Optimal Second-best Menu Design: Evidence from Residential Electricity Plans

Maria Garcia-Osipenko, Nicolai V. Kuminoff,

Spencer Perry, Nicholas Vreugdenhil*†

March 24, 2025

Abstract

Utilities increasingly sell electricity using complex menus of static and dynamic price schedules. We study how to design such a menu to maximize social welfare in a second-best environment where the marginal private and external costs of generating electricity vary over time, institutional constraints prevent mandating dynamic pricing, and consumer behavior is distorted by frictions. We develop a model of plan choice, consumption, and intertemporal substitution with time-varying marginal social costs, and estimate it using administrative data from a large utility. We provide evidence of substantial intertemporal substitution in response to dynamic price incentives, and selection across plans based on multidimensional heterogeneity. While the current menu's dynamic plans substantially shift consumption from high-price to low-price hours, we find that they reduce social welfare. This loss is mitigated by information frictions. We show how to redesign the menu to simultaneously improve outcomes for consumers, the utility, and the environment.

*Maria Garcia-Osipenko: Arizona State University. Nicolai V. Kuminoff: Arizona State University & NBER. Spencer Perry: University of Alaska, Anchorage. Nicholas Vreugdenhil, Arizona State University.

†This work received funding from the ASU-SRP Joint Research Project. We are grateful to the Salt River Project for sharing data and to Jed Cohen and Geno Malaty for helpful comments and suggestions. We are also grateful for comments and suggestions from Ulrich Wagner, Riley League, and audiences at the AERE Summer Meeting, AERNA conference, UC Berkeley Energy Camp, Arizona State University, Triangle Resource and Environmental Economics Seminar, University of Lausanne, and University of Mannheim / University of Heidelberg / ZEW.

1 Introduction

Electricity markets are inefficient. Consumers often pay a static price throughout the day as the marginal cost of generating electricity fluctuates. This inefficiency is compounded by externalities. For example, electricity generation accounts for 25% of US carbon emissions (EPA, 2024). A Pigouvian solution would set the real-time price of electricity to its marginal social cost of generation. However, consumers tend to be inattentive to frequent price changes (Fabra et al., 2021). Moreover, mandating real-time pricing is typically politically infeasible.

Many utilities have responded to these constraints through a form of second-degree price discrimination. They offer consumers a choice from a menu of static price plans and simple dynamic 'time-of-use' plans that charge a higher price during a predetermined interval when costs are usually high. Selling electricity using menus is widespread with half of the twenty largest US utilities — which serve approximately 50 million households comprising nearly 40% of the population — offering menus with multiple static and time-of-use plans.

In theory, when consumer types are private information, offering a menu can potentially match consumers to their most efficient plans through self-selection. For consumers who select time-of-use plans, within-day price changes provide incentives to 'load-shift' consumption from high-cost times to low-cost times, and to 'load-shave' by reducing consumption altogether. Despite the theoretical promise of this approach, little is known about how to design an efficient menu with time-varying social costs in practice. A key challenge is to anticipate how menu efficiency will be affected by the presence of multiple distortions, such as institutional constraints (e.g. a requirement to offer static plans) and behavioral constraints (e.g. consumers who are inattentive or misunderstand plan incentives). This challenge makes the menu design problem an important application of the theory of the second best, where correcting one distortion — for example, by offering a dynamic plan that approximates marginal social cost — may fail to maximize social welfare in the presence of other distortions (Lipsey and Lancaster, 1956).

In this paper we develop an empirical framework for optimizing menu design when there are time-varying social costs and realistic design constraints. Our new framework accommodates consumer selection into plans on multidimensional heterogeneity, information frictions,

and high-frequency decisions for utilization and intertemporal shifting. It also accommodates institutional constraints on menu design, such as mandates to offer static price plans or restrictions on the intermittency of high-price hours in dynamic plans.

We apply our model to data on a random sample of eight-thousand households from a large utility in Phoenix, Arizona. We observe each household's demographics, price plan enrollment and switching, and consumption in 15-minute intervals for five years.[1] This is an ideal setting in which to study menu design because of the richness of the menu, which includes three time-of-use (TOU) and two non-TOU plans. Using the estimated model, we (i) evaluate how TOU pricing affects welfare in the current menu, (ii) test counterfactual interventions designed to increase consumer responsiveness to plan incentives, and (iii) characterize the second-best optimal menu.

We begin by documenting four descriptive facts about consumer behavior. First, we show that a subset of consumers respond strongly to TOU incentives by load-shifting and load-shaving. This finding is based on a nonparametric double-difference estimator that leverages longitudinal and cross-sectional variation in consumption. Most of the load is shifted after prices decline, with a smaller shift before prices rise. The size of the shift declines with the high-price interval's duration. However, our second descriptive fact is that most consumers do not respond to marginal prices: 74% of consumers who selected TOU plans do not respond to within-day price changes of up to 400%.

Next, we analyze consumers' plan choice and switching decisions. Our third fact is that consumers select into and out of TOU plans based on both their within-day price sensitivity and the amount of money they can save by switching without adjusting consumption. In other words, there is selection into plans on both the level and slope of demand (Einav et al., 2013; Ito et al., 2023). However, our fourth fact is that switching appears to be inhibited by inattention. For example, consumers are more likely to switch plans just after receiving large monthly bills.

Based on this descriptive evidence, we develop an empirical model in which consumers

[1]These households that do not own solar panels or electric vehicles. This group best represents the US population since fewer than 5% of houses have rooftop solar (EIA, 2022) and fewer than 5% of vehicle registrations are electric or plug-in hybrid (DOE, 2022). Consumers who own solar panels or electric cars choose from different price schedules than the consumers we study.

repeatedly choose price plans and consumption profiles. Each month, a consumer has the option to switch plans, but the consumer only considers this option if they first receive an "attention shock". A consumer who considers switching makes a decision based on their annual utility from each plan and their decision may be influenced by inertia (i.e. an increase in utility for the default plan). Given their plan choice, the consumer next decides how much to consume in every 15-minute interval of a day. Their consumption preferences differ between weekdays and weekends, and across the year. We allow for heterogeneity in consumers' desired load shapes (modeled as a set of 15-minute bliss points throughout the day), their price sensitivities, and their disutility from shifting load within a day. We also allow for heterogeneity in how consumers perceive prices; specifically, we allow for the possibility that some consumers respond to the average prices they see on their monthly bills, which is a common behavioral heuristic (Ito, 2014; Shaffer, 2020).

We estimate the model via simulated method of moments. We exploit quasi-experimental variation from sharp within-day price changes in TOU plans, as well as consumer choice behavior, to identify model parameters. We find that consumers sort themselves over TOU plans based partly on load-shifting preferences. After consumers choose plans, we find that they become relatively inattentive to the menu, but that inertia toward the current plan is small.

We use our estimated model to ask three questions. First, how would social welfare be affected by removing existing TOU plans from the menu? Surprisingly, we find that this would *increase* welfare. Decomposing this result, we show that dramatic TOU price changes (which can be more than four times marginal social cost) cause consumers to make large shifts in consumption from high-price to low-price hours. This behavior is privately rational. However, the social benefit is relatively small because the load is shifted between periods where private and external generation costs are similar. This benefit is smaller than consumers' private utility costs of load-shifting.

Next, we ask whether social welfare would be improved by a counterfactual intervention that makes the subset of consumers who currently respond to average prices respond to marginal prices instead.[2] We find that this would *decrease* welfare by exacerbating ineffi-

[2]For example, the intervention could be an information treatment or a subsidy for 'smart thermostats'

cient load-shifting responses to current plan incentives. The outcome would be improved if inattention and inertia were eliminated — or if consumers could be assigned to the plans that maximize social welfare — but we show that either scenario would still result in a net welfare loss from the intervention. These results echo the theory of the second best: correcting one or more distortions when other distortions are present may reduce welfare.

Third, what is the optimal menu given the typical utility's constraints? To define a benchmark for comparison, we start with a hypothetical 'first-best' scenario: assign all TOU consumers to a single plan that best approximates marginal social cost by setting one high price and one low price in contiguous intervals. This increases social welfare relative to the status quo by $50.1 per consumer/year. While the environment and the utility benefit, consumer welfare declines by $218.9 per consumer/year, implying that such a scenario is likely infeasible. Further, if we add the realistic design constraint that consumers must be allowed to choose the non-TOU plans on the current menu, self-selection (out of TOU pricing) cuts the welfare gain from offering the 'first-best' plan by more than half.

Finally, we solve for the second-best optimal menu. We optimize over thousands of potential TOU plans and millions of potential menus. Our objective function is to solve for the peak hours and on- and off-peak prices of TOU plans that maximize social welfare, subject to the constraint that the two non-TOU plans are still offered on the menu, consumers choose optimally, and changes to the menu are at least budget neutral for the utility. We first consider the case where a single TOU plan is offered. We find that it is optimal to distort the off-peak price to below marginal social cost (to encourage selection into the TOU plan) and to extend peak hours to 1-10pm. This menu achieves almost 90% of the benefits of the 'first-best' plan. It is also more equitable: consumers, the utility, and the environment all benefit. Further, low-income consumers benefit relative to high-income consumers. More complicated menus with multiple time-of-use plans would slightly benefit consumers, but result in almost no additional social welfare gain.

Contributions and related literature. Our study advances knowledge in three ways. First, we develop a new framework to study optimal menu design with time-varying private costs and externalities. Importantly, we distinguish consumers' price sensitivity from their

or other technology that could lower the cost of responding to marginal prices.

preferences for intertemporal substitution (and allow selection into plans on both dimensions). This distinction is both crucial to solving the menu design problem and a substantial departure from prior work on electricity plan selection and consumption (e.g. Hanemann, 1984; Hausman, 1985; Reiss and White, 2005; Fowlie et al., 2021; Ito et al., 2023). It also differentiates our study from prior research on how consumers incrementally choose price schedules and consumption quantities in other markets, such as health insurance and cell phone service (e.g. Einav et al., 2013, 2021; Grubb and Osborne, 2015; Cabral, 2017; Lin and Sacks, 2019; Marone and Sabety, 2022). We evaluate the extent to which harnessing this knowledge in practice delivers on its theoretical promise.

In addition, our framework incorporates several aspects of consumer behavior that have been shown to be quantitatively important in RCTs and quasi-experimental studies of electricity markets, and we are able to assess the degree to which they jointly interact and matter for optimal menu design. These include information frictions and inattention (Sallee, 2014; Jessoe and Rapson, 2014; Hortaçsu et al., 2017), default effects (Fowlie et al., 2021), selection into plans based on potential savings and price sensitivity (Ito et al., 2023), and variation in whether consumers respond to average or marginal prices (Ito, 2014; Shaffer, 2020).

Second, we develop the first evidence on how electricity consumers respond to a menu comprised of multiple static and dynamic price schedules. These complex menus represent the way that many utilities implement dynamic pricing in practice. Consumers in our setting choose from five plans, where the three dynamic plans differ in when the high-price interval starts, how long it lasts, how much price increases, and how these features change seasonally. This variation allows us to identify how consumers sort themselves over dynamic plans, differentiating our study from prior work that examined binary choices between a static plan and a dynamic plan (Fowlie et al., 2021; Ito et al., 2023).

Finally, we develop a set of new results about designing socially efficient menus in the presence of time-varying externalities and consumer self-selection. Prior studies examined how price schedule design affects market outcomes, conditional on an assignment of consumers to schedules (e.g. Wolak, 2011; Jessoe and Rapson, 2014; Prest, 2019; Yang et al., 2020; Fabra et al., 2021; Blonz, 2022; Harding et al., 2023; Burkhardt et al., 2023; Bailey et al., 2024; Schittekatte et al., 2024; Hinchberger et al., 2024). Our results extend this line of

research to a more generic setting where consumers sort over multiple price schedules. We show that menus can be redesigned to yield welfare gains to consumers, utilities, and the environment, without having to modify how electricity is generated. This finding parallels evidence on the optimal incentive regulation of utilities in Abito (2020). More broadly, our results are also relevant for assessing the equity and efficiency of dynamic pricing in other markets with time-varying externalities such as transportation (Li, 2018; Yang et al., 2020).

2 Context

2.1 The Salt River Project Utility

The Salt River Project (SRP) is a vertically integrated public utility. It was established in 1903 and serves about 2.2 million people in the Phoenix, Arizona metropolitan area.[3] SRP's generation profile is similar to the US as a whole. Just over half its generation comes from natural gas and coal-fired plants. It also co-owns one of the nation's largest nuclear plants and has a small but growing share of generation from renewable sources.

SRP has been a pioneer in dynamic pricing. In 1980, it was among the first utilities to offer an optional TOU plan (Schwartz, 2012). This required installing advanced metering infrastructure, commonly known as "smart meters", to track consumption in real time. During our study period, more than 99% of SRP customers had smart meters and the option to enroll in multiple TOU plans.

SRP has several key objectives in designing its price plans and does not simply maximize profit (SRP, 2018). They include recovering generation costs, increasing consumer welfare, and increasing sustainability.[4] Its specific sustainability goals for 2035 include reducing the CO_2 intensity of its generation by 65% relative to 2005 (SRP, 2023). When SRP management proposes changing its plans, the changes must be approved by SRP's publicly-elected Board of Directors that is tasked with ensuring its menu of plans is consistent with SRP's stated goals. We return to these objectives in our optimal menu design problem in Section 8.2.

[3]Households in Phoenix are assigned to electricity providers based on their residential locations.

[4]For example, SRP's Pricing Principles include: "*Choice—to constantly improve customer satisfaction through the creative design of pricing structures that reflect customers' different desires or abilities to manage consumption, assume more price control, or demand differentiated products and services*" (SRP, 2018).

2.2 The Menu of Price Plans

The menu that a customer faces depends on whether they have solar panels, an electric vehicle, both, or neither. We focus on households that did not have solar panels or electric vehicles at any point during our study period. This group best represents the US population since fewer than 5% of US households own each technology (EIA, 2022; DOE, 2022).

Figure 1 summarizes the menu of plans. Each panel shows how a particular plan's prices vary from SRP's winter season (November through April) to the summer season (May, June, September, and October) to the peak summer season (July and August).

SRP offers three TOU plans that raise price during weekday on-peak hours.[5] We refer to each plan by its summer on-peak hours. Figure 1a shows the 2-8pm plan. In summer, its price rises 289% from 2pm-8pm (shown on the horizontal axis) and this differential increases to 330% in the peak summer season.[6] In winter, the price rises 38% from 5-9am and 5-9pm. Figure 1b shows the 3-6pm plan. It has the same peak hours year-round, with price increases of 44%, 349%, and 404% in winter, summer, and summer peak seasons. SRP also offers a 4-7pm plan (not shown in Figure 1) that is otherwise identical to the 3-6pm plan.

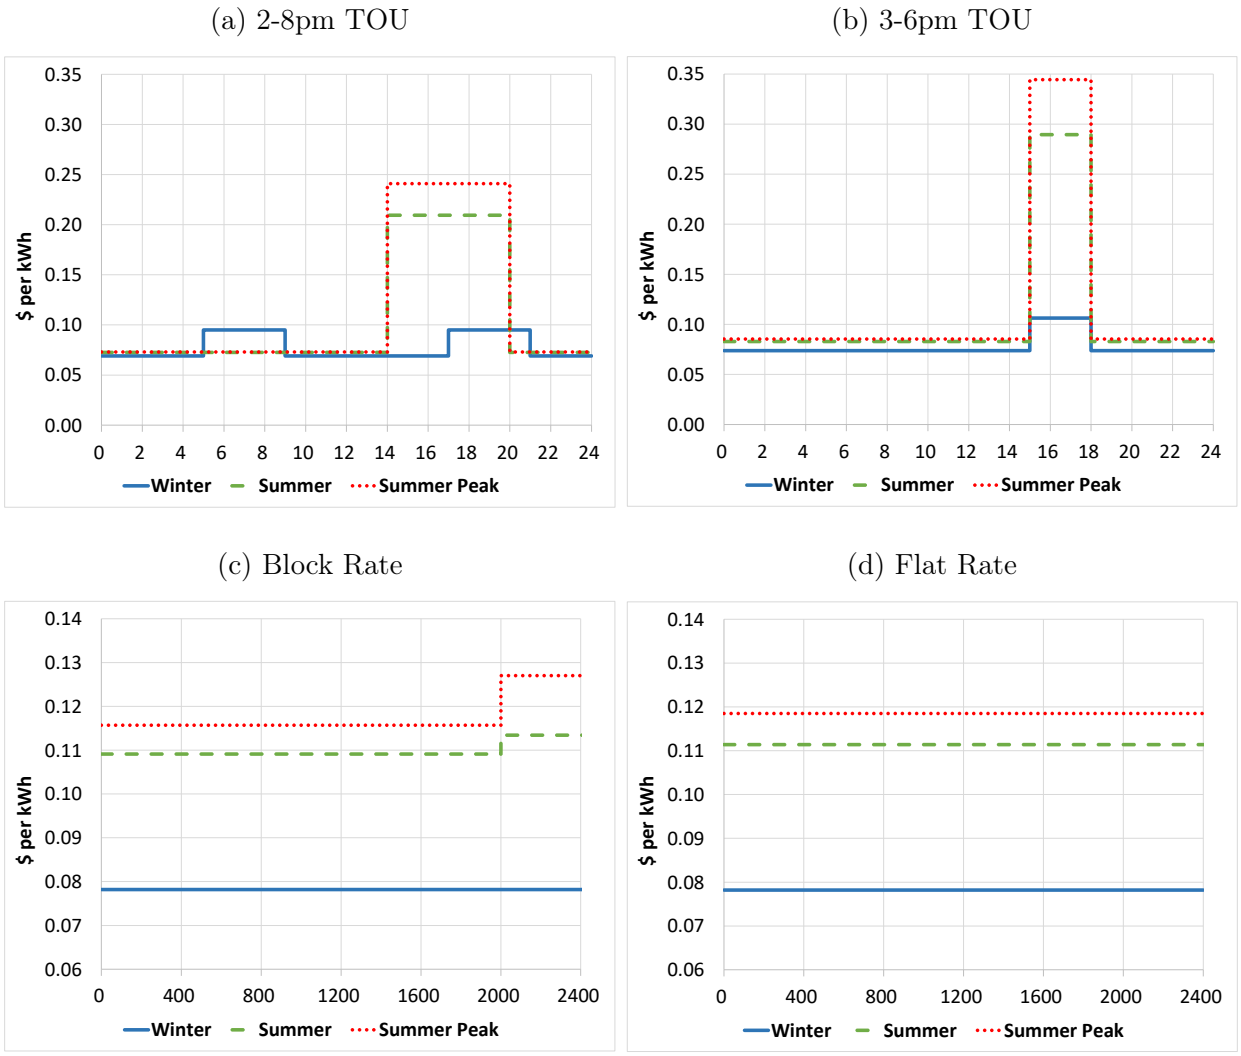
Figure 1c shows the block rate plan. Customers pay \$0.078 per kilowatt hour (kWh) in winter. In summer, the price rises to \$0.109 for the first 2,000 kWh during a monthly billing cycle (shown on the horizontal axis) and \$0.113 for each additional kWh. Those prices rise to \$0.116 and \$0.127 in the peak summer season. Figure 1d shows the flat rate plan. The winter price matches the block plan, and summer prices lie in between the block tiers. This plan is coupled with a prepay feature that requires customers to deposit money into an account that is drawn down to pay for consumption. It is designed for budget-minded customers.

This menu provides the first opportunity to study how consumers choose between multiple TOU and non-TOU plans. Prior observational studies of electricity demand focused on block-rate customers without access to TOU pricing (e.g. Reiss and White, 2005; Ito, 2014). Prior experimental studies allowed subjects to choose between one TOU plan and one non-TOU

[5] All weekends and the following holidays are excluded: New Year's Day, Memorial Day, Independence Day, Labor Day, Thanksgiving Day, and Christmas Day.

[6] The within-day variation in large compared to Spain's real-time pricing market, where pries varied by 23% on average (Fabra et al., 2021).

Figure 1: Plan-specific Prices by Hour of Day and Season



Note: In panels (a) and (b) the horizontal axis shows the hour-of-day. In panels (c) and (d) the horizontal axis shows total kilowatt hours consumed during a monthly billing cycle.

plan, after setting one of the two plans as the default (Fowlie et al., 2021; Ito et al., 2023).

2.3 Enrollment, Switching and Information

Importantly, there is no default plan on SRP's menu. New customers are prompted to choose a plan that matches their preferences by clicking a bubble on an online enrollment form or by talking to a customer service representative on the phone. Existing customers can follow the same process to switch plans at any time. SRP's website explains each plan's price schedule, along with an intuitive summary of how the plan works and how it differs from other plans.

Appendix A provides images of the choice process and additional background.

After choosing a plan, a customer can open an online account to monitor their hourly consumption and expenditures. The online accounts also report how much the customer can expect to save (or lose) by switching to other plans. These projections are based on one's consumption history and presented in terms of monthly and annual expenditures. Appendix Figures A.2 and A.3 show how the information is presented.

3 Data

3.1 Household Sample

We start with a random sample of 8,805 households who do not own solar panels or electric vehicles.[7] We observe the dates each household starts service, the zipcode of the address, and the plans they choose. For households who switch plans, we observe the dates they switch and the new plans they choose. Finally, we observe some demographics, including measures of household size and income.[8]

We use smart meter data on each household's consumption in 15 minute intervals from May 1, 2019 through April 30, 2023. We use these data to calculate monthly bills for each household. We also compute the monthly bill in each counterfactual plan holding consumption fixed; we leverage these counterfactual bills to develop descriptive evidence on plan choice in Section 4.2.[9] Appendix C provides supporting details for our bill calculator and evidence that it predicts actual expenditures almost perfectly (e.g. $\rho = 0.99$).

We make a few cuts to standardize the sample. First, we drop 0.3% of household-months with zero consumption. Second, we drop 0.3% of households in a pilot TOU plan that was closed to enrollment prior to our study period and eliminated in 2021.[10] Third, we drop 1.9% of households who are missing demographics and 4.3% whose primary residences are outside Phoenix.[11] Finally, we exclude March 2020 through April 2021 during which Arizona's social distancing policies temporarily changed residential electricity demand.

[7] As noted earlier, households with solar panels and/or electric vehicles chose from a different menu of plans and represent a small share of households.

[8] SRP obtained these data from an external contractor.

[9] Our model and counterfactuals allow for consumption to change when households switch plans.

[10] Peak hours were 2pm to 5pm. Dropping this plan simplifies construction of the choice set

[11] These households are labeled as "seasonal visitors" in the demographics shared by SRP. They face different incentives for plan choice and consumption than full-time residents.

Our trimmed sample includes 8,204 households. We observe them for 261,331 monthly billing cycles starting in May 2019 and ending in April 2023. Approximately 79% of households opened accounts before May 2019 and the remainder opened accounts during our study period. Their smart meter data comprise over 750 million 15-minute consumption intervals.

3.2 Summary Statistics

Table 1 summarizes household characteristics by price plan. The average household uses 1,313 kWh per month at a cost of $162. Consumption is driven by cooling due to Phoenix's desert climate. Mean consumption and expenditures are more than twice as large in peak summer months compared to winter months.[12] The last five columns show that the block rate plan has the largest market share (56%), followed by the 3-6pm TOU plan (23%) and the 2-8pm TOU plan (17%).

Table 1: Summary Statistics Overall and by Plan

	All plans	block rate	fixed rate	2-8pm TOU	3-6pm TOU	4-7pm TOU
market share (%)	100	56	3	17	23	1
# monthly bills	261,331	146,084	8,081	45,667	58,816	2,683
monthly kWh (mean)	1,313	1,202	1,229	1,609	1,373	1,240
monthly bill (mean $)	162	152	155	188	165	151
mean income ($1,000)	75	69	38	93	81	71
household size (mean)	2.1	2.0	1.8	2.4	2.0	1.8

Consumers are partially stratified across plans by household size and income. For example, mean income is lowest in the fixed-rate plan, consistent with that plan being designed for budget-minded customers. Likewise, income and household size are highest in the 2-8pm TOU plan.[13] In addition, regressing logged monthly consumption on household size and income reveals that adding a household member is conditionally associated with an 8% increase in consumption, which is equivalent to increasing income by about $27,000. These conditional associations, together with the stratification patterns, motivate us to include income and household size as potential preference shifters in our model in Section 6.

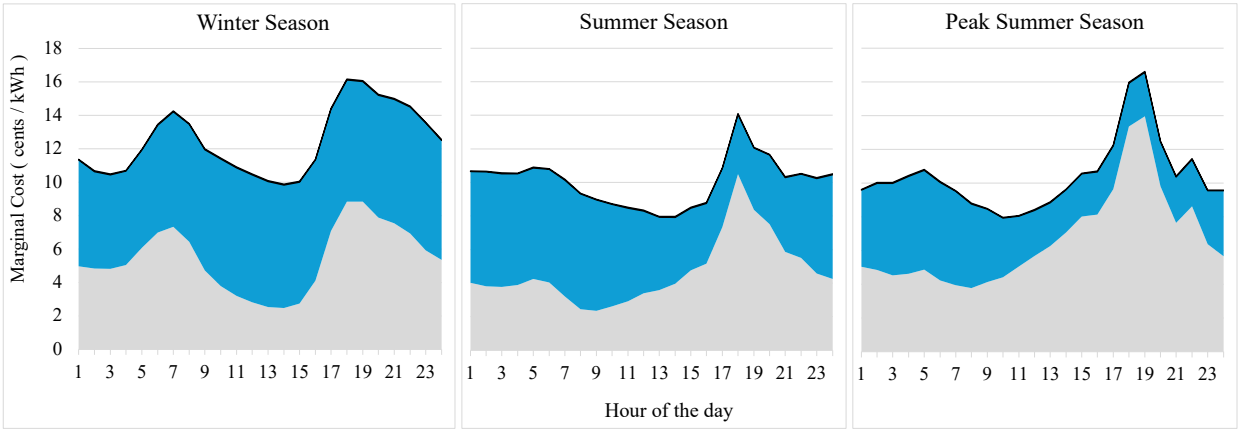
[12] Appendix Figure D.1 shows monthly variation in consumption and expenditures.
[13] Appendix D.1 provides additional information on household demographics.

3.3 Marginal Social Costs of Electricity Generation

We estimate the marginal social cost of electricity generation by applying methods from Borenstein and Bushnell (2022). This section provides a high-level summary of our approach, which we describe in detail in Appendix B.

We start by compiling locational marginal price data for SRP from the California Independent System Operator. These data describe wholesale prices at which SRP traded electricity with other utilities at 15-minute increments from 2021 through 2023. Wholesale prices reflect costs of generation and transmission on high-voltage power lines. Then we multiply the wholesale prices by inflation factors that Borenstein and Bushnell (2022) calculated to account for transmission costs on low-voltage power lines. This yields hourly estimates for the marginal private cost of generation.

Figure 2: Marginal Costs of Electricity Generation by Season and Hour



Note: The grey curve shows the marginal private cost of generation, averaged by hour within each pricing season. The blue curve shows the marginal external cost of pollution. Adding them yields the marginal social cost, shown as a black line at the top of each panel.

Next, we calculate hourly marginal costs of climate and health damages from burning fossil fuels. We use the AP3 integrated assessment model to estimate hourly damages from carbon dioxide, sulfur dioxide, nitrogen dioxide, and fine particulate air pollution emitted at each fossil fuel plant on the western interconnection of the US electricity grid (Clay et al., 2019). We aggregate hourly plant damages into four regions: SRP, the Southwest excluding SRP, the Northwest, and California. Then we regress regional damages for each pollutant on a spline function of hourly electricity load in each region. Finally, we use the coefficients to

predict the hourly external cost of electricity used by SRP customers. These costs account for damages from SRP plants as well as plants operated by other utilities that sell to SRP.

Figure 2 shows how marginal costs vary over the average day in each pricing season. Private costs (gray) reflect variation in wholesale prices and transmission costs. External costs (blue) are lower on summer afternoons when solar generation is higher and natural gas is the marginal fossil source of generation. Adding private and external costs yields the social cost curves shown in black at the top of each panel.

Note that the crests and troughs in Figure 2 do not align well with the peak period in the most popular 3-6pm TOU plan in Figure 1. We return to this observation in the counterfactuals. Moreover, TOU price levels diverge substantially from marginal social costs. For example, during the summer and peak summer seasons, TOU prices are two to three times higher than social costs during on-peak hours and slightly below social costs during most off-peak hours.[14]

4 Descriptive Evidence

4.1 How Do Consumers Respond to TOU Pricing?

We estimate the extent to which TOU households respond to TOU pricing by reducing consumption during high-price hours and shifting consumption to low-price hours. The challenge here is that we only directly observe the total consumption of each household, and this is jointly a combination of responses to plan incentives and preferences. Intuitively, and as we formally operationalize below, we isolate the effects of plan incentives in two steps. We first compare weekday and weekend consumption at the household level each week. The idea is that TOU incentives are switched off on weekends and so in principle this controls for preferences that vary across households and time. Moreover, we see that households sharply lower consumption in the TOU period on weekdays but not weekends in Appendix Figure 3b. Second, we use a non-parametric matching estimator to remove the effects of systematic differences in weekend versus weekday consumption.

More formally, let q_{itwj} denote consumer i's electricity use during 15-minute interval t on

[14] Appendix Figure B.3 illustrates this divergence by superimposing the enrollment-weighted average of TOU prices in Figure 1 on the cost curves in Figure 2.

day $w \in \{weekday(d), weekend(e)\}$ of a particular pricing season, given that the consumer is on price plan $j \in \{TOU, Block(B)\}$, where TOU denotes one of the three time-of-use plans. The difference between a TOU consumer's mean weekday and weekend consumption during t can be written as $E[q_d|i,t,TOU] - E[q_e|i,t,TOU] = \delta_{i,t,TOU} + \lambda_{i,t,TOU}$. The first term after the equality, $\delta_{i,t,TOU}$, measures how TOU pricing causes i to adjust consumption during t. This is the object of interest. The second term, $\lambda_{i,t,TOU}$, measures the remaining difference between weekday and weekend consumption that is unrelated to TOU pricing. For example, $\lambda_{i,t,TOU}$ may reflect differences in time spent outside the home for work or school on weekdays versus weekends.

The weekday-to-weekend consumption differential for a block-rate consumer can be written as $E[q_d|i,t,B] - E[q_e|i,t,B] = \lambda_{i,t,B}$.[15] Differencing average consumption differentials among TOU and block-rate consumers yields $\delta_{t,TOU} + \lambda_{t,TOU} - \lambda_{t,B}$. This statistic fails to identify $\delta_{t,TOU}$, however, if consumers sort themselves into plans such that $\lambda_{t,TOU} \neq \lambda_{t,B}$. For example, consumers who spend less time at home on weekdays may have more elastic demand during on-peak hours and, consequently, select into TOU pricing. This is an example of selection on the slope of demand (Einav et al., 2013; Ito et al., 2023).

Our estimator for $\delta_{i,t,TOU}$ mitigates selection bias by matching each TOU consumer to a composite block-rate consumer with a similar λ. The matching process embeds two assumptions. First, we assume there is a known internal, t', during which consumption is unaffected by TOU pricing: $\delta_{i,t',TOU} = 0$. Under this assumption, we can measure a TOU consumer's $\lambda_{t'}$ as: $\lambda_{i,t',TOU} = E[q_d|i,t',TOU] - E[q_e|i,t',TOU]$. Then we match the TOU consumer to block-rate consumers in the same quantile, v, of the unconditional distribution of $\lambda_{t'}$. Our second assumption is that these matches continue to be valid, on average, at all other times of the day. The two assumptions are stated formally as A1 and A2.

Assumption A1. *There exists a known t' s.t. $\delta_{i,t',TOU} = 0$ for all TOU consumers.*

Assumption A2. $(\lambda_{i,t,TOU}|v_{t'}) = E[\lambda_{i,t,B}|v_{t'}]$ *for all $t \notin t'$.*

Under assumptions A1 and A2, consumer i's response to TOU pricing can be recovered

[15] $\delta_{i,t,B} \equiv 0$ for all block-rate consumers because they have no price incentive to adjust consumption during the day.

from a nonparametric double-difference estimator:

$$\hat{\delta}_{i,t,TOU} = (E[q_d|i,t,v_{t'}] - E[q_e|i,t,v_{t'}]) - (E[q_{dB}|t,v_{t'}] - E[q_{eB}|t,v_{t'}]). \quad (1)$$

For TOU consumer i in quantile v of the λ distribution at t', the estimator subtracts the average weekday-to-weekend consumption differential among all block rate consumers in the $v_{t'}$ quantile from i's consumption differential. Repeating this calculation at all $t \notin t'$ recovers consumer i's average response to TOU pricing throughout the day.

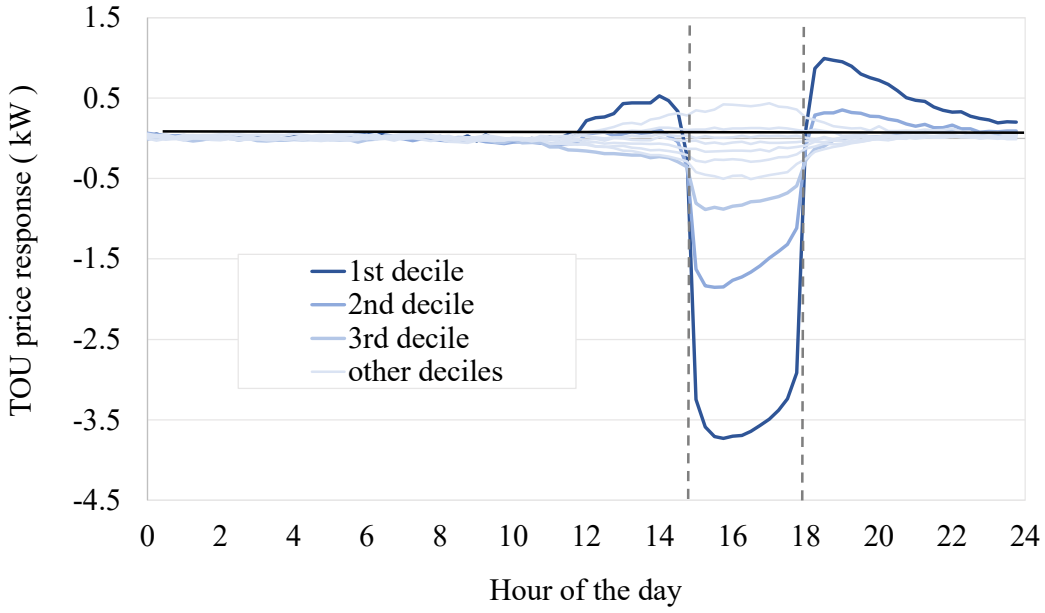
We implement this approach by setting $t' = [10am, 11am]$ during the summer and peak summer pricing seasons and dividing the $\lambda_{t'}$ distribution into ventiles. During the winter season we set $t' = [10am, 11am]$ for the 3-6pm and 4-7pm plans and t'= [12:30pm, 1:30pm] for the 2-8pm plan (which has winter peak hours from 5-9am and 5-9pm). We think these matching intervals are likely to satisfy Assumptions A1 and A2. First, they will satisfy A1 as long as TOU consumers do not shift consumption more than three hours from peak periods.[16] Second, the intervals capture differences in time spent outside the home on weekdays versus weekends due to regular commitments in late morning and early afternoon such as work and school. This will satisfy A2 as long as the time profile of these commitments is the same for matched TOU and block-rate consumers. Finally, while we cannot test A1 or A2 directly, some violations would produce "pre-trends". This provides an opportunity to falsify our assumptions. We return to this point after explaining Figure 3.

Evidence of load-shaving and load-shifting. Figure 3a presents a representative example of our estimates for $\hat{\delta}_{i,t}$. It shows the estimated causal response to TOU pricing during the peak summer season for consumers on the 3-6pm plan. The tinted curves depict heterogeneity in responsiveness. Each curve shows the mean response for a decile of consumers, ranked by the change in mean usage during on-peak hours (delineated by dashed vertical lines). Consumers in the three most responsive deciles reduce peak load by about 3.5 kW, 1.5 kW, and 0.75 kW. These reductions are substantial compared to the mean load during 3-6pm on weekends (about 4.5kW). Importantly, the figure also shows that the top two deciles shift

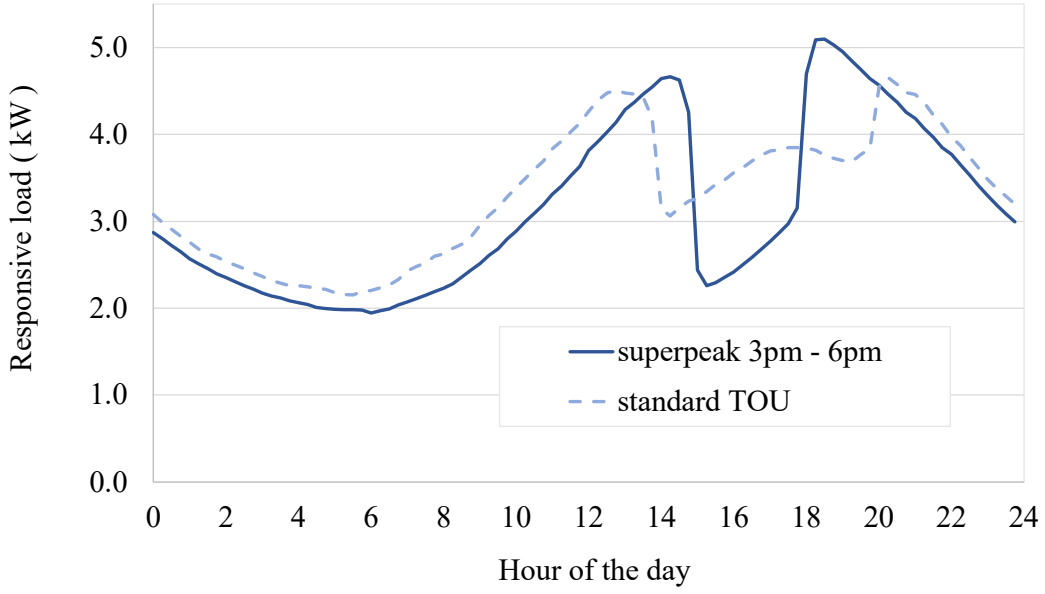
[16]This assumption would be stronger for households who own electric vehicles that require long charging periods. As noted earlier, our sample excludes such households.

Figure 3: Estimated Response to TOU Pricing

(a) Heterogeneity in load shifting and load shaving



(b) The mean load for responsive consumers is affected by peak hour incentives



Note: Panel (a) shows estimated consumer responses to TOU pricing in the 3-6pm plan. Each curve corresponds to a decile of consumers, ranked by the size of the estimated response during peak hours. Dotted vertical lines delineate peak hours. The solid line in Panel (b) shows the average weekday load curve for the subset of consumers that we find to be responsive to TOU pricing in the 3-6pm plan. The dashed line shows the analogous load curve for responsive consumers in the 2-8pm plan.

some of their on-peak load to adjacent off-peak hours, particularly 6pm to 8pm.[17] Analogous results for the other TOU plans and pricing seasons are shown in Appendix Figure D.2.

Figure 3a also suggests that most consumers have little or no response to TOU pricing. Peak usage changes by less than 0.5kW for the bottom seven deciles. Whether or not a particular consumer responds to TOU pricing could be driven by preferences, attention, or ownership of complementary products such as smart thermostats. Since we lack data to distinguish these hypotheses we treat responsiveness as a fixed characteristic. We use a one-sided paired sample t-test to divide the consumers we ever observe on TOU plans into two groups: "TOU-responsive" and "TOU-unresponsive". Our t-test is based on a donut discontinuity design comparing a consumer's mean usage 15-to-30 minutes before prices rise during summer and peak summer seasons with their mean usage 15-to-30 minutes after prices rise. We exclude the winter season to reduce statistical noise because consumption levels and changes are much smaller in winter, as shown in Appendix Figure D.2. Similarly, excluding 15-minute intervals on each side of the price increase reduces noise if some consumers adjust a few minutes early or late.[18]

The results indicate that 26% of consumers respond to hourly price changes.[19] Figure 3b shows peak summer load curves for these consumers in the two most popular TOU plans: 3-6pm and 2-8pm. Appendix Figure D.3 shows load curves for responsive and unresponsive households on each TOU plan and pricing season. The difference between the two curves motivates how intensive-margin adjustments to TOU pricing can potentially redirect some consumption to lower-cost hours.

[17] There are numerous ways to shift load. A leading example in Phoenix summers is to turn the thermostat down below one's bliss point to pre-cool living space from 12pm-3pm, then raise the thermostat above the bliss point from 3pm-6pm, before turning it back to the bliss point at 6pm. SRP's website advises customers that this pre-cooling strategy is the easiest way for TOU customers to save money.

[18] The test statistic is $(q_{i,2} - q_{i,-2})/[sd(q_{i,2} - q_{i,-2})/\sqrt{n}]$, where $t = -2$ and $t = 2$ refer to 15-to-30 minute intervals before and after prices increase at $t = 0$. An observation is a weekday during summer and peak-summer months. This yields 384 observations for a consumer that we always observe on TOU plans. The number of observations is smaller for consumers that switched between TOU and non-TOU plans, or that opened new accounts with SRP during our study period.

[19] In principle, it could be interesting to treat responsiveness as a state variable with transition probabilities that reflect opportunities for learning, technology adoption, and other factors driving changes in consumer decision-making. However, such transitions are rare. When we repeat the paired sample t-tests separately for 2019 and 2022 we find that only 2.2% of consumers switch from non-responsive to responsive and 1.8% switch in the opposite direction.

Validation and robustness. To judge the credibility of our approach to identifying the causal response to TOU pricing, it is helpful to note that Figure 3a shows no response to TOU pricing from midnight to 10am. All ten response curves approximately overlap the horizontal axis at zero. This provides an indirect validation check on Assumptions A1 and A2. If TOU and block-rate consumers are poorly matched (violating A2) or if TOU consumers shift load before 10am (implying that A1 is likely violated) then we would expect the curves to show "pre-trends" by diverging from zero before 10am. This is clearly not the case.

Habit formation poses another threat to identification. If TOU consumers were to habitually reduce electricity use during peak hours on weekends (when there is no price incentive to do so) then our estimator for the casual response would be attenuated. Fortunately, the scope for attenuation bias is minimal. Figure 3b shows that when weekday prices rise, usage declines by 31% for responsive consumers on the 2-8pm plan and by 51% for those on the 3-6pm plan. The corresponding changes on weekends are only 2% and 6%.[20]

4.2 How Do Consumers Choose Electricity Plans?

To investigate how consumers choose plans, we first focus on 395 consumers who switched between TOU and non-TOU plans. We regress an indicator for switching into TOU pricing on an indicator for whether the consumer is TOU-responsive, and on the average monthly amount the consumer would have saved on their non-TOU consumption profile had they purchased it on their cost-minimizing TOU plan. A caveat is that we compare bills in this exercise holding consumption fixed across plans; while this assumption is useful for the descriptive evidence, an important reason we later need a model is to relax this assumption.

The "TOU switcher" column of Table 2 shows that a $10 increase in monthly savings from switching to TOU pricing is associated with a 2% increase in the switching probability, consistent with selection on the level of cost savings. Moreover, TOU-responsive consumers are 13% more likely to select into TOU pricing, consistent with "selection on the slope" of hourly demand (Einav et al., 2013; Ito et al., 2023). In other words, consumers who respond to hourly price changes are more likely to choose TOU plans that reward that behavior.

Further, consumers who switch plans are more likely to minimize their annual costs after

[20] Appendix Figure D.4 provides additional visual evidence by showing weekday and weekend load curves for responsive and nonresponsive consumers on each TOU plan.

switching (and potentially adjusting their consumption to the new plan's incentives). This finding is based on 2,022 observations of annual costs for 761 consumers who switched plans exactly once between 2019 and 2023. We regress an indicator for whether the consumer minimized the cost of their annual consumption on an indicator for whether they switched previously, and an interaction between past switching and TOU-responsiveness.[21] We include fixed effects for consumers and years to identify the regression coefficients from within-consumer expenditure changes. The "cost minimizer" column in Table 2 shows that switchers are nearly 15 percentage points (or 50%) more likely to minimize costs after switching. The smaller interaction term suggests that switchers who do not respond to hourly price changes do respond to between-plan differences in price.

Finally, we observe that consumers are more likely to switch plans after receiving large monthly bills, consistent with inattention. Figure 4a shows that switchers are 2.5 times as likely to switch between May and August (when bills rise and peak) as between September and April.[22] Figure 4b shows that, within a given month, consumers are more than 3 times as likely to switch during the 5-day period after a billing cycle ends than during the rest of the month, despite being free to switch at any time.[23]

4.3 Summary of the Descriptive Evidence

In summary, the descriptive evidence suggests four main facts for our model to explain. First, the average TOU consumer shaves load during on-peak hours and shifts some load to off-peak hours (Figure 3b). Second, there is heterogeneity in load-shaving and load-shifting, and a substantial fraction of TOU consumers do not adjust hourly electricity use to hourly price changes (Figure 3a). Third, consumers select into plans based partly on the level of cost savings and the slopes of their hourly demand curves (Table 2). Finally, consumers are more likely to switch plans after receiving large monthly bills (Figure 4). Our model explains these facts, while also recognizing that TOU-unresponsive consumers may instead respond to average prices on their monthly bills (Ito, 2014).

[21] For consumers that were never on TOU plans the indicator is coded as zero.

[22] For the 1% of households who switched plans multiple times, we focus on their first switch

[23] We infer the end of the billing cycle from the day-of-month of each consumer's initial enrollment, which is normalized to zero in the figure. This provides a fuzzy measure for when a household would see its bill because of variation in: (i) the number of days per month, (ii) when households see bills sent by mail and email, and (iii) the relative timing of weekends and holidays, which affects when bills are sent.

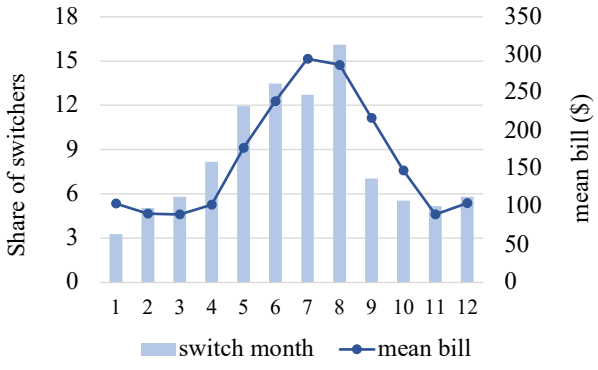
Table 2: Evidence on Plan Switching and Cost Savings

	TOU switcher	Cost minimizer
post switch		0.146
		(0.044)
post-switch x TOU-responsive		0.010
		(0.061)
TOU-responsive	0.129	
	(0.050)	
mean monthly savings from switching to TOU ($10)	0.023	
	(0.011)	
dependent variable mean	0.60	0.30
observation	account	account-year
# observations	395	2,022
R^2	0.03	0.72

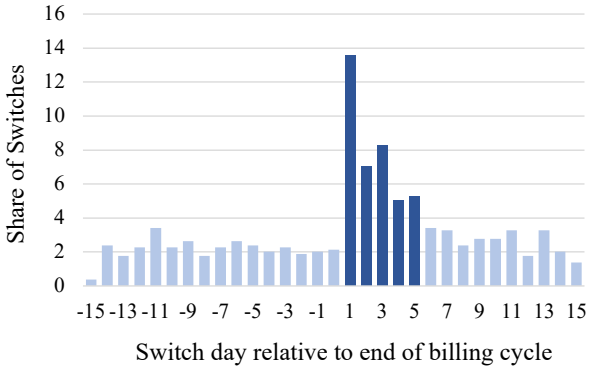
Note: The "TOU switcher" column shows results from a cross-section regression that uses data on consumers who switched between TOU and non-TOU plans. The "Cost minimizer" column shows results from a panel data regression that uses data on consumers who switched between any pair of plans. It includes fixed effects for consumer and year. Standard errors are clustered by consumer.

Figure 4: Timing of Plan Changes

(a) Switchers by month



(b) Switchers by day-of-month



Note: Panel (a) shows the share of switchers who switched each month alongside monthly average bills. Panel (b) shows when switches occurred relative to the end of the prior billing cycle.

5 Modeling Consumers' Plan Choices and Electricity Use

We model consumer behavior as a repeated two-stage process. Each month consumers choose whether to switch plans. Their plan choices are potentially subject to inertia and inattention. Then, given their current plans, consumers choose *how much* electricity to use each day as well as *when* to use it. We divide consumers into TOU-responsive and TOU-unresponsive groups based on their behavior, as explained in Section 4.1.

5.1 Daily Consumption Decisions Conditional on Plan Choice

5.1.1 TOU-responsive Consumer

Consider a TOU-responsive consumer i on plan j. For a representative weekday or weekend w in month m, the consumer has an ideal level of electricity consumption at each 15-minute interval t denoted by the bliss point ν_{imwt}. The bliss point is the amount of electricity the consumer would use if its price were zero.[24] Given these bliss points and the plan's prices, the consumer chooses how much electricity load to shift from each interval t to each other interval, $s_{ijmwtt'}$, where $t' \neq t$. The consumer also chooses how much electricity to directly consume in period t, q_{ijmwt}.[25] Overall, the consumer chooses $\mathbf{q}_{ijmw} = \{q_{ijmwt}\}_t$ and $\mathbf{s}_{ijmw} = \{s_{ijmwtt'}\}_{t,t'}$ to maximize the following indirect utility function:

$$\hat{V}_{ijmw} = \max_{\mathbf{q}_{ijmw}, \mathbf{s}_{ijmw}} \sum_t U(q_{ijmwt}, \mathbf{s}_{ijmw}) \tag{2}$$

$$\text{subject to: } q_{ijmwt} \geq 0 \text{ and } s_{ijmwtt'} \geq 0, \text{ where:}$$

$$U(q_{ijmwt}, \mathbf{s}_{ijmw}) = \underbrace{-(1/(2\omega_i)) \times (\nu_{imwt} - q_{ijmwt} - \Sigma_{t' \neq t} s_{ijmwtt'})^2}_{\text{Deviation from bliss point}}$$

$$- \underbrace{\Sigma_{t' \neq t} d_{imtt'}(s_{ijmwtt'})}_{\substack{\text{Disutility from} \\ \text{shifting consumption}}} - \underbrace{p_{jmwt} \times (q_{ijmwt} + \Sigma_{t' \neq t} s_{ijmwt't})}_{\text{Cost of load consumed at t}}$$

The first component of flow utility is a loss function that increases with the distance between the bliss point and the load that is either directly consumed at period t or shifted to another

[24] We restrict bliss points to be finite — that is, consumers have a satiation point of electricity consumption — to reflect the fact that they have finite appliances and preferences over their usage. For example, increasing air conditioning may reduce utility if the temperature is currently ideal.

[25] The term q_{ijmwt} can be thought of as unshifted load i.e. load that is desired to be consumed in period t that is also consumed in period t.

period. Here, the ω_i parameter indexes how painful it is for consumers to consume away from their bliss points.

The second component accounts for the disutility from shifting load from interval t to t'. For example, a consumer on a TOU plan with peak hours from 4pm to 7pm might shift cooking dinner from 6:00-6:30pm to 7:00-7:30pm when price is lower. Here, shifting consumption incurs disutility determined by $d_{imtt'}(.)$ but it can also reduce expenditures if load is shifted to periods with low prices, p_{jwmt}.

The third component of flow utility captures the total cost of load consumed at t, including load shifted to this period. Total load at period t, $q_{ijmwt} + \Sigma_{t' \neq t} s_{ijmwt't}$, is directly observed in the data. Note that we normalize the coefficient on price to -1, which implies that indirect utility can be measured in dollars.[26] The model allows consumers to have different price elasticities through differences in ω_i; this can be seen clearly in the model's first-order-conditions in Appendix Section E.1.

5.1.2 TOU-Unresponsive Consumers

TOU-unresponsive consumers could be price insensitive with $\omega_i \approx 0$, or they could respond to different measures of price such as the average prices they see on monthly bills. Responding to average price is a common behavioral heuristic among electricity consumers (Ito, 2014; Shaffer, 2020).[27] The average price is salient because it is reported directly on monthly bills whereas on- and off-peak prices are not. With this in mind, we nest both of these theoretical explanations and model a TOU-unresponsive consumer as choosing a load profile $\mathbf{q}_{ijmw}$ to

[26] Equivalently, we normalize the marginal utility of income to 1. This assumption is also common in the health insurance literature which uses related models (e.g. Einav et al., 2013; Marone and Sabety, 2022).

[27] For example, in a study of British Columbia consumers who were switched from a flat rate to a novel block rate plan, Shaffer (2020) finds that 92% responded to average or marginal prices and that 8% mistakenly perceived jumps in their new marginal block price as applying to all of their consumption. Our model allows for the first two responses. We think the third response is unlikely in our setting for two reasons. First, SRP's block rate plan is not novel; it had existed for more than 50 years prior to our study period (naturally, its prices changed over time). Second, SRP helps consumers avoid such mistakes by showing them how much they would have spent on their actual consumption over the past year had they purchased it on the block-rate plan (and each TOU plan). Appendix Figure A.3 shows how this information is presented.

maximize indirect utility given its bliss points and its monthly average price, $\bar{p}_{ijm}$.

$$\hat{V}_{ijmw} = \max_{\mathbf{q}_{ijmw}} \sum_t \Big(-(1/(2\omega_i)) \times (\nu_{imwt} - q_{ijmwt})^2 - \bar{p}_{ijm} q_{ijmwt} \Big) \tag{3}$$

$$\text{subject to: } q_{ijmwt} \geq 0, \quad \bar{p}_{ijm} = (5/7)\bar{p}_{ijm,weekday} + (2/7)\bar{p}_{ijm,weekend},$$

$$\text{where } \bar{p}_{ijmw} = \Sigma_t q_{ijmwt} p_{jwmt} / \Sigma_t q_{ijmwt}$$

Section 6.2.2 explains how we identify ω_i to determine whether TOU-unresponsive consumers actually respond to average price or are simply price-insensitive. Also note that the monthly average price is consumer-specific because it depends on how i's consumption profile interacts with plan j's price schedule on weekdays and weekends.

Equation (3) also differs from the TOU-responsive consumer's optimization problem in (2) in that there is no load shifting. This is not an assumption; it is the optimal response to load shifting being costly and consumers (potentially) perceiving price to be constant throughout the month.

5.1.3 Block-rate and Fixed-rate Plans

Non-TOU plans fit into the above framework. The fixed-rate plan can be viewed as a TOU plan with no on-peak period. The block-rate plan can be viewed as a modified fixed-rate plan, where the marginal price is determined by whether total consumption is above or below 2,000 kWh in that month.[28] Note that there is no incentive to load-shift on block-rate and fixed-rate plans because their prices are constant throughout the day.

5.2 Plan Choice

To calculate annual utility for plan j, we first aggregate flow utility in Equations (2) or (3) over weekdays and weekends in month m, and then aggregate over months to calculate:

$$V_{ij} = \sum_m (1/12) \times \Big((5/7) \times \hat{V}_{ijm,weekday} + (2/7) \times \hat{V}_{ijm,weekend} \Big) \tag{4}$$

At the start of each month m an existing customer i "pays attention" and considers

[28] When computing the indirect utility for the block-rate plan — which is relevant for the plan choice decision — we adjust for the fact that if a consumer is consuming more than 2,000 kWh in a month then consumption below 2,000 kWh has a lower price.

switching plans if they receive a positive draw from a Bernouilli distribution with parameter a_{im}.[29] If the consumer draws a value of 0 then they are inattentive and remain in their current plan. Conditional on paying attention, the consumer reconsiders their plan choice by maximizing annual utility plus ϵ_{ijm}, an i.i.d. logit error with scale parameter σ_ϵ:

$$\max_{j \in \mathcal{P}} \left\{ V_{ij} + \gamma_i 1[j = g(i)] + \sigma_\epsilon \epsilon_{ijm} \right\} \tag{5}$$

Here, γ_i is an "inertia" term that captures any other mechanisms, apart from attention, that could bias a consumer toward staying in their current plan $g(i)$.[30] While inertia and inattention are distinct mechanisms, both can contribute to a "default effect" in which consumers are unlikely to switch out of their default plans (Fowlie et al., 2021).

6 Estimation and Identification

To estimate the model we must parameterize the following components of Equations (2)-(5): the load shift disutility function $d_{imtt'}$, the bliss points ν_{imwt}, the loss function parameter ω_i, the attention parameter a_{im}, and the inertia parameter γ_i. We narrow our focus to estimating preferences for the subset of consumers who we ever observe on TOU plans. As we explain below, we rely on within-plan variation in prices for identification. Non-TOU plans do not contain enough variation to make this strategy feasible without adding relatively strong parametric assumptions.[31] Thus, focusing on TOU consumers improves internal validity. A caveat is that we need to consider whether non-TOU consumers will switch into TOU plans when we evaluate counterfactual policies. We discuss this in Section 8.

6.1 Parametric Forms

Loss function parameter ω_i. We set $\omega_i = \beta_{\omega 0 j(i)} + X_i \beta_{\omega 1}$, where X_i is a vector of consumer demographics incorporating household income and size. The intercept, $\beta_{\omega 0 j(i)}$, is specific to

[29] This is influenced by the descriptive evidence in Section 4.2 that switching plans is rare but follows a strong seasonal pattern. Note that, for new households, they make an active choice and so are attentive.

[30] New consumers make a similar choice except they do not have an inertia parameter because they do not have a current plan.

[31] For example, we would need to predict whether non-TOU consumers would respond to TOU pricing. We would also need to separately identify their loss function and load-shift disutility parameters from data on plan choice. In principle, this could be achieved via parametric restrictions on statistical distributions used to characterize each source of heterogeneity.

$j(i)$, the initial plan a consumer chose when they first opened their account.[32] These plan-specific intercepts address the "initial conditions" problem of disentangling preferences from inertia (Wooldridge, 2005). That is, $\beta_{\omega 0 j(i)}$ is designed to capture selection into plans based on unobserved preference heterogeneity. For example, TOU plans with larger within-day price differentials may attract consumers with more elastic demand.

We also allow the $\beta_{\omega 0 j(i)}$ and $\beta_{\omega 1}$ parameters to differ between TOU-responsive and TOU-unresponsive consumers. Thus, we estimate 10 loss function parameters: three initial plan parameters and two demographic parameters, separately for TOU-responsive and TOU-unresponsive consumers.

Bliss points. We parameterize the bliss points as $\nu_{imwt} = \bar{\nu}_{mwtj(i)} + \bar{\nu}_m + \bar{\nu}_i$. Here, $\bar{\nu}_i \sim N(X_i\beta_{\nu 1}, \sigma_\nu)$ is a consumer-specific random draw that shifts the bliss points up or down by a constant. There are three parameters to estimate in $\bar{\nu}_i$: the two parameters on household income and size in $\beta_{\nu 1}$, and the parameter σ_ν. The parameter vector $\bar{\nu}_m$ contains monthly vertical shifters of the bliss points. We estimate the $\bar{\nu}_m, \bar{\nu}_i$ parameters (except for σ_ν) separately for TOU-responsive and TOU-unresponsive consumers, so there are 29 parameters in total from these components.

The component $\bar{\nu}_{mwtj(i)}$ indexes how bliss points change, on average, across intervals of the day, weekdays versus weekends, and months of the year, all conditional on initial plan choice. We explain below how consumer-specific bliss points are identified at every 15 minute interval. However, estimation at this granularity would involve thousands of parameters. Therefore, we further parameterize the distribution of bliss points across each day. Concretely, we construct $\bar{\nu}_{mwtj(i)}$ using a $mwj(i)$-specific mean-preserving spread of the observed weekend consumption of the consumer's initial plan $j(i)$:

$$\bar{\nu}_{mwtj(i)} = \hat{c}_{mtj(i)} \times (\bar{c}_{mj(i)} / \bar{\hat{c}}_{mj(i)}) \tag{6}$$

$$\text{where:} \quad \hat{c}_{mwtj(i)} = \bar{c}_{mj(i)} + \beta_{\nu 2 wj(i)} \times (c_{mtj(i)} - \bar{c}_{mj(i)})$$

The bar notation denotes a daily mean, c denotes observed consumption, and $\beta_{\nu 2 wj(i)}$ indexes

[32] We treat this initial plan $j(i)$ effectively as a 'characteristic' of consumer i, which remains fixed even if consumer i subsequently switched to a different plan.

the level of the mean-preserving spread. For instance, at $\beta_{\nu 2wj(i)} = 1$, $\bar{\nu}_{mwtj(i)} = \hat{c}_{mtj(i)}$ so the bliss points are shaped similarly to observed consumption. At $\beta_{\nu 2wj(i)} = 0$, $\bar{\nu}_{mwtj(i)} = \bar{c}_{mj(i)}$ and the bliss points are constant throughout the day. We estimate $\beta_{\nu 2wj(i)}$ separately for each plan to capture selection into plans on load shape. To keep the number of parameters manageable, we set $\bar{\nu}_{mwtj(i)} = 1$ on weekends, and use the same initial plan-specific scale parameter for TOU-responsive and TOU-unresponsive consumers.[33]

Load shifting disutility. We set $d_{imtt'}(s_{ijmwtt'}) = \beta_{dimtt'} s_{ijmwtt'} + (1/2)\beta_{d4} s^2_{ijmwtt'}$, where (denoting t in a peak period and t' in an off-peak period):

$$\beta_{dimtt'} = \beta_{d0j(i)} + \beta_{d0m} + \beta_{d1}|t_{\text{midpoint}} - t'| + \beta_{d2} 1[t' \text{ in 9-5pm}] + X_i \beta_{d3} \tag{7}$$

Here, β_{d3} allows the disutility to vary with consumer-specific covariates and $\beta_{d0j(i)}$ captures selection into plans based on unobserved heterogeneity. β_{d0m} allows the disutility to vary each month to allow for seasonal changes in the use of appliances that can potentially be shifted (e.g. cooling). β_{d1} allows the disutility of load shifting to vary with how long it needs to be moved across time. Note that we measure the distance from the midpoint of the peak period $t_{midpoint}$. Finally, β_{d2} recognizes that it may be easier or harder to load shift during regular business hours when household members are less likely to be at home.

Inertia. We set $\gamma_i = \beta_{\gamma 0 j(i)} + X_i \beta_{\gamma 1}$. This specification allows inertia to vary with consumer demographics and initial plan choices, analogous to the loss function and bliss-point parameters.

Attention. As noted earlier, we model consumers as receiving attention shocks randomly drawn from a Bernoulli distribution with parameter a_{im}. We express a_{im} as a logistic function of X_i, indicators for initial plan choice, and month indicators. The month indicators allow the model to reproduce the seasonality in plan switching depicted in Figure 4.[34]

[33] Setting $\bar{\nu}_{mwtj(i)} = 1$ on weekends is motivated by the first-order-condition for consumption on weekends, which — given that there is a constant price on weekends in each plan — is: $q_{it} = v_{it} - \omega_i p_t$, i.e. consumption is a vertical shift of the bliss points.

[34] We experimented with allowing a_{im} to additionally vary with the percentage change in the consumer's last two monthly bills, following Hortaçsu et al. (2017). However, the associated parameter was imprecisely estimated because the percentage change in bills was nearly collinear with the month dummies, and including it did not substantially improve model fit.

6.2 Identification and Moment Construction

6.2.1 Consumption Parameters: TOU-responsive Consumers

The consumption parameters for TOU-responsive consumers are identified at the individual i level. We discuss intuition in this Section and provide formal details in Appendix E.

Identifying ω_i. A key identification challenge is price endogeneity. This is explicitly represented in the model primitives because the bliss points may increase in peak periods when TOU prices are higher. Therefore a naive regression of consumption on price could result in consumers appearing to prefer higher prices.

We solve this endogeneity problem by exploiting quasi-experimental variation that arises from the sharp price changes within TOU plans. Intuitively, under the assumption that bliss points 15 minutes before and after a cutoff are approximately equal, the consumption change at the cutoff is caused by the price change. This consumption change is a combination of load shifting and load shaving. The amount of load shifted is recovered from our descriptive analysis in Section 4.1. Therefore, we also observe directly — for each consumer — how a price change induces load shaving (i.e. induces the consumer to reduce consumption from their bliss point). This identifies ω_i at the consumer level (shown formally in Appendix E.3). Based on this idea, and our parameterization of ω_i, we include one moment for each TOU plan for the median difference in utilization left and right of the sharp jump in peak price, as well as moments for the covariance of this difference and each demographic characteristic.

Identifying bliss points. Using the identified ω_i, the bliss points are identified at the individual level for every 15-minute interval (shown formally in Appendix E.3). Intuitively, the bliss points can be recovered from the first-order condition for observed consumption once the quantity of load shifted and ω_i are both known.

We include three sets of moments to identify the bliss point parameters. First, we include three moments (one for each plan) for the average ratio of minimum to maximum consumption in August; this pins down the mean-preserving spread parameter for each plan. Second, we include average consumption moments to identify the plan-level and month-level bliss point shifters. These comprise three moments — one for each TOU plan in August — plus 11 more moments for average consumption on the 3-6pm TOU plan for each other month.

Finally, we include moments relating to the covariance of mean consumption and observable demographics to identify $\beta_{\nu 1}$, and the standard deviation of consumption across consumers in August for the 3-6pm TOU plan to identify σ_ν.

Identifying load-shifting parameters. Finally, the load-shift disutility parameters are identified by manipulating the first-order-conditions for the load shifted from peak period t to off-peak period t' to obtain $p_{jmwt} - p_{jmwt'} = d'_{imwtt'}(s_{ijmwtt'})$ (shown formally in Appendix E.1). This can be rewritten as:

$$p_{jmwt} - p_{jmwt'} = \beta_{d0j(i)} + \beta_{d0m} + \beta_{d1}|t_{\text{midpoint}} - t'| + \beta_{d2}1[t' \text{in 9-5pm}] + \beta_{d3}X_i + \beta_{d4}s_{ijmwtt'}$$

Thus, the β_d parameters are identified by how the shape of the shifted load varies over time.[35] To this end, we first include a moment for each plan measuring the mean load shifted 3 hours after the end of the peak period. Second, we include a moment for the 3-6pm TOU plan that load-shifting to 11am is 0.0, consistent with our descriptive results. Third, we include moments for the covariance of the mean load shifted and each observable demographic to identify heterogeneous effects. Fourth, we include a moment for the 3-6pm plan for the load shifted 3 hours before the start of the peak period. The difference between this moment and the shift after the peak pins down the curvature parameter β_{d4}. Finally, we include a moment for the mean shift in the 9am-5pm period to identify β_{d2}.

6.2.2 Consumption Parameters: TOU-unresponsive Consumers.

As noted earlier, a key challenge is to determine whether the TOU-unresponsive consumers are insensitive to TOU prices changes because they have very low ω_i's or because they respond to *average* price instead. We disentangle these hypotheses by exploiting plan switches.

Concretely, the difference in observed consumption in period t for a consumer that switches from plan j to plan j' is: $q_{ijmwt} - q_{ij'mwt} = \omega_i(\bar{p}_{ij'm} - \bar{p}_{ijm})$. Since the average prices before and after the switch are also directly observed, ω_i is identified at the consumer level for switchers (shown formally in Appendix E.4). We do not see all consumers switching

[35] A minor technicality is that, in practice, we do not observe temporal variation in X, so β_{d4} is identified by variation across households with different demographics. However, β_{d4} could be identified by changes in household demographics over time.

in our data, so for those that do not switch, ω_i is identified using the assumption that the price sensitivity is a function of observed demographics. With the price sensitivity identified, the bliss points are identified by similar arguments and an analogous set of moments as for the TOU-responsive consumers.

6.3 Plan Choice Parameters

The consumption parameters described above define the indirect utility of each plan for each consumer, placing no restrictions on how consumers choose their plans. With this information in hand, the remaining identification challenge is to disentangle inattention from inertia. Intuitively, suppose that there are three plans, and that we see the same consumer enrolled in two different plans at separate times. For example, in one period plan 1 is the consumer's default plan, and in another period plan 2 is their default. If they switch to plan 3 with the same probability in both scenarios, then there is no inertia and systematic failures to switch to higher-utility plans must be due to inattention. To the extent that their default plan distorts their choice probability, this then identifies inertia separately from inattention (shown formally in Appendix E.3). In practice we also exploit across-consumer variation (i.e. we can make the same argument for observationally equivalent consumers enrolled in different plans).

Overall, we follow Hortaçsu et al. (2017) who face a similar identification challenge. We include moments based on the choice probabilities of switching plans. We provide a full description of these moments in Appendix E.5.

6.4 Estimation

We estimate the model parameters in three steps, leveraging the fact that a consumer's daily consumption decision and plan choice decision are separable. First, we compute the marginal utility of income for TOU-nonresponsive consumers "offline" by exploiting switchers, and detail this procedure in Appendix Section E.4. Second, we estimate the parameters for within-day electricity consumption via simulated method of moments. We detail the computation of this step in Appendix Section E.2. Finally, we simulate each consumer's indirect annual utility for each price plan and estimate the plan choice model via general method of moments (as in Hortaçsu et al. (2017)), given the simulated indirect utilities.

7 Results

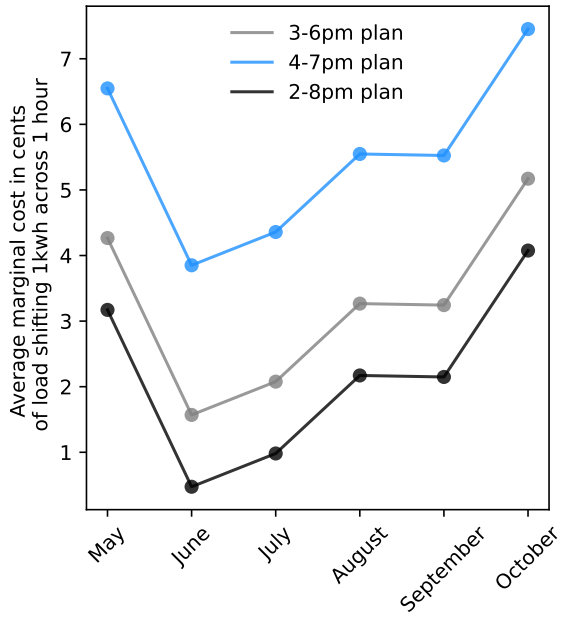
7.1 Daily consumption decision

We use the estimates to simulate consumption at 15-minute intervals for each plan, month, and weekday or weekend, for responsive and unresponsive TOU consumers. The model fits the targeted moments well (Appendix Tables E.4, E.5). It also reproduces untargeted moments describing monthly variation in plan-specific load curves (Appendix Figure E.1).

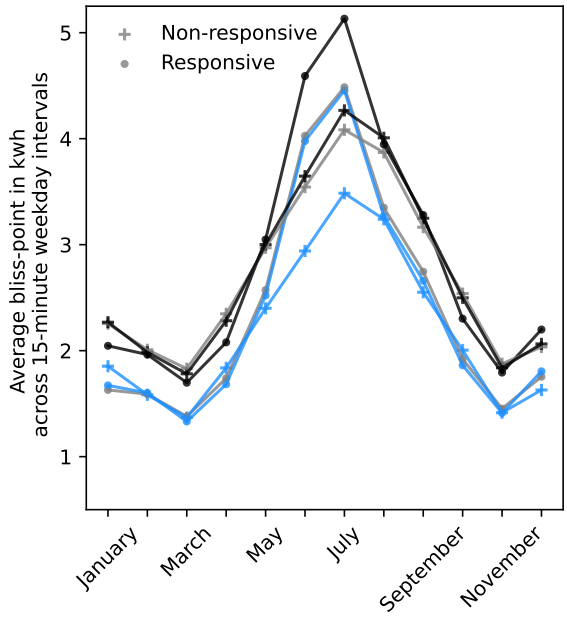
Panel A of Table 3 summarizes our estimates for the consumption parameters. While it is difficult to interpret their magnitudes in isolation, there are several interesting patterns. We see some sorting on price sensitivity across plans since there is heterogeneity in ω_i by initial plan choice. But this heterogeneity does not change much by income or household size.

Figure 5: Monthly average load-shifting disutility and bliss points

(a) Load-shifting disutility by month



(b) Avg. 15-min bliss point by month



Note: Panel (a) plots the load-shifting disutility by month, β_{d0m}, for TOU-responsive consumers in each plan. It is scaled as the (average) marginal cost to the consumer of shifting one kWh one hour to the right of the peak period, from a no-shift baseline where the quadratic load-shift term equals zero. We observe no load-shifting in the data for November-April, so our model does not allow consumers to load-shift in those months. Panel (b) plots the average weekday 15-minute bliss point by month, for TOU-responsive and TOU-unresponsive consumers in each plan, using the plan-specific color scheme from Panel (a). Each line is an average over the estimated bliss-points at each 15-minute interval of the day.

Table 3: Selected Parameter Estimates

Panel A: Daily Consumption Model

Parameter	Coef.	SE	**Parameter**	Coef.	SE
Load shifting disutility (β_d)			**Loss function (β_ω), TOU responsive**		
Duration (β_{d1})	0.057	(0.048)	*Demographics:*		
9am-5pm (β_{d2})	5.884	(2.210)	Household size	0.202	(0.146)
Demographics:			Income	0.786	(0.511)
Household size	0.0003	(0.00018)	*Initial plan choice:*		
Income	-0.003	(0.003)	3-6 pm plan	4.522	(2.099)
Initial plan choice:			4-7 pm plan	2.712	(1.430)
3-6 pm plan	0.008	(0.005)	2-8 pm plan	6.470	(2.605)
4-7 pm plan	0.014	(0.007)	**Loss function (β_ω), TOU non-responsive**		
2-8 pm plan	0.005	(0.003)	*Demographics:*		
Month intercepts:			Household size†	2.694	(0.666)
May	0.002	(0.002)	Income†	24.130	(7.762)
June	-0.004	(0.002)	*Initial plan choice:*		
July	-0.003	(0.001)	3-6 pm plan†	1.967	(0.934)
September	-0.0001	(0.00003)	4-7 pm plan†	-2.497	(0.732)
October	0.005	(0.003)	2-8 pm plan†	-0.389	(0.507)
			Bliss Points	See Appendix Table E.3	

Panel B: Plan Choice Model

Parameter	Coef.	SE	**Parameter**	Coef.	SE
Attention			**Inertia**		
Constant	-10.577	(1.032)	Incumbent plan dummy	4.343	(0.229)
Income	1.394	(3.869)	*Incumbent plan* $\times$:		
Household size	1.173	(0.175)	Income	0.370	(0.764)
Initial plan dummies:			Household size	0.659	(0.037)
4-7 pm plan	1.793	(0.521)	Init. 4-7 pm plan	4.547	(0.241)
2-8 pm plan	0.108	(0.555)	Init. 2-8 pm plan	2.380	(0.902)
Month dummies:	See Appendix Table E.2		Logit error scale (σ_ϵ)	0.007	(0.002)

Note.—Panel A reports results from the structural model of daily electricity consumption presented in Section 5.1 estimated via SMM. Panel B reports results from the plan choice model presented in Section 5.2 estimated via GMM. In both panels, parameter estimates are reported with bootstrapped standard errors in parentheses. † indicates externally calibrated parameters. See Appendix E.4 for more details.

We find that consumers dislike load shifting; that shifting is more costly during business hours (i.e., 9am-5pm); that income and household size do not significantly affect shifting preferences; and that consumers select into initial plans based on their capacity for shifting. To assess the overall magnitude of these preferences, Figure 5a plots the average marginal cost of shifting one on-peak kWh by one hour to the off-peak period from May through October.[36] This cost is negatively related to mean temperature, which can be explained by seasonal changes in load-shifting technology. For example, during hot summer months it is relatively easy to shift air conditioning just before or after peak hours. There is less need for air-conditioning in spring and fall, which limits the scope for shifting to appliances that may be harder to adjust (e.g. changing when dinner is cooked). The between-plan differences shown in the figure are consistent with selection into plans on load-shifting cost. Consumers with lower shifting costs tend to select the 2-8pm plan which has the longest peak period and so requires load-shifting across the largest amount of time.

Finally, Figure 5b summarizes the bliss point parameters (shown in Appendix Table E.3) by reporting monthly average bliss points by consumer type and plan. Intuitively, the bliss points increase in warmer months, reflecting the need for cooling. This pattern persists across all consumer types and plans.

7.2 Plan choice

Panel B of Table 3 reports the inattention and inertia parameters. Overall, the inattention parameters imply that the average consumer considers switching plans about once every 83 months. We find that attention increases with the number of household members and monthly bill size. For example, an average consumer that initially selected the 3-6pm TOU plan has a 0.03 percent chance of drawing an attention shock in February, when bills tend to be lowest, compared to a 0.40 percent chance in July when bills tend to peak (see Figure 4a). This is consistent with "bill shock" triggering consumers to reconsider their options, and similar to findings in Hortaçsu et al. (2017).

We find that inertia increases with household size and is also larger for consumers who

[36] We make this calculation from a no-shifting baseline where the quadratic load-shift term equals zero. Also recall that the model does not allow consumers to load-shift from November to April because we do not observe any load-shifting in the data during those months.

initially selected the 2-8pm and 4-7pm plans. To assess the relative importance of inertia and inattention for plan choice we use the model to predict the number of consumers who would switch plans over a one-year period if each mechanism were eliminated. Eliminating inertia would quadruple plan switching whereas eliminating inattention would increase it by a factor of 28. This difference is consistent with the hypothesis in Fowlie et al. (2021) that inattention drives the "default effect" in electricity plan choice.

8 Counterfactuals

8.1 Welfare implications of consumer behavior

We start by examining how the current menu of price plans incentivizes consumers to modify their behavior in ways that affect social welfare. We define the annual social welfare of consumer i in plan $j(i)$ as the sum of consumer surplus and producer surplus minus pollution damages:

$$\text{Welfare}_{i,j(i)} = V_{i,j(i)} + \text{PS}_{i,j(i)} - \text{Damages}_{i,j(i)} \tag{8}$$

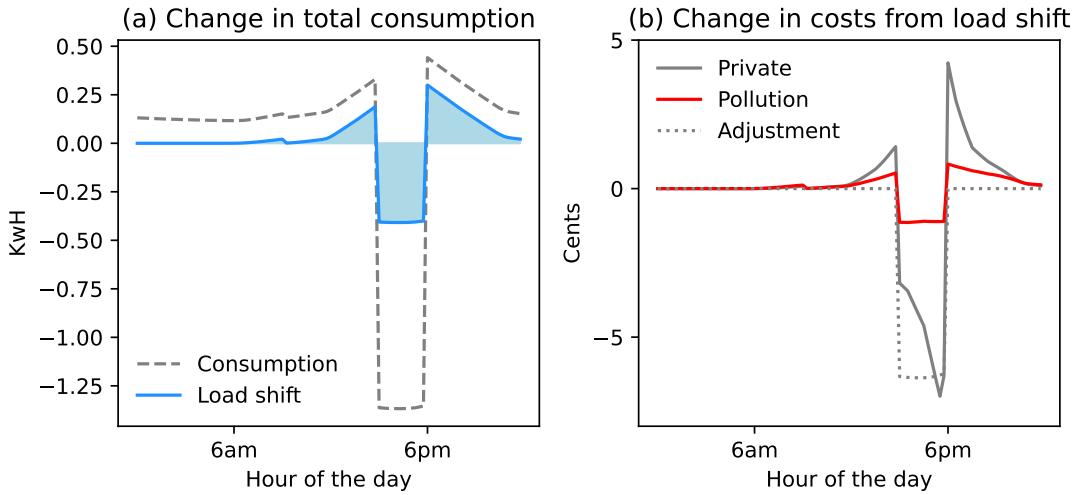
Therefore, total social welfare is $\Sigma_i \text{Welfare}_{i,j(i)}$.

Load-shifting and load-shaving. Figure 6 provides an example of how TOU pricing induces load-shifting and load-shaving behaviors that affect private and social costs. We illustrate these effects by reassigning all TOU-responsive consumers on the 3-6pm plan to the block rate plan. The dashed line in Panel (a) shows the mean change in consumption during September.[37] The shaded area is the load shifted from 3-6pm to adjacent off-peak hours due to TOU pricing. The difference between the shaded area and the dashed line from 3-6pm shows how much consumers load-shave when TOU prices are higher than block-rate prices. Conversely, consumption increases outside the 3-6pm window when TOU prices are lower than block prices, even after accounting for load-shifting.

Figure 6(b) shows how the load-shifting component of Panel (a) affects costs. The private cost of electricity generation and its pollution cost both decrease from 3-6pm. However, much

[37] We see similar patterns in all other months where we observe load-shifting behavior, specifically May through October.

Figure 6: Role of load-shifting and load-shaving: 3-6pm peak plan vs block-rate plan



Note: The figure shows the average change for TOU-responsive consumers reallocating them from the block rate plan to the 3-6pm TOU plan in September (we see similar patterns for other months when there is load shifting). The dashed line in Panel (a) shows the total change in consumption, and the shaded area shows the change from load-shifting. Panel (b) shows how load-shifting affects (i) the private costs of power generation, (ii) adjustment (load-shifting) costs borne by consumers, and (iii) costs from pollution damages.

Table 4: Welfare implications of consumer responses to the existing menu

	Change vs baseline menu (dollars/year/consumer)						
	Welfare	Consumer surplus	Producer surplus	Damages	Private costs	Expend.	Load shift costs
Existing menu							
No load shift	33.8	-13.4	45.1	-2.1	5.3	50.4	-37.1
Max. utility	23.5	22.6	4.1	3.2	-0.9	3.2	-16.0
Max. welfare	32.0	18.8	18.0	4.7	2.7	20.7	-24.6
Eliminate TOU	19.8	-62.8	73.7	-9.0	0.1	73.8	-9.1
↑ Responsiveness							
Current plans	-70.6	75.7	-117.1	29.2	-11.8	-128.9	32.8
Max utility	-55.3	90.5	-114.5	31.4	-14.5	-129.0	25.1
Max welfare	-47.1	86.2	-98.8	34.5	-9.9	-108.6	17.2

Note: The three rows below 'Existing menu' describe TOU-responsive consumers only (TOU-nonresponsive consumers are relatively unchanged). The 'max. utility' and 'max. welfare' rows reallocate consumers across the three TOU plans. 'Eliminate TOU' is the effect on both TOU-responsive and TOU-nonresponsive consumers from eliminating the three TOU plans in the existing menu. The rows below '↑ Responsiveness' describe TOU-nonresponsive consumers only, after a hypothetical policy causes them to respond to TOU prices similarly to TOU-responsive consumers.

of this mass is transferred to just after 6pm when marginal costs remain high. Thus, the 3-6pm plan induces consumers to shift load, but not necessarily to lower-cost hours. Further, consumers incur a significant utility cost of load-shifting, shown by the dotted line. Although load-shifting is a privately optimal response to TOU pricing, it is striking that the consumer adjustment costs exceed the private and pollution cost savings during peak hours.

To assess the broader implications of load-shifting we measure how social welfare would be affected by preventing all TOU-responsive consumers from load-shifting on their chosen plans. The first row of Table 4 reports results. Consumers use more electricity in peak hours, increasing their mean bill costs by \$50.4 per year. This is dampened by a \$37.1 reduction in load-shifting costs, so that mean consumer surplus declines by \$13.4 per year. However, producer surplus increases by \$45.1 per consumer/year due to higher revenue and a small increase in generation costs. The changes in private generation costs and environmental damages are small because the "sending" and "receiving" periods over which load is shifted have similar costs, as shown in Figure 6(b). Overall, eliminating load shifting in the current menu would *increase* social welfare by \$33.8 per consumer/year. This does not imply that load-shifting is intrinsically inefficient. Rather, it highlights the difficulty of designing a menu of price plans that succeeds in leveraging load-shifting to increase social welfare in the presence of time-varying private costs and externalities, choice frictions, self-selection, and consumers who differ in whether they respond to average or marginal prices.

Choice frictions. To measure how choice frictions affect welfare, we set the inertia and attention parameters to zero for TOU-responsive consumers and simulate their utility-maximizing plan choices. We find that they tend to move to the 2-8pm plan where they reduce load shifting. The "max utility" row of Table 4 shows that this increases social welfare by \$23.5 per consumer/year, mainly through higher consumer surplus.

Selection into plans. Next, we examine how self-selection affects welfare by assigning each TOU-responsive consumer to the current plan that would maximize social welfare. This assignment increases social welfare by \$32.0 per consumer/year. Much of the benefit again comes from moving consumers to the 2-8pm plan where they reduce load-shifting.

Does offering the existing time-of-use plans increase social welfare? To answer this question we simulate eliminating the three existing TOU-plans. This would *increase* social welfare by $19.8 per consumer/year. Thus, offering the existing TOU plans is inefficient from a social perspective.[38]

For TOU-unresponsive consumers, the inefficiency stems from the fact that the block-rate and fixed-rate plans yield average prices that are closer to marginal social cost. For TOU-responsive consumers, the inefficiency is exemplified by Figure 6: their utility cost of load-shifting far exceeds the reduction in generation costs. Overall, the current TOU plans make large transfers to consumers to induce load-shifting that is ineffective at reducing costs. In terms of distributional implications, eliminating the current TOU plans would increase producer surplus and reduce pollution damages, but it would also reduce mean consumer surplus by $62.8 per year.

Responding to marginal prices. Finally, we simulate a hypothetical policy that makes TOU-unresponsive consumers respond to TOU prices in same way as responsive consumers, conditional on demographics. The policy could be an information treatment, for example, or it could endow consumers with "smart thermostats" that lower the cost of adjusting to marginal prices. Examples of such interventions in the literature include Jessoe and Rapson (2014); Prest (2019); Blonz et al. (2024). If the treated consumers were to remain on their plans and not load shift then the intervention would lower social welfare by $41.6 per consumer/year. Load-shifting increases the loss to $70.6 per year. These results echo the theory of the second best: fixing one distortion — whether consumers respond to marginal prices — can reduce welfare when there are other distortions — here, prices that deviate from marginal social cost.

The last four rows of Table 4 repeat the experiment, first after removing choice frictions, and then after assigning consumers to their social welfare-maximizing plans. Removing choice frictions improves outcomes, as some consumers move to the 2-8pm plan and reduce load-shifting, but social welfare still declines relative to the pre-policy status quo. Assigning

[38]It is possible that offering a subset of the existing TOU plans could be socially preferable to eliminating them entirely. We tested this hypothesis by searching over menus for all permutations of existing TOU plans (plus the base plans). We found that no such subset exists.

consumers to their social welfare-maximizing plans yields little improvement.

8.2 The optimal second-best menu design problem: setup

We characterize the menu design problem as the social planner choosing a set of plans $\mathcal{P}$ from a class of potential menus $\mathcal{M}$ (which we vary across the counterfactuals) subject to constraints:

$$\max_{\mathcal{P}\in\mathcal{M}} \sum_i \text{Welfare}_{i,j(i)} \qquad \text{Subject to:} \tag{9}$$

$V_{i,j(i)} \geq V_{i,j'}$ for all $j' \in \mathcal{P}$ (Across-plan implementability)

Behavior of i in j determined by Equations 2-3 (Within-plan implementability)

$\Sigma_i PS_{i,j(i)} \geq \Sigma_i PS_{i,j_{\text{baseline}}(i)}$ (Budget non-negativity)

Fixed Rate, Block Rate $\in \mathcal{M}$ (Menu includes base plans)

The planner faces four constraints. First, consumers choose plans optimally.[39] Second, within-plan implementability: given plan j's incentives, consumer i chooses their 15-minute consumption optimally, potentially load-shifting and load-shaving. Third, the new plan menu must weakly increase the producer surplus of the utility. This captures the utility's pricing principles of 'Cost Relation' and 'Sufficiency' (i.e. "to recover the cost of... a system of assets" (SRP, 2018)): since we assume this principle is satisfied in the baseline menu this constraint implies it is also satisfied in counterfactual menus. Fourth, the block-rate and fixed-rate plans (which we refer to as the 'base plans') are in $\mathcal{M}$, consistent with the utility's preference for not mandating TOU pricing. This captures the utility's pricing principles of 'Choice' and 'Gradualism' (SRP, 2018): for example, in menu redesigns in 2019 and 2024 the utility kept these two conventional non-TOU plans, and instead redesigned the set of TOU plans.

Overall, the constraints in (9) explicitly capture four out of five of the pricing principles that guide the utility's approach to rate design; we later check that the fifth (consumer equity) is also satisfied. Furthermore, the objective function of maximizing social welfare —

[39] Since we are considering removing the current TOU plans and offering a new menu, this is an active choice and so consumers will not be subject to inertia or inattention. However, we also discuss whether inertia and inattention could be used as a policy tool.

as opposed to just profit — closely aligns with how changes to our non-profit utility's menu are evaluated in the regulatory process. For example, the SRP board of directors explicitly states the objective is "not to pursue the maximization of profit" (SRP, 2018), and more recently has considered carbon emissions in rate design (SRP, 2023).

Our characterization of the planner's problem embeds two main assumptions. First, we maintain the trivial participation constraint that every consumer must choose a plan. This is equivalent to assuming that changes in $\mathcal{P}$ do not induce consumers to move to residential locations outside the utility's service territory. However there is no TOU participation constraint. Since $\mathcal{P}$ always includes the block-rate and fixed-rate plans, they serve as outside options for existing TOU customers.

Second, we assume that non-TOU customers will not switch to newly offered TOU plans. The practical reason for this assumption is that, as mentioned in Section 6, it is difficult to identify these consumers' load-shifting preferences because we do not observe them in plans with within-day price changes. However, the assumption is arguably justified for two reasons. First, unlike current TOU consumers who would need to make active choices when TOU plans change, non-TOU customers would face inertia and inattention, which our estimates suggest are substantial. Second, since these consumers did not choose TOU plans in the existing menu, they may be unable to profitably respond to TOU incentives. As well, the counterfactual menus we consider below are, at most, slightly more generous than existing options for consumers who actively chose these plans in the baseline.

Computation. We consider single-peaked TOU plans with a connected set of on-peak hours in $\mathcal{M}$.[40] We allow the planner to vary the onset and duration of peak hours as well as on- and off-peak prices in 1¢/kwh increments. Overall, we optimize over thousands of potential TOU plans and millions of potential menus in the multiple-TOU case. In each case we compute each consumer's utilization in every 15-minute interval throughout the year.

Note that we do not optimize over the fixed fee for each plan (known as the 'monthly

[40]The single-peaked feature is shared by the existing TOU plans with the exception of the 2-8pm plan in Winter, which has two peak periods with relatively minor price changes. One may wonder about the potential efficiency gains of designing plans with multiple peaks. It would be computationally challenging to flexibly search over this dimension, given the large number of potential menus in the single-peaked case. Moreover, utilities typically employ single-peaked plans. In the limit, offering multiple peaks approaches real-time-pricing, which utilities have chosen not to employ.

service charge'). This is constant across plans in the empirical menu, and so allowing it to vary by plan does not appear to conform with the utility's decision making. Furthermore, as we discuss below, modifying the fixed fee would likely not improve outcomes. An interesting topic for future research — albeit one that we do not consider here because it would require additional assumptions on how consumers interact with plan designs that are not in our sample — would be to consider optimizing over more complicated plan designs that, for example, price the maximum load in any interval (known as a 'demand charge') or that raise price when load approaches grid capacity (known as 'critical peak pricing').[41]

8.3 The 'first-best TOU' benchmark

We consider the benefits of the 'first-best TOU' benchmark. By 'first-best TOU' we mean both (i) a menu of plans, $\mathcal{P}_{first}$, and (ii) an allocation of consumers to plans that maximizes social welfare. Equivalently, this is the solution to the menu design problem if the planner knows consumer types and does not face the 'across-plan implementability' and the 'budget non-negativity' constraints. The resulting first-best menu and allocation is simple. It consists of the base plans plus a single TOU plan with prices and peak hours that best approximate real-time-pricing at marginal social cost. The planner then assigns all TOU consumers to the new plan.

Figure 7(a) plots the first-best prices against marginal social cost during the summer peak season in July and August, and Table 5 shows that the plan increases social welfare by \$50.1 per consumer/year. However, the plan's distributional implications could undermine its feasibility. While the utility and the environment benefit, mean consumer surplus declines by \$218.9 per year. Note that even if the utility could modify the fixed part of the tariff to compensate consumers at the expense of the utility — i.e. implementation with non-distortionary transfers — consumers would be still worse off, since the decrease in consumer surplus is greater than the increase in producer surplus. Further, mandating consumers switch to plans where they are hundreds of dollars worse off violates our utility's regulated objective of allowing consumer choice.

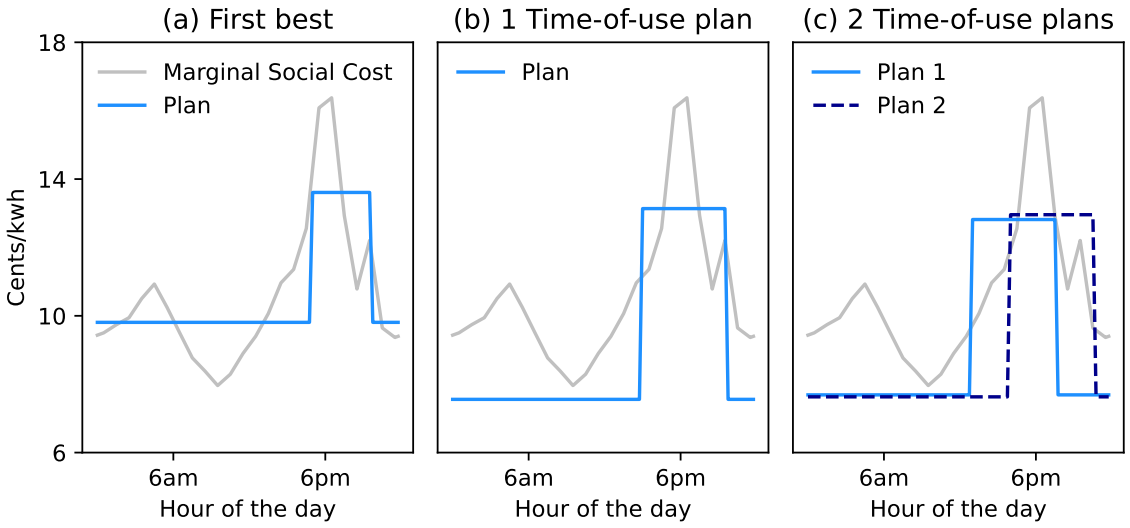
Therefore, we next assign TOU consumers to the first-best plan, but allow them to opt

[41] For evaluation of critical peak pricing schemes see, for example, Jessoe and Rapson (2014), Fowlie et al. (2021), Blonz (2022), and Hinchberger et al. (2024).

Table 5: Welfare effects of counterfactual menu designs

	Change vs baseline menu (dollars/year/consumer)					
	Welfare	Consumer surplus	Producer surplus	Damages	Private costs	Expend.
First best:						
No choice	50.1	-218.9	191.7	-77.3	-49.9	141.8
Choice + inattention, inertia	44.4	-184.9	165.4	-63.8	-40.0	125.4
Choice + inertia	26.7	-60.5	65.4	-21.8	-9.2	56.2
Max. utility	23.1	-57.6	67.3	-13.3	-2.9	64.4
Menu designs:						
Base plans + optimize:						
1 peak/off TOU plan	44.2	1.3	6.2	-36.7	-14.6	-8.4
2 peak/off TOU plans	44.7	8.5	0.2	-36.1	-13.8	-13.7

Figure 7: Counterfactual menu designs (time-of-use component)



Note: These figures show the design of the TOU components in various classes of menus. All menus include the 'base plans' of the fixed-rate plan and the block-rate plan. Since the marginal social cost and plan design change over the year, we plot this figure for weekdays in SRP's defined 'summer peak' season of July and August. In Panel (a) the peak hours are 3-10pm. In Panel (b) the peak hours are 1-10pm. In Panel (c), plan 1's peak hours are 1-8pm and plan 2's peak hours are 4-11pm.

out by switching to the block-rate or fixed-rate plans. Table 5 shows that if the inattention and inertia that we observe in the data were to persist after consumers are assigned to plans, then social welfare would increase by $44.4 per consumer/year. This is only a slight decline from the first-best assignment rule because our choice model parameters imply that few consumers would opt out of the default (first-best) TOU plan. On the other hand, this policy would be a dramatic change from the status quo and, as such, it would arguably eliminate (or substantially lower) inertia and likely cause consumers to be attentive. In this case, Table 5 shows that self-selection would cut the benefits of offering $\mathcal{P}_{first}$ by more than half to $23.1 per consumer/year.

Overall, these results illustrate an important caveat to the conventional wisdom that offering a plan that prices at marginal social cost is socially optimal. When consumers have access to other options — here, the base plans — the planner may be able to improve welfare by accounting for self-selection and *not* pricing at marginal social cost. This motivates our final counterfactual exercise in solving for the optimal menu under self-selection and the other design constraints in (9).

8.4 Optimal second-best menu design: results

We re-solve the menu optimization problem for a scenario where existing TOU consumers are required to actively choose between the base plans and a new single-peaked TOU plan. We find that the optimal second-best menu would increase social welfare by $44.2 per consumer/year. Further, this plan achieves 88.3 percent of the welfare gain from assigning all TOU consumers to the first-best plan.

Figure 7(b) shows that the off-peak TOU price is substantially lower than in the first-best plan. While this distortion provides socially beneficial incentives to induce consumers to select into the plan, it causes consumption to exceed socially optimal levels during off-peak hours, generating a slight welfare loss relative to first-best assignment. Table 5 also shows that consumers, the utility, and the environment all do better on average, with most of the benefits coming from a reduction in environmental damages. We also show in Appendix Table E.6 that the optimal menu is progressive, in the sense that low-income consumers do better than high-income consumers. This suggests that equity is unlikely to be a concern.

Next, we consider the benefits of offering more than one TOU plan. Specifically, we solve for the optimal menu under self-selection with two TOU plans.[42] In theory, moving from one TOU plan to two could increase social welfare by leveraging heterogeneity in bliss points, load-shifting disutility, and the price sensitivity. In practice, however, we find that the added benefit is minimal—less than one dollar per consumer/year. Part of the reason is that the optimal second-best menu with one TOU plan achieves social welfare near the 'first-best TOU' upper bound. This illustrates a policy-relevant lesson: a simple optimally-designed menu can perform very well. In contrast, more complicated menus of TOU plans may add little benefit, and risk harming welfare through self-selection.

9 Conclusion

We developed and estimated a novel model of electricity plan choice, consumption, and intertemporal substitution. We used the model to design a second-best optimal menu of electricity price schedules, given realistic design constraints and consumer self-selection. Our findings demonstrate that menus can be redesigned to simultaneously increase consumer welfare, reduce private generation costs, and improve environmental outcomes. These findings are based on data from the hottest metropolitan area in North America: Phoenix, Arizona. Its residents have adapted to extreme heat through energy-intensive air conditioning. Therefore, our findings may inform how other regions can optimize menu design to manage the growing demand for cooling as global temperatures rise.

References

Abito, Jose Miguel, "Measuring the welfare gains from optimal incentive regulation," *The Review of Economic Studies*, 2020, *87* (5), 2019–2048.

Bailey, Megan R., David P. Brown, Blake C. Shaffer, and Frank A. Wolak, "Show me the money! Incentives and nudges to shift electric vehicle charge timing," *Forthcoming in American Economic Journal: Economic Policy*, 2024.

Blonz, Joshua A., "Making the best of the second-best: Welfare consequences of time-

[42]Note that our framework does permit us to look at three or more TOU plans in the menu. However, the computational burden also scales exponentially with the number of plans. Since we ultimately find that two TOU plans does not add much value compared to one TOU plans, we do not also consider this case.

varying electricity pricing," *Journal of the Association of Environmental and Resource Economists*, 2022, *9* (6), 1087–1126.

Blonz, Joshua, Karen Palmer, Casey J. Wichman, and Derek C. Wietelman, "Automation, smart thermostats, and time-varying pricing," *Forthcoming in American Economic Journal: Applied Economics*, 2024.

Borenstein, Severin and James B. Bushnell, "Do Two Electricity Wrongs Make a Right? Cost Recovery, Externalities, and Efficiency," *American Economic Journal: Economic Policy*, 2022, *14* (4), 80–110.

Burkhardt, Jesse, Kenneth T. Gillingham, and Praveen K. Kopalle, "Field Experimental Evidence on the Effect of Pricing on Residential Electricity Conservation," *Management Science*, 2023, *69* (12), 7784–7798.

Cabral, Marika, "Claim timing and ex post adverse selection," *Review of Economic Studies*, 2017, *84* (1), 1–44.

Clay, Karen, Akshaya Jha, Nicholas Muller, and Randall Walsh, "External costs of transporting petroleum products: Evidence from shipments of crude oil from North Dakota by pipelines and rail," *The Energy Journal*, 2019, *40* (1), 55–72.

DOE, "TransAtlas data files," Technical Report, United States Department of Energy, Alternative Fuels Data Center 2022.

EIA, "Residential Energy Consumption Survey," Technical Report, United States Energy Information Administration 2022.

Einav, Liran, Amy Finkelstein, and Neale Mahoney, "The IO of selection markets," *Handbook of Industrial Organization*, 2021, *5* (1), 389–426.

_ , _ , Stephen P. Ryan, Paul Schrimpf, and Mark R. Cullen, "Selection on moral hazard in health insurance," *American Economic Review*, 2013, *103* (1), 178–219.

EPA, "Inventory of U.S. Greenhouse Gas Emissions and Sinks: 1990-2022," Technical Report, United States Environmental Protection Agency, EPA 430R-24004 2024.

Fabra, Natalia, David Rapson, Mar Reguant, and Jingyuan Wang, "Estimating the elasticity to real-time pricing: Evidence from the Spanish electricity market," *American Economic Review, Papers and Proceedings*, 2021, *111*, 425–429.

Fowlie, Meredith, Catherine Wolfram, Patrick Baylis, C. Anna Spurlock, Annika Todd-Blick, and Peter Cappers, "Default effects and follow-on behaviour: Evidence from an electricity pricing program," *Review of Economic Studies*, 2021, *88* (6), 2886–2934.

Grubb, Michael D and Matthew Osborne, "Cellular service demand: Biased beliefs, learning, and bill shock," *American Economic Review*, 2015, *105* (1), 234–271.

Hanemann, W. Michael, "Discrete/Continuous models of consumer demand," *Econometrica*, 1984, *52* (3), 541–561.

Harding, Matthew, Kyle Kettler, Carlos Lamarche, and Lala Ma, "The (alleged) environmental and social benefits of dynamic pricing," *Journal of Economic Behavior and Organization*, 2023, *205*, 574–593.

Hausman, Jerry A., "The econometrics of nonlinear budget sets," *Econometrica*, 1985, *53* (6), 1255–1282.

Hinchberger, Andrew J, Mark R Jacobsen, Christopher R Knittel, James M Sallee, and Arthur A van Benthem, "The Efficiency of Dynamic Electricity Prices," *NBER Working Paper #32995*, 2024.

Holland, Stephen P., Erin T. Mansur, Nicholas Z. Muller, and Andrew J. Yates, "Are There Environmental Benefits from Driving Electric Vehicles? The Importance of Local Factors," *American Economic Review*, 2016, *106* (12), 3700–3729.

Hortaçsu, Ali, Seyed Ali Madanizadeh, and Steven L. Puller, "Power to Choose? An Analysis of Consumer Inertia in the Residential Electricity Market," *American Economic Journal: Economic Policy*, November 2017, *9* (4), 192–226.

Ito, Koichiro, "Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing," *American Economic Review*, 2014, *104* (2), 537–563.

_ , Takanori Ida, and Makoto Tanaka, "Selection on welfare gains: experimental evidence from electricity plan choice," *American Economic Review*, 2023, *113* (11), 2937–2973.

Jessoe, Katrina and David Rapson, "Knowledge is (less) power: experimental evidnece from residential energy use," *American Economic Review*, 2014, *104* (4), 1417–1438.

Li, Shanjun, "Better lucky than rich? Welfare analysis of automobile licence allocations in Beijing and Shanghai," *The Review of Economic Studies*, 2018, *85* (4), 2389–2428.

Lin, Haizhen and Daniel W. Sacks, "Intertemporal substitution in health care demand: evidence from the RAND health insurance experiment," *Journal of Public Economics*, 2019, *175*, 29–43.

Lipsey, Richard G and Kelvin Lancaster, "The general theory of second best," *The review of economic studies*, 1956, *24* (1), 11–32.

Marone, Victoria R. and Adrienne Sabety, "When should there be vertical choice in health insurance markets?," *American Economic Review*, January 2022, *112* (1), 304–342.

Prest, Brian C., "Peaking interest: how awareness drives the effectiveness of time-of-use electricity pricing," *Journal of the Association of Environmental and Resource Economists*, 2019, *7* (1), 103–143.

Reiss, Peter C. and Matthew W. White, "Household electricity demand, revisited," *Review of Economic Studies*, 2005, *72* (3), 853–883.

Sallee, James M., "Rational inattention and energy efficiency," *Journal of Law and Economics*, 2014, *57* (3), 781–820.

Schittekatte, Tim, Dharik S. Mallapragada, Paul L. Joskow, and Richard Schmalensee, "Electricity retail rate design in a decarbonized economy: an analysis of time-of-use and critical peak pricing," *The Energy Journal*, 2024, *45* (3), 25–56.

Schwartz, Judith, "Case Study: Salt River Project–the Persistence of Consumer Choice," Technical Report, US Department of Energy and Federal Energy Regulatory Commission 2012.

Shaffer, Blake, "Misunderstanding Nonlinear Prices: Evidence from a Natural Experiment on Residential Electricity Demand," *American Economic Journal: Economic Policy*, August 2020, *12* (3), 433–61.

SRP, "Proposed Adjustments to SRP's Standard Electric Price Plans Effective with the May 2019 Billing Cycle," Technical Report, Salt River Project Agricultural Improvement and Power District 2018.

_ , "SRP 2035 Sustainability Goals, FY23 Progress Update," Technical Report, Salt River Project Agricultural Improvement and Power District 2023.

US EPA, "Report on the Social Cost of Greenhouse Gasses: Estimates Incorporating Re-

cent Scientific Advances, Docket ID No. EPA-HQ-OAR-2021-0317," Technical Report, US Environmental Protection Agency 2023.

Wolak, Frank A., "Do residential customers respond to hourly prices? Evidence from a dynamic pricing experiment," *American Economic Review, Papers and Proceedings*, 2011, *101* (3), 83–87.

Wooldridge, Jeffrey M., "Simple solutions to the initial conditions problem in dynamic, nonlinear panel data models with unobserved heterogeneity," *Journal of Applied Econometrics*, 2005, *20* (1), 39–54.

Yang, Jun, Avralt-Od Purevjav, and Shanjun Li, "The marginal cost of traffic congestion and road pricing: evidence from a natural experiment in Beijing," *American Economic Journal: Economic Policy*, 2020, *12* (1), 418–453.

Supplemental Appendices

A Additional Background on Plan Choice

A.1 Enrollment

SRP did not have a default plan for new customers during our study period. It encouraged consumers to choose plans that matched their preferences.[43] Figure A.1 shows the enrollment form that an existing customer could use to switch plans.[44] New customers were presented with a similar form, or they could choose a plan by talking with a customer service representative on the phone.

Enrolling in the Customer Generation Plan or the M-Power Plan required calling customer service to coordinate installation of additional metering equipment. A customer who clicked a "learn more" link would be presented with images explaining the plan's price function and intuitive summaries of the plan and how it differed from other options. For example, the linked description of the "EZ-3 Plan (3-6pm)" shown in Figure 1b included the following text:

How it works: *At certain times of the day, many people and businesses in the Valley are using energy, which puts pressure on the power grid. The SRP EZ-3 Price Plan offers price breaks for customers who can commit to shifting their energy use outside of those high-demand, or "on-peak," hours. Save money by limiting your energy use during on-peak hours, from either 3 to 6 p.m. or 4 to 7 p.m. on weekdays, year-round. Your reward? You'll pay lower off-peak prices all other hours, including weekends and six observed holidays. On-peak and off-peak pricing changes based on the season.*

How is EZ-3 different from other time-of-day plans? *The EZ-3 plan requires you*

[43] SRP has a stated goal of promoting consumer choice. The "Pricing Principles" of its rate design process include: "Choice—to constantly improve customer satisfaction through the creative design of pricing structures that reflect customers' different desires or abilities to manage the consumption, assume more price control, or demand differentiated products and services" (SRP, 2018).

[44] SRP offered a "risk-free" 90-day trial period for customers who switched from the block rate plan to a TOU plan. Customers were told: "*If your first three bills on [TOU plan name] aren't lower than what you would have paid with the [block plan name], we'll credit you the difference and switch you back to the [block plan name]*". While interesting, this is not an empirically important feature of our data. Of the 87 households in our main estimation sample that switched from the block rate to a TOU plan, only 1 switched back within 120 days.

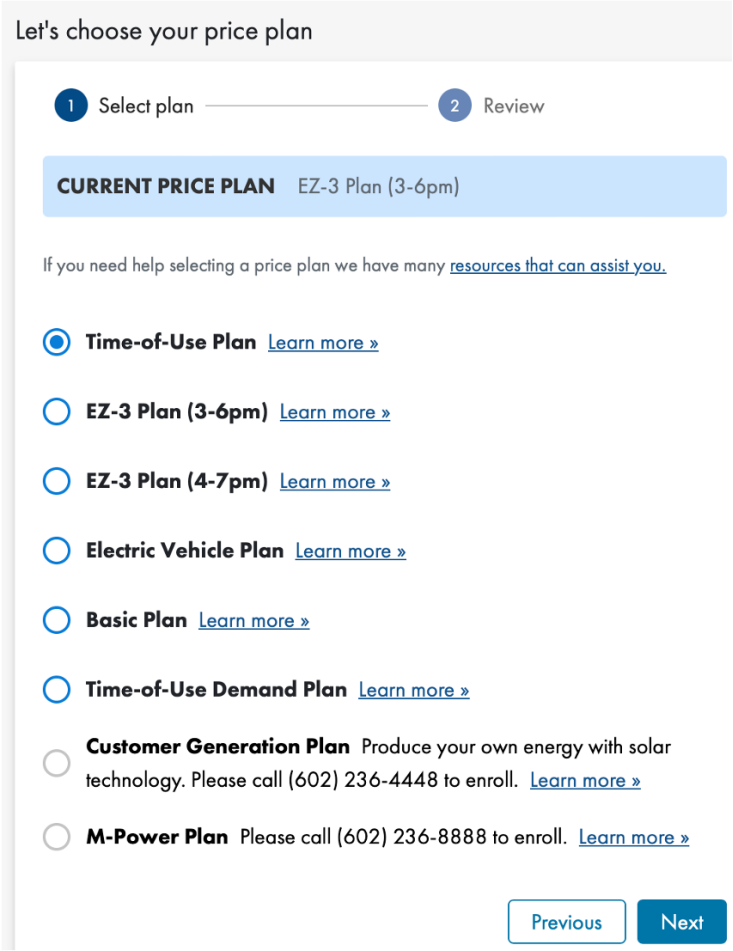


Figure A.1: Screen Shot of SRP's Online Enrollment Form

Note: This figure shows the enrollment form that existing customers could use to switch plans. New customers were presented with a similar form.

to limit your energy use during just three on-peak hours on weekdays. All other hours of the day, including weekends and six observed holidays, are billed at the lower off-peak rate. Other SRP plans, like our Time-of-Use plan, offer lower prices during off-peak hours. This can translate to greater cost savings, but to realize those savings, you'll need to limit your energy use during a larger window of on-peak hours.

A.2 Information about consumption and expenditures

After creating an online account, a customer could view their hourly history of consumption and expenditures. This information was posted within 48 hours. Figure A.2 provides an example of how this information could be viewed for a customer on the Superpeak 3pm-6pm

TOU plan. The left panel shows a particular day's usage and the right panel shows the associated hourly expenditures. Red and blue colors distinguish high price and low price hours. It is important to note that this screen shot is not derived from the confidential account data shared by SRP. Rather, it is based on one of our own accounts (Kuminoff).

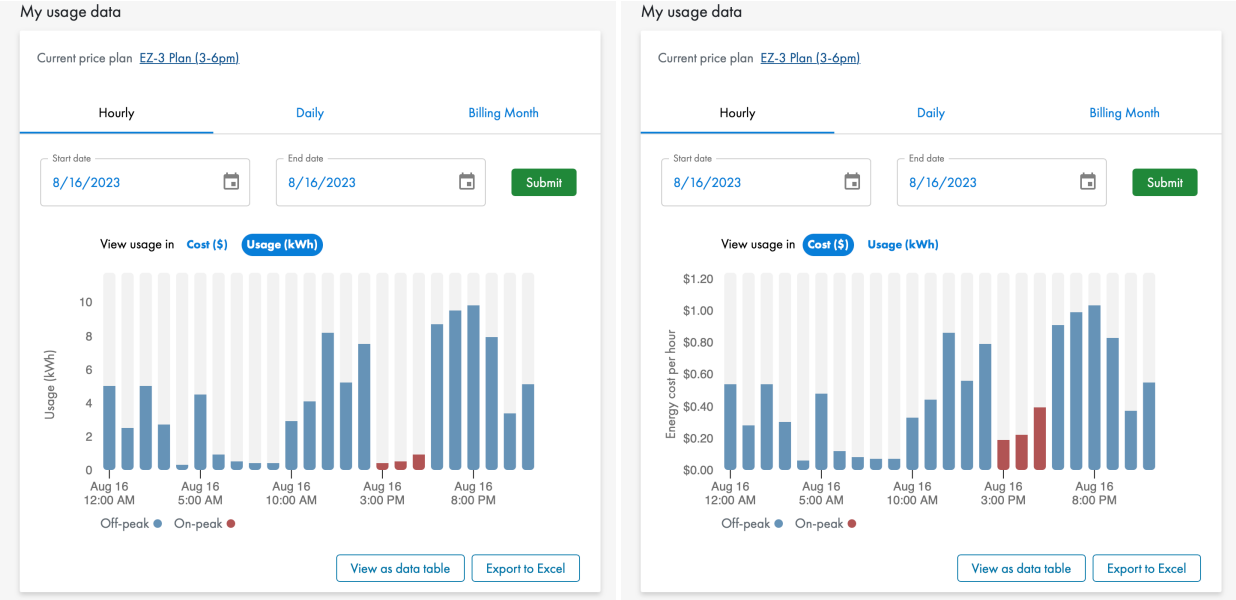


Figure A.2: Screen Shot of a Customer's Usage Data

Note: This figure shows hourly consumption on a particular date for a particular SRP customer (Kuminoff). It was not generated from confidential data provided by SRP.

Customers could also use their online accounts to see how much they would be projected to save (or lose) if they were to switch from their current plan to a different TOU or block-rate plan. These projections were based on the account's consumption history and presented as monthly and annual expenditures. Hence, the projections effectively measured ex post realized savings (or losses) from customers' prior enrollment decisions. Figure A.3 provides an example. Again, it is important to note that this screen shot is based on one of our own accounts (Kuminoff) and not derived from the confidential data provided by SRP.

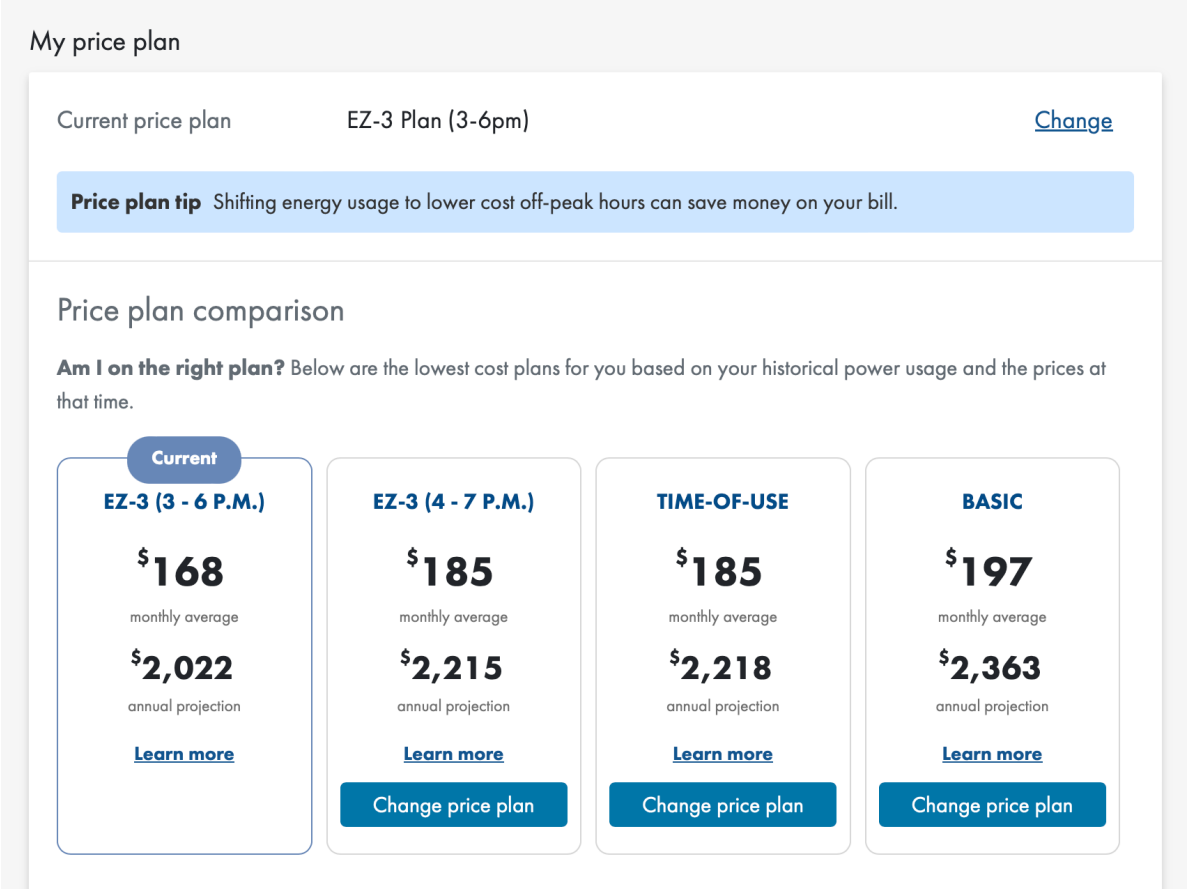


Figure A.3: Information about a Customer's Potential Savings from Switching Plans

Note: This figure shows how much a particular SRP customer (Kuminoff) would lose by switching plans given his household's consumption history. It was not generated from confidential data provided by SRP.

B Marginal Social Cost of Electricity Generation

Our approach to calculating the marginal social cost of electricity generation is based on Borenstein and Bushnell (2022) [henceforth BB2022]. We followed the methods from that study as closely as possible, drawing on the article's supplemental appendix and replication package. The main differences in our approach are that we collected more recent data and down-scaled their regional analysis to focus on SRP's service territory within the Southwestern region of the western interconnection section of the US electricity grid. We summarize our procedures in this section and direct readers to BB2022 for additional background on the underlying data sources, methods, and institutional features of the US electricity grid.

B.1 Marginal Private Cost

Our approach to calculating the marginal private cost (MPC) of generation started by compiling locational marginal price (LMP) data for SRP from the California Independent System Operator. We used the SRP-specific node DGAP_SRP_APN from the report PRC_RTPD_LMP for January 1, 2021 through December 31st, 2023. These data describe the wholesale prices at which SRP traded electricity on the grid at 15-minute increments. SRP's locational marginal price reflects its marginal cost of generating electricity plus the cost of transmission congestion and losses on high-voltage power lines.

Then we aggregated the LMPs by calendar month and hour and multiplied them by an inflation factor from BB2022 that is designed to adjust for the additional cost of distributing electricity on low-voltage power lines. Specifically, we extracted the month-by-hour inflation factors that BB2022 calculated for SRP using data for 2014-2016. We assume that SRP's month-by-hour percentage losses on low-voltage power lines were the same during our study period.

Figure B.2 shows our estimates for MPC by hour and month. Averaging these monthly data over SRP's winter, summer, and summer peak pricing seasons produces the top row of Figure 2.

B.2 Marginal External Cost

Burning fossil fuels to generate electricity emits air pollution as a byproduct of production. The pollutants include carbon dioxide (CO_2), sulfur dioxide (SO_2), nitrogen dioxide (NO_2), and fine particulate matter ($PM_{2.5}$). CO_2 contributes to climate change, whereas elevated exposures to SO_2, NO_2, and $PM_{2.5}$ are believed to increase human mortality risk. We estimate the marginal external costs (MEC) of these pollution externalities by down-scaling the regional analysis in BB2022 to focus on emissions generated from producing the electricity used by SRP customers.

A key feature of the BB2022 methodology is to account for the spatial transmission of electricity and air pollution externalities. SRP trades electricity on the western interconnection section of the United States electricity grid. The western interconnection includes Arizona, Washington, Oregon, California, Idaho, Nevada, Wyoming, Colorado, Utah, and

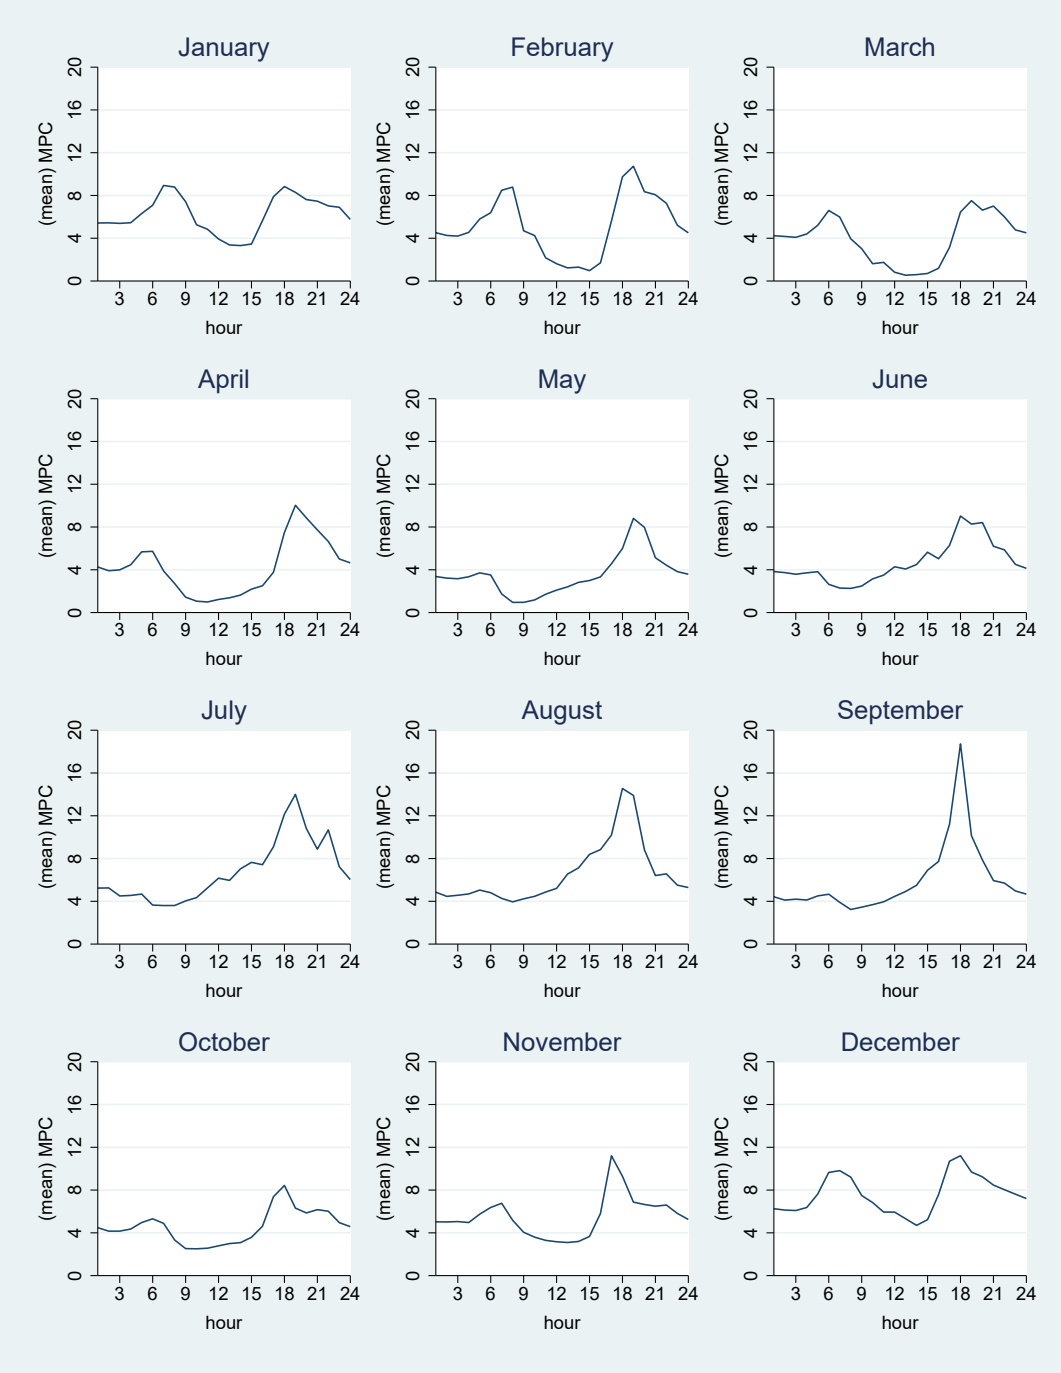


Figure B.1: Marginal Private Cost of Generation by Month and Hour (cents/kWh)

Note: Each panel shows our estimate for the marginal private cost of SRP's electricity generation, averaged by hour within the corresponding month. The vertical axis is measured as cents per kilowatt hour of generation.

parts of Montana, New Mexico, and South Dakota. This means that some of the electricity purchased by SRP customers is produced in other states, where it generates local pollution

externalities.

Our procedure for calculating pollution externalities started by using the 2024 release of the EPA's Emissions & Generation Resource Integrated Database (eGRID) database to determine the county in which each fossil fuel power plant is located. For each of these plants, we compiled data on hourly emissions of CO_2, SO_2, and NO_2 for 2018 through 2023 from the EPA's Continuous Emissions Monitoring System (CEMS) data, which we downloaded from the EPA's Clean Air Markets Program Data website: http://campd.epa.gov/.

$PM_{2.5}$ is not reported in the CEMS data. We addressed this data limitation by following a procedure that BB2022 adapted from Holland et al. (2016) to estimate plant-specific $PM_{2.5}$ emissions. First we collected plant-specific data on annual $PM_{2.5}$ emissions from the EPA's 2020 National Emissions Inventory. Next, we merged these data with plant-specific measures of heat input reported by CEMS. Then we calculated an average emissions rate for each plant by dividing annual $PM_{2.5}$ emissions by annual heat input. Finally, we multiplied a plant's annual average emissions rate by its hourly heat input to estimate its hourly $PM_{2.5}$ emissions.

We calculated the monetary damages from each plant's hourly emissions by feeding data on emissions, plant locations, and smokestack heights into the Pollution Emission Experiments and Policy Analysis Model (release AP3) summarized in BB2022. AP3 converts emissions into monetary damages in two steps. First, it uses an assumed social cost of carbon to monetize climate damages from CO_2. We set the social cost of carbon to \$224 per metric ton (in 2024 dollars) based on the EPA's latest estimates (US EPA, 2023). Second, AP3 uses an air dispersion model to predict how emissions of SO_2, NO_2, and $PM_{2.5}$ in one county contribute to ambient pollution in other counties, combined with an epidemiological model to predict how pollution exposure affects mortality rates, and data on population sizes and an assumption for the value per statistical life to convert a change in mortality rates into dollars. We followed BB2022 in using AP3's default values for NOx, SO2, and PM2.5 damages, converted to 2024 dollars using the CPI.

Next, we aggregated hourly plant-level damages into four geographic regions. We started by defining three regions that align with the US Energy Information Administration's (EIA) subdivision of the western interconnection into California, Northwest, and Southwest areas. SRP owns power plants in Arizona and New Mexico (which are part of Southwest) and in

Colorado (which is part of Northwest). We extracted data on all of the SRP-owned plants from the Southwest and Northwest areas and grouped them into a fourth region that we call SRP. The purpose for doing this was to account for the potentially tighter integration of generation at these plants with SRP's load. We compiled data on hourly load for utilities located in each of the three EIA regions and for SRP alone.[45] These data were originally collected from EIA Form 930.

To illustrate our approach to estimating external damages, let D_{prt} denote monetary damages from aggregate emissions of pollutant p in region r during hour t. Equation (B.1) shows how we model damages as a function of hourly load in that region, hourly load in the other three regions, $s \neq t$, and time fixed effects, ϕ. The parametric functions $f(\cdot)$ and $g(\cdot)$ are spline functions that allow the marginal effect of hourly load to vary by quantile of the load distribution in a given region.

$$D_{prt} = f(l_{r,t}; \beta) + g(l_{s,t}; \gamma) + \phi_{prt} + \epsilon_{prt} \tag{B.1}$$

We follow BB2022 in stacking the data over all pollutants and regions, taking 24-hour differences prior to estimation, and clustering by date. Taking 24-hour differences purges the fixed effects, and stacking the data to estimate all model parameters simultaneously ensures that standard errors account for correlations across regions and pollutants. Finally, we parameterized the spline functions by using quintiles of the load distribution, and verified that our results were broadly robust to marginally increasing or reducing flexibility.

Table B.1 shows our pollutant-specific estimates for the marginal external damage per megawatt hour of SRP load. We applied these coefficients to hourly data on SRP load to measure hourly damages throughout our study period. Then we applied the inflation factors from BB2022 to account for distribution losses. Finally, we calculated the month-by-hour specific measures of MEC shown in Figure B.2. Averaging these monthly data over SRP's winter, summer, and summer peak pricing seasons produces the middle row of Figure 2.

[45] We subtract SRP's load from the Southwest region where its service territory is located.

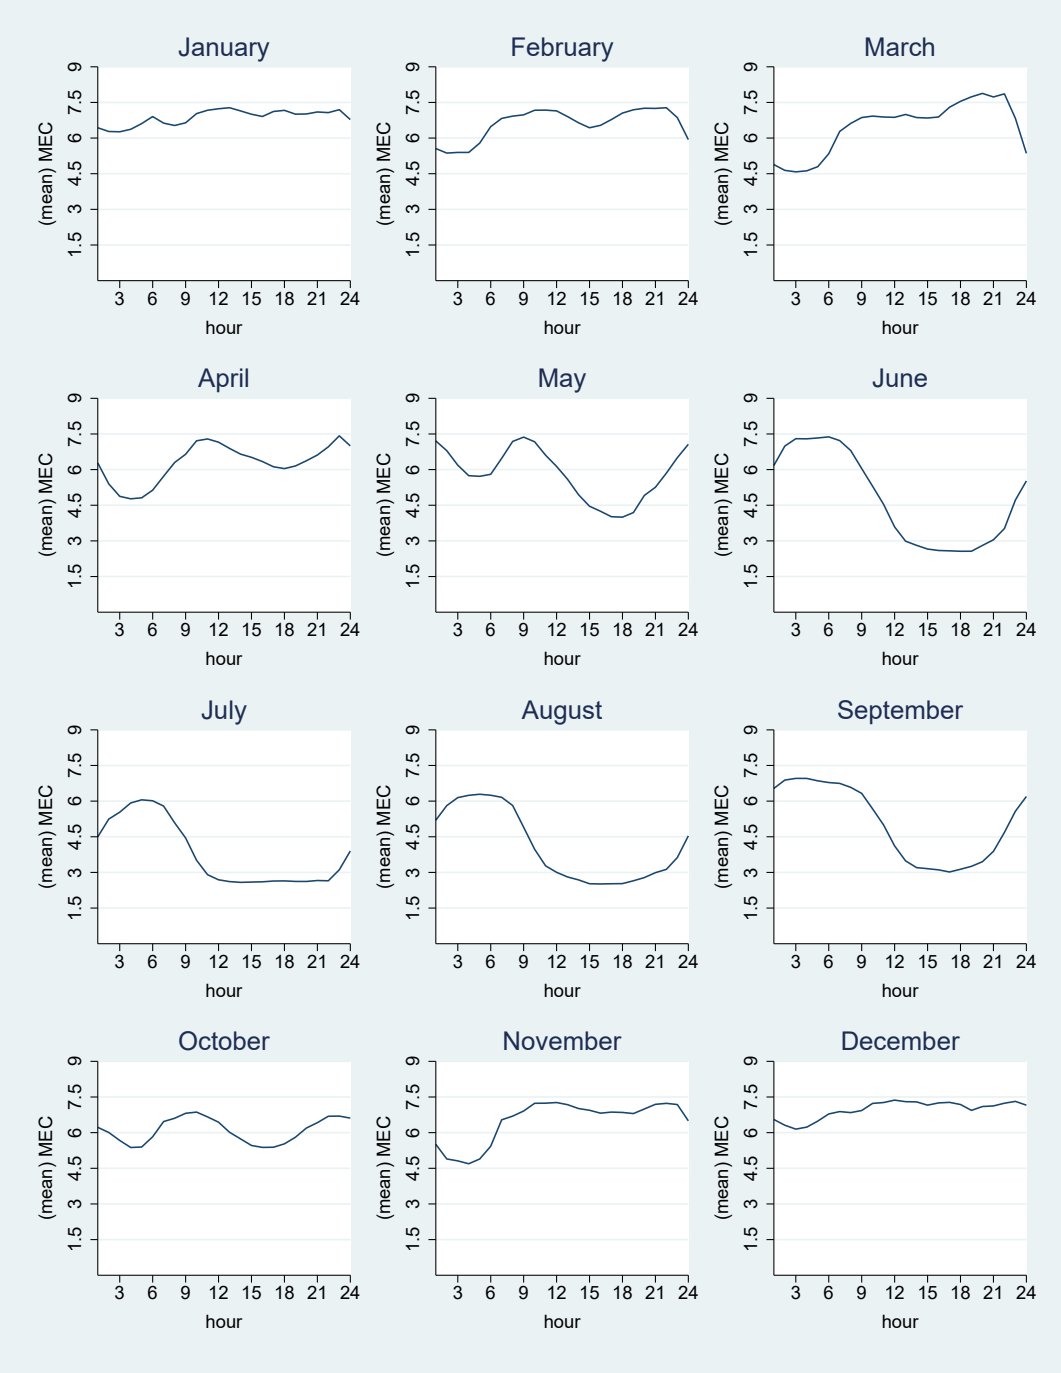


Figure B.2: Marginal External Cost of Generation by Month and Hour (cents/kWh)

Note: Each panel shows our estimate for the marginal external cost of SRP's electricity generation, averaged by hour within the corresponding month. The vertical axis is measured as cents per kilowatt hour of generation.

Table B.1: Marginal External Damages from SRP Load ($ per mWh)

load quintile	Carbon Dioxide	Nitrogen Dioxide	Sulfer Dioxide	Fine Particulate Matter
1	37.42	1.65	0.58	2.56
	(12.76)	(0.45)	(0.19)	(0.95)
2	72.49	1.03	0.86	0.48
	(10.36)	(0.54)	(0.15)	(1.14)
3	70.29	0.80	0.87	1.37
	(8.95)	(0.33)	(0.12)	(1.01)
4	52.61	0.82	0.77	0.61
	(6.25)	(0.28)	(0.09)	(0.73)
5	20.75	0.64	0.37	0.50
	(1.45)	(0.05)	(0.03)	(0.06)

Note: The table reports regression coefficients and standard errors from local linear regression of pollutant-specific damages on quintiles of SRP load. The parameters are estimated simultaneously with three other generation areas: California, the Northwest, and the Southwest excluding SRP. Coefficients for non-SRP areas are supressed for brevity. Standard errors are clustered by date.

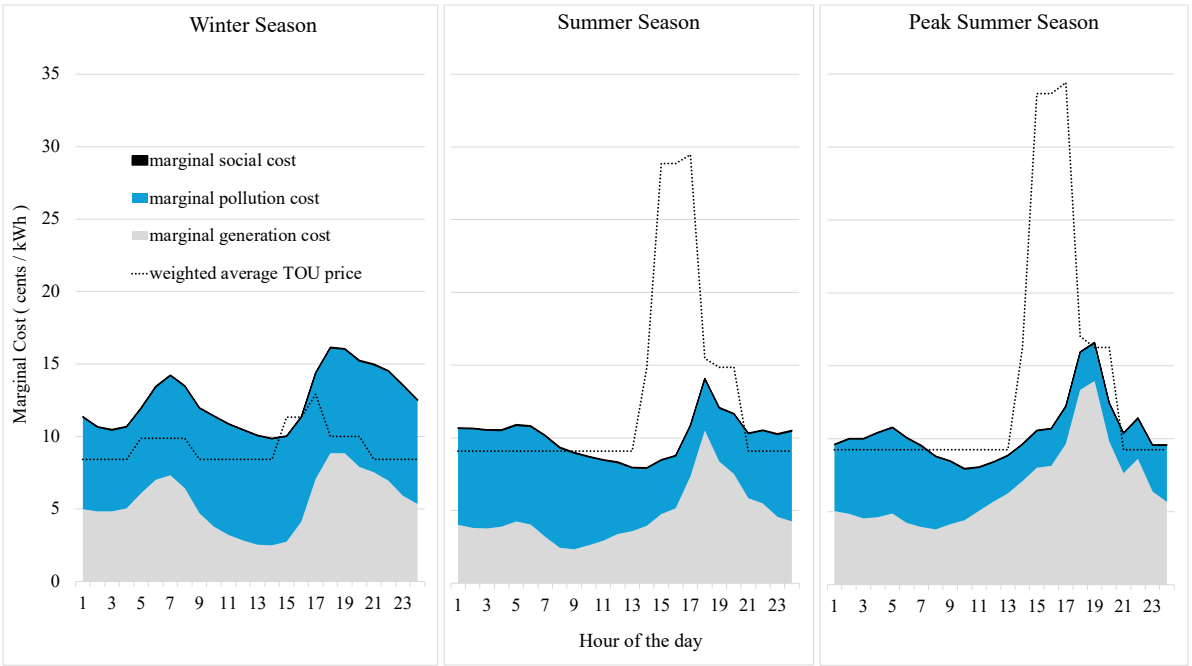


Figure B.3: Marginal Social Costs of Generation and Average TOU Price (cents/kWh)

Note: The grey area shows the marginal private cost of generation, averaged by hour within each pricing season. The blue area shows the marginal external cost of pollution. Adding them yields the marginal social cost, shown as a black line at the top of each panel. The dotted line is the average hourly time-of-use price, calculated over the 2-8pm, 3-6pm, and 4-7pm plans, weighted by enrollment.

C Counterfactual Bill Calculator

We use the smart meter data, plus the rate books and tax data, to compute counterfactual bills for households if they switched to other plans. We describe in this section the data sources for these components and how we validate our bill calculator.

SRP Rate Tables We determined the fixed cost and marginal cost of consumption on each electricity price plan using SRP's rate books. The first set of rates was active from April 2015 through April 2019. The Second set of rates was active from May 2019 through the end of our study period.

Salt River Project Agricultural Improvement and Power District, "Standard Electric Price Plans: Prices Effective with the April 2015 Billing Cycle". Technical Report. 2015.

Salt River Project Agricultural Improvement and Power District, "SRP's Standard Electric Price Plans Effective with the May 2019 Billing Cycle". Technical Report. 2019.

Tax Rates for Electricity Data on state, county, and city tax rates for electricity were drawn from the Arizona Department of Revenue's "Transaction Privilege and Other Tax Rate Tables" effective April 1, 2023. The combined tax rate for the state of Arizona and Maricopa county was 6.3% throughout our study period. The vast majority of households in our data live in Maricopa county, which includes almost all of the Phoenix Metropolitan Area. Tax rates for other counties varied by no more than four tenths of a percentage point. City taxes varied from 0% to 4%. The data can be found here: https://azdor.gov/business/transaction-privilege-tax/tax-rate-table

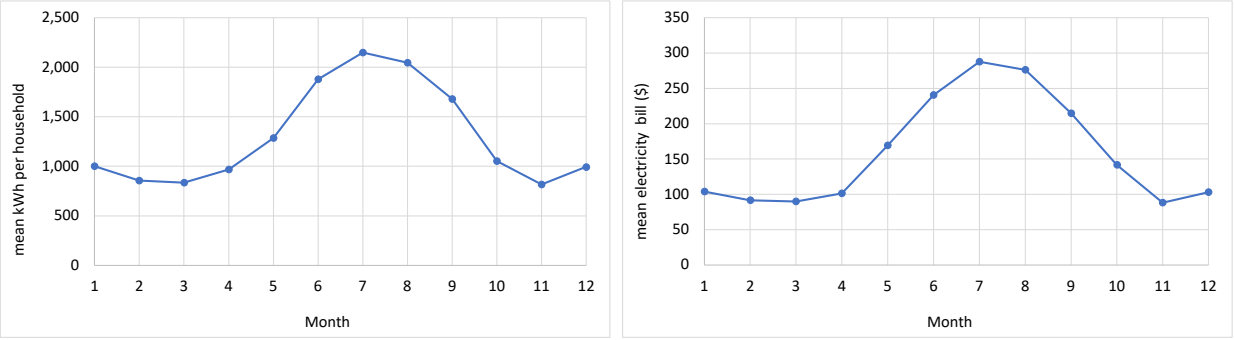
Bill calculator validation We validate our bill calculator by comparing its predictions against data from SRP on each household's consumption and expenditures from May 2022 through April 2023. Correlation coefficients between predicted and observed measures of consumption and expenditures are 0.991 and 0.985 respectively.[46]

[46] A univariate regression of actual expenditures on predicted expenditures yields an R^2 of 0.97 and a slope coefficient of 1.05. It makes sense for our slope coefficient to exceed one because our predicted bills exclude ad hoc fees that we do not observe, including fees for service establishment, returned payments, field visits, high-bill audits, theft investigation, replacement of customer-damaged equipment, and late payments.

D Additional Descriptive Evidence

D.1 Summary Statistics

Figure D.1a shows average monthly electricity consumption per household and Figure D.1b shows the associated expenditures. The shapes of the curves reflect cooling costs. The average household uses 2.6 times as much electricity in July as in November, and its July bill is 3.3 times its November bill. The July-to-November billing ratio is higher than the consumption ratio becuase prices are also higher in the summer.



(a) average monthly consumption (kWh) (b) average monthly bill ($)

Figure D.1: Average monthly electricity consumption and expenditures

SRP obtained household demographic data from an external contractor. Table D.1 reports additional summary statistics for the households enrolled in each plan. Households in the pre-pay plan have substantially lower credit scores.[47] This is consistent with their relatively lower incomes (Table 1). They are also less likely to be over age 65 and more likely to have created online accounts with SRP.

Overall, 72% of households have online accounts. This enables them to get monthly billing statements by email, pay bills online, and log in to check their consumption history. Households on the pre-pay plan can also use their online accounts to monitor their account balances and deposit money as needed using a plan-specific app.

The last row of Table D.1 shows the fractions of households in each plan who opted in to SRP's "bring your own thermostat" demand response program (BYOT). Enrollees receive

[47] The 85% of households that we code as having satisfactory credit all having ratings of "preferred" or "satisfactory" in the credit score shared by SRP. The other 14% of households have ratings of "slow", "unsatisfactory", "cash only", or "new".

	All plans	block rate	fixed rate	2-8pm TOU	3-6pm TOU	4-7pm TOU
market share (%)	100	56	3	17	23	1
# monthly bills	261,331	146,084	8,081	45,667	58,816	2,683
satisfactory credit (%)	85	87	21	91	85	88
over 65 (%)	31	34	9	35	21	29
online account (%)	72	66	91	72	85	78
BYOT (%)	6	4	2	9	11	8

Table D.1: Summary Statistics by Plan

Note: The table reports summary statistics by price plan. An observation is a household-month.

an initial lump sum payment for participating ($50 per smart thermostat) plus a bill credit of $25 per thermostat at the end of each summer season. In exchange, SRP is able to adjust their thermostat settings for a few hours during up to 15 peak demand days during the summer months. Customers receive advance notice of each event and have the ability to override temporary changes that SRP makes to their thermostats if the changes cause them discomfort. Thus, the program allows SRP to reduce peak demand by raising thermostat settings of households who are not inconvenienced by the adjustments, for example, because they are not at home.[48]

We exclude the BYOT program from our bill calculators because we lack data on the history of peak demand events. There is minimal scope for this to affect our findings because peak demand events are rare, the fraction of households participating is low (6%), and the program incentives are neutral to plan choice. Nevertheless, the heterogeneity in participation rates across plans is notably consistent with higher rates of smart thermostat ownership among households in TOU plans.

D.2 Additional Descriptive Evidence

Figure D.2 summarizes our estimates for how households who selected into each TOU plan responded to hourly price changes on peak pricing days of each pricing season.[49] As in Figure 3a, we highlight heterogeneity by reporting mean responses for deciles of consumers, ranked by the average estimated response during peak hours. The figures show substantial

[48]Blonz et al. (2024) evaluates a similar program that automates thermostat settings to adjust to time-of-use pricing.

[49]The noisier estimates for the 4pm-7pm plan in the middle row reflect its smaller enrollment share.

heterogeneity in responsiveness. Most households do not adjust consumption substantially. However, the top deciles reduce consumption during peak hours. During summer months, some consumption is shifted to off-peak hours, especially after peak pricing periods end.

In comparison to the two superpeak plans, consumers in the standard TOU plan shift less load to off-peak periods. This may be partly explained by the fact that it is driven by a relatively smaller price increase, as shown in Figure 1. For the most responsive decile in the Standard TOU plan, consumption drops sharply at the start of the on-peak period and then trends up as the period continues. This suggests that the disutility of load-shifting may increase with the distance over which load is shifted. For example, as indoor temperatures rise further above a household's bliss point we expect its marginal utility of air conditioning to increase.

Figure D.3 plots plan-by-season mean load curves for households that we define as TOU-responsive and TOU-unresponsive based on paired-sample t-tests described in section 4.1. Each panel shows kilowatts on the vertical axis and hour of the day on the horizontal axis. The curves are based on data for peak pricing days only (i.e. weekdays, excluding holidays).

Figure D.4 plots plan-specific mean load curves during the peak summer pricing season for TOU-responsive and TOU-unresponsive households. Each panel shows kilowatts on the vertical axis and hour of the day on the horizontal axis. The curves in the left column are based on days when peak pricing rates were applied (i.e. weekdays, excluding holidays). The curves in the right column are based on days when peak rates were not applied (i.e. weekends and holidays).

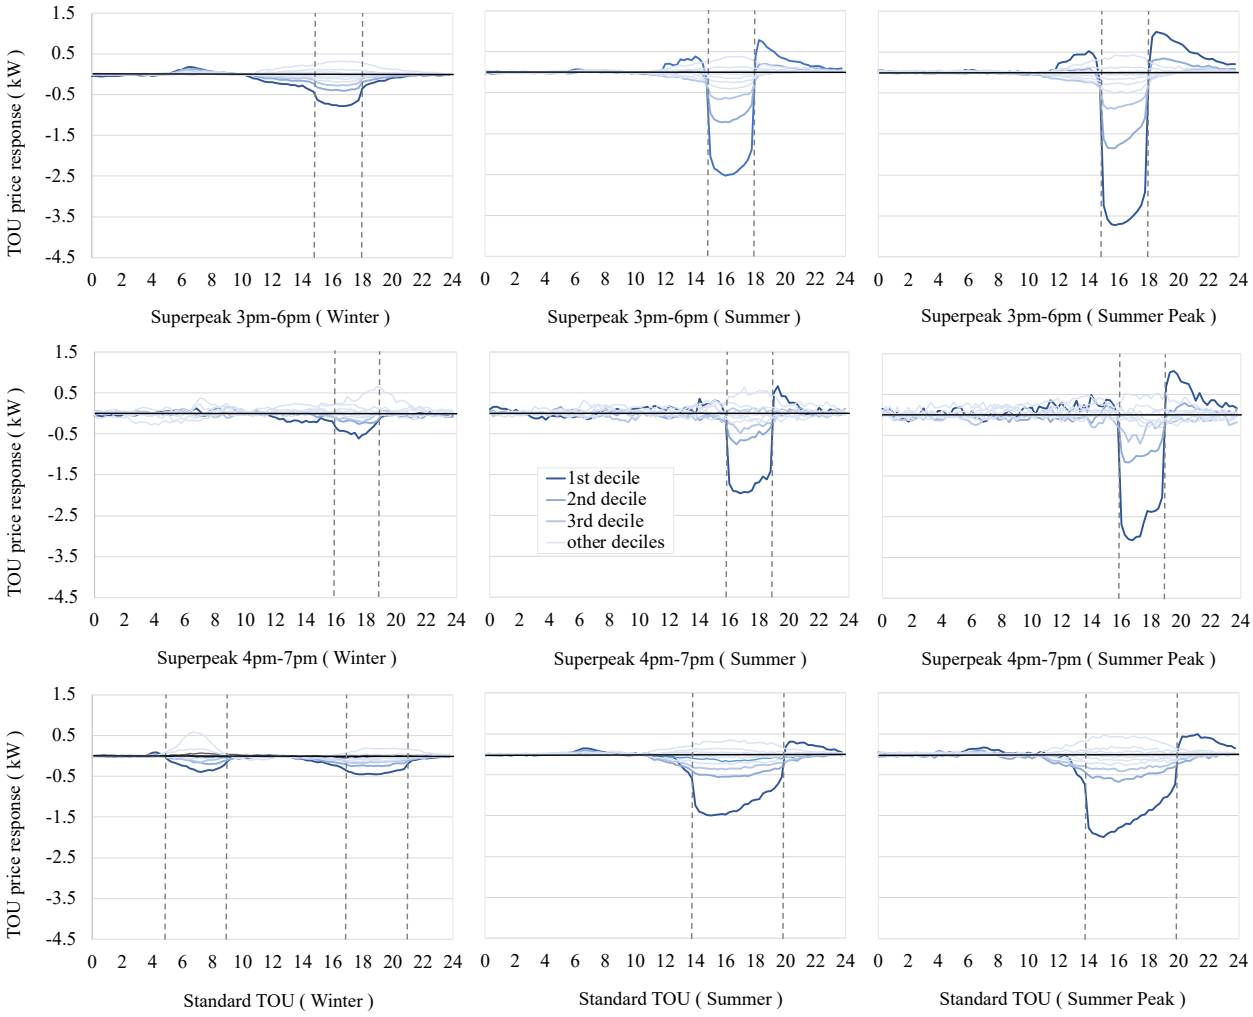


Figure D.2: TOU responsiveness by plan and season

Note: Each panel shows the estimated household response to TOU pricing for a particular plan and pricing season. Each curve corresponds to a decile of households, ranked by the size of the estimated shift during peak hours. Dotted vertical lines delineate peak hours.

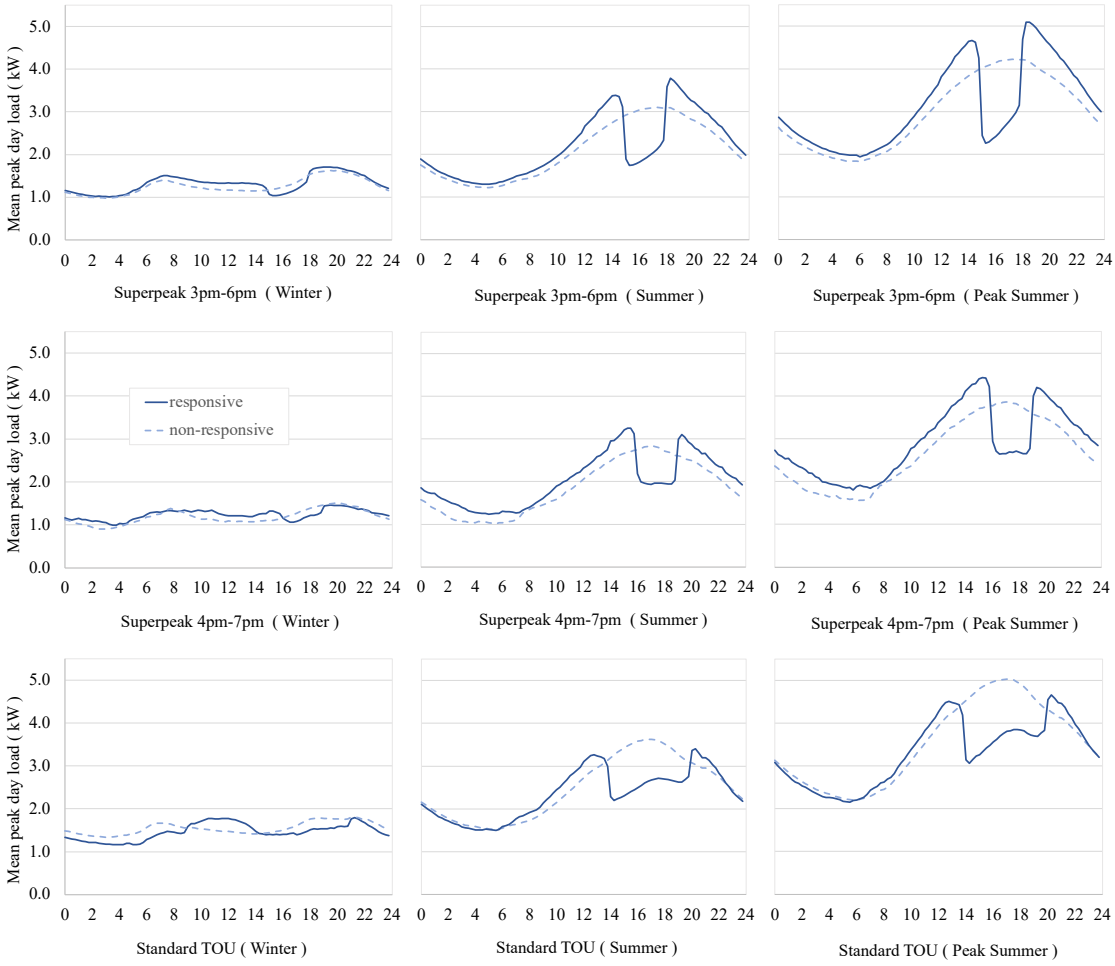


Figure D.3: Load curves by season and responsiveness to TOU pricing

Note: Each panel shows weekday load curves for a particular TOU plan and pricing season. Vertical axes measure the load in kilowatts. Horizontal axes measure the hour of the day. Solid lines show the mean load for TOU-responsive households. Dashed lines show the mean load for TOU-unresponsive households.

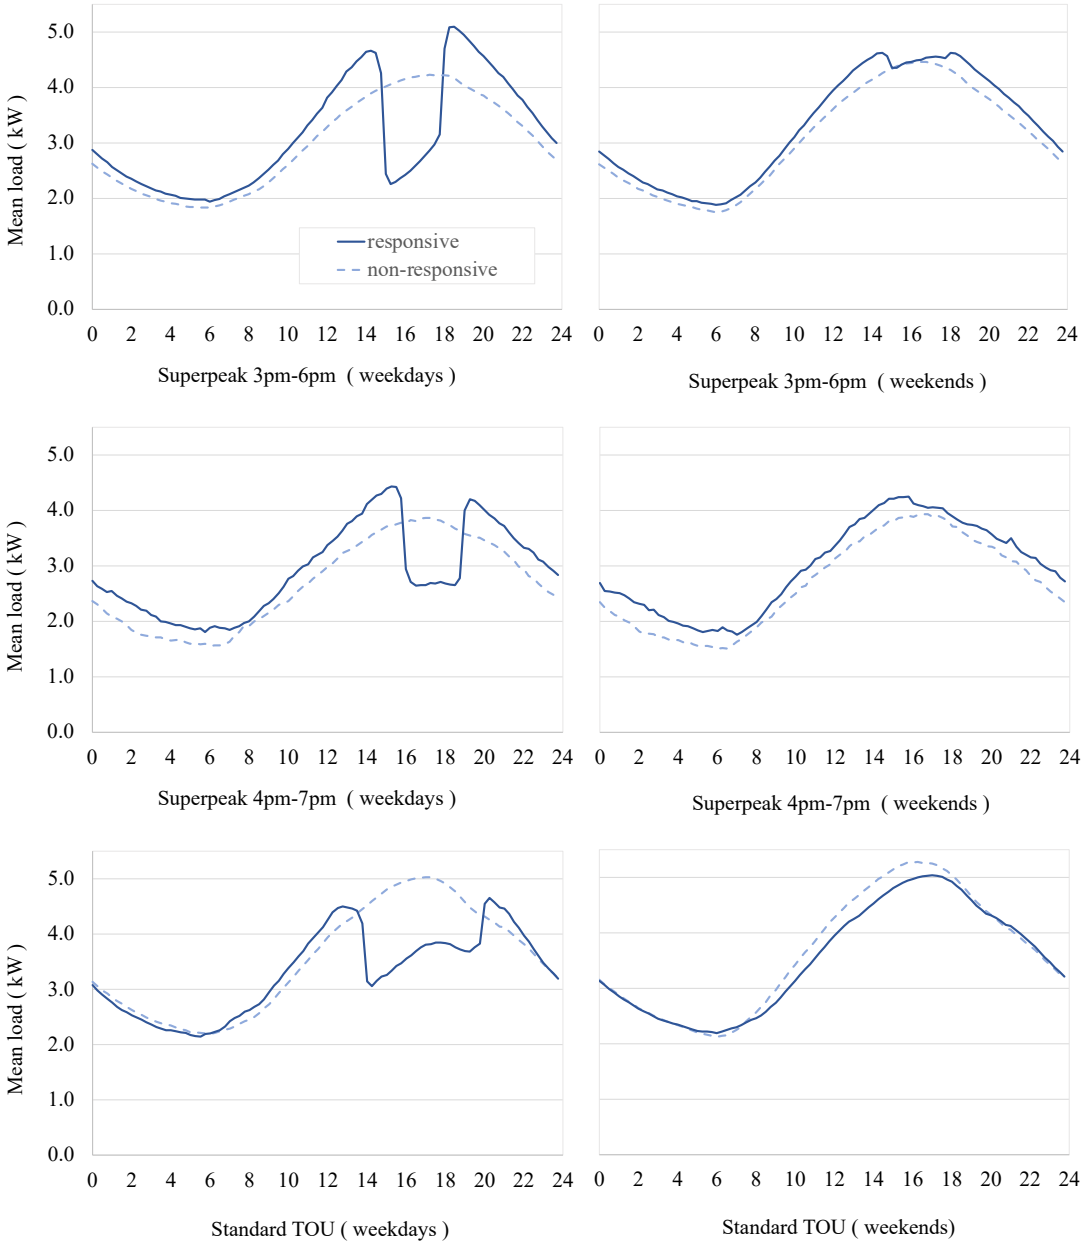


Figure D.4: Load curves by peak day and responsiveness to TOU pricing

Note: Each panel shows load curves for a particular TOU plan during the peak summer pricing season. The left column panels show days when peak rates were applied. The right column panels show days when peak rates were not applied. Vertical axes measure the load in kilowatts. Horizontal axes measure the hour of the day. Solid lines show the mean load for TOU-responsive households. Dashed lines show the mean load for TOU-unresponsive households.

E Additional Model Details

E.1 First-order-conditions

TOU-responsive households For the TOU-responsive households, the first-order-conditions are (noting that we replace t' with h as the index of the summations to differentiate it from the period t' the load is being shifted to):

$$\nu_{imwt} - q_{ijmwt} - \Sigma_{h \neq t} s_{ijmwth} = \omega_i p_{jmwt} \qquad [q_{ijmwt}] \qquad \text{(E.1)}$$

$$\nu_{imwt} - q_{ijmwt} - \Sigma_{h \neq t} s_{ijmwth} = \omega_i p_{jmwt'} + \omega_i d'_{imtt'}(s_{ijmwtt'}) \qquad [s_{ijmwtt'}] \qquad \text{(E.2)}$$

Note that the above equations have different prices on the right-hand-side (the first equation uses the price at time t, the second equation uses the price at time t'). Manipulating these first order conditions and comparing t within a TOU period and t' outside it provides the following equations which we later use for identification and estimation:

$$(\nu_{imwt} - \nu_{imwt'}) - (q_{ijmwt} - q_{ijmwt'}) - \Sigma_{h \neq t} s_{ijmwth} = \omega_i (p_{jmwt} - p_{jmwt'}) \qquad \text{(E.3)}$$

$$d'_{imtt'}(s_{ijmwtt'}) = p_{jmwt} - p_{jmwt'} \qquad \text{(E.4)}$$

TOU-nonresponsive households The first order condition for TOU-nonresponsive households is:

$$\nu_{imwt} - q_{ijmwt} = \omega_i \bar{p}_{ijm} \qquad [q_{ijmwt}] \qquad \text{(E.5)}$$

Note that we do not incorporate feedback effects on the average price in this equation (i.e. $\bar{p}_{ijm}$ is technically a function of q_{ijmwt}). The reason is that these effects are small and so do not have a first order impact on consumer choice. Concretely, they correspond to how a change in consumption in a particular 15 minute interval affects the average price in a month (which is constructed from numerous (2880) such intervals).

E.2 Computational algorithm

TOU-responsive households Here we set out how we compute the equilibrium consumption q_{ijmwt} and load shifting $s_{ijmwtt'}$. We compute separately weekdays and weekends for

each month, since the household's problem is separable across days. A challenge is ensuring that $q_{ijmwt} \geq 0$ given the total amount of load shifting from that period i.e. ensuring that the amount of load shifted is not so high that consumption in that period is negative.[50] For each household i:

1. Initialize a guess at iteration $k = 0$ of q^k_{ijmwt} and $s^k_{ijmwtt'}$.

2. Within the TOU period, get the total load shifted at each t in a peak period to a non-peak period t' in the following way (using the first-order-conditions):

$$\begin{aligned}\hat{x}^k_{ijmwtt'} &= p_{jmwt} - p_{jmwt'} \\ &\quad - \left(\beta_{d0j(i)} + \beta_{d0m} + \beta_{d1}|t_{\text{midpoint}} - t'| + \beta_{d2}1[\text{t' in 9-5pm}] + \beta_{d3}X_i\right) \qquad \text{(E.6)} \\ s^{k+1}_{ijmwtt'} &= \frac{1}{\beta_{d4}}\left(\hat{x}^k_{ijmwtt'} - \text{Penalty}^k_{ijmwt}\right) \qquad \text{(E.7)}\end{aligned}$$

   Here $\text{Penalty}^k_{ijmwt} = \mu(\epsilon_1/q^k_{ijmwt}) + (1-\mu)(\epsilon_1/q^{k-1}_{ijmwt})$. This term reduces load shifting smoothly as $q^k_{ijmwt} \to 0$ i.e. when consumption in a period is close to its bound. We also include a dampening weight term $\mu = 0.05$ to assist with convergence across iterations.

   Note: we set ϵ_1 to be very small and so the penalty term is typically second order. As $\epsilon_1 \to 0$ the penalty term is eliminated completely. Therefore, we check what happens if we set it to an order of magnitude lower $(1/10)\epsilon_1$ and the results are almost unchanged, illustrating that our results are robust to the choice of this tuning parameter. The penalty term is most useful for ensuring numerical convergence in (i) regions of the parameter space far away from the current parameters (ii) occasional cases where one of our thousands of consumers get very low draws of bliss points.

3. Get the update of consumption (using the first-order-conditions):

$$q^{k+1}_{ijmwt} = \text{SmoothBound}\left(\nu_{imwt} - \omega_i p_{jmwt} - \Sigma_{h\neq t} s^k_{ijmwth}\right) \qquad \text{(E.8)}$$

   Here, $\text{SmoothBound}(x) = 0.5x + 0.5\sqrt{x^2 + \epsilon_2}$. Note that as $\epsilon_2 \to 0$ then $\text{SmoothBound}(x) \to$

[50] Typically, this constraint is far from binding. However, there are edge cases that can cause numerical issues: regions of the parameter space far away from the estimated parameters, and occasional cases where one of our thousands of consumers get very low draws of bliss points.

$0.5x + 0.5|x| = \max\{0, x\}$. Therefore, we check what happens if we set ϵ_2 to an order of magnitude lower $(1/10)\epsilon_2$ and the results are almost unchanged, illustrating that our results are robust to the choice of this tuning parameter. Again, this is most useful for ensuring numerical convergence in (i) regions of the parameter space far away from the current parameters (ii) occasional cases where one of our thousands of consumers get very low draws of bliss points.

Note that $\Sigma_{h \neq t} s^k_{ijmwth} = 0$ if t is outside a peak period. Also note that we use a 'dampened' update to q^{k+1}_{ijmwt} in practice for numerical reasons.

4. Compute the convergence criterion $\max_t |q^{k+1}_{ijmwt} - q^k_{ijmwt}|$. If it is too large, continue again from Step 2.

TOU-non-responsive households The algorithm for these households is similar to the above, except: (i) we skip Step 2 (since there is no load-shifting i.e. $s_{ijmwtt'} = 0$) (ii) we solve for weekdays and weekends jointly in each month, and include an extra outer loop where we update the monthly average price for each consumer given the current iterations of their consumption choices on weekdays and weekends.

E.3 Formal details on identification

Identifying the marginal value of income ω_i More formally, consider the difference in consumption just before and after a TOU price increase. Denote the pre-period by l and the post-period by h to indicate a low or high price. Then, focusing on a single individual and suppressing the model's i, w, m subscripts to simplify notation, the difference in consumption can be written as:

$$(q_l + \Sigma_{l \neq l'} s_{l'l}) - (q_h + \Sigma_{h' \neq h} s_{h'h}) = (\nu_l - \nu_h) + \omega p_h - \omega p_l + \Sigma_{l \neq l'} s_{l'l} \tag{E.9}$$

$$= \omega(p_h - p_l) + \Sigma_{l \neq l'} s_{l'l}, \tag{E.10}$$

The right-hand-side of Equation (E.9) follows from inverting the first-order optimality condition that $\nu_t - q_t - \Sigma_{t' \neq t} s_{tt'} = \omega p_t$.[51] Equation (E.10) follows from our assumption that

[51] For the l period, since no load is shifted away, this optimality condition is $\nu_l - q_l = \omega p_l$. For the h period, this optimality condition is $\nu_h - q_h - \Sigma_{h' \neq h} s_{hh'} = \omega p_h$. For expositional clarity we assume that the nonnegativity conditions do not bind in these periods.

$\nu_l = \nu_h$ just before and after the sharp price change. Note that $\Sigma_{l \neq l'} s_{l'l}$, the total load shifted to period l, can be directly constructed from the data using the load-shifting estimator in Section 4.1. Therefore, the only unknown component is ω.

Bliss points can then be identified at the consumer level for every 15 minute interval during the day. Formally, rewrite the first-order-condition with respect to q_t yields $\nu_t = q_t + \Sigma_{t' \neq t} s_{tt'} + \omega p_t$. All three terms on the right-hand-size are known or can be constructed from the data.

Identifying the plan choice parameters Consider a consumer who chooses between their current (default) plan, j, and two alternatives, j' and j''. The probability of switching to j' equals the probability of receiving an attention shock (triggering the consumer to reconsider their prior plan choice) times the conditional probability of choosing to switch:

$$P(\text{switch } j \to j') = a \frac{\exp(V_{j'}/\sigma_\epsilon)}{\Sigma_k \exp((V_k + 1[j=k]\gamma)/\sigma_\epsilon)} \tag{E.11}$$

Comparing the probabilities of switching to j' and j'' identifies the σ_ϵ parameter such that $\ln(P(\text{switch } j \to j')) - \ln(P(\text{switch } j \to j'')) = (1/\sigma_\epsilon)(V_{j'} - V_{j''})$. Intuitively, since we know $V_{j'}$ and $V_{j''}$, and observe switching probabilities in the data, σ_ϵ indexes the degree to which consumers choose their best option conditional on switching.

Comparing the probabilities of switching out of different default plans identifies inertia (γ) separately from attention (a). This variation could arise at the consumer level from the data's panel structure, after controlling for factors that shift the consumer's attention over time (e.g. exploiting variation across years within a particular month).[52] Intuitively, if we observe the same consumer choosing from the same plans when they are enrolled in different default plans — all else equal — that induces variation in how choices are distorted by inertia, while keeping the probability of an attention shock fixed. More formally, consider the same

[52] Although these arguments are at the consumer i level, in practice the data also contain observationally identical consumers who are enrolled in different default plans in the same month.

consumer in different plans, and compare the switching probabilities in the following way:

$$\frac{P(\text{switch } j \to j) - P(\text{switch } j' \to j')}{P(\text{switch } j' \to j)}$$

$$= \exp\left(\frac{V_j - V_{j'} + \gamma}{\sigma_\epsilon}\right) \cdot \frac{\Sigma_k \exp((V_k + 1[j' = k]\gamma)/\sigma_\epsilon)}{\Sigma_k \exp((V_k + 1[j = k]\gamma)/\sigma_\epsilon)} + \exp\left(\frac{V_{j'} - V_j + \gamma}{\sigma_\epsilon}\right) \tag{E.12}$$

The expression to the left of the equality is data, and the expression to the right is known up to γ, so it is identified. The attention parameter can then be recovered by matching the observed probability of any choice, given the remaining parameters. Intuitively, we can think of lowering a from $a = 1$ (full attention) until the model matches the probability of switching between j and any other $j' \neq j$.

E.4 Externally Estimated Parameters

We estimate the parameters governing price sensitivity for the average price responsive consumers outside the model. To do so, we focus on households that switched plans across years and determine how their monthly consumption varies for a change in average price. The rational for focusing exclusively on plan switchers is that, in a given calendar month, households that do not change plans face very little variation in average prices across years since the price schedules are fixed throughout our study. Therefore, only households that actively switch plans will face meaningful changes in monthly average prices.

Our estimation strategy is based on the observation that, manipulating the first order conditions in Appendix Section E.1 for TOU non-responsive consumers and assuming that the bliss points are comparable before and after the switch (which is justified by time and household controls):

$$q_{ijmwt} - q_{ij'mwt} = \omega_i(\bar{p}_{ij'm} - \bar{p}_{ijm}) \tag{E.13}$$

Where j is the plan before the switch, and j' is the plan after the switch.

After limiting the sample to households that are not responsive to TOU pricing, that have switched plans in the same calendar month across years, and were initially in a TOU

Table E.1: TOU non-responsive consumer price sensitivity

	(1)
$\bar{P}_{i,m,y}$	-1.967**
	(0.934)
$\bar{P}_{i,m,y}$ × 4-7 pm plan	4.464***
	(0.954)
$\bar{P}_{i,m,y}$ × 2-8 pm plan	2.357***
	(0.589)
$\bar{P}_{i,m,y}$ × income	-22.163***
	(7.642)
$\bar{P}_{i,m,y}$ × HH size	-0.727*
	(0.392)
Constant	0.909***
	(0.042)
Overall adjusted R-squared	0.92
N	2693

Note.— Standard errors clustered at the household level.
* $(p < .1)$,** $(p < .05)$,*** $(p < .01)$

plan, we estimate the following regression:

$$\begin{aligned} C_{i,m,y} = \beta_0 + \beta_1 \bar{P}_{i,m,y} + \beta_2 \bar{P}_{i,m,y} \times Superpeak\ 4pm - 7pm_i \\ + \beta_3 \bar{P}_{i,m,y} \times Standard\ TOU_i + \beta_4 \bar{P}_{i,m,y} \times Income_i \\ + \beta_5 \bar{P}_{i,m,y} \times Household\ size_i + \rho_i + \psi_m + \phi_y + \delta_{i,m} \\ + \pi_{m,y} + \varepsilon_{i,y,m}, \end{aligned} \tag{E.14}$$

where $C_{i,y,m}$ is household i's electricity consumption (in kWh) for calendar month m in year y rescaled to correspond to a 15-minute interval of the model (i.e., divided by 30 days × 24 hours × 4 intervals); $\bar{P}_{i,m,y}$ is the average price for a kWh of consumption; *Superpeak* $4pm - 7pm$ and *Standard TOU* are indicators for whether household i was enrolled in either of those plans at the start of the sample period (with the Superpeak 3pm-6pm plan being the base plan); *Income* is household i's income level, which is divided into bins; *Household size* is the number of members in the household; ρ_i, ψ_m, ϕ_y, $\delta_{i,m}$, and $\pi_{m,y}$ are household, calendar month, year, household-by-calendar month, and calendar month-by-year fixed-effects, respectfully.[53] Table E.1 presents the results.

[53]To make sure that the elasticities generated from Equation (E.14) are comparable to the model estimated

The coefficient estimates presented in Table E.1 are converted into the relevant model parameters, which are included in Tables 3 and E.3 under "Loss function (β_ω), TOU non-responsive".

E.5 Additional moment construction details

To estimate the plan choice parameters we follow Hortaçsu et al. (2017) in minimizing a vector of moments of the following form

$$\eta_{jt}^{(k)} = \frac{N_{jt}^{(k)} - \hat{N}_{jt}^{(k)}}{N_t^{(k)}},$$

where $N_{jt}^{(k)}$ is the number of households that switch from plan k to plan j in month t, $\hat{N}_{jt}^{(k)}$ is the predicted number of switchers, and $N_t^{(k)}$ is the total number of households enrolled in plan k in month t. This denominator aims to down-weight moments with larger variance.

We also include moments designed to help match aggregate switching trends and to identify heterogeneity in attention and inertia. The first additional set of moments we try to match are the total number of plan switches each month.[54] Next, to identify heterogeneity in the first stage of the model, we also match overall switching patterns across plans by observable characteristics. Specifically, we construct a vector of moments of the following form:

$$\eta_{Cj}^{(k)} = \frac{\bar{C}_j^{(k)} - \hat{\bar{C}}_j^{(k)}}{\bar{N}^{(k)}},$$

where C denotes a specific observational household characteristic (e.g., income, household size, or initial plan); $\bar{C}_j^{(k)}$ is the mean of C for households that switched from plan k to plan j over the entire sample period; $\hat{\bar{C}}_j^{(k)}$ is the model predicted mean of C for k to j switchers; and $\bar{N}^{(k)}$ is the average number of households enrolled in plan k over the sample period.

We limit the window of observation to January 2020 and April 2023 for all choice plan moments.

counterparts for TOU-responsive consumers, we treat income bins as continuous and divide them by 100000. Similarly, we treat the household size variable as continuous.

[54] Though the inclusion of these moments is somewhat redundant, we view it as important to effectively replicate these aggregate switches for our counterfactual simulations. For some reason, relying only on the η vector of moments pushes the model to severely under-predict monthly switches.

E.6 Model outputs

Table E.2: First-stage Model Results

Panel A: Attention Parameters			**Panel B: Choice Parameters**		
Parameter	Coef.	SE	**Parameter**	Coef.	SE
Constant	-10.577	(1.032)	Incumbent plan dummy	4.343	(0.229)
Income	1.394	(3.869)	*Incumbent plan* ×:		
Household size	1.173	(0.175)	Income	0.370	(0.764)
Initial plan dummies:			Household size	0.659	(0.037)
4-7 pm plan	1.793	(0.521)	Init. 4-7 pm plan	4.547	(0.241)
2-8 pm plan	0.108	(0.555)	Init. 2-8 pm plan	2.380	(0.902)
Month dummies:			Logit error scale (σ_ϵ)	0.007	(0.002)
February	-0.245	(0.645)			
March	-0.059	(0.214)			
April	0.449	(0.490)			
May	1.545	(0.478)			
June	2.232	(0.480)			
July	2.428	(0.482)			
August	2.354	(0.474)			
September	2.284	(0.471)			
October	1.363	(0.410)			
November	0.965	(0.517)			
December	1.126	(0.527)			

Note: This table reports results from the plan choice model presented in Section 5.2 estimated via GMM. Parameter estimates are reported with bootstrapped standard errors in parentheses.

Table E.3: Daily Consumption Model Results

Parameter	Coef.	SE	**Parameter**	Coef.	SE
Load shifting disutility (β_d)			**Loss function (β_ω), TOU responsive**		
Duration (β_{d1})	0.057	(0.034)	*Demographics:*		
9am-5pm (β_{d2})	5.884	(1.908)	Household size	0.202	(0.107)
Demographics:			Income	0.786	(0.523)
Household size	0.0003	(0.00016)	*Initial plan choice:*		
Income	-0.003	(0.002)	3-6 pm plan	4.522	(1.654)
Initial plan choice:			4-7 pm plan	2.712	(1.043)
3-6 pm plan	0.008	(0.004)	2-8 pm plan	6.470	(2.471)
4-7 pm plan	0.014	(0.005)	**Loss function (β_ω), TOU non-responsive**		
2-8 pm plan	0.005	(0.002)	*Demographics:*		
Month intercepts:			Household size†	2.694	(0.666)
May	0.002	(0.001)	Income†	24.130	(7.762)
June	-0.004	(0.002)	*Initial plan choice:*		
July	-0.003	(0.001)	3-6 pm plan†	1.967	(0.934)
September	-0.0001	(0.00002)	4-7 pm plan†	-2.497	(0.732)
October	0.005	(0.002)	2-8 pm plan†	-0.389	(0.507)
TOU-responsive bliss points:			**TOU-unresponsive bliss points:**		
Draw std. dev. (σ_ν)	1.257	(0.033)	*Demographics:*		
Demographics:			Household size	-0.018	(0.006)
Household size	0.254	(0.083)	Income	12.667	(0.777)
Income	8.739	(3.803)	*Initial plan choice intercepts:*		
Initial plan choice intercepts:			3-6 pm plan	-0.067	(0.015)
3-6 pm plan	-0.987	(0.427)	4-7 pm plan	-0.297	(0.138)
4-7 pm plan	-1.030	(0.468)	2-8 pm plan	-0.278	(0.110)
2-8 pm plan	-0.817	(0.549)	*Month intercepts:*		
Month intercepts:			January	0.056	(0.076)
January	-0.227	(0.104)	February	0.018	(0.023)
February	-0.142	(0.060)	March	0.019	(0.005)
March	-0.338	(0.088)	April	0.230	(0.059)
April	-0.206	(0.078)	May	-0.044	(0.043)
May	0.175	(0.074)	June	-0.607	(0.214)
June	0.864	(0.214)	July	-0.632	(0.220)
July	1.090	(0.222)	September	-0.106	(0.032)
September	-0.247	(0.094)	October	0.283	(0.057)
October	-0.083	(0.036)	November	0.188	(0.009)
November	-0.275	(0.076)	December	-0.178	(0.067)
December	-0.063	(0.028)			
Mean-preserving spread parameters:					
3-6 pm plan, flatten	0.922	(0.252)			
4-7 pm plan, flatten	0.990	(0.370)			
2-8 pm plan, flatten	0.944	(0.335)			

Note: The table reports results from the structural model of daily electricity consumption presented in Section 5.1 estimated via SMM. Parameter estimates are reported with bootstrapped standard errors in parentheses. † indicates externally calibrated parameters. See Appendix E.4 for more details.

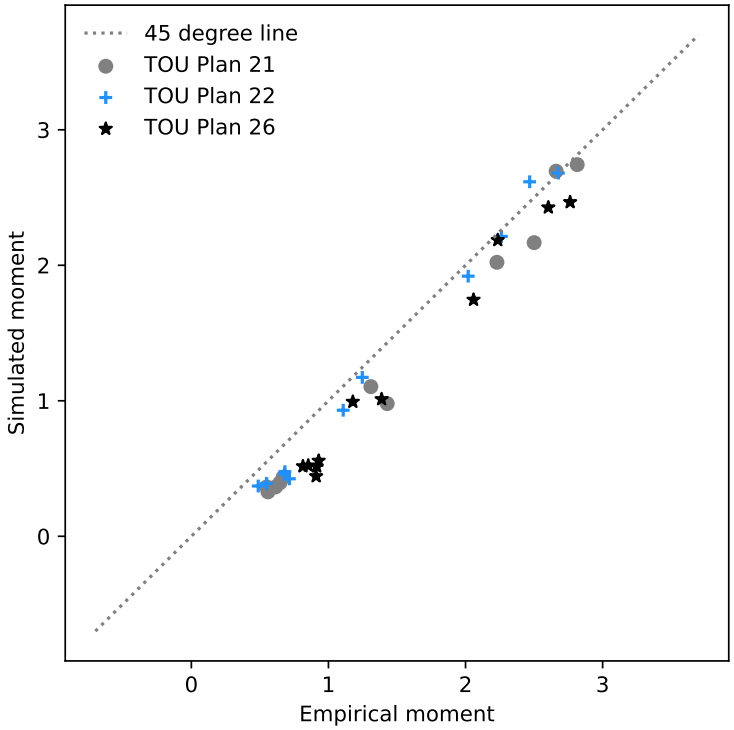
Table E.4: Moments: second stage, time-of-use responsive consumers

	Model	Data		Model	Data
Δ Consump. peak→off			**Av. consump./hr**		
Mean Δ:			*Mean*:		
2-8pm plan, August	1.34	1.35	2-8pm plan, August	3.16	3.14
4-7pm plan, August	1.27	1.33	4-7pm plan, August	2.90	2.91
3-6pm plan, August	2.16	2.22	3-6pm plan, January	1.36	1.31
Corr(Δ, income):			3-6pm plan, February	1.32	1.28
3-6pm plan, August	0.30	0.30	3-6pm plan, March	1.18	1.09
Corr(Δ, hhold size):			3-6pm plan, April	1.44	1.42
3-6pm plan, August	0.20	0.20	3-6pm plan, May	1.98	2.06
Load-shift 3hr post-peak			3-6pm plan, June	2.78	3.49
Mean:			3-6pm plan, July	3.12	3.90
2-8pm plan, August	3.35	3.35	3-6pm plan, August	2.97	2.77
4-7pm plan, August	2.25	2.25	3-6pm plan, September	2.48	2.23
3-6pm plan, May	1.71	1.71	3-6pm plan, October	1.52	1.46
3-6pm plan, June	4.96	3.85	3-6pm plan, November	1.20	1.16
3-6pm plan, July	5.69	4.75	3-6pm plan, December	1.37	1.43
3-6pm plan, August	3.75	3.74	*Std. dev:*		
3-6pm plan, September	2.40	2.46	3-6pm plan, August	1.40	1.40
3-6pm plan, October	0.91	0.91	*Corr(cons., income):*		
Corr(shift, income):			3-6pm plan, August	0.47	0.41
3-6pm plan, August	0.15	0.15	*Corr(cons., hhold size):*		
Corr(shift, hhold size):			3-6pm plan, August	0.31	0.35
3-6pm plan, August	0.03	0.02	**Max-min consump.**		
Load-shift 3hr pre-peak			*Mean:*		
Mean:			2-8pm plan, August	2.53	2.54
3-6pm plan, August	2.09	2.38	4-7pm plan, August	2.60	2.61
Load-shift to 11am			3-6pm plan, August	2.64	2.58
Mean:					
3-6pm plan, August	0.00	0.00			

Table E.5: Moments: second stage, time-of-use non-responsive consumers

	Model	Data		Model	Data
Av. consump./hr			**Av. consump./hr**		
Mean:			*Mean:*		
2-8pm plan, August	3.16	3.14	3-6pm plan, October	1.52	1.46
4-7pm plan, August	2.90	2.91	3-6pm plan, November	1.20	1.16
3-6pm plan, January	1.36	1.31	3-6pm plan, December	1.37	1.43
3-6pm plan, February	1.32	1.28	*Std. dev:*		
3-6pm plan, March	1.18	1.09	3-6pm plan, August	1.40	1.40
3-6pm plan, April	1.44	1.42	**Max-min consump.**		
3-6pm plan, May	1.98	2.06	*Mean:*		
3-6pm plan, June	2.78	3.49	2-8pm plan, August	2.53	2.54
3-6pm plan, July	3.12	3.90	4-7pm plan, August	2.60	2.61
3-6pm plan, August	2.97	2.77	3-6pm plan, August	2.64	2.58
3-6pm plan, September	2.48	2.23			

Figure E.1: Model fit on untargeted moments



Note: This figure shows the model's fit to the max-min consumption moment for every month, across the three TOU plans, except August. These moments are not targeted in estimation. We use them to test whether the model is able to capture how daily consumption patterns evolve during the year. The results show a close fit between the model-based predictions and the untargeted empirical moments.

Table E.6: Menu design counterfactuals: equity implications

	Consumer surplus change vs baseline menu (dollars/year/consumer)	
	Low-income	High-income
First best:		
No choice	-152.2	-306.3
Choice + inattention, inertia	-138.0	-246.3
Choice + inertia	-34.5	-94.5
Max. utility	-31.4	-91.8
Menu designs:		
Base plans + optimize:		
1 peak/off TOU plan	27.8	-33.3
2 peak/off TOU plans	35.0	-26.2

Note: This table splits up the aggregate consumer surplus into high-income (above median income) versus low-income (below median income). Note that the split is not exact (i.e. there is not exactly 50% of consumers in each type) because income is discretized. Therefore, the mean of the low-income and high-income numbers does not quite equal the aggregate change numbers, but it is still informative. Overall, the table shows that the optimal menu design is also slightly progressive with low-income consumers doing better, alleviating equity concerns.