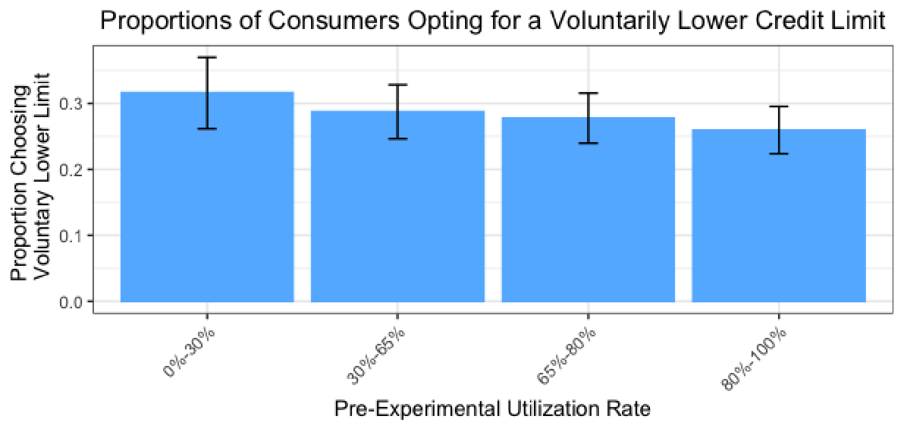

Figure 2.13: No Significant Differences in Use of the Commitment Device

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 2.5 shows variation between groups in consumers' confidence in their prediction. 86.5% of consumers with the lowest pre-experimental demand believe that there is a less than 10% chance that they will use the line of credit one fewer time than they predict. Consumers in other stratification bins are less certain, with 60.4%-71.7% believing that there is a less than 10% chance that they will use the line of credit one fewer time than predicted.

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

Discussion

Taken together, the evidence supports a model where consumers buying a perishable good with transaction costs respond to the line of credit differentially depending on the severity of liquidity constraints they face. Transaction costs and liquidity constraints are common market frictions in low-income countries. Although the strict non-storability of PAYGo solar

Table 2.5: Differences in Consumer Expectations and Realizations of Credit Use

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

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

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

access time is an extreme case, any perishable good with transaction costs will force consumers to make similar trade-offs. My results point to a range of concerns for policymakers seeking to promote more efficient and equitable markets.

Typically, market frictions negatively impact both firms and consumers. The experimental results show that the solar firm collects less revenue from certain groups of consumers as a result of market frictions while receiving more from others. When I re-weight treatment effects on the utilization rate to be representative of the distribution of consumers in figure 2.3, I find that offering the line of credit does not significantly increase revenue collection for the firm. Appendix figure 2.20 shows that the line of credit creates similarly heterogeneous impacts on repossession. Consumers with the lowest pre-experimental demand are less likely to default but offering the line of credit to consumers with the highest pre-experimental demand increases the risk of repossession, although estimates are not statistically significant. In a partial equilibrium sense, offering the line of credit is not profitable for the firm.

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

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

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

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

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

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

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

2.6 Welfare

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

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

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

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

and

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

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

$$\frac{dQ(F)}{dF} = \frac{dQ_p(F)}{dF} + \frac{dQ_b(F)}{dF}. \tag{2.13}$$

$$\frac{dP(F)}{dF} = \frac{\frac{dQ_P}{dF} + \frac{dQ_b}{dF}(1+F) + Q_b(F)}{Q_p(F) + Q_b(F)} - \frac{Q_p(F) + Q_b(F)(1+F)}{(\frac{dQ_p(F)}{dF} + \frac{dQ_b(F)}{dF})^2}. \tag{2.14}$$

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

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

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

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

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

Figure 2.14 shows the estimated slopes for each stratification bin. I cannot reject that demand is perfectly inelastic across the distribution of pre-experimental demand. To be conservative, I take the bottom of the 99% confidence interval around my estimates of the slope for each group in my estimation of consumer surplus.

I use the estimated demand curves along with observed demand to calculate a conservative lower bound on consumer surplus. I anchor the estimated demand curves at the total quantity demanded when the effective price is 1 for each stratification bin. To convert into monetary terms, I use the median daily rate, RWF 190 when the effective price is 1. I form a lower bound by only considering the bottom of the resulting Marshallian welfare triangle where I have empirical support for the price variation, as illustrated in figure 2.15.

Table 2.6 provides two lower bounds on consumer surplus. Column (1) shows a lower bound for consumer surplus at current prices, or a bound on the consumer surplus that solar

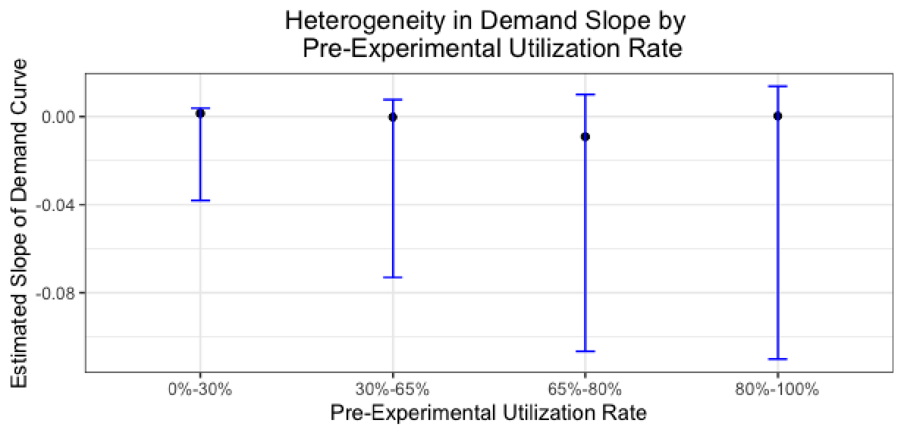

Figure 2.14: Estimated Slopes of the demand curve for solar

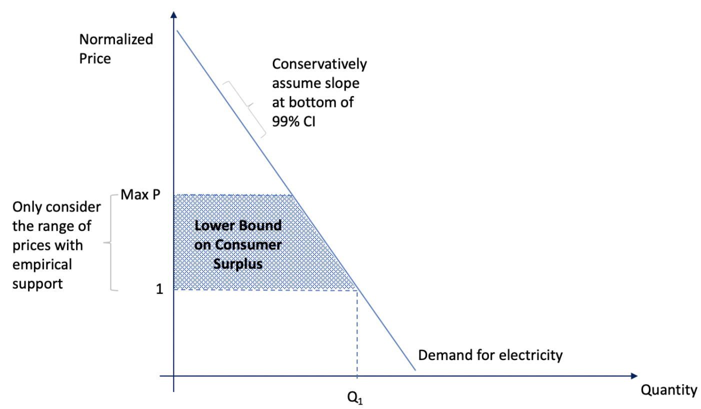

Figure 2.15: Estimated Slopes of the demand curve for solar

customers would enjoy if they had access to the line of credit and paid current daily rates for solar. Inframarginal consumers enjoy the most consumer surplus, at least $55 per household per year. As I move along the demand curve to increasingly marginal consumers, the lower bound on surplus drops under $7 per year.

Column (2) shows a lower bound on consumer surplus if solar access time were fully subsidized so that consumers paid a price of zero. The lower bound in column (2) allows me to assess which groups of consumers have a willingness to pay exceeding the cost of the solar home system. The total value of the median PAYGo contract in my setting is $230 paid over approximately 3.75 years, which includes the solar home system, basic appliances, maintenance, financing costs, and the labor associated with administering the PAYGo system. For the firm to break even, consumers need to have a willingness to pay of at least $61 per year. Column (2) in table 2.6 shows that 76% of current consumers have a lower

Table 2.6: Consumer Surplus from Solar for Different Levels of Utilization

Pre-Experimental Utilization Rate	Estimated Slope (99% CI Lower Bound)	Q_1 (weighted)	Max % Δ P	CS Lower Bound (hh/year), current prices (1)	CS Lower Bound (hh/year), fully subsidized (2)
0%-30%	-0.038	15,487	2.5%	$6.99	$13.84
30% - 65%	-0.073	56,473	5.1%	$24.76	$48.35
65% - 80%	-0.107	54,693	6.2%	$34.51	$67.03
80% - 100%	-0.110	709,445	4.5%	$55.33	$108.27
Weighted Mean			4.5%	$44.31	$86.67

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

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

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

The lower bounds in table 2.6 additionally facilitate comparisons to other recent measures of consumer surplus from electrification in the literature. In table 2.7, I take the weighted average of my estimated lower bound on consumer surplus and compare it to three recent estimates in the literature: Grimm et al (2020), Lee et al (2020), and Burgess et al (2020). My lower bound on consumer surplus is equal to or larger than estimates of total consumer surplus in all three papers, with the exception of the upper range of estimates in Lee et al (2020). If I extend the range of prices included in my demand curve or use my point estimates of the slope of the demand curve rather than the bottom of the 99% confidence interval, I obtain estimates for consumer surplus that are substantially higher than other estimates in the literature. My results suggest that consumer surplus from electricity is likely higher than previously believed because market frictions are distorting demand.

Table 2.7: Consumer Surplus from Electricity: Comparisons to the Literature

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

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

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

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

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

2.7 Conclusion

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

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

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

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

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

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

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

2.8 Appendix

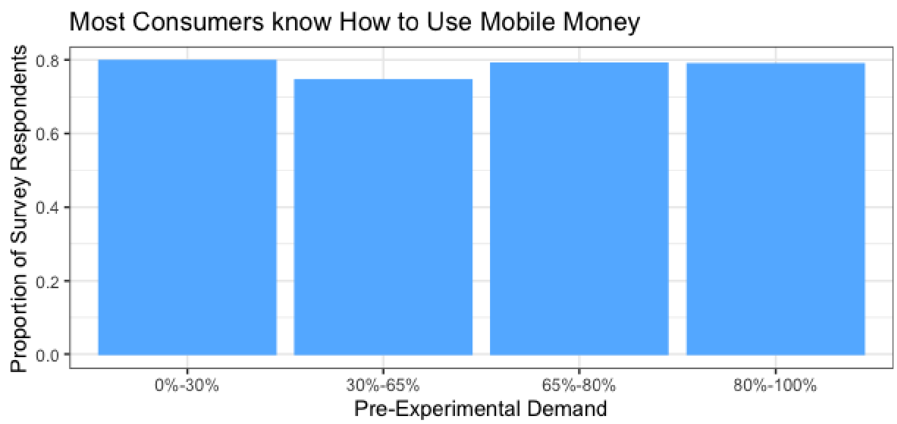

Note: The figure shows the proportion of consumers in each stratification bin who answer ”yes” to the question, “If you had mobile money already in your account and you wanted to use it to pay for solar, do you know how you would do that?”

Figure 2.16: Self-Reported Knowledge of Using Mobile Money to Buy Solar

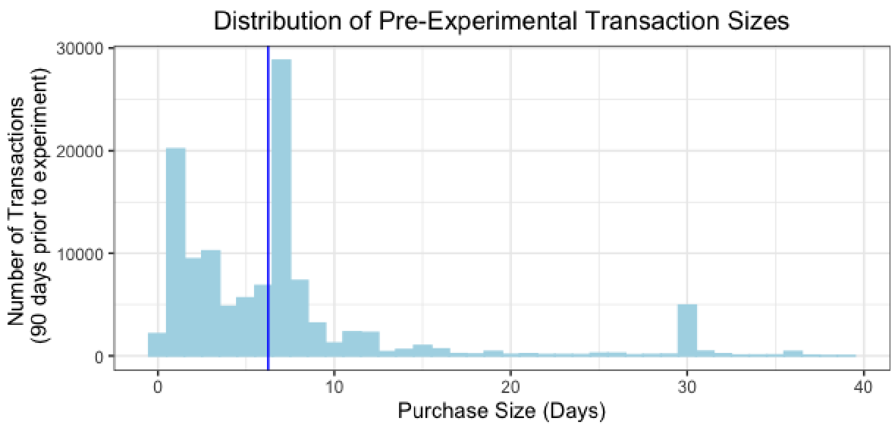

Figure 2.17: Distribution of Payment Sizes 90 Days Prior to the Experiment

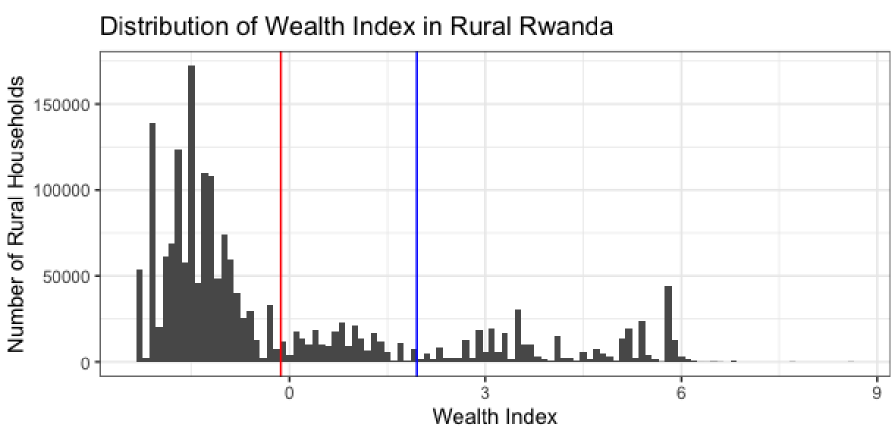

Note: The figure shows the distribution of wealth indices for a nationally representative sample of rural households in Rwanda. The red line indicates the mean level of wealth for the nationally representative sample and the blue line indicates the mean level of wealth for the consumers in my experimental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

Figure 2.18: Distribution of Wealth Among Rural Households in Rwanda

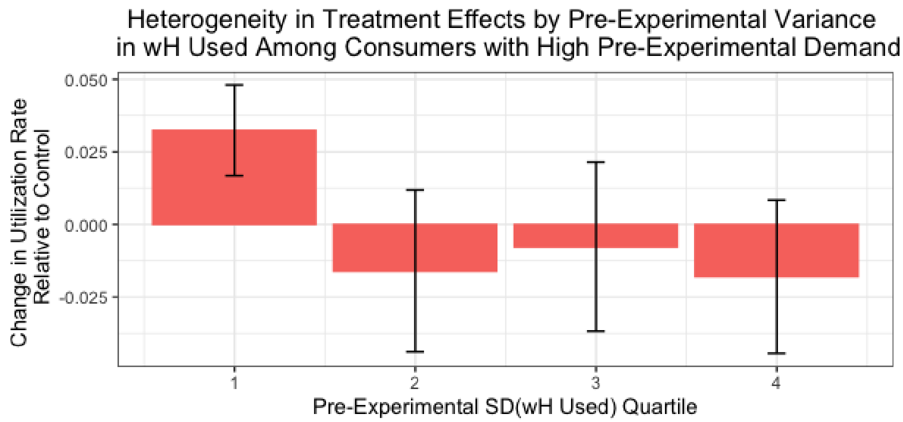

Note: The figure shows heterogeneous treatment effects based on pre-experimental variance in wH used only for consumers with the highest pre-experimental demand. The general that I find in the entire sample holds here, among consumers who have access to their solar home systems on most days in the pre-experimental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

Figure 2.19: Heterogeneous Treatment Effects: High Pre-Experimental Demand

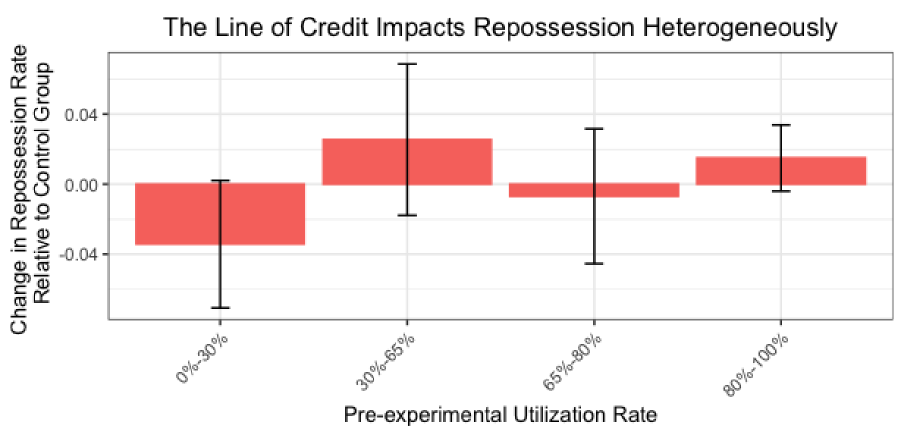

Note: I measure the repossession rate as the proportion of consumers eligible for repossession one month after the experiment ends. 95% confidence intervals calculated using White robust standard errors. Estimates are not statistically significant at the 5% level.

Figure 2.20: Heterogeneous Impacts on Default by Pre-Experimental Demand

Table 2.8: Balance Table

	Dependent variable:					
	Pre UR	Daily Rate	Mean wH	SD wH	Tenure	Pmt Size
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.0002	1.798	0.407	0.491	−2.194	0.284
	(0.002)	(2.292)	(0.665)	(0.417)	(5.371)	(0.368)
Bin 1	−0.901***	30.931***	−18.581***	1.149	38.890***	−8.767***
	(0.002)	(2.627)	(1.128)	(0.718)	(6.055)	(0.417)
Bin 2	−0.466***	18.141***	−9.840***	1.556***	29.601***	−3.225***
	(0.002)	(2.546)	(0.686)	(0.431)	(5.943)	(0.407)
Bin 3	−0.230***	13.721***	−6.631***	1.747***	12.471*	−3.732***
	(0.002)	(3.008)	(0.812)	(0.509)	(7.058)	(0.485)
Constant	0.957***	200.303***	53.705***	16.436***	451.822***	13.157***
	(0.001)	(0.962)	(0.254)	(0.159)	(2.256)	(0.154)
Observations	11,605	11,201	10,448	10,427	11,695	11,605
R^2	0.954	0.017	0.045	0.003	0.005	0.042
Adjusted R^2	0.954	0.017	0.044	0.003	0.005	0.042
Note:						*p<0.1; **p<0.05; ***p<0.01

Table 2.9: Differences in Wealth and Energy Spending between Stratification Bins

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

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

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

Chapter 3

Extensive Margin Responses to Intensive Margin Prices: the Case of Subsidies for Solar in Togo

3.1 Introduction

The Sustainable Development Goals aim to achieve universal access to clean, reliable electricity by 2030; however, recent academic work suggests that the returns to electrification may be small or negative for currently non-electrified populations due to limited demand (see Lee, Miguel, and Wolfram 2020, Burgess et al. 2020, Grimm, Lenz, et al. 2020, Burlig and Preonas 2016, Chaplin et al. 2017, Lenz, Anciet Munyehirwe, and Sievert 2017, Bayer et al. 2020, among others). Understanding the causes of low observed demand for electricity is critical to accurately evaluating the welfare impacts of electrification and formulating effective strategies for universal electrification. In this paper, I contribute to the literature on the demand for electricity in rural, low income settings by exploring the relationship between adoption of electricity and the price of electricity on the intensive, or use, margin.

Studying the relationship between the electrification decision and the intensive margin price of electricity speaks to the fundamental question of how large the expected net benefits of electrification would need to be for currently non-electrified households to adopt. It also provides further evidence on the welfare impacts of electrification for households who have already adopted, or would adopt at market prices. As I show in chapter two, some households face short-term liquidity constraints that can significantly reduce their demand for electricity at market prices, limiting the benefits of adoption. When prices are lower, such constraints are less likely to bind, leading to larger benefits for households that have already adopted and making it optimal for a greater share of households to adopt.

The relationship between adoption and the intensive margin price of electricity is difficult to study in low income settings because utilities often sign customer contracts with fixed prices or change prices for all customers uniformly. Even in settings where utilities change the intensive margin price of electricity, it is unclear how potential adopters form expectations over future prices. I leverage the phased rollout of a major subsidy for pay as you go (PAYGo) solar electricity in Togo to study the impact of changes in the intensive margin price of electricity on the adoption decision. The phased implementation of the subsidy allows for a clean identification of the impact of price reductions on the adoption of PAYGo solar. The subsidy arguably shifts consumer expectations in a more straightforward manner than a price change on the part of a utility because it is a change in price that consumers can expect to persist into the future. Finally, the magnitude of the subsidy ranges from 17.7% to 41.7%, allowing me to study adoption responses to price reductions that are larger than would naturally occur in the course of business for a typical utility.

Pay as you go solar is an ideal setting to study electricity adoption decisions because the upfront costs are small relative to alternative forms of electrification and systems are bundled with appliances. Each system also includes a battery to store electricity. While PAYGo solar home systems provide more limited access to electricity than a grid connection could provide, they represent an important first step onto the energy ladder for non-electrified households (Grimm, Lenz, et al. 2020). Combined, these features of the setting ameliorate concerns about credit constraints, lack of adoption because of limited access to appliances, and low demand stemming from poor supply-side reliability.[1]

Understanding the relationship between adoption and the intensive margin price is particularly relevant for PAYGo solar, and prepaid electricity more generally. Prepaid arrangements allow utilities to enforce contracts at a fraction of the cost of postpaid contracts (K. Jack and Smith 2015, K. Jack and Smith 2020). Pay as you go solar contracts additionally leverage the prepaid arrangement to allow consumers to pay for the durable goods included in the solar home system over time, providing a model for electrification of low-income, credit constrained households. A recent report estimates that PAYGo contracts could expand access to multi-light, or Tier 1, systems from 476 million to 670 million people (Lighting Global 2020). The PAYGo model has proved popular so far: it is estimated to be the fastest growing off-grid sector with sales increasing by 38% between 2019 and 2020 (Lighting Global 2020). However, projections of future growth often consider only how feasible it is for households to save for the down payment on a PAYGo system without accounting for the level of access time that households will be able to afford on an ongoing basis. Given that the structure of PAYGo contracts cause most costs to accrue in the future, it is critical to go beyond considerations of the upfront costs of adopting a PAYGo system and study consumer responses to intensive margin prices.

This paper broadly contributes to the literature on the impacts of rural electrification in low-income countries. One strand of this literature assesses the impacts of rural electrification by estimating reduced-form effects from electrification on a range of outcomes such as income, consumption, and labor force participation. Reduced-form impacts of electrification are

[1]Concerns about poor supply-side reliability may be particularly important, as Deutschmann, Postepka, and Sarr 2021 and Alberini, Steinbuks, and Timilsina 2020 have both documented that households are willing to pay significantly more for reliable electricity access.

varied. Dinkleman 2011 finds positive effects on female employment, Lipscomb, Mobarak, and Barham 2013 find that electrification improves economic development in Brazil, and Walle et al. 2017 and Khandker, Samad, et al. 2014 both document higher consumption from electrification in India. Conversely, Burlig and Preonas 2016, Bensch, Kluve, and Peters 2011, and Chaplin et al. 2017 find few positive impacts from electrification.

The other strand of the literature on rural electrification uses revealed preference measures to evaluate welfare impacts. Lee, Miguel, and Wolfram 2020 follow this approach by randomizing the cost of grid connections in Kenya, Grimm, Lenz, et al. 2020 trace out the demand curve using willingness to pay surveys in Rwanda, and Burgess et al. 2020 implement willingness to pay measures in India. Revealed preference approaches avoid some of the measurement issues associated with reduced-form impacts. By simply observing willingness to pay for electricity, researchers can evaluate the total value of electricity to consumers. The primary limitation is that market frictions that impact demand may distort revealed preference measures, leading to estimates of welfare that are biased downward (see chapter two for a more detailed discussion).

This paper adds to the evidence base that uses revealed preference measures to quantify the welfare impacts of electricity. To my knowledge, it is the first to use intensive margin price variation to estimate the demand curve for electricity on the extensive margin. It further adds a new geographic context to the literature on rural electrification by evaluating the impacts of a national subsidy in Togo.

Next, I provide a detailed description of the implementation of the subsidy in Togo. Section 3 describes the data I use, my empirical strategy, and validates the assumptions necessary for causal identification. In section 4, I present results in terms of the impacts of the subsidy and the implied elasticities. Section 5 uses my empirical results to estimate a demand curve and calculate consumer surplus from electrification. In section 6 I discuss the policy implications of my results and propose a range of explanations for discrepancies between my welfare estimates in Togo and others I have obtained using intensive margin demand in Rwanda. In section 7 I conclude by offering directions for future research.

3.2 Background

In 2017, 35% of Togo's population had access to electricity: 74% of urban households and 5% of rural households. The government started a new initiative called "CIZO" to increase rural electrification rates to 40% by 2022 with the goal of achieving universal access to electricity by 2030 (USAID 2017).

Pay as you go (PAYGo) solar home systems are a critical component of Togo's electrification strategy. PAYGo solar contracts are explicitly designed to lower financial barriers to electrification for low-income households. They do so by requiring down payments that are small relative to the cost of buying a solar home system outright or connecting to the grid, and by allowing for some flexibility in the amount and timing of ongoing payments. PAYGo contracts allow consumers to adopt a solar home system (SHS) consisting of solar panels, a battery for storing electricity generated by the solar panels, and a variety of high-efficiency appliances such as light bulbs, rechargeable radios, portable torches, phone chargers, or televisions. The consumer chooses which of these appliances to include in their SHS, with each additional appliance increasing the value of the overall contract. Once a consumer has selected their bundle, they make a down payment and have the system installed in their home.

After the solar home system has been installed, consumers "pay as they go." The solar company sets a daily rate, which is the price for one day of solar access time. The more appliances are included in the SHS, the higher the daily rate. In periods when consumers have purchased access time, they enjoy unlimited use of their solar home system. Consumers prepay for access time using mobile money. 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 an extended period, the solar company may repossess the SHS. Remote lockout and a credible threat of repossession render PAYGo solar contracts highly enforceable.

The government of Togo is collaborating with PAYGo providers to dramatically increase access to solar home systems. In March 2019, the government of Togo announced that it would start subsidizing solar in an effort to increase electrification rates, reduce the use of kerosene, and enhance economic development (Oteng 2019). The subsidy entitles each solar customer to 2,000 CFA per month; however, the subsidy is structured as a match on each qualifying purchase of a day of solar access time rather than being administered as a flat subsidy at the start of the month.

The structure of the subsidy is ideal for understanding the relationship between adoption of a solar home system and the expected intensive margin price of using the system, as it provides a constant discount on each intensive margin purchase. If the total amount that a consumer pays plus the subsidy is enough to buy solar access time for every day in a month, then the consumer receives the full 2,000 CFA subsidy. Otherwise, the consumer will receive a subsidy commensurate with the quantity of access time that they purchased in the month.

Figure 3.1 shows the distribution of daily rates. Around 95% of consumers select one of three daily rates: 160 CFA, 220 CFA, or 375 CFA. Given the popularity of these three system sizes, I limit my analysis to only those customers selecting one of these three daily rates. For the remainder of the paper, I refer to systems with a daily rate of 160 CFA as small, those with a daily rate of 220 CFA as medium, and those with a daily rate of 375 CFA as large. Given the 2,000 CFA per month subsidy, consumers with small systems receive a discount of 41.7%, those with medium systems get a discount of 30.3%, and consumers with large systems get a discount of 17.8%. Regardless of the system a consumer chooses, the subsidy represents a meaningful reduction in the cost of purchasing access time.

The government phased in the subsidy over the course of five months in 2019. The government rolled out the subsidy in an initial eleven prefectures starting on February 28, 2019. On April 30, 2019, the government added an additional thirteen prefectures to the subsidy program. The government added the remaining twelve prefectures on July 4, 2019. I leverage this phased rollout of the subsidy to empirically estimate consumer responses to

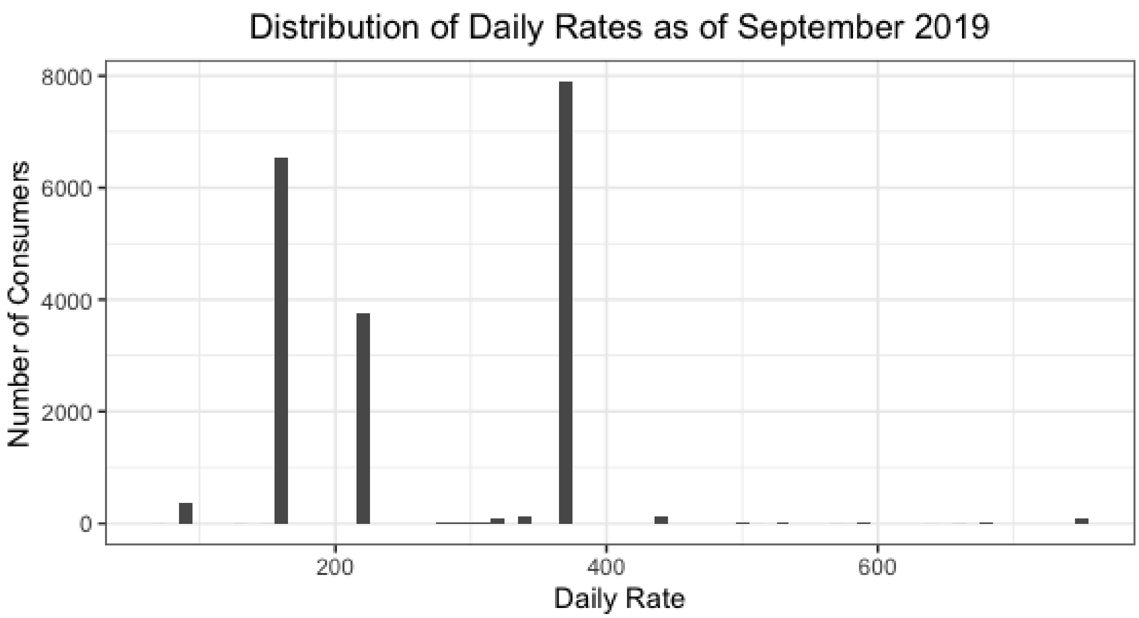

Figure 3.1: Distribution of Daily Rates

reductions in the intensive margin price of electricity.

3.3 Empirical Strategy

Data

I use administrative data from one of the largest solar companies in Togo to measure the effect of the subsidy on adoption of solar home systems. I observe the date each customer signed a PAYGo contract and had a SHS installed in their home, the size of the system that they chose and the associated daily rate that they pay for a day of access time, and the prefecture where the consumer lives. Using this information, it is straightforward to calculate the intensive margin price that the consumer expected to be facing at the time they adopted by combining information on the daily rate, the date of adoption, and the prefecture of residence.

Given that I have limited covariates on individual consumers, I aggregate the consumer-level data to create a prefecture by month dataset. The aggregated dataset contains the total number of new customers opting into each daily rate in each prefecture during every month in the dataset. I generate dummy variables that indicate when the subsidy came into effect in each prefecture as well as a dummy variable indicating when the subsidy was first announced nationally.

Taken together, the information in my dataset allows me to estimate increases in adoption of solar home systems as a result of the subsidy. In the next section, I outline three identification strategies to estimate the impact of the subsidy on adoption. For each approach, I describe the necessary assumptions, provide evidence to validate the assumptions, and discuss limitations.

Causal Identification

My first identification strategy is a simple event study. I estimate the impact of the subsidy using the following specification

$$IHS(New_{pt}) = \alpha + \beta_1 Sub_{pt} \times Sm_{pt} + \beta_2 Sub_{pt} \times Med_{pt} + \beta_3 Sub_{pt} \times Lg_{pt} + \beta_4 Med_{pt} + \beta_5 Lg_{pt} + \gamma_p + \delta M_t + \epsilon_{pt}. \quad (3.1)$$

Here, p indexes prefectures and t indexes months. The dependent variable is the inverse hyperbolic sine of the number of new customers, Sub_{pt} is a dummy variable indicating whether the subsidy had been rolled out in prefecture p in month t, Sm_{pt} is a dummy variable indicating whether the customer count is for customers adopting small systems, Med_{pt} is the analogous variable for medium systems, and Lg_{pt} is the analog for large systems. γ_p is a prefecture fixed effect and M_t is a linear time trend to control for overall growth in adoption of solar home systems over time.

Using this specification, β_1 is the percent change in adoption of small systems, β_2 is the percent change in adoption of medium systems, and β_3 is the percent change in adoption of large systems. I can easily translate these reduced form estimates into elasticities by dividing by the subsidy for each daily rate: $\frac{\beta_1}{41.7\%}$, $\frac{\beta_2}{30.3\%}$, and $\frac{\beta_3}{17.8\%}$.

The identifying assumption for the event study is that the growth in adoption of solar home systems would have followed a linear trend absent the introduction of the subsidy within each group of prefectures. Figure 3.2 shows the number of new customers in each group of prefectures over time, while figure 3.3 shows the cumulative number of customers in each group of prefectures over time. Informed by these figures, I only use a pre-period starting in July, 2018 when the growth in the number of customers starts to appear roughly linear for each group of prefectures.

My second identification strategy leverages the phased rollout of the subsidy to estimate a two-way fixed effects model. As shown in figures 3.2 and 3.3, prefectures in the first two phases of the subsidy rollout display parallel trends in the months prior to the introduction of the subsidy. Prefectures in the third phase of the rollout display slightly slower growth in the number of new customers. Figure 3.4 suggests one possible explanation: of the five regions in Togo, the coastal ("Maritime") region makes up a much higher proportion of prefectures in the third phase than in either of the first two phases. The capital city Lomé is in the Maritime region and urban locales have much higher electrification rates than rural areas, potentially explaining lower growth rates of PAYGo adoption.

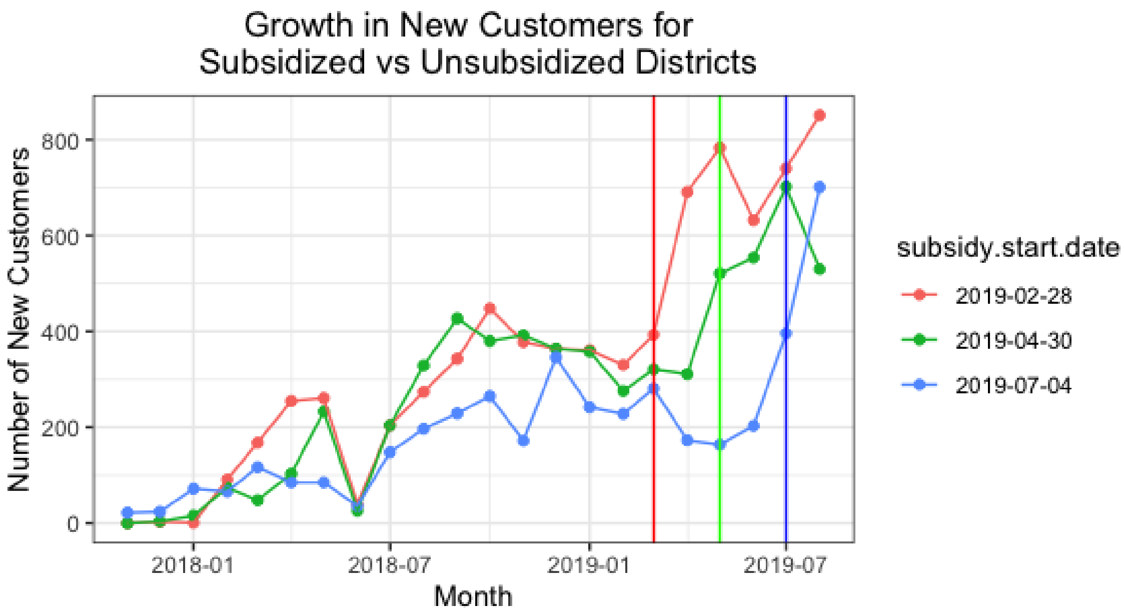

Figure 3.2: Counts of New Customers by Month

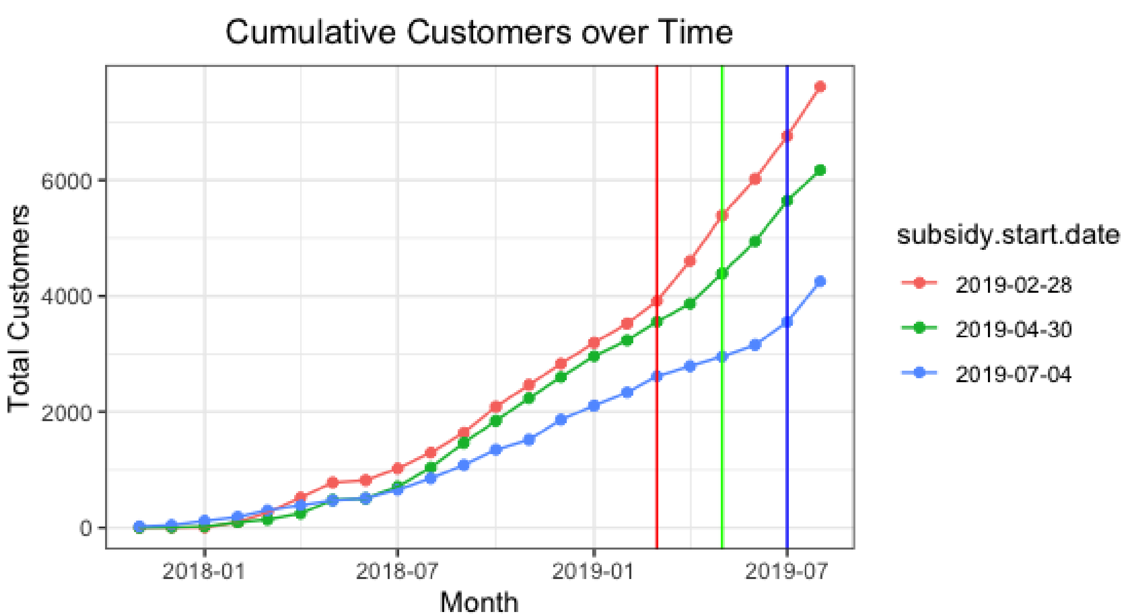

Figure 3.3: Cumulative Customers over Time

Given concerns about differences between the prefectures in the third wave of the subsidy

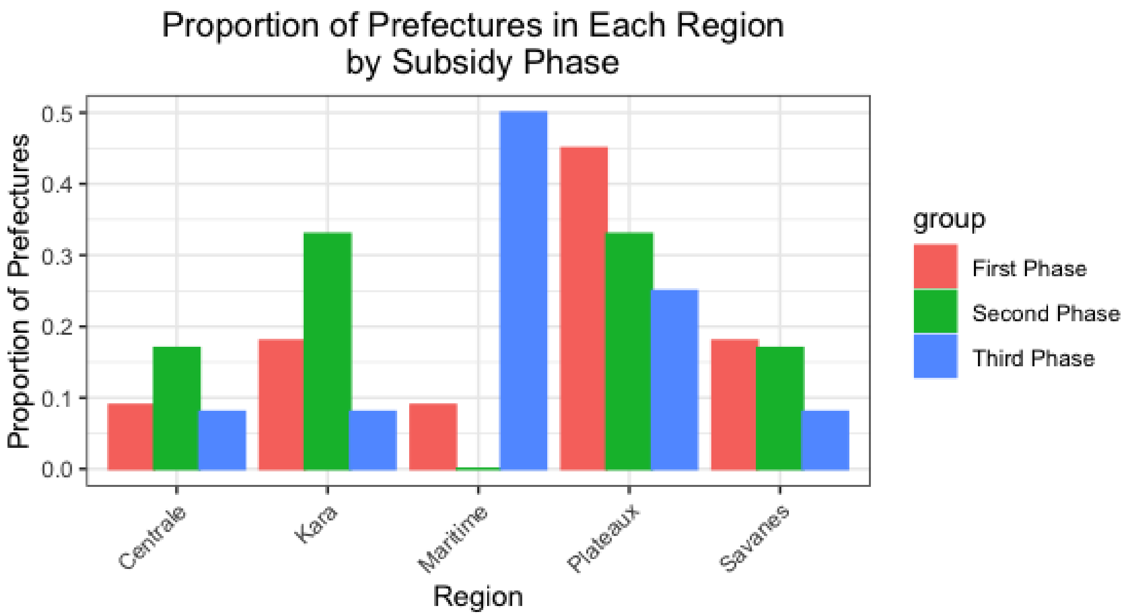

Figure 3.4: Prefectures by Region in each Phase of Subsidy

rollout versus those in the first two waves, I propose two different two-way fixed effects specifications. My first fixed effects specification is

$$IHS(New_{pt}) = \alpha + \sum_t \beta_t P1_p \times Month_t + \sum_t \delta_t P2_p \times Month_t + \gamma_p + \gamma_t + \epsilon_{pt}. \quad (3.2)$$

Here, $P1$ is a dummy variable equal to one if prefecture p is in the first phase of the subsidy rollout and $P2$ is a dummy variable equal to one if prefecture p is in the second phase of the subsidy rollout, leaving the third group of prefectures to act as the control group. I interact these phase dummies with month dummies to estimate impacts of the subsidy separately for each group of prefectures in each treated month. I average across all of these estimated effects to estimate a single increase in adoption attributable to the subsidy to avoid the weighting concerns with two-way fixed effects designs highlighted in Goodman-Bacon 2021. I estimate equation (3.2) separately for small, medium, and large systems to obtain estimates that are comparable to my other empirical specifications.

My final identification strategy takes seriously the concern that the prefectures in the third phase of the subsidy rollout does not follow close enough parallel trends to the first and second groups to act as a valid counterfactual. Instead, I drop the third group of prefectures and instead estimate a simple two-way fixed effects model using only the first two prefectures in the first two months of the subsidy rollout:

$$IHS(New_{pt}) = \alpha + \beta_1 Sub_{pt} \times Sm_{pt} + \beta_2 Sub_{pt} \times Med_{pt} + \beta_3 Sub_{pt} \times Lg_{pt} + \beta_4 Med_{pt} + \beta_5 Lg_{pt} + \gamma_p + \gamma_t + \epsilon_{pt}. \quad (3.3)$$

Anticipation Effects

The two-way fixed effects specifications allow for a weaker identifying assumption than the event study, as they only requires that all groups of prefectures follow parallel trends prior to the introduction of the subsidy and that there is no anticipation. However, there is reason to be concerned about anticipation effects for the second and third groups of prefectures. If consumers in prefectures that only received the subsidy in April or July are delaying adoption, then those prefectures are no longer valid counterfactuals for subsidized prefectures. Significant anticipation effects would lead me to overestimate the impact of the subsidy.

In figure 3.2, it does appear that rates of adoption may have fallen slightly for prefectures in the second and third phases after the initial rollout of the subsidy and before the subsidy was rolled out in those prefectures, but it is difficult to say for certain. In figures 3.5 and 3.6, I disaggregate to the daily level to better determine whether consumers appear to reduce adoption in anticipation of receiving the subsidy.

Figure 3.5 shows daily totals of new adopters in the second group of prefectures in event time, with day 0 being the day the subsidy became available for the second group of prefectures. The blue vertical line shows the date the first phase of the subsidy was announced. Figure 3.6 shows the analogous figure for the third group of prefectures. If consumers in either group of prefectures are reducing adoption in response to the initial phase of the subsidy, I should see a drop in the number of new adopters between the pre-period (left of the blue vertical line) and the period between the initial rollout and the time the subsidy was enacted in the second and third prefectures (the area between the blue and black vertical lines). I see no such reduction in figure 3.5 for consumers in the second phase of the subsidy rollout: trends in adoption remain relatively flat after the first phase of the subsidy and appear to continue at levels that are not statistically different from levels prior to the first phase of the subsidy. The same is true for consumers in the third phase in figure 3.6: adoption appears to follow a relatively constant, smooth trend even after the subsidy is first announced.

Smooth trends around the initial subsidy announcement could mask growth in the number of new adopters that would have happened had the subsidy not been announced. The second check I implement considers whether there is bunching around the time the subsidy is implemented in each group of prefectures. Intuitively, if consumers who would have adopted after the initial subsidy announcement but before the subsidy is implemented in their prefecture instead choose to wait, then I should see this pent up demand expressed as a discontinuous jump in demand immediately after the subsidy is implemented.

Turning again to figure 3.5, no such jump appears to occur for prefectures in the second phase. The overall level of adoption is higher, but it is not driven by a large increase in

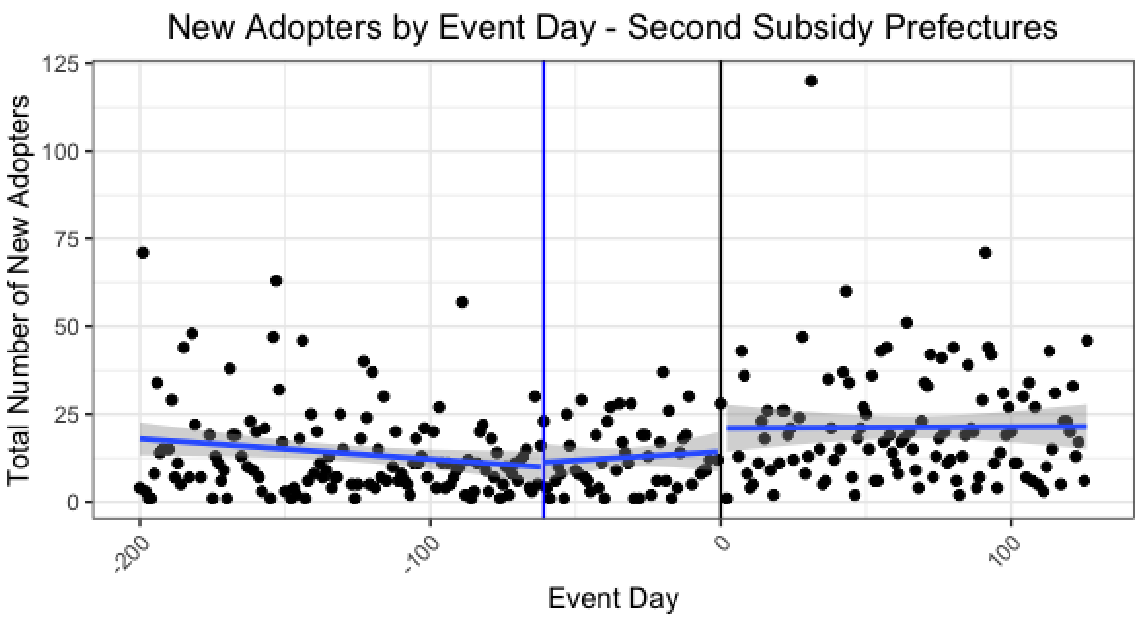

Note: The blue vertical line denotes the day (in event time) that the subsidy went into effect in the first group of prefectures, while the black vertical line denotes the day that the subsidy was available to prefectures in the second phase.

Figure 3.5: Daily Adoption in Second Phase Prefectures

new adopters in the days immediately following the introduction of the subsidy. Similarly, there is no discontinuous jump in the number of new adopters after the subsidy rollout for consumers in the third phase. Figure 3.6 shows that there is an increase in the rate of new adoption, but again the pattern is not one that indicates pent up demand.

Taken together, patterns in daily adoption totals suggest that there are not anticipation effects. While this may seem surprising, it could also speak to the nature of the subsidy. Even consumers who adopt prior to the subsidy going into effect are eligible for the discount once the subsidy is rolled out in their prefecture. If a consumer has already saved for the down payment, the cost of waiting is simply the discount from the subsidy multiplied by the number of months until the subsidy comes into effect, which may not outweigh the benefits of having a solar home system sooner.

My final check for anticipation effects is a simple event study. Rather than estimating effects on the subsidy rollout, I estimate effects on the subsidy announcement for prefectures in the second and third phase, including only the period before March 2019 through the time when each group actually received the subsidy. Significantly negative effects here would be indicative of anticipation effects.

Table 3.1 shows the results of my event study on anticipation effects for each system size, with the first column showing effects for prefectures in the second phase and the third column showing effects for prefectures in the third phase. I find no evidence of anticipation

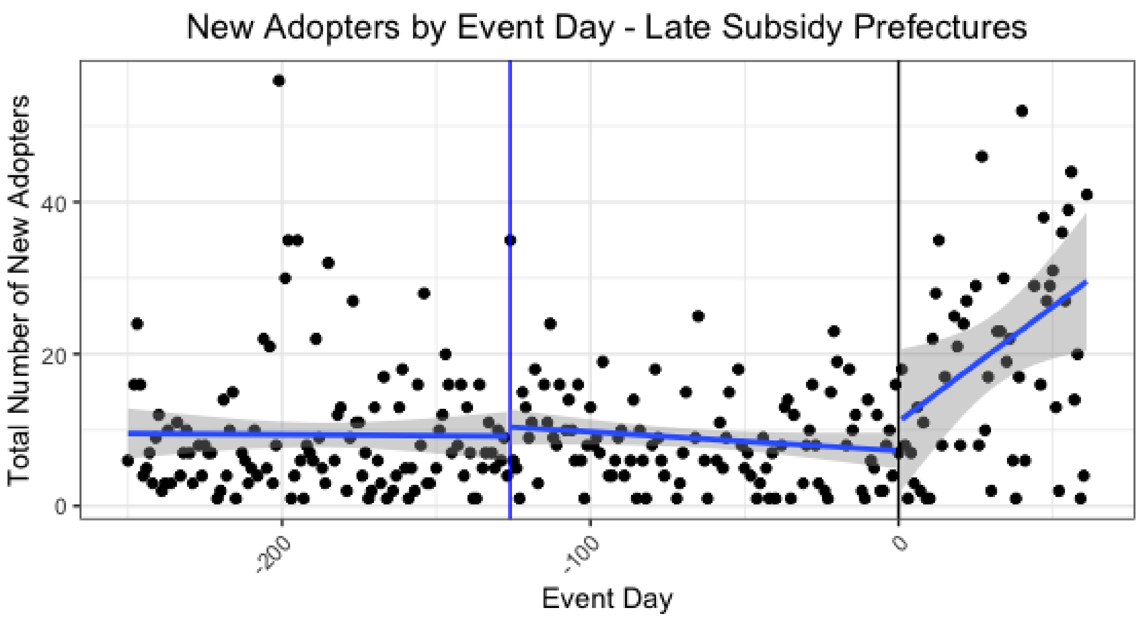

Note: The blue vertical line denotes the day (in event time) that the subsidy went into effect in the first group of prefectures, while the black vertical line denotes the day that the subsidy was available to prefectures in the second phase.

Figure 3.6: Daily Adoption in Third Phase Prefectures

effects for consumers opting for small systems. There is some evidence of anticipation effects for consumers choosing medium systems, although the effects are only significantly at the 10% level. Consumers choosing large systems in the second phase exhibit no anticipation effects but those in the third phase do.

Overall, this evidence suggests that the identifying assumptions for my two-way fixed effects design hold when examining the impact of the subsidy rollout on demand for small solar home systems. I cannot completely rule out anticipation effects for consumers opting for medium and large systems, suggesting that my estimated impacts will likely be an upper bound on the true impact of the subsidy for these consumers.

3.4 Results

I first present raw impacts of the subsidy under my three specifications. Figure 3.7 shows the percentage increase in adoption as a result of the subsidy using the naive event study, the two way fixed effects design with all prefectures, and the two way fixed effects design dropping the prefectures in the third phase. I cannot reject that the impacts are the same across the three specifications for all three system sizes, though at times the magnitudes are meaningfully different.

Table 3.1: Test for Anticipation Effects

	Dependent variable:	
	IHS(New Customers)	
	May Start	July Start
Medium	−0.055	−0.204
	(0.220)	(0.149)
Large	0.659***	1.085***
	(0.209)	(0.224)
Linear Time Trend	0.006	0.070**
	(0.055)	(0.028)
Announced x Small	0.152	−0.281
	(0.259)	(0.243)
Announced x Medium	−0.322*	−0.458*
	(0.181)	(0.253)
Announced x Large	−0.142	−0.426**
	(0.261)	(0.175)
Observations	390	432
R^2	0.584	0.480
Adjusted R^2	0.564	0.459
Note:	*p<0.1; **p<0.05; ***p<0.01	

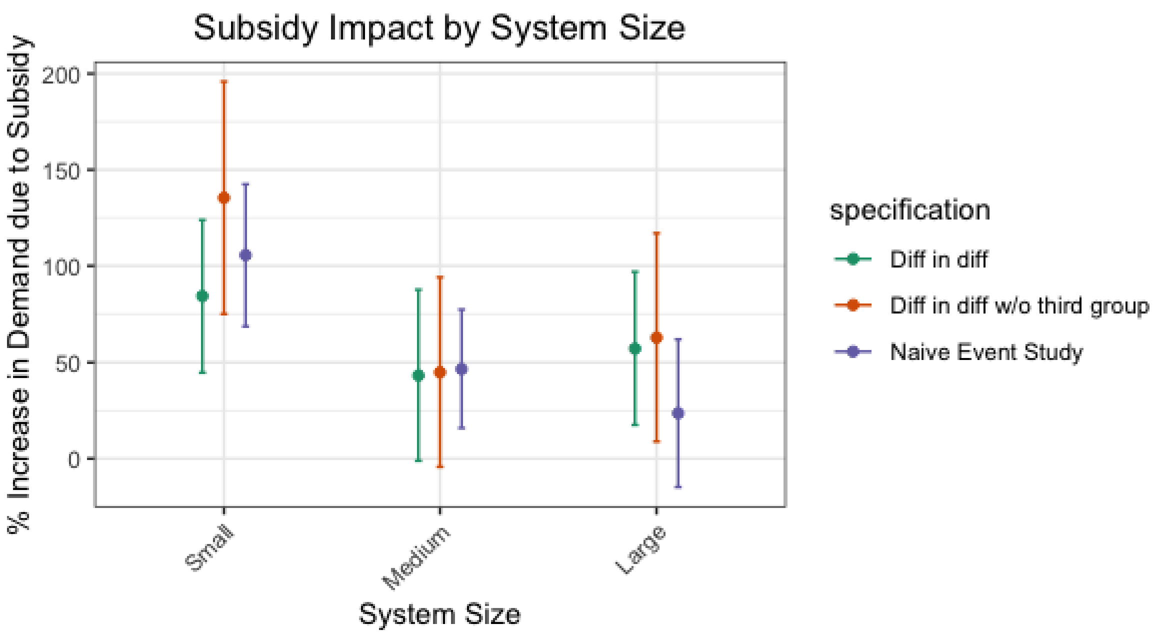

Figure 3.7: Raw Subsidy Impact by System Size

The impact of the subsidy are largest for the small systems in all three specifications, as I would expect given that the small systems receive the largest subsidy. For small systems using my three preferred specifications, my estimates range from an 84.3% increase to a 135.4% increase. Impacts are slightly smaller for medium systems, with estimates ranging from 43.2% to 46.6%. The impacts for large systems are similar to those for medium systems: 23.6% to 62.9%. Again, my estimates for medium and large systems are upper bounds given that I cannot rule out anticipation effects.

Figure 3.8 shows my raw estimates translated into elasticities of adoption with respect to the intensive margin price: $\frac{\%\Delta Adoption}{\%\Delta p}$. These elasticities are relatively high for small systems, allowing me to reject even unit elasticity for small systems using all three specifications. For small systems, the elasticity ranges from 2 to 3.25.

Estimated elasticities vary more for medium and large systems. Using my most conservative estimate yields an elasticity of 1.43 for medium systems and 1.33 for large systems, though both are estimated so imprecisely that I cannot rule out perfectly inelastic demand. These slightly less elastic estimates for medium and large systems relative to small systems make sense given that households who choose to adopt larger systems are powering more appliances with their systems, potentially causing them to rely on the systems more. Households with larger systems may also have higher incomes, allowing them to purchase access time more consistently.

A critical question is whether changes in adoption persist over time. The raw data in figure 3.2 suggests that this may be the case: the number of new customers each month

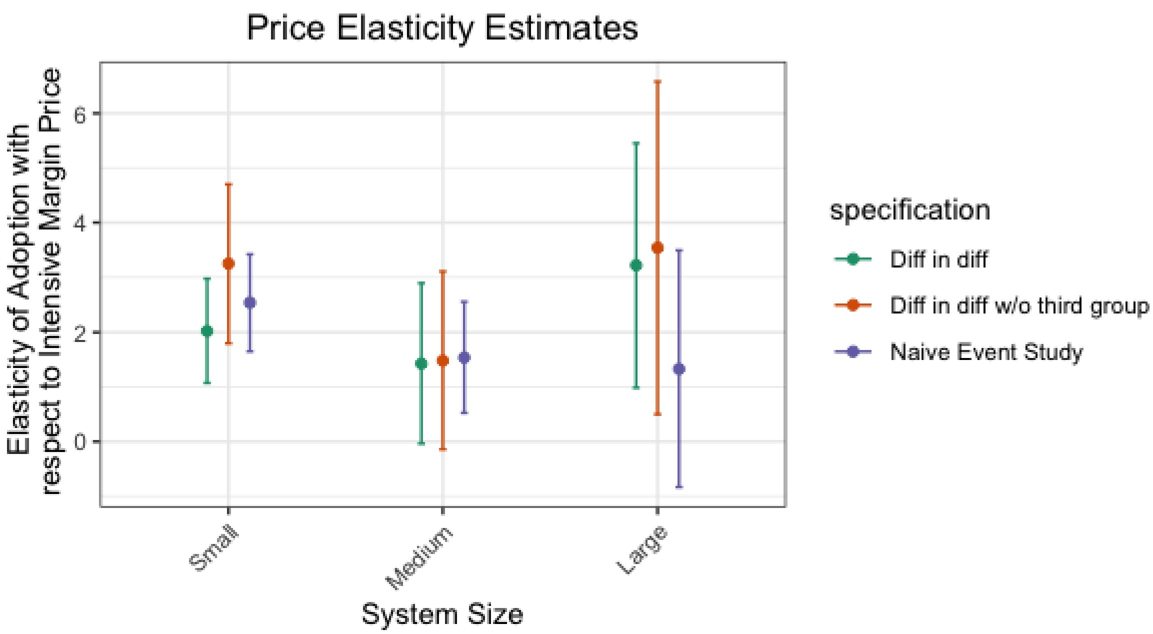

Figure 3.8: Implied Elasticities by System Size

dramatically increases after the initial rollout of the subsidy and appears to remain high during the subsequent five months in the data. To formally explore trends over time, I estimate the impacts of the subsidy month by month using my naive event study.

Figure 3.9 shows raw impacts of the subsidy by month. Examining the months prior to the introduction of the subsidy shows that trends appear to be relatively flat, consistent with the limited evidence of anticipation that I found previously. Second, the increase in adoption after the introduction of the subsidy persists for around four months. There is some variation in the magnitude of effects month to month after the subsidy begins, but I cannot reject that effects are equal within each size category in the first four months of the post period. Both specifications also suggest that effects may dissipate after five months, although estimates here are noisy and driven entirely by the first group of prefectures to receive the subsidy.

3.5 Welfare

I can use observed demand under the subsidy and my estimated elasticities for different system sizes to estimate consumer surplus from electricity overall as well as the additional consumer surplus generated by the subsidy. For each system size, I observe the total number of adopters as of September, 2019 when my data end. I take this as the quantity demanded under the subsidized price, Q_s. I then calculate the counterfactual adoption under the non-

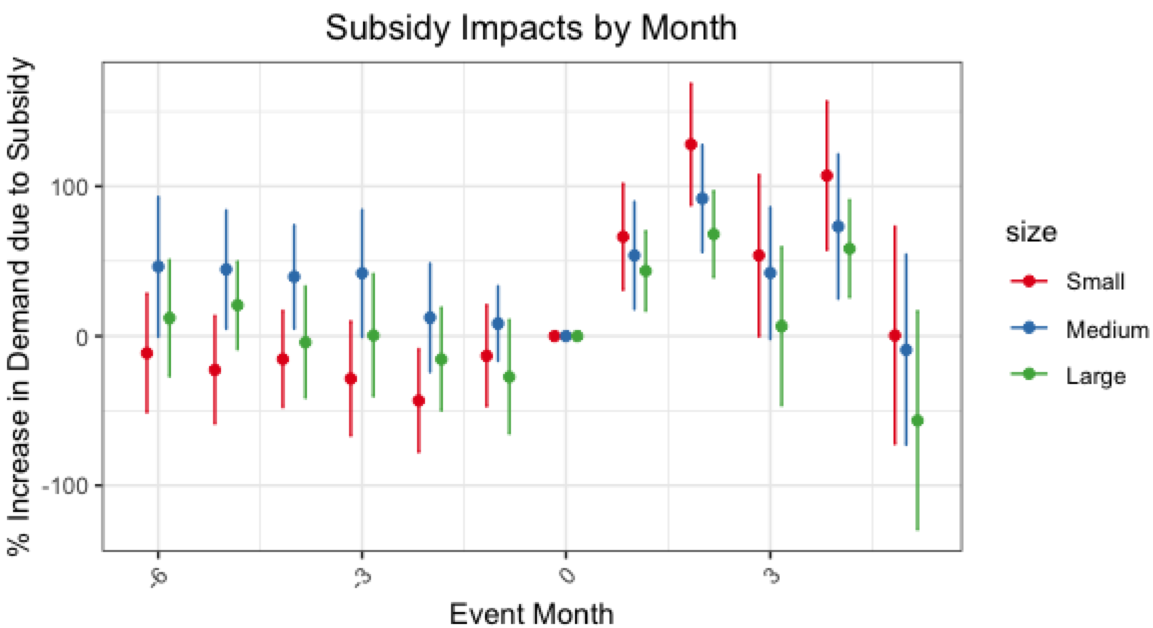

Figure 3.9: Monthly Subsidy Impacts

subsidized price, Q_n, by assuming that the demand curve is linear over the range of prices that I observe in the subsidy. Formally, I re-arrange the arc elasticity equation to solve for Q_n in terms of my estimated elasticity ϵ, the subsidized price P_s, and the non-subsidized price P_n:

$$Q_n = \frac{Q_s(1 - \epsilon\frac{P_s - P_n}{P_n + P_s})}{1 + \epsilon\frac{P_s - P_n}{P_n + P_s}}. \tag{3.4}$$

I further assume that the demand curve is linear outside the range of prices that I observe under the subsidy, which is likely a conservative assumption in terms of calculating welfare.

For ease of comparison with other studies in the literature, I compute welfare on a yearly basis. To do so, I assume that households purchase access time for 80% of days in the year and have monthly discount rates of around 0.3% for annual discount rates of 10%. For small systems, this suggests annual expenditures of 41,500 CFA at the market price and 24,200 CFA at the subsidized price. For medium systems, annual expenditures are 57,060 CFA at the market price and 39,770 CFA at the subsidized price and for large systems expenditures are 97,260 CFA at the market price and 80,000 CFA at the subsidized price.

I first consider the consumer surplus generated by the subsidy only for those consumers who would not have adopted in the absence of the subsidy. Column (1) in table 3.2 shows the total consumer surplus generated by the subsidy for each system size. The gains in consumer welfare are highest for the small systems at \$63,122, as those saw the greatest increase in adoption as a result of the subsidy. The gains for medium and large systems are slightly smaller due to smaller increases in adoption as a result of the subsidy, but are still substantial at $20,187 and $53,677, respectively. Adding up all of the numbers in column (1), the subsidy generated a total of $136,986 in consumer surplus for new adopters.

Table 3.2: Consumer Surplus from Solar

System Size	Total CS generated by subsidy for new adopters (1)	Total CS generated by subsidy for inframarginal adopters (2)	Total CS (hh/year) - average among market price adopters at P=0 (3)	Total CS (hh/year) - average among market price + subsidy adopters at P=0 (4)
Small	$63,122	$54,945	$81.66	$66.07
Medium	$20,187	$58,661	$125.37	$109.82
Large	$53,677	$116,186	$192.07	$176.52

The subsidy also generates benefits for consumers who would have adopted even in the absence of the subsidy. I do not have estimates of the magnitude of changes in demand on the intensive margin as a result of the subsidy, but my work in chapter 2 on Rwanda suggests that intensive demand may be fairly inelastic. If I simply assume constant intensive margin demand at 80% utilization, it indicates that the subsidy generated $230,614 in consumer surplus for these consumers, with each of these consumers saving around $31 per year.

Columns (3) and (4) of table 3.2 turn to estimating the total direct benefits of electrification. Column (3) estimates total benefits only for consumers who would adopt at the market price, meaning that I compute consumer surplus for consumers who would have adopted at the market price were the price to be set to zero. As expected, these benefits are highest for consumers with large systems and smallest for consumers with small systems.

Column (3) facilitates straightforward comparisons with the estimates in Grimm, Lenz, et al. 2020 and Lee, Miguel, and Wolfram 2020. For each of those studies, I use their estimated demand curves to calculate total consumer surplus for those consumers who would adopt at the market price. In Grimm, Lenz, et al. 2020, this yields total benefits of $89.50 per household per year, which is in line with my estimate for small systems and well below my estimate for medium and large systems. To compare with Lee, Miguel, and Wolfram 2020, I only consider consumer surplus above the market price. Lee, Miguel, and Wolfram 2020 provide a range of estimates from $23–$330 for on grid connections. My comparable estimates are slightly smaller, ranging from $7 to $23 per year, potentially because I am assuming linearity of the demand curve even at small quantities and potentially because the expected value of benefits from solar home systems are smaller in magnitude than the expected benefits of a grid connection. Given the differences in contract structures, technologies, and geographic locations, I view my welfare estimates as largely in line with other estimates of welfare from rural electrification that use extensive margin responses to trace out the demand curve for electricity.

3.6 Discussion

My reduced form results show that the intensive margin subsidy for PAYGo solar in Togo was highly effective at expanding the number of electrified households, at least during the period for which I have data. The strong consumer response to the intensive margin subsidy confirms that long-term affordability is a critical consideration for households choosing whether or not to opt into a PAYGo contract. From a more general policymaker's perspective, offering a subsidy to all consumers who adopt a solar home system is beneficial if the costs are less than the total positive externalities from adoption. While I cannot provide direct measures of the total positive externalities that would result from widespread adoption of solar home systems in Togo, I can take estimates from the literature to give a sense of the cost effectiveness of the subsidy.

Fetter and Phillips 2019 estimate that the yearly benefits of climate-affecting emissions for each household with a solar home system in East Africa is $13.70. If similar estimates hold in Togo, this would imply that the reduction in climate-affecting emissions alone is equal to approximately half of the subsidy cost. Grimm, Anicet Munyehirwe, et al. 2017 and Stojanovski et al. 2017 further document that access to solar significantly reduces the use of non-rechargeable batteries, pointing to positive externalities on environmental quality. Other positive externalities like reduced local pollution, shared benefits with non-electrified neighbors, and improved dissemination of information, among others, may easily lead to total externalities that exceed the cost of the subsidy. A key area for future work will be better identifying and measuring these types of positive externalities, particularly for technologies like solar home systems that allow low-income households to start working their way up the energy ladder.

While my estimates of the welfare impacts of electricity are broadly in line with other extensive margin estimates in the literature, they are somewhat at odds with intensive margin estimates that I obtain from work with PAYGo solar customers in Rwanda, documented in chapter two. In my work on demand for PAYGo solar on the intensive margin in Rwanda, I calculate a lower bound on total consumer surplus from solar (comparable to column (3) in table 3.2) that ranges from $14 to $108 per household per year, depending on how much access time households purchase. Taking a weighted average yields a conservative lower bound of $86.67 per household per year; however, I find that demand is inelastic on the intensive margin, which would suggest that total welfare is substantially higher.

My intensive margin estimate of consumer surplus may imply higher consumer surplus from electricity than my extensive margin estimate in Togo and other extensive margin estimates in the literature for multiple reasons. If credit constraints limit adoption, then extensive margin demand will appear to be more price elastic than it truly is. While credit constraints likely explain some of the differences between my intensive margin estimates and those put forward by Lee, Miguel, and Wolfram 2020 and Grimm, Lenz, et al. 2020, the down payments required to adopt a PAYGo system in Togo are a small fraction of the upfront costs in those papers. As such, I would expect credit constraints to play a much smaller distortionary role when considering my results from Togo. Beyond credit constraints, it is unclear how consumers evaluate the net benefits of adoption for a relatively new technology. Incomplete or inaccurate information could lead consumers to form incorrect expectations about the net benefits of adopting a PAYGo system. If so, traditional revealed preference measures based on observed demand would provide inaccurate estimates of the welfare impacts from electrification.

Conversely, my intensive margin estimates may underestimate price elasticity under a variety of conditions. If consumers face frictions in switching between energy sources once they adopt a PAYGo system, then inelastic demand on the intensive margin may simply reflect these switching costs. From a behavioral perspective, studies in the U.S. have found that consumers may be inattentive to the price of electricity on the intensive margin (Ito 2014). Inattention seems less likely in the PAYGo setting because the price of access time is not variable. Electricity expenditures also form a larger share of household budgets for low-income, rural consumers. Beyond inattention, the sunk cost fallacy has been demonstrated across a range of firm-level and individual-level decisions (see Roth, Robbert, and Straus 2015 for an overview). Consumers engaged in the sunk cost fallacy may demonstrate inelastic demand on the intensive margin if they are trying to get the most out of their PAYGo systems after incurring the initial cost of adoption. While I cannot measure the importance of the sunk cost fallacy for the consumers in my intensive margin work, Roth, Robbert, and Straus 2015 find evidence from multiple studies showing that the impacts of the sunk cost fallacy decline over time. Given that PAYGo consumers in my intensive margin sample all have at least three months of experience with their systems, sunk costs may not play a major role.

The preceding discussion demonstrates that welfare from consumer durables may not be accurately captured using either intensive or extensive margin demand. On a positive note, many of the factors that may distort intensive margin demand point to underestimating elasticities while non-price factors shaping extensive margin demand are more likely to overestimate elasticities. In some cases, this could allow researchers to bound elasticities in ranges that are usefully narrow. Disentangling true consumer welfare from such a complex decision environment is non-trivial, and requires a deeper understanding of consumer behavior than simple purchasing data can provide.

3.7 Conclusion

I show that adoption of PAYGo solar home systems is strongly responsive to reductions in the intensive margin price of access time using the phased implementation of a solar subsidy in Togo. Increases in adoption appear to be persistent and generate substantial welfare gains both for new adopters and those consumers who would have adopted at market prices. Estimates of the positive externalities associated with solar home systems from the literature further suggest that the subsidy is likely cost effective in the sense that the per-household cost is lower than the potential positive externalities generated by increased adoption of solar home systems.

My results add evidence to the literature on rural households' willingness to pay for electricity and the potential welfare impacts of rural electrification. I offer the first estimate of the extensive margin response to price changes on the intensive margin. In doing so, I confirm the importance of long-term affordability in the electrification decision for low-income households even with the low upfront adoption costs of PAYGo systems.

I ultimately conclude that the precise magnitude of the welfare effects from rural electrification are likely not well captured using revealed preference measures due to a range of non-price factors that shape observed demand on both the intensive and the extensive margin. However, the differences between my results and others in the literature highlights the need for further study on the intensive margin of energy demand in rural, low-income settings. How does consumer demand on the intensive margin compare to consumer expectations about demand when they are making the adoption decision? Does inelastic demand on the intensive margin reflect high value to consumers, or is it driven by behavioral factors that would not conventionally be considered welfare-enhancing?

Apart from these questions about the factors driving intensive versus extensive margin demand for electricity, what are socially optimal levels of electrification and intensive margin consumption of solar electricity? Intensive margin subsidies like the one in Togo are canonical tools for increasing demand for goods with positive externalities to socially optimal levels, but it is difficult to assess the correct level for such subsidies without a better understanding of both the private and external benefits of solar adoption and consumption.

Bibliography

[1] Anna Alberini, Jevgenijs Steinbuks, and Govinda Timilsina. "How Valuable is the Reliability of Residential Electricity Supply in Low-Income Countries? Evidence from Nepal". 2020.

[2] Hunt Allcott and Dmitry Taubinsky. "Evaluating behaviorally motivated policy: experimental evidence from the lightbulb market". In: *American Economic Review* 105 (2015), pp. 2501–38.

[3] Nava Ashraf, Dean Karlan, and Wesley Yin. "Deposit Collectors". In: *The B.E. Journal of Economic Analysis and Policy* 5.2 (2006), pp. 1–24.

[4] Diego Aycinena, Claudia Martinez, and Dean Yang. "The Impact of Transaction Fees on Migrant Remittances: Evidence from a Field Experiment Among Migrants from El Salvador". 2010.

[5] Patrick Bayer et al. "The need for impact evaluation in electricity access research". In: *Energy Policy* 137 (2020), p. 111099.

[6] Thornsten Beck, Demirguc-Kunt Asli, Maria Soledad, and Peria Martinez. "Banking Services for Everyone? Barriers to Bank Access and Use around the World". In: *World Bank Research Observer* 22.3 (2008), pp. 397–430.

[7] Thornsten Beck, Demirguc-Kunt Asli, Maria Soledad, Peria Martinez, and Maria Soledad. "Reaching Out: Access to and use of banking services across countries". In: *Journal of Financial Economics* 85.1 (2007), pp. 234–266.

[8] Gunther Bensch, Jochen Kluve, and Jorg Peters. "Impacts of Rural Electrification in Rwanda". In: *Journal of Development Effectiveness* 3.4 (2011), pp. 567–588.

[9] Douglas Bernheim and Antonio Rangel. "Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics". In: *Quarterly Journal of Economics* 124 (2009), pp. 51–104.

[10] Robin Burgess et al. "Demand for Electricity on the Global Electrification Frontier". 2020.

[11] Fiona Burlig and Louis Preonas. "Out of the Darkness and into the Light? Development Effects of Rural Electrification". EI @ Haas WP 268. 2016.

[12] Duncan Chaplin et al. *Grid Electricity Expansion in Tanzania by MCC: Findings from a Rigorous Impact Evaluation*. Tech. rep. Mathematica Policy Research, 2017.

[13] Daryl Collins et al. *Portfolios of the Poor: How the World's Poor Live on $2 a Day*. Princeton University Press, 2009.

[14] Angus Deaton. "Savings and Liquidity Constraints". In: *Econometrica* 59.5 (1991), pp. 1221–1248.

[15] Tatyana Deryugina, Alexander MacKay, and Julian Reif. "The Long-Run Dynamics of Electricity Demand: Evidence from Municipal Aggregation". In: *Forthcoming, American Economic Journal: Applied Economics* (2019).

[16] Joshua Deutschmann, Agnieszka Postepka, and Leopold Sarr. "Measuring Willingness to Pay for Reliable Electricity: Evidence from Senegal". In: *World Development* 138 (2021), p. 105209.

[17] Taryn Dinkleman. "The Effects of Rural Electrification on Employment: New Evidence from South Africa". In: *American Economic Review* 101.7 (2011), pp. 3078–3108.

[18] Pascaline Dupas et al. "Banking the Unbanked? Evidence from Three Countries". In: *American Economic Journal: Applied Economics* 10.2 (2018), pp. 257–297.

[19] Rob Fetter and Jonathan Phillips. "The True Cost of Solar Tariffs in East Africa". 2019.

[20] Eilin Francis, Joshua Blumenstock, and Jonathan Robinson. "Digital Credit: A Snapshot of the Current Landscape and Open Research Questions". 2017.

[21] Andrew Goodman-Bacon. "Difference-in-differences with Variation in Treatment Timing". In: *Journal of Econometrics* (2021).

[22] Michael Grimm, Luciane Lenz, et al. "Demand for Off-Grid Solar Electricity: Experimental Evidence from Rwanda". In: *Journal of the Association of Environmental and Resource Economists* 7 (2020), pp. 417–454.

[23] Michael Grimm, Anicet Munyehirwe, et al. "A first step up the energy ladder? Los cost solar kits and households' welfare in rural Rwanda". In: *World Bank Economic Review* 31 (2017), pp. 631–649.

[24] Beth Hayes. "Competition and Two-Part Tariffs". In: *The Journal Of Business* 60 (1987), pp. 41–54.

[25] Koichiro Ito. "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing". In: *American Economic Review* 104.2 (2014), pp. 537–563.

[26] Kelsey Jack and Grant Smith. "Charging ahead: Prepaid electricity metering in South Africa". In: *American Economic Journal: Applied Economics* 12.2 (2020), pp. 134–168.

[27] Kelsey Jack and Grant Smith. "Pay as you go: Pre-paid metering and electricity expenditures in South Africa". In: *American Economic Review Papers and Proceedings* 105.5 (2015), pp. 237–241.

[28] William Jack and Tavneet Suri. "Risk Sharing and Transactions Costs: Evidence from Kenya's Mobile Money Revolution". In: *American Economic Review* 104.1 (2014), pp. 183–223.

[29] *Kenya - Data at a Glance*. URL: `http://finclusion.org/country/africa/kenya.html#dataAtAGlance`.

[30] Shahidur Khandker, Douglas Barnes, and Hussain Samad. "Welfare Impacts of Rural Electrification: A Case Study from Bangladesh". WPS 4859. 2009.

[31] Shahidur Khandker, Hussain Samad, et al. "Who Benefits Most from Rural Electrification? Evidence in India". In: *The Energy Journal* 35.2 (2014), pp. 76–96.

[32] David Laibson. "Golden Eggs and Hyperbolic Discounting". In: *Quarterly Journal of Economics* 112 (1997), pp. 443–478.

[33] Gregory Lane. "Credit Lines as Insurance: Evidence from Bangaldesh". 2020.

[34] Kenneth Lee, Edward Miguel, and Catherine Wolfram. "Experimental Evidence on the Economics of Rural Electrification". In: *Journal of Political Economy* 128.4 (2020), pp. 1523–1565.

[35] Luciane Lenz, Anciet Munyehirwe, and Maximiliane Sievert. "Does Large-Scale Infrastructure Investment Alleviate Poverty? Impacts of Rwanda's Electricity Access Roll-Out Program". In: *World Development* 89 (2017), pp. 88–110.

[36] ESMAP Lighting Global GOGLA. *Off-Grid Solar Market Trends Report 2020*. Tech. rep. International Finance Corporation, 2020.

[37] Molly Lipscomb, Mushfiq Mobarak, and Tania Barham. "Development Effect of Electrification: Evidence from the Topographic Placement of Hydropower in Brazil". In: *American Economic Journal: Applied Economics* 5.2 (2013), pp. 200–231.

[38] Ted O'Donoghue and Matthew Rabin. "Doing it Now or Later". In: *American Economic Review* 89 (1999), pp. 103–124.

[39] Eric Oteng. "Togo subsidises off-grid solar panels to ease access to electricity". In: *Africa News* (Mar. 3, 2019). URL: `https://www.africanews.com/2019/03/03/togo-subsidises-off-grid-solar-panels-to-ease-access-to-electricity/`.

[40] Stefan Roth, Thomas Robbert, and Lennart Straus. "On the sunk-cost effect in economic decision-making: a meta-analytic review". In: *Business Research* 8 (2015), pp. 99–138.

[41] *Rwanda - Data at a Glance*. URL: `http://finclusion.org/country/africa/rwanda.html#dataAtAGlance`.

[42] Ognen Stojanovski et al. "Rural energy access through solar home systems: Use patterns and opportunities for improvement". In: *Energy for Sustainable Development* 37 (2017), pp. 33–50.

[43] Tavneet Suri, William Jack, and Thomas Stoker. "Documenting the birth of a financial economy". In: *Proceedings of the National Academy of Sciences* 109.26 (2012), pp. 10257–10262.

[44] *The Little Data Book on Financial Inclusion 2018*. Tech. rep. International Bank for Reconstruction and Development, 2018.

[45] USAID. *Togo Power Africa Fact Sheet*. Tech. rep. United States Agency for International Development, 2017.

[46] Dominique van de Walle et al. "Long-Term Gains from Rural Electrification in Rural India". In: *World Bank Economic Review* 31.2 (2017), pp. 385–411.

[47] Julie Zollman et al. *Escaping Darkness: Understanding Consumer Value in PAYGo Solar*. Tech. rep. CGAP Access to Finance Forum, 2017.