5-2024

Design of off-grid lighting business models to serve the poor: Field experiments and structural analysis

Bhavani Shanker UPPARI
Singapore Management University, bhavaniu@smu.edu.sg

Serguei NETESSINE
University of Pennsylvania

Ioana POPESCU
INSEAD

Rowan P. CLARKE
Harvard University

Citation

UPPARI, Bhavani Shanker; NETESSINE, Serguei; POPESCU, Ioana; and CLARKE, Rowan P.. Design of off-grid lighting business models to serve the poor: Field experiments and structural analysis. (2024). *Management Science*. 70, (5), 3038-3058.
Available at: https://ink.library.smu.edu.sg/lkcsb_research/7142

Design of Off-Grid Lighting Business Models to Serve the Poor: Field Experiments and Structural Analysis

Bhavani Shanker Uppari
Singapore Management University, 50 Stamford Road, 178899 Singapore, bhavaniu@smu.edu.sg

Serguei Netessine
The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104, netessin@wharton.upenn.edu

Ioana Popescu
INSEAD, 1 Ayer Rajah Avenue, 138676 Singapore, ioana.popescu@insead.edu

Rowan P. Clarke
Harvard Business School, Harvard University, Boston, Massachusetts 02163, rclarke@hbs.edu

A significant proportion of the world's population has no access to grid-based electricity and so relies on off-grid lighting solutions. Rechargeable lamp technology is gaining popularity as an alternative off-grid lighting model in developing countries. In this paper, we explore consumer behavior and the operational inefficiencies that result under this model. Specifically, we are interested in (i) measuring the impact of inconvenience (of travelling to recharge the lamp) along with the impact of liquidity constraints (due to poverty) on lamp usage, and (ii) evaluating the efficacy of strategies that address these factors. We build a structural model of consumers' recharge decisions that incorporates several operational features of the low-income regions. We conducted large-scale field experiments in Rwanda in partnership with a local rechargeable lamp operator and use the resultant data to estimate and test our model.

We find that the complete removal of inconvenience and liquidity constraints from the current business model results in 73% and 126% increases in both recharges and revenue, thereby suggesting that these constraints are major sources of inefficiency. By implementing simple *operations-based strategies* – such as starting more recharge centers, visiting consumers periodically to collect their lamps for recharge, and allowing consumers to partially recharge their lamps and pay flexibly for the recharge – more than half the benefit of completely eliminating the inefficiencies can be attained. By contrast, the *price- and capacity-based strategies* that vary the economic variables (i.e., the amount paid per recharge and the amount of light obtained in return) but not the operational model perform far worse than the aforementioned strategies. Overall, our analysis emphasizes the importance of managing operations effectively even in markets with cash-constrained consumers, where firms may have a natural tendency to focus more on reducing prices.

1. Introduction

One tenth of humankind still has no access to electricity. More than 80% of this population inhabits countries in Africa, a continent that is half unelectrified (IEA 2020). Not surprisingly, countries with low electrification rates are those in which most citizens live on less than \$2 (US) per day (often referred to as the *bottom of the pyramid*, or BoP for short). Grid-based models of electricity supply have been unsuccessful in these countries because they require substantial capital investment, and in many cases it may neither be technically feasible nor economical to extend grid electricity to these regions. Hence, there is a huge market for off-grid energy in these countries. For example, Rwanda's current national electrification rate is 11% off-grid and 30% on-grid, and its 2024 objective is to electrify 48% off-grid and 52% on-grid (USAID 2018).

Currently, the predominant sources of lighting for poor households are either flame-based (e.g., kerosene, candles) or battery-based (e.g., flashlights), mainly because these solutions are easily accessible in local retail stores. However, they are expensive to consumers in the long run and pose a threat to health and the environment due either to the harmful smoke they generate or the improper disposal of replaced batteries.

Solar-based off-grid solutions (such as basic solar home systems) are cleaner and cheaper in the long run, but they require a high upfront investment that places them well beyond the reach of liquidity-constrained consumers in countries like Rwanda.

An alternative off-grid lighting model that is becoming prominent in low-income countries is rechargeable lamp technology. Under this model, instead of selling lamps to consumers at full price, firms either rent them or sell them at a subsidized price. Continued use of such lamps requires that they be recharged at a (usually village-level) recharge center for a small recharging fee. The revenue stream from repeated recharges makes it possible for the firm to subsidize the upfront price by financing it through ongoing payments. Sunlabob in Laos, Shidhulai in Bangladesh, and Nuru Energy in Rwanda are examples of companies that operate based on this model. In this paper, we explore, in close collaboration with Nuru Energy, the consumer behavior and the operational inefficiencies that result from the rechargeable lamp-based off-grid lighting models.

Because the consumers in our study are poor, their liquidity constraints naturally play a role in determining their usage of lamps. Moreover, the rechargeable lamp-based model requires that the consumers travel to a dedicated recharge center to get their lamps recharged. Many villages in East African countries, for example, are located in hilly areas and typically have neither efficient public transportation nor good-quality roads for walking. From our surveys in Rwanda, we know that most consumers walk to the recharge center, and for some of them, a round trip can take up to an hour. Therefore, the time required to recharge lamps is a significant inconvenience, which impacts lamp usage.[1] We are specifically interested in examining, both theoretically and empirically, the impact of the above two factors – liquidity constraints and recharge inconvenience – on the usage of lamps.

Furthermore, our broader objective is to evaluate the efficacy – in terms of increasing the number of recharges – of different strategies that address those two factors and thereby identify better business models to serve the poor. (We use the terms strategy, business model, and policy interchangeably.) Such an undertaking has implications for both firm-level operational decisions and government-level policy decisions. Firms such as Nuru aim to serve the poor population while making a profit. Therefore, better business models enable higher usage rates of cleaner lighting sources and less use of harmful sources (such as kerosene and candles), which in turn would benefit firms, consumers, and the environment. Countries such as Rwanda, whose objective is to provide off-grid lighting to a large proportion of their populations, are currently collaborating with multiple off-grid companies operating under different business models (World Bank 2017). Understanding the benefits and limitations of using rechargeable lamp technologies enables policymakers to weigh them against the benefits and limitations of other technologies, such as solar home systems.

However, it is not easy to identify better business models to serve the poor. The lean startup philosophy (Ries 2011) taught in entrepreneurship classes advocates constant experimentation to achieve rapid improvements in the business models. But there are two issues with such an approach in a context like ours.

[1] This feature is peculiar to the rechargeable lamp-based lighting model. The grid-based and solar-based lighting solutions do not require any travel. The consumers purchase lighting sources such as candles and flashlight batteries at local retail stores, where they also purchase other retail goods (e.g., food-related items, cigarettes, toothpaste); thus, the overall inconvenience in making the purchase is split across several goods that are purchased. In contrast, when a consumer visits a recharge center, all that she gets in return is light for her lamp.

First, the direct implementation of the candidate policies in the field to assess their performance may not be practical because of the remote location of the villages and the budget constraints of firms. Second, even if we have sufficient resources for constant experimentation, as there are multiple ways we could address liquidity constraints and inconvenience, it is often unclear which strategies deserve most attention, and hence should be tested first. Therefore, we need a method to predict the performance of policies *before* they are implemented. Since a policy change could vary from a simple change in a parameter to a sophisticated change in a process within the firm, the prediction framework must be flexible enough to generate a wide variety of *counterfactual* policies to facilitate their evaluation. We also need data to build such a framework, but because there is a dearth of reliable datasets in the BoP markets, the required data may not even be readily available. The methodological approach used in this paper responds to these concerns.

Methodological approach. We build a structural model of consumers' recharge decisions and estimate that model using the field data. A primary benefit of estimating a structural model of behavior is the ability to calculate outcomes under economic environments not observed in the data. Our model therefore provides a framework for predicting counterfactual policies. The consumer in our model is forward looking, and her decision process is represented by a stochastic dynamic program over a finite horizon. In a time period in which the consumer's lamp is discharged and she has sufficient money for a recharge, by choosing not to recharge her lamp, she incurs a blackout cost for going without lamplight in that period. However, if she chooses to recharge, then she experiences no blackout cost as long as the lamp remains charged, but incurs an inconvenience cost of traveling to the recharge center and pays the recharge price in that period.[2] In contrast, in a period when the lamp is discharged but the consumer has no money for a recharge, she simply incurs a blackout cost. The liquidity constraints, which determine whether the consumer has money for a recharge in a given period, are captured in our setting through a model for disposable income for lamplight that follows a Markov process. The assumed structure can generate a variety of counterfactuals.

The policy evaluation using the structural model requires us to estimate consumer sensitivity to inconvenience and blackouts, as well as the parameters of the consumption and disposable income processes. We argue that, under some mild structural assumptions, the intertemporal variation in consumers' recharge decisions, along with exogenous variations in recharge price, lamp capacity (hours of light obtained per recharge), and consumer inconvenience (proxied by distances from the recharge center), are sufficient to identify each component of the model separately. Consequently, we need the actual data on lamp usage at different prices, capacities, and inconvenience levels. Due to the lack of any such data, we conducted field experiments in collaboration with Nuru in 29 villages of Rwanda, where we randomly assigned the recharge price and lamp capacity to consumers. We implemented an automated system that records recharge timestamps along with the identifiers of the lamps being recharged. We also recorded the global positioning system (GPS) coordinates of households and recharge centers to calculate the distances between them.

[2] In our paper, we use the word "lamp" to refer to the whole package consisting of a light-emitting diode (LED) as the light source, a plastic housing for the LED, a strap, a battery, and a switch to turn the light on and off. The terms "lamp recharge" and "lamp capacity" actually refer to a recharge of the battery inside the lamp and the capacity of that battery, respectively.

To rely on the predictions made by a structural model in a counterfactual setting, we must establish that the model is empirically consistent. For this purpose, we first conduct a simple reduced-form regression analysis and show that the theoretical predictions made by the model are directionally consistent with the experimental data. Thereafter, we test the ability of the model and its variants (e.g., including village- and individual-level heterogeneities, discounting) to predict the number of recharges both in-sample and out-of-sample. We observe that the best-fitting model predicts the number of recharges both in- and out-of-sample reasonably well.

Results. We examine the performance of business model changes that target liquidity constraints, inconvenience, recharge price, and lamp capacity. We find that the complete removal of liquidity constraints and inconvenience from the current business model results in up to a 126% and 73% increase in recharges respectively, suggesting that they are major sources of inefficiency. Although these benchmark cases may not be achievable in practice, significant improvements – and sometimes more than half the benefits to be gained from the complete elimination of those inefficiencies – can be achieved by implementing some simple strategies that (i) alleviate liquidity constraints, e.g., allowing consumers to partially recharge their lamps (20% increase), prepay for the recharge (43% increase), and recharge on credit 1–2 times (84% increase); and (ii) alleviate inconvenience, e.g., starting 2–3 more recharge centers per village (26% increase), visiting households door-to-door once a week to collect the lamps for recharge (34% increase), and visiting just five locations per village twice a week to collect the lamps in those localities (32% increase).

The above-mentioned strategies only vary the operational model of the firm by addressing the sources of inefficiencies and by changing the recharge and payment processes within the firm. They do not affect the recharge price or lamp capacity, and hence they increase both recharges and revenue simultaneously. We also examine strategies that vary price and capacity without affecting the current operational model. We find that the recharges are relatively inelastic with respect to price; thus, although it increases the number of recharges, reducing price simply decreases revenue. Consequently, scaling up by offering subsidies to consumers may not be a sustainable strategy for the firm in the long run. If the firm also varies lamp capacity along with price, then we find that both recharges and revenue improve and that it is optimal for the firm to reduce both price and capacity. However, (i) the resulting improvements are quite limited (7% increase in recharges and 2% increase in revenue), and (ii) the resulting ratio of capacity to price (i.e., bang for the buck) is lower than that at the status quo. In other words, the consumers pay a poverty premium when the light is provided in a smaller, affordable package. Overall, we find that operations-based strategies perform better than price- and capacity-based strategies.

The firms and policymakers in poor countries generally lean toward price-based strategies because poverty, by definition, entails a lack of money. However, our analysis reveals that even when operating under poverty, where monetary constraints are overpowering and may limit technology adoption, the operational inefficiencies embedded in the business model (e.g., constraints such as making the consumer travel to a single village-level recharge center and making her pay only when she recharges her lamp) may also be major hindrances to adoption and addressing them results in significantly more benefits.

We believe that our model and some of its insights have the potential to encompass services other than light, such as access to clean water, or fertilizers, or even consumer goods, where liquidity-constrained consumers frequently incur an inconvenience to access the service. Moreover, the template that we use for policy evaluation, combining structural modeling and field experiments, can be applied to settings that go beyond our context. Such template has been advocated by development economists, but it is yet to penetrate the operations management literature. The approach in this paper can be used to generate hypotheses on business operations – grounded in both theory and data – for experimentation, and thereby to arrive at appropriate business models that deliver life-improving goods and services to poor consumers.

Organization of the paper. Section 2 provides an overview of our approach: it describes the relationship between our research objectives, our structural model, and the design of our field experiments. Thereafter, Section 3 gives the details of the field experiments, Section 4 describes the model and the estimation procedure, Section 5 presents the performance of the counterfactuals that we examine, and Section 6 makes some concluding remarks. After reading Section 2, Sections 3–5 can be read in any order depending on the interests and priorities of the reader. For example, an entrepreneur or a policy official interested mainly in our findings and an outline of the methodology used to arrive at those findings can read only Sections 2 and 5 without becoming burdened by the technicalities, whereas an academic or a practitioner interested in the implementation details can additionally read Sections 3 and 4. We briefly review the related literature in the remainder of this section.

Related literature. Our paper is positioned at the intersection of two streams of literature: sustainable operations and the economics of poverty.

The challenges of devising sustainable business models are discussed extensively in Kleindorfer et al. (2005), Drake and Spinler (2013), Girotra and Netessine (2013), and Plambeck (2013). We broadly relate to the growing body of research that studies the operational issues in the business models that serve the BoP markets. Balasubramanian et al. (2017) study the inventory issues arising in the context of mobile money agents; Jonasson et al. (2017) develop models to improve the capacity allocation of laboratories for the early diagnosis of the human immunodeficiency virus (HIV) among infants; Gui et al. (2019) examine the efficacy of purchasing cooperatives and non-profit wholesalers in terms of replenishing goods for microretailers; de Zegher et al. (2018) propose payment strategies to curb illegal deforestation by smallholder farmers; Guajardo (2019) explores the relationship between consumer usage and payment behaviors in a rent-to-own business model for the distribution of solar lamps; and Kundu and Ramdas (2019) investigate the impact of timely after-sales service on the adoption of solar home systems by first-time users in developing countries.

Our paper is closely related to Uppari et al. (2019), who use an analytical model to study why some consumers may prefer using kerosene to rechargeable lamps even when the latter cost less money than the former. They find that the consumers who face either high inconvenience costs or high blackout costs tend to prefer kerosene to rechargeable lamps because the former's flexibility, with regard to quantity, helps reduce whichever cost is dominating. That paper also discusses some strategies (e.g., allowing partial recharges, encouraging consumers to pool) that would improve lamp adoption. The empirical work in this paper complements the theoretical work of Uppari et al. (2019). Our structural model includes inconvenience and blackout costs along with liquidity constraints much in the spirit of Uppari et al. (2019), and our field data allows us to estimate those costs and the magnitude of the liquidity constraints. Furthermore, our paper quantifies the impact of the strategies discussed in Uppari et al. (2019), while also examining several additional strategies.

The method of deploying field experiments to analyze consumer behavior in the context of poverty has been used in the development economics literature. This literature mainly investigates various behavioral impacts of prices on the adoption of a technology. For example, Cohen and Dupas (2010) examine whether consumers waste goods that are distributed freely to them; Ashraf et al. (2010) analyze the role of higher prices in increasing product use through screening and sunk-cost effects; Dupas (2014b) studies the extent to which consumers' anchoring on subsidized prices impacts the long-run adoption after those subsidies are removed; and Duflo et al. (2011) study the impact of seasonal income on the purchase of fertilizers. In most of these papers, consumers make no more than two purchase decisions at different price levels (e.g., subsidized and non-subsidized). In contrast, consumers in our setting make multiple recharge decisions over time at a given price level, but those price levels differ across consumers. In such a setting with *repeated purchases*, inconvenience cost – which is incurred with every purchase – plays an important role in the long-run usage of the technology. Therefore, our interest in this paper lies in examining the impact of price as well as the impact of inconvenience on lamp usage.

Another emerging field in development economics deploys field experiments to evaluate the economic, health, and educational impacts of using cleaner lighting sources on the lives of the poor. Barron and Torero (2017) and Lee et al. (2020a,b) conduct impact analyses of providing access to grid-based electricity, Aklin et al. (2017) investigate the benefits of connecting to solar microgrids, and the welfare impacts of using solar lamps are explored by Furukawa (2014), Grimm et al. (2017), Rom et al. (2017), and Stojanovski et al. (2018). Overall, these studies find that access to clean light reduces indoor air pollution, increases productive time, and reduces the expenditure on harmful sources; the evidence for savings and children's study times is mixed. Our paper complements these impact studies by focusing on firm-level policies that increase lamp usage, which thereby also increases the welfare of consumers. In other words, the number of recharges done by the consumers acts as a proxy measure for welfare in our paper. A full-blown impact analysis in our context is ongoing (Clarke et al. 2022), and it is beyond the scope of this paper.

Banerjee et al. (2017) advocate the addition of a structural model to experimental research as it facilitates "structured speculation"; that is, it can lead to a fully-specified set of falsifiable predictions in external environments. We use a dynamic programming (DP)-based structural model in our paper. The introduction of the framework to estimate discrete-choice DP models is associated with independent contributions by Miller (1984), Pakes (1986), Rust (1987), and Wolpin (1984, 1987). For the methodology of this framework, we refer the reader to Rust (1994) and Aguirregabiria and Mira (2010). The framework has been applied in several contexts to evaluate economic policies that are particularly relevant to developing countries; see Todd and Wolpin (2010b) for a review of such applications. We use that framework in conjunction with field experiments for an *extensive business model analysis* in the context of poverty. Moreover, since the DP model in our paper incorporates several operational features of the low-income regions that are not part of the existing DP models in the structural estimation literature, it presents a unique set of identification challenges that are resolved through field experiments. Finally, to rely on the predictions made by a structural model in a counterfactual scenario, Keane (2010) and Todd and Wolpin (2010b) emphasize the necessity of validating its predictive ability. In this study, we examine the predictive ability of our DP model – estimated using the data from treatment conditions – on the data from both the control condition and a test treatment condition. Such validation methods are also applied in Todd and Wolpin (2006) and Duflo et al. (2012).

2. Overview of Our Approach to the Problem

Our research objectives, our structural model of consumer behavior, and the design of our field experiments are intimately related to each other. Therefore, instead of delving into a deeper discussion on these three topics sequentially, we first provide an overview in this section to shed light on their relationship with one another. The technical details of our model and the implementation details of our field experiments are provided in the later sections.

2.1. Experimental Context and Research Objectives

The research in this paper was conducted in collaboration with Nuru Energy in Rwanda. Nuru is a for-profit social enterprise, with operations in Rwanda, Burundi, and Kenya, that aims to address the issue of energy poverty through the provision of rechargeable lamps and lamp-recharging centers to poor off-grid rural communities. The lamps are sold below cost to make them affordable. (Each lamp costs Nuru 6 USD to manufacture but is sold to consumers at 1–1.5 USD. The lamps are made to last for 250 recharges.) Continued use of lamps requires that they be recharged at a centralized pedal-and-solar-powered recharge center operated by village-level entrepreneurs (VLEs) who charge lamps. The recharge centers are usually located in the homes of those VLEs. Under the current business model, a lamp recharge costs 100 Rwandan Francs (RWF) for a lamp capacity of 18 hours. The VLEs earn 50 RWF per recharge. The revenue stream from recharges makes it possible for Nuru to subsidize the upfront price by financing it through ongoing payments; thus, in order to turn a profit, it is important for Nuru to have a steady stream of recharges.

The quantity of interest for Nuru and other firms running a rechargeable lamp business is the number of recharges they can expect based on a given policy $\mathcal{P}$. We use the term *policy* quite broadly here, and it may encompass several business-related decisions such as the price and the capacity of the lamps, the payment schemes, and the location and the number of recharge centers in a village. For instance, under the status quo policy, Nuru charges 100 RWF per recharge for the 18-hour capacity lamps, and the consumer travels to a village-level recharge center to recharge her lamp and pay the recharge fee to the VLE.

We are interested in assessing how the number of recharges varies across different policies.[3] Given our broad definition of policy, it is important to note that we are not interested in the *optimal* policy of any particular type. Instead, we focus on evaluating the efficacy of some implementable policies that address the liquidity constraints and inconvenience of consumers. Such policy changes may vary from simple ones such

[3] Because rechargeable lamps are sold at a low (subsidized) price, upfront purchase cost is unlikely to be a barrier to adoption; therefore, we do not consider its impact in this paper. Accordingly, we set the purchase price to zero in our field experiments. Since 2019, Nuru also sets the purchase price to zero. (For a closer examination of the impact of upfront price in our context, we refer the reader to the parallel work by Clarke et al. 2020.)

as changing a parameter (e.g., price and capacity) to more sophisticated ones such as decoupling payments from recharges through mobile payment schemes.

One way to evaluate candidate policies is by directly implementing them in the field to see how they perform. However, this approach may not always be feasible, especially in a context like ours, wherein the firms operate under tight budget constraints and the policies must be implemented in remote villages requiring nontrivial investments in both time and money. Therefore, we need to be able to assess the effectiveness of a policy *before* it is implemented in the field, i.e., we need to perform *ex-ante policy evaluation* (Todd and Wolpin 2010a). Such evaluation necessitates a formal framework for the consumers' recharge decisions.

Formally, assuming a discrete time domain, we denote the recharge decision of consumer j in period t as r_{jt}, where $r_{jt}=1$ if that consumer recharges her lamp in period t, and it is 0 otherwise. In an arbitrarily given duration $\{1,\ldots,T\}$, we denote the sequence of a consumer's recharge decisions as $\mathfrak{r}_j=(r_{j1},\ldots,r_{jT})$, which we assume is a realization of the random variable $\tilde{\mathfrak{r}}_j=(\tilde{r}_{j1},\ldots,\tilde{r}_{jT})$. Correspondingly, the total number of recharges is denoted by $R_j=\sum_{t=1}^T r_{jt}$ and $\tilde{R}_j=\sum_{t=1}^T \tilde{r}_{jt}$. Then, under a policy $\mathcal{P}$, the expected number of recharges by consumer j is

$$\mathbb{E}\tilde{R}_j(\mathcal{P})=\sum_{\mathfrak{r}_j\in\mathfrak{R}} R_j\times \Pr\left(\tilde{\mathfrak{r}}_j=\mathfrak{r}_j \mid \mathcal{P}\right), \tag{1}$$

where $\Pr\left(\tilde{\mathfrak{r}}_j=\mathfrak{r}_j \mid \mathcal{P}\right)$ is the probability of observing the recharge sequence $\mathfrak{r}_j$ under policy $\mathcal{P}$, and $\mathfrak{R}$ is the set of all 2^T possible recharge sequences.

To evaluate different policies, we need to know what the aforementioned probability would be under each one. We take the following approach to estimate that probability in *counterfactual* settings. First, we build a model of consumer behavior wherein its structural components interact to generate the distribution of recharge decisions. Second, we estimate the parameters of the model components; this requires data with the necessary set of variations that are obtained through our field experiments. Finally, under a counterfactual policy, the model components would interact in a different manner, and because we know the parameter estimates of the components, we can also estimate the distribution of recharges under that policy. In the next two subsections, we give an overview of our structural model and discuss the relationship between its components and the variations required in the data to be able to identify those components.

2.2. Structure of the Decision Process

Let P be the recharge price of the lamp, Q be the lamp's capacity, and I be the inconvenience experienced by the consumer in recharging the lamp. Figure 1 pictorially represents the decision process that we assume for our focal consumer. (For notational simplicity, we suppress the subscript j representing consumers.) The consumer recharges her lamp in period t if and only if the following three conditions are satisfied: (i) her lamp is *discharged* in period t, (ii) she has sufficient *money* for the recharge in period t, and (iii) it is better to recharge *sooner* (i.e., in period t) than later (i.e., in a period $t'>t$). Conditions (i), (ii), and (iii) are, respectively, represented by the indicator variables $\mathbb{D}_t$, $\mathbb{M}_t$, and $\mathbb{S}_t$ in Figure 1.

Condition (i) is motivated by both empirical and anecdotal evidence: from the data recorded at the Nuru recharge centers, we find that there is no charge remaining in the lamp when it is plugged in to be recharged.

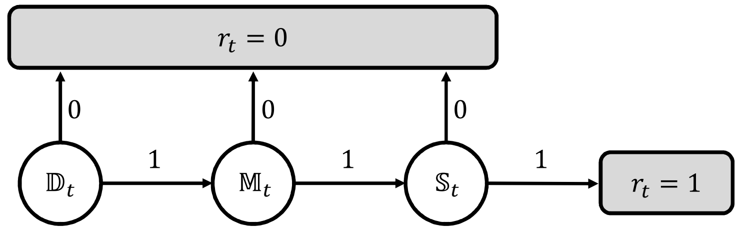

Figure 1 Decision process of the consumer.

This finding is further supported by the survey data wherein consumers said that they do not recharge their lamps until they are completely discharged.[4] Condition (ii) represents the liquidity constraints of the consumers; even if a consumer wants to recharge her lamp, she may not have enough money to do so, and hence she does not have the option to recharge.

Condition (iii) captures the trade-off between a consumer's *inconvenience cost* and *blackout cost*. Assuming conditions (i) and (ii) are satisfied in period t, a consumer who chooses not to recharge her lamp in that period incurs a disutility called the blackout cost for not having lamplight; the magnitude of this disutility could vary over time depending on the consumer's valuation of lamplight (e.g., the consumer might value it less when she has a stock of alternative lighting sources and more when her children have exams). Alternatively, if the consumer chooses to recharge her lamp in period t, then she incurs an inconvenience cost of traveling to the recharge center in that period *but* then experiences no blackout cost in the following periods for as long as the lamp remains charged. Consequently, the consumer may make a strategic decision to delay her recharge in period t if the blackout costs avoided in the current and next few periods by recharging the lamp are relatively lower than the inconvenience cost of recharging.

The indicator variables $\mathbb{D}_t$, $\mathbb{M}_t$, and $\mathbb{S}_t$ may be random, and their collections over time are modeled as (nonstationary) stochastic processes in our paper. (The model specifications are presented in Section 4.) The three processes together determine the probability $\Pr\left(\tilde{\mathfrak{r}} = \mathfrak{r} \mid \mathcal{P}_0\right)$; here, $\mathcal{P}_0$ is Nuru's status quo policy.

The structure in Figure 1 is simple, yet it allows the evaluation of the counterfactual policies that are of interest to us. For instance, any consumption-related policy interventions affect the decision process through $\mathbb{D}_t$, the policies that alleviate liquidity constraints affect the decision process through $\mathbb{M}_t$, and the strategies that target inconvenience affect through $\mathbb{S}_t$. To evaluate such policies, we first need to separately identify and estimate each of the component processes $\mathbb{D}_t$, $\mathbb{M}_t$, and $\mathbb{S}_t$. We next discuss what is further required to identify the components.

2.3. Motivation for the Field Experiments

Assuming the decision process in Figure 1, the events between recharges are illustrated in Figure 2. The time between two successive recharges, hereafter called the *interarrival time* (IAT), consists of two components: (i) the *consumption time* (CT) – when the lamplight is being consumed, and (ii) the *blackout time* (BT) – when the consumer waits to get her lamp recharged. After the consumer recharges her lamp, she consumes

[4] This could be because with the current design of the lamps, consumers cannot directly know the charge remaining in their lamp, so they do not consider recharging until it is completely discharged.

the lamplight in the next few periods, when $\mathbb{D}_t = 0$. After the lamp is discharged ($\mathbb{D}_t = 1$), the consumer might not recharge her lamp immediately due to a lack of funds. Even if she has the money for a recharge, she may choose not to do so and instead experience a few days of blackout in order to minimize frequent visits, thereby balancing the recharge inconvenience against the blackouts. Thus, the blackout time may arise either due to the consumer's liquidity constraints ($\mathbb{M}_t = 0$) or to her strategic behavior to balance the inconvenience and blackout costs ($\mathbb{S}_t = 0$).

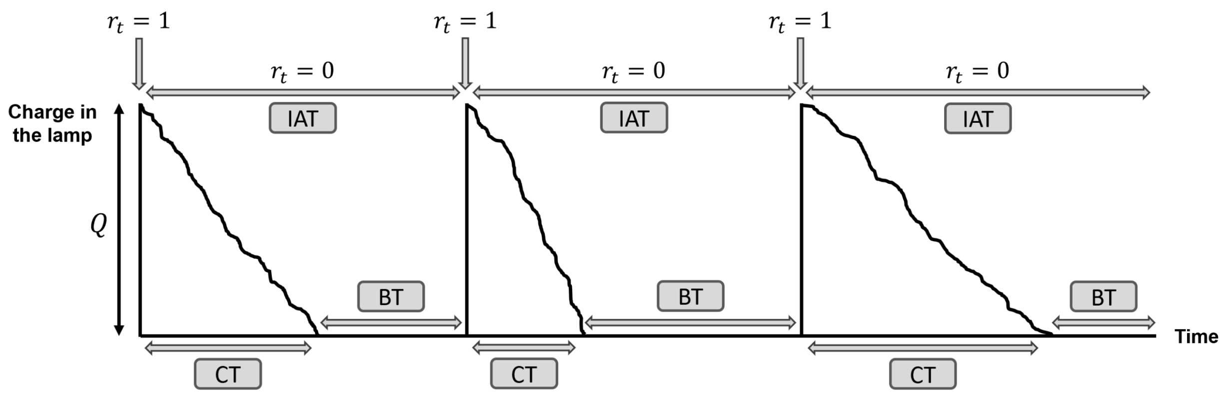

Figure 2 Consumption cycles, showing interarrival times (IATs), consumption times (CTs) and blackout times (BTs). The event $\{r_t = 1\} = \{\mathbb{D}_t = 1 \wedge \mathbb{M}_t = 1 \wedge \mathbb{S}_t = 1\}$, and CT and BT, respectively, consist of the events $\{\mathbb{D}_t = 0\}$ and $\{\mathbb{D}_t = 1 \wedge (\mathbb{M}_t = 0 \vee \mathbb{S}_t = 0)\}$.

Since consumers visit recharge centers to recharge their lamps, the firm can record the corresponding timestamps and thereby keep track of IATs. Unfortunately, it should be evident from Figure 2 that the data on IATs alone is not sufficient to separately identify the underlying stochastic processes $\mathbb{D}_t$, $\mathbb{M}_t$, and $\mathbb{S}_t$. A longer IAT could be because of either a longer CT or a longer BT, but if we do not separately observe CT and BT, then we cannot solely attribute the IAT's length either to $\mathbb{D}_t$ or to $\mathbb{M}_t$ and $\mathbb{S}_t$. Similarly, a longer BT could be either because of a lack of funds to recharge the lamp or because of the willingness to face a few extra days of blackout. Without data on the consumer's disposable income for the lamplight we cannot attribute the length of BT solely either to $\mathbb{M}_t$ or to $\mathbb{S}_t$.

However, if we perfectly observe (i) the instances when a consumer's lamp is discharged, (ii) the instances when she has sufficient disposable income to recharge, and (iii) the blackout costs experienced by her in every period when the lamp is discharged, then we know the CTs, BTs, and the liquidity constraints of that consumer, thereby allowing us to identify the components $\mathbb{D}_t$, $\mathbb{M}_t$, and $\mathbb{S}_t$. However, observing (iii) is almost infeasible, whereas observing (i) and (ii) requires recording the events at the consumer's household *between* recharges (in contrast to the recharge data, which is recorded at the recharge center). Given the rural context and the lack of advanced technologies that closely monitor consumer behavior, automated recording of such detailed instances is extremely difficult. Asking a consumer to regularly self-report (i) and (ii) may be costly to her, and the resultant data may not be reliable. We therefore need to use other methods to disentangle the different IAT components.

If there exists a variable that affects only one of the component processes without affecting the others, then *all else equal*, when we vary this variable exogenously, the resulting variation in IATs can be attributed solely to the corresponding process, which consequently identifies that process.

Accordingly, we make the following structural assumptions: (A1) lamp capacity Q affects only $\mathbb{D}_t$, (A2) recharge price P affects only $\mathbb{M}_t$, and (A3) inconvenience I affects only $\mathbb{S}_t$. Then, under A1–A3, exogenous variations in respectively Q, P, and I identify $\mathbb{D}_t$, $\mathbb{M}_t$, and $\mathbb{S}_t$. In other words, we need the recharge data at different exogenously-assigned values of Q, P, and I in order to estimate our structural model. Since it would be difficult to create these variations at the individual level (and hence we may not be able to identify the processes at the individual level), we resort to field experiments, wherein we create those variations across individuals and use the recharge data of those individuals to estimate the model. We return to the discussion on identification in Section 4 in relation to the parameters of $\mathbb{D}_t$, $\mathbb{M}_t$, and $\mathbb{S}_t$ after we describe the models for these processes.

Remark 1. The model presented in Figure 1 is in the spirit of a *partial equilibrium model*, i.e., we model optimal behavior only with respect to recharge decisions. Although consumers could be making several monetary allocation decisions across their various needs and consumption decisions across different lighting sources, we do not explicitly incorporate those decisions.

Moreover, the structural assumptions A1–A3 are necessary to separately identify $\mathbb{D}_t$, $\mathbb{M}_t$, and $\mathbb{S}_t$ processes. Although these assumptions seem reasonable, they may break down if the consumer's decision process is more sophisticated. For example, she may start saving for the recharges when the lamp gets closer to the discharge point, or her consumption and blackout times may be intertwined (in contrast to Figure 2, wherein BT appears only after CT) based on some of her strategic decisions. These possibilities introduce direct dependencies between Q and $\mathbb{M}_t$ or $\mathbb{S}_t$. A1–A3 rule out such dependencies.

Introducing any of these features into the decision process comes at the cost of requiring richer data (e.g., consumption, income, and expenditure decisions made across several needs, risks faced by and precautionary behavior exhibited by the consumers) for cleaner identification and model estimation. As discussed earlier, such data is difficult to record in our rural context. Moreover, as we will see in Section 4, the assumed decision model displays good out-of-sample predictive ability, suggesting that it captures the most important factors that determine consumers' recharge decisions.

3. Field Experiments

In collaboration with Nuru, we developed a purpose-designed data collection technology to record lamp recharges remotely using Global System for Mobile (GSM) communications technology. The lamps were programed to be able to communicate with the recharging device, and the recharging device was engineered to communicate via GSM to our cloud-based database. When a lamp is attached to the recharging device, it records the lamp's serial number and the date-time stamp of the recharge (accurate to a minute) and transmits this information via GSM to our database.[5] We use this data-recording mechanism in our experiments, which are described next.

3.1. Experimentally Varying the Recharge Price and the Lamp Capacity

The field experiments focused on randomly varying the recharge price and lamp capacity faced by consumers. There were ten treatment conditions in total: (i) seven conditions with seven different price levels (0, 50, 60, 70, 80, 100, and 120 RWF) and lamp capacity of 18 hours, (ii) two conditions with two price levels (80 and 100 RWF) and lamp capacity of 14 hours, and (iii) a test treatment condition where every fourth recharge was free (with a regular recharge price of 100 RWF and lamp capacity of 18 hours).

We were restricted to using only two lamp capacity conditions because of budget constraints and technological restrictions imposed by Nuru (varying the capacity requires changes at the hardware level). The above price levels for the 14-hour lamps were chosen to keep the same price per hour rate in two conditions ($14/80 \approx 18/100$). The main purpose of including the test treatment condition in our experiments was to test the predictive ability – in a counterfactual setting – of the models that are estimated using the data from the first nine conditions (see Section 4.4 for details). The purchase price of the lamps in all treatment conditions was set to zero so that there were no selection effects with regard to consumers' purchasing decisions. The VLEs remained unaffected by these treatment conditions as they were always reimbursed 50 RWF after a recharge.

The experiments were conducted in 29 villages of the Ruhango district of Rwanda. These villages are representative of rural Rwanda (and East Africa in general).[6] They had no grid connection, there were no plans to extend the grid to those villages in the near future, and Nuru had no prior operations in those villages. Before the experiments began, a list of all households in each village was obtained. Thereafter, a total of 2500 households, with around 80–90 households per village, were randomly selected and assigned to one of the above ten conditions. The random assignment of households to the treatment conditions was stratified at the village level in order to achieve balance, with around 8–9 households in a village per treatment condition.[7]

The treatments ran for a total of three months from the beginning of December 2016 to the end of February 2017 (after which business continued under Nuru's regular business model). The experimental conditions in those three months were operationalized through a coupon system. When the lamps were handed over to the

[5] The inbuilt hardware mechanisms ensure that Nuru's lamps can be charged only by Nuru's proprietary charging devices. Enabling the lamps to be charged using alternative sources (called unlocking the lamps) requires tampering with the hardware by cutting through the plastic covering of the lamp, which is usually impossible for typical consumers. None of the lamps used in our experiments were unlocked. This was verified in a survey conducted at the end of the experiments, when we asked consumers to show their lamps to examine if they had been tampered with.

[6] All of our experiments and surveys were conducted by the organization Innovations for Poverty Action, which uses standard protocols for data sampling and collection.

[7] The design with multiple treatment arms per village is common in the development economics literature; see e.g., Ashraf et al. (2010), Cohen and Dupas (2010), Meredith et al. (2013), Dupas (2014a,b), Barron and Torero (2017). By contrast, a design with only one treatment arm per village requires hundreds of villages to detect treatment effects with reasonable statistical power, thereby making such an experiment extremely costly.

consumers, they were also given a coupon card containing 15 perforated coupons, with each coupon having its own identifier (ID).[8] The coupon card explicitly mentioned the ID of the lamp assigned to the household, the recharge price and lamp capacity assigned to that household, and the names and birthdates of up to two household heads. The consumers were aware that the coupon cards they received were the result of a lottery and that the coupons would expire after three months.

To recharge during the experiments, households had to bring their lamp, coupon card, and their government-issued personal ID to the recharge center. The VLE recharged the lamp only if the ID printed on the lamp matched the lamp ID on the coupon card *and* the name and date of birth on the personal ID presented to the VLE matched the ones on the coupon card. If the details matched, then the VLE recharged the lamp, tore a coupon from the coupon card, collected the price written on the coupon, and sent the coupon ID through a mobile message to Nuru's reimbursement system. When the lamp was recharged, the lamp's ID and the recharge timestamp were automatically recorded by the recharging device and sent to Nuru's database. The VLE was then reimbursed only if the coupon ID sent by the VLE belonged to the lamp that was actually recharged (whose ID was automatically recorded).

This experimental design ensured that the households and VLEs deviated little from the protocol. As long as the VLE performed the required checks, it would not be possible for a consumer to bring another consumer's lamp or coupon and get a recharge unless they also brought the personal ID of that consumer. However, Rwandans are usually uncomfortable sharing their personal IDs with others, so this was considered to be an unlikely event. The VLE was also incentivized to perform the checks because an inconsistent pairing of coupon and lamp ID would not result in any reimbursement.

Remark 2. The aforementioned checks can maintain consistency between a lamp and its assignee at the time of recharge but *not* after it. We briefly discuss the implications of this limitation here. Between two successive recharges, a consumer could share her lamp with her neighbors for a short period. For instance, because the treatment conditions differed across neighbors in our experiment, a low-price neighbor may share her lamp with a high-price neighbor by charging her some intermediate price such that both the parties can make a gain from the trade. If such sharing was prevalent in the field, then it must reflect in the data too: a household's recharges when its neighbor was assigned the same treatment condition must be significantly different from that household's recharges when its neighbor was assigned a different treatment condition. We see from a regression analysis that there is no statistically significant relationship between the number of recharges done by a household and the treatment conditions that are (randomly) assigned to its four nearest neighbors in our sample. (The detailed analysis can be obtained from the authors upon request.) Accordingly, we ignore trading of lamps in our analysis.

Even if trading was absent, neighboring households may share their lamps because of their time-varying blackout costs. If a household's lamp is discharged but its neighbor's lamp is not, and the household's blackout cost in the upcoming periods is exceedingly high but its neighbor's blockout cost is not, then the neighbor may

[8] Fifteen coupons are sufficient if a consumer recharges once a week throughout the three months of the experiment's duration. In reality, however, the average number of recharges per household in those three months was only 2.93 (3.02), and only four households in our dataset recharged 15 times.

be willing to share her lamp. It is difficult to prevent such sharing behavior through experimental protocols – a limitation of our study. Also, we cannot detect such behavior using the data at hand, as we neither observe lamps' discharge points nor consumers' time-varying needs. However, the qualitative evidence from our field visits and surveys suggests that the households are usually possessive of their lamps and tend not to share them with others, mainly because each household had only one lamp, and its light (when available) was used on a daily basis. Accordingly, we ignore this type of lamp sharing too, and we attribute the observed usage of a lamp only to its owner in our analysis. Moreover, to the extent that lamp sharing happens randomly and not through some strategic decision-making behavior of consumers, it is implicitly accounted for in our structural model through a time-varying blackout cost (see Section 4.1): the blackout cost is low when the consumer can get her hands on one of her neighbors' lamps and high when she cannot. Therefore, we do not further complicate our model by explicitly incorporating lamp-sharing dynamics.

3.2. Variation in Inconvenience

To measure recharge inconvenience, we recorded the (three-dimensional) GPS coordinates of recharge centers and of all households in the sample. We quantify the inconvenience faced by a consumer as the distance between the GPS coordinates of her house and the recharge center. Unlike the variation in recharge price and lamp capacity, the variation in inconvenience measured in this manner is not exogenously created. But as we argue below, we consider this variation to be as good as random.

The recharge centers were established in villages at the very beginning of the experiments. If the location of the recharge center in a village was randomly selected, then the distances measured from that location are also random. However, the location of a recharge center may not be completely random because it is located in the home of the VLE – who is not randomly selected – in that village. Nevertheless, we attribute the emergence of VLEs more to their entrepreneurial inclination and less to their wealth, occupation, or any other characteristics that are correlated with their neighborhood. Hence, it is not unreasonable to assume that the location of the VLE is not systematically correlated with the characteristics of consumers, such as income and family size, that may determine lamp usage.

To further test this claim, we collected information on income, occupation, family composition, and other characteristics of consumers in twelve randomly-selected villages in our sample and found no significant correlation between consumer characteristics and distance to the recharge center. (In contrast to the survey that collected GPS coordinates, the consumer survey required multiple questions, and was therefore more expensive. Because of our budgetary constraints, we could not conduct the consumer survey in all sampled villages.) The detailed analysis is available in Appendix A. We extrapolate this observation to other villages in our sample. Our main results, presented in Section 5, remain qualitatively unaffected if we use the data from only these 12 villages instead of the full sample in our structural analysis.

3.3. Data

Table 1 reports the total number of recharges recorded in the 29 sampled villages. We see from the last column that the villages with IDs 15, 17, 18, 19, 21, 23, and 29 recorded a much lower number of recharges than the others in the sample. This is because the GSM component of the recharging device broke down in these villages during the experiments, thereby disrupting the experiment and causing partial loss of data. We therefore removed these 7 villages from our sample, and report the results using the data from the remaining 22 villages.

Village ID	Total number of households	Number of households in the sample	Number of recharges recorded
1	138	88 (64%)	115
2	151	77 (51%)	372
3	170	78 (46%)	360
4	138	92 (67%)	158
5	158	92 (58%)	392
6	120	92 (77%)	206
7	120	88 (73%)	293
8	158	75 (47%)	146
9	104	75 (72%)	259
10	159	93 (58%)	345
11	200	83 (42%)	139
12	122	90 (74%)	281
13	103	86 (84%)	326
14	149	93 (62%)	283
15	160	91 (57%)	1
16	148	89 (60%)	192
17	252	84 (33%)	1
18	127	87 (68%)	61
19	125	77 (62%)	36
20	151	84 (56%)	392
21	138	79 (57%)	61
22	156	87 (56%)	248
23	150	88 (59%)	59
24	126	89 (71%)	247
25	156	87 (56%)	248
26	136	93 (68%)	276
27	137	90 (66%)	134
28	103	84 (82%)	165
29	146	89 (61%)	60

Table 1 **Village-level statistics.**

Figure 3(a) shows the histogram of interarrival times of recharges. The modal value of IAT observed is 7 days. The histogram is also long tailed, with some mass beyond 60 days. The average value of IAT is 14.2 days, with a standard deviation of 12.4 days. Overall, IATs in the sample display sizeable variation. As we discussed in Section 2.3, this variation is crucial to identifying and estimating our model. Moreover, the minimum IAT observed in the data is 3 days. Henceforth, we combine 3 days into one unit called a *period*. It will be evident in Section 4 that the computational complexity of our estimation procedure scales with the number of time periods. Therefore, making the timeline coarser reduces computational time without losing any information on recharges.

Figure 3(b) plots the number of observed recharges across periods for three price conditions. We see that across all price conditions, the recharges are spread over the time horizon and seem to stabilize as time progresses. This suggests there is no reason to believe that consumers are trying the lamps in the first few

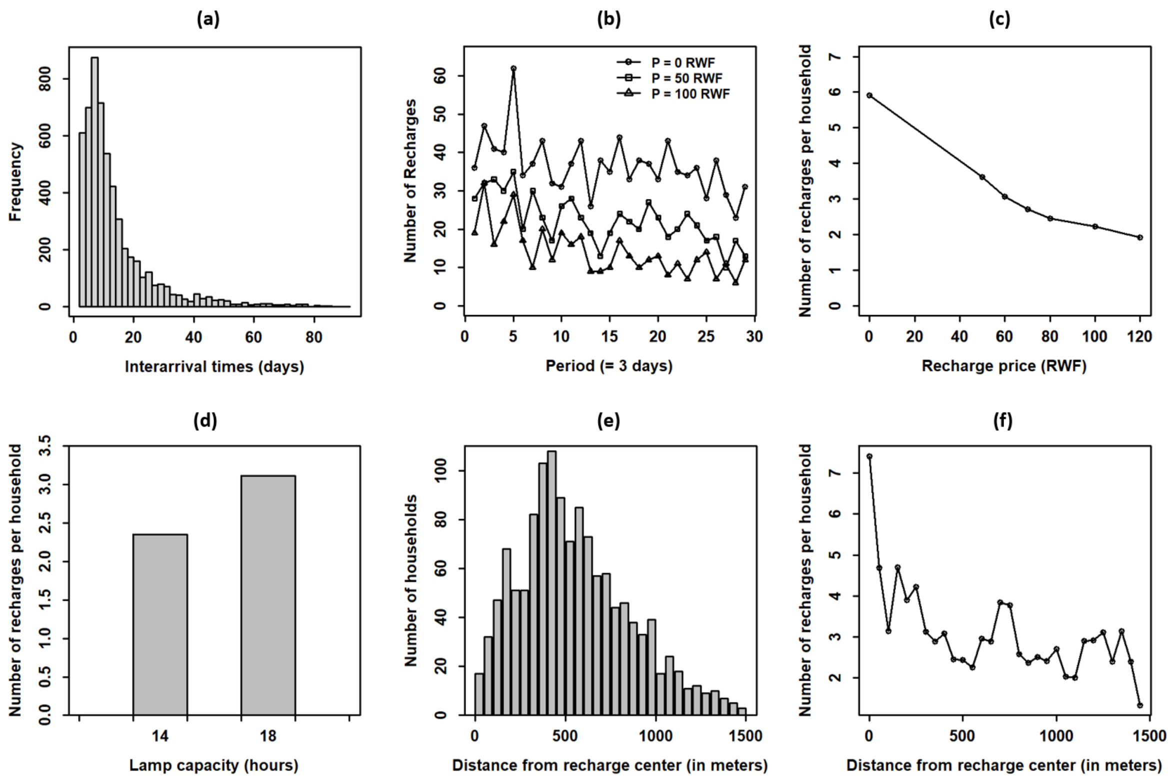

Figure 3 Some patterns observed in the recharge data: (a) histogram of interarrival times, (b) recharge patterns over time for $P \in \{0, 50, 100\}$ RWF, (c) recharges per household as a function of recharge price, (d) recharges per household as a function of lamp capacity, (e) histogram of distances from the recharge center, (f) recharges per household as a function of distance from the recharge center.

weeks (out of a desire to try a new technology) and then not using them at all. Figure 3(c) plots the average number of recharges recorded *per household* against the recharge price. Both Figure 3(b) and Figure 3(c) show that the number of recharges decreases as the price increases. It is noteworthy that in Figure 3(c), when price is zero, the average number of recharges per household is 6 in the three months of experimental duration, implying (on average) only 2 recharges per month. This shows that there are frictions beyond price at play in this context that hinder more frequent lamp usage.

Figure 3(d) plots the average number of recharges per household for the two lamp capacity conditions in our experiments. We may expect the number of recharges to be lower for the lamps with higher capacity because they last longer. However, we record 2.35 recharges per household when $Q = 14$ hours and 3.11 recharges per household when $Q = 18$ hours. We will make sense of this observation in Section 4, where we explore in more detail the theoretical relationship between the number of recharges and the lamp capacity.

Figure 3(e) plots the histogram of the distances (in buckets of length 50 meters) of households from the respective recharge centers. The modal distance between the VLE and consumers is 450 meters, and the average distance is 550 meters (with a standard deviation of 350 meters). The distribution is right-skewed, and some consumers live more than 1200 meters away from the VLE. Because of the hilly character of Rwandan villages, our surveys reveal that it takes 25–30 minutes (on average) for a consumer to travel a kilometer.

An average round trip to the recharge center could last longer than half an hour, and for some consumers it could take more than an hour. Figure 3(f) plots the average number of recharges per household as a function of distance from the VLE. The number of recharges decreases as the distance increases, suggesting that traveling to the recharge center is indeed an inconvenience that hinders lamp usage.

4. Model Formulation and Structural Estimation

4.1. Model of Recharge Decisions

The decision process in Figure 1 spans from the beginning of our experiment's duration (denoted as period $t=1$) to the end of that duration (denoted as period $t=T$). We assume that our consumer is forward-looking and that her recharge decisions emerge from a stochastic dynamic program. As in Section 2.1, the sequence of the consumer's recharge decisions is denoted as $\mathfrak{r}=(r_1,\dots,r_T)$. We denote its subsequence from period 1 to period t as $\mathfrak{r}\langle t\rangle$, such that $\mathfrak{r}\langle T\rangle=\mathfrak{r}$. Next, we construct the models for the monetary process $\mathbb{M}_t$ and the discharge process $\mathbb{D}_t$, and then formulate our consumer's DP problem. Since $\mathbb{S}_t$ captures the consumer's dynamics of cost considerations, the model for $\mathbb{S}_t$ stems from the Bellman equations of that DP.

Model of $\mathbb{M}_t$. In period t, $\mathbb{M}_t$ indicates whether a consumer has sufficient money for a recharge in that period. The money under consideration is the consumer's disposable income for lamp recharge. In other words, this is the money that the consumer could use to recharge her lamp after accounting for all her other needs. If this disposable income is greater than or equal to recharge price P, then $\mathbb{M}_t=1$; otherwise $\mathbb{M}_t=0$.

We model $\mathbb{M}_t$ as a Markov process. The probability that the process jumps from $\mathrm{m}'\in\{0,1\}$ in period $t-1$ to $\mathrm{m}\in\{0,1\}$ in period t is denoted by

$$\check{v}(t,\mathrm{m},\mathrm{m}'\,;\mathfrak{r}\langle t-1\rangle,P)\equiv\Pr(\mathbb{M}_t=\mathrm{m}\,|\,\mathbb{M}_{t-1}=\mathrm{m}'\,;\mathfrak{r}\langle t-1\rangle,P).$$

The serial dependence in the Markov chain reflects the possibility that the consumer's disposable income in the current period may depend on what she had in the previous period; the strength of this dependence may vary over time, which is reflected here by the non-stationarity of transition probabilities.

The transition probabilities may also depend on the history $\mathfrak{r}\langle t-1\rangle$ of the consumer's recharge decisions until time $t-1$ because every recharge is funded out of disposable income, and hence the subsequent probabilities of transition are affected by that recharge. We assume that the Markov chain renews after every recharge. This simplifies the relationship between $\mathfrak{r}\langle t-1\rangle$ and $\check{v}(t,)$ as shown below:

$$\check{v}(t,\mathrm{m},\mathrm{m}'\,;\mathfrak{r}\langle t-1\rangle,P)=\check{v}(t,\mathrm{m},\mathrm{m}'\,;l_t,P)\equiv v(t-l_t,\mathrm{m},\mathrm{m}'\,;P),\tag{2}$$

where l_t is defined as the latest time period before t in which the lamp was recharged. Under the renewals assumption, l_t is the only piece of information from $\mathfrak{r}\langle t-1\rangle$ that affects the transition probabilities. Moreover, the transition probabilities at any time t depend only on the relative time period $\tau_t=t-l_t$, which is period t relative to the last recharge period. In (2), the transition probability function $\check{v}$ is defined using the absolute time period t, whereas the transition probability function v is defined using the relative time period τ.

The renewals assumption can be interpreted as the setting where the consumer spends all her disposable income when the recharge is done. We believe this assumption is not unrealistic because the consumer is poor, so her disposable income for the lamp recharge would never significantly exceed the recharge price. Such renewals in disposable income are also consistent with the mental accounting model (Thaler 1985, 1999) of managing income, as discussed in Uppari et al. (2019). Accordingly, we assume that the Markov process starts from $\mathbb{M}_0 = 0$ at the outset and after every recharge.

Finally, with a slight abuse of notation, we denote by v_q the probability that the consumer has sufficient money for a recharge q periods after the recent lamp recharge, for $q \in \{1, 2, \ldots\}$. From the renewals assumption, it follows that

$$v_q(P) = \Pr(\mathbb{M}_q = 1 \mid \mathbb{M}_0 = 0\,;\, P).$$

The expression for v_q can be computed from the transition probabilities of $\mathbb{M}_t$.

Model of $\mathbb{D}_t$. We also model $\mathbb{D}_t$ as a Markov process. In period t, $\mathbb{D}_t = 0$ indicates that the lamp is not yet discharged, whereas $\mathbb{D}_t = 1$ indicates that it is discharged. Given the recharge history $\mathfrak{r}\langle t-1\rangle$, the probability of a jump from $\mathbb{d}' \in \{0, 1\}$ in period $t-1$ to $\mathbb{d} \in \{0, 1\}$ in period t is denoted by $\check{u}(t, \mathbb{d}, \mathbb{d}'\,;\, \mathfrak{r}\langle t-1\rangle, Q)$. As before, the serial dependence in the Markov chain and the non-stationarity of transition probabilities reflect the possibility that the lamp's discharge status in period t may depend on the status in period $t-1$.

Because the charge in the lamp is reset to Q hours after every recharge, the Markov chain of $\mathbb{D}_t$ renews after every recharge. Consequently, the Markov process starts from $\mathbb{D}_0 = 0$ at the outset and after every recharge. Moreover,

$$\check{u}(t, \mathbb{d}, \mathbb{d}'\,;\, \mathfrak{r}\langle t-1\rangle, Q) \equiv \Pr(\mathbb{D}_t = \mathbb{d} \mid \mathbb{D}_{t-1} = \mathbb{d}'\,;\, \mathfrak{r}\langle t-1\rangle, Q) = \check{u}(t, \mathbb{d}, \mathbb{d}'\,;\, l_t, Q) \equiv u(t - l_t, \mathbb{d}, \mathbb{d}'\,;\, Q).$$

Once the lamp is discharged, it remains in that state until it is recharged again. Therefore, we have $u(\tau, 1, 1\,;\, Q) = 1$ for any given relative time period τ.

Let $\mathcal{Q}$ be the set containing the possible number of periods the lamp could last. Then, for $q \in \mathcal{Q}$, the probability that the lamp lasts exactly q periods after the recent lamp recharge (denoted as u_q) is given by

$$u_q(Q) = \Pr(\mathbb{D}_1 = 0, \ldots, \mathbb{D}_{q-1} = 0, \mathbb{D}_q = 1 \mid \mathbb{D}_0 = 0\,;\, Q) = \prod_{\tau=1}^{q-1} u(\tau, 0, 0\,;\, Q) \times u(q, 1, 0\,;\, Q).$$

Model of $\mathbb{S}_t$. The consumer makes her recharge decisions by dynamically trading off her cost associated with recharging the lamp against the cost of not recharging it. On the one hand, if the consumer chooses to recharge in period t, then she has lamplight for the next few periods but incurs an inconvenience cost of recharging the lamp in period t. We denote this cost by αI, where I is the distance between the consumer's household and the recharge center and the coefficient α encapsulates both the physical and psychological costs associated with traveling to the recharge center (α converts distance into RWF).

On the other hand, if the consumer chooses not to recharge her lamp in period t, then she incurs a blackout cost in that period. The blackout cost arises either from having no light at all or from switching to (inferior) alternative sources such as candles, flashlights or firewood. This cost might fluctuate across periods because of inherent variations in consumer preferences and the availability of alternative lighting sources. Therefore, we model the blackout cost in period t as a random variable $\tilde{\beta}_t$ (in RWF) and assume that $\tilde{\beta}_t$ is independently and identically distributed (i.i.d.) over time, with a finite mean β and a cumulative distribution function (CDF) F that is continuous and differentiable.[9]

We assume that the consumer minimizes her total cost from period 1 to T. The state space for the consumer's DP constitutes (i) the current time period t, (ii) the blackout cost $\tilde{\beta}_t = b$ realized in period t, (iii) the indicator $\mathbb{M}_t = \mathbb{m}$ indicating whether the consumer has sufficient money for the recharge in period t, and (iv) the last period l in which the lamp was recharged. The Bellman equation for cost $C(t,b,\mathbb{m},l)$ when $\mathbb{m} = 1$ is as follows:

$$C(t,b,1,l) = \min\Big\{ \overbrace{\alpha I}^{\text{inconvenience cost}} + \sum_{q\in\mathcal{Q}} u_q\big[\overbrace{v_q\bar{C}(t+q,1,t)}^{\text{enough money after } q \text{ periods}} + \overbrace{(1-v_q)\bar{C}(t+q,0,t)}^{\text{not enough money after } q \text{ periods}} \big], \underbrace{b}_{\text{blackout cost}} + \underbrace{v(t-l+1,1,1)\bar{C}(t+1,1,l)}_{\text{enough money in the next period}} + \underbrace{v(t-l+1,0,1)\bar{C}(t+1,0,l)}_{\text{not enough money in the next period}} \Big\}, \tag{3}$$

where $\bar{C}(t,\mathbb{m},l) = \mathbb{E}C(t,\tilde{\beta},\mathbb{m},l)$. The expectation is taken with respect to the distribution of $\tilde{\beta}_t$. (Here we suppressed the argument P in v and v_q and argument Q in u_q.)

The first term in braces in (3) corresponds to the decision $r_t = 1$, while the second corresponds to $r_t = 0$. If the consumer chooses to recharge, then she incurs the inconvenience cost αI and jumps (say) q periods ahead without experiencing any cost in those periods. The lamp is again discharged after those q periods, and the consumer additionally experiences the expected cost of either $\bar{C}(t+q,1,t)$ (with probability v_q) or $\bar{C}(t+q,0,t)$ (with probability $1-v_q$). The last recharge period is set to t if the lamp is recharged. The exact number of periods that the lamp lasts (q) is uncertain at the point of recharge, so the consumer takes an expectation over its possible realizations (the realization q happens with probability u_q).

Instead, as shown by the second term in braces in (3), the consumer may opt not to recharge her lamp in period t and thereby incur the blackout cost b and an additional expected cost of either $\bar{C}(t+1,1,l)$ or $\bar{C}(t+1,0,l)$. The former expected cost is incurred if the income process remains in state 1 in period $t+1$, which happens with probability $\check{v}(t+1,1,1;l) = v(t-l+1,1,1)$, whereas the latter expected cost is incurred if the income process jumps to state 0 in period $t+1$. The variable $\mathbb{S}_t$ in Figure 1 simply indicates whether the cost of recharging is lower than the cost of not recharging:

$$\mathbb{S}_t = \mathbb{1}\Big\{\alpha I + \sum_{q\in\mathcal{Q}} u_q\big[v_q\bar{C}(t+q,1,t) + (1-v_q)\bar{C}(t+q,0,t)\big] - b - \sum_{\mathbb{m}\in\{0,1\}} v(t-l+1,\mathbb{m},1)\,\bar{C}(t+1,\mathbb{m},l) < 0\Big\}. \tag{4}$$

In the event that the consumer lacks funds for the recharge (i.e., when $\mathbb{m} = 0$), the Bellman equation is given by

$$C(t,b,0,l) = b + v(t-l+1,1,0)\bar{C}(t+1,1,l) + v(t-l+1,0,0)\bar{C}(t+1,0,l). \tag{5}$$

[9] Similar to blackout cost, the consumer's inconvenience cost may also vary over time; call it $\tilde{\iota}_t$. It will be evident from the analysis that follows that the decision about whether to recharge or not is a function of $\tilde{\iota}_t - \tilde{\beta}_t$, not $\tilde{\iota}_t$ and $\tilde{\beta}_t$ separately. If we assume that $\tilde{\iota}_t = \alpha I + \tilde{\xi}_t^{(1)}$ and $\tilde{\beta}_t = \beta + \tilde{\xi}_t^{(2)}$, where $\tilde{\xi}_t^{(1)}$ and $\tilde{\xi}_t^{(2)}$ are zero-mean random variables, then $\tilde{\iota}_t - \tilde{\beta}_t = \alpha I - \beta - (\tilde{\xi}_t^{(2)} - \tilde{\xi}_t^{(1)})$. Therefore, using the data on recharge decisions, we cannot separately identify $\tilde{\xi}_t^{(1)}$ and $\tilde{\xi}_t^{(2)}$, so we attribute time-varying nature of the inconvenience–blackout trade-off only to blackout cost.

Here, the consumer experiences blackout cost b and the expected cost of $\bar{C}(t+1,\mathbb{m},l)$ depending on the realization of $\mathbb{m}$ in period $t+1$. Finally, we assume that the cost incurred after the terminal period T is zero, i.e., $C(T+n,b,\mathbb{m},l)=0$ for all $n\geq 1$ and all feasible $(b,\mathbb{m},l)$.[10]

Remark 3. Consistent with the structural assumptions A1–A3 discussed in Section 2.3, (i) of Q, P, and I, u is a function of only Q, and v is a function of only P; and (ii) I appears only in the expression for $\mathbb{S}_t$ and does not affect either u or v. Although it may seem that Q and P affect $\mathbb{S}_t$ in (4), they do so only through u and v respectively, and hence they do not *directly* influence $\mathbb{S}_t$.

Remark 4. One may question at this point why we do not incorporate the money-in-hand (call it m_t) directly into the decision model (i.e., use m_t, instead of $\mathbb{M}_t$, as a state variable in our DP). As we mentioned in Remark 1, it is difficult to observe the consumer's disposable income for lamplight at any point in time. Therefore, m_t will be an unobserved and serially-correlated state variable in the DP. While estimating the model using recharge data, we need to account for all the possible paths the money process could have followed. Under a continuous state income process such as m_t, the estimation would involve T integrals per consumer, thereby imposing a huge computational burden (see Stinebrickner (2000) for the econometric issues associated with unobserved serially-correlated state variables in dynamic programs).

However, if we discretize the income model, then it replaces the integrals with the summations over discrete states, which lowers the computational burden. Under a setting where no explicit information on disposable income is available, only the recharge timestamps are informative about the income process: the consumer has sufficient money (i.e., $\mathbb{M}_t=1$) at the time of recharge. Therefore, if we discretize the process all the way down to a binary random variable ($\mathbb{M}_t$), we do not lose any information, but at the same time we also significantly reduce the computation time. For similar reasons, we do not model the charge remaining in the lamp; instead, we model only the corresponding indicator variable $\mathbb{D}_t$.

4.2. Structural Results

When the consumer's lamp is discharged and she has sufficient money for a recharge, she recharges her lamp if and only if $\mathbb{S}_t=1$. It follows from (4) that this condition can be rewritten as $\tilde{\beta}_t > k_t$ for some threshold k_t. The following result shows how these *blackout-cost thresholds* can be computed. (All the proofs are in Appendix G.)

Proposition 1. *In a period* $t\in\{1,\dots,T\}$ *and for a last-recharge period* $l\in\{0,\dots,t-1\}$*, the blackout-cost threshold* $k(t,l)$ *is found recursively as follows:*

$$\begin{aligned} k(t,l) &= \alpha I - \kappa(t+1,l,1) \\ &\quad + \mathbb{1}\{t+q\leq T\}\sum_{q\in\mathcal{Q}} u_q\big[v_q\big(\mathbb{E}\min\{k(t+q,t),\tilde{\beta}\}+\kappa(t+q+1,t,1)\big)+(1-v_q)\big(\beta+\kappa(t+q+1,t,0)\big)\big], \end{aligned}$$

where, for feasible $(l,\mathbb{m})$*, the function* κ *is given by* $\kappa(t,l,\mathbb{m})=0$ *for* $t\geq T+1$*, and for* $1\leq t\leq T$ *it is*

$$\kappa(t,l,\mathbb{m}) = v(t-l,1,\mathbb{m})\big[\mathbb{E}\min\{k(t,l),\tilde{\beta}\}+\kappa(t+1,l,1)\big]+v(t-l,0,\mathbb{m})\big[\beta+\kappa(t+1,l,0)\big].$$

[10] If we assume $C(T+n,b,\mathbb{m},l)$ to be some arbitrary constant $\mathfrak{c}$ for all $n\geq 1$, then the right hand sides of (3) and (5) increase by $\mathfrak{c}$, but (4) remains the same. Since the recharge decisions in our model are determined only by $\mathbb{S}_t$, and not by the constituent cost functions, $\mathfrak{c}$ does not feature in the estimation process. Thus, we simply set $\mathfrak{c}=0$.

In period t, since the consumer compares $\tilde{\beta}_t$ with $k(t,l)$ and recharges if the former is greater than the latter, we can interpret k as the effective cost (or a shadow cost) of recharging in period t, whereas $\tilde{\beta}_t$ is the cost of not recharging in that period. A lamp recharge in period t involves not only incurring an inconvenience cost of αI in that period, but also jumping a few periods ahead without experiencing any cost in those periods. Therefore, the effective cost k also accounts for the potential cost savings in those interim periods.

Using the threshold structure characterized in Propostion 1, we next write the probability of observing a recharge sequence under the decision model from Section 4.1. As in Section 2.1, we denote the recharge sequence that is a random variable as $\tilde{\mathfrak{r}}$ and the recharge sequence that is an instance of that random variable as $\mathfrak{r}$. Furthermore, we denote by (i) $\Theta = \{\alpha, u, v, F\}$ the set of all model parameters consisting of the inconvenience-sensitivity parameter and (with a slight abuse of notation) the parameters of the probability functions; (ii) $\Gamma = (I, P, Q)$ the treatment condition of the consumer comprised of her inconvenience, recharge price, and lamp capacity; (iii) $\mathfrak{l} = (l_1, \ldots, l_T)$ the sequence of observed last-recharge points, which can be computed using the recursion $l_t = (1 - r_{t-1})l_{t-1} + r_{t-1}(t-1)$ for $t > 1$ and $l_1 = 0$; and (iv) $\bar{F} = 1 - F$. The following result forms the basis for writing the likelihood function for the observed recharge sequences.

Proposition 2. *The probability of observing the recharge sequence* $\mathfrak{r}$ *is given by*

$$\Pr(\tilde{\mathfrak{r}} = \mathfrak{r}\,;\, \Theta, \Gamma, \mathfrak{l}) = \sum_{\mathbb{d} \in \{0,1\}} \sum_{\mathbb{m} \in \{0,1\}} \Omega(T, \mathbb{d}, \mathbb{m}).$$

The function $\Omega(t, \mathbb{d}, \mathbb{m})$*, wherein* $t \in \{1, \ldots, T\}$*,* $\mathbb{d} \in \{0,1\}$ *and* $\mathbb{m} \in \{0,1\}$*, can be computed recursively as*

$$\begin{aligned}
\Omega(t, \mathbb{d}, \mathbb{m}) &= \Pr(\tilde{\mathfrak{r}}\langle t\rangle = \mathfrak{r}\langle t\rangle, \mathbb{D}_t = \mathbb{d}, \mathbb{M}_t = \mathbb{m}\,;\, \Theta, \Gamma, \mathfrak{l}) \\
&= \sum_{\mathbb{d}' \in \{0,1\}} \sum_{\mathbb{m}' \in \{0,1\}} \Omega(t-1, \mathbb{d}', \mathbb{m}') \times u(t - l_t, \mathbb{d}, \mathbb{d}') \times v(t - l_t, \mathbb{m}, \mathbb{m}') \\
&\qquad \times \left[\mathbb{d}\mathbb{m}\bar{F}(k(t, l_t))\right]^{r_t} \left[1 - \mathbb{d}\mathbb{m}\bar{F}(k(t, l_t))\right]^{1-r_t} \quad \textit{for } 2 \le t \le T, \textit{ and} \\
\Omega(1, \mathbb{d}, \mathbb{m}) &= u(1, \mathbb{d}, 0) \times v(1, \mathbb{m}, 0) \times \left[\mathbb{d}\mathbb{m}\bar{F}(k(1,0))\right]^{r_1} \left[1 - \mathbb{d}\mathbb{m}\bar{F}(k(1,0))\right]^{1-r_1}.
\end{aligned}$$

It is noteworthy that both the thresholds in Proposition 1 and the probability of recharge in Proposition 2 can be expressed as recursive functions. This is important in our estimation exercise because recursivity allows for efficient computation in polynomial time using the *memoization* technique.[11]

To provide some empirical validity for our model, we examine the relationships, as predicted by our model, between the expected number of recharges and inconvenience, recharge price, and lamp capacity. For the sake of brevity, the details of that exercise are presented in Appendix B. Here, we present only our main findings. Using a simpler version of the model that is amenable to formal analysis, we arrive at the following theoretical predictions:

(Π_1) The expected number of recharges decreases in inconvenience.

(Π_2) The expected number of recharges decreases in recharge price.

[11] Memoization is a computational technique wherein computer programs are sped up by storing the results of expensive function calls and returning the cached results when the same inputs occur again. Without recursivity and memoization, computing the probability of a recharge sequence would require accounting for all the possible paths of (serially dependent) $\mathbb{D}_t$ and $\mathbb{M}_t$. Such path enumeration will result in an exponential time complexity.

(Π_3) The expected number of recharges decreases in lamp capacity for relatively smaller values of inconvenience, and it is unimodal for relatively larger values of inconvenience.

Of the above, perhaps Π_3 is most surprising because one may intuit that as Q decreases, the expected number of recharges should increase because consumers get less light per recharge, and hence the time between successive recharges will decrease. However, the model predicts that this will only be the case for relatively smaller values of I. When the consumer's inconvenience of recharging is relatively high, a decrease in capacity results in a larger number of highly inconvenient trips to the recharge center, which in turn negatively affects the overall number of recharges, as stated in Π_3. (Indeed, this is the effect that we observe on average in the data as evident in Figure 3(d).) Through regression analysis in Appendix B, we find that the recharge data from the field supports Π_1–Π_3.

4.3. Empirical Models and Parameter Identification

For the purpose of both simulating and estimating the decision model, we use the empirical models described in this section. Without these models, we need to resort to nonparametric methods to estimate the probability functions $v(\tau, \mathbb{m}, \mathbb{m}'; P)$, $u(\tau, \mathbb{d}, \mathbb{d}'; Q)$, and $F(\cdot)$. That would result in too many parameters to be estimated and require too much variation in IATs (we would need to observe every possible value of IAT a significant number of times in each condition of (I, P, Q)). The absence of such massive variation in the IATs in our sample necessitates a parametric approach to modeling the aforementioned probability functions. This of course raises the issue of whether such restrictions are valid for the particular application at hand. We take a formal approach to validating our empirical models using out-of-sample testing in Section 4.4.

Empirical model of v. We impose a structure on the transition probabilities in (2) by explicitly modeling the underlying disposable income. We assume that the log of disposable income, denoted as m_t, follows an AR(1) process, i.e., $m_t = \rho m_{t-1} + \epsilon_t$. Here, $\rho \in [0, 1)$ represents the strength of serial correlation in the AR(1) process, and the innovation $\epsilon_t \sim \mathrm{N}(\mu, \sigma^2)$ is an i.i.d. normal random variable.

Because the consumer starts to consider recharging the lamp only after she is given the lamp, we assume that this process starts afresh with initial state $m_0 = 0$. To reflect renewals, we assume that m_t is reset to 0 after every recharge. Then, we have the following result, wherein Φ is the CDF of standard normal distribution and $\bar{\Phi} = 1 - \Phi$.

Lemma 1. *Let G_z represent the CDF of normal distribution $\mathrm{N}\left(\frac{\mu(1-\rho^z)}{1-\rho}, \frac{\sigma^2(1-\rho^{2z})}{1-\rho^2}\right)$ and $\bar{G}_z = 1 - G_z$. Under the above model of disposable income for the lamplight, the following statements hold:*

(i) $v(1, 1, \cdot; P) = \bar{G}_1(\log P; \mu, \sigma, \rho)$. *For relative time period $\tau > 1$,*

$$v(\tau, 1, 0; P) = \frac{1}{G_{\tau-1}(\log P; \mu, \sigma, \rho)} \int_{-\infty}^{\log P} \bar{\Phi}\left(\frac{\log P - \rho x - \mu}{\sigma}\right) dG_{\tau-1}(x; \mu, \sigma, \rho), \tag{6}$$

$$v(\tau, 1, 1; P) = \frac{1}{\bar{G}_{\tau-1}(\log P; \mu, \sigma, \rho)} \int_{\log P}^{\infty} \bar{\Phi}\left(\frac{\log P - \rho x - \mu}{\sigma}\right) dG_{\tau-1}(x; \mu, \sigma, \rho). \tag{7}$$

(ii) $v_q(P) = \bar{G}_q(\log P; \mu, \sigma, \rho)$.

Empirical model of u. We build the model by imposing a structure on the number of periods $\tilde{N}$ that the lamp lasts after a recharge. We assume that $\tilde{N}-1$ is distributed as Poisson$(Q\lambda)$, such that $\tilde{N} \geq 1$ always. λ is the fraction of a period that is served, on average, by one hour of lamplight (or alternatively, assuming a sufficiently small λ, $Q\lambda$ is the probability that the lamp lasts longer than a period). As expected, the lamp lasts longer with higher Q and higher λ. Moreover, since $u(\tau,1,0\,;Q)$ is the probability that the lamp discharges in the τ^{th} period, given that it has not discharged in the $(\tau-1)^{\text{th}}$ period, it is simply the hazard rate of the random variable $\tilde{N}$. The following lemma formalizes these statements.

Lemma 2. *Let H represent the probability mass function of* Poisson$(Q\lambda)$*. Under the above model of lamplight consumption, the following statements hold:*

(i) $u_q(Q) = H(q-1\,;Q\lambda)$.

(ii) $u(1,1,\cdot\,;Q) = H(0\,;Q\lambda)$. *For relative time period $\tau > 1$,*

$$u(\tau,1,0\,;Q) = \frac{H(\tau-1\,;Q\lambda)}{1-\sum_{s=1}^{\tau-1} H(s-1\,;Q\lambda)} \quad \textit{and} \quad u(\tau,1,1\,;Q) = 1.$$

Empirical model of F. In our setting, $\tilde{\beta}_t$ is the disutility that the consumer experiences when she does not have lamplight. In other words, it captures the valuation that the consumer places on the lamplight *relative* to the valuation of her alternative options such as relying on kerosene and candles or simply not using a light source at all (and thereby experiencing a blackout). Because we do not model the consumption of and preference for alternative lighting solutions, we model $\tilde{\beta}_t$ over the real line such that it can take both positive values (e.g., when the consumer strongly values the lamplight) and negative values (e.g., when the consumer has a stock of alternative sources). We assume that $\tilde{\beta}_t = \beta + \tilde{\xi}_t$ where $\tilde{\xi}_t \sim \mathrm{N}(0,\sigma_\xi^2)$ and that F is the CDF of $\tilde{\beta}_t$.

Identification of parameters. With the empirical models just discussed, $\{\alpha,\beta,\sigma_\xi\}$ is the set of parameters of $\mathbb{S}_t$, $\{\mu,\sigma,\rho\}$ is the set of parameters of $\mathbb{M}_t$, and λ is the parameter of $\mathbb{D}_t$. The set of all the parameters, which we denoted as Θ in Section 4.2, reduces to $\Theta = \{\alpha,\beta,\sigma_\xi,\mu,\sigma,\rho,\lambda\}$. We now write the likelihood function, which will be maximized later to estimate Θ. We use the index set $\mathfrak{V}$ for villages and the index set function $\mathfrak{J}(v)$ for the consumers in village $v \in \mathfrak{V}$. For a consumer $j \in \mathfrak{J}(v)$, we denote by (i) $\mathfrak{r}_{jv}$ the sequence of observed recharge decisions, (ii) Γ_{jv} the treatment condition, and (iii) $\mathfrak{l}_{jv}$ the sequence of observed last-recharge points. The collections of (i)–(iii) for the entire consumer pool are written simply as $\{\mathfrak{r}_{jv}\}$, $\{\Gamma_{jv}\}$, and $\{\mathfrak{l}_{jv}\}$ respectively. The log-likelihood function for the consumer pool in village v is then given by

$$L(\Theta\,;v,\mathfrak{J}(v),\{\mathfrak{r}_{jv}\},\{\Gamma_{jv}\},\{\mathfrak{l}_{jv}\}) = \log \prod_{j\in\mathfrak{J}(v)} \Pr\left(\tilde{\mathfrak{r}} = \mathfrak{r}_{jv}\,;\Theta,\Gamma_{jv},\mathfrak{l}_{jv}\right) \quad \text{for } v \in \mathfrak{V}, \tag{8}$$

where $\Pr(\tilde{\mathfrak{r}})$ is given by Proposition 2. The complex nature of the likelihood function in (8) renders intractable the mathematical analysis of the identifiability of model parameters (e.g., showing that L is unimodal in Θ). Therefore, we resort to intuitive arguments on the identification of parameters that parallel the arguments made in Section 2.3. A detailed discussion of which sources of variation in the data identify which parameters is in Appendix C. Here, we present a summary of that discussion.

Broadly, the variation in IATs due to variation in I (resp., P and Q) helps identify $\mathbb{S}_t$ (resp., $\mathbb{M}_t$ and $\mathbb{D}_t$). If we assume that λ is known, then we can control for consumption time, and the residual variation in IATs is due to the variation in blackout times. For consumers facing the zero-price condition, the BTs only consist of the strategic component. The variation in IATs *across* those consumers (with varying inconvenience levels) identifies α, whereas the variation in IATs *within* consumers identifies β. As is common in the discrete choice literature, we cannot identify the variance of the error term, and so we normalize σ_ξ to 1. With these parameters identified, we can control for the strategic component of BT. The residual variation in IATs across individuals with varying price levels can be attributed to the liquidity component of BT, and that variation identifies μ and σ. We cannot identify ρ because of the assumed renewal structure for the disposable-income process. We treat ρ as a *hyperparameter*, i.e., we set it exogenously and then examine the sensitivity of other parameter estimates by varying ρ. All the aforesaid parameters are identified as a function of λ, and they control for blackout times. The remaining variation in IATs across individuals with varying capacity levels is attributed to the variation in CTs, which thereby identifies λ.

4.4. Model Estimation and Parameter Interpretation

Estimation procedure. As the necessary variations just discussed are available to us at the village level, we estimate the model parameters separately for each village, thereby fully incorporating heterogeneity across villages. In addition to village-level heterogeneity, there could be heterogeneity within a village across consumers. We account for this in our estimation by modeling blackout cost as a random effect.[12] Instead of estimating parameter β, we assume that β for a consumer in a village is drawn from the distribution Normal$(\mu_\beta, \sigma_\beta^2)$, and we estimate the parameters μ_β and σ_β. Since the variation in IATs *within* consumers identifies β at the individual level, it is overidentified. Thus, we can identify more parameters that represent heterogeneity in β. Furthermore, we have assumed so far in the DP formulation that the consumer does not discount her future costs. We relax this assumption in our estimation by incorporating a discount factor $\delta \in [0,1]$ into the Bellman equations (3) and (5). We, however, cannot estimate the parameter δ using the variations in our data (this is a common problem in several DP-based structural models; see, e.g., Rust 1994), and so we treat δ also as a hyperparameter.

Our estimation proceeds as follows. For each village $v \in \mathfrak{V}$, we split $\mathfrak{I}(v)$ (exclusively and exhaustively) into three sets: (i) *training set* $\mathfrak{I}^{tr}(v)$, which constitutes the data from the experimental conditions $(P,Q) \in \{\{0,50,60,70,80,120\} \times \{18\}\} \cup \{\{80,100\} \times \{14\}\}$; (ii) *cross validation set* $\mathfrak{I}^{cv}(v)$, consisting of the data from consumers facing $(P,Q) = (100,18)$, which is the current business model; and (iii) *test set* $\mathfrak{I}^{ts}(v)$, consisting of the data from our experimental condition where every fourth recharge was offered free.

We fix the set of hyperparameters $\varkappa = \{\delta, \rho\}$ exogenously and estimate the set of remaining parameters $\vartheta = \Theta \setminus \varkappa$ by maximizing the log-likelihood on training set: $\hat{\vartheta}(\varkappa, v) = \arg\max_\vartheta L(\vartheta \cup \varkappa;\, v, \mathfrak{I}^{tr}(v), \cdot)$. Thereafter, we select the (discretized) set $\varkappa$ that maximizes a score constructed by assigning half the weight to the log-likelihood on cross validation set and the other half to the log-likelihood on training set: $\hat{\varkappa}(v) =$

[12] We estimated a few other specifications with heterogeneity incorporated in income and consumption processes as well; however, they do not improve upon the performance displayed by the model presented in this section.

$\arg\max_{\varkappa}\left\{L(\hat{\vartheta}(\varkappa,v)\cup\varkappa;\, v,\mathfrak{J}^{cv}(v),\cdot)+L(\hat{\vartheta}(\varkappa,v)\cup\varkappa;\, v,\mathfrak{J}^{tr}(v),\cdot)\right\}$. We adopt this *cumulative* likelihood score as opposed to the likelihood score only on cross validation set (as in the traditional method of cross validation) to avoid overfitting to the cross validation set. We demonstrate in Appendix D that by following this method, the parameter estimates are robust to the choice of cross validation set.[13]

Overall, the set of parameter estimates that maximize the likelihood is given by $\hat{\Theta}(v)=\hat{\vartheta}(\hat{\varkappa}(v),v)\cup\hat{\varkappa}(v)$. Table 2 presents $\hat{\Theta}(v)$ for all the villages. Given $\hat{\varkappa}(v)$, Table 2 also presents the standard error estimates for parameters obtained by inverting the Fisher information matrix computed on training set.[14] Before we interpret these parameter estimates, we briefly discuss the predictive ability of the estimated model.

Goodness of fit. In addition to the primary model specification discussed so far, we estimate few alternative specifications and discuss their performance in Appendix E. We investigate some special cases, wherein the consumers (i) do not discount their future costs, (ii) are fully myopic about their future costs, and (iii) are homogenous in their blackout costs. We also examine a specification wherein the consumers are partially forward-looking and account for the costs only in the next few periods, which introduces dynamically inconsistent present-biased preferences (O'Donoghue and Rabin 1999). Comparison with these cases reveals the importance of various structural features incorporated in the model. We also compare the performance of our structural approach with some *atheoretical* approaches that do not explicitly account for the underlying decision-making processes, to see if there are any benefits to assuming a theoretical structure. (The terminology of structural vs. atheoretical is borrowed from Keane 2010.) Specifically, we study the specifications where the observed recharges are modeled as the realizations of a Poisson process or of a Bernoulli process, rather than being the realizations of a controlled decision process.

To evaluate these model specifications, we use as our goodness-of-fit criterion the mean absolute percentage error (MAPE) of the *predicted* number of recharges with respect to the *actual* number of recharges on the test set. As we use only training and cross validation sets to estimate the models, evaluation on test set is out of sample. Moreover, this evaluation criterion represents our overall objective of predicting counterfactuals:

[13] We first deployed the traditional method of cross validation, where we use the variation in the log-likelihood score on the cross validation set to select the hyperparameters. We estimated the model multiple times with $(P,Q)\in\{(100,18),(80,18),(50,18)\}$ as the cross validation sets. However, we found that, although the overall predictive results largely remained the same, the parameter estimates changed (and were statistically different) as we changed the cross validation set. This was because we use a particular treatment condition as the cross validation set and its size is small; thus, selecting (hyper)parameters using such a set resulted in fitting more to that treatment condition. To resolve this issue, we deviate slightly and use the variation in the *cumulative* likelihood score, which is the sum of the log-likelihood scores on training and cross validation sets, to select hyperparameters. The first method is equivalent to grading students based on only a final exam, whereas the second method is equivalent to grading them based on a cumulative score constructed from all the components (e.g., assignments, class participation, and the final exam). There should not be any difference between the two – in terms of who turns out to be the class topper – if the final exam alone was representative of overall performance of the students thus far. However, if the final exam is conducted only on a part of the syllabus or tests only a particular skill, then that score alone is not an ideal metric. Indeed, with the modified score, we find that the parameter estimates are robust to the choice of cross validation set, see Appendix D. The final model selection happens using the performance on test set, which we do not use in any manner while estimating the models (see Appendix E).

[14] Most of the parameters are estimated precisely. Out of $6\times 22=132$ estimates in Table 2, the ratio of standard error to mean is above 1 for only 20 parameters and above 0.5 for 43 parameters. The impreciseness of parameter estimates in some villages is a result of lack of sufficient variation in recharges in those villages.

Village ID	$\hat{\delta}$	$\hat{\rho}$	$\hat{\alpha}$	$\hat{\mu}_\beta$	$\hat{\sigma}_\beta$	$\hat{\mu}$	$\hat{\sigma}$	$\hat{\lambda}$
1	0.0	0.30	1.5055	−1.2540	0.8869	4.7302	1.6275	0.0289
			(0.3477)	(0.1805)	(0.1290)	(2.0856)	(2.7236)	(0.0269)
2	0.0	0.20	0.4566	−0.2014	1.1068	3.5760	2.3751	0.0572
			(0.5396)	(0.3248)	(0.1901)	(0.3149)	(1.2978)	(0.0092)
3	0.4	0.00	2.6597	0.8735	0.9426	4.1654	0.9505	0.0467
			(0.6349)	(0.3051)	(0.1216)	(0.1717)	(0.3574)	(0.0061)
4	0.5	0.35	0.0309	−0.6509	0.4565	0.6584	3.3411	0.0391
			(0.3192)	(0.3140)	(0.1445)	(2.3633)	(2.6643)	(0.0276)
5	0.0	0.10	1.1127	0.1032	0.8147	4.1199	0.5743	0.0559
			(0.3529)	(0.2453)	(0.1336)	(0.1517)	(0.1827)	(0.0090)
6	0.1	0.00	0.9713	−0.4916	0.4450	3.9314	1.2868	0.0310
			(0.2715)	(0.2334)	(0.1045)	(0.4133)	(0.6572)	(0.0148)
7	0.0	0.10	0.1175	−0.6420	1.2595	3.7451	0.6466	0.0656
			(0.3600)	(0.2180)	(0.5123)	(0.1509)	(0.1776)	(0.0124)
8	0.0	0.15	2.2479	−0.3595	0.6632	2.9158	1.8955	0.0199
			(0.6590)	(0.2597)	(0.1980)	(0.6767)	(1.2372)	(0.0133)
9	0.9	0.00	2.0519	0.2881	0.5753	4.2720	1.5615	0.0618
			(0.4108)	(0.1725)	(0.0746)	(0.3581)	(1.3477)	(0.0102)
10	0.9	0.05	1.0684	0.3601	0.4376	3.9018	1.1187	0.0651
			(0.2472)	(0.1820)	(0.0859)	(0.2302)	(0.4362)	(0.0091)
11	0.7	0.00	0.2340	−1.5404	0.2006	8.6952	0.1580	0.0003
			(0.2965)	(0.1976)	(0.0782)	(0.0000)	(0.0000)	(0.0082)
12	0.0	0.15	1.0358	−0.3319	0.6626	3.5978	0.9985	0.0414
			(0.5100)	(0.3916)	(0.2168)	(0.4008)	(0.4367)	(0.0140)
13	0.9	0.05	1.6954	0.4397	0.6054	3.0098	2.0606	0.0466
			(0.7334)	(0.3108)	(0.0920)	(0.7324)	(1.1813)	(0.0085)
14	0.0	0.20	0.6066	−1.0638	0.4112	4.1306	0.2308	0.0032
			(0.2539)	(0.1295)	(0.0601)	(0.0960)	(0.0749)	(0.0119)
16	0.9	0.05	2.5716	0.4926	0.7026	2.4371	1.9946	0.0802
			(2.0618)	(0.5881)	(0.2986)	(5.2275)	(6.1563)	(0.0153)
20	0.4	0.20	0.0756	−0.1505	0.8166	3.3165	1.2684	0.0572
			(0.3246)	(0.1726)	(0.1525)	(0.1686)	(0.3067)	(0.0089)
22	0.0	0.20	0.4816	−0.7718	0.3954	3.6135	0.7245	0.0302
			(0.3076)	(0.2080)	(0.1033)	(0.1861)	(0.1948)	(0.0135)
24	0.0	0.20	0.3857	−0.5956	1.1044	3.1494	1.8277	0.0863
			(0.2962)	(0.1877)	(0.3398)	(0.4014)	(1.4254)	(0.0165)
25	0.0	0.25	0.1772	−1.1752	0.5865	4.2838	0.7596	0.0363
			(0.4103)	(0.1862)	(0.0964)	(0.2862)	(0.3376)	(0.0126)
26	0.0	0.05	1.4031	−0.6974	0.5193	4.3436	1.3133	0.0098
			(0.4848)	(0.2239)	(0.0835)	(0.3944)	(0.7986)	(0.0096)
27	0.0	0.15	0.0321	−1.8766	0.6870	4.1345	0.2596	0.0237
			(0.2308)	(0.1793)	(0.0846)	(0.1132)	(0.1342)	(0.0189)
28	0.1	0.00	0.7453	−0.7044	0.6175	2.5094	2.692	0.0074
			(0.5038)	(0.1682)	(0.2883)	(2.2126)	(3.0814)	(0.0123)
Average	0.3	0.13	0.9848	−0.4523	0.6772	3.7835	1.3484	0.0406
	(0.4)	(0.10)	(0.8377)	(0.6975)	(0.2627)	(1.4082)	(0.8290)	(0.0242)

Table 2 Maximum likelihood estimates of the model parameters for all villages, along with their standard errors in parentheses. The last row presents the averages and the standard deviations of the mean parameter estimates.

the consumer's decision process in the experimental condition of test set is *structurally different* from the decision process in Figure 1, because the price faced by a consumer in the test set varies dynamically depending on her recharge decisions, and she needs to keep track of an additional state variable, which is the number of recharges done since the previous free recharge. (The corresponding decision process is presented in Figure 9(a) and the Bellman equations are given in Appendix F.)

We resort to simulations to obtain the predicted number of recharges. For each model specification, we generate the recharge sequence $\hat{\mathfrak{r}}_{jv,n}$ for consumer j in village v in simulation round n. The sum of the recharge decisions in $\hat{\mathfrak{r}}_{jv,n}$ is denoted by $\hat{R}_{jv,n}$ and the actual sum of recharges observed in the sample is represented by R_{jv}. Our MAPE measure compares $\hat{R}_{jv,n}$ with R_{jv} aggregated over consumers in all the villages, averaged across N_s simulation rounds. We use $N_s = 100$ in all the simulations in our paper.

We find that the primary model specification, with the estimates in Table 2, displays a MAPE of 7.5% (2.1%) on the test set – the lowest of all the specifications that we have examined (see Appendix E). It performs better than its special cases *out of sample*, thereby implying that the consumers are heterogenous in their valuations of lamplight and they discount their future costs but are not fully myopic. Moreover, as the estimated discount rates in Table 2 are already quite small, we see no benefits in artificially limiting consumers' planning horizons. The atheoretical approaches too perform worse on the test set, suggesting that it is important to account for the decision-making aspects while modeling a counterfactual setting. This feature particularly becomes important in Section 5, where the counterfactuals are structurally more complex and even representing them is nearly impossible without imposing a theoretical structure.

The goodness-of-fit of our model is further demonstrated in Figure 4, which plots the simulated recharges $\hat{R}_{jv,n}$ against the actual recharges R_{jv} by slicing and dicing the data in different ways. As we mentioned in Section 4.3, the variations in I, P, and Q in conjunction with the variation in IATs identify the model parameters. Figures 4(a)–(d) show how the model fits back to those identifying variations. When plotted against the recharge price, the simulated mean number of recharges per household closely tracks the actual number of recharges per household in Figure 4(a), and the actual recharges are within the 2.5 percentile and the 97.5 percentile of the simulated recharges. The same is true with recharges across the two capacity conditions in Figure 4(b). In Figure 4(c), except at the extreme of $I \to 0$, the actual number of recharges per household at different inconvenience levels is encapsulated within the 2.5 and 97.5 percentiles of the simulated number of recharges per household, and the mean numbers track the actual numbers well. Figure 4(d) plots the cumulative histograms of simulated and actual IATs. A majority portion of the distribution of the actual IATs is within the 2.5 of 97.5 percentiles of distributions of the simulated IATs. The distribution of actual IATs seems to mildly second order stochastically dominate the mean distribution of simulated IATs, suggesting that the actual IATs have slightly more variation than the simulated IATs.

In Figure 4(d), we split the data as training, cross validation, and test sets. In all the three data sets, the actual number of recharges per household are within the 2.5 and 97.5 percentiles of the simulated numbers. The model's mean and the actual mean coincide in the training set, whereas they deviate to a small degree (7%–9%) in the cross validation and test sets. Figure 4(e) splits the data at the village level. In some villages, the mean number of simulated recharges per household deviates from the actual number of recharges per household; yet, these deviations are small, and in all the villages, the actual numbers are encapsulated within the 2.5 and 97.5 percentiles of the simulated numbers.

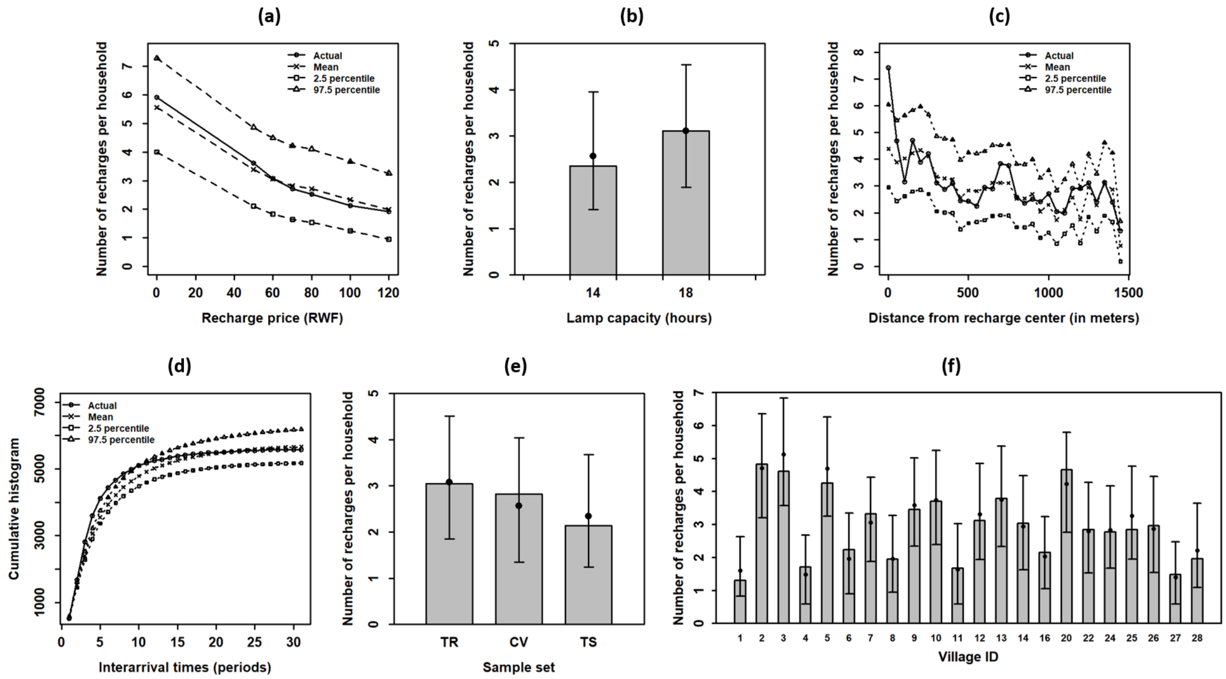

Figure 4 Goodness-of-fit of the model when the data is sliced by (a) recharge price, (b) lamp capacity, (c) inconvenience, (d) interarrival times, (e) data sets used for estimation (TR = training set, CV = cross validation set, and TS = test set), and (f) villages in the sample. In the line plots (a, c, and d), the dashed lines with crosses, squares, and triangles represent, respectively, the mean, the 2.5 percentile, and the 97.5 percentile of the simulated data. In the bar plots (b, e, and f), the black dot represents the mean of the simulated data, and the top and the bottom ends of the error bar represent the 2.5 and 97.5 percentiles of the simulated data.

Interpretation of parameters. For the sake of brevity, here we interpret in detail the average parameter estimates presented in the last row of Table 2. The interpretable measures discussed in the remainder of this section are presented for all villages later in Table 3.

The estimate $\hat{\delta} = 0.3$ suggests that a representative consumer in our sample is somewhat myopic about weighting the future costs and benefits associated with using lamps. Specifically, while making the decision in period t, compared to a weight of 1 placed on the cost in the current period, the consumer places a weight of only 0.09 ($= \hat{\delta}^{t+2}/\hat{\delta}^{t}$) on the cost that she expects to incur at the start of the next week and a weight of 0.0081 ($= \hat{\delta}^{t+4}/\hat{\delta}^{t}$) on the cost expected toward the end of the next week.[15]

We modeled the number of periods that a lamp lasts as $\tilde{N} - 1 \sim \text{Poisson}(Q\lambda)$ in Section 4.3. Using the estimate $\hat{\lambda} = 0.0406$ period/hour, Figure 5(a) presents the probability that the lamp is discharged within n periods for 14-hour and 18-hour lamps and for $n \in \{1, \dots, 5\}$. As expected, an 18-hour lamp lasts longer than a 14-hour lamp. The charge of an 18-hour (resp., a 14-hour) lamp is fully consumed within one period with probability 48% (resp., 57%), and within two periods ($\approx$ a week) with probability 83% (resp., 89%). Both types of lamp are almost certainly discharged within 4 periods.

[15] The wide range of $\hat{\delta}$ values in Table 2 calls to be viewed in relationship with the horizon length: a range of $[0.00, 0.90]$ for a 3-day horizon discount factor becomes $[0.00, 0.59]$ with a 15-day horizon, and $[0.00, 0.04]$ with a 3-month horizon. These values are consistent with the largely myopic behavior of our consumers when it comes to light.

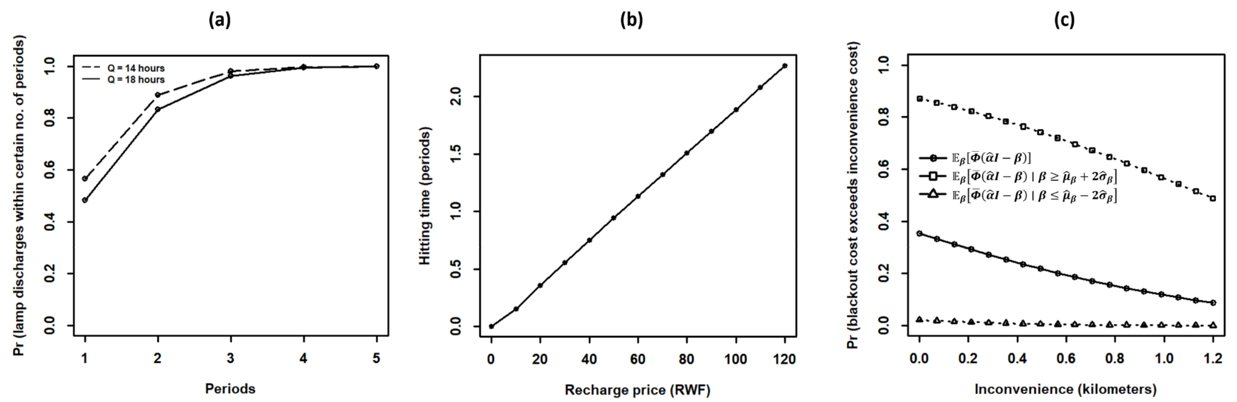

Figure 5 Interpretation of mean parameter estimates: (a) probability that the lamp is discharged within certain number of periods for $Q \in \{14, 18\}$ hours, (b) average time it takes for the disposable income to hit different price levels, and (c) probability that the blackout cost exceeds inconvenience cost at different levels of I.

Under the disposable income model discussed in Section 4.3, the estimates $\hat{\mu}$, $\hat{\sigma}$, and $\hat{\rho}$ imply that the log of the consumer's disposable income for the lamplight in period t grows as $m_t = 0.13 \times m_{t-1} + 3.7835 + 1.3484 \times \tilde{z}_t$, where $\tilde{z}_t$ is the standard normal random variable. The low value of $\hat{\rho}$ suggests that the disposable income of consumers is almost independent across periods – a feature that is expected because the poor in general do not save much, and most of the time there are enough needs in their lives that absorb the residual amounts. An alternative way to interpret the income process is in terms of its hitting times. For recharge price P, if we denote the corresponding (first) hitting time by $\tilde{h}(P)$, then

$$\Pr\left\{\tilde{h}(P) = k\right\} = \Pr\left\{\mathbb{M}_1 = \ldots = \mathbb{M}_{k-1} = 0, \mathbb{M}_k = 1\,;\, P\right\} \quad \text{for } k \in \{1, 2, \ldots\}. \tag{9}$$

Using the above probability mass function of $\tilde{h}(P)$, we compute $\mathbb{E}\tilde{h}(P)$ and plot the average hitting times for different recharge price values in Figure 5(b). We see that it takes (i) almost no time to accrue 10 RWF, (ii) 3.39 days to accrue 60 RWF, and (iii) 6.81 days on average for the disposable income to hit 120 RWF.

Because σ_ξ is normalized to 1, the coefficients α and β (thereby μ_β and σ_β) become unitless quantities. Therefore, we interpret the estimates $\hat{\alpha}$, $\hat{\mu}_\beta$, and $\hat{\sigma}_\beta$ through their impact on inconvenience and blackout costs. In any given period, the probability that the blackout cost in that period exceeds inconvenience cost is given by

$$\Pr\left\{\alpha I - \tilde{\beta}_t \leq 0\right\} = \Pr\left\{\alpha I/\sigma_\xi - \beta/\sigma_\xi \leq \tilde{z}_t\right\} = \bar{\Phi}(\alpha I - \beta), \tag{10}$$

where $\tilde{z}_t$ is the standard normal random variable. Since we model β as a random variable with distribution $\text{Normal}(\mu_\beta, \sigma_\beta^2)$, we compute three versions of (10) and plot them as a function of I in Figure 5(c). First, we compute (10) averaged across all the possible values of β: $\mathbb{E}_\beta[\bar{\Phi}(\hat{\alpha}I - \beta)]$. Recall that β in our model captures the disutility that the consumer experiences when she does not have lamplight. The aforementioned measure captures (on average) the fraction of times a representative consumer (or equivalently, the fraction of consumers in the sample) needs the lamplight enough that its value exceeds the inconvenience of recharging. As expected, this measure is decreasing in I, and it is equal to 22% and 12% when I is equal to 0.5 kms and 1.0 kms, respectively. It must be noted that this measure is not 100%, but 35%, as $I \to 0$. Although the households that are close to the recharge center face almost no inconvenience, they may not always need the lamplight (and hence their blackout cost may not exceed inconvenience cost) because some of those households may not have much activity in the night time (e.g., those with older people), or on some days they may not mind the lack of lamplight because they have a stock of alternative lighting sources.

Second, we compute (10) conditional on β being relatively high: $\mathbb{E}_\beta[\bar{\Phi}(\hat{\alpha}I-\beta) \mid \beta \geq \hat{\mu}_\beta + 2\hat{\sigma}_\beta]$. This measures the proportion of instances when the blackout cost exceeds inconvenience cost for consumers who usually value the lamplight highly (e.g., the ones with school-going children at home), or equivalently for a representative consumer on the days when she has a strong need for lamplight (e.g., on days when children have exams or guests visit the home). It is equal to 74% and 56% when I is equal to 0.5 kms and 1.0 kms respectively, and it is 87% as $I \to 0$. Third, we compute $\mathbb{E}_\beta[\bar{\Phi}(\hat{\alpha}I-\beta) \mid \beta \leq \hat{\mu}_\beta - 2\hat{\sigma}_\beta]$, which measures the probability that blackout cost is greater than inconvenience cost conditional on β being relatively low (for the reasons discussed in the previous paragraph). It is equal to 2%, 1%, and 0% when I is 0.0, 0.5, and 1.0 kilometers respectively.

Village ID	Weights on next week's costs	Avg. no. of days an 18-hour lamp lasts	Avg. no. of days to hit 100 RWF	Pr{Blackout cost ≥ Inconvenience cost}	Recharges recorded per household
1	$0.00 \to 0.00$	4.56	1.86	$0.00 \to 0.06 \to 0.50$	1.31
27	$0.00 \to 0.00$	4.28	3.97	$0.00 \to 0.06 \to 0.40$	1.49
11	$0.49 \to 0.24$	3.02	0.00	$0.02 \to 0.05 \to 0.12$	1.67
4	$0.25 \to 0.06$	5.11	20.65	$0.04 \to 0.27 \to 0.66$	1.72
8	$0.00 \to 0.00$	4.08	9.96	$0.00 \to 0.18 \to 0.69$	1.95
28	$0.01 \to 0.00$	3.40	10.75	$0.01 \to 0.22 \to 0.71$	1.96
16	$0.81 \to 0.66$	7.33	17.22	$0.00 \to 0.20 \to 0.73$	2.16
6	$0.01 \to 0.00$	4.67	6.99	$0.02 \to 0.16 \to 0.49$	2.24
24	$0.00 \to 0.00$	7.66	7.60	$0.00 \to 0.30 \to 0.96$	2.78
22	$0.00 \to 0.00$	4.63	7.34	$0.03 \to 0.18 \to 0.48$	2.85
25	$0.00 \to 0.00$	4.96	2.47	$0.00 \to 0.14 \to 0.55$	2.85
26	$0.00 \to 0.00$	3.53	3.67	$0.01 \to 0.16 \to 0.55$	2.97
14	$0.00 \to 0.00$	3.17	3.15	$0.01 \to 0.11 \to 0.36$	3.04
12	$0.00 \to 0.00$	5.24	7.98	$0.01 \to 0.28 \to 0.80$	3.12
7	$0.00 \to 0.00$	6.54	11.81	$0.00 \to 0.33 \to 0.98$	3.33
9	$0.81 \to 0.66$	6.34	4.22	$0.01 \to 0.19 \to 0.64$	3.45
10	$0.81 \to 0.66$	6.51	6.94	$0.04 \to 0.27 \to 0.64$	3.71
13	$0.81 \to 0.66$	5.52	9.91	$0.03 \to 0.36 \to 0.84$	3.79
5	$0.00 \to 0.00$	6.02	5.42	$0.01 \to 0.32 \to 0.90$	4.26
3	$0.16 \to 0.03$	5.52	6.32	$0.01 \to 0.41 \to 0.97$	4.62
20	$0.16 \to 0.03$	6.09	8.22	$0.02 \to 0.44 \to 0.95$	4.67
2	$0.00 \to 0.00$	6.09	4.87	$0.00 \to 0.39 \to 0.98$	4.83

Table 3 Interpretation of parameter estimates for all villages: the second column presents $\hat{\delta}^2 \to \hat{\delta}^4$, the third column presents $\mathbb{E}[\tilde{N}; Q=18]$, the fourth column presents $\mathbb{E}[\tilde{h}(100)]$, and the fifth column presents $\mathbb{E}_\beta[\bar{\Phi}(\hat{\alpha}I_v-\beta) \mid \beta \leq \hat{\mu}_\beta - 2\hat{\sigma}_\beta] \to \mathbb{E}_\beta[\bar{\Phi}(\hat{\alpha}I_v-\beta)] \to \mathbb{E}_\beta[\bar{\Phi}(\hat{\alpha}I_v-\beta) \mid \beta \geq \hat{\mu}_\beta + 2\hat{\sigma}_\beta]$, where I_v is the average inconvenience in village v.

Table 3 restates the estimates from Table 2 in an interpretable manner. It presents for each village (i) the weights that a representative consumer places on the costs of recharging that are expected in the next week while making decisions in the current period, (ii) the expected number of days an 18-hour lamp lasts, (iii) the expected number of days it takes for the disposable income process to hit 100 RWF, (iv) the fraction of instances when an average consumer needs the lamplight enough that its value exceeds the inconvenience of recharging, both on average and in the extreme conditions, and (v) the recharges recorded per household in our experiments (which is equal to fourth column of Table 1 divided by its third column). The table is sorted such that the values in the last column appear in ascending order.

In Table 3 we see that in almost all villages, the charge in the lamp is consumed within a week. In villages with relatively low recharge rates, consumers either place a low valuation on lamplight (e.g., villages 1, 27, and 11) or take too long to accrue money for the recharge (e.g., villages 4, 8, 28, and 16). When consumers look far ahead while accounting for costs, they tend to recharge relatively more often (e.g., villages 9 and 10), even when it takes longer for them to accumulate money for the recharge (e.g., villages 16 and 13).

5. Counterfactual Analysis

We broadly classify the counterfactual policies that we study here based on the factor(s) that they address: (i) liquidity-based, (ii) inconvenience-based, and (iii) price-and-capacity-based. We distinguish between liquidity-based policies and price-based policies based on whether they affect the recharge price. The latter directly modify the distribution of $\mathbb{M}_t$ by varying the price, whereas the former do not vary the price but create an *option* to recharge even when $\mathbb{M}_t = 0$.

For any given counterfactual policy $\mathcal{P}$, we are interested in the expected sum of recharges across all villages $R(\mathcal{P}) = \sum_{v\in\mathfrak{V}} \sum_{j\in\mathfrak{J}(v)} \mathbb{E}\tilde{R}_{jv}(\mathcal{P})$ resulting under that policy and in the corresponding expected revenue $V(\mathcal{P}) = P \times R(\mathcal{P})$. As in Section 4.4, we approximate the expectations with sample averages of recharge decisions in the simulations: $\mathbb{E}\tilde{R}_{jv} \approx \sum_{n=1}^{N_s} \hat{R}_{jv,n}/N_s$. We generate $\hat{R}_{jv,n}$ using (i) the decision process corresponding to the counterfactual policy $\mathcal{P}$, (ii) the condition $(I,P,Q)_{jv}$ as determined by the policy $\mathcal{P}$, and (iii) the probability models discussed in Section 4.3 with the estimated parameters $\hat{\Theta}(v)$ in Table 2.

Throughout this section, we assume that the current business model – the decision process in Figure 1 and $(I,P,Q)_{jv} = (I_{jv}, 100\text{ RWF}, 18\text{ hours})$ – is the base case to which the performance of other counterfactuals is compared. It should be noted that this base case is also a counterfactual policy because not every consumer in our sample was subjected to the aforementioned price-and-capacity conditions. We find that the expected number of recharges in this base case is 4514 (27). (Unsurprisingly, this number is lower than the actual number of recharges (5577) observed in the sample because our experiments included many consumers who faced price values lower than 100 RWF and therefore recharged more often.) Hereafter, we measure the impact of alternative policies in terms of the resultant (i) percentage increase in the number of recharges and (ii) percentage increase in the firm's revenue, when compared to the base case. Both (i) and (ii) are equivalent if the recharge price remains unchanged.

5.1. Liquidity-based Counterfactual Policies

The current business model imposes two constraints that critically interact with consumers' liquidity constraints. First, the consumer is required to always *fully* recharge her lamp and accordingly pay the full-recharge price. Sometimes however, she may have a strong need for light but may not have sufficient money to fully recharge the lamp. In that case, she could partially recharge the lamp and use that light to satisfy her need, but in the current business model, she is compelled to experience a blackout. Therefore, providing the option to partially recharge the lamp may bring benefits to both the firm and the consumer. Second, payments are *coupled* with recharges in the existing model: payment for a recharge happens at the same time as the recharge. Because of the stochastic nature of the consumer's needs and disposable income, even if a consumer has sufficient money for the recharge today, she may not have it later when she desires to recharge the lamp. Hence, it may be beneficial to offer some flexibility in payments by decoupling payments from recharges, e.g., by providing the option to prepay for the upcoming recharge or to recharge on credit.

Before we discuss the efficacy of the aforementioned policies, we consider a benchmark decision process shown in Figure 6(a). Here, the money process $\mathbb{M}_t$ does not act as a constraint in the consumer's decision process. In a given period, if the lamp is discharged and if it is convenient to recharge, then the consumer simply proceeds with the recharge. It is assumed under this benchmark case that the consumer always has sufficient money for the recharge and pays when she recharges the lamp. This case thus removes liquidity constraints from the model. For brevity, we present the Bellman equations for this decision process, and for the others in this section, in Appendix F.

Figure 6(g) presents the performance of all the liquidity-based counterfactual policies discussed in this section, by simulating the recharge decisions with respective decision processes as the data-generating processes. We see that under the benchmark model in Figure 6(a), the recharges increase by 126% over the base case.[16] Such a high increase may not be that surprising, given that our context is poverty, the defining feature of which is a lack of liquidity, a constraint that is a major contributor to the inefficiency in the business model. We also now know that the best that any policy addressing consumers' liquidity constraints can achieve is a 126% increase in recharges.

Allow partial recharges. Uppari et al. (2019) theoretically show that the firm can benefit by allowing partial recharges. We now quantify those benefits using our experimental data. We consider only two options: half recharge (denoted as $r_t^{(0.5)}$) and full recharge (denoted as $r_t^{(1)}$). Although our model can be extended to an arbitrary number of options, we note that offering a variety of partial recharge levels could negatively impact the lamp's battery life while also requiring recharge centers to upgrade their technology to track the charge level in the lamp. Therefore, the assumed setup offers both analytical and practical simplicity.

Figure 6(b) shows how the provision of partial recharges changes the consumer's decision process. The variable $\mathbb{M}_t^{(0.5)}$ (resp., $\mathbb{M}_t^{(1)}$) indicates whether there is enough money for a half (resp., full) recharge. By comparing Figure 1 with Figure 6(b), we see that in the former case, the consumer could proceed only when $\mathbb{M}_t = 1$, whereas in the latter case, she can proceed even when there is no money for a full recharge. Thus, the decision process branches at the $\mathbb{M}_t$ node, thereby alleviating the liquidity constraint to some extent. Since the trade-offs and the future costs (and hence the Bellman equations) differ across the alternative branches of the decision process, we represent $\mathbb{S}_t$ separately for each branch.

[16] The standard errors of the mean number of recharges under *all* the counterfactual policies discussed in this and subsequent sections are between 25 and 40 recharges. Therefore, the reported improvements in recharges are highly statistically significant.

We simulate the recharge decisions with Figure 6(b) as the data-generating process. Since the consumer only pays $P/2$ for a half recharge, to retain the equivalence between the increase in recharges and the increase in revenue, we count a half recharge as $r = 0.5$ in our simulations. Figure 6(g) shows that providing the partial recharging option increases the number of recharges (and revenue) by 20% over the base case.

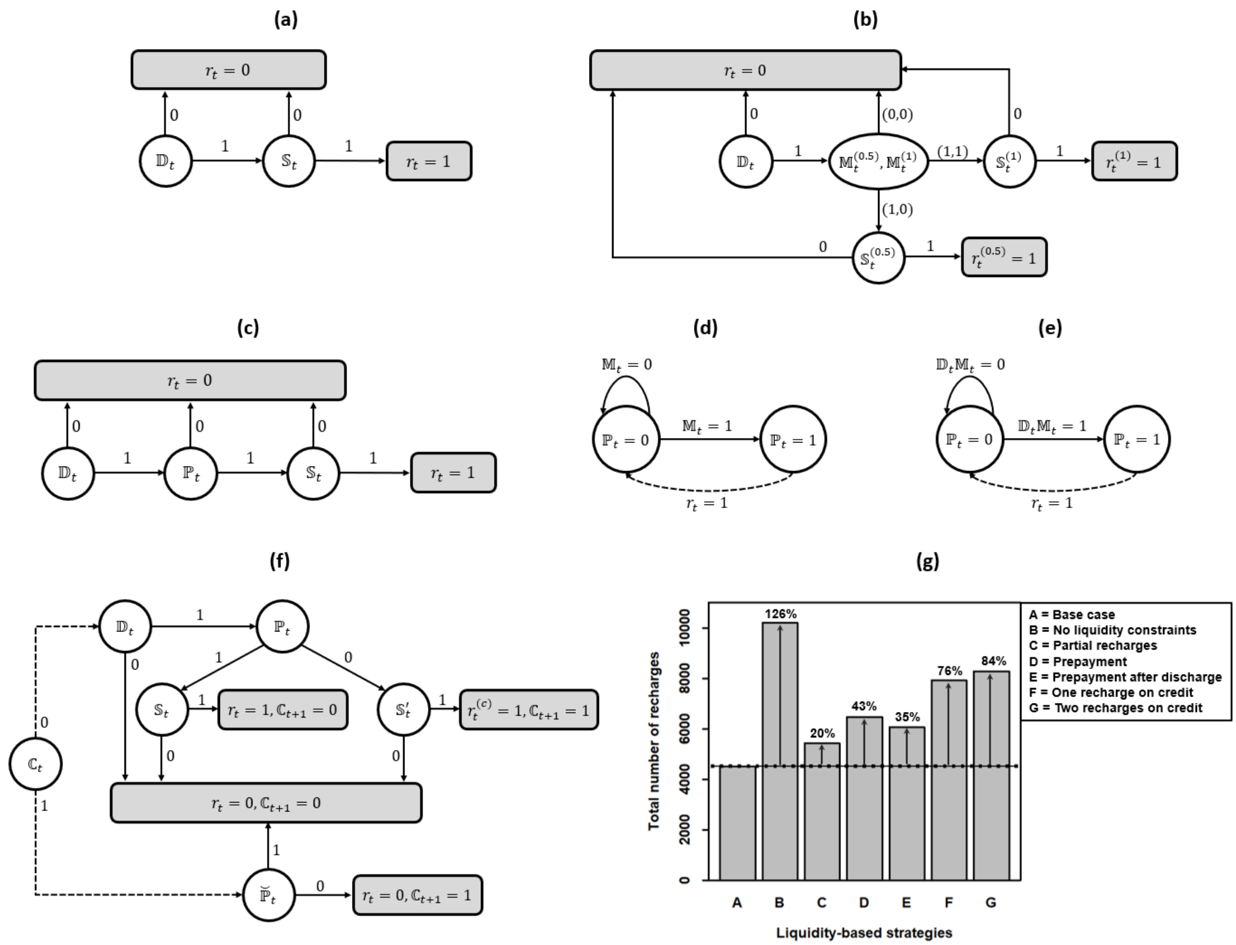

Figure 6 Liquidity-based counterfactuals: (a), (b), (c), and (f) present the decision processes, respectively, for the cases with no liquidity constraints (benchmark), the option to partially recharge, the option to prepay, and the option to recharge on credit once; (d) and (e) show the evolution of $\mathbb{P}_t$ for the case of prepayment and prepayment after discharge, respectively; and (g) shows the performance of all the liquidity-based counterfactuals.

Decouple payments from recharges. We now turn our attention to decoupling payments from recharges and note two important points. First, when payments and recharges can occur at separate points in time, there must be a mechanism in place for the consumer to make the payment without traveling to the recharge center. Otherwise, decoupling may not bring any benefits to the consumer. The flexible payment schemes can be implemented in practice by a mobile payment mechanism, wherein the consumer transfers money through her mobile phone without any need to travel. Such mobile transfers have become fairly common in sub-Saharan Africa and in developing Asia with the increased market penetration of mobile technology in those regions (GSMA 2020).

Second, as we mentioned earlier, we analyze decoupling using two mechanisms: the option to prepay and the option to recharge on credit. Given that the consumers are cash constrained, one might ask whether they will be inclined to prepay for the recharges. Using the payment data for solar lamps in sub-Saharan Africa, Guajardo (2019) demonstrates that consumers sometimes bundle their payments instead of paying a fixed amount weekly – such bundles constitute both the payments that were skipped in the past weeks as well as the advance payments for the upcoming weeks. We infer from this finding that, because of stochastic income and liquidity constraints, there is a demand for both prepayments and on-credit recharges on the consumer side. Moreover, evidence in the microfinance literature demonstrates that introducing such flexibility in payment schemes reduces financial stress and enables consumers to manage their income better (Laureti and Hamp 2011, Field et al. 2012, Barboni 2017).[17]

Figure 6(c) shows the decision process when the consumer is given the option to prepay for the recharge. We replace the variable $\mathbb{M}_t$ with the variable $\mathbb{P}_t$ in the decision process. $\mathbb{M}_t$ indicates whether there is enough money in hand to pay for the recharge, whereas $\mathbb{P}_t$ indicates whether the payment for the upcoming recharge has already been made. $\mathbb{P}_t$, naturally, is a function of $\mathbb{M}_t$, and Figures 6(d)–(e) show two plausible evolutions of $\mathbb{P}_t$. Figure 6(d) shows the setting where as soon as the consumer has enough money for the recharge, she pays for it with mobile money – she neither waits for the lamp to discharge nor considers inconvenience–blackout trade-offs. Thus, when $\mathbb{M}_t = 1$, $\mathbb{P}_t$ immediately transits from 0 to 1. Assuming that the consumer pays at the first instance when she has money may be optimistic, so we consider an alternative case in Figure 6(e). In this case, the consumer pays at the first instance when she has enough money *after* the lamp is discharged (i.e., to transit to $\mathbb{P}_t = 1$, we need both $\mathbb{M}_t = 1$ and $\mathbb{D}_t = 1$). It is straightforward to show that $\mathbb{P}_t$ in Figures 6(d)–(e) stochastically dominates $\mathbb{M}_t$. Therefore, the liquidity constraints are milder under the prepayment option. The two prepayment models result, respectively, in 43% and 35% increases in recharges over the base case.

The decision process when the firm allows only one recharge on credit is shown in Figure 6(f). The consumer can recharge her lamp once without paying, and then she must pay off that debt before the next recharge. The variable $\mathbb{C}_t$ in Figure 6(f) indicates whether the consumer previously recharged on credit. $\mathbb{C}_t$ evolves together with the recharge decisions of the consumer. When $\mathbb{C}_t$ is 0, the consumer can choose to recharge even when the payment for that recharge is not done yet (i.e., the decision process branches at $\mathbb{P}_t$). If the recharge is done without payment (denoted by $r_t^{(c)} = 1$), then $\mathbb{C}_{t+1}$ is set to 1, indicating for future reference that the consumer has recharged earlier on credit. When $\mathbb{C}_t$ is 1, the consumer has no choice other than to pay for the previous recharge. The variable $\breve{\mathbb{P}}_t$ indicates whether the payment for the previous recharge has been made. If $\breve{\mathbb{P}}_t = 1$, then $\mathbb{C}_{t+1}$ is reset to 0, indicating that the consumer is no longer in debt. Otherwise, $\mathbb{C}_{t+1}$ remains at 1. We can similarly consider the case where the firm allows two recharges on credit. The corresponding decision process is shown in Figure 9(b) in Appendix F. The simulation results show a 76% increase in recharges when the firm allows one recharge on credit and a 84% increase with the option of two recharges on credit.

[17] One can also think of the benchmark in Figure 6(a) as the setting where the consumer is given full flexibility to pay for the recharges without any constraints, which completely decouples payment and recharge processes.

To summarize, because the consumers are poor and have erratic cash flows, there are significant gains to be made by providing flexibility in payments. The maximum benefit is seen when the firm offers the option to recharge on credit, which is followed by the option to prepay and the option to partially recharge the lamp. It is worth noting that just by providing 1–2 on-credit recharges, the firm can attain more than half the benefit arising from eliminating liquidity constraints altogether.

5.2. Inconvenience-based Counterfactual Policies

Start more recharge centers. Figure 7(a) plots the impact of reducing inconvenience in various counterfactual settings. As a benchmark, we first examine the case wherein we set $I = 0$ for all the consumers. We see in Figure 7(a) that the expected number of recharges increases by 43% under this benchmark case over the base case. This reaffirms our observation from previous sections that recharge inconvenience is a significant contributor to the inefficiency in this business model, leading to low recharge rates.

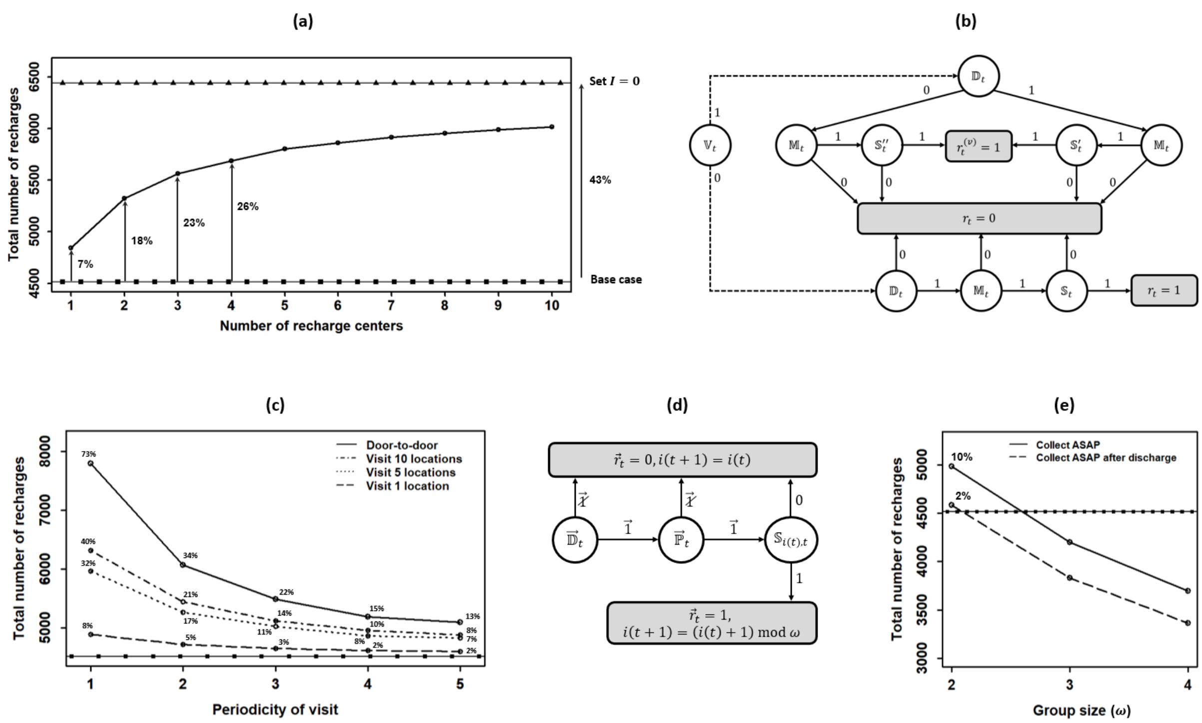

Figure 7 Inconvenience-based counterfactuals: (a) impact of increasing the number of recharge centers on total expected number of recharges; (b) decision process under the periodic-visit model; (c) performance as a function of the periodicity of VLE's visits when the visits are door-to-door, or to 1, 5, and 10 locations per village; (d) decision process under the pooling model; and (e) performance as a function of group size ω presented for two versions, one where the coordinator collects money from a group member as soon as possible (ASAP) and the other where the coordinator collects money ASAP but only after the member's lamp is discharged.

One way to reduce consumer inconvenience is by opening more recharge centers per village. Figure 7(a) shows the expected number of recharges under settings with different numbers of (optimally-located) recharge centers. (In simulations, we cluster the households in a village through a k-means clustering algorithm and then measure the impact of locating the recharge centers at the centroids of those clusters.) We find that moving the recharge center to a more central location in respective villages improves the expected recharges by 7%. Moreover, by establishing 1–3 more (optimally-located) recharge centers, the firm can improve the number of recharges by 18–26%. Although there are gains to be made beyond adding three additional recharge centers, the marginal value that they bring decreases.

There are two main practical insights from this analysis: First, there are significant gains to be made by moving the recharge center to an *inconvenience-minimizing location*. In reality, this location may or may not be the center of the village. It could also be a school or a retail store or any other place that is frequented by consumers. Because we do not have information on visit patterns to different locations in villages, we cannot measure the potential impact of moving the recharge center to one of those alternative locations. However, one would choose to move to these locations instead of the village center only if they are more convenient than the latter. The above estimate can therefore be thought of as a lower bound on the resulting improvement in recharges.

Second, it may seem obvious that adding more recharge centers reduces consumer inconvenience and increases the number of recharges, but it is often not clear how many additional recharge centers are required to bring a significant increase in the number of recharges. The above analysis shows that just two or three additional recharge centers are sufficient to capture half the benefits of reducing inconvenience to zero. It is possible that as we increase the number of recharge centers, VLEs at those centers may be less incentivized (or may even drop out) because of their reduced market share. So instead of adding the recharge centers, the firm could add “drop-off points” in a village, where consumers could safely drop off their lamps and collect them later once they are recharged.

Make periodic visits to consumers. An alternative way to reduce consumer inconvenience is by delivering a door-to-door recharge service. Rather than the consumer traveling to the recharge center, the VLE or another firm-representative can travel to the consumer. This service can also be implemented in collaboration with firms that employ agents who visit consumers frequently. For example, the agents hired by Living Goods go door-to-door to sell health products, the agents of Vision Spring conduct regular eye tests and sell eyeglasses, and Shakti agents sell fast-moving consumer goods (FMCG) products by Unilever. Such firms sometimes tend to act as distribution platforms in the BoP markets, and their services can be leveraged if feasible.

Here, we consider a simpler version wherein the firm’s VLE visits the households in her village once every n periods. If $n = 1$, then the VLE visits every period, whereas if $n = 2$ (resp., $n = 5$), then the VLE visits (approximately) once a week (resp., once every two weeks). On the days of the VLE’s visit, consumers experience zero inconvenience as they can hand over the lamp to the VLE for recharging. On such days, since the consumer must decide whether to give the (plausibly not-yet-discharged) lamp to the VLE, the consumer’s decision process under this business model will be different from the one in Figure 1. Figure 7(b) shows the consumer’s decision process under the *periodic-visit model* just described.

The indicator variable $\mathbb{V}_t$ in Figure 7(b) is equal to 1 in period t if the VLE visits the consumer in that period. $\mathbb{V}_t$ evolves deterministically: it is equal to 1 once every n periods, and it is 0 in the remaining periods.

As expected, the decision process coincides with that shown in Figure 1 when $\mathbb{V}_t$ is 0. When $\mathbb{V}_t$ is 1, the decision process branches at $\mathbb{D}_t$. Even when the lamp is not completely discharged (i.e., $\mathbb{D}_t = 0$), the consumer may decide to recharge the lamp because she can give the lamp to the VLE and thereby experience zero inconvenience. Otherwise, the consumer herself needs to visit the recharge center in a later period after the lamp discharges. In Figure 7(b), $r_t^{(v)}$ indicates the recharges where the lamps were handed over to VLE. (The Bellman equations for this model are given in Appendix F.)

Using the decision process in Figure 7(b) as the data-generating process, we simulate the recharge decisions for visit frequencies $n \in \{1, \dots, 5\}$. The corresponding expected numbers of recharges are plotted in Figure 7(c). It is noteworthy that for $n = 1$, wherein the consumer experiences zero inconvenience in all periods, we observe a 73% increase in recharges, which is higher than the improvement achieved by setting $I = 0$ in Figure 7(a). This is because the consumer decision process differs across these two cases: under the latter case, only the discharged lamps arrive for recharge, whereas under the periodic-visit model, some partially-discharged lamps are recharged too, thereby resulting in a higher number of recharges. As n increases, the increase in the number of recharges declines steeply. If the VLE visits once a week, the recharges increase by 34%, whereas if she visits once every two weeks, we observe a 13% increase in recharges.

The periodic-visit model considered so far is a door-to-door service requiring the VLE to visit each household in the village. Such door-to-door visits may be costly for a VLE, and hence in practice, she may choose to periodically visit only some select locations in the village (at some pre-decided and publicly known time), and the consumers must visit those locations to hand over the lamps to VLE and then to collect them back later. Figure 7(c) also shows the increase in recharges when the VLE visits one, five, and ten (optimally-located) points in the village at varying frequencies. We do not see much improvement when the visits are to only one location, and the benefits marginally decrease as the number of visit locations increases, as there is little difference between visiting five or ten locations. However, it is worth noting that just by visiting 5–10 locations per village, half the benefits from visiting door-to-door are captured (for all n) even though the latter option requires visiting ∼90 households per village.

Encourage consumers to pool. In the current business model, each consumer recharges only her lamp and experiences inconvenience cost with every recharge. However, as suggested by Uppari et al. (2019), if the consumers in a village form groups and the members of the groups take turns to recharge the lamps of *all* the group members, then each consumer experiences inconvenience cost only periodically when it is her turn to recharge the lamps. We now investigate this *pooling* strategy. Although the firm can only encourage but cannot enforce pooling, to shed light on the efficacy of pooling, we simply assume that the consumers adhere to pooling strategy while recharging their lamps. Furthermore, we assume that the members in a group neither pool their money nor exchange their lamps for light consumption; they form the groups only to take turns for recharging.

In this setting, we treat groups, as opposed to individuals, as the decision-making entities. Alternatively, there could be a member in each group who coordinates with all the other members of that group to make the recharge decisions. The decision process of a group (or of a *coordinating member*) is presented in Figure 7(d). We use the index set $\{0, \dots, \omega - 1\}$ for the members of a group of size ω. The index of the member who is responsible for recharge in period t is denoted as $i(t)$. Since the group is circular in terms of taking turns to recharge, the next member responsible for recharge is given by $(i(t)+1) \bmod \omega$ (where mod is the modulus operator). The vectors $\overrightarrow{\mathbb{D}_t} = [\mathbb{D}_{t,0}, \ldots, \mathbb{D}_{t,\omega-1}]$ and $\overrightarrow{\mathbb{P}_t} = [\mathbb{P}_{t,0}, \ldots, \mathbb{P}_{t,\omega-1}]$, respectively, indicate the consumption and payment statuses of the group members. In Figure 7(d), if all the lamps in the group are discharged (i.e., $\overrightarrow{\mathbb{D}_t} = \overrightarrow{1}$) *and* all the members have paid for the recharge (i.e., $\overrightarrow{\mathbb{P}_t} = \overrightarrow{1}$), then the member $i(t)$'s inconvenience cost $I_{i(t)}$ is weighed against the total blackout cost of the *entire* group (along with the future expected costs), and the lamps are taken for the recharge only if the former cost is lower than the latter cost. This inconvenience–blackout trade-off is represented by $\mathbb{S}_{t,i(t)}$ in Figure 7(d), and the corresponding Bellman equations are in Appendix F.

An important trade-off emerges under the pooling model. On the one hand, because the members of a group take turns to recharge, each member experiences lower inconvenience cost in the long run, and at any point in time, only one member's inconvenience cost is weighed against the blackout cost of the entire group. Thus, taking turns has a positive impact on the group's probability of recharge. On the other hand, individual members do not completely control when their respective lamps are recharged, as they need to wait for all the lamps in the group to discharge and for all the members to pay for the recharge. The need-to-coordinate extends the interarrival times, and ineffective coordination may negatively impact the group's recharge rates. These two counteracting forces make the overall impact of pooling ambiguous.

In Figure 7(d), we use $\mathbb{P}_t$ instead of $\mathbb{M}_t$, which reiterates the importance of effective coordination under pooling: Using $\mathbb{P}_t$ in the model is equivalent to assuming that the coordinating member proactively starts collecting money from each other member either immediately after the previous recharge ($\mathbb{P}_t$ in Figure 6(d)) or immediately after the member's lamp is discharged ($\mathbb{P}_t$ in Figure 6(e)). In contrast, using $\mathbb{M}_t$ results in a model with no proactive money collection by the coordinating member; hence, there might be cases in which all the members have sufficient money for recharge in a period, but if the lamps are not recharged in that period (because of the inconvenience–blackout trade-off), then some members may consume their money and the group reverts to a state in which it again needs to wait for all the members to accrue sufficient money. We see in our simulations that the model with $\mathbb{M}_t$ performs far worse than the status quo model, suggesting that pooling can hurt if the groups do not coordinate effectively to accrue money. Therefore, the pooling model in Figure 7(d) not only addresses consumers' inconvenience, but it also alleviates their liquidity constraints – a feature that is necessary for pooling to improve upon the status quo.

To simulate pooling, we set the group size ω exogenously and vary it from 2 to 4. For a given ω, we determine the groups by clustering the households in a village, such that each cluster has ω number of households (to the extent possible). In other words, we examine the impact of pooling when a household pools only with its nearest neighbor, or with its nearest 2 or 3 neighbors. The performance results are presented in Figure 7(e). As expected, the version of pooling model in which the coordinating member collects money from a group member as early as possible performs better than the version in which the coordinating member collects money as early as possible but only after the member's lamp is discharged. In both these versions, however, pooling performs better than status quo only when $\omega = 2$. In larger group sizes, the costs due to coordination dominate the improvements brought in by alleviating inconvenience and liquidity constraints.

Therefore, if the firm chooses to deploy pooling in the field, then it must encourage consumers to simply pair up with their neighbors and not build groups of size larger than two.

Another important thing to note is that both the versions of pooling model perform worse than the prepayment models discussed in Section 5.1, which also have $\mathbb{P}_t$ as the underlying monetary model. The latter models increase recharges by 35% and 43%, which is achieved purely by addressing liquidity constraints. In constrast, the former models increase recharges by 2% and 10%, but in this case, the increase in recharges achieved by addressing inconvenience and liquidity constraints is counteracted by the decrease in recharges because of coordination. Consequently, pooling strategy performs worse than prepayment strategy. Nevertheless, pooling can be a useful strategy as it can be readily implemented in the field, whereas implementing prepayment models first requires surmounting the barrier of deploying mobile payment mechansims.

Remark 5. The consumers may naturally tend toward the pooling strategy even if the firm does not explicitly encourage it. In that case, the benchmark for the inconvenience-based strategies examined earlier will not be our base model, but it will be the model with pooling. Moreover, because pooling already reduces consumers' inconvenience to some extent, the percentage increases in recharges presented in Figures 7(a) and 7(c) are overestimated if pooling was the benchmark. From the above analysis of pooling, we can also assess *indirectly* a counterfactual's performance if consumers pooled by default in the field (which, as we discussed above, is likely to happen only in small group sizes). For example, if pooling results in 10% increase, whereas going door-to-door results in 34% increase, then the latter strategy results in $(0.34-0.10)\times 100/1.10 \approx 22\%$ improvement relative to pooling. Alternatively, depending on the extent of pooling in the field, the door-to-door strategy may result in an improvement that is between 22% and 34%.

5.3. Price- and Capacity-based Counterfactual Policies

In all the counterfactuals that we have discussed thus far, the recharge price P was fixed at 100 RWF and the lamp capacity Q at 18 hours. We now consider the counterfactual settings wherein the firm changes P and Q. These changes are assumed to be in the status quo business model, and so the decision process in these price/capacity-based counterfactuals is the same as in Figure 1. We rewrite expected recharges $R(\mathcal{P})$ and expected revenue $V(\mathcal{P})$, respectively, as $R(P,Q)$ and $V(P,Q)$ to make it explicit that we now view them as functions of price and capacity.

Set price optimally. We first fix the lamp capacity at $Q = 18$ hours and examine the impact of price changes. Unlike the strategies in Sections 5.1 and 5.2, which increase both R and V simultaneously, varying price has different impacts on R and V. Figure 8(a) shows the total number of recharges (left axis) and total revenue (right axis) expected at different price levels. The number of recharges falls steeply as we increase the price. Revenue increases as price increases because of the relatively inelastic nature of recharges with respect to price.[18] (This is not an artefact of any modeling assumptions. All price elasticities calculated from the raw data in Figure 3(c) are also greater than -1.) Consequently, if the firm focuses solely on maximizing

[18] At a given Q, the derivative of $V(P,Q)$ with respect to P is $R(P,Q)(1+e(P\,;Q))$, which is positive if the price elasticity $e(P\,;Q) > -1$ (the condition for relative inelasticity). Although in Figure 8(a) we show revenue only for prices between 0 RWF and 120 RWF (which was our experimental range), we numerically verified that for prices up to 200 RWF and for capacity levels between 6 hours and 30 hours, the revenue continues to increase in price.

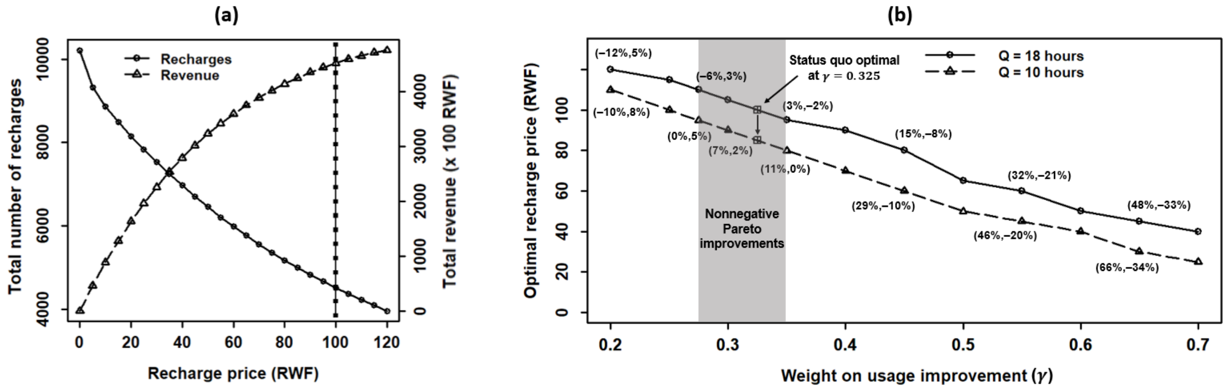

Figure 8 **Price- and capacity-based counterfactuals: (a) shows for $Q = 18$ hours, the expected number of recharges and expected revenue at different price levels. (b) presents optimal price $P^*(Q\,;\gamma)$ along with percentage improvements in usage and revenue for $Q = 18$ hours (status quo capacity) and $Q = 10$ hours (optimal capacity) at different values of γ.**

revenue, then the optimal price (i.e., $\arg\max_P V(P,Q)$) will not be an interior solution. By increasing price, the firm gains revenue, but it also loses the market penetration of its product.

However, several firms operating in the BoP markets are social enterprises with a dual objective that values both the usage of their products by poor consumers *and* revenue for the firms' sustenance, even when improving the former may hurt the latter. Assuming that the firm has a Cobb-Douglas-type preference for improvements in usage and revenue relative to the status quo, we model this dual objective as

$$W(P,Q\,;\gamma) = \left[\frac{R(P,Q)}{R(P_0,Q_0)}\right]^{\gamma}\left[\frac{V(P,Q)}{V(P_0,Q_0)}\right]^{1-\gamma},$$

where $(P_0,Q_0) = (100 \text{ RWF}, 18 \text{ hours})$ is the status quo, and $\gamma \in [0,1]$ is the weight given to usage relative to revenue. Depending on the financial status of the firm and the usage requirements imposed either by investors or the regulator, the firm may place different weights on usage at different points in time. Here, we exogenously vary the values of γ and examine the sensitivity of optimal gains in revenue and usage.

The optimal price $P^*(Q\,;\gamma) = \arg\max_P W(P,Q\,;\gamma)$ balances the two conflicting objectives of increasing both R and V at a given value of Q. The solid line in Figure 8(b) plots $P^*(18\,;\gamma)$ for different values of γ. It also presents the corresponding percentage improvements in usage and revenue. (For these values of γ, we see numerically that $W(P,Q\,;\gamma)$ is unimodal in P; thus, $P^*(Q\,;\gamma)$ exists and is unique.) When the firm places an equal weight on usage- and revenue-improvements (i.e., $\gamma = 0.5$), it is optimal to reduce the price to 65 RWF. This leads to an increase in recharges of 28% and a drop in revenue of 17%. As we move leftward from $\gamma = 0.5$, the firm becomes more revenue focused, making it optimal to *increase* the price. The reverse is true if we move rightward from $\gamma = 0.5$; the firm then becomes more usage focused and it becomes optimal to *lower* the price. In fact, the status quo price and capacity are optimal if the firm places a 32.5% weight on usage.

Set both price and capacity optimally. It is clear from Figure 8(a) and the solid line in Figure 8(b) that any movement from the status quo either decreases usage or revenue: the higher the increase in usage, the higher the corresponding drop in revenue, and vice versa. However, Pareto improvements along both the dimensions are plausible if the firm also varies capacity along with price. We define $Q^*(\gamma) = \arg\max_Q W(P^*(Q\,;\gamma), Q\,;\gamma)$. As we discussed under Π_3 in Section 4.2, as Q decreases, the number of recharges increase but only for relatively large values of capacity. Reducing Q to relatively smaller values results in too many trips and too much inconvenience for consumers, which in turn negatively affects usage. The optimal value $Q^*(\gamma)$ balances this trade-off. We find that $Q^*(\gamma) = 10$ hours for all $\gamma \in [0.2, 0.8]$.

The dashed line in Figure 8(b) plots $P^*(10\,;\gamma)$. As expected, the optimal price values are lower when capacity is reduced. We see in general that setting both price and capacity optimally leads to an increase in recharges that is higher than, and a drop in revenue that is lower than, the case wherein the firm sets only price optimally. Nonnegative Pareto improvements are indeed possible when the firm's γ is between 0.275 and 0.350. For the status quo weight on usage (i.e., $\gamma = 0.325$), we find that the optimal price is 85 RWF, and the resultant increases in usage and revenue are 7% and 2% respectively. Overall, the firm benefits by *reducing both price and capacity*. We next note two points regarding this finding.

First, Prahalad and Hart (2002) exemplify FMCG companies that have adapted some of their products (e.g., shampoo, tea, and cold medicines) for the BoP market by repackaging goods in smaller volumes to make them more affordable. Delivering light through a rechargeable lamp is equivalent to selling light in a small package. Our analysis suggests that – given consumer behavior in the market – it is better for the firm to sell light in a package of even smaller size by reducing price and capacity. Prahalad and Hart (2002) also argue that the consumers at the top of the pyramid (ToP) have enough disposable income, buy in bulk, and shop less frequently, i.e., they use their spending money to "inventory convenience", whereas the consumers at the BoP have limited cash, shop every day, and look for smaller packages. Although this may seem true on the surface, our analysis conclusively shows that (in)convenience is an important factor not only for the ToP consumers, but also for the BoP consumers, and that it must be considered when deciding on package size in BoP markets.

Second, the ratio of capacity and price at the status quo is $18/100 = 0.18$ hours/RWF, whereas it is $10/85 = 0.12$ hours/RWF at the optimum at $\gamma = 0.325$: the *bang for the buck* is lower at the reduced price and capacity levels. Because of the consumer's cash constraints, lowering the price increases her ability to pay for the recharges. But to generate enough revenue at the lowered price, the firm must also then substantially reduce capacity to induce a higher recharge frequency. In other words, the consumer pays a *poverty premium* when the light is provided in a smaller, more affordable package. Mendoza (2011) calls this the *size effect* – the penalty that the poor pay for being served in smaller portions. For a comprehensive discussion on poverty premia in other contexts, we refer the reader to Mendoza (2011) and Davies et al. (2016).

5.4. Discussion

Sections 5.1–5.3 evaluated the efficacy of several counterfactual policies that target liquidity constraints, inconvenience, recharge price, and lamp capacity; our analysis is summarized in Table 4. From now on, we refer to the inconvenience- and liquidity-based strategies together as *operations-based* strategies because they improve performance, not by varying the economic variables (namely, price and lamp capacity – the amount paid and the amount obtained in return) but by addressing the sources of inefficiencies (namely, liquidity

Counterfactual policy	Usage-improvement	Revenue-improvement
Price/Capacity-based		
Set P optimally, $\gamma: 0.2 \to 0.5$	$-12\% \to 28\%$	$5\% \to -17\%$
Set both P and Q optimally, $\gamma: 0.2 \to 0.5$	$-10\% \to 40\%$	$8\% \to -16\%$
Set both P and Q optimally, $\gamma = 0.325$	7%	2%
Liquidity-based		
Benchmark: no liquidity constraints	126%	
Allow partial (half) recharges	20%	
Allow prepayments	35%–43%	
Allow 1–2 recharges on credit	76%–84%	
Inconvenience-based		
Benchmark: {set $I = 0$, visit door-to-door every period}	43%, 73%	
Start 2–3 more recharge centers	23%–26%	
Visit door-to-door once in {1 week, 2 weeks}	34%, 15%	
Visit 5 locations once in {every period, 1 week, 2 weeks}	32%, 17%, 7%	
Encouraging consumers to pool	2%–10%	

Table 4 Summary of the performance of counterfactual policies.

constraints and inconvenience) in the current business model and by changing the recharge and payment processes within the firm.

We discuss three important implications of our analysis and findings. First, we see from Table 4 that *simple strategies can achieve good performance*. For example, just by allowing the consumers to recharge on credit 1–2 times or by starting 2–3 more recharge centers/drop-off points per village, the firm can reap half the benefits from *completely* removing liquidity constraints and inconvenience from the business model. It is important to recognize that while evaluating a counterfactual strategy, we have incorporated only those structural changes in the model that are required by that strategy. That way, we isolate the impact of implementing a particular strategy. Therefore, when combined together, the strategies in Table 4 may result in an even stronger impact on lamp usage and the firm's revenue. Analyzing the performance of any combinatorial strategy is a straightforward extension of the analysis conducted in our paper.

Moreover, Banerjee et al. (2017) note that in "all external decisionmaking problems, inference is unavoidably subjective. In structural modeling, the source of subjectivity is the model itself". Admittedly, the decision models in counterfactual settings need not coincide with the ones that we assumed in our paper. Nevertheless, because our process of external extrapolation is transparent, one can further enrich the analysis by introducing additional, plausibly subjective, components to those decision models. For example, in the periodic-visit model, the consumer may not always be at home when the VLE visits to collect her lamp, which may in turn affect the performance of that strategy; one can augment the decision process in Figure 7(b) with a probabilistic node representing the presence of consumer at home; the corresponding probability cannot be estimated from the data that we have, and hence it must be introduced either through subjective beliefs or through market research. Such extensions are also straightforward.

Second, *operations-based strategies tend to perform better than price/capacity-based strategies.* To increase the penetration of its product, a firm operating in the BoP market may have a natural tendency to reduce the price because (economic) poverty, by definition, entails a lack of money, and hence the inability to pay a higher price. Furthermore, several blog posts by entrepreneurs, technical reports by policy organizations (e.g., Bates et al. 2012, Girardeau and Pattanayak 2018), and much of the development economics literature cited in Section 1 emphasize pricing strategies, thereby making them more salient. Of course, price is an important lever in determining product adoptions under poverty conditions, but as we saw in Section 5.3, reducing price in our context simply decreases revenue. Therefore, operating at a lower price without the support from either donors or investors to fulfill any financial deficit (due to subsidies) may not be a sustainable strategy in the long run. Changing capacity along with price improves matters, but the resulting smaller packaging makes consumers pay a poverty premium, and the improvements are also quite small scale. In contrast, operations-based strategies increase usage and revenue simultaneously without consumers paying a poverty premium. This observation underscores the importance of removing, to the extent possible, the inefficiencies embedded in the business model by design – e.g., constraints such as making the consumer travel to a single village-level recharge center and allowing her to pay only when she recharges her lamp – that critically interact with the consumer behavior and limit product adoption.

We acknowledge that the firm may need to incur some costs to implement the operations-based strategies, and it is a limitation that we do not consider costs in this paper. Some strategies, such as encouraging pooling and allowing half-recharges, can be implemented without incurring any significant costs on top of the status quo, while some other strategies may involve both fixed and variable costs: To implement flexible payment schemes, the primary requirement is to collaborate with a mobile payment provider and build the necessary software that integrates the payment options that the firm desires to provide. Depending on the contract between the firm and the provider, the software costs incurred by the firm could be either fixed or variable. Starting new recharge centers, on the one hand, entails hiring new VLEs and installing recharge equipment for them – which are mainly fixed costs incurred at the beginning, and on the other hand, necessitates adjustments to the incentives provided to all the VLEs because of reduced market share per VLE – which are mainly variable costs. The periodic-visit models require a firm's representative to regularly travel to households in a village, and hence the costs involved are mostly variable and will depend on whether the firm uses its own VLEs or it piggybacks on the distribution network of other firms. Since these costs of implementing counterfactuals are context-specific and depend on various factors, they, unlike revenues, are difficult to estimate in a theoretically consistent manner. Nonetheless, the revenue-improvement estimates in Table 4 give upper bounds on costs for different strategies to be sustainable, which in turn can guide firm's cost-investment efforts. A detailed analysis of implementation costs is out of the scope of the current paper.

Third, *our research template can facilitate better experimentation with strategies.* Several firms operating in the BoP markets are budding startups. Although there is some emphasis on setting up the right business model to deliver life-improving goods and services to BoP consumers (e.g., IFC 2012), there is no formal method developed to arrive at one. The lean startup philosophy (Ries 2011) advocates constant experimentation for rapid improvements in the business model. The practitioners can adopt the methodology in this paper for the purpose of generating hypotheses. After a firm has (i) a *minimum viable product* (MVP), (ii) a plausible theory on consumer behavior arising under MVP, and (iii) a model that formalizes that behavior, it can conduct experiments – with the MVP – consisting of the minimal set of treatment conditions that are required to estimate the model. Unlike the ToP markets, there is a dearth of reliable datasets in the BoP markets; therefore, the data arising from such experiments can be used for a variety of analyses, including structural estimation. The estimated model can then guide the firm on what to do next and what is expected in return, thereby resulting in hypotheses that are grounded in both theory and data, which can later be tested by further experimentation.

6. Concluding Remarks

In this paper, we carried out a rigorous analysis of consumer behavior and the operational inefficiencies that result under a rechargeable-lamp-based off-grid lighting model. Our work has implications for firms, policymakers, academics, and consumers. We build a model of consumer recharge behavior and estimate it using field data from Rwanda. The estimated model can serve as a decision support system to assess the potential revenue opportunities in any alternative strategies ex-ante. Such analysis can also guide the treatment conditions for future field experiments. Firms operating rechargeable lamp businesses in other countries can also collect data and fit it to our model and thereafter use it for their own decision making.

Although we do not analyze all off-grid solutions currently in the market, we delve deep in terms of analyzing the rechargeable lamp model. Such analysis can help policymakers in Rwanda and similar countries that want to scale up their off-grid connectivity to weigh rechargeable lamp technologies against the benefits and limitations of other technologies. For instance, although solar home systems largely remove inconvenience (a main source of inefficiency in rechargeable lamp technologies), they are unaffordable for poor consumers and hence need to be heavily subsidized either by the government or donor organizations. In contrast, it is much cheaper to get rechargeable lamps into the hands of consumers, but their usage is limited by the inconvenience of recharging. The current study estimates consumer usage under the latter model. By also estimating usage under the solar home systems model, and by accounting for the corresponding costs, policymakers can assess which model works better at offering light to consumers.

In terms of academic relevance, we contribute to the nascent literature that studies operational issues in poor countries. Our dynamic model of off-grid light consumption incorporates consumer inconvenience, income and consumption uncertainties, and liquidity constraints. We characterize the optimal solution of consumer's DP and examine some properties of our model. The theoretical analysis in the paper is used only to generate predictions that can be tested using the field data. Future research can further explore our base model and its counterfactual versions. For example, one could investigate the theoretically optimal method of providing off-grid light to consumers and the optimally flexible payment plans to alleviate liquidity constraints. One could also conduct experiments with more sophisticated treatment conditions, collect detailed data, and employ a more elaborate model to estimate the additional effects – such as the impacts of consuming alternative lighting sources and consumers' behavioral biases – that we could not because of our budgetary constraints and the difficulty of collecting data from remote villages.

Finally, better business models and energy policies that account for consumer inconvenience, liquidity constraints, and actual usage data should result in higher usage of cleaner and cheaper lighting sources, which in turn contributes to increasing consumer productivity, improving their health, alleviating poverty, and promoting economic growth.

Acknowledgments

We are grateful to (i) Nuru Energy, and in particular Sameer Hajee (CEO), for partnering on this investigation and discussions which helped us better understand the Rwandan context and challenges; (ii) Innovations for Poverty Action, especially Phillip Okull, for managing our experiments and surveys in Rwanda; (iii) Martine Visser and Manuel Barron, for helpful conversations, and for collaboration in coordinating data collection efforts for our respective projects; and (iv) the Environment and Energy Partnership, Grand Challenges Canada (Stars in Global Health), INSEAD's Emerging Markets Institute, INSEAD's Randomized Controlled Trials Lab, the INSEAD-Wharton Alliance, and International Growth Centre for funding our experiments.

We gratefully acknowledge the department editor, the associate editor, and three anonymous referees for their numerous comments that significantly improved this paper. For valuable feedback we thank Kashish Arora, Neil Bearden, Tian Heong Chan, Zhi Chen, Vivek Choudhary, Sameer Hasija, Saravanan Kesavan, Anup Walvekar, Sasa Zorc, and the seminar participants at the China India Consumer Insights Conference, Columbia University, Early-Career Sustainable OM Workshop, ESMT Berlin, Harvard University, Indian School of Business, INFORMS Annual Meetings, Innovations for Poverty Action Policy Workshop, INSEAD, MSOM Conference, Nanyang Technological University, POMS Annual Conference, University of California at Los Angeles, University of North Carolina, Singapore Management University, Singapore University of Technology and Design, and Stanford University.

References

Aguirregabiria, V., P. Mira. 2010. Dynamic discrete choice structural models: A survey. *Journal of Econometrics* **156**(1) 38–67.

Aklin, M., P. Bayer, S.P. Harish, J. Urpelainen. 2017. Does basic energy access generate socioeconomic benefits? A field experiment with off-grid solar power in India. *Science Advances* **3**(5).

Ashraf, N., J. Berry, J.M. Shapiro. 2010. Can Higher Prices Stimulate Product Use? Evidence from a Field Experiment in Zambia. *American Economic Review* **100**(5) 2383–2413.

Balasubramanian, K., D. Drake, D. Fearing. 2017. Inventory Management for Mobile Money Agents in the Developing World. Working paper, Harvard University.

Banerjee, A.V., S. Chassang, E. Snowberg. 2017. Decision Theoretic Approaches to Experiment Design and External Validity. A.V. Banerjee, E. Duflo, eds., *Handbook of Field Experiments*, vol. 1. 141–174.

Barboni, G. 2017. Repayment flexibility in microfinance contracts: Theory and experimental evidence on take up and selection. *Journal of Economic Behavior & Organization* **142** 425–450.

Barron, M., M. Torero. 2017. Household electrification and indoor air pollution. *Journal of Environmental Economics and Management* **86** 81–92.

Bates, M.A., R. Glennerster, K. Gumede, E. Duflo. 2012. The Price is Wrong. Special issue 4, Field Actions Science Reports.

Clarke, R.P., M. Barron, M. Visser. 2020. Short-run subsidies, take-up, and long-run demand for off-grid solar for the poor: Evidence from large-scale randomized trials in Rwanda. Efd discussion paper, University of Cape Town.

Clarke, R.P., S. Delecourt, M. Barron, M. Visser. 2022. Affirmative action, business performance, and welfare: Experimental evidence from a solar microenterprise program. Working paper, Harvard Business School.

Cohen, J., P. Dupas. 2010. Free Distribution or Cost-Sharing? Evidence from a Randomized Malaria Prevention Experiment. *The Quarterly Journal of Economics* **125**(1) 1–45.

Davies, S., A. Finney, Y. Hartfree. 2016. *Paying to be poor: Uncovering the scale and nature of the poverty premium*. University of Bristol, United Kingdom.

de Zegher, J.F., D.A. Iancu, E. Plambeck. 2018. Sustaining Rainforests and Smallholders by Eliminating Payment Delay in a Commodity Supply Chain–It takes a Village. Working paper, Stanford University.

Drake, D.F., S. Spinler. 2013. OM Forum–Sustainable Operations Management: An Enduring Stream or a Passing Fancy? *Manufacturing & Service Operations Management* **15**(4) 689–700.

Duflo, E., R. Hanna, S.P. Ryan. 2012. Incentives Work: Getting Teachers to Come to School. *American Economic Review* **102**(4) 1241–1278.

Duflo, E., M. Kremer, J. Robinson. 2011. Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya. *American Economic Review* **101**(6) 2350–2390.

Dupas, P. 2014a. Getting essential health products to their end users: Subsidize, but how much? *Science* **345**(6202) 1279–1281.

Dupas, P. 2014b. Short-Run Subsidies and Long-Run Adoption of New Health Products: Evidence From a Field Experiment. *Econometrica* **82**(1) 197–228.

Field, E., R. Pande, J. Papp, Y.J. Park. 2012. Repayment Flexibility Can Reduce Financial Stress: A Randomized Control Trial with Microfinance Clients in India. *PLoS ONE* **7**(9) 1–7.

Furukawa, C. 2014. Do Solar Lamps Help Children Study? Contrary Evidence from a Pilot Study in Uganda. *The Journal of Development Studies* **50**(2) 319–341.

Girardeau, H., S.K. Pattanayak. 2018. Household Solar Adoption in Low- and Middle-Income Countries: A Systematic Review. Discussion paper, Environment for Development.

Girotra, K., S. Netessine. 2013. OM Forum–Business Model Innovation for Sustainability. *Manufacturing & Service Operations Management* **15**(4) 537–544.

Grimm, M., A. Munyehirwe, J. Peters, M. Sievert. 2017. A First Step Up the Energy Ladder? Low Cost Solar Kits and Household's Welfare in Rural Rwanda. *World Bank Economic Review* **31**(3) 631–649.

GSMA. 2020. The Mobile Economy 2020. Tech. rep., GSMA Intelligence.

Guajardo, J.A. 2019. How Do Usage and Payment Behavior Interact in Rent-to-Own Business Models? Evidence from Developing Economies. *Production and Operations Management* **28**(11) 2808–2822.

Gui, L., C. Tang, S. Yin. 2019. Improving Microretailer and Consumer Welfare in Developing Economies: Replenishment Strategies and Market Entries. *Manufacturing & Service Operations Management* **21**(1) 231–250.

IEA. 2020. *World Energy Outlook*. OECD Publishing.

IFC. 2012. From Gap to Opportunity: Business Models for Scaling Energy Access. Tech. rep., Washington, D.C.: The World Bank.

Jonasson, J.O., S. Deo, J. Gallien. 2017. Improving HIV Early Infant Diagnosis Supply Chains in Sub-Saharan Africa: Models and Application to Mozambique. *Operations Research* **65**(6) 1479–1493.

Keane, M.P. 2010. Structural vs. atheoretic approaches to econometrics. *Journal of Econometrics* **156** 3–20.

Kleindorfer, P.R., K. Singhal, L.N. Van Wassenhove. 2005. Sustainable Operations Management. *Production and Operations Management* **14**(4) 482–492.

Kundu, A., K. Ramdas. 2019. Timely After-Sales Service and Technology Adoption: Evidence from the Off-Grid Solar Market in Uganda. Working paper, London Business School.

Laureti, C., M. Hamp. 2011. Innovative flexible products in microfinance. *Savings and Development* **35**(1) 97–129.

Lee, K., E. Miguel, C. Wolfram. 2020a. Does Household Electrification Supercharge Economic Development? *Journal of Economic Perspectives* **34**(1) 122–144.

Lee, K., E. Miguel, C. Wolfram. 2020b. Experimental Evidence on the Economics of Rural Electrification. *Journal of Political Economy* **128**(4) 1523–1565.

Mendoza, R.U. 2011. Why do the poor pay more? Exploring the poverty penalty concept. *Journal of International Development* **23**(1) 1–28.

Meredith, J., J. Robinson, S. Walker, B. Wydick. 2013. Keeping the doctor away: Experimental evidence on investment in preventative health products. *Journal of Development Economics* **105** 196–210.

Miller, R.A. 1984. Job Matching and Occupational Choice. *Journal of Political Economy* **92** 1086–1120.

O'Donoghue, T., M. Rabin. 1999. Doing It Now or Later. *American Economic Review* **89**(1) 103–124.

Pakes, A. 1986. Patents as Options: Some Estimates of the Value of Holding European Patent Stocks. *Econometrica* **54**(4) 755–784.

Plambeck, E.L. 2013. OM Forum–Operations Management Challenges for Some "Cleantech" Firms. *Manufacturing & Service Operations Management* **15**(4) 527–536.

Prahalad, C.K., S.L. Hart. 2002. The Fortune at the Bottom of the Pyramid. *Strategy+Business* **26**.

Ries, E. 2011. *The Lean Startup: How Today's Entrepreneurs Use Continuous Innovation to Create Radically Successful Businesses*. Crown Business, New York.

Rom, A., I. Günther, K. Harrison. 2017. The Economic Impact of Solar Lighting: Results from a randomised field experiment in rural Kenya. Working paper, ETH Zurich.

Rust, J. 1987. Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher. *Econometrica* **55**(5) 999–1033.

Rust, J. 1994. Structural Estimation of Markov Decision Processes. Elsevier, 3081–3143.

Stinebrickner, T.R. 2000. Serial Correlated Variables in Dynamic, Discrete Choice Models. *Journal of Applied Econometrics* **15**(6) 595–624.

Stojanovski, O., M. Thurber, G. Muwowo, F.A. Wolak, K. Harrison. 2018. Assessing Opportunities for Solar Lanterns to Improve Educational Outcomes in Rural Off-Grid Regions: Challenges and Lessons from a Randomized Controlled Trial. Working paper, Asian Development Bank Institute.

Thaler, R.H. 1985. Mental Accounting and Consumer Choice. *Marketing Science* **4**(3) 199–214.

Thaler, R.H. 1999. Mental Accounting Matters. *Journal of Behavioral Decision Making* **12**(3) 183–206.

Todd, P.E., K.I. Wolpin. 2006. Assessing the Impact of a School Subsidy Program in Mexico: Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility. *American Economic Review* **96**(5) 1384–1417.

Todd, P.E., K.I. Wolpin. 2010a. Ex Ante Evaluation of Social Programs. *Annales d'Économie et de Statistique* (91) 259–286.

Todd, P.E., K.I. Wolpin. 2010b. Structural Estimation and Policy Evaluation in Developing Countries. *Annual Review of Economics* **2**(1) 21–50.

Uppari, B.S., I. Popescu, S. Netessine. 2019. Selling Off-Grid Light to Liquidity-Constrained Consumers. *Manufacturing & Service Operations Management* **21**(2) 308–326.

USAID. 2018. *Power Africa Rwanda Fact Sheet*.

Wolpin, K.I. 1984. An Estimable Dynamic Stochastic Model of Fertility and Child Mortality. *Journal of Political Economy* **92**(5) 852–874.

Wolpin, K.I. 1987. Estimating a Structural Search Model: The Transition from School to Work. *Econometrica* **55**(4) 801–817.

World Bank. 2017. *Rwanda – Renewable Energy Fund Project (English)*. World Bank Group.

Appendix A: A Note on Variation in Inconvenience

Because no strategic choices were made either by the firm or by the VLEs while placing a recharge center in our experimental villages, we hypothesize that the inconvenience faced by a consumer may not be systematically related to any relevant observable characteristics of that consumer. To test this claim, we collected data on households from twelve villages with IDs $3, 4, 5, 6, 10, 12, 13, 14, 20, 24, 26$, and 27. This set of villages were randomly selected from the set of all sampled villages.

Our dependent variable is the distance (measured in meters) between the GPS coordinates (recorded in three dimensions) of a household in a village and the recharge center in that village. We regress this variable against some key characteristics of households and the economic activities of household members (listed in

Table 5). We use village-level fixed effects to account for the time-invariant unobservables at the village level. The results of our regression analysis are presented in Table 5. All p-values are greater than 0.05, and thus we find no significant relationship between the dependent and independent variables.

Variable	Coefficient	Standard Error	p-value
Number of children in the household (HH)	0.055	0.032	0.088
Maximum education level of members in the HH	−0.008	0.014	0.543
Number of rooms in the dwelling	−0.002	0.007	0.820
$\mathbb{1}$\{HH uses kerosene\}	0.076	0.073	0.294
$\mathbb{1}$\{HH uses flashlight\}	0.039	0.026	0.137
Number of HH members who own small businesses	−0.025	0.036	0.487
Number of HH members who own farming land	0.025	0.030	0.395
Number of HH members who have regularly-paid jobs	0.016	0.051	0.762
Number of HH members who are farm laborers	0.031	0.019	0.100
Number of HH members who engage in part-time jobs	−0.037	0.027	0.162

Table 5 Regression of distance (in meters) on consumer covariates.

A few things should be noted with regard to the regressors in Table 5. First, the education level is coded as 0 if the consumer has no schooling, 1 if she is educated up to primary level, 2 if she is educated up to secondary level, 3 if she completed secondary education, and 4 if she received some vocational training. Second, the variance inflation factors of all the covariates are below 1.5; thus, there is no multicollinearity problem in our regression specification. Third, we collected a variety of other variables, such as the number of females in the household, the ages of the household members, the materials used to build the dwelling, the assets owned by the household, the income earned by the household members from various sources, and the list of all the lighting sources used by the household. We tested several other regression specifications with these variables and found that they also do not have a significant relationship with the distance variable. Therefore, to be concise, and to avoid including too many variables in the regression, we here present results only with the key variables. Finally, although none of the regressors in Table 5 are significant at 0.05 level of significance, some variables have p-values close to 0.1. To see if these weak relationships have any impact on our results, we re-estimated our structural model by including the regressors in the expression for $\mathbb{S}_t$ in (4) and found that the results in Section 5 largely remain the same; especially, the ordering of the counterfactuals in terms of their performance remains unaffected.

Because of randomized assignment to the experimental conditions, we expect the variables P and Q to be neither (significantly) correlated with the variable I, nor with the covariates in Table 5. We validate this claim in Tables 6 and 7. Table 6 regresses the recharge price assigned to consumers on their covariates. We use the same set of covariates as in Table 5, we include village-level fixed effects, and we include the distance to the recharge center as an additional control. (The results remain same when the distance is excluded.) Table 7 regresses the variable $\mathbb{Q}$, which indicates whether the assigned lamp capacity is 18 hours (as opposed to 14 hours), on the same set of covariates as in Table 6. We use linear regression in Table 7 even though $\mathbb{Q}$ is a binary variable for the sake of consistency with Tables 5 and 6. Using logistic regression instead of linear regression does not affect the findings.

Variable	Coefficient	Standard Error	p-value
Distance to the recharge center	−0.253	3.512	0.943
Number of children in the household (HH)	−0.778	3.296	0.813
Maximum education level of members in the HH	0.543	1.412	0.700
Number of rooms in the dwelling	0.017	0.677	0.979
$\mathbb{1}\{$HH uses kerosene$\}$	1.075	7.400	0.885
$\mathbb{1}\{$HH uses flashlight$\}$	4.702	2.702	0.082
Number of HH members who own small businesses	−0.134	3.680	0.971
Number of HH members who own farming land	1.895	3.033	0.532
Number of HH members who have regularly-paid jobs	1.218	5.230	0.816
Number of HH members who are farm laborers	0.291	1.918	0.880
Number of HH members who engage in part-time jobs	−5.095	2.714	0.061

Table 6 Regression of assigned recharge price (in RWF) on consumer covariates.

Variable	Coefficient	Standard Error	p-value
Distance to the recharge center	−0.009	0.046	0.837
Number of children in the household (HH)	0.040	0.043	0.354
Maximum education level of members in the HH	0.013	0.018	0.489
Number of rooms in the dwelling	0.017	0.009	0.053
$\mathbb{1}\{$HH uses kerosene$\}$	0.099	0.097	0.307
$\mathbb{1}\{$HH uses flashlight$\}$	−0.010	0.035	0.769
Number of HH members who own small businesses	−0.007	0.048	0.887
Number of HH members who own farming land	−0.021	0.040	0.598
Number of HH members who have regularly-paid jobs	−0.017	0.068	0.798
Number of HH members who are farm laborers	−0.013	0.025	0.591
Number of HH members who engage in part-time jobs	−0.024	0.035	0.503

Table 7 Regression of assigned lamp capacity ($\mathbb{Q} = \mathbb{1}\{\text{capacity} = 18 \text{ hours}\}$) on consumer covariates.

We can see from Tables 6 and 7 that the randomized assignment was mostly successful. There are mildly significant relationships between (i) P and the number of household members who do part-time jobs and whether the household uses flashlights, and (ii) $\mathbb{Q}$ and the number of rooms in the household. However, we have no reason to believe that there is a theoretical basis for these relationships, and they exist because of imperfections in randomization. Nevertheless, we explicitly included the covariates in both the reduced-form and structural analyses for the sake of robustness and the results remain unaffected.

Appendix B: Testable Predictions and Reduced-Form Analysis

B.1. A Simpler Model and Testable Predictions

We are interested in understanding how inconvenience, recharge price and lamp capacity affect the expected number of recharges $\mathbb{E}\tilde{R}$ (defined in (1)) under our model to test the validity of those relationships vis-á-vis data. Although using Proposition 2, we can obtain the expression for $\mathbb{E}\tilde{R}$, analytically characterizing it is difficult because there are 2^T possible combinations of recharge sequences, and the time-varying thresholds along with the uncertainties in liquidity and consumption make analysis of recharge probabilities cumbersome.

Instead, we analyze a version of the model that incorporates all factors of interest, but in the simplest possible manner, and is amenable to formal analysis. To build this model, we assume that (i) the consumption time is deterministic, and upon recharge, the lamp lasts for exactly $q\,(\geq 1)$ periods; and (ii) the disposable income process is stochastic but i.i.d. over time periods, and the probability that the consumer has sufficient money for recharge in a period is given by $v_\perp(P)$. With these assumptions, the blackout-cost threshold given in Proposition 1 (now denoted simply as k_t) simplifies to

$$k_t = \alpha I - \sum_{i=1}^{\check{q}-1} \Big[v_\perp \mathbb{E} \min\{k_{t+i}, \tilde{\beta}\} + (1 - v_\perp)\beta \Big], \tag{11}$$

where $\check{q} = \min\{q, T - t + 1\}$. To understand the threshold structure, recall that k_t is the effective cost of recharging in period t, whereas $\tilde{\beta}_t$ is the cost of not recharging in that period. As we discussed in Section 4.2, the effective cost of recharging in a period must account for both the inconvenience cost and the potential cost savings from jumping q periods ahead. The cost that the consumer would have incurred in period $t+i$ (for $1 \le i \le q-1$, assuming no end of horizon) if she does not recharge in period t is equal to (i) the (expected) minimum of k_{t+i} and $\tilde{\beta}_{t+i}$ if she has sufficient money for recharge in period $t+i$ (and has the option to recharge); and (ii) $\mathbb{E}\tilde{\beta} = \beta$ if she does not have enough money in that period (and has no option to recharge). Since the probability of the former event is $v_\perp$ and of the latter is $1 - v_\perp$, the expected cost saving in an interim period is $v_\perp \mathbb{E} \min\{k_{t+i}, \tilde{\beta}\} + (1 - v_\perp)\beta$. These expected cost terms for the interim periods are deducted from αI in the expression for k_t in (11).

We simplify the model further by assuming that the thresholds are stationary (which is equivalent to having an infinite horizon) to obtain a simple expression for interarrival times. We denote the time-invariant threshold by k^*, which is the (unique) solution to the following equation:

$$k = \alpha I - (q-1)\big\{ v_\perp \mathbb{E} \min\{k, \tilde{\beta}\} + (1 - v_\perp)\beta \big\}. \tag{12}$$

To see why, we compare (12) with (11). The former is obtained by using the time-invariance property of the thresholds on the latter. We formalize the above arguments in the following result:

Lemma 3. *The sequence* $\{k_T, k_{T-1}, \ldots\}$ *is convergent. The limit of this sequence is* k^**.*

Now we characterize the expected IAT under this model. After $q-1$ periods of consumption, the consumer waits until she has sufficient money for recharge *and* her blackout cost is above the threshold. The probability of this event in any period is $v_\perp \bar{F}(k^*)$, and the wait time is geometrically distributed. Hence, the expected blackout time is the mean of this geometric distribution, and the expected interarrival time Ψ is given by

$$\underbrace{\Psi}_{\text{IAT}} = \underbrace{q-1}_{\text{CT}} + \underbrace{\underbrace{1/v_\perp}_{\text{BT}_\text{L}} \times \underbrace{1/\bar{F}(k^*)}_{\text{BT}_\text{S}}}_{\text{BT}}. \tag{13}$$

The above expression for IAT parallels the structural pattern presented in Figure 2. The time between successive recharges (IAT) is the sum of consumption time and blackout time. The blackout time has two components: one because of liquidity constraints (BT_L) and the other because of strategic behavior (BT_S). As evident from Figure 2, BT ends when both constraints $\mathbb{M}_t = 1$ and $\mathbb{S}_t = 1$ are satisfied together (i.e., $\mathbb{M}_t = 1 \wedge \mathbb{S}_t = 1$), and parallelly, BT is given by $\text{BT}_\text{L} \times \text{BT}_\text{S}$.

Because the recharges follow a renewal process in the above formulation with mean inter-renewal interval Ψ, by the elementary renewal theorem, the expected number of recharges $\mathcal{R}$ in a duration T that is large enough is given by $\mathcal{R} \approx T/\Psi$. The following result establishes the relationship between $\mathcal{R}$ and the variables I, P, and Q (which is proxied here by q). Although q is a discrete variable, in the result below, we treat it as a continuous variable (satisfying $q \ge 1$) to simplify the analysis.

Proposition 3. *If we assume that $v_\perp(P)$ is monotonically decreasing in P and that F has increasing hazard rate, then the following statements hold:*

(i) *$\mathcal{R}$ is decreasing in I.*

(ii) *$\mathcal{R}$ is decreasing in P.*

(iii) *There exists a threshold $\hat{I} \geq 0$ such that $\mathcal{R}$ is unimodal in q for $I \geq \hat{I}$, and it is decreasing in q for $I < \hat{I}$.*

The three subparts of Proposition 3 parallel the three predictions stated in Section 4.2. As we have only two lamp capacity conditions in our experiments, we restate Π_3 as follows:

(Π_3) The difference between the number of recharges at $Q = 18$ hours and at $Q = 14$ hours is (a) negative for low values of inconvenience and (b) positive for high values of inconvenience.

B.2. Testing the Model's Predictions

To test Π_1–Π_3, we use the data from nine experimental conditions (i.e., $(P, Q) \in \{\{0, 50, 60, 70, 80, 100, 120\} \times \{18\}\} \cup \{\{80, 100\} \times \{14\}\}$) and run the following Poisson regression:

$$\tilde{R}_{jv} \sim \text{Pois}(\lambda_{jv}), \quad \text{where} \quad \log \lambda_{jv} = a_0 + a_1 I_{jv} + a_2 P_{jv} + a_3(I_{jv})\mathbb{Q}_{jv} + e_v. \tag{14}$$

Here, subscript j corresponds to an individual consumer and subscript v to a village. $\tilde{R}_{jv}$ is the total number of recharges of consumer j in village v in the three months of the experimental duration, P_{jv} is the recharge price (in RWF) assigned to this consumer, I_{jv} is her inconvenience (in kilometers), $\mathbb{Q}_{jv}$ is a dummy variable indicating whether the consumer's lamp capacity is 18 hours (as opposed to 14 hours), and e_v indicates a village-level fixed effect. As per Π_1 and Π_2, we expect a_1 and a_2 to be negative. Because the impact of capacity on recharges depends on consumer inconvenience, the coefficient of $\mathbb{Q}_{jv}$ in (14) is a function of I_{jv}. According to Π_3, $a_3(I)$ should be negative for smaller values of I and positive for larger values of I.

The results of our regression analysis are presented in Table 8, wherein we analyze three specifications of (14). In specification I, function a_3 is assumed to be constant. The resultant regression model presents the average effects of variables I, P, and $\mathbb{Q}$ in the sample. The coefficients of I and P in specification I are negative and statistically significant. The coefficient of $\mathbb{Q}$ is positive yet lacks significance, perhaps because the capacity conditions used in the experiments were 14 hours and 18 hours, which are not far apart from each other. The economic interpretation of these coefficients is as follows: a 10 RWF increase in recharge price, all else equal, decreases the expected number of recharges by 9% ($= \exp(-0.0097 \times 10) - 1$). An equivalent decrease in recharges is obtained by increasing inconvenience by 165 meters (i.e., $-0.5866 \times 0.165 = -0.0097 \times 10$). This 9% decrease in the expected number of recharges can be compensated for by increasing the lamp capacity by 11.5 hours (because $\exp(0.0308 \times 11.5/4) - 1 = 9\%$).

In specification II, we set $a_3(I) = a_{30} + a_{31} I$. However, this specification is not informative to either support or reject the hypothesized structure of $a_3(I)$ as both the coefficients related to $\mathbb{Q}$ lack significance. This may be because of some nonlinearity in $a_3(I)$; therefore, we consider specification III, wherein we set $a_3(I) = a_{30} + a_{31} I^2$. Here, the coefficients of I and P have the same signs as in specification I. The sign of the coefficient of $\mathbb{Q}$ flips (but the coefficient lacks significance), and the interaction term is positive (and significant). We see that, under specification III, a 4-hour increase in lamp capacity results in a 5% *drop* in

Variable	(I)	(II)	(III)
I	-0.5866^{***}	-0.6159^{***}	-0.8374^{***}
	(0.0484)	(0.1019)	(0.0848)
P	-0.0097^{***}	-0.0097^{***}	-0.0097^{***}
	(0.0004)	(0.0004)	(0.0004)
$\mathbb{Q}$	0.0308	0.0133	-0.0482
	(0.0393)	(0.0661)	(0.0444)
$\mathbb{Q} \times I$		0.0366	
		(0.1117)	
$\mathbb{Q} \times I^2$			0.2507^{***}
			(0.0679)
N	1709	1709	1709
Pseudo-R^2	0.2596	0.2597	0.2622

Table 8 Impact of inconvenience, recharge price, and lamp capacity. Superscript '*' is used when $p \leq 0.001$ and no superscript is used when $p > 0.1$.**

the expected number of recharges for consumers living close to the recharge center, whereas it results in a 1% and 22% *rise* for consumers living, respectively, 500 meters and a kilometer away from the recharge center.

Together, these findings provide support for Π_1, Π_2, and Π_{3b}. To comment on Π_{3a} with some confidence, we resort to a subsample analysis, wherein we run the Poisson regression with $\log \lambda_{jv} = a_0 + a_1 I_{jv} + a_2 P_{jv} + a_3 \mathbb{Q}_{jv}$ over a subsample with $I < \mathcal{I}$, where $\mathcal{I} \in \{10, 20, 30, 50, 100\}$ meters. We find that the coefficient of $\mathbb{Q}$ is negative for all values of $\mathcal{I}$ and that it is statistically significant for smaller values of $\mathcal{I}$ (thereby providing support for Π_{3a}) and loses significance as $\mathcal{I}$ increases.[19] Our findings remain unaffected under negative binomial and zero-inflated Poisson regression specifications. (The details are available upon request.) We conclude that recharge data from the field provides support for Π_1–Π_3 (albeit mildly for Π_{3a}).

Finally, note that in (14) we did not include any consumer covariates, because as we discussed in Appendix A, the variables P, Q, and I are not systematically correlated with covariates. For the sake of robustness, we reran the regression analysis, but with only 12 villages mentioned in Appendix A for which the covariates data is available, and found that the inferences drawn in this section are unaffected.

Appendix C: Identifiability of Model Parameters

In this section, we present some intuitive arguments on the identification of parameters. We discuss which sources of variation in the data and which theoretical assumptions in the model drive the quantitative values of the parameter estimates. We first examine identification in simpler models wherein only a subset of parameters appears. If a parameter cannot be identified in simpler models, then nor can it be identified in more general models. Thereafter, we move to the general versions and examine whether the parameters identifiable in simpler models continue to be identifiable in the general ones.

Identifiability of $(\alpha, \beta, \sigma_\xi)$. The simplest model for understanding the identifiability of these parameters is the one that assumes no consumption uncertainty or liquidity constraints: the consumer always has

[19] The coefficients of $\mathbb{Q}$, along with the p-values in parentheses, for the aforementioned values of $\mathcal{I}$ are as follows: $-0.669(0.020), -0.702(0.014), -0.375(0.081), -0.177(0.328), -0.166(0.189)$.

sufficient money for the recharge and the lamp lasts for exactly $q\,(\geq 1)$ periods upon recharge. Using the normal distribution assumption of $\tilde{\xi}_t$, we can rewrite (3) as

$$C(t,\tilde{z}_t) = \min\left\{\alpha I + \bar{C}(t+q), \beta + \sigma_\xi \tilde{z}_t + \bar{C}(t+1)\right\}, \tag{15}$$

where $\bar{C}(t) = \int_{-\infty}^{\infty} C(t,z)\, d\Phi(z)$. We cannot identify σ_ξ because equation (15) is only identified up to scale; multiplying both sides by a positive constant does not change the recharge decision. This remains the case in more complex models too. Therefore, we follow a common standard in the discrete choice literature and normalize the variance of the error term to 1.

Nor does the recharge decision change in (15) when we increase αI and β by the same amount. Thus, at the individual level, we cannot identify α and β separately; we can only identify the difference $\alpha I - \beta$. Nevertheless, because we have recharge data for multiple consumers and the value of inconvenience I varies across them, we can estimate α using the variation in inconvenience across individuals. Assuming that the value of α is given, we investigate whether we can identify β at the individual level.

We assumed under this model that we know consumption time perfectly and that it is equal to q periods. To be consistent with this assumption, we also assume that all interarrival times of the focal consumer are greater than or equal to q periods; otherwise, this model assigns zero likelihood. Then, the BT is given by the difference between IAT and q. Because there are no liquidity constraints under this model, BT here is purely due to the strategic behavior of the consumer. Specifically, the blackout time here corresponds to the hitting time of the blackout cost, i.e., the time it takes (after the lamplight is consumed) for the realized blackout cost to go beyond the (time-varying) threshold. Since these times are purely a function of β, the variation in these blackout times can identify β.

To summarize, we cannot identify σ_ξ, so we normalize it to 1. Hence, α and β become unitless quantities. At the individual level, we cannot identify α and β separately. Using the variation in inconveniences across individuals, we can estimate α. Given a value of α, using the variation in blackout times of an individual, we can estimate β at the individual level. Therefore, β is overidentified when we jointly estimate α and β using the recharge data of multiple individuals. This means that we can identify more parameters that represent heterogeneity in β (e.g., through random effects).

Identifiability of $(\boldsymbol{\mu}, \boldsymbol{\sigma}, \boldsymbol{\rho})$. The blackout times in the model just discussed were purely due to the strategic behavior of the consumer. In this section, we discuss identification under a model that is the other extreme. We assume that after the lamp is discharged, the consumer recharges her lamp whenever she has sufficient money for a recharge. She is neither sensitive to inconvenience nor to blackouts and hence does not seek to balance out inconvenience and blackout costs. So in this case, BTs are purely due to liquidity constraints.

We first assume that $\rho \geq 0$ is given. Recall that the consumer's disposable income starts growing immediately after a recharge. To keep the discussion simple, we assume that she never has sufficient money for a recharge in the first period when her lamp is discharged. In that case, all hitting times are greater than q. (The arguments can be easily extended to the cases where this assumption is relaxed.) Consequently, BT is the time that it takes beyond the consumption time for the income process to hit the threshold P. As we argue next, the variation in these BTs forms the source of identification for μ and σ.

The variation in BTs helps us identify transition probabilities $v(t,1,0)$ (for some t) of the Markov chain. Moreover, the estimate of $v(t,1,0)$ is equal to the sample hazard rate, i.e., it is the proportion of instances in the sample where the income did not hit the threshold P in $t-1$ periods but hits it at t. We illustrate this point using the following example: assume that $q=1$, and that we observe the following hitting times: $\{3,2,3,4,2\}$. Then the corresponding likelihood function is $[v(2,0,0)v(3,1,0)][v(2,1,0)][v(2,0,0)v(3,1,0)][v(2,0,0)v(3,0,0)v(4,1,0)][v(2,1,0)]$. If we treat each $v(t,\mathbb{m},\mathbb{m}')$ as a separate variable, then by noting that $v(t,1,0)=1-v(t,0,0)$, we obtain the estimate $v(2,1,0)=2/5$ by maximizing the above likelihood with respect to $v(2,1,0)$. This estimate is exactly equal to the proportion of instances in which the income did not hit the threshold in one period but in two. Similarly, $v(3,1,0)=2/3$. We cannot identify $v(4,1,0)$ from the above hitting time data. However, since $v(2,1,0)$ and $v(3,1,0)$ are two distinct functions of μ and σ, we can identify them from the estimates of those transition probabilities. In a general setting, to identify μ and σ, we should be able to estimate at least two transition probabilities, and hence we need to observe at least three distinct hitting time values in the data with two of them occurring more than once.

The identification of μ and σ is also possible when we use the data from multiple consumers with at least two distinct price levels. What we need is multiple distinct equations, corresponding to multiple distinct estimates of transition probabilities, to estimate μ and σ. Because the transition probabilities are functions of recharge price too, we obtain distinct equations (corresponding to distinct price levels) in this case as well. This feature is important (i) when the Markov chain is either serially independent or stationary, wherein $v(t,\mathbb{m},\mathbb{m}')$ becomes independent of t, and (ii) when we estimate parameters in more complex models wherein BTs may not be purely due to liquidity constraints.

Now we argue that ρ is non-identifiable. ρ represents serial correlation in the income process. Therefore, its source of identification is the variation in the number of times the process transitions from 0 to 0, 0 to 1, 1 to 0, and 1 to 1 in the successive periods. To be able to estimate three parameters jointly, we need to observe at least three types of transition. However, because we assume a renewal structure for the income process, we never observe the last two types of transition. The first two types of transition are useful in estimating only μ and σ (for a given ρ).

In summary, for a given ρ, the variation in blackout times – either within an individual or across individuals facing different price levels – identifies μ and σ. We cannot identify ρ either at the individual level or at the aggregate level because of the assumed renewal structure; therefore, we treat it as a hyper-parameter.

Joint identifiability of (α,β) and (μ,σ). To understand the joint identifiability of these parameters, we combine the features of the two simplified models discussed above. Now the consumer is liquidity constrained and accounts for inconvenience–blackout trade-offs, yet her consumption time is deterministic (and equal to q periods). In addition, we assume that (i) σ_ξ is normalized to 1, (ii) ρ is exogenously specified, and (iii) all the observed IATs are greater than q such that BTs are interarrival times minus q.

We previously identified β and (μ,σ) separately at the individual level by attributing the variation in BTs either to strategic behavior or to liquidity constraints. Now, both the strategic behavior of the consumer and her liquidity constraints contribute to her blackout times; they cannot be disentangled at the individual level. The variation in recharge price across individuals plays an important role in disentangling the two components. We have recharge data for consumers who faced zero recharge price. The liquidity constraints play no role in the decision-making process of these consumers; their decision model reduces to the one that we discussed at the beginning of this section. Therefore, using the data of zero-price consumers, we can estimate α and β. Given these estimates of α and β, we can control for the strategic component of BTs for the consumers facing nonzero price. The residual variation in BTs – across consumers facing different price levels – can then be purely attributed to liquidity constraints, which then identifies μ and σ.

The above procedure of estimating parameters sequentially is mentioned only to demonstrate that the variation in recharge price allows us to identify $(\alpha, \beta, \mu, \sigma)$. This procedure, however, is inefficient because it only uses partial data to estimate each set of parameters. We can jointly estimate all four parameters by maximizing the likelihood for zero-price and nonzero-price consumers together.

Identifiability of λ along with other parameters. Until now, we have assumed that q is known, which allowed us to compute blackout times directly from the observed interarrival times. If q is uncertain, but the capacity Q and the parameter λ are known, then we can still identify all of $(\alpha, \beta, \mu, \sigma)$ because given a realization of q, we can disentangle CT from BT, and the variation in BTs identifies the parameters. However, if λ is also a parameter that needs to be estimated, then there is no source of variation that disentangles consumption and blackout times.

To identify λ, we need something that affects CT but not BT. The variation in lamp capacity serves this purpose. If we assume that λ is known, then all the aforesaid parameters are identified as a function of λ, and they can control for blackout times. The remaining variation in IATs across individuals with varying capacity levels is attributed to the variation in CTs, which thereby identifies λ.

Appendix D: Robustness to the Choice of Cross Validation Set

We chose the experimental condition ($P = 100$ RWF, $Q = 18$ hours) as the cross validation set in our estimation method described in Section 4.4. To avoid overfitting the parameters to the cross validation set while selecting the hyperparameters, we used a cumulative likelihood score that accounts for both the likelihood on the cross validation set as well as the likelihood on the training set. In this section, we demostrate that, indeed, this cumulative scoring rule makes the parameters robust to choice of cross validation set.

We re-estimated the model with ($P = 80$ RWF, $Q = 18$ hours) and ($P = 50$ RWF, $Q = 18$ hours) as the cross validation sets. The corresponding parameter estimates are, respectively, given in Tables 9 and 10. We see that there is no meaningful difference between the estimates in these tables and those in Table 2. Especially, the average parameter estimates presented in the last rows of Tables 2, 9, and 10 are quite close to each other and are statistically indistinguishable. (We verified that, without the cumulative scoring rule, the parameter estimates are sensitive to the cross validation set.)

Additionally, in Table 11, we replicate Table 3, the interpretable versions of the parameter estimates, but with the cross validation sets ($P = 80$ RWF, $Q = 18$ hours) and ($P = 50$ RWF, $Q = 18$ hours). For each village, the first row presents the estimates for the former, whereas the second row presents them for the latter. These values too remain unaffected. Therefore, we conclude that our set of results is not an artefact of choosing a particular condition as the cross validation set.

Village ID	$\hat{\delta}$	$\hat{\rho}$	$\hat{\alpha}$	$\hat{\mu}_{\beta}$	$\hat{\sigma}_{\beta}$	$\hat{\mu}$	$\hat{\sigma}$	$\hat{\lambda}$
1	0.0	0.30	1.3551	−1.2574	0.7843	4.2585	0.5954	0.0186
			(0.4764)	(0.2043)	(0.1290)	(2.1556)	(1.2349)	(0.0202)
2	0.0	0.20	0.3957	−0.2098	1.1708	3.4765	2.5316	0.0563
			(0.5121)	(0.3074)	(0.2144)	(0.3210)	(1.6101)	(0.0090)
3	0.5	0.00	2.4564	0.7218	0.9335	4.0336	1.1015	0.0430
			(0.7326)	(0.2594)	(0.1565)	(0.2165)	(0.4336)	(0.0053)
4	0.6	0.15	0.0341	−0.7500	0.3210	1.2145	4.0580	0.0166
			(0.1869)	(0.1944)	(0.1069)	(0.6591)	(2.6643)	(0.0121)
5	0.0	0.15	1.3038	0.0459	0.7546	3.9895	0.8248	0.0498
			(0.5103)	(0.3694)	(0.2079)	(0.2458)	(0.1635)	(0.0117)
6	0.1	0.00	1.0702	−0.4219	0.3990	3.9550	1.1695	0.0213
			(0.2708)	(0.2425)	(0.1015)	(0.3465)	(0.4826)	(0.0140)
7	0.0	0.10	0.1880	−0.6600	1.0534	3.7412	0.6212	0.0608
			(0.3777)	(0.1749)	(0.3068)	(0.0938)	(0.1352)	(0.0243)
8	0.1	0.20	2.6701	−0.2422	0.7758	2.8364	2.1065	0.0194
			(0.6797)	(0.2226)	(0.0633)	(0.4664)	(1.2372)	(0.0116)
9	0.9	0.00	1.6956	0.0842	0.5626	4.1476	1.8178	0.0441
			(0.5582)	(0.0935)	(0.0370)	(0.3791)	(1.4138)	(0.0106)
10	0.9	0.05	1.0079	0.3333	0.4646	3.8317	1.1139	0.0604
			(0.2484)	(0.1851)	(0.0920)	(0.2583)	(0.4378)	(0.0091)
11	0.7	0.10	0.5444	−1.4035	0.2615	6.0021	0.0731	0.0024
			(0.3369)	(0.2019)	(0.0754)	(5.9584)	(0.1369)	(0.0087)
12	0.0	0.20	1.1179	−0.4624	0.6100	3.4180	1.2026	0.0292
			(0.4534)	(0.7057)	(0.4019)	(0.2972)	(0.5585)	(0.0204)
13	0.9	0.05	2.4525	0.6890	0.6485	2.8313	2.5032	0.0532
			(0.6449)	(0.2496)	(0.0952)	(1.1386)	(2.4526)	(0.0083)
14	0.0	0.15	0.4456	−1.0641	0.3852	4.1832	0.1510	0.0047
			(0.2358)	(0.1202)	(0.0572)	(0.0201)	(0.0749)	(0.0103)
16	0.9	0.00	4.0804	0.8936	0.8489	0.5173	4.1735	0.0947
			(1.6165)	(0.4642)	(0.2410)	(1.2275)	(7.9160)	(0.0139)
20	0.3	0.15	0.0136	−0.1508	0.8599	3.5261	1.2971	0.0593
			(0.3752)	(0.1987)	(0.1850)	(0.1827)	(0.5302)	(0.0090)
22	0.0	0.20	0.2106	−0.9131	0.3606	3.6325	0.6982	0.0269
			(0.2928)	(0.1844)	(0.0879)	(0.1676)	(0.1866)	(0.0135)
24	0.0	0.25	0.4810	−0.5497	1.2813	2.6208	3.7861	0.0874
			(0.2610)	(0.1783)	(0.2970)	(0.5117)	(3.0533)	(0.0250)
25	0.0	0.40	0.5064	−1.0985	0.5622	4.2487	0.8990	0.0359
			(0.3161)	(0.1599)	(0.0682)	(0.3401)	(0.3491)	(0.0137)
26	0.1	0.05	1.3109	−0.7322	0.5296	4.4204	2.1258	0.0067
			(0.5089)	(0.2361)	(0.0906)	(0.7403)	(4.4340)	(0.0106)
27	0.0	0.15	0.1231	−1.7911	0.6870	4.1449	0.2288	0.0250
			(0.2348)	(0.1811)	(0.0861)	(0.1300)	(0.1290)	(0.0198)
28	0.1	0.00	0.9481	−0.7042	0.6235	2.8501	2.2108	0.0015
			(0.4921)	(0.1841)	(0.3923)	(1.7467)	(1.9770)	(0.0117)
Average	0.3	0.13	1.1096	−0.4383	0.6763	3.5400	1.6041	0.0371
	(0.4)	(0.11)	(1.0332)	(0.7061)	(0.2734)	(1.1263)	(1.2186)	0.0262)

Table 9 **Maximum likelihood estimates of the model parameters for all villages, along with their standard errors in parentheses when the cross validation set is ($P = 80$ RWF, $Q = 18$ hours).**

Appendix E: Alternative Model Specifications

In this section, we elaborate on the alternative model specifications discussed in Section 4.4. For any given model specification, we denote the sample-level sum of recharges simulated using that model in round n by

Village ID	$\hat{\delta}$	$\hat{\rho}$	$\hat{\alpha}$	$\hat{\mu}_\beta$	$\hat{\sigma}_\beta$	$\hat{\mu}$	$\hat{\sigma}$	$\hat{\lambda}$
1	0.0	0.35	1.2620	−1.4081	0.8064	4.1666	0.6940	0.0183
			(0.2839)	(0.2619)	(0.1303)	(1.8608)	(1.0647)	(0.0290)
2	0.0	0.20	0.1911	−0.3348	1.0669	3.5338	2.3191	0.0564
			(0.5454)	(0.3371)	(0.1781)	(0.3039)	(1.1580)	(0.0094)
3	0.5	0.00	2.2507	0.6667	0.8803	4.1180	0.9804	0.0454
			(0.4913)	(0.5489)	(0.1646)	(0.1597)	(0.3339)	(0.0075)
4	0.5	0.05	0.0339	−0.8910	0.3884	3.2131	1.4733	0.0089
			(0.1971)	(0.1835)	(0.1081)	(0.3386)	(1.1083)	(0.0116)
5	0.1	0.10	1.3744	0.0239	0.7402	4.1551	0.7769	0.0537
			(0.2714)	(0.0960)	(0.2625)	(0.1698)	(0.2617)	(0.0074)
6	0.2	0.00	0.9888	−0.4576	0.3447	3.8308	1.3029	0.0106
			(0.2478)	(0.2349)	(0.0983)	(0.5056)	(0.6790)	(0.0134)
7	0.0	0.10	0.0239	−0.6989	1.4981	3.6878	0.7086	0.0650
			(0.3509)	(0.1580)	(0.3594)	(0.2785)	(0.1617)	(0.0037)
8	0.2	0.20	2.2421	−0.3492	0.6002	1.8137	4.7102	0.0194
			(0.6585)	(0.2598)	(0.1744)	(0.4637)	(5.9266)	(0.0118)
9	0.9	0.00	1.6949	0.0755	0.5628	4.2540	1.7518	0.0531
			(0.3770)	(0.0944)	(0.0866)	(0.5216)	(1.7956)	(0.0237)
10	0.9	0.05	1.0826	0.3552	0.3782	3.5743	1.7692	0.0590
			(0.2240)	(0.1390)	(0.0648)	(0.0003)	(0.0003)	(0.0091)
11	0.8	0.10	0.5207	−1.4369	0.2417	8.1383	0.1730	0.0012
			(0.3284)	(0.2338)	(0.0763)	(0.0000)	(0.0000)	(0.0091)
12	0.0	0.20	0.7583	−0.4117	0.7140	3.3203	1.0937	0.0382
			(0.3635)	(0.4273)	(0.4338)	(0.3304)	(0.6590)	(0.0250)
13	0.8	0.05	2.0372	0.5750	0.6655	2.2000	3.4536	0.0497
			(0.6120)	(0.2481)	(0.0882)	(0.3184)	(1.6441)	(0.0079)
14	0.0	0.15	0.4001	−1.1173	0.4266	4.2020	0.1690	0.0070
			(0.2669)	(0.1414)	(0.0580)	(0.0732)	(0.0522)	(0.0115)
16	0.9	0.15	3.1975	0.6643	0.8119	2.3681	1.9770	0.0744
			(5.9806)	(1.9005)	(0.8324)	(0.8670)	(1.8361)	(0.0204)
20	0.3	0.15	0.1125	−0.0550	0.7896	3.5249	1.2497	0.0605
			(0.3923)	(0.2260)	(0.1351)	(0.1686)	(0.4398)	(0.0089)
22	0.0	0.20	0.8116	−0.7374	0.3010	3.7399	0.7981	0.0345
			(0.2990)	(0.1800)	(0.0757)	(0.1686)	(0.3073)	(0.0132)
24	0.0	0.20	0.3505	−0.6278	1.2068	2.7576	3.8403	0.0820
			(0.2401)	(0.2372)	(0.4686)	(0.2419)	(2.5014)	(0.0154)
25	0.0	0.40	0.5128	−1.1358	0.5202	4.2448	0.8225	0.0340
			(0.2609)	(0.1409)	(0.0600)	(0.2394)	(0.2395)	(0.0143)
26	0.0	0.05	1.6626	−0.6368	0.5665	4.4333	1.4866	0.0132
			(0.5033)	(0.2613)	(0.0889)	(0.7267)	(1.2303)	(0.0101)
27	0.0	0.15	0.0789	−1.8144	0.6794	4.1410	0.2321	0.0252
			(0.2306)	(0.1787)	(0.0854)	(0.1245)	(0.1308)	(0.0194)
28	0.3	0.00	0.7109	−0.7466	0.6840	3.4492	1.2255	0.0054
			(0.5095)	(0.1986)	(0.4060)	(0.6420)	(0.5182)	(0.0134)
Average	0.3	0.13	1.0135	−0.4772	0.6761	3.7667	1.5003	0.0371
	(0.4)	(0.11)	(0.8660)	(0.6906)	(0.3035)	(1.2110)	(1.1779)	(0.0243)

Table 10 Maximum likelihood estimates of the model parameters for all villages, along with their standard errors in parentheses when the cross validation set is ($P = 50$ RWF, $Q = 18$ hours).

$\hat{R}_n^{\mathfrak{s}} = \sum_{v\in\mathfrak{V}} \sum_{j\in\mathfrak{J}^{\mathfrak{s}}(v)} \hat{R}_{jv,n}^{\mathfrak{s}}$ for $\mathfrak{s} \in \{tr, cv, ts\}$. We represent the actual sum of recharges observed in the sample as $R^{\mathfrak{s}} = \sum_{v\in\mathfrak{V}} \sum_{j\in\mathfrak{J}^{\mathfrak{s}}(v)} R_{jv}^{\mathfrak{s}}$. Then, $\text{MAPE}^{\mathfrak{s}} = 1/N_s \times \sum_{n=1}^{N_s} \left|(R^{\mathfrak{s}} - \hat{R}_n^{\mathfrak{s}})/R^{\mathfrak{s}}\right|$ is the mean absolute percentage error of the model on set $\mathfrak{s}$, and $\hat{R}^{\mathfrak{s}} = 1/N_s \times \sum_{n=1}^{N_s} \hat{R}_n^{\mathfrak{s}}$ is the average number of recharges predicted by

Village ID	Weights on next week's costs	Avg. no. of days an 18-hour lamp lasts	Avg. no. of days to hit 100 RWF	Pr{Blackout cost $\geq$ Inconvenience cost}	Recharges recorded per household
1	$0.00 \to 0.00$	4.00	2.35	$0.00 \to 0.06 \to 0.44$	1.31
	$0.00 \to 0.00$	3.99	2.45	$0.00 \to 0.05 \to 0.42$	
27	$0.00 \to 0.00$	4.35	3.81	$0.00 \to 0.06 \to 0.41$	1.49
	$0.00 \to 0.00$	4.36	3.85	$0.00 \to 0.06 \to 0.40$	
11	$0.49 \to 0.24$	3.13	0.00	$0.01 \to 0.06 \to 0.16$	1.67
	$0.64 \to 0.41$	3.06	0.00	$0.01 \to 0.06 \to 0.14$	
4	$0.36 \to 0.13$	3.90	11.87	$0.06 \to 0.23 \to 0.49$	1.72
	$0.25 \to 0.06$	3.48	12.52	$0.03 \to 0.20 \to 0.50$	
8	$0.01 \to 0.00$	4.05	8.87	$0.00 \to 0.19 \to 0.76$	1.95
	$0.04 \to 0.00$	4.05	7.89	$0.01 \to 0.18 \to 0.64$	
28	$0.01 \to 0.00$	3.08	11.04	$0.01 \to 0.21 \to 0.69$	1.96
	$0.09 \to 0.01$	3.29	14.36	$0.01 \to 0.22 \to 0.75$	
16	$0.81 \to 0.66$	8.12	15.33	$0.00 \to 0.13 \to 0.69$	2.16
	$0.81 \to 0.66$	7.01	15.64	$0.00 \to 0.17 \to 0.75$	
6	$0.01 \to 0.00$	4.15	7.38	$0.02 \to 0.16 \to 0.45$	2.24
	$0.04 \to 0.00$	3.57	7.86	$0.03 \to 0.16 \to 0.41$	
24	$0.00 \to 0.00$	7.72	6.48	$0.00 \to 0.31 \to 0.98$	2.78
	$0.00 \to 0.00$	7.43	6.15	$0.00 \to 0.30 \to 0.97$	
22	$0.00 \to 0.00$	4.45	7.21	$0.03 \to 0.17 \to 0.44$	2.85
	$0.00 \to 0.00$	4.86	5.81	$0.04 \to 0.15 \to 0.35$	
25	$0.00 \to 0.00$	4.94	2.20	$0.00 \to 0.12 \to 0.49$	2.85
	$0.00 \to 0.00$	4.84	2.19	$0.00 \to 0.11 \to 0.44$	
26	$0.01 \to 0.00$	3.36	3.26	$0.01 \to 0.16 \to 0.55$	2.97
	$0.00 \to 0.00$	3.71	3.28	$0.01 \to 0.16 \to 0.58$	
14	$0.00 \to 0.00$	3.25	3.27	$0.02 \to 0.12 \to 0.36$	3.04
	$0.00 \to 0.00$	3.38	3.26	$0.01 \to 0.12 \to 0.39$	
12	$0.00 \to 0.00$	4.58	7.52	$0.01 \to 0.23 \to 0.71$	3.12
	$0.00 \to 0.00$	5.06	9.08	$0.01 \to 0.29 \to 0.84$	
7	$0.00 \to 0.00$	6.28	12.48	$0.00 \to 0.30 \to 0.95$	3.33
	$0.00 \to 0.00$	6.51	12.16	$0.00 \to 0.35 \to 1.00$	
9	$0.81 \to 0.66$	5.38	4.49	$0.01 \to 0.19 \to 0.63$	3.45
	$0.81 \to 0.66$	5.87	4.13	$0.01 \to 0.19 \to 0.63$	
10	$0.81 \to 0.66$	6.26	7.70	$0.04 \to 0.28 \to 0.67$	3.71
	$0.81 \to 0.66$	6.19	7.02	$0.06 \to 0.26 \to 0.58$	
13	$0.81 \to 0.66$	5.87	9.09	$0.02 \to 0.32 \to 0.83$	3.79
	$0.64 \to 0.41$	5.68	9.20	$0.02 \to 0.35 \to 0.86$	
5	$0.00 \to 0.00$	5.69	4.84	$0.01 \to 0.27 \to 0.84$	4.26
	$0.01 \to 0.00$	5.90	4.60	$0.01 \to 0.25 \to 0.81$	
3	$0.25 \to 0.06$	5.32	6.94	$0.01 \to 0.39 \to 0.96$	4.62
	$0.25 \to 0.06$	5.45	6.69	$0.01 \to 0.40 \to 0.95$	
20	$0.09 \to 0.01$	6.20	7.34	$0.02 \to 0.45 \to 0.97$	4.67
	$0.09 \to 0.01$	6.27	7.54	$0.03 \to 0.46 \to 0.96$	
2	$0.00 \to 0.00$	6.04	5.08	$0.00 \to 0.40 \to 0.99$	4.83
	$0.00 \to 0.00$	6.04	5.04	$0.00 \to 0.39 \to 0.98$	

Table 11 Interpretation of parameter estimates for all villages. The first and second rows for each village interpret the parameters when the cross validation sets are, respectively, ($P = 80$ RWF, $Q = 18$ hours) and ($P = 50$ RWF, $Q = 18$ hours). The second column presents $\hat{\delta}^2 \to \hat{\delta}^4$, the third column presents $\mathbb{E}[\tilde{N}\,;\,Q = 18]$, the fourth column presents $\mathbb{E}[\tilde{h}(100)]$, and the fifth column presents $\mathbb{E}_\beta[\bar{\Phi}(\hat{\alpha}I_v - \beta) \mid \beta \leq \hat{\mu}_\beta - 2\hat{\sigma}_\beta] \to \mathbb{E}_\beta[\bar{\Phi}(\hat{\alpha}I_v - \beta)] \to \mathbb{E}_\beta[\bar{\Phi}(\hat{\alpha}I_v - \beta) \mid \beta \geq \hat{\mu}_\beta + 2\hat{\sigma}_\beta]$, where I_v is the average inconvenience in village v.

the model on set $\mathfrak{s}$. Table 12 reports MAPE$^{\mathfrak{s}}$ and $\hat{R}^{\mathfrak{s}}$ along with their standard errors for all the model specifications that we examine.

		Training set		Cross validation set		Test set	
		Actual = 4640		Actual = 429		Actual = 560	
		Pred.	MAPE	Pred.	MAPE	Pred.	MAPE
Structural	S_1	4627.66	0.9%	463.13	8.1%	504.57	10.9%
		(50.58)	(0.7%)	(21.03)	(3.2%)	(23.48)	(3.9%)
	S_2	4556.92	1.8%	469.14	9.0%	476.84	15.9%
		(54.21)	(1.1%)	(22.67)	(4.1%)	(20.36)	(4.0%)
	S_3	4619.21	0.9%	457.91	7.1%	514.42	8.6%
		(48.72)	(0.7%)	(15.75)	(3.3%)	(15.71)	(2.3%)
	S_4	4600.72	1.3%	468.54	9.7%	506.52	9.8%
		(66.00)	(1.0%)	(19.92)	(4.3%)	(24.32)	(4.0%)
	S_5	4669.92	0.7%	456.43	7.2%	517.45	7.5%
		(25.21)	(0.4%)	(10.53)	(2.4%)	(10.93)	(2.1%)
Atheoretical	A_1	4600.58	1.2%	436.92	5.5%	451.29	19.4%
		(59.12)	(0.9%)	(23.81)	(3.0%)	(22.91)	(4.1%)
	A_2	4601.32	1.2%	440.08	5.3%	459.41	18.0%
		(55.03)	(0.8%)	(22.52)	(2.7%)	(23.37)	(4.2%)

Table 12 Goodness of fit of various model specifications. Actual = actual number of recharges; Pred. = predicted number of recharges; and MAPE = mean absolute percentage error.

In Table 12, specification S_1 is exactly that described in Sections 4.1–4.3: the consumers do not discount their future costs and the blackout cost β is homogenous. S_1 fits the data in the training and cross validation sets reasonably well. The out-of-sample MAPE on the test set is (on average) 10.9%.[20] However, the assumption made by S_1 that the consumers are forward-looking with a discount factor of $\delta = 1$ may be unreasonable. To understand the relevance of this assumption, we next investigate the performance under a specification that lies on the other extreme. S_2 assumes that the consumers are myopic with $\delta = 0$. The fit of S_2 to training and cross validation datasets is almost indistinguishable from that of S_1. However, the limitation of assuming myopia is evident from the poor performance of S_2 on the test set. Because every fourth recharge is free in the test set, a consumer's current recharge decision has an impact on her future recharge price, and hence in reality, consumers plausibly account for future costs while making the recharge decisions in such settings. Since S_2 rules out forward-looking behavior, it poorly predicts recharge decisions in the test set.

It is possible that the consumers are neither completely forward looking (as in S_1) nor completely myopic (as in S_2). Their discount factor δ could be between 0 and 1. Specification S_3 incorporates discounting in the Bellman equations of the DP model in Section 4.1. As we mentioned in Section 4.4, we treat δ as a hyperparameter and estimate it along with ρ using cross validation. The predictive ability of S_3 is better than that of both S_1 and S_2, thereby suggesting that our consumers partly discount future costs.

[20] In village v, we generate the recharge sequence $\hat{\mathfrak{r}}_{jv,n}$ for consumer j in simulation round n using (i) the decision process in Figure 1 for $j \in \mathfrak{J}^{tr}(v) \cup \mathfrak{J}^{cv}(v)$ and the decision process in Figure 9(a) for $j \in \mathfrak{J}^{ts}(v)$; (ii) the treatment condition of that consumer Γ_{jv}; and (iii) the probability models of $\mathbb{D}_t$, $\mathbb{M}_t$, and $\mathbb{S}_t$ along with their estimated parameters from $\hat{\Theta}(v)$.

Yet another way to account for partially forward-looking behavior is by assuming that the consumer in period t looks forward only $\mathfrak{T} < T - t$ periods, instead of $T - t$ periods (as in the fully forward-looking case) or 0 periods (as in the fully myopic case). We call this specification S_4, and we estimate it with $\mathfrak{T}$ as a hyperparameter. As the consumer's planning horizon changes in every period under S_4, she exhibits dynamically inconsistent present-biased preferences. In our implementation, we assume that the consumer is *naive* about her dynamic inconsistency (O'Donoghue and Rabin 1999). We see that S_4 does not improve upon S_3. This is because the discount rates estimated under our exponential discounting specifications are already quite small. (In Table 2, on average, the weight on the cost that is 3 periods ahead is $0.3^3 = 0.027$, which is minuscule.) Therefore, we see no benefits in artificially shrinking consumers' planning horizon. Moreover, S_4 not only introduces dynamic inconsistency – a theoretically unappealing feature, but it also precludes the usage of memoization techniques, thereby increasing computational complexity.

Building on S_3, in addition to village-level heterogeneity, there could be heterogeneity within a village across consumers. We incorporate this by assuming that the blackout cost is heterogenous (and model it as a random variable – i.e., a random effect); we call the corresponding specification S_5, which is the specification that is examined in detail in Section 4.4. We assume that β for a consumer is drawn from the distribution Normal$(\mu_\beta, \sigma_\beta^2)$ and estimate the parameters μ_β and σ_β. S_5 improves further on S_3 and has an out-of-sample MAPE of 7.5%. We estimate a few other specifications with heterogeneity incorporated in income and consumption processes as well; however, they do not improve upon the performance displayed by S_5 and are also computationally intensive (because of more random effect terms), so we do not consider them further.

All the above specifications are different versions of a basic theoretical structure that we impose on the data. To see if there are any benefits to assuming a theoretical structure, we now compare its performance with two atheoretical approaches that do not explicitly account for the underlying decision-making processes. In Table 12, A_1 is a Poisson regression model, similar to the one that we use in our reduced-form analysis in Appendix B. This model assumes that the recharges observed in the experimental duration are the realizations of a Poisson process, rather than being the realizations of the controlled decision processes in Figure 1 or Figure 9(a). The arrival rate of the Poisson process is estimated as a function of I, P, and Q using the training data. (The Poisson regression model is estimated separately for each village, and it thereby completely absorbs the village-level heterogeneity.) A_1 fits the training data well, and its performance on the cross validation set is slightly better than that of our structural models.

However, computing the performance of A_1 on the test set, where every fourth recharge is free, is not straightforward. One plausible way to model this counterfactual is by assuming that the recharges observed in the test set are the realizations of a non-homogenous Poisson process whose instantaneous arrival rate depends on the number of recharges done so far after the previous free recharge. As we see in Table 12, this model performs far worse on the test set than our structural models.[21]

[21] Under model A_1, the Poisson regression estimates the aggregate arrival rate of recharges $\Lambda(I, P, Q)$ over a duration of T periods as a function of I, P, and Q using the data from the training set. The per-period arrival rate, or approximately the probability of recharge in a period, is given by Λ/T. To simulate recharges for a consumer j in test set, we model the probability of recharge as $\Lambda(I_{jv}, P_{jv}, Q_{jv})/T$ in the periods where the upcoming recharge is not free and as $\Lambda(I_{jv}, 0, Q_{jv})/T$ in the periods where the upcoming recharge is free. In contrast, model A_2 directly estimates the probability of recharge as a function of I, P, and Q using logistic regression on the period-level data from the training set. The logic for simulating test-set recharges is the same as that of A_1.

An alternative atheoretical approach is to assume that the recharge decisions follow a Bernoulli process, which is equivalent to assuming that the consumer decides whether to recharge in a period by tossing a (plausibly unfair) coin. This model is called A_2. We determine the probability of recharge as a function of I, P, and Q using logistic regression on the training data. Similar to A_1, this model also fits the training and cross validation sets well but performs badly on the test set.

Several important points follow from the above analysis. Ignoring the consumer decision-making process may not hurt when the counterfactuals of interest are merely changes to the values of parameters (as in our cross validation set), because the predictions in such settings can be made by fitting and extrapolating curves in the parameter space (as in the Poission and logistic regression models). However, when a counterfactual involves structural changes to the interactions between firm and consumer (as in our test set), (i) we must be able to formally represent that counterfactual using the components that were estimated in the original model, and (ii) that representation must reflect the reality reasonably well. Although (i) was achieved in A_1 and A_2, their representations of the counterfactual ignores some crucial decision-making aspects (e.g., forward-looking behavior), which results in poor performance. As the counterfactuals become structurally more complex (see Section 5), even representing them becomes nearly impossible without imposing a theoretical structure. In such settings, structural models become the natural candidates for making predictions.

Appendix F: Bellman Equations

Here, we present the Bellman equations for all the counterfactuals discussed in the main text. We only write the equations for $t \le T$, and in all cases the cost $C(T+n;\cdot) = 0 \,\forall\, n \ge 1$. Moreover, we assume that the discount factor $\delta = 1$ in the equations; they can be easily extended to any arbitrary $\delta \in [0,1)$. For notational simplicity, we denote $\tau = t - l$ in the equations.

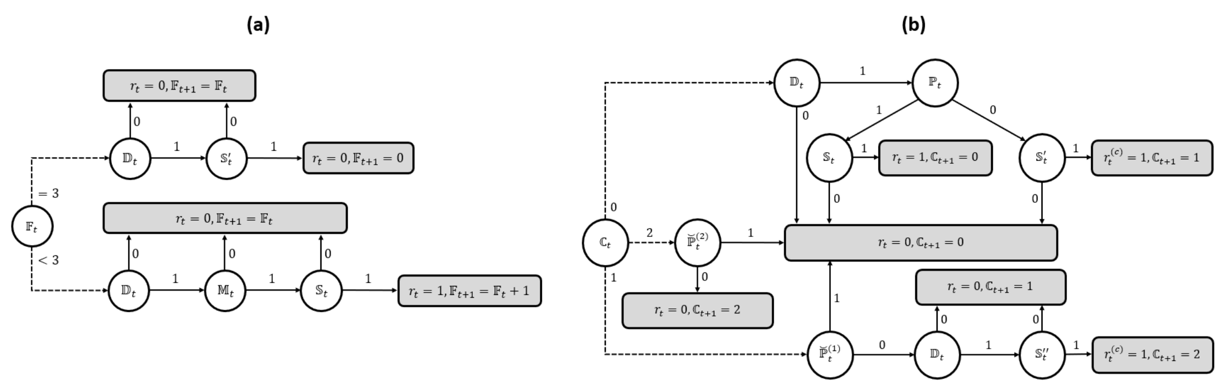

Figure 9 Decision process when (a) every fourth recharge is free, and (b) two recharges are allowed on credit.

Every fourth recharge free: The decision process when the firm offers every fourth recharge for free is shown in Figure 9(a). The variable $\mathbb{F}_t$ keeps track of the number of recharges done by the consumer *after* availing of the previous free recharge. Therefore, $\mathbb{F}_t \in \{0,1,2,3\}$. If $\mathbb{F}_t < 3$, then the decision process coincides with that in Figure 1. When $\mathbb{F}_t$ hits 3, the consumer becomes eligible for a free recharge, and hence her liquidity constraint disappears (i.e., $\mathbb{M}_t = 1$ until the next recharge). Before the free recharge is availed, $\mathbb{F}_t$

increments by 1 with every recharge, and $\mathbb{F}_t$ resets to 0 after the free recharge. The Bellman equations for this decision process are given below. The state space now includes an extra variable i to keep track of $\mathbb{F}_t$. Below, $i \in \{0,1\}$.

$$\begin{aligned}
C(t,b,0,l,i) &= b + v(\tau+1,0,0)\bar{C}(t+1,0,l,i) + v(\tau+1,1,0)\bar{C}(t+1,1,l,i),\\
C(t,b,1,l,i) &= \min\Big\{\alpha I + \sum_{q\in\mathcal{Q}} u_q\Big[v_q\bar{C}(t+q,1,t,i+1) + (1-v_q)\bar{C}(t+q,0,t,i+1)\Big],\\
&\qquad b + v(\tau+1,0,1)\bar{C}(t+1,0,l,i) + v(\tau+1,1,1)\bar{C}(t+1,1,l,i)\Big\},\\
C(t,b,0,l,2) &= b + v(\tau+1,0,0)\bar{C}(t+1,0,l,2) + v(\tau+1,1,0)\bar{C}(t+1,1,l,2),\\
C(t,b,1,l,2) &= \min\Big\{\alpha I + \sum_{q\in\mathcal{Q}} u_q\,\bar{C}(t+q,1,t,3),\\
&\qquad b + v(\tau+1,0,1)\bar{C}(t+1,0,l,2) + v(\tau+1,1,1)\bar{C}(t+1,1,l,2)\Big\},\\
C(t,b,0,l,3) &= b + \bar{C}(t+1,1,l,3),\\
C(t,b,1,l,3) &= \min\Big\{\alpha I + \sum_{q\in\mathcal{Q}} u_q\Big[v_q\bar{C}(t+q,1,t,0) + (1-v_q)\bar{C}(t+q,0,t,0)\Big], b + \bar{C}(t+1,1,l,3)\Big\}.
\end{aligned}$$

No liquidity constraints benchmark: Because the consumer does not experience any liquidity constraints, she need not keep track of her monetary status $\mathbb{m}$ or the last-recharge point l in the state space. The Bellman equations simplify to

$$C(t,b) = \min\Big\{\alpha I + \sum_{q\in\mathcal{Q}} u_q\,\bar{C}(t+q),\ b + \bar{C}(t+1)\Big\}.$$

Allowing partial recharges: Instead of $\mathbb{M}_t$, the consumer's cost function now keeps track of the state variable $(\mathbb{M}_t^{(0.5)}, \mathbb{M}_t^{(1)})$, whose possible values are $(0,0)$, $(1,0)$, and $(1,1)$. Analogous to $v(\tau,\mathbb{m},\mathbb{m}')$ and v_q, we define $w(\tau,(\mathbb{m}^{(0.5)},\mathbb{m}^{(1)}),(\mathbb{m}'^{(0.5)},\mathbb{m}'^{(1)}))$ and $w_q(\mathbb{m}^{(0.5)},\mathbb{m}^{(1)})$ as the probability transition functions for $(\mathbb{M}_t^{(0.5)}, \mathbb{M}_t^{(1)})$. The expressions for w and w_q can be derived from the empirical model for $\mathbb{M}_t$ discussed in Section 4.3, and are given as follows: $w_q(0,0) = G_q(\log P/2)$ and $w_q(1,0) = G_q(\log P) - G_q(\log P/2)$; $w(1,(0,0),\cdot) = G_1(\log P/2)$ and $w(1,(1,0),\cdot) = G_1(\log P) - G_1(\log P/2)$; for relative time period $\tau > 1$,

$$\begin{aligned}
w(\tau,(0,0),(0,0)) &= \frac{1}{G_{\tau-1}(\log P/2)}\int_{-\infty}^{\log P/2} \Phi\left(\frac{\log P/2 - \rho x - \mu}{\sigma}\right) dG_{\tau-1}(x),\\
w(\tau,(1,0),(0,0)) &= \frac{1}{G_{\tau-1}(\log P/2)}\int_{-\infty}^{\log P/2} \left\{\Phi\left(\frac{\log P - \rho x - \mu}{\sigma}\right) - \Phi\left(\frac{\log P/2 - \rho x - \mu}{\sigma}\right)\right\} dG_{\tau-1}(x),\\
w(\tau,(0,0),(1,0)) &= \frac{1}{G_{\tau-1}(\log P) - G_{\tau-1}(\log P/2)}\int_{\log P/2}^{\log P} \Phi\left(\frac{\log P/2 - \rho x - \mu}{\sigma}\right) dG_{\tau-1}(x),\\
w(\tau,(1,0),(1,0)) &= \frac{1}{G_{\tau-1}(\log P) - G_{\tau-1}(\log P/2)}\int_{\log P/2}^{\log P} \left\{\Phi\left(\frac{\log P - \rho x - \mu}{\sigma}\right)\right.\\
&\qquad\left. - \Phi\left(\frac{\log P/2 - \rho x - \mu}{\sigma}\right)\right\} dG_{\tau-1}(x),\\
w(\tau,(0,0),(1,1)) &= \frac{1}{\bar{G}_{\tau-1}(\log P)}\int_{\log P}^{\infty} \Phi\left(\frac{\log P/2 - \rho x - \mu}{\sigma}\right) dG_{\tau-1}(x),\\
w(\tau,(1,0),(1,1)) &= \frac{1}{\bar{G}_{\tau-1}(\log P)}\int_{\log P}^{\infty} \left\{\Phi\left(\frac{\log P - \rho x - \mu}{\sigma}\right) - \Phi\left(\frac{\log P/2 - \rho x - \mu}{\sigma}\right)\right\} dG_{\tau-1}(x).
\end{aligned}$$

In addition, we denote by $u_q^{(0.5)}$ the probability that a half-charged lamp lasts for q periods and $u_q^{(1)} = u_q$ as expressed in Section 4.3. Then, the Bellman equations are given as follows:

$$
\begin{aligned}
C(t,b,(0,0),l) = {} & b + w(\tau+1,(0,0),(0,0))\bar{C}(t+1,(0,0),l) + w(\tau+1,(1,0),(0,0))\bar{C}(t+1,(1,0),l) \\
& + w(\tau+1,(1,1),(0,0))\bar{C}(t+1,(1,1),l), \\
C(t,b,(1,0),l) = {} & \min\Big\{\alpha I + \sum_{q\in\mathcal{Q}} u_q^{(0.5)}\Big[w_q(0,0)\bar{C}(t+q,(0,0),t) + w_q(1,0)\bar{C}(t+q,(1,0),t) \\
& \quad + w_q(1,1)\bar{C}(t+q,(1,1),t)\Big], \\
& b + w(\tau+1,(0,0),(1,0))\bar{C}(t+1,(0,0),l) + w(\tau+1,(1,0),(1,0))\bar{C}(t+1,(1,0),l) \\
& \quad + w(\tau+1,(1,1),(1,0))\bar{C}(t+1,(1,1),l)\Big\}, \\
C(t,b,(1,1),l) = {} & \min\Big\{\alpha I + \sum_{q\in\mathcal{Q}} u_q^{(1)}\Big[w_q(0,0)\bar{C}(t+q,(0,0),t) + w_q(1,0)\bar{C}(t+q,(1,0),t) \\
& \quad + w_q(1,1)\bar{C}(t+q,(1,1),t)\Big], \\
& \alpha I + \sum_{q\in\mathcal{Q}} u_q^{(0.5)}\Big[w_q(0,0)\bar{C}(t+q,(0,0),t) + w_q(1,0)\bar{C}(t+q,(1,0),t) \\
& \quad + w_q(1,1)\bar{C}(t+q,(1,1),t)\Big], \\
& b + w(\tau+1,(0,0),(1,1))\bar{C}(t+1,(0,0),l) + w(\tau+1,(1,0),(1,1))\bar{C}(t+1,(1,0),l) \\
& \quad + w(\tau+1,(1,1),(1,1))\bar{C}(t+1,(1,1),l)\Big\}.
\end{aligned}
$$

Prepayment: The differences between the Bellman equations under Figure 1 and the prepayment model are: (i) v_q is replaced by $\breve{w}_q$, where $\breve{w}_q = \Pr(\mathbb{P}_t = 1 \mid \mathbb{P}_0 = 0) = \Pr(\tilde{h}(P) \leq q)$, i.e., the probability that the hitting time, as defined in (9), is less than or equal to q periods; and (ii) the second term within braces in the first equation below does not incorporate the possibility of losing the accrued money for the recharge, because once the income hits the threshold P, the payment is made using mobile money. Correspondingly the Bellman equations are given by

$$
\begin{aligned}
C(t,b,1,l) &= \min\Big\{\alpha I + \sum_{q\in\mathcal{Q}} u_q\big[\breve{w}_q\bar{C}(t+q,1,t) + (1-\breve{w}_q)\bar{C}(t+q,0,t)\big],\ b + \bar{C}(t+1,1,l)\Big\}, \\
C(t,b,0,l) &= b + v(t-l+1,1,0)\bar{C}(t+1,1,l) + v(t-l+1,0,0)\bar{C}(t+1,0,l).
\end{aligned}
$$

Prepayment after discharge: Unlike the prepayment model discussed above, the consumer here pays for the recharge only after the lamp is discharged; therefore, v_q is not replaced by $\breve{w}_q$ under this model. Consequently, the cost functions are written as

$$
\begin{aligned}
C(t,b,1,l) &= \min\Big\{\alpha I + \sum_{q\in\mathcal{Q}} u_q\big[v_q\bar{C}(t+q,1,t) + (1-v_q)\bar{C}(t+q,0,t)\big],\ b + \bar{C}(t+1,1,l)\Big\}, \\
C(t,b,0,l) &= b + v(t-l+1,1,0)\bar{C}(t+1,1,l) + v(t-l+1,0,0)\bar{C}(t+1,0,l).
\end{aligned}
$$

One recharge on credit: Here, the cost function is denoted as $C(t,b,\mathbb{d},l_r,l_p,s)$, where l_r is the last *recharge* point and l_p is the last *payment* point; because the recharges and payments are decoupled under this model, we need to keep track of l_r and l_p separately, which previously coincided with a single variable l in other settings. The additional state variable s indicates the debt status of the consumer: (i) $s=1$ indicates that there is a debt of one recharge, (ii) $s=0$ indicates that there is no debt and the money for the next recharge is not paid, and (iii) $s=0'$ indicates that there is no debt and the money for the next recharge is paid. The cost functions are given by

$$\begin{aligned}
C(t,b,1,l_r,l_p,1) &= b+v(t-l_p+1,0,0)\bar{C}(t+1,1,l_r,l_p,1)+v(t-l_p+1,1,0)\bar{C}(t+1,1,l_r,t+1,0),\\
C(t,b,0,l_r,l_p,1) &= \sum_{\mathbb{d}'\in\{0,1\}} u(t-l_r+1,\mathbb{d}',0)v(t-l_p+1,0,0)\bar{C}(t+1,\mathbb{d}',l_r,l_p,1)\\
&\qquad + \sum_{\mathbb{d}'\in\{0,1\}} u(t-l_r+1,\mathbb{d}',0)v(t-l_p+1,1,0)\bar{C}(t+1,\mathbb{d}',l_r,t+1,0),\\
C(t,b,1,l_r,l_p,0) &= \min\Big\{b+v(t-l_p+1,0,0)\bar{C}(t+1,1,l_r,l_p,0)+v(t-l_p+1,1,0)\bar{C}(t+1,1,l_r,t+1,0'),\\
&\qquad \alpha I+\sum_{\mathbb{d}'\in\{0,1\}} u(1,\mathbb{d}',0)v(t-l_p+1,0,0)\bar{C}(t+1,\mathbb{d}',t,l_p,1)\\
&\qquad + \sum_{\mathbb{d}'\in\{0,1\}} u(1,\mathbb{d}',0)v(t-l_p+1,1,0)\bar{C}(t+1,\mathbb{d}',t,t+1,0)\Big\},\\
C(t,b,0,l_r,l_p,0) &= \sum_{\mathbb{d}'\in\{0,1\}} u(t-l_r+1,\mathbb{d}',0)v(t-l_p+1,0,0)\bar{C}(t+1,\mathbb{d}',l_r,l_p,0)\\
&\qquad + \sum_{\mathbb{d}'\in\{0,1\}} u(t-l_r+1,\mathbb{d}',0)v(t-l_p+1,1,0)\bar{C}(t+1,\mathbb{d}',l_r,t+1,0'),\\
C(t,b,1,l_r,l_p,0') &= \min\Big\{\alpha I+\sum_{\mathbb{d}'\in\{0,1\}} u(1,\mathbb{d}',0)v(1,0,0)\bar{C}(t+1,\mathbb{d}',t,t,0)\\
&\qquad + \sum_{\mathbb{d}'\in\{0,1\}} u(1,\mathbb{d}',0)v(1,1,0)\bar{C}(t+1,\mathbb{d}',t,t+1,0'), b+\bar{C}(t+1,1,l_r,l_p,0')\Big\},\\
C(t,b,0,l_r,l_p,0') &= \sum_{\mathbb{d}'\in\{0,1\}} u(t-l_r+1,\mathbb{d}',0)\bar{C}(t+1,\mathbb{d}',l_r,l_p,0').
\end{aligned}$$

Two recharges on credit: The decision process when the consumer is allowed to recharge two times on credit is shown in Figure 9(b). As in the previous model, the cost function is denoted as $C(t,b,\mathbb{d},l_r,l_p,s)$, where l_r and l_s are the last recharge and payment points respectively, and s indicates the debt status of the consumer: (i) $s=2$ indicates that there is a debt of two recharges, (ii) $s=1'$ indicates that there is a debt of two recharges, of which one is paid for, (iii) $s=1$ indicates that there is a debt of one recharge, (iv) $s=0$ indicates that there is no debt and the money for the next recharge is not paid, and (v) $s=0'$ indicates that there is no debt and the money for the next recharge is paid. The corresponding Bellman equations are:

$$\begin{aligned}
C(t,b,1,l_r,l_p,2) &= b+v(t-l_p+1,0,0)\bar{C}(t+1,1,l_r,l_p,2)+v(t-l_p+1,1,0)\bar{C}(t+1,1,l_r,t+1,1'),\\
C(t,b,0,l_r,l_p,2) &= \sum_{\mathbb{d}'\in\{0,1\}} u(t-l_r+1,\mathbb{d}',0)v(t-l_p+1,0,0)\bar{C}(t+1,\mathbb{d}',l_r,l_p,2)\\
&\qquad + \sum_{\mathbb{d}'\in\{0,1\}} u(t-l_r+1,\mathbb{d}',0)v(t-l_p+1,1,0)\bar{C}(t+1,\mathbb{d}',l_r,t+1,1'),\\
C(t,b,1,l_r,l_p,1') &= b+v(t-l_p+1,0,0)\bar{C}(t+1,1,l_r,l_p,1')+v(t-l_p+1,1,0)\bar{C}(t+1,1,l_r,t+1,0),
\end{aligned}$$

$$
\begin{aligned}
C(t,b,0,l_r,l_p,1') = &\sum_{\mathbb{d}'\in\{0,1\}} u(t-l_r+1,\mathbb{d}',0)v(t-l_p+1,0,0)\bar{C}(t+1,\mathbb{d}',l_r,l_p,1')\\
&+\sum_{\mathbb{d}'\in\{0,1\}} u(t-l_r+1,\mathbb{d}',0)v(t-l_p+1,1,0)\bar{C}(t+1,\mathbb{d}',l_r,t+1,0),\\
C(t,b,1,l_r,l_p,1) = \min\Big\{&b+v(t-l_p+1,0,0)\bar{C}(t+1,1,l_r,l_p,1)+v(t-l_p+1,1,0)\bar{C}(t+1,1,l_r,t+1,0),\\
&\alpha I+\sum_{\mathbb{d}'\in\{0,1\}} u(1,\mathbb{d}',0)v(t-l_p+1,0,0)\bar{C}(t+1,\mathbb{d}',t,l_p,2)\\
&+\sum_{\mathbb{d}'\in\{0,1\}} u(1,\mathbb{d}',0)v(t-l_p+1,1,0)\bar{C}(t+1,\mathbb{d}',t,t+1,1')\Big\},\\
C(t,b,0,l_r,l_p,1) = &\sum_{\mathbb{d}'\in\{0,1\}} u(t-l_r+1,\mathbb{d}',0)v(t-l_p+1,0,0)\bar{C}(t+1,\mathbb{d}',l_r,l_p,1)\\
&+\sum_{\mathbb{d}'\in\{0,1\}} u(t-l_r+1,\mathbb{d}',0)v(t-l_p+1,1,0)\bar{C}(t+1,\mathbb{d}',l_r,t+1,0),\\
C(t,b,1,l_r,l_p,0) = \min\Big\{&b+v(t-l_p+1,0,0)\bar{C}(t+1,1,l_r,l_p,0)+v(t-l_p+1,1,0)\bar{C}(t+1,1,l_r,t+1,0'),\\
&\alpha I+\sum_{\mathbb{d}'\in\{0,1\}} u(1,\mathbb{d}',0)v(t-l_p+1,0,0)\bar{C}(t+1,\mathbb{d}',t,l_p,1)\\
&+\sum_{\mathbb{d}'\in\{0,1\}} u(1,\mathbb{d}',0)v(t-l_p+1,1,0)\bar{C}(t+1,\mathbb{d}',t,t+1,0)\Big\},\\
C(t,b,0,l_r,l_p,0) = &\sum_{\mathbb{d}'\in\{0,1\}} u(t-l_r+1,\mathbb{d}',0)v(t-l_p+1,0,0)\bar{C}(t+1,\mathbb{d}',l_r,l_p,0)\\
&+\sum_{\mathbb{d}'\in\{0,1\}} u(t-l_r+1,\mathbb{d}',0)v(t-l_p+1,1,0)\bar{C}(t+1,\mathbb{d}',l_r,t+1,0'),\\
C(t,b,1,l_r,l_p,0') = \min\Big\{&\alpha I+\sum_{\mathbb{d}'\in\{0,1\}} u(1,\mathbb{d}',0)v(1,0,0)\bar{C}(t+1,\mathbb{d}',t,t,0)\\
&+\sum_{\mathbb{d}'\in\{0,1\}} u(1,\mathbb{d}',0)v(1,1,0)\bar{C}(t+1,\mathbb{d}',t,t+1,0'),b+\bar{C}(t+1,1,l_r,l_p,0')\Big\},\\
C(t,b,0,l_r,l_p,0') = &\sum_{\mathbb{d}'\in\{0,1\}} u(t-l_r+1,\mathbb{d}',0)\bar{C}(t+1,\mathbb{d}',l_r,l_p,0').
\end{aligned}
$$

Periodic-visit model: The VLE visits once every n days. Instead of $\mathbb{V}_t$, we keep track of the variable k – the number of days left until the VLE's visit – in the state space. The cost function is given by $C(t,b,\mathbb{d},\mathbb{m},l,k)$, where $\mathbb{d}\in\{0,1\}$ is the lamp's discharge status in period t, and the other variables in the state space are as in the rest of the paper. In the equations below, $k\in\{1,\dots,n-1\}$.

$$
\begin{aligned}
C(t,b,0,\mathbb{m},l,k) = &\sum_{(\mathbb{d}',\mathbb{m}')\in\{0,1\}^2} u(\tau+1,\mathbb{d}',0)v(\tau+1,\mathbb{m}',\mathbb{m})\bar{C}(t+1,\mathbb{d}',\mathbb{m}',l,k-1),\\
C(t,b,1,0,l,k) = b+&\sum_{\mathbb{m}'\in\{0,1\}} v(\tau+1,\mathbb{m}',0)\bar{C}(t+1,1,\mathbb{m}',l,k-1),\\
C(t,b,1,1,l,k) = \min\Big\{&\alpha I+\sum_{(\mathbb{d}',\mathbb{m}')\in\{0,1\}^2} u(1,\mathbb{d}',0)v(1,\mathbb{m}',0)\bar{C}(t+1,\mathbb{d}',\mathbb{m}',t,k-1),\\
&b+\sum_{\mathbb{m}'\in\{0,1\}} v(\tau+1,\mathbb{m}',1)\bar{C}(t+1,1,\mathbb{m}',l,k-1)\Big\},\\
C(t,b,0,0,l,0) = &\sum_{(\mathbb{d}',\mathbb{m}')\in\{0,1\}^2} u(\tau+1,\mathbb{d}',0)v(\tau+1,\mathbb{m}',0)\bar{C}(t+1,\mathbb{d}',\mathbb{m}',l,n-1),\\
C(t,b,1,0,l,0) = b+&\sum_{\mathbb{m}'\in\{0,1\}} v(\tau+1,\mathbb{m}',0)\bar{C}(t+1,1,\mathbb{m}',l,n-1),
\end{aligned}
$$

$$C(t,b,0,1,l,0) = \min\Big\{ \sum_{(\mathbb{d}',\mathbb{m}')\in\{0,1\}^2} u(1,\mathbb{d}',0)v(1,\mathbb{m}',0)\bar{C}(t+1,\mathbb{d}',\mathbb{m}',t,n-1), \\ \sum_{(\mathbb{d}',\mathbb{m}')\in\{0,1\}^2} u(\tau+1,\mathbb{d}',0)v(\tau+1,\mathbb{m}',1)\bar{C}(t+1,\mathbb{d}',\mathbb{m}',l,n-1)\Big\},$$

$$C(t,b,1,1,l,0) = \min\Big\{ \sum_{(\mathbb{d}',\mathbb{m}')\in\{0,1\}^2} u(1,\mathbb{d}',0)v(1,\mathbb{m}',0)\bar{C}(t+1,\mathbb{d}',\mathbb{m}',t,n-1), \\ b + \sum_{\mathbb{m}'\in\{0,1\}} v(\tau+1,\mathbb{m}',1)\bar{C}(t+1,1,\mathbb{m}',l,n-1)\Big\}.$$

Encouraging consumers to pool: We consider a group of size ω. The state space for the group's DP constitutes (i) the current time period t, (ii) the cumulative number of lamps in the group that are discharged in period t, denoted as d, (iii) the cumulative number of group members who paid for the upcoming recharge in period t, denoted as z, (iv) the total blackout cost experienced by the members whose lamps are discharged, denoted as B, which is a realization of $\tilde{\beta}_{t,1} + \ldots + \tilde{\beta}_{t,d}$, (v) the last period l in which all the lamps in the group were recharged, and (vi) the index i of the member in the group who is responsible for recharge in period t. We denote by $\mathfrak{B}(n,r,p)$ the binomial probability density function, which gives the probability of observing r successes in n trials with a success probability of p.

If we assume the model of $\mathbb{P}_t$ in Figure 6(d), then z is always behind d. The Bellman equation when $d \le \omega$ and $z \le \omega$, but $(d,z) \neq (\omega,\omega)$ is given by

$$C(t,B,d,z,l,i) = B + \sum_{z_o=0}^{d-z} \mathfrak{B}(d-z, z_o, v(\tau+1,1,0)) \times \\ \sum_{d_e=0}^{\omega-d} \mathfrak{B}(\omega-d, d_e, u(\tau+1,1,0)) \sum_{z_e=0}^{d_e} \mathfrak{B}(d_e, z_e, v_{\tau+1})\ \bar{C}(t+1, d+d_e, z+z_o+z_e, l, i).$$

When either d or z is less than ω, the group does not have the option to recharge. It experiences blackout cost B and moves to the next state. If in the next period d_e lamps discharge, then the state d moves to $d+d_e$. Given that there are $\omega - d$ lamps yet to discharge, and the probability that a lamp discharges in relative time period $\tau+1$ is $u(\tau+1,1,0)$, the probability that d_e lamps discharge is given by $\mathfrak{B}(\omega-d, d_e, u(\tau+1,1,0))$. Of the d_e members whose lamps discharged, if z_e members have sufficient money for the recharge, then they transfer it to the coordinating member. The probability of this event is given by $\mathfrak{B}(d_e, z_e, v_{\tau+1})$ as the probability that a member has sufficient money in relative period $\tau+1$ is $v_{\tau+1}$. Additionally, of the $d-z$ members who did not already pay for the recharge, if z_o accrue sufficient money in period $\tau+1$ (with probability of success $v(\tau+1,1,0)$), then they also transfer their money to the coordinating member, and the probability of this event is $\mathfrak{B}(d-z, z_o, v(\tau+1,1,0))$. Accordingly, the state z then moves to $z+z_o+z_e$. In any period t, $\bar{C}(t,d,z,l,i) = \mathbb{E}C(t,\tilde{B},d,z,l,i)$; the expectation is taken with respect to the distribution of $\tilde{B}$, which is given the convolution of d i.i.d. $\tilde{\beta}_t$ random variables.

When $(d,z) = (\omega,\omega)$, the group has the option to recharge. The corresponding Bellman equation is

$$C(t,B,\omega,\omega,l,i) = \min\Big\{\alpha I_i + \sum_{d_e=0}^{\omega} \mathfrak{B}(\omega, d_e, u(1,1,\cdot)) \sum_{z_e=0}^{d_e} \mathfrak{B}(d_e, z_e, v_1)\ \bar{C}(t+1, d_e, z_e, t, (i+1) \bmod \omega), \\ B + \bar{C}(t+1,\omega,\omega,l,i)\Big\}.$$

If the group chooses to recharge, then because only the i^{th} member of the group travels to recharge all the lamps, only the inconvenience cost I_i is accounted for in the cost expressions. The continuation cost upon recharge is written using the same logic described in the previous paragraph. If the group chooses not to recharge, then because the coordinating member has already collected the required amount, the state (d,z) remains at (ω,ω). The index i remains the same if the lamps are not recharged; otherwise, it updates to $(i+1) \bmod \omega$, which is the index of the next member – in a circular group – responsible for lamp recharges.

Instead, if we assume the model of $\mathbb{P}_t$ in Figure 6(e), then the evolution of z is independent of the evolution of d. Following the same logic as in the previous case, the Bellman equations are written as

$$C(t,B,d,z,l,i) = B + \sum_{d_e=0}^{\omega-d} \mathfrak{B}(\omega-d,d_e,u(\tau+1,1,0)) \sum_{z_e=0}^{\omega-z} \mathfrak{B}(\omega-z,z_e,v(\tau+1,1,0))\ \bar{C}(t+1,d+d_e,z+z_e,l,i),$$

$$C(t,B,\omega,\omega,l,i) = \min\Big\{\alpha I_i + \sum_{d_e=0}^{\omega} \mathfrak{B}(\omega,d_e,u(1,1,\cdot)) \sum_{z_e=0}^{\omega} \mathfrak{B}(d_e,z_e,v(1,1,\cdot))\ \bar{C}(t+1,d_e,z_e,t,(i+1) \bmod \omega),$$
$$B + \bar{C}(t+1,\omega,\omega,l,i)\Big\}.$$

Appendix G: Proofs

Proof of Proposition 1. From (4), we know that the threshold $k(t,l)$ is given by

$$k(t,l) = \alpha I \ - \sum_{\mathbb{m}\in\{0,1\}} v(t-l+1,\mathbb{m},1)\,\bar{C}(t+1,\mathbb{m},l) \ + \sum_{q\in\mathcal{Q}} u_q\big[v_q\bar{C}(t+q,1,t)+(1-v_q)\bar{C}(t+q,0,t)\big]. \quad (16)$$

We then obtain the following expected cost functions using (3), (5), and (16):

$$\bar{C}(t,1,l) = \mathbb{E}\min\{k(t,l),\tilde{\beta}\} + v(t-l+1,1,1)\bar{C}(t+1,1,l) + v(t-l+1,0,1)\bar{C}(t+1,0,l), \quad (17)$$

$$\bar{C}(t,0,l) = \beta + v(t-l+1,1,0)\bar{C}(t+1,1,l) + v(t-l+1,0,0)\bar{C}(t+1,0,l). \quad (18)$$

If we define $\kappa(t,l,\mathbb{m}) = v(t-l,1,\mathbb{m})\bar{C}(t,1,l) + v(t-l,0,0)\bar{C}(t,0,l)$ and substitute it back in (16), (17), and (18), then we obtain the expression for the threshold as given in the statement of the proposition. □

Proof of Proposition 2. The following equalities are obtained from the definition of $\Omega(t,\mathbb{d},\mathbb{m})$. (We suppressed the arguments Θ and Γ in the probability expressions.)

$$\begin{aligned}
\Omega(t,\mathbb{d},\mathbb{m}) &= \Pr(\tilde{\mathfrak{r}}\langle t\rangle = \mathfrak{r}\langle t\rangle, \mathbb{D}_t=\mathbb{d}, \mathbb{M}_t=\mathbb{m}\,;\,\mathfrak{l}) \\
&= \sum_{\mathbb{d}'\in\{0,1\}} \sum_{\mathbb{m}'\in\{0,1\}} \Pr(\tilde{\mathfrak{r}}\langle t\rangle = \mathfrak{r}\langle t\rangle, \mathbb{D}_t=\mathbb{d}, \mathbb{M}_t=\mathbb{m}, \mathbb{D}_{t-1}=\mathbb{d}', \mathbb{M}_{t-1}=\mathbb{m}'\,;\,\mathfrak{l}) \\
&= \sum_{\mathbb{d}'\in\{0,1\}} \sum_{\mathbb{m}'\in\{0,1\}} \Pr(\tilde{\mathfrak{r}}\langle t-1\rangle = \mathfrak{r}\langle t-1\rangle, \mathbb{D}_{t-1}=\mathbb{d}', \mathbb{M}_{t-1}=\mathbb{m}'\,;\,\mathfrak{l}) \\
&\qquad \times \Pr(\tilde{r}_t = r_t, \mathbb{D}_t=\mathbb{d}, \mathbb{M}_t=\mathbb{m} \mid \tilde{\mathfrak{r}}\langle t-1\rangle = \mathfrak{r}\langle t-1\rangle, \mathbb{D}_{t-1}=\mathbb{d}', \mathbb{M}_{t-1}=\mathbb{m}'\,;\,\mathfrak{l}) \\
&= \sum_{\mathbb{d}'\in\{0,1\}} \sum_{\mathbb{m}'\in\{0,1\}} \Omega(t-1,\mathbb{d}',\mathbb{m}') \times \Pr(\mathbb{D}_t=\mathbb{d} \mid \mathbb{D}_{t-1}=\mathbb{d}'\,;\,l_t) \\
&\qquad \times \Pr(\mathbb{M}_t=\mathbb{m} \mid \mathbb{M}_{t-1}=\mathbb{m}'\,;\,l_t) \times \Pr(\tilde{r}_t = r_t \mid \mathbb{D}_t=\mathbb{d}, \mathbb{M}_t=\mathbb{m}\,;\,l_t) \\
&= \sum_{\mathbb{d}'\in\{0,1\}} \sum_{\mathbb{m}'\in\{0,1\}} \Omega(t-1,\mathbb{d}',\mathbb{m}') \times u(t-l_t,\mathbb{d},\mathbb{d}') \\
&\qquad \times v(t-l_t,\mathbb{m},\mathbb{m}') \times \big[\mathbb{d}\mathbb{m}\bar{F}(k(t,l_t))\big]^{r_t}\big[1-\mathbb{d}\mathbb{m}\bar{F}(k(t,l_t))\big]^{1-r_t}.
\end{aligned}$$

Here, the second equality follows from the law of total probability, the third equality from Bayes' law, and the fourth equality from recognizing that – conditional on $\mathfrak{l}$, other parameters and the previous period's states $\mathbb{d}'$ and $\mathbb{m}'$ – the evolution of processes $\mathbb{D}_t$ and $\mathbb{M}_t$ and the decision $\tilde{r}_t$ are independent of previous recharge decisions $\tilde{\mathfrak{r}}\langle t-1\rangle$. Finally, the realized recharge decision r_t is equal to 1 if and only if $\mathbb{d}=1$, $\mathbb{m}=1$, and the blackout cost is above the threshold $k(t,l_t)$; thus, $\Pr(\tilde{r}_t=1 \mid \mathbb{D}_t=\mathbb{d}, \mathbb{M}_t=\mathbb{m}\,; l_t)=\mathbb{d}\mathbb{m}\bar{F}(k(t,l_t))$. The expression of Ω for period $t=1$ can be obtained in a similar manner. □

Proof of Lemma 1. From the definition of m_τ, we obtain $m_\tau=\rho m_{\tau-1}+\epsilon_\tau=\epsilon_\tau+\rho\epsilon_{\tau-1}+\cdots+\rho^{\tau-1}\epsilon_1+m_0$. Since $\epsilon_t\sim \mathrm{N}(\mu,\sigma^2)$, the CDF of m_τ is given by G_τ as defined in the statement of Lemma 1. Then

$$
\begin{aligned}
v(\tau,1,0) &= \Pr(\mathbb{M}_\tau=1 \mid \mathbb{M}_{\tau-1}=0)=\Pr(m_\tau\geq \log P \mid m_{\tau-1}<\log P)\\
&= \Pr(m_\tau\geq \log P, m_{\tau-1}<\log P)/\Pr(m_{\tau-1}<\log P)\\
&= \Pr(\epsilon_\tau\geq \log P-\rho m_{\tau-1}, m_{\tau-1}<\log P)/\Pr(m_{\tau-1}<\log P)\\
&= \frac{1}{G_{\tau-1}(\log P)}\int_{-\infty}^{\log P}\Pr(\epsilon_\tau\geq \log P-\rho x)\,dG_{\tau-1}(x),
\end{aligned}
$$

which is the same as the expression in (6). The expression in (7) can be derived in a similar manner. Moreover, $v_q=\Pr(\mathbb{M}_q=1 \mid \mathbb{M}_0=0)=\Pr(m_q\geq \log P)=\bar{G}_q(\log P)$. □

Proof of Lemma 2. The result trivially follows from the definitions of u and u_q, and the assumption $\tilde{N}-1\sim \mathrm{Poisson}(Q\lambda)$. □

Proof of Lemma 3. We split the sequence $\{k_T,k_{T-1},\ldots\}$ as $\{k_T,\ldots,k_{T-q+1}\}\cup\{k_{T-q},k_{T-q-1},\ldots\}$. Since the first subsequence is of finite length, it suffices to show that the latter subsequence is convergent. We do that by showing that it is a bounded sequence. It then follows from the Bolzano-Weierstrass theorem that this sequence is convergent.

We first note that when $t\leq T-q$, we see from (11) that threshold k_t is a function of the next $q-1$ thresholds. We denote this function as ζ; hence,

$$
k_t=\zeta(k_{t+1},\ldots,k_{t+q-1})=\alpha I-\sum_{i=1}^{q-1}\big\{v_\perp\mathbb{E}\min\{k_{t+i},\tilde{\beta}\}+(1-v_\perp)\beta\big\}=\alpha I-(q-1)\beta+v_\perp\sum_{i=1}^{q-1}\mathbb{E}[\tilde{\beta}-k_{t+i}]^+.
$$

Because $\mathbb{E}\min\{k_{t+i},\tilde{\beta}\}\leq\beta$, $k_t\geq \alpha I-(q-1)\beta\equiv \underline{k}$, and because $\mathbb{E}[\tilde{\beta}-k_{t+i}]^+\leq \mathbb{E}[\tilde{\beta}-\underline{k}]^+$, $k_t\leq \underline{k}+v_\perp(q-1)\mathbb{E}[\tilde{\beta}-\underline{k}]^+\equiv\overline{k}$. Therefore, the sequence is bounded below by $\underline{k}$ and above by $\overline{k}$.

Given that the sequence is convergent, we denote its limit as $k_{-\infty}$. Then, because ζ is continuous in all its arguments, it follows that

$$
\begin{aligned}
k_{-\infty}=\lim_{n\to\infty}k_{T-q-n} &= \lim_{n\to\infty}\zeta(k_{T-q-n+1},\ldots,k_{T-q-n+q-1})\\
&= \zeta(\lim_{n\to\infty}k_{T-q-n+1},\ldots,\lim_{n\to\infty}k_{T-q-n+q-1})=\zeta(k_{-\infty},\ldots,k_{-\infty}).
\end{aligned}
$$

The fixed-point equation $k_{-\infty}=\zeta(k_{-\infty},\ldots,k_{-\infty})$ is the same as (12). Therefore, the limit of the sequence $k_{-\infty}$ is the same as the k^* that solves (12). It only remains to show that k^* exists and is unique. For that purpose, we define the following function:

$$
\mathfrak{K}(k,I,v_\perp,q)=k-\alpha I+(q-1)\big\{v_\perp\mathbb{E}\min\{k,\tilde{\beta}\}+(1-v_\perp)\beta\big\}, \tag{19}
$$

such that k^* is the solution to the implicit equation $\mathfrak{K}(k,\cdot)=0$. The function $\mathfrak{K}$ is increasing in k because $\partial\mathfrak{K}/\partial k = 1+(q-1)v_\perp \bar{F}(k)\geq 0$. Moreover, $\mathfrak{K}(\underline{k},\cdot) = -v_\perp(q-1)\mathbb{E}\big[\tilde{\beta}-\underline{k}\big]^+\leq 0$ and $\mathfrak{K}(\overline{k},\cdot)=v_\perp(q-1)\big\{\mathbb{E}\big[\tilde{\beta}-\underline{k}\big]^+ - \mathbb{E}\big[\tilde{\beta}-\overline{k}\big]^+\big\}\geq 0$. Thus, there exists a unique k^* that satisfies $\mathfrak{K}(k^*,\cdot)=0$. □

Proof of Proposition 3. By applying the implicit function theorem to (19), we obtain

$$\frac{\partial k^*}{\partial I} = -\frac{\partial\mathfrak{K}/\partial I}{\partial\mathfrak{K}/\partial k} = \frac{\alpha}{1+(q-1)v_\perp\bar{F}(k^*)}\geq 0, \tag{20}$$

$$\frac{\partial k^*}{\partial v_\perp} = -\frac{\partial\mathfrak{K}/\partial v_\perp}{\partial\mathfrak{K}/\partial k} = \frac{(q-1)\mathbb{E}[\tilde{\beta}-k^*]^+}{1+(q-1)v_\perp\bar{F}(k^*)}\geq 0, \tag{21}$$

$$\frac{\partial k^*}{\partial q} = -\frac{\partial\mathfrak{K}/\partial q}{\partial\mathfrak{K}/\partial k} = \frac{v_\perp\mathbb{E}[\tilde{\beta}-k^*]^+-\beta}{1+(q-1)v_\perp\bar{F}(k^*)} = \frac{k^*-\alpha I}{(q-1)(1+(q-1)v_\perp\bar{F}(k^*))}. \tag{22}$$

We see from (20) that $\bar{F}(k^*)$ is decreasing in I, and from (13) that Ψ is increasing in I; hence, $\mathcal{R}$ is decreasing in I. Next, using (21), we see that $v_\perp\bar{F}(k^*)$ is increasing in $v_\perp$:

$$\begin{aligned}\frac{\partial(v_\perp\bar{F}(k^*))}{\partial v_\perp} &= \bar{F}(k^*)-f(k^*)v_\perp\frac{(q-1)\mathbb{E}[\tilde{\beta}-k^*]^+}{1+(q-1)v_\perp\bar{F}(k^*)}\geq \bar{F}(k^*)-f(k^*)\frac{\mathbb{E}[\tilde{\beta}-k^*]^+}{\bar{F}(k^*)}\\ &= \int_{k^*}^{\infty}\bar{F}(s)\left[\frac{f(s)}{\bar{F}(s)}-\frac{f(k^*)}{\bar{F}(k^*)}\right]ds\geq 0.\end{aligned}$$

The last inequality is because the function $f/\bar{F}$ is increasing, where f is the density function of $\tilde{\beta}$. Thus, from (13), Ψ is decreasing in $v_\perp$. As $v_\perp$ is decreasing in P, $\mathcal{R}$ is also decreasing in P. To examine the behavior of Ψ wrt q, we note that

$$\frac{\partial\Psi}{\partial q} = 1+\frac{f(k^*)}{v_\perp\bar{F}(k^*)^2}\frac{\partial k^*}{\partial q} \tag{23}$$

$$= \left\{\frac{v_\perp\bar{F}(k^*)^2}{f(k^*)}(1+(q-1)v_\perp\bar{F}(k^*))-\beta+v_\perp\mathbb{E}[\tilde{\beta}-k^*]^+\right\}\times\frac{f(k^*)}{v_\perp\bar{F}(k^*)^2(1+(q-1)v_\perp\bar{F}(k^*)^2)}. \tag{24}$$

We see from (22) that whether k^* is increasing or decreasing in q depends on the sign of $k^*-\alpha I$. It is easy to verify from (20) that $k^*-\alpha I$ is decreasing in I and $\lim_{I\to\infty}k^*-\alpha I=-\infty$. Therefore, there exists a threshold $\hat{I}_1\geq 0$ such that for $I\geq\hat{I}_1$, k^* is decreasing in q and for $I<\hat{I}_1$, k^* is increasing in q. In the latter case, it follows from (23) that Ψ is increasing in q. Now it remains to examine what happens when k^* decreases in q.

The term outside the braces in (24) is positive; therefore, the sign of $\partial\Psi/\partial q$ depends only on the term inside the braces, which we denote as $\mathfrak{b}(q)$. We note that, when $I\geq\hat{I}_1$, $\mathfrak{b}(q)$ is increasing in q because (i) $\bar{F}$ is decreasing, (ii) $\bar{F}/f$ is decreasing, (iii) $\mathbb{E}[\tilde{\beta}-k]^+$ is decreasing in k, and (iv) k^* is decreasing in q. Moreover, $\lim_{q\to\infty}\mathfrak{b}(q)=\infty$ and $\lim_{q\to 1}\mathfrak{b}(q)=v_\perp\bar{F}(\alpha I)^2/f(\alpha I)-\beta+v_\perp\mathbb{E}[\tilde{\beta}-\alpha I]^+\equiv\mathfrak{z}(I)$. Furthermore, $\mathfrak{z}(I)$ is decreasing in I and $\lim_{I\to\infty}\mathfrak{z}(I)=-\beta<0$ (because when k^* is decreasing in q, from (22), $\beta > v_\perp\mathbb{E}[\tilde{\beta}-k^*]^+>0$). It follows that there exists a threshold $\hat{I}_2\geq 0$ such that (a) for $I\geq\hat{I}_2$, $\mathfrak{z}(I)=\lim_{q\to 1}\mathfrak{b}(q)$ is negative, and hence $\partial\Psi/\partial q$ single crosses the horizontal axis from below. In other words, Ψ is U-shaped in q, and $\mathcal{R}$ is unimodal in q, and (b) for $I<\hat{I}_2$, $\mathfrak{z}(I)=\lim_{q\to 1}\mathfrak{b}(q)$ is positive, Ψ is increasing in q and $\mathcal{R}$ is decreasing in q.

Overall, since $\hat{I}_2\geq\hat{I}_1\geq 0$, we define the threshold $\hat{I}=\hat{I}_2$, and conclude that $\mathcal{R}$ is unimodal in q when $I\geq\hat{I}$ and $\mathcal{R}$ is decreasing in q when $I<\hat{I}$. □