Socially Efficient Menu Design for Residential Electricity Plans

Link to most recent version here.

Maria Garcia-Osipenko, Nicolai V. Kuminoff,

Spencer Perry, Nicholas Vreugdenhil*†

January 16, 2025

Abstract

This paper considers how to design a socially efficient menu of electricity price plans subject to realistic design constraints. We develop a model of plan choice, consumption, and intertemporal consumption-shifting, and estimate it using a dynamic panel of eight thousand consumers who choose between five static and dynamic price plans and then consume in 15-minute intervals for five years. We find evidence of substantial consumption-shifting among a subset of consumers, and selection into plans based on multidimensional heterogeneity. Counterfactual simulations reveal that the menu can be redesigned to simultaneously increase consumer welfare, reduce private generation costs, and improve environmental outcomes.

*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 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

Residential electricity pricing is inefficient. The vast majority of US households pay a constant marginal price throughout the day as supply and demand oscillate. This inefficiency is compounded by externalities: electricity generation accounts for 25% of the nation's greenhouse gas emissions (EPA, 2024). Ideally, a planner would set the instantaneous price of electricity equal to its marginal social cost of generation. While real-time pricing is rare, utilities increasingly allow consumers to choose from a menu of time-constant and time-varying ('time-of-use') price plans.[1]

In theory, offering a menu of time-constant and time-varying price plans can improve social welfare relative to charging a constant marginal price. Time-of-use incentives steer consumers in those plans to reduce electricity use in high-cost intervals and shift use to low-cost intervals. When important features of consumer heterogeneity are private information — such as their price sensitivity and preferences for intertemporal substitution — offering a menu of plans has potential to leverage self-selection to match consumers to their most efficient plans.

However, the socially efficient menu may set time-of-use prices above or below marginal social cost. Intuitively, if individual utility is not aligned with social welfare then price incentives may be needed to induce consumers to select into time-of-use pricing. Additional incentives may be needed to steer consumers who do not fully understand time-of-use pricing, or who are subject to inattention or inertia. This presents a tradeoff. Decoupling time-of-use prices from marginal social cost may increase selection into time-of-use plans, while simultaneously generating misaligned incentives about when to consume. This tradeoff may produce a socially efficient outcome, or it may degrade social welfare, depending on how the menu is designed.[2]

This paper considers how to design a socially efficient menu of electricity plans.

[1] A typical "time-of-use" schedule charges a high price during a predetermined interval when demand is usually high and a low price at all other times.

[2] Indeed, even if heterogeneous consumers (i) select voluntarily into plans with prices that are correlated with social marginal cost (ii) respond to the steering and load-reduction incentives, this is not a sufficient condition for a menu of time-of-use plans to improve upon a fixed-rate plan.

We model plan choice, consumption, and intertemporal substitution. The model accommodates consumer selection into plans on multidimensional heterogeneity and information frictions, as well as heterogeneity in within-plan decisions for utilization and intertemporal shifting across 15-minute intervals each day, throughout the year. After estimating the model, we use our results to characterize the optimal menu.

We estimate the model using administrative data from a large utility that promotes consumer choice and time-of-use pricing. We use a random sample of eight thousand households that do not own solar panels or electric vehicles.[3] We observe their demographics, price plan enrollment, and consumption in 15-minute intervals for five years.

Three features of our setting are particularly important for advancing knowledge of menu design and electricity demand. First, the menu is rich. Consumers choose from five price schedules, three of which employ time-of-use (TOU) pricing. Each TOU plan divides weekdays into high- and low-price intervals. They differ in when the high-price interval starts, how long it lasts, how much price increases, and how these features change seasonally. This variation can help to identify consumer preferences for shifting consumption through much of the day, and across the year, differentiating our study from prior work that examined binary choices between a TOU plan and a non-TOU plan (Fowlie et al., 2021; Ito et al., 2023).

Second, the market environment is information-rich. Consumers can monitor their hourly consumption, hourly expenditures, and the savings (or losses) that result from their price plan enrollment decisions. This should mitigate information frictions and strengthen the link from consumer preferences to market choices (Jessoe and Rapson, 2014; Prest, 2019).

Finally, we study the hottest metropolitan area in North America: Phoenix, Arizona. Its residents have adapted to extreme heat through energy-intensive air

[3]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.

conditioning. Therefore, our findings may inform how other regions can use menu design to manage the growing demand for cooling as global temperatures rise.

Our analysis starts by developing model-free evidence on plan selection and consumption. We measure if and how TOU consumers shift electricity load from high- to low-price intervals using a nonparametric double-difference estimator that leverages longitudinal and cross-sectional variation in consumption to characterize the full distribution of load-shifting behavior. We find substantial load-shifting among consumers who respond to TOU pricing. Most of the shift occurs after prices decline, with a smaller shift before prices rise. The size of the shift declines with the high-price interval's duration. However, we also find that load-shifting is concentrated among 26% of TOU consumers. The remaining 74% do not respond to within-day price changes of up to 400%.

Four additional facts guide how we model plan choice. First, holding the timing of consumption fixed, consumers leave relatively little money on the table when choosing plans. The average consumer could have reduced annual expenditures by just 2% had they chosen their ex-post cost-minimizing plan. Second, consumers are more likely to switch plans after receiving large monthly bills. Third, switchers save money by switching. Finally, consumers select into (and out of) TOU pricing 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).

Our structural model reproduces these empirical facts, and the estimates yield several additional insights. For example, we find that consumers sort themselves over TOU plans based partly on load-shifting preferences. Those with greater shift disutility tend to select plans with shorter high-price intervals. We also find the disutility of load-shifting *decreases* with the prevalence of extreme heat. This may be due to seasonal variation in load-shifting technology. For example, when the temperature is high, consumers can shift when they use air conditioning during the day, whereas the lack of air conditioning in mild temperatures requires other

adjustments that may be more costly (e.g. changing when they cook dinner). After consumers choose a plan, we find they become relatively inattentive to the menu, and that when they are triggered to consider switching plans (e.g. by a large bill) there is still inertia toward the current plan.

Our counterfactual simulations reveal that electricity price schedules can be redesigned to simultaneously increase consumer welfare, reduce private generation costs, and improve environmental outcomes. This conclusion is based on using our structural model estimates to solve for a socially optimal menu, subject to realistic design constraints. Specifically, we retain the existing non-TOU plans on the menu and restrict counterfactual TOU plans to have one high-price interval, with prices discretized in increments of 1 cent per kilowatt/hour. Offering one or two TOU plans yields millions of potential menus. For each potential menu, we compute each consumer's utility-maximizing plan choice and utilization for every 15-minute interval throughout the year. We then calculate the associated changes in generation costs and environmental damages by adapting methods in Borenstein and Bushnell (2022). We find that the optimal menu increases social welfare by approximately \$75 per person per year. Further, this constrained-optimal menu achieves over 93% of the social welfare benefits of an unconstrained-optimal menu in which all TOU consumers are assigned to a single TOU plan that provides the closest discretized approximation to real-time-pricing at marginal social cost.

Our study contributes to three distinct literatures. First, it adds to literature on the design of markets where consumers incrementally choose price schedules and consumption quantities (Einav et al., 2021). Examples include electricity, health insurance, cell phone and internet service, and online education (Einav et al., 2013; Cabral, 2017; Lin and Sacks, 2019; Fabra et al., 2021; Fowlie et al., 2021; Marone and Sabety, 2022; Ito et al., 2023). Our study differs from prior work through our focus on estimating preferences for high-frequency intertemporal substitution. One of our utility's stated objectives in offering time-of-use plans is to steer consumers to shift load from high to low-cost time periods, and so understanding this channel is central

to menu design. In principle, if consumers with lower load-shifting costs select TOU plans that pay them to shift consumption to intervals with lower marginal generation costs, then this selection is *advantageous* because it reduces both private costs and environmental damages. This contrasts with evidence of *adverse* selection on intertemporal substitution in certain health care plans (Cabral, 2017; Lin and Sacks, 2019). Using our framework we can evaluate the extent to which harnessing this feature delivers on its theoretical promise.

Second, our paper relates to literature on the equity and efficiency of dynamic pricing in markets with externalities, especially electricity and transportation. Prior studies examined how price schedule design affects market outcomes when consumers are assigned to schedules (Wolak, 2011; Jessoe and Rapson, 2014; Prest, 2019; Yang et al., 2020; Schittekatte et al., 2024; Harding et al., 2023; Burkhardt et al., 2023; Bailey et al., 2024; Blonz, 2022). We extend this line of research to a more generic setting where consumers freely sort over multiple price schedules.

Finally, this paper relates to an extensive literature on electricity demand. Our structural model is a substantial departure from the seminal discrete-continuous models of consumer behavior in electricity markets (Hanemann, 1984; Hausman, 1985; Reiss and White, 2005). Importantly, we are the first to formally model intertemporal substitution, which is the key margin for targeting pollution externalities. Equally important is the fact that our model incorporates several aspects of consumer behavior that have been shown to be quantitatively important in randomized controlled trials and quasi-experimental studies. These include inattention (Sallee, 2014), inertia in plan choice (Fowlie et al., 2021), bill shock (Hortaçsu et al., 2017), 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). Incorporating all of these features yields a rich framework to investigate menu design.

In the remainder of the paper, Sections 2 and 3 summarize the empirical context and data. Section 4 presents descriptive evidence on plan choice and load-shifting.

Section 5 develops the structural model. Section 6 explains how its parameters are identified and estimated using simulated GMM. Sections 7 and 8 summarize results and counterfactual simulations before offering brief conclusions.

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 1.1 million customers in the Phoenix, Arizona metropolitan area.[4] SRP's generation profile is similar to the US as a whole. Just over half the electricity it generates 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 solar and other 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.

2.2 The Menu of Price Plans

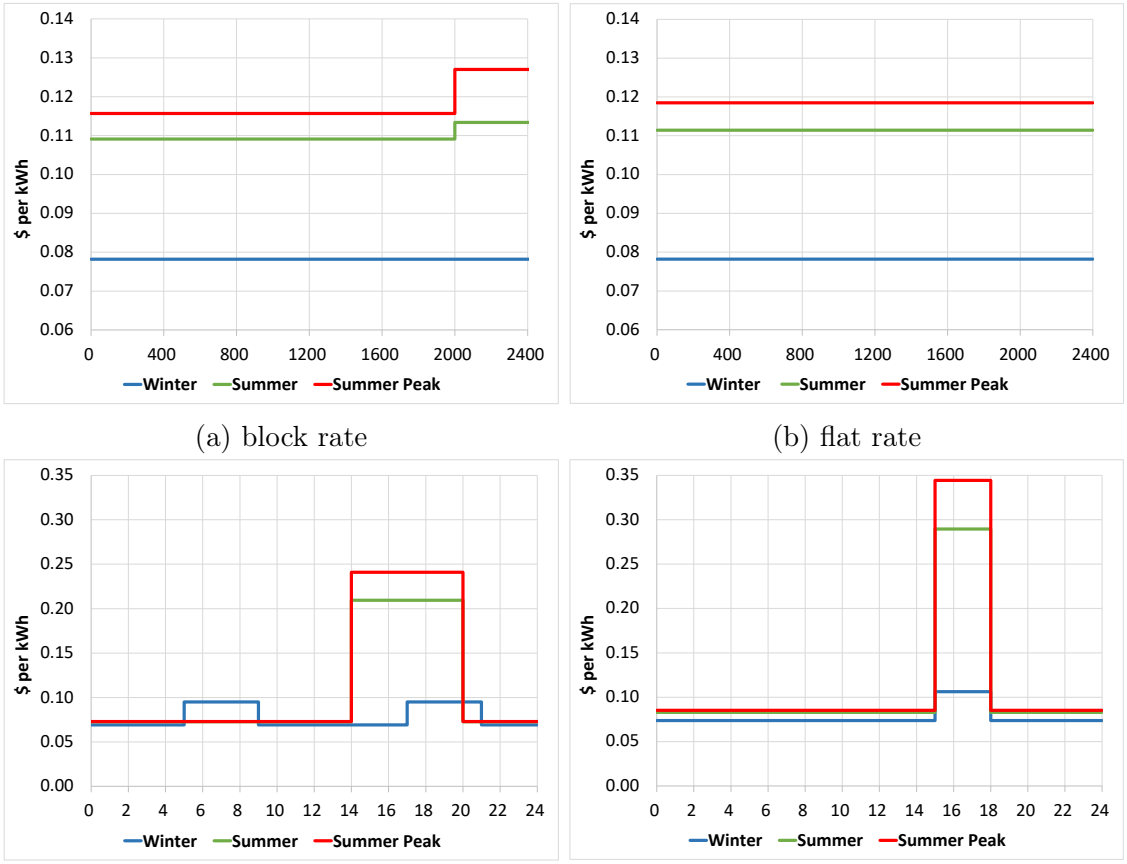
The menu that a customer faces depends on whether they have solar panels, an electric vehicle, both, or neither. The menus available to each group present distinct incentives for shifting consumption that, in turn, create distinct challenges for modeling behavior. In this study, we focus on households 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 households own each technology.[5]

Figure 1 shows the four types of plans on the menu. Within each panel, the colored lines show how prices vary from SRP's winter season (November through

[4]Phoenix is divided into two electric service territories, served by the SRP and APS utilities, so that a household's utility is determined by its residential location.

[5]As of 2022, fewer than 5% of houses in the US had rooftop solar (EIA, 2022) and fewer than 5% of vehicle registrations were electric or plug-in hybrid (DOE, 2022).

April) to the summer season (May, June, September, and October) to the peak summer season (July and August).



(a) block rate

(b) flat rate

(c) standard TOU

(d) superpeak TOU

Figure 1: Plan-specific prices by hour of day and season

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

Figure 1a 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 1b 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.[6] It is designed for budget-minded customers.

Figure 1c shows the standard TOU plan. In winter, the price rises 38% during peak hours from 5am-9am and 5pm-9pm (shown on the horizontal axis). In summer, the price rises 289% from 2pm-8pm, and this differential increases to 330% in peak summer months. Peak pricing only applies on weekdays.[7]

Finally, Figure 1d shows a superpeak TOU plan. Its peak hours are 3pm-6pm year-round, with price increases of 44%, 349%, and 404% on weekdays in winter, summer, and summer peak months. SRP also offers an alternate version of this plan with peak hours from 4pm-7pm (not shown in Figure 1). We henceforth refer to the five plans on the menu as: (1) block rate, (2) flat rate, (3) standard TOU, (4) superpeak 3pm-6pm, and (5) superpeak 4pm-7pm.

The menu spans the plan types examined in prior research on electricity demand, and it presents a unique opportunity to learn about consumer preferences for intertemporal substitution. Prior non-experimental studies focused on households that were assigned to block plans (e.g. Reiss and White, 2005; Ito, 2014). Prior studies that used controlled experiments to study plan selection forced subjects to choose between one TOU plan and one non-TOU plan (Fowlie et al., 2021; Ito et al., 2023). This is the first opportunity to study a market where consumers are free to choose between block rate, flat rate, and multiple TOU plans.[8] Moreover, the TOU plans differ in when peak pricing starts, how long it lasts, how much price increases, and how these features change during the year.[9] This variation offers the potential to identify preferences for shifting consumption through much of the day.

[6]Funds can be deposited in person or via an app that tracks consumption and remaining credit.

[7]The following holidays are also excluded: New Year's Day, Memorial Day, Independence Day, Labor Day, Thanksgiving Day, and Christmas Day.

[8]Lang et al. (2023) find that survey respondents say they would be more likely to enroll in a TOU plan if there were more TOU plans, larger price differentials between high-price and low-price periods, and shorter high-price periods.

[9]The within-day variation in TOU prices (38% to 404%) is also large compared to Spain's real-time pricing market, where pries varied by 23% on average (Fabra et al., 2021).

2.3 Enrollment, Switching and Information

Importantly, there is no default plan. 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.[10]

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.[11] This information has potential to mitigate information frictions and strengthen the link from consumer preferences to enrollment and consumption decisions (Jessoe and Rapson, 2014; Prest, 2019).

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.[12] Each household is identified by a unique code linking their billing account to their residential address.[13] We observe the cities where they live, the dates they started service, and the plans they chose. For households who subsequently switched plans, we observe the dates they switched and the new plans they chose. Finally, we observe some demographics that SRP obtained from an

[10] Appendix A provides screen shots of the choice process and additional background.

[11] Hence, the projections measured ex post savings (or losses) from customers' prior enrollment decisions. Appendix Figures A.2 and A.3 show how the information is presented.

[12] 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.

[13] We are unable to track customers who move across addresses. If a customer changes addresses, they need to open a new account at the new address. A new resident at the old address would also need to create a new account.

external contractor, including measures of household size and income.[14]

3.2 Smart Meter Data and Predicted Bills

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 in their actual plan and in each counterfactual plan.[15] We accomplish this by mapping smart meter data onto the price schedules in Figure 1 and additional information on fixed costs and taxes.[16]

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.[17]

3.3 Marginal Social Costs of Electricity Generation

Evaluating welfare implications of counterfactual menus requires knowing the hourly marginal social costs of generating electricity used by SRP households. We estimate these costs by applying methods from Borenstein and Bushnell (2022) to data for our study period. 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

[14]The demographics also include several variables that we do not use in our main analysis: homeownership, dwelling type, account holder age, credit score, whether a household had an online account, and whether a household ever participated in programs that subsidize energy efficiency investments such as shade screens, shade trees, and multi-stage pool pumps.

[15]Calculating counterfactual bills requires predicting how consumers would adjust to counterfactual price schedules. Section 5 presents our structural model of consumption.

[16]Appendix C provides supporting details for our monthly bill calculator.

[17]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 fees for ad hoc services that we cannot observe. Examples include fees for service establishment, returned payments, field visits, high-bill audits, theft investigation, replacement of customer-damaged equipment, and late payments.

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.

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.

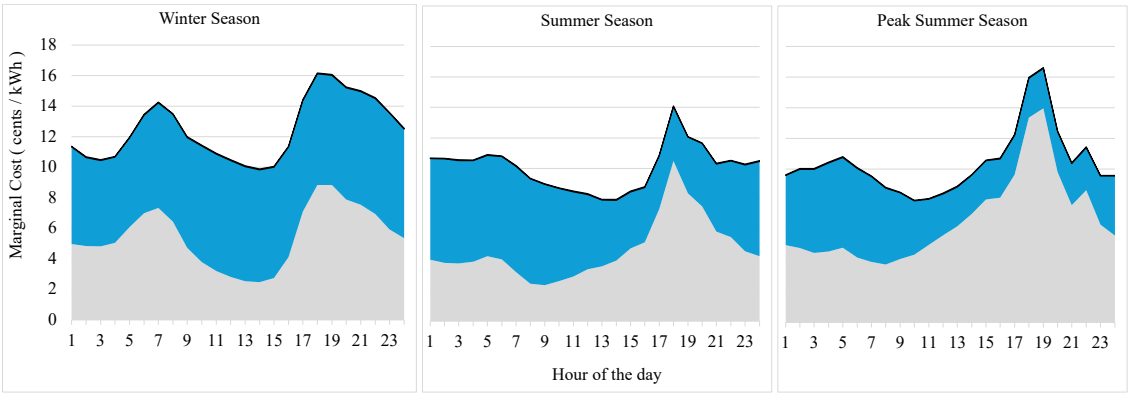


Figure 2: Marginal Costs of Electricity Generation by Season and Hour

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.

Figure 2 shows how marginal costs vary over the average day in each pricing season. Private costs (grey) mirror the between-season and within-day variation in retail prices shown in Figure 1. 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.

4 Descriptive Evidence

4.1 Estimation Sample

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.[18] Third, we drop 1.9% of households who are missing demographics and 4.3% who spend winters in Phoenix but live elsewhere most of the year.[19] Finally, we exclude March 2020 through April 2021 during which Arizona's social distancing policies temporarily increased residential electricity demand.[20]

Our final estimation 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 household-specific consumption intervals.

4.2 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.[21] The last five columns show that the block rate plan has the largest market share (56%), followed by the Superpeak 3pm-6pm plan (23%) and the standard TOU plan (17%).

Households are partially stratified across plans by size and income. For example,

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

[19] 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.

[20] Mean hourly consumption increased temporarily during the Covid-19 pandemic, especially during normal work hours and school hours (Appendix Figure D.1).

[21] Appendix Figure D.2 shows monthly variation in consumption and expenditures.

	All plans	block rate	fixed rate	standard time-of-use	superpeak 3pm-6pm	superpeak 4pm-7pm
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 $)	162	152	155	188	165	151
mean income ($)	75,250	69,446	38,063	92,792	81,346	71,037
household size (mean)	2.1	2.0	1.8	2.4	2.0	1.8

Table 1: Summary Statistics Overall and by Plan

mean income is lowest among households in the fixed-rate plan, consistent with that plan being designed for budget-minded customers. Likewise, income and household size are highest in the standard TOU plan.[22] 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.

4.3 How Do Households Respond to TOU Pricing?

We estimate the extent to which households respond to TOU pricing by reducing consumption during high-price hours and shifting consumption to low-price hours. To fix ideas, let q_{itwj} denote the electricity consumed by household i during 15-minute interval t on day $w \in \{weekday(d), weekend(e)\}$ of a particular pricing season, given that the household is on price plan $j \in \{TOU, Block(B)\}$. The difference between a TOU household's average 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 household 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 consump-

[22] Appendix Table D.1 summarizes other household characteristics.

tion 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 household can be written as $E[q_d|i,t,B] - E[q_e|i,t,B] = \lambda_{i,t,B}$.[23] Differencing average consumption differentials among TOU and block-rate households yields $\delta_{t,TOU} + \lambda_{t,TOU} - \lambda_{t,B}$. This statistic fails to identify $\delta_{t,TOU}$, however, if households sort themselves into plans such that $\lambda_{t,TOU} \neq \lambda_{t,B}$. For example, households that 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 household to a composite block-rate household 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 household'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 household to all block-rate households 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 households.*

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, household i's response to TOU pricing can be recovered from a nonparametric double-difference estimator:

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

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

For TOU household i in quantile v of the λ distribution at t', the estimator subtracts the average weekday-to-weekend consumption differential among all block rate households in the $v_{t'}$ quantile from the TOU household's consumption differential. Repeating this calculation at all $t \notin t'$ recovers household 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 superpeak plans and t'= [12:30pm, 1:30pm] for the standard TOU plan. We think these matching intervals are likely to approximately satisfy Assumptions A1 and A2. First, they will satisfy A1 as long as TOU households do not shift consumption more than three hours from peak periods.[24] Second, the matching intervals capture the effect of 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 households. 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 Superpeak 3pm-6pm households. The tinted curves depict heterogeneity in responsiveness. Each curve shows the mean response for a decile of households, ranked by the change in mean consumption during on-peak hours (delineated by dashed vertical lines). Households 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 3pm-6pm on weekends (about 4.5kW). Importantly, the figure also shows that the top two

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

deciles shift some of their on-peak load to adjacent off-peak hours, particularly 6pm to 8pm.[25] Analogous results for the other TOU plans and pricing seasons are shown in Appendix Figure D.3.

Figure 3a also suggests that most households have little or no response to TOU pricing. Peak consumption changes by less than 0.5kW for the bottom seven deciles. Whether or not a particular household 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 households 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 household's mean consumption 15-to-30 minutes before prices rise during summer and peak summer seasons with its mean consumption 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.[26] Similarly, excluding 15-minute intervals on each side of the price increase reduces noise if some households adjust a few minutes early or late.[27] The results indicate that 26% of households respond to hourly price changes.[28] Figure 3b shows peak summer load curves for these consumers in the two most popular plans: Superpeak 3pm-6pm and standard TOU.[29]

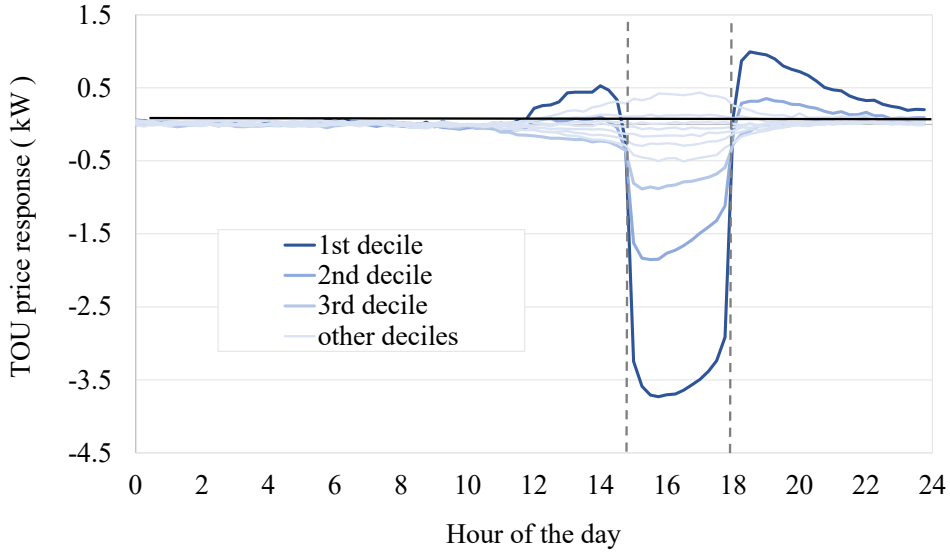
[25]There are numerous ways to shift load. A leading example in Phoenix summers is to turn the thermostat down below one's bliss point to pre-cool living space from 12pm-3pm, then raise the thermostat above the bliss point from 3pm-6pm, before turning it back to the bliss point at 6pm. SRP's website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

[26]This is shown in Figure D.3.

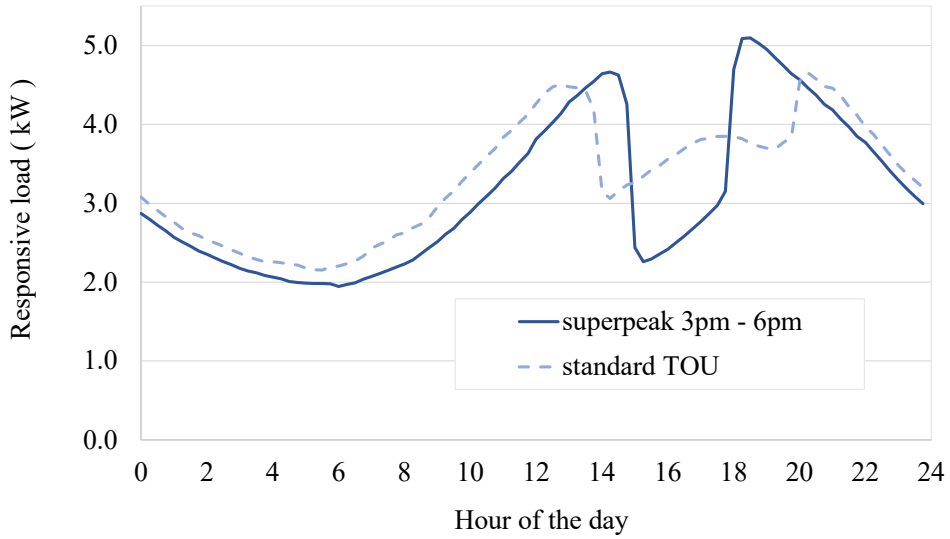
[27]The test statistic is $(q_{i,2} - q_{i,-2})/[sd(q_{i,2} - q_{i,-2})/\sqrt{n}]$, where $t = -2$ and $t = 2$ refer to 15-to-30 minute intervals before and after prices increase at $t = 0$. An observation is a weekday during summer and peak-summer months. This yields 384 observations for a household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

[28]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 household 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 households switch from non-responsive to responsive and 1.8% switch in the opposite direction.

[29]Appendix Figure D.4 shows load curves for responsive and unresponsive households on each



(a) response heterogeneity example



(b) mean load for responsive households

Figure 3: Estimated Response to TOU Pricing

Note: Panel (a) shows estimated household responses to TOU pricing in the superpeak 3pm-6pm plan. Each curve corresponds to a decile of households, 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 households that we find to be responsive to TOU pricing in the superpeak 3pm-6pm plan. The dashed line shows the analogous load curve for responsive households in the standard TOU plan.

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.

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 households are poorly matched (violating A2) or if TOU households shift load before 10am (implying that A1 is very 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 households were to habitually reduce consumption 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, consumption declines by 31% for responsive households on the Standard plan and by 51% for those on the Superpeak plan. The corresponding changes on weekends are only 2% and 6%.[30]

Summary. Overall, Figure 3 presents four facts for our model to explain. First, the average TOU household shaves load during on-peak hours. Second, the average household shifts some load to off-peak hours. Third, there is substantial heterogeneity in load-shaving and load-shifting. Finally, the majority of TOU households do not adjust hourly consumption to hourly price changes during peak-pricing days.[31]

[30] Appendix Figure D.5 provides additional visual evidence by showing weekday and weekend load curves for responsive and nonresponsive households on each TOU plan.

[31] While 74% of households have no statistically significant response to *hourly* price changes, the model in Section 5 still allows them to respond to average prices on monthly bills as in Ito (2014) and/or to between-plan differences in expected expenditures at the time of enrollment as in Hortaçsu et al. (2017).

4.4 Plan Selection

Data on enrollment and expenditures exhibit four patterns that guide how we model plan selection. First, households leave little money on the table when choosing plans. The average household could have saved just $35 per year (2% of annual expenditures) had it chosen the plan that would have minimized the cost of its actual load.[32] Of course, even if households had chosen one of the five plans randomly, 20% could have minimized expenditures without adjusting load. Yet the actual share is much higher: 36%. This suggests households are attentive to price when choosing plans and/or that they subsequently adjust load as in Figure 3.[33]

Second, households are more likely to switch after receiving large monthly bills. 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.[34] Figure 4b shows that, within a given month, households 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.[35]

Third, switchers save money by switching. This finding is based on 761 households that switch plans exactly once between 2019 and 2023. We regress an indicator for whether the household minimizes the annual cost of its load on an indicator for whether it switched previously, and an interaction between past switching and an indicator for whether it is TOU-responsive.[36] We add fixed effects for households and years to identify the coefficients from within-household expenditure changes. The "cost minimizer" column in Table 2 shows results. Switchers are nearly 15 per-

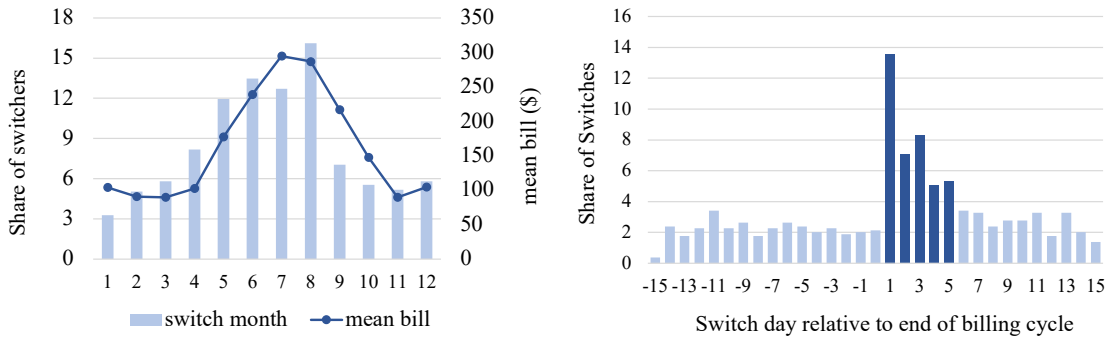
[32]For the average household, annual expenditures differ by $95 between their most and least expensive plans. This range varies substantially across households; its standard deviation is $88.

[33]This is consistent with evidence that providing information about prices and potential savings (see Section 2.3) can increase consumers' price sensitivity (Jessoe and Rapson, 2014; Prest, 2019).

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

[35]We infer the end of the billing cycle from the day-of-month of each household'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.

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



(a) Switchers by month (b) Switchers by day-of-month

Figure 4: Timing of Plan Changes

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.

centage points (or 50%) more likely to minimize costs after switching. The smaller interaction term suggests that switchers are attentive to between-plan differences in price, even if they do not respond to TOU price incentives.

	Cost minimizer	TOU switcher
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	0.30	0.60
observation	account-year	account
# observations	2,022	395
R^2	0.72	0.03

Table 2: Evidence on Plan Switching and Cost Savings

Note: The "Cost minimizer" column shows results from a panel data regression with fixed effects for household and year. The "TOU switcher" column shows results from a cross-section regression. Both columns use data on households that switched between different price plans. Standard errors are clustered by household.

Finally, TOU-responsive households are more likely to select TOU plans. This

finding is based on 395 households that switched between TOU and non-TOU plans. We observe whether each household is TOU-responsive, as well as the average monthly amount it could have saved by switching to TOU *without responding to TOU prices* when it was on a non-TOU plan. We regress an indicator for switching to TOU on both variables. The "TOU switcher" column in Table 2 reports results. A $10 increase in monthly savings from switching is associated with a 2% increase in the switching probability. Moreover, TOU-responsive households are 13% more likely to select TOU pricing. The first result is consistent with "selection on the level" of savings from switching and the second result is consistent with "selection on the slope" of hourly demand (Einav et al., 2013; Ito et al., 2023).

5 Model

We model how households choose price plans, how much electricity they choose to use, and when they choose to use it. The model is designed to capture the load shaving, load shifting, and plan selection behaviors described in Section 4. For heuristic purposes, the choice process can be depicted as a repeated two-stage problem. At the start of each month a household first determines its optimal daily consumption in each of a finite number of price plans and then chooses the particular plan that maximizes its utility. The structure of the daily consumption decision depends on whether the household responds to hourly price changes. We divide households into *TOU-responsive* and *TOU-unresponsive* groups based on their behavior, as explained in Section 4.3, and model the consumption decision for each group.

5.1 Daily consumption decision

5.1.1 TOU-responsive households

Consider a TOU-responsive household i on price plan j. For a representative weekday or weekend w in month m, the household chooses how much electricity to use in each 15-minute interval, q_t, and how much to shift from that interval to every other interval, $s_{tt'}$, where $t' \neq t$. We denote the household's consumption and load-shifting

profiles throughout the day by $\mathbf{q} = \{q_t\}$ and $\mathbf{s} = \{s_{tt'}\}$. These decisions are specific to i, j, w, m, but we suppress their subscripts to simplify notation. The consumer chooses $\mathbf{q}$ and $\mathbf{s}$ to maximize the following indirect utility function:

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

$$\text{subject to: } q_t \geq 0 \text{ and } s_{tt'} \geq 0, \text{ where}$$

$$U(q_t, \mathbf{s}) = \underbrace{u\Big(\nu_{imwt} - q_t - \sum_{t' \neq t} s_{tt'}\Big)}_{\text{Deviation from bliss point}} - \underbrace{\sum_{t' \neq t} d_{imtt'}(s_{tt'})}_{\substack{\text{Disutility from} \\ \text{shifting consumption}}} - \underbrace{p_{jmwt} \times \left(q_t + \sum_{t' \neq t} s_{t't}\right)}_{\text{Cost of load consumed at t}}$$

The first component of flow utility, $u(.)$, is a loss function that increases with the distance between the household's actual consumption and its bliss point, ν_{imwt}. The bliss point is the amount of electricity the household would use if its price were zero.[37] Consumption may occur during interval t or be shifted to t'. For example, a household on a TOU plan with peak hours from 4pm to 7pm might shift cooking dinner from 6:15-6:30pm to 7:00-7:15pm when price is lower.

Shifting consumption over time incurs disutility $d_{imtt'}(s_{tt'})$ but it can also reduce expenditures on a TOU plan if load is shifted to periods with low prices, p_{jwmt}.[38] The ways in which households respond to this tradeoff can help to explain the load-shifting behavior in Section 4.3. Heterogeneity in the willingness to pay for electricity and to avoid load-shifting is reflected in the parameters of $u(.)$ and $d_{imtt'}(s_{tt'})$ because we normalize the marginal utility of income to one (Einav et al., 2013).[39] Finally, note that there is no incentive to load-shift on block-rate and fixed-rate plans because

[37] We restrict bliss points to be finite to reflect the fact that households have finite electrical appliances and preferences over their optimal usage. For example, air conditioning may reduce utility if the air temperature is ideal.

[38] We do not include a separate discount factor for temporal substitution, due to the short duration of a single day.

[39] We present parametric forms of $u(.)$ and $d_{imtt'}(s_{tt'})$ in Section 6.1.

their prices are constant throughout the day.[40]

5.1.2 TOU-unresponsive households

Daily consumption decisions are simpler for households that are unresponsive to TOU pricing. They choose a consumption profile $\mathbf{q}$ to maximize indirect utility given their bliss points and their average monthly price, $\bar{p}_{ijm}$.

$$\hat{V}_{ijmw} = \max_{\mathbf{q}} \sum_{t} \Big(u(\nu_{imwt} - q_t) - \bar{p}_{ijm} q_t \Big) \tag{3}$$
$$\text{subject to: } q_t \geq 0, \quad \bar{p}_{ijm} = \sum_{w,t} q_t p_{jmt} / \sum_{w,t} q_t$$

Equation (3) differs from the TOU-responsive problem in Equation (2) in two ways. First, TOU-unresponsive consumers (potentially) use their average price to make decisions. This price is household-specific because it depends on how i's consumption profile interacts with plan j's price schedule. Second, there is no load shifting. This is not an assumption; it is the optimal response to load shifting being costly and households perceiving price to be constant throughout the month.[41]

5.2 Plan choice

A new customer considers its annual utility in each plan $j \in \mathcal{P}$. To calculate annual utility for plan j, we first aggregate flow utility in (2) or (3) over weekdays and weekends in month m, and then aggregate months to calculate: $V_{ij} = \sum_{m} (1/12) \times \Big((5/7) \times \hat{V}_{ijm,weekday} + (2/7) \times \hat{V}_{ijm,weekend} \Big)$. The household chooses the plan that

[40] Technically, a household on a block-rate plan could have an incentive to shift load if it approached a consumption threshold for a price increase on the last day of a monthly billing cycle. In this case, the household could potentially reduce expenditures by shifting load to the next month's bill. We ignore this possibility because it is statistically improbable and would require the household to exert considerable effort to track its hourly and monthly consumption.

[41] Put differently, for a TOU-unresponsive household, the optimization problem in Equation (2) simplifies to Equation (3).

maximizes annual utility plus ϵ_{ijm}, an i.i.d. logit error with scale parameter σ_ϵ.

$$\max_{j \in \mathcal{P}} \left\{ V_{ij} + \sigma_\epsilon \epsilon_{ij} \right\} \tag{4}$$

To enable the model to reproduce the fact that switching plans is rare but follows a strong seasonal pattern (see Section 4.4), we model plan choice for existing customers as being contingent on receiving an attention shock. Specifically, at the start of month m, household i receives a random draw from a Bernoulli distribution with parameter a_{im}. If the household draws a value of 0 then it is not attentive and it remains in its current plan. If the household draws a value of 1 then it is attentive and reconsiders its plan choice by solving:

$$\max_{j \in \mathcal{P}} \left\{ V_{ij} + \gamma_{ij} + \sigma_\epsilon \epsilon_{ijm} \right\} \tag{5}$$

Here, γ_{ij} is an "inertia" term that captures default bias toward the household's current plan.

6 Estimation and Identification

To estimate the model we must first parameterize the following components of Equations (2)-(5): the loss function $u(.)$, the load shift disutility function $d_{imtt'}(s_{tt'})$, the bliss points ν_{imwt}, the attention parameter a_{im}, and the inertia parameter γ_{ij}. This section presents our parametric forms, identification arguments, and econometric approach. Importantly, we narrow our focus to estimating preferences for the subset of households that we ever observe on TOU plans.

Our focus on ever-TOU households is strategic. As we explain below, we rely on within-plan variation in prices to identify key parameters. Non-TOU plans do not contain enough variation to make this strategy feasible without adding relatively strong parametric assumptions.[42] Thus, focusing on TOU households improves in-

[42]For example, we would need to predict whether non-TOU households would respond to TOU pricing. We would also need to separately identify their loss function and load-shift disutility

ternal validity. A caveat is that it limits us to considering counterfactual experiments that make the set of TOU plans weakly worse. This will cause existing TOU households to substitute across and away from TOU plans, and toward non-TOU plans. However, it will not disturb the welfare of households who revealed that they prefer non-TOU plans by virtue of never selecting a TOU plan.[43] Since many interesting and policy-relevant counterfactuals make the set of TOU plans worse, we think this approach is constructive.

6.1 Parametric Forms

Loss function. We define the loss function as $u(x) = -(1/(2\omega_i))x^2$, where ω_i is a parameter indexing curvature in the disutility of consuming away from a bliss point. The ω_i parameter also implicitly defines the consumer's marginal willingness to pay for electricity, conditional on distance from the bliss point. This interpretation follows from the fact that the marginal utility of income is normalized to one (Einav et al., 2013). Thus, ω_i provides a measure of price sensitivity.

We set $\omega_i = \beta_{\omega 0 j(i)} + \beta_{\omega 1} X_i$, where X_i is a vector of demographics incorporating household income and household size.[44] The intercept, $\beta_{\omega j(i)}$, is specific to $j(i)$, the initial plan a household chose when it first opened its account.[45] These plan-specific intercepts address the "initial conditions" problem of disentangling preferences from inertia (Wooldridge, 2005). That is, $\beta_{\omega j(i)}$ is designed to capture selection into plans based on unobserved preference heterogeneity. For example, TOU plans with larger within-day price differentials may attract households with more elastic demand.

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.

[43] Technically, one could argue that making TOU plans weakly worse necessarily reduces expected welfare for all consumers within our model because the iid Type I EV shocks imply a positive probability that every consumer chooses every plan. This argument assumes that the Type I EV shocks are best interpreted as latent preference heterogeneity. Further, even under this assumption, as the number of choice occasions grows, the probability that a consumer who never chooses a TOU plan would ever do so, conditional on the choice set and their characteristics, shrinks toward zero.

[44] Ideally we would include the square footage of the house, since previous work has shown this to be a good predictor of electricity consumption, but these data are not systematically available to us so we use household size (number of people) to proxy for house size.

[45] 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.

We also allow the $\beta_{\omega j(i)}$ and $\beta_{\omega 1}$ parameters to differ between TOU-responsive and TOU-unresponsive households. Thus, for each household type there is an intercept for each plan and two heterogeneity parameters. This stratification allows the shape of the loss function to vary flexibly with TOU-responsiveness.

Bliss points. We parameterize the bliss points as $\nu_{imwt} = \bar{\nu}_{mwtj(i)} + \bar{\nu}_i$. Here, $\bar{\nu}_i \sim N(\beta_{\nu 0} + \beta_{\nu 1} X_i, \sigma_\nu)$ is a household-specific random draw that shifts the bliss points up or down by a constant amount. 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. Specifically, we parameterize $\bar{\nu}_{mwtj(i)}$ using a $mwj(i)$-specific mean-preserving spread of the observed weekend consumption of the block-rate plan. For example, this would allow the bliss points for TOU consumers to be flatter across the day relative to block-rate consumers due to selection into plans. It also allows these bliss points to change across weekends, weekdays, and months.

Load shifting disutility. We set $d_{imtt'}(s_{tt'}) = \beta_{dimtt'} s_{tt'} + (1/2)\beta_{d4} s_{tt'}^2$, where:

$$\beta_{dimtt'} = \beta_{d0j(i)} + \beta_{d0m} + \beta_{d1}|t_{\text{midpoint}} - t'| + \beta_{d2} 1[t' \text{ in 9-5pm}] + \beta_{d3} X_i \qquad (6)$$

Similar to our specification for the loss function, β_{d3} allows the disutility to vary with household-specific covariates and $\beta_{d0j(i)}$ captures selection into plans based on unobserved heterogeneity. β_{0m} allows the disutility to shift each month. This allows for seasonal changes in load-shifting technology. For instance, choosing when to run an air conditioner is thought to be a key way that Arizona households load-shift in summer, but not in winter. β_{d1} allows the disutility of load shifting to vary with its duration.[46] This is suggested by the data where load-shifting is concentrated just before and after peak hours (e.g. Figure 3a). Without loss of generality, we depict the case where t is in a peak period and t' is in an off-peak period. Note that we

[46]For instance, shifting oven use from 6pm to 8pm may have a higher disutility than shifting it from 6pm to 7pm.

measure the distance from the midpoint of the peak period $t_{midpoint}$.[47] 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 $\gamma_{ij} = \beta_{\gamma 0 j(i)} + \beta_{\gamma 1} X_i$. This specification allows inertia to vary with household demographics and initial plan choices, analogous to the loss function and bliss-point parameters.

Attention. As noted earlier, we model households 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.[48]

6.2 Identification and Moment Construction

6.2.1 Consumption Parameters

TOU-responsive consumers. The consumption parameters for TOU-responsive households are identified at the individual i level. A key identification challenge is price endogeneity. This is explicitly represented in the primitives in our model because the bliss points may increase in peak periods when TOU price incentives are higher. Therefore a naive regression of consumption on price would result in consumers appearing to prefer higher prices.

We solve the price endogeneity problem by exploiting the sharp price changes within TOU plans. 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

[47] If we instead measured it via $|t - t'|$ this would result in a peak period utilization that is bowed inwards, which does not match the data.

[48] We experimented with allowing a_{im} to additionally vary with the percentage change in the household'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.

can be written as:

$$(q_l + \Sigma_{l \neq l'} s_{l'l}) - (q_h + \Sigma_{h' \neq h} s_{h'h}) = (\nu_l - \nu_h) + (u')^{-1} p_h - (u')^{-1} p_l + \Sigma_{l \neq l'} s_{l'l} \quad (7)$$

$$= \omega(p_h - p_l) + \Sigma_{l \neq l'} s_{l'l}, \quad (8)$$

The right-hand-side of Equation (7) follows from inverting the first-order optimality condition that $u'(\nu_t - q_t - \Sigma_{t' \neq t} s_{tt'}) = p_t$.[49] Equation (8) follows from the functional form for $u(.)$ and our assumption that $\nu_l = \nu_h$ 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 4.3. Therefore, the only unknown component is ω. In practice, the components of ω are identified off of one moment for each TOU plan for the median difference in utilization left and right of sharp jump in peak price, as well as moments for the covariance of this difference and each observable demographic characteristic.

Using this ω, the bliss points are identified at the individual level for every 15-minute interval. Rewriting 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. In estimation, ν_t contains several components as discussed in Section 6.1. The first is the term $\nu_{mwtj(i)}$, which is the 'average' bliss point for plan j at each time period. To estimate this we construct transformations of the block-rate average consumption at each mwt. Specifically, we estimate a parameter for a mean-preserving spread transformation of the block rate consumption across the time periods of the day, which allows consumers in the TOU plans to potentially have a flatter (or sharper) consumption schedule than consumers in the block-rate plan. We include a corresponding moment for the minimum consumption as a share of the maximum consumption in August.

We also include initial plan-specific and month-specific parameters that shift

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

bliss points up and down. Therefore, we include a corresponding average consumption moment for each TOU plan for the month of August, and 11 more moments for the average consumption on the Superpeak 3pm-6pm plan for each other month to estimate the month shifters. Finally, we include moments relating to the covariance of mean consumption and observable demographics to estimate $\beta_{\nu 1}$, and the standard deviation of consumption across consumers in August for the Superpeak 3pm-6pm plan to estimate σ_ν.

Finally, the load-shifting 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 $p_t - p_{t'} = d'_{tt'}(s_{tt'})$, which can be rewritten as:

$$p_t - p_{t'} = \beta_d + \beta_{d0m} + \beta_{d1}|t_{\text{midpoint}} - t'| + \beta_{d2}1[t' \text{in 9-5pm}] + \beta_{d3}X + \beta_{d4}s_{tt'}, \quad (9)$$

where we reintroduce the m subscript for clarity. Thus, the β_d parameters are identified by how the shape of the shifted load varies over time.[50] To this end, we include one moment for each plan for the mean load shifted 3 hours right of the end of the peak period.[51] We also include moments for the covariance of the mean load shifted and each observable demographic to identify heterogeneous effects. We include a moment for the Superpeak 3pm-6pm plan (the most common TOU plan) for the load shifted 3 hours left of the start of the peak period. The difference of this moment and the shift to the right of the peak pins down the curvature parameter β_{d4}. Finally, we include a moment for the mean shift in the 9am-5pm period to estimate β_{d2}.

TOU-unresponsive consumers. While one might conclude that TOU-unresponsive consumers have a very low (or zero) price sensitivity, we are more flexible. We accommodate the possibility that these consumers are sensitive to price, but instead

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

[51] For example, if the peak period ended at 6pm, the relevant hours here would be 6-9pm.

of using the marginal price they respond to the average price as in Ito (2014). The identification challenge is to recover the price sensitivity (as well as the other parameters) when there is a subset of consumers who *potentially* respond to average prices. We emphasize that we are *not* ruling out the possibility that these consumers are insensitive to both marginal and average prices. Rather, we view this subset of consumers as potential candidates for average-price responsiveness and nest the possibility of non-responsiveness.

The price sensitivity for a particular consumer can be identified using plan switches. The difference in observed consumption in period t for a consumer who switches from plan j to plan j' is: $q_{jt} - q_{j't} = \omega(\bar{p}_{j'} - \bar{p}_j)$. Note that the average prices before and after the switch are also directly observed, so ω is identified at the consumer level for switches. We do not see all consumers switching in our dataset, so for those that do not switch, ω 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.

6.2.2 Plan choice parameters

The consumption parameters identify the indirect utility of each plan for each TOU-responsive consumer.[52] Given this, we now show that the plan choice parameters are also identified at the consumer level. As above, we suppress the i, m subscript to simplify notation.

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

[52]For the block-rate plan, we can also construct a counterfactual utility given the parameters of consumer i. We use the empirical utilization to determine the relevant block for the consumer. We then compute the counterfactual load on the block-rate plan and the corresponding total utility as if that was the consumer's rate. We then adjust total utility if the consumer is in the upper block, to account for the lower price in the lower block.

plan choice) times the conditional probability of choosing to switch:

$$P(\text{switch } j \to j') = a\frac{\exp(V_{j'}/\sigma_\epsilon)}{\Sigma_k \exp((V_k + 1[j=k]\gamma)/\sigma_\epsilon)} \tag{10}$$

Comparing the probabilities of switching to j' and j'' identifies the σ_ϵ parameter: $\ln(P(\text{switch } j \to j')) - \ln(P(\text{switch } j \to j'')) = (1/\sigma_\epsilon)(V_{j'} - V_{j''})$. 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).[53] 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 consumer in different plans, and compare the switching probabilities in the following way:

$$\begin{aligned}&\frac{P(\text{switch } j \to j) - P(\text{switch } j' \to j')}{P(\text{switch } j' \to j)} \\ &= \exp\left(\frac{V_j - V_{j'} + \gamma}{\sigma_\epsilon}\right) \cdot \frac{\Sigma_k \exp((V_k + 1[j'=k]\gamma)/\sigma_\epsilon)}{\Sigma_k \exp((V_k + 1[j=k]\gamma)/\sigma_\epsilon)} + \exp\left(\frac{V_{j'} - V_j + \gamma}{\sigma_\epsilon}\right)\end{aligned} \tag{11}$$

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 $j' \neq j$.

To construct the relevant switching probabilities outlined above, we follow Hor-

[53] 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.

taçsu et al. (2017). We include additional moments designed to help better match aggregate switching trends and to identify heterogeneity in attention and inertia. Additional details on the construction of these moments are provided in Appendix E.2.

6.3 Estimation

We estimate the parameters of the model in two-steps, leveraging the fact that a household's daily consumption decision and plan choice decision are separable. First, we estimate the parameters for within-day electricity consumption via simulated method of moments. With that set of parameters in hand, we simulate each household'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)), treating the simulated indirect utilities as data. We then separately iterate on these two procedures to calculate bootstrap standard errors.

7 Results

7.1 Within-day consumption decision

We simulate the model at a 15-minute daily frequency, across plans, months, and weekday/weekend consumption patterns. Table 3 presents the point estimates and bootstrapped standard errors of the parameters of the consumption choice model. Although it is difficult to interpret each parameter estimate in isolation, there are several noteworthy patterns that emerge. In terms of load shifting disutility, we see that shifting is substantially more costly for households during business hours (i.e., 9am-5pm); that income and household size do not meaningfully affect the cost of load shifting; and that households seem to select into their initial plans based on their capacity for load shifting, which is reflected in the substantial heterogeneity in the plan choice intercepts.

To facilitate the interpretation of the additional load shifting results, Panel (a) of Figure 5 plots the load-shifting disutility by month. The load-shifting disutility

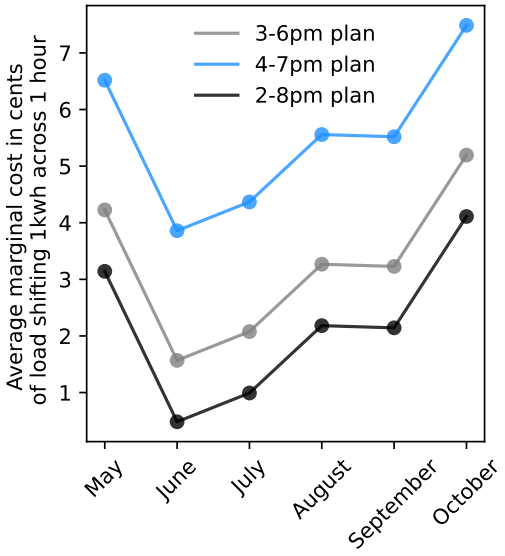
Table 3: Daily consumption model results

Parameter	Coef.	SE	**Parameter**	Coef.	SE
Load shifting disutility (β_d)			**Load shifting disutility, contd.**		
Duration (β_{d1})	0.057	(0.015)	*Month intercepts:*		
9am-5pm (β_{d2})	5.884	(1.125)	May	0.002	(0.001)
Demographics:			June	-0.004	(0.001)
Household size	0.000	(0.000)	July	-0.003	(0.001)
Income	-0.003	(0.002)	September	-0.000	(0.000)
Initial plan choice:			October	0.005	(0.001)
3-6 pm plan	0.008	(0.004)	**Loss function (β_ω)**		
4-7 pm plan	0.014	(0.003)	*Demographics:*		
2-8 pm plan	0.005	(0.001)	Household size	0.202	(0.089)
TOU-responsive bliss points:			Income	0.786	(0.138)
Draw std. dev. (σ_ν)	1.257	(0.036)	*Initial plan choice:*		
Demographics:			3-6 pm plan	4.522	(0.854)
Household size	0.254	(0.138)	4-7 pm plan	2.614	(0.566)
Income	8.739	(1.724)	2-8 pm plan	6.353	(1.428)
Initial plan choice intercepts:			**TOU-unresponsive bliss points:**		
3-6 pm plan	-0.987	(0.211)	Draw std. dev. (σ_ν)	3.963	(0.124)
4-7 pm plan	-1.097	(0.334)	*Demographics:*		
2-8 pm plan	-0.854	(0.300)	Household size	-0.018	(0.004)
Month intercepts:			Income	12.667	(0.656)
January	-0.197	(0.056)	*Initial plan choice intercepts:*		
February	-0.257	(0.045)	3-6 pm plan	-0.067	(0.020)
March	-0.246	(0.047)	4-7 pm plan	-0.297	(0.099)
April	-0.182	(0.087)	2-8 pm plan	-0.278	(0.047)
May	0.114	(0.042)	*Month intercepts:*		
June	0.839	(0.136)	January	0.056	(0.016)
July	1.079	(0.132)	February	0.018	(0.004)
September	-0.261	(0.150)	March	0.019	(0.005)
October	-0.060	(0.011)	April	0.230	(0.095)
November	-0.255	(0.084)	May	-0.044	(0.010)
December	-0.065	(0.017)	June	-0.607	(0.142)
Mean-preserving spread parameters:			July	-0.632	(0.097)
3-6 pm plan, flatten	0.922	(0.184)	September	-0.106	(0.032)
4-7 pm plan, flatten	1.755	(0.349)	October	0.283	(0.088)
2-8 pm plan, flatten	1.236	(0.308)	November	0.188	(0.067)
			December	-0.178	(0.049)
			Mean-preserving spread parameters:		
			3-6 pm plan, flatten	0.805	(0.208)

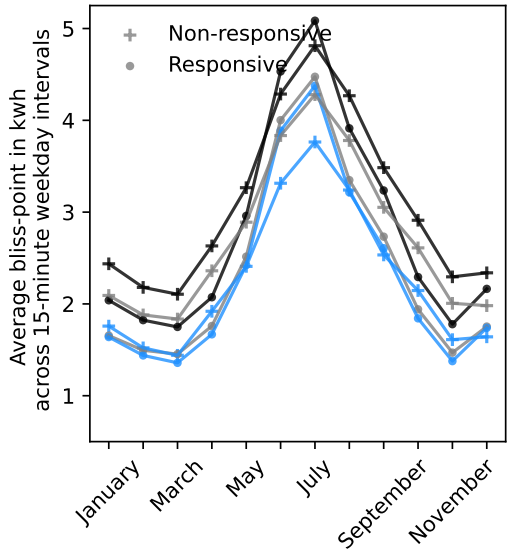
Note: The table reports results from the structural model of daily electricity consumption presented in Section 5.1 estimated via SMM. Parameter estimates are reported with bootstrapped standard errors in parentheses.

Figure 5: Estimates across months

(a) Load-shifting disutility by month



(b) Av. 15-minute bliss point by month



Note: Panel (a): plots the load-shifting disutility by month β_{d0m}, split across the distribution of consumers who self-select into each plan. Precisely, it is the (average) marginal cost to the consumer of load-shifting 1kwh across 1 hour of time to the right of the peak hour, when there is currently no load-shifting (i.e. the quadratic load-shifting term is at 0.0). There is no observed load-shifting November-April in our data, so we do not allow consumers to load-shift in those months in the model. These consumers are only TOU-responsive consumers, since TOU non-responsive consumers do not perceive any incentives to load-shift. Panel (b): plots the average 15-minute bliss point in weekdays by month split out by the distribution of consumers who self-select into each plan, using the same color scheme for plans as Panel (a). Recall that our model outputs an individual bliss point for each 15-minute interval in the day, so this panel is designed to show changes across the year rather than within-day.

decreases in the hottest months of the year in Arizona (especially June and July) and is lower in the other months as the weather gets cooler.[54] These patterns are consistent with consumers using different load-shifting technologies across the year. Concretely, at the height of the summer, consumers are able to load-shift by changing when they switch on their air conditioning. In months with cooler weather, consumers use less air-conditioning and so are forced to load-shift in ways that generate a higher disutility (e.g. by changing when they use cook dinner using

[54]Note that we do not observe and load-shifting outside these months and so we do not allow consumers to load-shift in November-April.

their oven).

Panel (a) of Figure 5 also clearly shows that consumers sort across the plans by their ability to load-shift. Consumers with the lowest load-shifting disutility sort into the 2-8pm peak plan. This is intuitive because the 2-8pm peak plan requires the furthest load-shift across time to profitably move consumption to off-peak periods.

The bliss point estimate for income for both TOU-responsive and -unresponsive households indicates that richer households desire more electricity at a given point in time; however, only TOU-responsive households see their bliss points increasing with household size. The remainder of the bliss point parameters are graphically depicted in Panel (b) of Figure 5. The figure demonstrates that for both types of households, bliss points are highest in the hottest months of the year which makes sense given the extreme summer temperatures in Phoenix. Further, Panel (b) shows that there is likely selection into initial price plan based on a household's bliss point, particularly for households in the 2-8pm plan.

Intuitively, the loss function parameters on household demographics presented in Table 3 suggest that larger and higher income households are willing to pay more in order to be closer to their bliss point consumption levels. We also see substantial heterogeneity across initial plan choice that indicates households with a higher desire to consumer near their bliss points are most likely to select into the 2-8pm peak plan. Recall that the 2-8pm peak plan has a narrower price differential between on-peak and off-peak hours. Thus, this result suggests households with higher willingness to pay to be near their bliss points are selecting into the plan that requires them to make smaller but more frequent deviations from their bliss point throughout the day.

7.2 Plan choice

Table 4 reports the parameter estimates from the plan choice model presented in Section 5.2. Panel A of the table includes the model of attention and Panel B the parameters of the choice model.

Table 4: Plan choice model results

Panel A: Attention parameters			**Panel B: Choice parameters**		
Parameter	Coef.	SE	**Parameter**	Coef.	SE
Constant	-7.045	(0.299)	Incumbent plan	0.151	(0.061)
Income	-0.200	(0.087)	*Incumbent plan ×:*		
Household size	-0.070	(0.033)	Income	0.996	(0.476)
Initial plan dummies:			Household size	0.157	(0.056)
4-7 pm plan	-2.520	(0.936)	Init. 4-7 pm plan	5.990	(2.430)
2-8 pm plan	0.950	(0.506)	Init. 2-8 pm plan	1.674	(0.305)
Month dummies:			σ_ϵ	0.013	(0.005)
February	-0.041	(0.017)			
March	0.098	(0.040)			
April	0.685	(0.280)			
May	1.227	(0.245)			
June	1.723	(0.251)			
July	1.832	(0.232)			
August	1.782	(0.240)			
September	1.603	(0.239)			
October	1.077	(0.282)			
November	0.813	(0.310)			
December	0.787	(0.343)			

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.

The parameters of the attention model indicate that households do not frequently search for new plans. In fact, when converting the reported parameters into actual probabilities of search, we see that households in our sample only search roughly 0.37 percent of months. Interestingly, we find that income and household size are not particularly strong predictors of search. However, there is a strong seasonal trend in the likelihood of receiving an attention shock. For example, a household with average income and household size that was initially in plan E21 (the most common plan, and the base plan in the model) has 0.07 percent chance of drawing an attention shock in February, but that probability rises to nearly 0.47 percent in July.

The plan choice model results presented in panel B of the table indicate that,

conditional on deciding to choose a plan in a given month, households still face some inertia. However, for the mean household, monthly inertia is equivalent to only about \$1.45. Additionally, the small value of $\sigma = 0.013$ suggests that households that reach the choice stage are typically choosing the plan that will most benefit them. Overall, there are very few switchers (both in the data and the model), which the results in Table 4 suggest are mostly driven by a lack of attention/deciding to choose, rather than inertia in the choosing stage.

8 Counterfactuals

8.1 Overview: the menu design problem

We characterize the menu design problem in the following way:

$$\max_{\mathcal{P}\in\mathcal{M}} \sum_i \underbrace{V_{i,j(i)} - \text{Private costs}_{i,j(i)} - \text{Damages}_{i,j(i)} + \text{Expenditure}_{i,j(i)}}_{\text{Total Annual Welfare of Consumer } i \text{ Selecting Plan } j(i)} \tag{12}$$

Subject to:

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

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

Here, the social planner offers a set of plans $\mathcal{P}$ from a class of potential menus $\mathcal{M}$. We vary this class $\mathcal{M}$ across the counterfactuals. For example, in one counterfactual $\mathcal{M}$ is all permutations of existing plans. In all of our counterfactuals we incorporate the two non-TOU plans in $\mathcal{M}$, consistent with how our electricity utility redesigns their plan menus.

Overall, the objective of the planner is to maximize total social welfare. This incorporates: the indirect utility of consumer i's optimal choice of plan $j(i)$, $V_{i,j(i)}$, in dollars; the private costs of electricity generation; damages to the environment; and, an adjustment for consumer expenditure. Consumer expenditure is also embedded in $V_{i,j(i)}$ and so this simply nets out expenditure (which is just a transfer from the utility to consumers) in social welfare.

The social planner faces two constraints. First, consumers choose their plan optimally, and we denote $j(i)$ as the optimal plan for consumer i. The standard rationale for this across-plan implementability constraint in contract theory is that the planner does not have information about the private types of each consumer. In our setting, an additional reason for this constraint is that the planner cannot contract directly on private types, even if they have some information about consumers. The second constraint is within-plan implementability: given plan j's incentives consumer i chooses their 15-minute interval consumption optimally across weekdays, weekends, and different months of the year, potentially load-shifting and load-shaving in response to incentives.

Discussion of assumptions We now justify three assumptions we make in the social planner's problem. First, we do not incorporate attention and inertia considerations in the 'across-plan implementability' constraint. However, all the plan design problems we consider involve eliminating existing TOU plans and therefore forcing these consumers to make an 'active' choice.

The second assumption is that we do not allow non-TOU consumers to switch their plans to the new TOU plans. This is because, as mentioned in the estimation section, it is difficult to estimate these consumers' ability to load-shift since we do not observe them in plans with within-day price changes.

Our second assumption is arguably justified for several reasons. First, unlike consumers on TOU plans who would need to make an active choice when the new menu is offered, non-TOU customers would face inertia and inattention. Our empirical estimates suggest these behavioral choice frictions are substantial. Second, since these consumers are not choosing the TOU plans in the baseline menu, they may be of a type who cannot profitably respond to the TOU incentives and so find it valuable to switch. Indeed, we actually see existing TOU customers switching out of TOU plans in all the counterfactual menus. This suggests that they are not particularly generous compared to the empirical baseline menu. Third, it is always

Table 5: Summary of counterfactuals

	Change vs baseline menu (dollars/year/consumer)				
	Welfare	Consumer surplus	Private costs	Damages	Expenditure
Eliminate TOU plans	25.6	-53.7	2.4	-7.9	73.7
First best:					
With planner's allocation	79.6	-210.3	-61.7	-88.6	139.7
With self-selection	34.7	-44.5	-4.9	-17.0	57.3
Menu designs:					
Base plans + optimize:					
1 peak/off TOU plan	74.5	42.6	-23.8	-45.8	-37.8
2 peak/off TOU plans	74.6	45.6	-22.7	-45.2	-39.0

possible to explicitly prevent non-TOU from customers switching by offering only TOU customers a special 'deal' of the new menu.

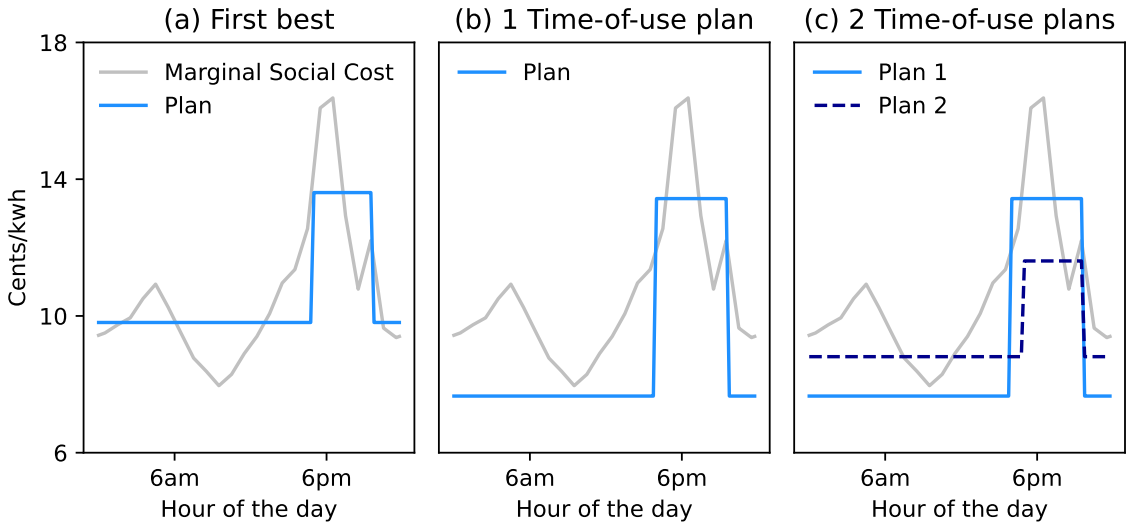
Our third assumption is that we do not allow consumers to not choose a plan at all (i.e. a participation constraint). However, since we always keep the two non-TOU plans available in our counterfactual menus, this acts effectively as an outside option that consumers can switch to if the new TOU menu is not generous enough.

8.2 Does offering the existing time-of-use plans increase social welfare?

We summarize the counterfactuals in Table 5. We first consider the effects of eliminating the three TOU-plans entirely. Here, $\mathcal{P}$ just consists of the two non-TOU plans (a block-rate plan and a fixed-rate plan). We find that eliminating the TOU plans would *increase* social welfare by 25.6 dollars per year per consumer. Therefore, overall, the existing menu of TOU plans does not appear to be well-designed. In other words, a poorly design menu that includes TOU plans — even when these plans correlate price with marginal social cost — can be worse than no TOU plans at all.

Table 5 illustrates where the inefficiency is coming from. While consumers select into the TOU plans based on their total utility (and so their consumer surplus would decrease if we eliminate these plans) the plans do not align their private incentives

Figure 6: 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 5-10pm. In Panel (b) the peak hours are 4-10pm. In Panel (c), plan 1's peak hours are 4-10pm and plan 2's peak hours are 5-10pm.

with social incentives. This is because of two reasons. First, in terms of the TOU-nonresponsive consumers, the base plans have an average price that is closer to the social marginal cost. So they are better aligned with social incentives on average. Second, although the TOU-responsive consumers load shift a lot, the hours are not well-designed. For example, compare the most common TOU plan (3-6pm) with the social marginal costs around 6pm. Consider a consumer who shifts their load from 5pm to 7pm: in some cases, it is actually *more* costly to produce electricity at 7pm than 5pm. In other words, it is not always the case that load-shifting improves social welfare.

It is possible that offering some subset of the TOU plans could improve welfare beyond just eliminating them entirely. Therefore, we also experimented by searching over a set over menus $\mathcal{M}$ that includes all combinations of the empirical TOU plans

(plus the base plans). However, we found that no such subset exists.

8.3 The 'first-best' benchmark

Our next set of counterfactuals considers the benefits of the 'first-best' benchmark. By 'first-best' we mean both (i) a menu $\mathcal{P}_{first}$ and (ii) an allocation of consumers to plans in this menu that would maximize social welfare if the planner knew the exact type of each consumer. Equivalently, this is the solution to the menu design problem if the planner did not face the 'across-plan implementability' constraint.

Computation When searching for the first-best menu and allocation, we restrict $\mathcal{M}$ to only contain combinations of single-peak TOU plans plus the base plans. So, it is best interpreted as a 'constrained' first best. These reflect the TOU plans in our data (with one exception): they contain an on-peak period for a connected set of hours, and an off-peak period otherwise.[55]

We allow the planner to vary peak hours as well as on and off peak prices in 1c/kwh discrete intervals.[56] We discuss further in the Appendix exactly how we construct these increments. Overall, we optimize over thousands of potential plan designs and millions of different potential menus of these plans. In each case we individually compute the utilization of each consumer i for every 15-minute interval across weekdays/weekends/months throughout the year.

Results The first-best menu $\mathcal{P}_{first}$ and allocation is quite simple. It consists of a single TOU plan where the prices and peak hours on the plan provide the closest approximation to real-time-pricing at the marginal social cost (as well as the base

[55]The exception is the 2-8pm TOU plan in Winter, which has two peak period with relatively minor price changes. One may ask about more general $\mathcal{M}$ which contain multiple peaks. Given the number of potential menus that arise when consider the single-peak case, it would be computationally challenging to flexibly search over this dimension as well. As well, the regulator is our setting appears to constrain itself to these kinds of plans. Furthermore, increasing the number of peaks in the plan increases the complexity for consumers and it is not clear how predictive our model would be about their responses: in the limit, offering multiple peaks approaches real-time-pricing which is arguably infeasible.

[56]For example, one potential plan has peak hours at 5-8pm with an on-peak price of 11c/kwh and an off-peak price of 9c/kwh. Another plan has peak hours at 4-9pm with an on-peak price of 10c/kwh and an off-peak price of 8c/kwh.

plans). The planner then allocates all consumers to the TOU plan.

We plot the design of the first-best plan against the marginal social cost in Figure 6. This plan increases social welfare by 79.6 dollars per consumer/year on average. However, this allocation would result in a very large decrease in consumer surplus of -210.3 dollars per consumer/year on average. Such a decrease would likely be politically infeasible for our publicly-regulated utility.

Therefore, we next consider offering $\mathcal{P}_{first}$ with consumer choice, by also reintroducing the 'across-plan implementability constraint'. We find that the efficacy of this menu is undermined by consumer self-selection, cutting the benefits of $\mathcal{P}_{first}$ by over one half.

Overall, the results indicate an important caveat to the textbook solution that offering a plan that prices at marginal social cost is socially optimal. Concretely, when consumers have access to other options — here, the base plans — the social planner may be able to improve welfare by accounting for self-selection and *not* pricing at marginal social cost. This motivates our analysis in the next section, where we make this point quantitatively by solving for the optimal menu under self-selection.

8.4 Could alternative menu designs further improve social welfare?

Offering base plans plus a single optimally-designed TOU plan Next, we optimize over a class of menus $\mathcal{M}$ that contains the base plans and a single TOU plan. We use the same procedure for determining the various TOU plans to the first-best menu (single-peak TOU plans, varying peak hours and on/off peak prices in discrete intervals of 1c/kwh). This results in thousands of potential menus in $\mathcal{M}$. We solve the problem in Equation 12 including both implementability constraints.

We find that the optimal menu would increase social welfare by 74.5 dollars per consumer/year. Moreover, we find that all stakeholders – the environment, the electricity utility, and consumers — do better on average from offering this new menu of TOU plans. Furthermore, this plan achieves over 93% of the social welfare

benefits of the first-best plan.

This menu requires distortions to the TOU plan to provide "information rents" to get consumers to select into the plan. These rents are illustrated in Figure 6: the off-peak price is lower than in the first-best plan. As a result, while these provide socially beneficial across-plan incentives, they also undermine within-plan incentives. Specifically, since consumers in the plan face a price that is lower than social marginal cost off-peak, they consumer too much during these periods.

Offering base plans plus multiple optimally-designed TOU plans We next consider the benefits of offering more than one TOU plan in the menu (plus the base plans), similar to our electricity utility's current design. We consider the case of two TOU plans.[57]

We do not find much extra benefit to adding an additional TOU plan to the menu. Part of the reason is that the optimally-designed menu with one TOU plan is extremely effective at increasing social welfare towards the first-best 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 selection if they are not well-designed (as is the case in our empirical setting).

9 Conclusion

We demonstrated how to design a socially efficient menu of electricity price schedules, subject to realistic design constraints. We started by using administrative data from a large utility to develop descriptive evidence on how consumers in an information-rich market environment choose from a relatively large menu of static and dynamic price schedules and respond to hourly price changes. This evidence guided our development of a novel model of plan choice, consumption, and intertem-

[57]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.

poral substitution. We used the model estimates to evaluate millions of counterfactual menus. Our findings suggest that the current menu can be redesigned to simultaneously increase consumer welfare, reduce private generation costs, and improve environmental outcomes.

The potential for using these insights to improve market efficiency is strong. While nearly three-quarters of US households have "smart" electricity meters that are equipped for time-of-use pricing, only nine percent are enrolled in time-of-use plans (EIA, 2024). Moreover, the benefits of socially efficient menu design are likely to grow in the future as climate change increases electricity demand. Our results for the Arizona desert preview how menus can be designed efficiently in areas where extreme heat creates a strong demand for cooling.

References

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-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.

_ , "EIA-861, annual electric power industry report, detailed data files," Technical Report, United States Energy Information Administration 2024.

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.

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.

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.

Lang, Corey, Yeuming (Lucy) Qiu, and Luran Dong, "Increasing voluntary enrollment in time-of-use electricity rates: Findings from a survey experiment," *Energy Policy*, 2023, *173*, 113410.

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

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.

US EPA, "Report on the Social Cost of Greenhouse Gasses: Estimates Incorporating Recent 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 Shanjin 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.[58] Figure A.1 shows the enrollment form that an existing customer could use to switch plans.[59] 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 1d 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*

[58] 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).

[59] 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.

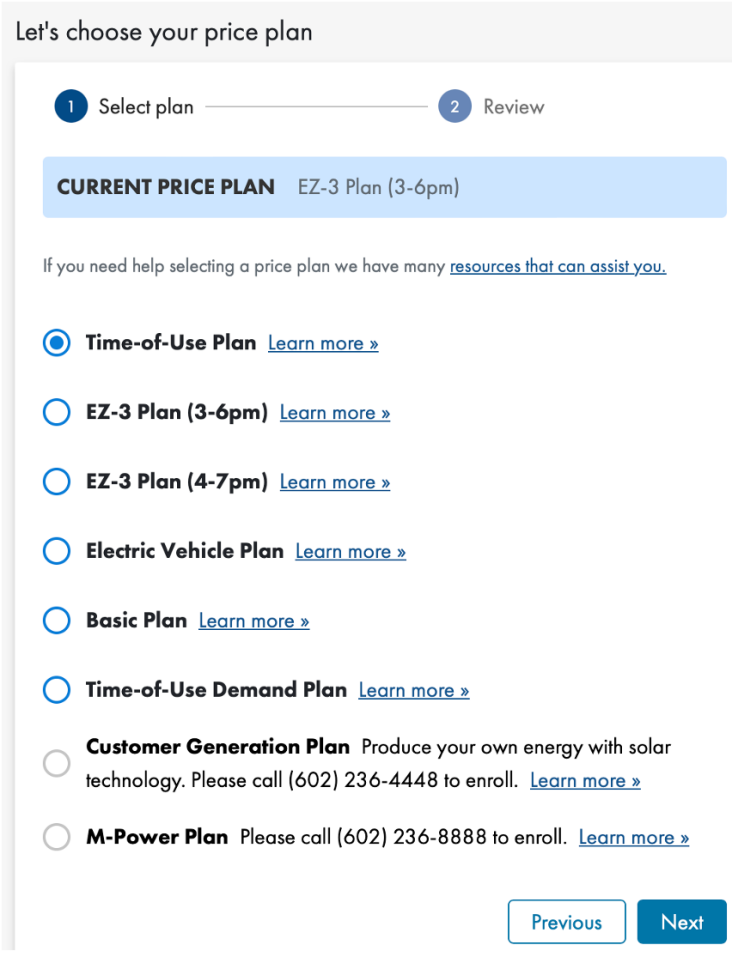


Figure A.1: Screen Shot of SRP's Online Enrollment Form

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

the season.

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

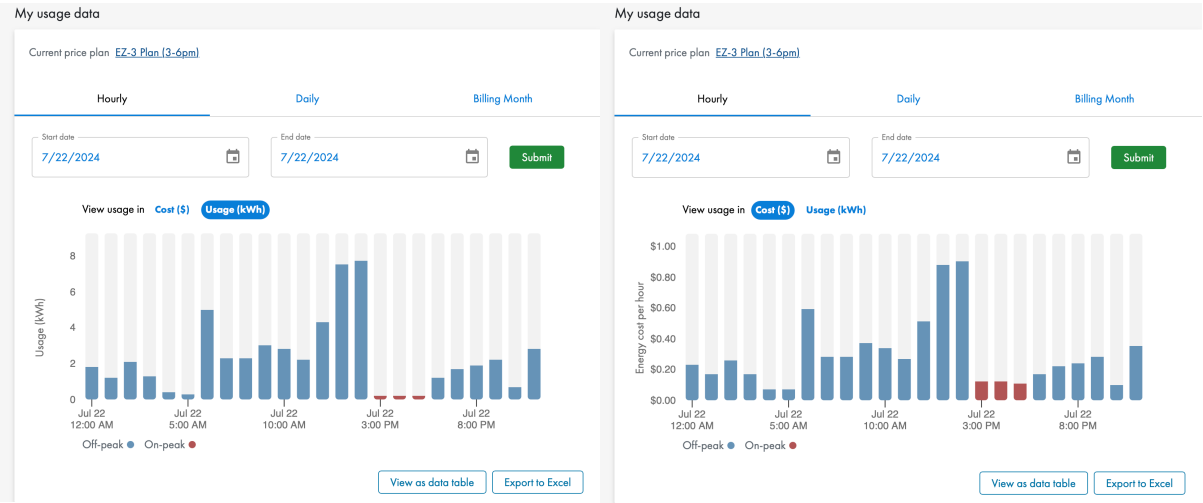


Figure A.2: Screen Shot of a Customer's Usage Data

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.

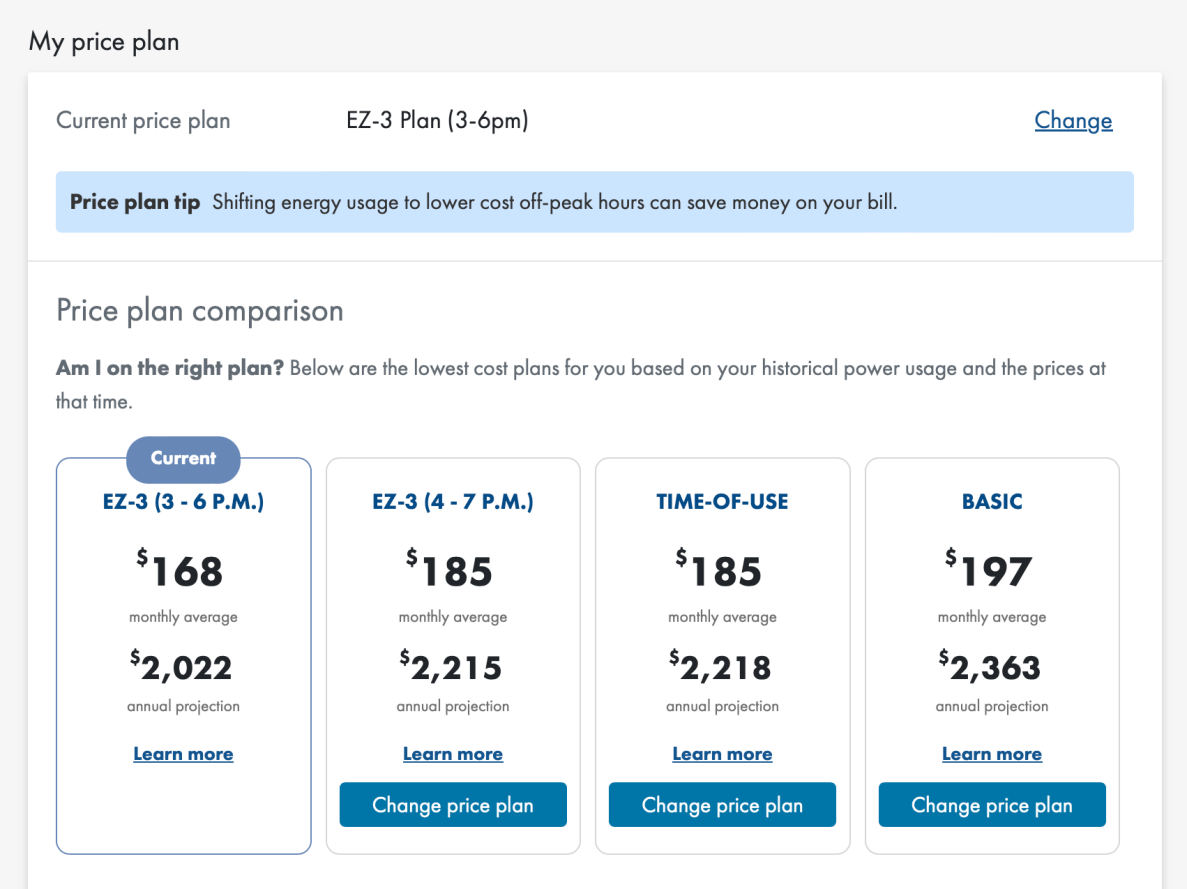


Figure A.3: Information about a Customer's Potential Savings from Switching Plans

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 sec-

tion 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

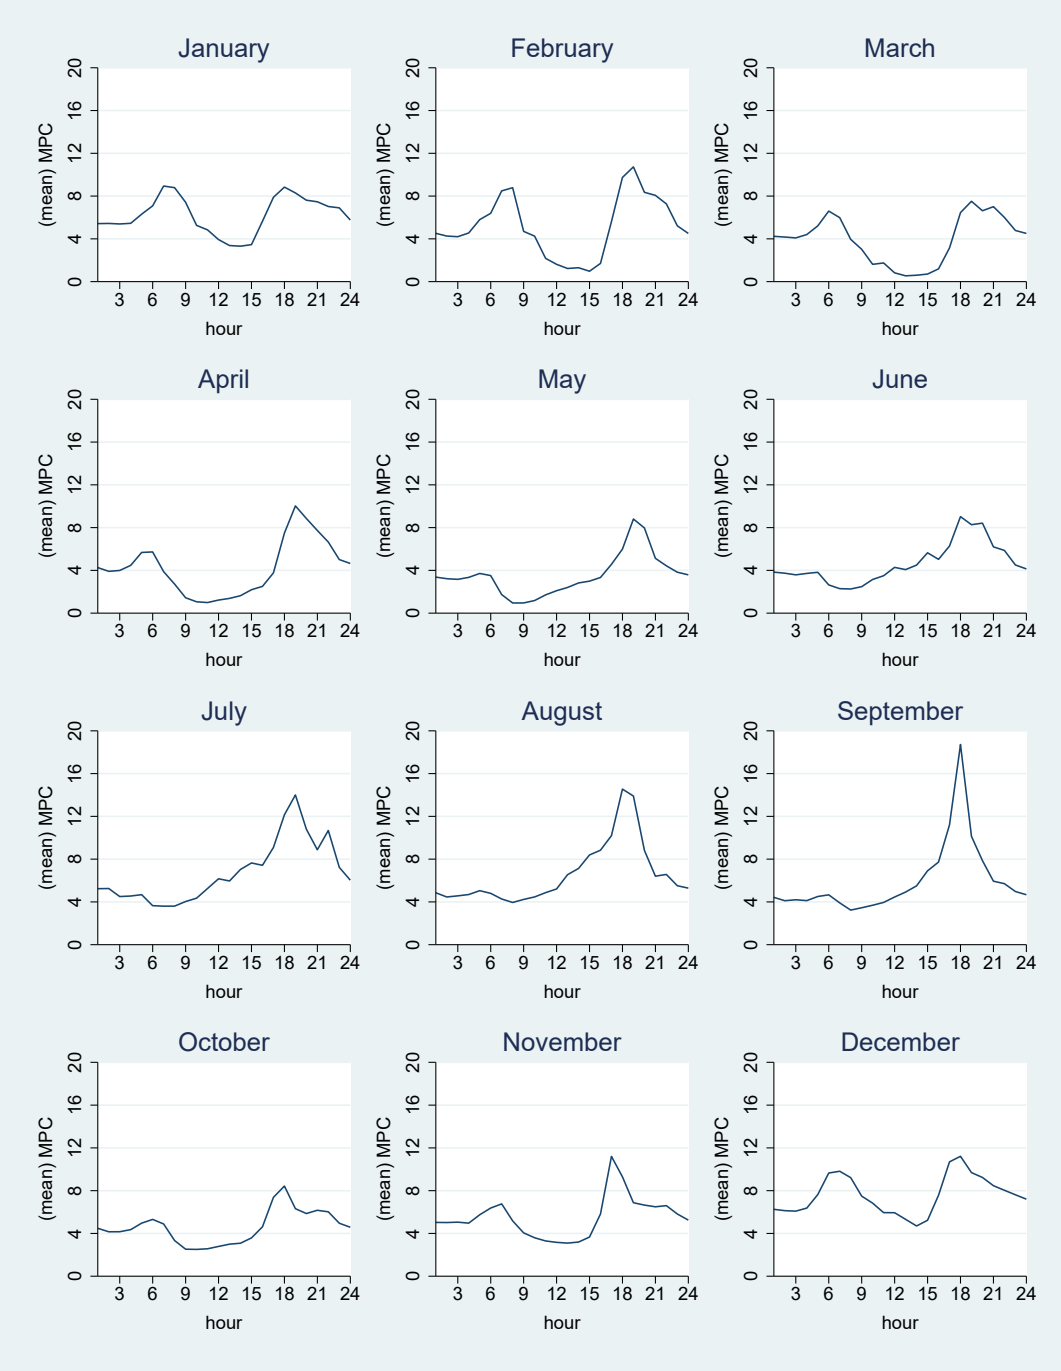


Figure B.1: Marginal Private Cost of Generation by Month and Hour (cents/kWh)

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.

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 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.[60] 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

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

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: 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.

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.

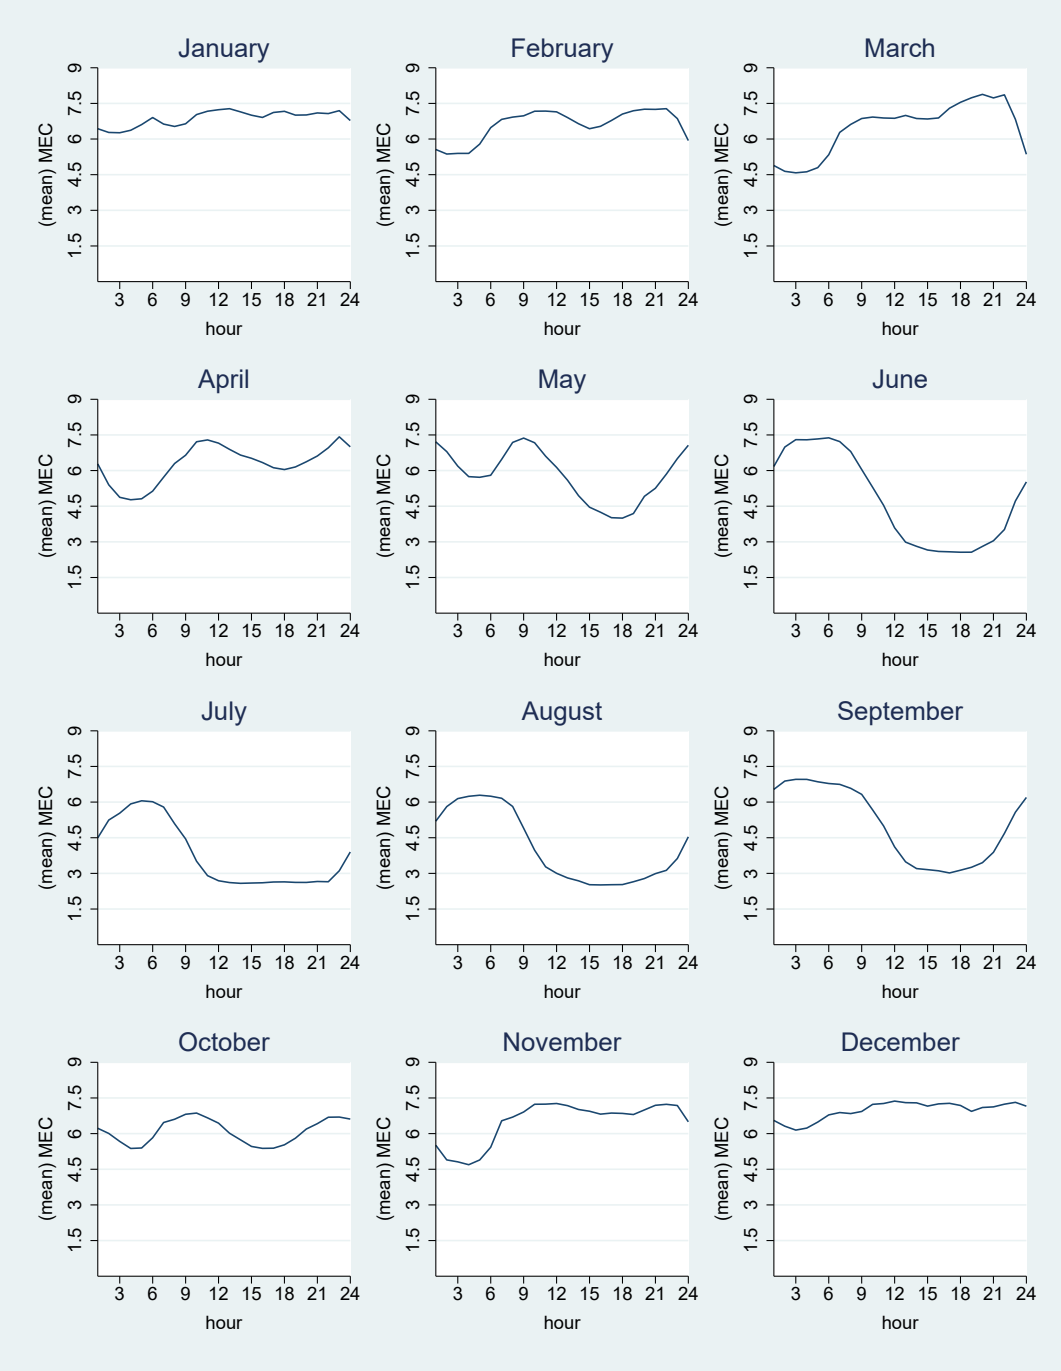


Figure B.2: Marginal External Cost of Generation by Month and Hour (cents/kWh)

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.

C Other Data

C.1 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.

C.2 Tax Rates for Electricity

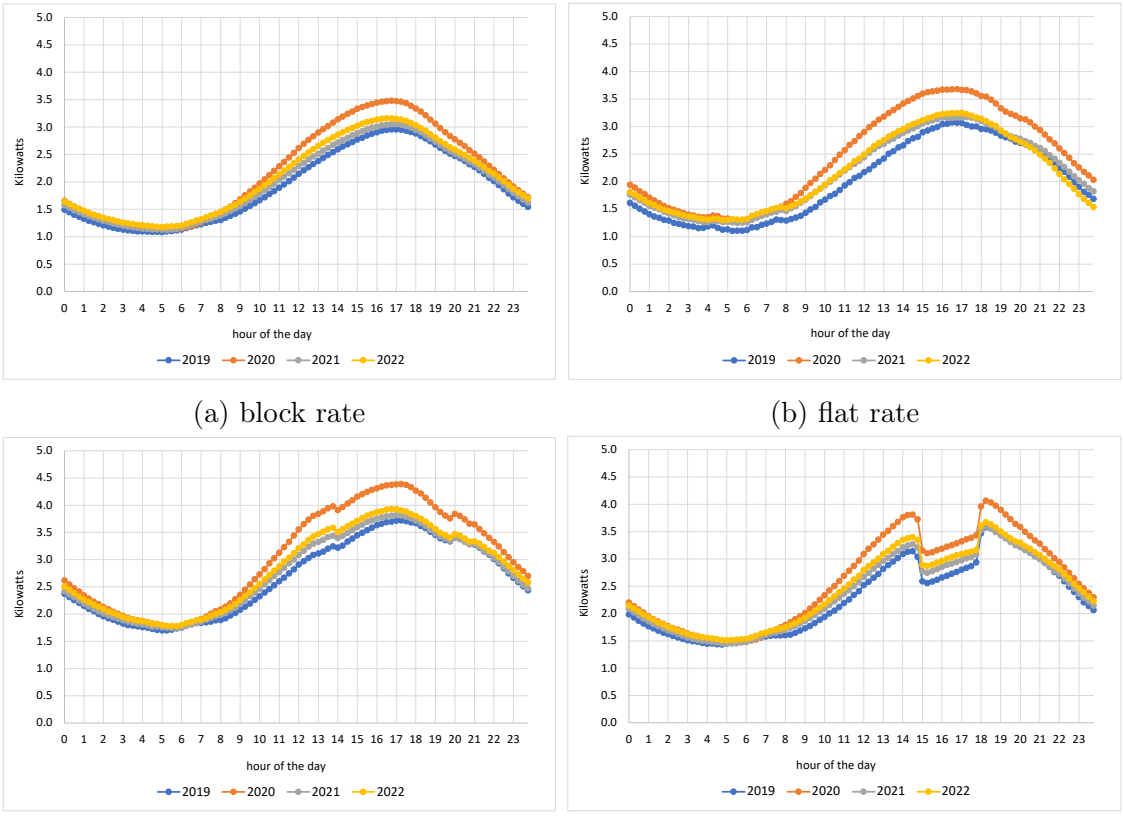
Data on state, county, and city tax rates for electricity were drawn from the Arizona Department of Revenue's "Transaction Privilege and Other Tax Rate Tables" effective April 1, 2023. The combined tax rate for the state of Arizona and Maricopa county was 6.3% throughout our study period. The vast majority 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

D Additional Descriptive Evidence

D.1 Electricity Use During the Covid-19 Pandemic

Figure D.1 shows year-specific average load curves from May through October for households in four of the five plans. The fifth plan is the Superpeak TOU plan from 4pm-7pm (not shown). It looks similar to Figure D.1d but noisier due to its

small enrollment share. The orange lines at the top of each panel correspond to May through October of 2020 when SRP households faced various federal, state, and local social distancing polices due to the Covid-19 pandemic and there was a surge in remote work and remote schooling.



(a) block rate

(b) flat rate

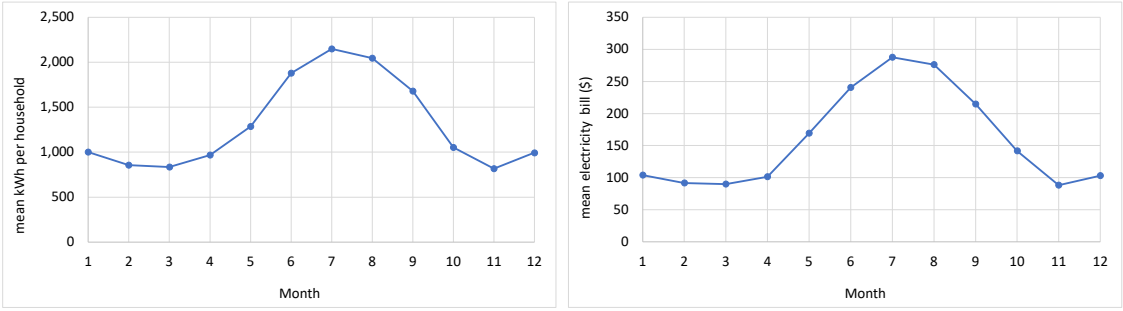
(c) standard time-of-use

(d) superpeak time-of-use

Figure D.1: Average load curves for May through October, by Year

D.2 Summary Statistics

Figure D.2a shows average monthly electricity consumption per household and Figure D.2b 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.2: Average monthly electricity consumption and expenditures

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.[61] 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.

	All plans	block rate	fixed rate	standard time-of-use	superpeak 3pm-6pm	superpeak 4pm-7pm
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.

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.

[61] 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".

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 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.[62]

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.3 Additional Descriptive Evidence

Figure D.3 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.[63] 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 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

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

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

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.

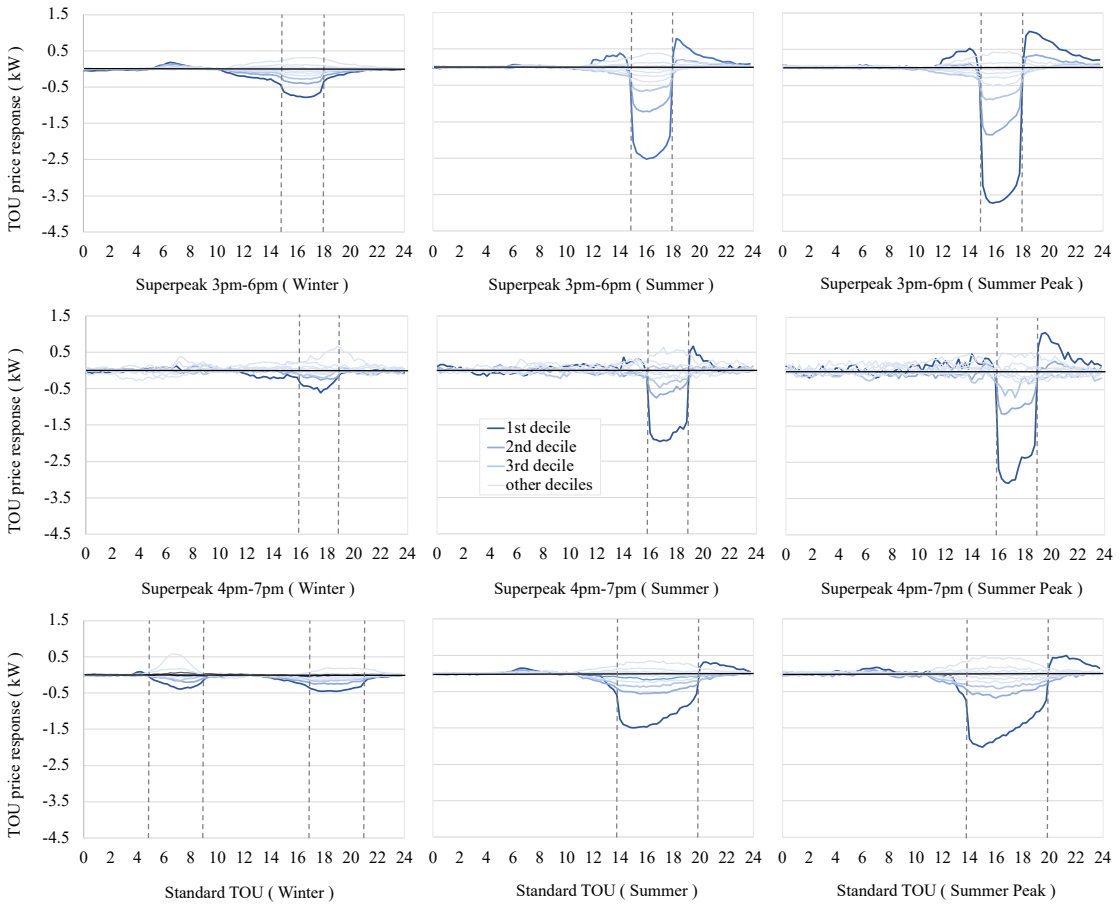


Figure D.3: TOU responsiveness by plan and season

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.4 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 4.3. 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).

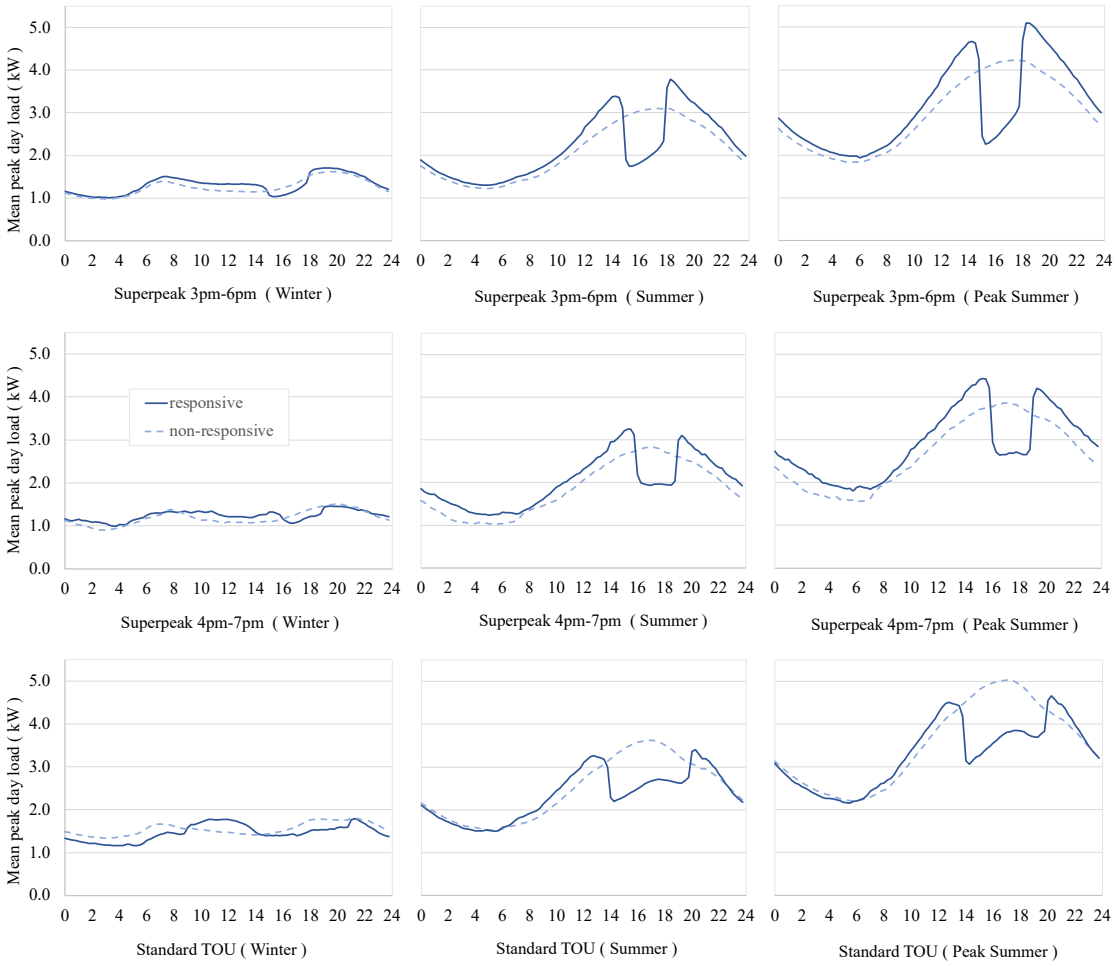


Figure D.4: Load curves by season and responsiveness to TOU pricing

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.5 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).

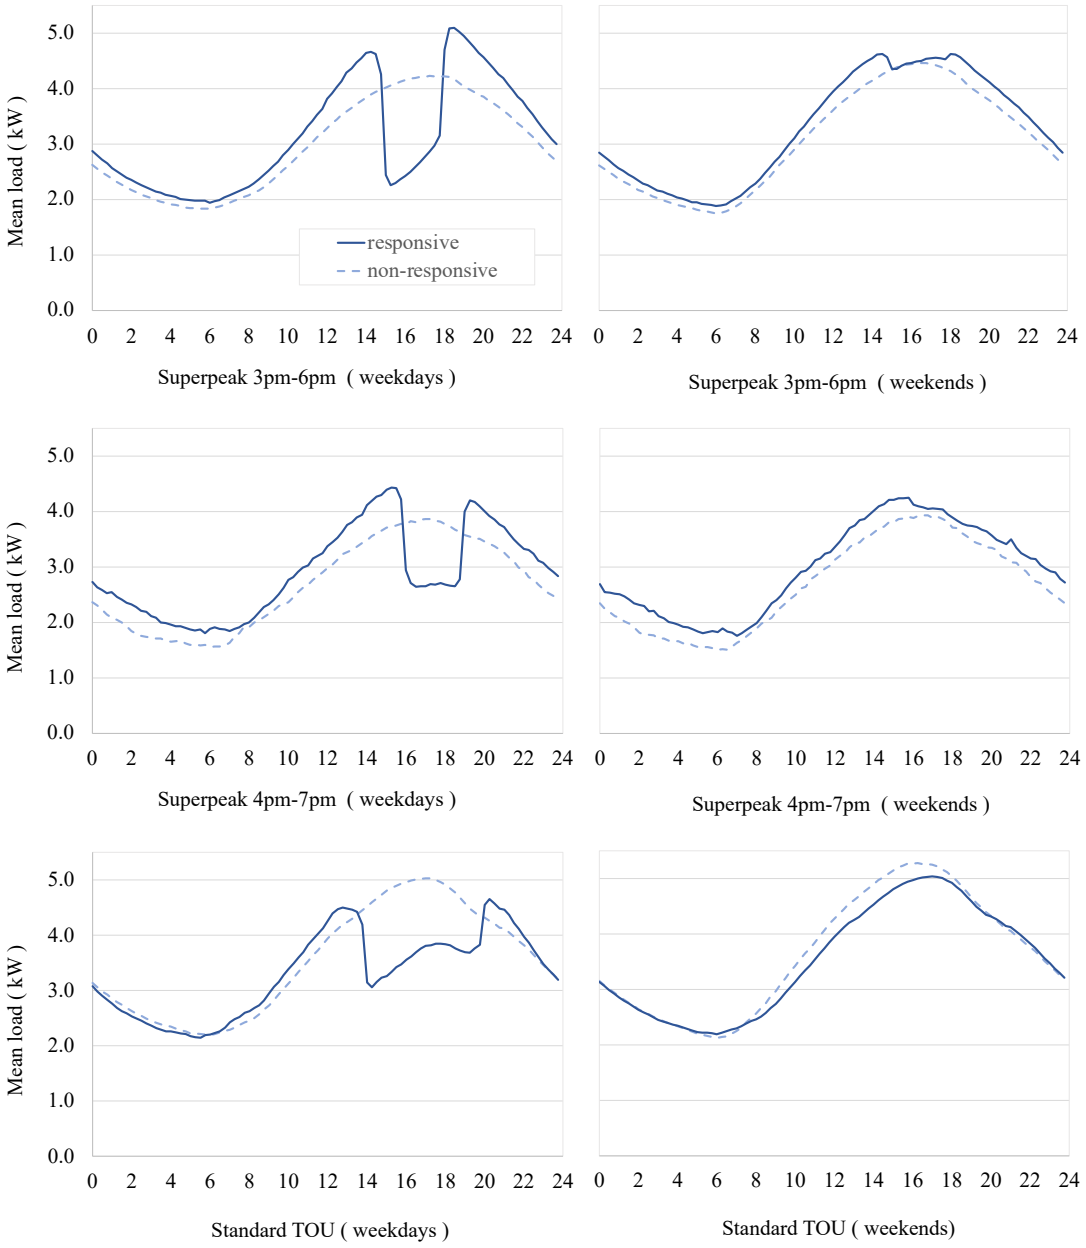


Figure D.5: Load curves by peak day and responsiveness to TOU pricing

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 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.

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 plan, we estimate the following regression:

$$\begin{aligned} C_{i,m,y} = \beta_0 + \beta_1 \bar{P}_{i,m,y} + \beta_2 \bar{P}_{i,m,y} \times Superpeak\ 4pm-7pm_i \\ + \beta_3 \bar{P}_{i,m,y} \times Standard\ TOU_i + \beta_4 \bar{P}_{i,m,y} \times Income_i \\ + \beta_5 \bar{P}_{i,m,y} \times Household\ size_i + \rho_i + \psi_m + \phi_y + \delta_{i,m} \\ + \pi_{m,y} + \varepsilon_{i,y,m}, \end{aligned} \tag{E.1}$$

where $C_{i,y,m}$ is household i's electricity consumption (in kWh) for calendar month m in year y; $\bar{P}_{i,m,y}$ 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.[64] Table E.1 presents the results.

[64]To make sure that the elasticities generated from Equation (E.1) are comparable to the model estimated counterparts for TOU-responsive consumers, we treat income bins as continuous and

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

$\bar{P}_{i,m,y}$	-5665.42**
	(2688.73)
$\bar{P}_{i,m,y}$ × Plan 22	12856.67***
	(2748.11)
$\bar{P}_{i,m,y}$ × Plan 26	6787.12***
	(1695.42)
$\bar{P}_{i,m,y}$ × income	-63829.67***
	(22007.82)
$\bar{P}_{i,m,y}$ × HH size	-2092.45*
	(1129.39)
Constant	2617.24***
	(120.26)
Overall adjusted R-squared	0.92
N	2693

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

In order to map the coefficient estimates in Table E.1 to the model parameters, we add the base level price response to the relevant interaction term. Then we rescale the value so that it corresponds to a 15-minute interval of the model rather than at the monthly level (i.e., divide by 30 days × 24 hours × 4 intervals). For example, the β_2 parameter for TOU non-responsive consumers would be:

$$\beta_2 = \frac{-5665.42 + 12856.67}{30 \times 24 \times 4} \approx 2.5.$$

divide them by 100000. Similarly, we treat the household size variable as continuous.

E.2 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.[65] 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.

[65] 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.3 Model outputs

Table E.2: 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	1.28	2-8pm plan, August	3.16	3.13
4-7pm plan, August	1.27	1.34	4-7pm plan, August	2.90	2.87
3-6pm plan, August	2.16	2.22	3-6pm plan, January	1.36	1.34
*Corr(*Δ*, income)*:			3-6pm plan, February	1.32	1.19
3-6pm plan, August	0.30	0.30	3-6pm plan, March	1.18	1.16
*Corr(*Δ*, hhold size)*:			3-6pm plan, April	1.44	1.44
3-6pm plan, August	0.20	0.20	3-6pm plan, May	1.98	2.01
Load-shift 3hr post-peak			3-6pm plan, June	2.78	3.46
Mean:			3-6pm plan, July	3.12	3.89
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	2.22
3-6pm plan, May	1.71	1.72	3-6pm plan, October	1.52	1.47
3-6pm plan, June	4.96	3.85	3-6pm plan, November	1.20	1.17
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	2.47	3-6pm plan, August	1.40	1.40
3-6pm plan, October	0.91	0.90	*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	2.53
3-6pm plan, August	2.09	2.38	4-7pm plan, August	2.60	2.87
Load-shift to 11am			3-6pm plan, August	2.64	2.76
Mean:					
3-6pm plan, August	0.00	0.00			

Table E.3: 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	3.13	3-6pm plan, October	1.52	1.47
4-7pm plan, August	2.90	2.87	3-6pm plan, November	1.20	1.17
3-6pm plan, January	1.36	1.34	3-6pm plan, December	1.37	1.43
3-6pm plan, February	1.32	1.19	*Std. dev:*		
3-6pm plan, March	1.18	1.16	3-6pm plan, August	1.40	1.40
3-6pm plan, April	1.44	1.44	**Max-min consump.**		
3-6pm plan, May	1.98	2.01	*Mean:*		
3-6pm plan, June	2.78	3.46	2-8pm plan, August	2.53	2.53
3-6pm plan, July	3.12	3.89	4-7pm plan, August	2.60	2.87
3-6pm plan, August	2.97	2.77	3-6pm plan, August	2.64	2.76
3-6pm plan, September	2.48	2.22			