It's About Time: Transitioning to Time-of-Use Pricing and Consumer Demand for Electricity

Click Here for current draft

Greg Kirwin

November 4, 2024

Abstract

This paper examines the impact of a state-mandated transition of electric utility rates from increasing-block pricing to time-of-use on consumer demand. Consumers were given the option to opt out of the transition, and selection out of the program is correlated with consumers' perceived gains or losses. After the program is implemented, I find limited evidence that consumers on time-of-use rates make significant changes to their consumption patterns, though their billing costs fluctuate in accordance with the new rates. These results appear to be impacted by household self-selection out of the program, leading to only minor changes by the remaining households.

1 Introduction

Despite an abundance of scholarship on the efficiency of time-based pricing, residential electric utilities in the United States are most often metered via flat or increasing-block pricing[1]. Though time-of-use pricing has only recently seen broader implementation by utilities, papers describing its efficiency extend back to at least the 1940s. The pivotal work Boiteux (1960) describing the advantages of peak load pricing was originally published in the *Revue générale de l'électricité*, though it was not translated into English until 1960. Other early works include Houthhaker (1951), Hirshleifer (1958), and Steiner (1957), which described the issue of electricity pricing in the presence of time-dependent load. These papers were followed by others like Williamson (1966), Aigner and Leamer (1984), and Pressman (1970), which built on the foundation for capacity constraints and optimal electricity pricing with further theoretical solutions. In broad terms,

[1] According to Faruqui and Tang (2023), in 2021 only 8.7% of US households were on a time-of-use rate. While specific estimates for the US market are not easily accessible, block pricing and flat tariffs almost surely comprise the majority of the remaining plans.

these authors recognized that utility companies face a problem of needing significantly more capacity during short periods of the day, which then goes unused during the remaining hours. Under time-independent pricing, such as the common "block" pricing schemes, consumers who use less electricity during peak periods also effectively subsidize their neighbors, as they do not contribute to the peak load problem but face prices that account for it. Being able to charge customers based on the timing of their consumption thus allows utilities to adjust marginal prices based on the timing of their demand.

While these advantages have been known for decades, the main difficulty appears to have been the ability to implement such rates. Public utilities around the US began experimenting with peak pricing schemes during the 1970s and 1980s. Train and Mehrez (1994) discusses one such experiment in which consumers are given the option of choosing from a menu of time-of-use (TOU) rate. While they find that consumers do reduce their peak consumption, the authors cite the cost of implementing TOU metering, which requires new devices capable of measuring demand during specific hours, as costing the utility an additional \$2.17 per month in each household, while their experimental results indicated that three of the available plans would yield a total surplus increase less than \$2.00, and only one would yield a net gain in total surplus. Though they estimate that the utility's profits would increase under this plan as well, they conclude that implementing it at scale may not yield such results, and that further experimentation is required.

In this paper, I study the implementation of default TOU rates at scale, whereby consumers were automatically transitioned to the new rate structure unless they opted out before their rollout date. This presents two major questions. Did consumers exhibit advantageous selection in their decision to be transitioned to the new plan? What was the impact of default TOU enrollment on consumer demand amongst the transitioned households? While previous experimental results have indicated a significant household response to a change in rate structure, I find limited evidence that this was the case during the course of my sample. However, advantageous selection seems to have muted the potential response to the program, with households that were more likely to have faced higher post-switch costs instead electing to stay on the original plan.

The paper is structured as follows. First, I give an overview of similar literature. Next, I describe the program and provide relevant background on changes in electricity rate structures. I then investigate the existence of advantageous selection and subsequent results from the program using two different types of treatment effect estimators.

2 Relevant Literature

The primary contribution of this paper is to the literature on time-based pricing. Much of the literature to date has focused on the efficiency of time-based pricing relative to flat or block pricing through experiments of various sizes and scales. In Wolak (2010), for example, a utility sought to understand how their customers would respond to the implementation of a new rate structure. In the former, consumers were assigned to one of three different dynamic pricing schemes: hourly pricing, critical peak pricing (CPP), or critical peak pricing with a rebate (CPR). Under CPP, prices may temporarily increase when there are anticipated increases in demand, often due to warm weather, and consumers are notified ahead of time of the increase so that they can adjust their demand accordingly. The paper finds that CPP impacted hourly demand more than the other two rate types, particularly when combined with a smart thermostat that allowed customers to see real-time consumption. A more recent work by Hinchberger et al. (2024) quantify the economic efficiency of various dynamic pricing schemes using wholesale pricing data across several electricity markets in the United States, and find that TOU pricing reduces deadweight loss in the electricity market by approximately 10%, though it is not as efficient as real-time pricing, which is harder to implement.

Other papers like Ito et al. (2023) and Fowlie et al. (2021) have also modeled the social benefits and costs of switching consumers from one pricing scheme to another using additional experimental evidence. Fowlie et al. (2021) in particular thoroughly details the differences between opt-in and opt-out plan switching for consumers, concluding that defaulting consumers into a pricing switch is more effective at inducing desired reductions in peak demand than allowing consumers to opt in on their own, in large part due to issues with consumer awareness.

Finally, there are two contemporaneous works specifically on transitioning consumers to TOU pricing that are more directly related to this paper. Enrich et al. (2024) covers TOU implementation with an opt-out program in Spain. Similarly to this paper, the public utilities were required to default their consumers into time-of-use pricing by a government authority; where this paper differs is in the plan switch implementation. First, most Spanish consumers were on a flat tariff by default rather than increasing-block pricing. Second, the switch occurred for all consumers at the same time. Lastly, the Spanish government appears to have advertised the plan switch more aggressively than in California [2]. As a result, the authors find an immediate and lasting impact on electricity demand from the switch, with consumers reducing their consumption by 5.7% during the "mid-peak" period and 8.9% during the peak period under one specification, but only 1-2% under an alternative estimator. Bernard et al. (2024) assesses the impact of a staggered rollout of heat pumps in residential customers' homes in conjunction with a time-of-use tariff. Since heat pumps increase electricity usage in exchange for decreasing gas usage, the additional tariff is intended to account for anticipated increases in peak load. The authors find that not only did the heat pump installations decrease household greenhouse gas emissions, but the tariff successfully incentivized households to shift some of their electricity consumption during peak hours to off-peak.

3 Data

3.1 Data Acquisition

Data was obtained from Pacific Gas & Electric (PG&E) in California under the state's Energy Data Request Program. A random sample was chosen from each zip code in their service area that represented a fixed percentage of their customers in that zip. Zips with fewer than 100 active

[2] While not described in the paper itself, there was broader coverage of the transition, including explanations of optimal usage in the country's largest newspaper *El País*. See this link for more information

customers were excluded. To be eligible for my sample, the household must be zoned as a single-family home and must have been an active PG&E customer at that address for the entire four years. Multifamily homes are often "master-metered", meaning the meter is read for the entire building's usage and not individual units, and hence I wanted to exclude these households from my sample. Staying in one home four the entire four years is also vital because houses will have different base energy needs depending on their size; requiring static addresses ensures that consumers' baseline consumption is also more likely to be consistent across billing cycles. The data are anonymized to the zip code level, so while I observe the household's demand during this period, I cannot observe the house's characteristics or customer demographics.

PG&E provided the data in sets of one-hour increments over the four years, spanning 2018 through 2021, representing 75,000 households in the initial sample. For each customer, I also observe their monthly bill amount and total usage from 2018 through the end of 2022, plus any modifiers to their costs, such as participation in a subsidy program or solar interconnection. Customers have an assortment of plans that they may select during the course of the sample, and I observe any changes in their plan during the sample via a record indicating the beginning and end dates of their service agreement.

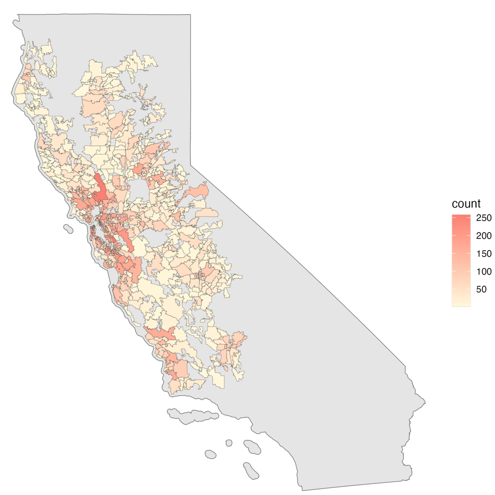

Figure 1: The number of households by zip code, removing alternative plans, solar, EVs, and the top and bottom percentiles of usage.

3.2 Electric Rate Transition

In PG&E's service, customers can select from two types of rates- time-invariant "block" pricing and "time-of-use" (TOU) pricing. Through the end of 2020, block pricing comprises the overwhelming majority of plans in my sample, about 90% of bills in a given month. Most consumers have "smart" meters, which enables the measurement of real-time consumption [3], at the beginning of the sample, and this coverage increases to nearly 100% by the end of the sample. Under block pricing, consumers face a uniform tariff rate that increases at set usage thresholds. Under TOU pricing, consumers face a constant price during "off-peak" hours, but the price increases at "peak" hours of the day, which typically correspond with times of high aggregate consumption across the utility's service area. Some of PG&E's TOU plans eliminate the additional "penalty" for high monthly usage, but others do not.

Starting in mid-2021, PG&E begins to transition nearly all of its customers to TOU pricing from the original block price at the direction of the state. Based on contemporaneous documentation, PG&E had anticipated switching consumers beginning in October of 2020, but this was likely delayed due to Covid [4]. While a small number of consumers do transition to the new plans before the end of 2020, in my sample, the first "wave" of consumers do not appear to change plans until April of 2021 [5]. Rollout is done on a county-by-county basis, with the last group switching in April of 2022. Crucially, customers had the option to opt out of the transition and either stay on their original rate or choose one of the alternative TOU rates. Consumers were given the option to notify PG&E up until the month of their county's transition, and could do so online, where they also had the ability to compare rate plans for their usage history and decide what their optimal rate plan should be. According to available documents, customers received emails up to four months prior to the transition notifying them of the change.

Additionally, the company provided "risk-free bill protection" that would reimburse customers on the TOU plan for any additional cost over the block pricing plan for their first 12 months on the TOU plan. Whether this involved monthly credits during each billing cycle, one lump-sum reimbursement at the end of the 12-month period, or some alternative method is unclear. In the former possibility, consumers that were automatically transitioned to the TOU plan are likely to not have changed their consumption behavior during the bill protection period, as there would be little incentive to do so based on their previous bills. In the latter, we would likely see consumers adjust their consumption behavior during the bill protection period, and then keep their behavior constant in the period afterwards. I test for both of these possibilities in the next section.

Some consumers are on alternative rate plans that they opt into, such as specialized rates for electric vehicles or other TOU plans, but the vast majority are not. The new default, "TOU-C", features peak pricing periods from 4PM to 9PM every day of the year, including holidays and weekends. As with block pricing, consumers also face increased tariff rates when they exceed their climate zone's monthly allotment; the penalty is uniform for both the peak and off-peak, so both rates are raised by the same amount. Additionally, there are now rate seasons; peak periods are about 5% higher than off-peak periods during the winter (October-May) but about 15% higher than off-

[3] Note that "smart" meters here pertain to grid-connected devices that allow for the utility company to constantly check real-time consumption, and not "smart" thermostats that allow for consumers to manage their in-home temperature, though the latter can connect to the former.

[4] See this link for more information

[5] See this link for more information.

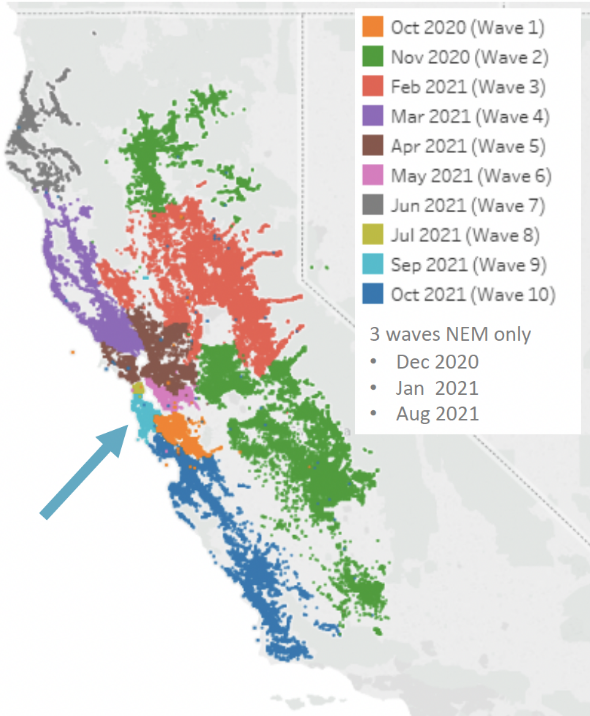

Figure 2: Initial draft of the rollout for PG&E's TOU-C transition program, per their official documentation. Listed dates are for the original plan and not the actual dates.

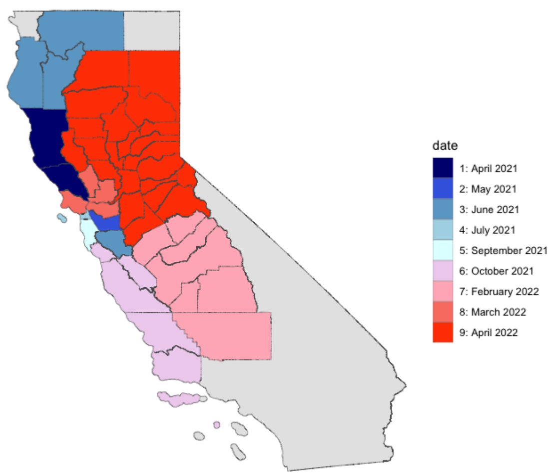

Figure 3: Actual transition dates by county as seen in the data.

peak in the summer (June-September). This is presumably to account for anticipated increases in electricity consumption from air conditioning; more than 90% of the households in my sample have gas heating and thus do not use as much electricity in the winter.

As alluded to above, consumers in California are divided into "climate zones" that dictate their

Wave	Old Wave Number	Area	Wave Date	Households
1	4	North Coast	Apr-21	1,792
2	6	Oakland	May-21	3,277
3	7	Far North Coast	Jun-21	4,251
4	8	San Francisco	Jul-21	1,340
5	9	San Mateo	Sep-21	2,256
6	10	Southern Coast	Oct-21	3,338
7	3	North Central	Feb-22	2,315
8	5	Sonoma Valley	Mar-22	5,231
9	2	Central Valley	Apr-22	4,914

Table 1: Listed dates based on observed transition month in the sample.

allotment of electricity at the lowest price tier each month. Since PG&E's service area has a wide variation in climate, customers in areas that face similar weather conditions are grouped, and the region's "allowance" of consumption at the lowest price is calibrated to match the daily usage of households between the 50th and 60th percentiles [6]. This allotment is then multiplied by the number of days in the billing cycle to obtain the total billing cycle allotment; for example, a daily baseline allowance of 20 kWh for a 30-day billing cycle yields 600 kWh. This applies to the cumulative consumption during the billing cycle and does not reset at the beginning of the day. Baseline allowances change in accordance with the winter/summer cycles above, and consumers on both the block and TOU plans face the same allowance within their respective climate zones. Consumers with gas heating have much lower allowances in the winter than consumers with electric heating. The variation across zones can be substantial, and consumers that are quite close geographically can face drastically different climates. For example, the amount of electricity that Boonville consumers in California's mountainous north can use without incurring an increased cost in the summer of 2021 is 10.3 kWh for living in zone X, while those in the beach town of Manchester an hour away can only use 6.8 kWh as residents of zone T. In the zone W city of Bakersfield to the far south, where summertime temperatures regularly exceed 100 degrees Fahrenheit, the baseline allotment during this same time period was 20.2 kWh. Figure 4 shows a map of these zones in 1990, and Figure 5 shows a comparison of how baselines change across seasons. The zone borders and allowances are updated infrequently. Borders were adjusted in 2020 but largely stayed the same; allowances are updated every three years but also do not tend to change much in magnitude.

3.3 Prices

Prices during the sample period are relatively stable for the lower tiers of block pricing, but vary widely for both the third block pricing tier and both of the TOU price periods. Figure 6 shows

[6] See this link, under "Allowances are determined as follows", for more information.

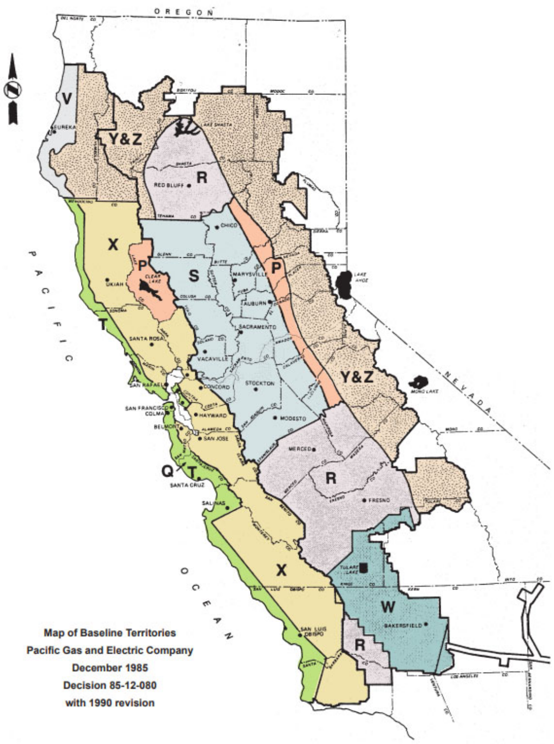

Figure 4: Climate zones for PG&E's service area starting in 1990.

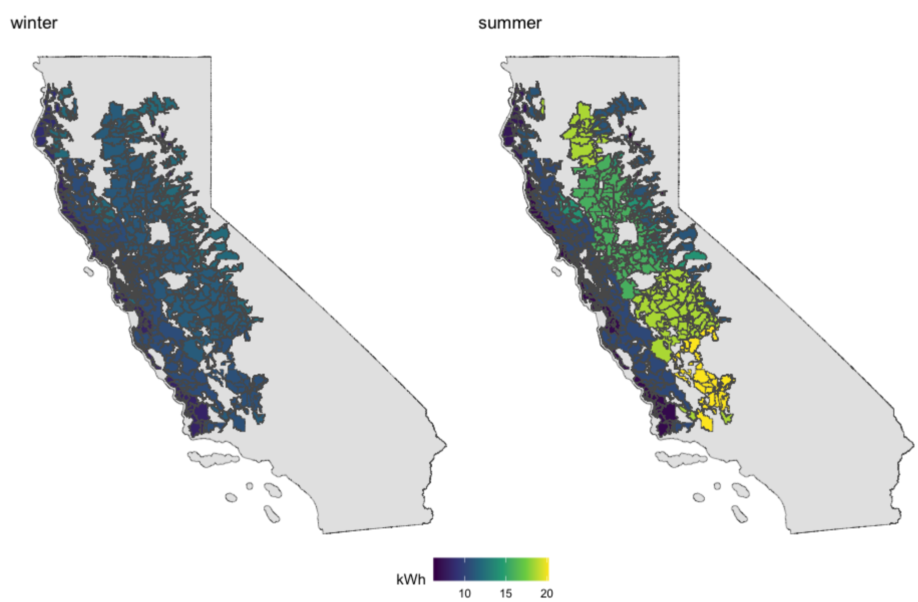

Figure 5: Baselines by season. Scale indicates the estimated average daily consumption for the climate zone.

a comparison between the block price and TOU price over the entirety of my sample. The most severe change in price occurs for the highest block price tier in June of 2020; this price cut was ordered by the governor to combat concerns that Californians' electricity consumption would skyrocket because they were staying inside during the Covid pandemic [7]. However, I find that in practice, the percentage of households hitting this price tier did not change significantly. The "jumps" in TOU prices reflect price increases during the summer, with rates within the TOU price tiers being significantly higher than their respective tier of block pricing. During the winter periods, however, TOU rate tiers are actually lower than the block tier.

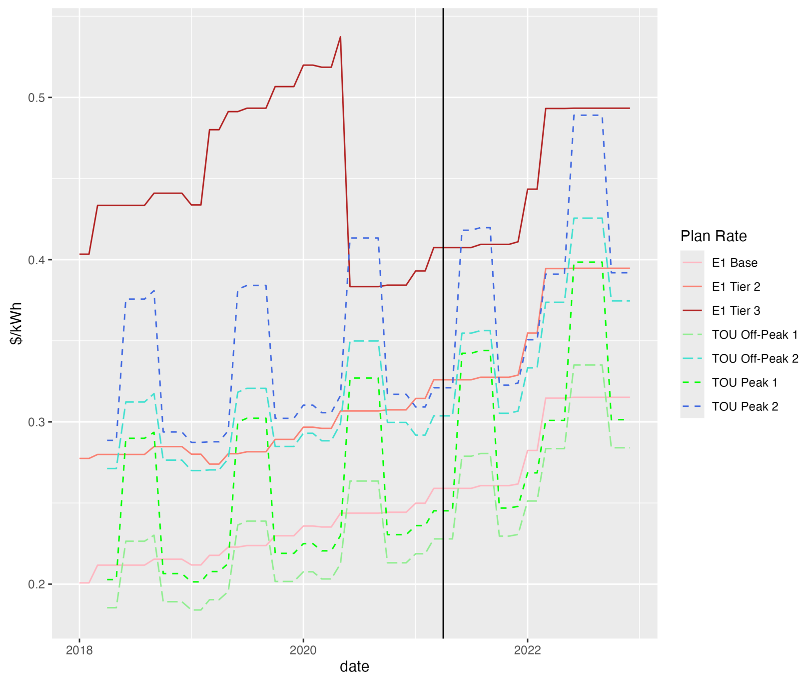

Figure 6: Pricing for TOU and block (E1) plans over time. The black line indicates the first month of the TOU transition rollout. "Peak 1" and "Off-Peak 1" refer to the first tier of the TOU plan, while "Peak 2" and "Off-Peak 2" refer to the second tier.

Figures 7 and 8 plot examples of comparisons between the block pricing and new default TOU plan in April and June of 2021 to illustrate how prices changed for consumers. April is the first month of the transition program and June is the third. Prices did not necessarily increase for consumers due to seasonal differentiation in addition to the aforementioned time-of-day differentiation. In the winter months of 2021, prices are better for all consumers relative to the block plan at all tiers of usage, with the most savings for consumers at the high end of consumption. However, during the summer, prices are strictly worse for consumers in the first two tiers; consumers that breach the third usage tier only face higher prices for their peak period consumption.

Across the sample, TOU-C peak prices are an average of 8% higher than the Tier 1 block prices, and off-peak prices are an average of 5% cheaper than the Tier 1 block prices. When compared to the Tier 2 block prices, TOU-C's peak is 12% higher on average, and the off-peak is only 1% more expensive. When accounting for the differentiation in seasons, the summer peak prices are

[7] See Advice Letter 5831E at this link.

an average of 33% higher than Tier 1 and 31% higher than Tier 2, while off-peak prices are 7% and 11% higher, respectively. In winter, peak prices are an average of 5% cheaper in Tier 1 and 2% more expensive in Tier 2. Off-peak prices are an average of 12% cheaper for Tier 1 in winter, and 4% cheaper than Tier 2.

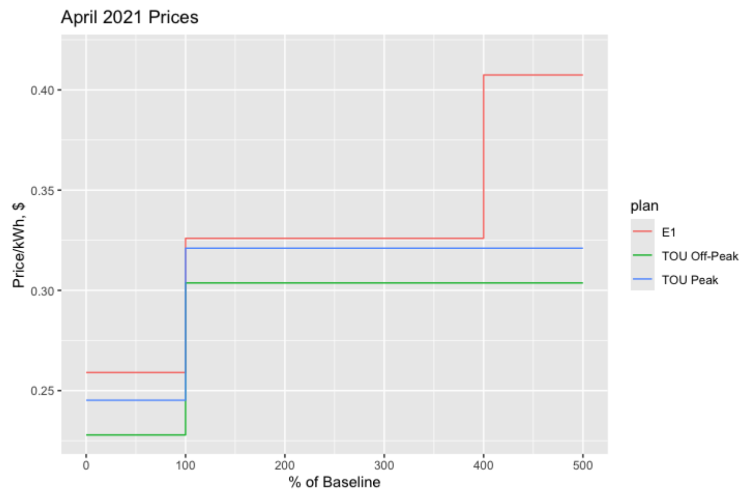

Figure 7: Price comparison, first month of transition (April 2021).

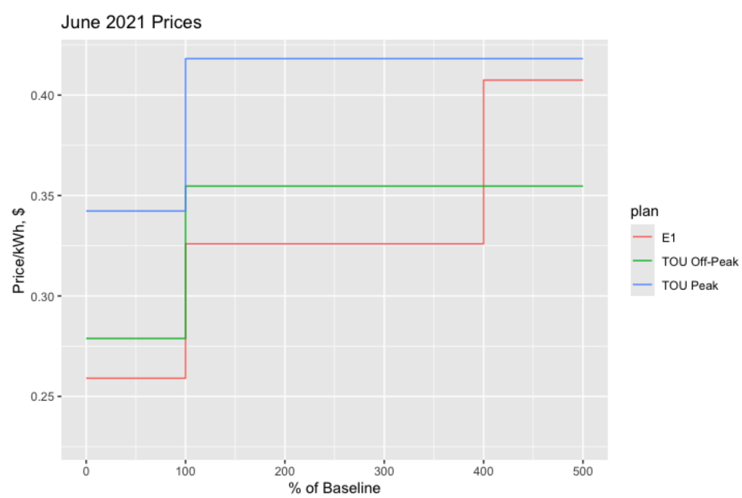

Figure 8: Price comparison, third month of transition (June 2021).

3.4 Constructing the Dataset

From the raw dataset, I remove customers that are ever on an electric vehicle plan, subsidy program (CARE and FERA), or have solar. I also remove any customer that ever records an entire month with less than 15 kWh or more than 3600 kWh. Finally, I remove households that switch to TOU-C earlier or later than their county's wave. This leaves 28,714 households for the analysis from the original sample of 75,000.

Select summary statistics are presented in table 2, split by pre- and post-rollout and whether the household transitioned to the new program or not. In comparing the two, it appears that in the pre-transition period, staying households use approximately the same amount of electricity as switching households, with no decisive trend between waves. However, in the post-transition period, switching households used more electricity on average in every wave. In comparing the two, households that opt out of the program (the "stayers") use slightly more electricity on average, and have higher bills as well. They also use more electricity on the peak period. These trends continue in the post-period. It is also worth noting that both demand and bill amounts declined across the board. This may be due to the fact that most restrictions from the Covid pandemic had been removed by the start of the transition, or that prices rose over the last two years of the sample, or a combination of both.

The opt-out rates for seven of the nine waves are well below half. It is worth noting that the two outlier waves were delayed from the revised transition schedule that I have found available online [8] and thus it is plausible that customers had more time to opt out. Nonetheless, most consumers were transitioned to the new rate plans. Some consumers were exempt from the transition and did not need to opt out: those with subsidized plans, electric vehicle plans, have a "medical" baseline to accommodate electricity-heavy medical equipment, or are already on a different TOU plan. However, given that I have removed these customers save for the medical baseline [9], they should not impact the estimated switching rate. PG&E's "FAQ" for the TOU program also mentions that customers in "hot climate zones" were exempt from the transition; however, I do not find evidence of this in the data. See 9 for switching rates by wave. While customers in the hottest climate zones—P, R, and W, which are mostly consumers in waves 7 and 9—do have lower switch rates than the more temperate regions, it does not appear that customers in any specific climate zone were completely exempt from the rate transition program, nor were there specific zip codes that were opted out[10]. Thus, customers being automatically exempt based on geography does not seem to be a plausible explanation for why the switch rates in the latter waves are conspicuously low. As such, I currently believe that the reason is a combination of consumers in these waves being more attentive to their utility billing, and the aforementioned delay in rate transition.

4 Selection Out of the Program

Given the known parameters of the rate transition program, the main threat to achieving the goal of decreasing peak-load consumption appears to be advantageous selection out of the program by households that stand to gain less from the transition. To test for this, we can use the observable

[8] See this press release for more information.

[9] I do not have an indicator in the data for consumers on a medical baseline; however, these customers are likely filtered out due to very high electricity usage or high rates of error when re-constructing their bills from the hourly level. Medical baseline eligibility requires that the customer use specific equipment such as an electric wheelchair or life support, which also likely represents a small portion of the sample population.

[10] By checking the transition rate in each zip code, we can see if there were specific zips that were exempted from the transition, accounting for zips with very few consumers. In my sample, no zip code with more than 35 customers has a switch rate of 0%, implying that no zip was exempt from transition

	Switch			Stay			Mean Delta
	Mean	Median	S.D.	Mean	Median	S.D.	
Pre-Rollout							
kWh	522.91	440	346.38	567.12	478	368.22	-44.21
$	139.65	112.06	107.17	145.92	118.75	106.18	-6.27
Peak kWh	146.26	118.32	109.27	165.92	132.34	123.53	-19.66
Post-Rollout							
kWh	476.99	407	304.36	489.77	428	292.93	-12.78
$	168.64	132.13	133.58	177.58	142.68	133.13	-8.94
Peak kWh	119.21	98.67	88.46	131.3	110.41	94.78	-12.09

Table 2: Summary statistics split by pre- and post-rollout.

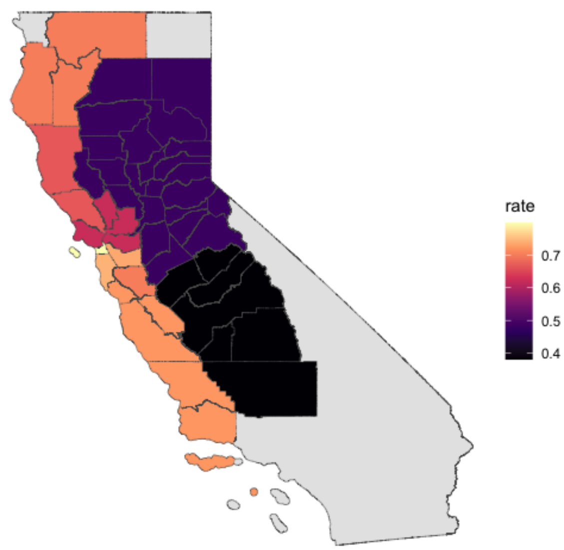

Figure 9: Percentage of households that switched by December 2022, by wave. The final three waves are in warmer regions, and waves 7 and 9 had additional time to opt-out due to a delay in their transition dates.

characteristics of each household's usage and attempt to identify which ones most closely predict opting out. Throughout this section, I use the following probit regression specification to evaluate how the observables impact the consumer's probability of deciding to switch to the TOU plan:

$$Y_{ijk} = X\beta + \gamma_j + \gamma_k + \epsilon_{ijk}$$

where Y_{ij} takes value 1 if household i in climate zone j and wave k is automatically opted in to the TOU program, and 0 if they opt out, and X are the covariates explored in the following set of regressions.

4.1 Selection on Consumption Uncertainty

If households are attentive to their power usage and concerned about whether it will negatively impact their bill in the future, they may consider changing their consumption patterns to re-optimize their usage. They could do this by, for example, delaying their laundry or dish-washing to be later in the evening, after the peak hour period. However, this may not be possible if there is a high hassle cost, especially if the household is less able to predict their energy needs on a given day.

Household "predictability" in energy needs may differ due to work schedule, weather, or other factors that engender more variable electricity usage. To proxy for this, I calculate a set of household energy volatility measures that focus on variation across different periods for each household. Specifically, I calculate the mean and standard deviation of the household's peak and off-peak consumption for both weekdays and weekends. Table 3 shows probit results using these constructed measures from 2019 and 2020, respectively. Consumers' mean and standard deviation of consumption in these variables does not change much from 2018 to 2019, though the distribution in 2020 shows much lower concentration at the mean than in the prior two years. As such, I decided to run these probit models by separating 2019 and 2020, since consumers may be concerned about their consumption after a return to pre-pandemic trends in peak usage. I find that none of the 2020 variables are statistically significantly correlated with the decision to switch or stay on the TOU plan. In the 2019-only model, a higher weekday peak standard deviation, which indicates more volatility in peak-hour consumption during the work week, is both negative and statistically significant. This implies that consumers with higher weekday volatility were more likely to opt out of the switch.

As a second set of variables, I also consider the possibility that work-from-home changes caused by the Covid pandemic factored into household decisions to opt out. Even though the standard "workday hours" from 9AM to 5PM only overlap with the peak period by an hour, customers may be concerned that sustained energy usage during the day will exacerbate the added costs to their bill, particularly during the summer period. While California's stay-at-home order expired in January 2021, I observe that household consumption during workday hours stayed elevated until the spring of 2021. This appears to be the case because some schools in California had partial or total online learning until the fall of 2021[11].

In running the regression, I find that only pre-Covid work-hour consumption and off-hour consumption during Covid are significantly correlated with switching to the TOU plan. The sign on work-hour consumption is negative while the sign on off-hour consumption is positive. Though both are very small in magnitude, the mean of the former is 130 kWh and the mean of the latter is 265 kWh, so these two variables take values approximately as large as the fixed effects included in the regression. Some consumers may not have been concerned about at-home consumption during work hours while the pandemic was in full force, but if they anticipated continuing to stay at home during the day once schools and offices reopened, then they may have been wary of increasing

[11] According to the LA Times, approximately 90% of the state's primary and secondary students had an in-person or hybrid schedule by April 2021. See this link for more information.

	P(Switch to TOU)	
(kWh)	**Using 2019**	**Using 2020**
Weekday, Peak Mean	-0.014	-0.050
	(0.042)	(0.049)
Weekday, Peak S.D.	-0.127	-0.016
	(0.036)	(0.042)
Weekend, Peak Mean	-0.005	0.064
	(0.042)	(0.049)
Weekend, Peak S.D.	0.093	-0.052
	(0.036)	(0.041)
Wave FE	Yes	Yes
CZ FE	Yes	Yes
Clustering	Zip	Zip
N	25,258	25,258

Table 3: Probit results for select household volatility measurements. Variables are standardized to mean 0 and standard deviation 1. Models run use data from only one year in each column: 2019 in the left, and 2020 in the right.

their bills.

4.2 Selection on Bill-Shock Sensitivity

If households are instead more responsive to changes in their bill cost than uncertainty in their usage, they may be concerned about the implications of changing their plan for their future bills. While most people likely pay their utility bills each month without closely auditing their past usage[12], it is possible that the households opting out of the transition are doing so out of concern for its impact on their month-to-month rates. Like many household utilities, electricity is consumed during a billing cycle without knowing the cumulative monthly cost, and then a bill is sent at the end of the month. If a household receives a bill that is larger than their expected cost for the previous month's consumption, they may respond by decreasing their consumption for the next billing cycle. Households that respond more strongly to shocks may be more attentive to the transition program, and thus choose to opt out. To test this, I regress the difference between the previous two bills on each household's current consumption:

$$\log(kWh_{ijt}) = \beta_0 + \beta_1(bill_{ij,t-1} - bill_{ij,t-2}) + \gamma_t + \gamma_j + \gamma_i + \epsilon_{ijt}$$

where the γ are fixed effects are for the calendar month, climate zone, and household, respectively.

[12]The author included.

	P(Switch to TOU)
During Covid	
Work Hour kWh	0.035
	(0.032)
Off-Work Hour kWh	0.126
	(0.053)
Weekend kWh	-0.097
	(0.063)
Pre-Covid	
Work Hour kWh	-0.067
	(0.028)
Off-Work Hour kWh	0.012
	(0.052)
Weekend kWh	0.009
	(0.058)
Wave FE	Yes
CZ FE	Yes
N	25,258

Table 4: Probit regression of switching plans on work-hour consumption. "Work Hour kWh" counts average monthly consumption only during 9am to 5pm on week days; "off-work hour kWh" is consumption during all other hours on weekdays. "Weekend kWh" is all consumption on Saturdays and Sundays. Variables are standardized to mean 0 and standard deviation 1.

The hypothesis is that if households with higher average differences between bills respond more strongly by decreasing their current-month consumption, then they may be more attentive to the TOU program and more likely to opt out.

The initial regression, shown in Table 5, appears to show a statistically significant but small positive, rather than negative, relationship between monthly bill deviations and consumption. This indicates that consumers who have larger month-to-month differences in their bill costs have slightly higher consumption in the following month. Using the mean of these bill differences for each household, I then perform a probit regression to predict each household's likelihood of opting out of the program, and find that households with larger average month-to-month differences in cost are slightly more likely to have switched to the TOU rate. Given the small magnitude of the coefficient, this likely indicates that households do not use the magnitude of monthly differences in their bill to inform their decision to opt in or out. An alternative regression specification that used detrended monthly bills found a largely similar result, indicating that this is not driven by seasonality in consumer billing.

	P(Switch to TOU)		
	Lag Difference	**Bill Deviation**	**Bill Deviation Pct**
	0.058	-0.050	-0.013
	(0.009)	(0.009)	(0.008)
Wave FE	Yes	Yes	Yes
CZ FE	Yes	Yes	Yes
Cluster	Zip	Zip	Zip
N	25,258	25,258	25,258

Table 5: Regression results for bill shocks. "Lag difference" is the difference in cost between the previous two months' bills. "Bill deviation" is the previous month's bill as a deviation from the consumer's average calendar month. All variables are averaged over the pre-rollout period and standardized to mean 0 and standard deviation 1.

4.3 Selection on Pre-Rollout Costs

Per PG&E's documents about notifying residents of the TOU transition, residents would be shown a comparison of their previous consumption under the "E1" block pricing plan to what the bill would have cost under the new TOU plan when choosing whether to opt out of the transition, though how many months is not known. Consumers may thus decide whether to opt out depending on their expected net value from their previous bills. This may be derived from one of several similar metrics, such as the consumer's total net cost, the number of total negative-cost months that they observe, or simply their worst-case observed cost.

	Mean	**S.D.**	**Median**	**25th**	**75th**
April Net Cost ($)	7.75	11.30	7.48	6.02	8.38
May	-7.99	22.55	-4.60	-18.37	2.93
June	-23.47	22.11	-19.32	-34.62	-10.27
July	-26.20	25.10	-21.39	-39.58	-10.93
August	-24.40	21.94	-20.92	-35.93	-11.22
September	-33.98	23.74	-29.87	-47.94	-16.05
October	7.41	9.60	7.21	5.87	8.24

Table 6: Net cost summary statistics by month, averaged over 2018 through 2020. Defined as the true cost under E1 block plan minus the TOU-C plan in the same month. Units are in dollars.

If consumers decide whether to opt out based on this information, then we should find a relationship between more negative estimated costs in the pre-transition period and the probability of opting out from the transition. To do this, I estimate "net cost" as the difference between the consumer's true bill cost under the block pricing plan and what they would have paid in that month under TOU pricing. Given that TOU pricing varies by season, I then take the mean of these costs by calendar month over 2018 through 2020 and estimate a probit of these net costs on switching. As such, when this variable takes a negative value, it means that TOU pricing cost more in that month than the consumer's original plan. If the variable is positive, then the consumer would have saved money by being on TOU pricing in that month. As I cannot know which specific months a consumer was able to see when making their opt-out decision, I take the average over winter and summer months, and then use this in a probit, seen in Table 7. The summer coefficient is much larger in magnitude than in winter, indicating that its impact on opt-out decision is larger. While the coefficient for the summer months is positive, the raw data in this period takes overwhelmingly negative values, and thus the positive coefficient indicates that the likelihood of switching decreases, rather than increases, as the variable gets larger in magnitude. The regression thus indicates that consumers whose expected value from the TOU plan was worse were much more likely to have opted out of the transition program. A rational consumer who selects on their expected net value across months may take into account the savings that they experience during winter months compared to their losses in summer months, but the significance and sign on these coefficients indicates that consumers are probably biased by what they observe specifically in the summer months.

	P(Switch to TOU)
Summer Avg Net	0.182
	(0.009)
Winter Avg Net	-0.070
	(0.009)
Wave FE	Yes
CZ FE	Yes
Cluster	Zip
N	25,258

Table 7: Regression for "net cost", defined as the block plan's cost minus the TOU plan's cost in that same month, holding the consumer's usage fixed. Variables are averaged over the summer and winter season, and then standardized to mean 0 and standard deviation 1.

As a final test, I use the statistically significant variables from each of the previous sections in a final probit to test whether any variables stand out relative to the others. In addition, I have added a set of average county temperature variables by calendar month, and interacted the wave fixed effects with summer net cost in order to incorporate variation in billing costs that are inherent to the weather in each rollout wave. I also added the household's average total kWh per month as a proxy for household size and general demand. Lastly, I added an indicator for heating type, which is significant in this regression, though the overwhelming majority of consumers are on gas and not electric heat. The negative coefficient implies that households with gas heat are substantially more likely to opt out than electric-heating households, which is a surprising result given that electric heat portends a larger electric bill in the winter-price months.

The only coefficients that are not statistically significant are for average total kWh, and weekend peak mean consumption, which is significant at the 10% level instead of 5%. The previously mentioned household volatility measure for weekend peak consumption also flips its sign, so now an increase in that measure increases the likelihood that the household switches plans. The coefficient on consumption during work hours before Covid also changes sign.

	P(Switch to TOU)
Net Cost, Summer	0.535
	(0.062)
Net Cost, Winter	-0.198
	(0.019)
Total kWh	0.041
	(0.082)
Bill Deviation	-0.059
	(0.011)
Weekday Peak Mean kWh, 2019	-0.284
	(0.044)
Weekend Peak Mean, 2019	0.079
	(0.041)
Work-Hour kWh, Pre-Covid	0.064
	(0.031)
Off-Work kWh, Pre-Covid	0.334
	(0.053)
Gas Heating	-0.086
	(0.030)
Wave FE	Yes
CZ FE	Yes
Cluster	Zip
Average Monthly Temperature	Yes
Wave FE x Net Cost	Yes
N	23,475

Table 8: Probit using previous results, combined into one regression. ”Average monthly temperature” is a set of variables for average temperature by month, and “Wave FE x Net Cost” is an interaction between “summer net cost” and wave fixed effects. All non-indicator variables standardized to mean 0 and standard deviation 1.

In using these results to predict a household’s TOU enrollment decision, I estimate the model in Table 8 separately for each wave, and then predict the household’s likelihood. To evaluate the results, I plot the estimated probability density for both switching and non-switching households. The results for waves 1 and 5 are available in Figure 10. The degree of separation between the densities indicates that I am accurately assessing whether households in either group are actually likely to switch or not. It currently appears that I routinely estimate many household to be likely

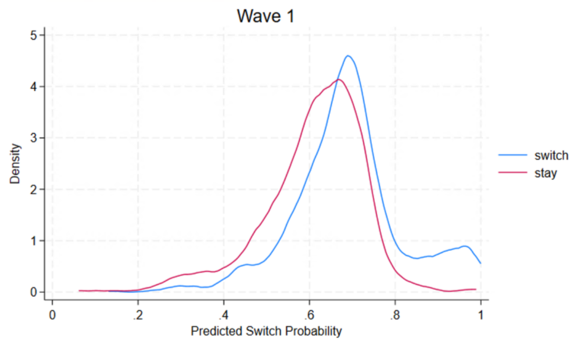

(a) Wave 1 switching probability density.

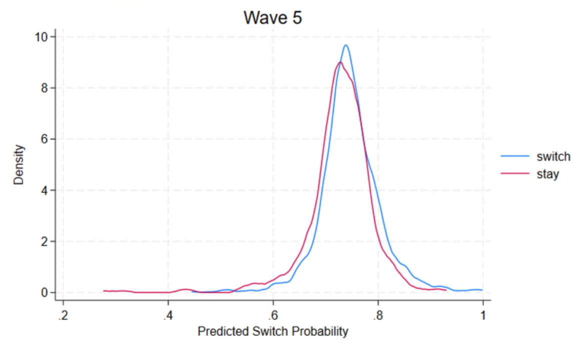

(b) Wave 5 switching probability density.

Figure 10: Comparison of estimated switching probabilities for two waves.

to switch, even though they do not do so; this may indicate that there is an unobservable factor is highly relevant to some households' decisions to opt out of the program. Using a simple rule-of-thumb that a predicted probability less than 0.5 indicates likely opt-out, and greater than or equal to 0.5 indicates a likely switch to TOU, I show the percentage of "correct" predictions in Table 9.

Wave	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**
% Correct Prediction	67.1%	73.4%	70.3%	79.9%	74.6%	72.3%	65.1%	63.5%	60.0%

Table 9: Estimated percent of correct predictions by wave, using the probit model.

For four of the waves, I have an estimated accuracy below 70%, while for the other five, I estimate accuracy between 70 and 80%. This is a surprising result, given that I have estimated warm months to be strongly correlated with deciding to stay on block pricing, and the accuracy is still lower than in the other regions.

5 Treatment Effects

For the households that do not opt out of the program and are transitioned to TOU pricing, it is vital to evaluate their post-transition behavior to understand how the new rate plan impacted their consumption patterns. Households may have responded in several ways. The seasonal pricing may induce households to increase their overall consumption in the winter relative to the summer, since both peak and off-peak periods in winter are less expensive than in the summer. Depending on their schedule flexibility, households may not respond to specific inter-day prices and merely decrease or increase their consumption across all periods, rather than forming new habits to offload some of their consumption to the peak period.

5.1 Two-Way Fixed Effects

A basic way to evaluate the impact of the program on the households that have switched is to use two-way fixed effects. Though selection out of the program is nonrandom—and in fact appears to be driven by consumers making decisions based on their cost expectations under the program—I anticipate that the households that do not switch should not change their consumption patterns after the program rollout. These households retain the same block pricing plan that they have been on since the beginning of the sample, and thus are unaffected by the program, so they should serve as valid controls for the "treated" households that do change their plan.

The two-way fixed effect model that I use will have the form

$$Y_{ijkt} = \beta_0 + \beta_2 post_{jt} + \beta_3 switched_{ik} \times post_{jt} + \gamma_i + \gamma_j + \gamma_t + \gamma_k + \epsilon_{ijkt}$$

where Y_{ijkt} is the outcome variable of choice for household i living in climate zone k that becomes subject to the program as part of wave j, with time period t. The typical indicator for treatment in this specification has been removed due to collinearity with the household fixed effects.

I first run this model for the log of total consumption and bill price, seen in table 10. I omit 2020 and 2021 from the regression due to the outsized impact of Covid on consumption and that the rollout period does not begin for the latter waves until 2022, respectively. While this specification estimates that both bill cost and consumption declined by around 1%, this assumes that summer and winter had the same treatment effect on the outcome variable. In Figure 11, I explore the differential treatment effect by month to observe whether the effects with change significantly between months on winter and summer pricing. In addition to log consumption and bill cost, I also report the change in percentage of consumption on peak hours and total peak-hour consumption. The latter two variables do not have data from April through March due to the progam rollout not beginning until April 2021. The overall impact of the treatment on each month appears small, with the exception of August and September 2022, which seem to have been outliers in terms of billing, and persist even when adding county-level weather effects. Overall, households switching to the TOU rate plan do not seem to have responded in a significant or consistent manner to the plan. However, in the next section, I explore the importance of separating the regressions by both season and the "bill protection" period.

5.2 Risk-free Bill Protection And Seasons

From the previous section, I have established that the treatment effect of the TOU transition does not appear to have accomplished its goals of decreasing peak load among households that switched. In the data section, I had mentioned that PG&E promised temporary "bill protection" to switching households by offering to automatically credit their account the difference in their bill between the

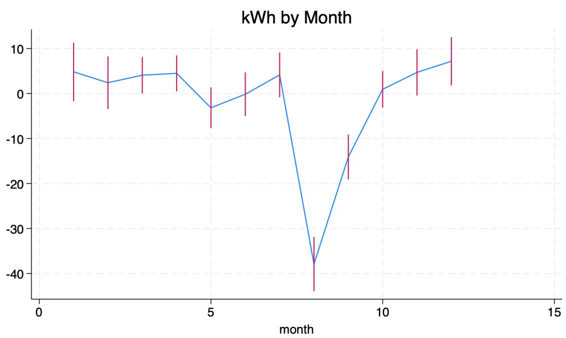

(a) log kWh.

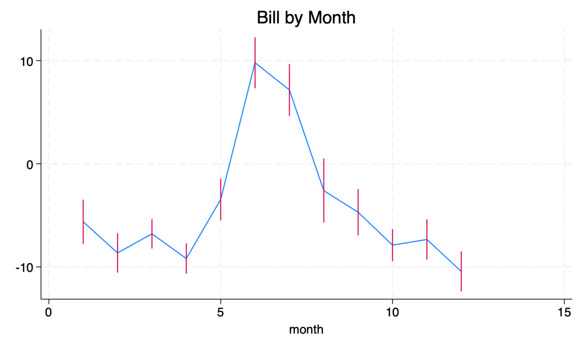

(b) log bill cost.

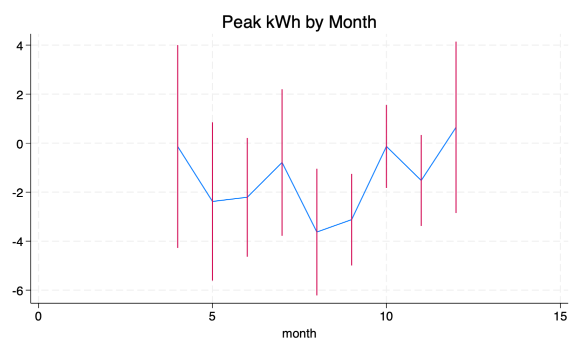

(c) Peak kWh.

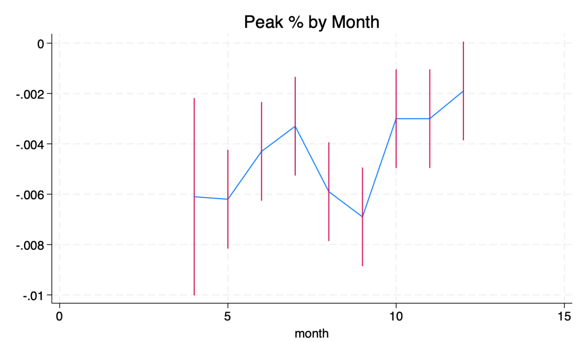

(d) Peak percentage.

Figure 11: Two-way FE models by month.

	log(kWh)	log(bill $)
Post x Switch	-0.010	-0.018
	(0.003)	0.004
Post	0.004	0.332
	(0.003)	(0.004)
HH FE	Yes	Yes
Wave FE	Yes	Yes
CZ FE	Yes	Yes
Cluster	Zip	Zip
N	868560	868560
R2	0.752	0.748
Mean of Dependent	6.085	4.762

Table 10: Two-way fixed effects estimation of logged consumption and bill cost. Omits 2021 and 2020 from estimation.

block and TOU plan if their bill would be higher during that month than on the original plan, for their first 12 months on the plan. It is possible, then, that attentive customers on the TOU plan would continue to use electricity without altering their habits, undercutting the program's goals. Additionally, the impact of the treatment appears to be counter to economic theory—consumption decreases *and* bills decrease. Since we know from the prices laid out in Figures 7 and 8 that prices are mechanistically lower in winter and higher in summer on TOU, it stands to reason that we should anticipate opposing signs on the treatment effect estimator for billing costs, depending on the season. Both of these factors should be considered.

While I do not have more than 12 months of post-transition data for every wave, the first five waves have between 15 and 20 months of data for us to see if consumers began to adjust their usage after the bill protection period expired. To do this, I split the post-transition data into "short difference" and "long difference" periods, with the former comprising the first 12 months of the rollout, and the latter the periods afterwards. To account for seasonal variation, I also split the regressions into summer and winter periods. The results are shown in Tables 11 and 12.

Contrary to the previous section, it is more apparent that the program did not achieve its goals in the first 12 months post-rollout, and there is weak evidence that they did so after the bill protection program expired. I estimate that while bill costs increased by around 7 to 10% in summer, they decreased by around 5% in winter. However, consumption did not follow; only in winter after the bill protection program do I estimate a statistically significant change in consumption.

	log(kWh)			
	Summer, Short	**Summer, Long**	**Winter, Short**	**Winter, Long**
Post x Switch	0.001	-0.001	0.007	0.016
	(0.005)	(0.006)	(0.005)	(0.006)
Post	0.005	-0.004	-0.100	0.034
	(0.006)	(0.006)	(0.006)	(0.006)

Table 11: Two-way fixed effects estimates of log consumption. Uses only waves 1 through 5. "Short" difference is the wave's first 12 months on the plan. "Long" difference is the period afterwards.

	log(Bill $)			
	Summer, Short	**Summer, Long**	**Winter, Short**	**Winter, Long**
Post x Switch	0.109	0.072	-0.051	-0.045
	(0.006)	(0.007)	(0.006)	(0.007)
Post	0.227	0.375	0.247	4.555
	(0.007)	(0.007)	(0.006)	(0.001)

Table 12: Two-way fixed effects estimates of log bill cost. Uses only waves 1 through 5. "Short" difference is the wave's first 12 months on the plan. "Long" difference is the period afterwards.

5.3 Matching Estimator

As a final method of estimation, I use a matching estimator approach in order to accommodate the compositional differences across consumers. In matching, I can estimate a propensity score using the previous probit exercise so that consumers with similar features are matched to each other before estimating the average treatment effect. Rather than pooling across consumers with vastly different observables, as in the two-way fixed effects method, matching can estimate the average treatment effect for households relative to others that share their characteristics. I implement this using a "matching difference-in-differences" approach, whereby I first take each household's average difference in their pre- and post-rollout consumption and bills. The average treatment effect (ATE) of a household that switches from block to TOU pricing can be estimated as

$$E[Y_{it}(1) - Y_{it}(0)] = \Delta Y_{it}$$

where $Y_{it}(1)$ is the outcome of interest for unit i in time t when treated, and $Y_{it}0)$ the outcome when not treated. Assuming that the difference between treated and untreated units would have been constant absent the program, the average treatment effect on the treated (ATT) can be estimated as

	kWh				Billing ($)			
	Summer		**Winter**		**Summer**		**Winter**	
	Short	**Long**	**Short**	**Long**	**Short**	**Long**	**Short**	**Long**
ATE	-2.093	-7.440	-7.823	-2.726	-1.929	14.813	-13.262	-9.141
	(2.794)	(3.051)	(2.489)	(3.172)	(1.018)	(1.273)	(0.931)	(1.284)
Neighbors	1	1	1	1	1	1	1	1
N	6,124	8,504	10,308	10,308	6,124	8,504	10,308	10,308
Mean of Y	-9.761	8.088	-40.265	34.069	9.055	64.394	6.146	55.107

Table 13: Propensity score matching estimation for households, using only the first five waves.

$$E[\Delta Y_{it}(1) - \Delta Y_{it}(0)]$$

While traditional difference-in-differences imposes an assumption of random treatment, in this estimation approach we exploit the selection of households into—or out of—the program by their observable characteristics to "match" them so that we estimate the treatment effect under the assumption of conditional independence. That is, for two households whose characteristics are highly similar, their outcomes without treatment are equivalent:

$$E[Y_{1,0}(1)|X] = E[Y_{2,0}(0)|X]$$

In this setting, I match on the consumers' observable characteristics prior to rollout, which includes their consumption habits and the geographical components used in the previous discussion of selection.

The results are shown in Table 13. While I tested a number of different "nearest neighbors" in the implementation, there does not appear to be much variation across a wide array of neighbors in either the estimated coefficients or standard errors, and thus I present merely the single-matched case. Per the results, we see that households now appear to have responded to the program by slightly decreasing their consumption in the summer months, and having no significant change in winter months. These changes are most notable in the post-bill protection period. As a percentage of consumption in the pre-period, consumers switching to TOU pricing decrease their summer consumption by approximately 1.5% after 12 months. While winter consumption is estimated to decrease as well by about 0.6%, this estimate is also not statistically significant.

6 Discussion

Having estimated the treatment effect of customers being switched to the TOU rate structure, I consider the implications of switching all residential customers onto this plan and disallowing opt-out. If customers that had previously opted out are doing so to avoid anticipated cost increases due to their usage, they may also be more elastic in their usage around the peak period.

Prior to the program and Covid, the top marginal price faced by consumers in my sample in 2019 was \$0.27/kWh. Recall that prices were not differentiated by season at this time. In the post-rollout period, prices are differentiated by both season and rate plan. Consumers that stayed on block pricing faced average marginal prices of \$0.37 in the winter months of 2022, and \$0.38 in the summer, while switching consumers faced prices of \$0.369 and \$0.44, respectively. We can use the treatment effects estimated by matching in conjunction with these prices to estimate separate elasticities for both winter and summer TOU pricing. Elasticity of demand in the summer is approximately -0.011, while winter elasticity is approximately -0.007. Various experiments in dynamic pricing with electricity cataloged by Harding and Sexton (2017) have produced estimated elasticities ranging from -0.06 to -0.20. The aforementioned contemporaneous TOU rollout in Enrich et al. (2024) estimated that elasticities were between -0.09 and -0.11. If opt-out were instead disallowed in the current setting, it is possible that demand response would be more drastic, though the average elasticity estimate would likely still be substantially lower than the other literature. At present, it is difficult to assess whether this indicates the literature over-estimates elasticity under dynamic pricing, or whether the extended rollout of this program combined with the bill protection plan hindered the accomplishment of its goal of reducing aggregate consumption in the first two years.

7 Conclusion

While TOU pricing has long been known in the literature to be effective at internalizing the differences in consumer preferences across time periods, large-scale implementation of this rate structure has not been done in the United States until recently. The program studied in this paper show that the specific manner of transitioning consumers is imperative to success in the short run, measured as a decrease in aggregate peak consumption. Without an effective method of communicating to consumers how their costs are shifting, they may not respond as expected until price increases are salient. Additionally, allowing consumers the opportunity to opt out presents the problem of self-selection out of the program by households that use more kWh on peak hours. Compared to the previously discussed results in Enrich et al. (2024), my results are less pronounced, and are particularly muted for several months after the program's initial rollout. Future implementations of TOU pricing should consider these factors beforehand.

References

D. J. Aigner and E. E. Leamer. Estimation of time-of-use pricing response in the absence of experimental data: An application of the methodology of data transferability. *Journal of Econometrics*, 26(1):205–227, 1984.

L. Bernard, A. Hackett, R. D. Metcalfe, and A. Schein. Decarbonizing heat: The impact of heat pumps and a time-of-use heat pump tariff on energy demand. *NBER Working Paper Series*, October 2024.

M. Boiteux. Peak-load pricing. *The Journal of Business*, 33(2):157–179, April 1960.

J. Enrich, R. Li, A. Mizrahi, and M. Reguant. Measuring the impact of time-of-use pricing on electricity consumption: Evidence from spain. *Journal of Environmental Economics and Management*, 123(1):102901, 2024.

A. Faruqui and Z. Tang. Time-varying rates are moving from the periphery to the mainstream of electricity pricing for residential customers in the united states. *Handbook on Electricity Regulation*, Forthcoming, 2023.

M. Fowlie, C. Wolfram, P. Baylis, C. A. Spurlock, A. Todd-Blick, and P. Cappers. Default effects and follow-on behaviour: Evidence from an electricity pricing program. *Review of Economic Studies*, 88:2886–2934, 2021.

M. Harding and S. Sexton. Household response to time-varying electricity prices. *Annual Review of Resource Economics*, 9:337–359, 2017.

A. J. Hinchberger, M. R. Jacobsen, C. R. Knittel, J. M. Sallee, and A. A. van Benthem. The efficiency of dynamic electricity prices. *NBER Working Paper Series*, September 2024.

J. Hirshleifer. Peak loads and efficient pricing: Comment. *The Quarterly Journal of Economics*, 72(3):451–462, August 1958.

H. S. Houthhaker. Electricity tariffs in theory and practice. *The Economic Journal*, 61(241):1–25, March 1951.

K. Ito, T. Ida, and M. Tanaka. Selection on welfare gains: Experimental evidence from electricity plan choice. *American Economic Review*, 113(11):2937–2973, 2023.

I. Pressman. A mathematical formulation of the peak-load pricing problem. *The Bell Journal of Economics and Management Science*, 1(2):304–326, 1970.

P. O. Steiner. Peak loads and efficient pricing. *The Quarterly Journal of Economics*, 71(4):585–610, November 1957.

K. Train and G. Mehrez. Optional time-of-use prices for electricity: Econometric analysis of surplus and pareto impacts. *The RAND Journal of Economics*, 25(2):263–283, 1994.

O. E. Williamson. Peak-load pricing and optimal capacity under indivisibility constraints. *American Economic Review*, 56(4):810–827, September 1966.

F. Wolak. An experimental comparison of critical peak and hourly pricing: The powercentsdc program. *Working Paper*, September 2010.