Are consumers more responsive to prices in the long run? Evidence from electricity markets

Jesse Buchsbaum*

JOB MARKET PAPER

[Click here for most recent version]

October 20, 2023

One fundamental question of economics is how consumers respond to price variation in the long run, with applications across a variety of fields. But there is a dearth of causally-identified long-run elasticity estimates, due to challenging empirical conditions. In this paper, I leverage a novel source of exogenous and persistent price variation to estimate the long-run price elasticity of demand in the setting of residential electricity. In this setting, I find that consumers are sixteen times as responsive to prices in the long run compared to the short-run, with elasticity estimates of -2.24 and -0.14 respectively. I explore mechanisms and find that in the long run, consumers respond differently to temperature across price regimes, with these differences accounting for 34% of the observed consumption differences. These findings highlight the potential impacts of price-based policies on demand and emphasize the importance of setting prices to reflect social marginal costs.

Keywords: Long-run elasticity, demand, retail electricity markets, durable goods, income heterogeneity, energy policy, carbon tax

JEL Codes: Q4, D1, D4, Q5, H2

*University of Chicago, Energy & Environment Lab, Becker Friedman Institute for Economics; email: jbuchsbaum@uchicago.edu. I thank Meredith Fowlie, Catie Hausman, Michael Greenstone, James Sallee, and Severin Borenstein for invaluable advice, feedback, and mentorship throughout this project. I am also grateful to Fiona Burlig, Lucas Davis, Koichiro Ito, Louis Preonas, David Zilberman, and participants at various seminars for helpful comments. Finally, thank you to Robert Lucadello and PG&E for invaluable assistance accessing and understanding the utility data used for this project. I do not have any financial relationships that relate to this research.

1 Introduction

How do consumers respond to a permanent difference in price? Across economics, researchers have grappled with this question from a variety of perspectives, whether to evaluate the welfare effects of a change in taxes, demand curves under technological innovation, the impact of changes in market competition, or any number of other settings where permanent changes to prices occur. Yet despite their importance and their applicability across fields, there are few causally identified empirical estimates of long-run price elasticities of demand.

The dearth of causally identified long-run elasticity estimates is a product of challenging empirical conditions – to empirically estimate a long-run price elasticity, one must leverage a persistent source of price variation that gives consumers time to adjust their behaviors and investments and reach a new equilibrium. Sources of persistent exogenous price variation are rare, however, and many estimates therefore rely heavily on structural modeling and assumptions surrounding the underlying utility function (Kamerschen and Porter, 2004; Storchmann, 2005; Dergiades and Tsoulfidis, 2008). Researchers and policymakers are left to either use these structural estimates, or to rely on short-run estimates, even though extensive research has shown that consumers learn and adjust their demand over time (Deryugina, MacKay and Reif, 2019; Karlan and Zinman, 2019).

In this paper, I leverage a novel source of persistent spatial price variation to estimate the differences between short- and long-run price elasticities of demand in the setting of residential electricity. There is perhaps no setting is more important than residential electricity to study the difference between short- and long-run consumption, with its unique combination of applicability to other important settings and direct relevance to some of the most important policy questions that exist today. In particular, the mechanisms and dynamics of price responsiveness in residential electricity markets have direct analogs to many other contexts in economics. Other utility customers settings, such as natural gas and water, are similar derived demand settings where prices can drive long-run capital intensive investments with significant consumption impacts. Transportation, and particularly gasoline consumption, exhibits similar characteristics where demand is derived and gasoline prices play an important role in vehicle investments (Busse, Knittel and Zettelmeyer, 2013; Sallee, West and Fan, 2016). Even consumer goods that are less directly analogous to electricity exhibit some similar properties that might drive differences in short-run and long-run behavior, such as habit formation and learning over time.

Furthermore, understanding how consumers respond differently in the short and long run has a number of direct implications for critical energy economics and policy questions. First, numerous policy proposals reduce emissions and mitigate climate change require significant changes to consumer behaviors. Residential electricity consumption accounts for nearly 1 billion tons in annual emissions of greenhouse gases in the US(EPA, 2023) and makes up almost 40% of US power generation.[1] Inducing changes to the consumption of residential electricity requires an understanding of how consumers respond to incentives. Second, numerous energy policies impact retail electricity prices either directly or indirectly.[2] Evaluating the impacts of these policies depends on understand-

[1]https://www.eia.gov/energyexplained/electricity/use-of-electricity.php

[2]Policies that directly impact retail electricity prices include frequent rate changes for all households, rate re-

ing how consumers will respond to retail prices in the long run. Third, long-run electricity demand forecasts are critical for grid planners and utilities to make long-term planning and infrastructure investment decisions.

To estimate short- and long-run price elasticities of demand, I leverage price variation driven by a subtle feature of California's non-linear pricing regime. In the increasing block pricing rate structure used throughout California, marginal prices increase when electricity usage exceeds a threshold. Because of differences in heating and cooling needs for households across the different climates of California, these thresholds are set to different levels that depend on where a consumer lives. Within Pacific Gas & Electric's (PG&E's) service territory, there are ten such territories, called "baseline territories", that were established in 1982. The boundaries for these territories are typically determined according to discontinuities in a household's elevation, lines drawn between two points on a map, or geopolitical demarcations. These boundaries have led to long-lasting persistent price variation, with one side of the border consistently facing higher prices than the other. I leverage these price discontinuities to estimate elasticities in both the short and long run.

Estimating a causally identified long-run elasticity that spans more than a few years is typically empirically challenging, and the panel methods commonly used in the literature can miss important margins of response. Standard panel methods compare consumption before and after a price shock, relying on counterfactual data on a consumer both before and after the change in price. Notably, they miss all energy-intensive decisions that consumers make that aren't a direct response to a price shock, including when homes are built, when new tenants move in (often a time in which home renovations occur), and when appliances break.

The persistence in cross-sectional price variation in the setting of this paper provides empirical leverage to estimate long-run price responses of 30 to 40 years, driven by cross-sectional price variation that was established in 1982. I use administrative household-level utility data leverage this cross-sectional price variation driven by the levels of the baselines across baseline territory boundaries to estimate a long-run price elasticity of demand. By leveraging this cross-sectional variation, I capture a more comprehensive measure of how consumption responds to price variation. I estimate a long-run elasticity of -2.25.

To anchor this result within the existing literature, I use standard methods to estimate a short-run price elasticity of demand.[3] In the short run, I follow the methods of Ito (2014), using a simulated instrument to isolate exogenous variation in the price schedule over time. I find that electricity consumption in the short run is relatively price inelastic, with an elasticity of -0.14 in my preferred specification,[4] indicating that consumers are 16 times as responsive to permanent price

ductions for low-income households, special rates for adopters of technology like rooftop solar or electric vehicles, and carbon taxes. Other energy policies and technologies have been shown to indirectly impact electricity prices, including electricity market deregulation(Borenstein and Bushnell, 2015), plant closures(Davis and Hausman, 2016), and renewable energy generation(Kyritsis, Andersson and Serletis, 2017).

[3]In this context, short-run refers to a timescale of one year, in contrast with the 30 to 40 year time scale of the long run esimates.

[4]This estimate is slightly larger than Ito (2014), which found a short-run elasticity estimate of -0.05, but is within the range of typical estimates in the literature – a meta-analysis from Zhu et al. (2018) finds an average short-run elasticity of -0.22 across the literature.

changes in the long run than to short-run price fluctuations.

I then explore several possible mechanisms that would drive such large price responses in the long run. First, I estimate the relationship between outdoor heat and electricity consumption across the pricing boundary. Consumers facing low prices are significantly more responsive to temperature than those facing higher prices, and the difference in temperature gradients can explain 34% of the differences in consumption driving the observed long-run elasticity. Second, I estimate how adoption of observable energy-intensive durable goods, finding no change in adoption of rooftop solar or energy efficiency across the pricing boundary. Finally, I explore the trajectory of electricity consumption over time for new builds. There are immediate and persistent differences in consumption for new builds across the pricing border, as well as changes in consumption over time that are consistent with learning and differential appliance adoption over time. That said, I do not find evidence that property characteristics vary with price, though these results are noisy.

Furthermore, substantial heterogeneity exists in price responsiveness according to a consumer's income level. There are several conflicting mechanisms that make it difficult to predict what types of consumers will be more responsive – electricity bills may be more salient to low-income households as they have less discretionary income than higher income households. However, higher income households likely have more energy-intensive appliances, leading to more margins for response to price changes. Furthermore, durable goods that reduce consumption often have high capital costs, leading to potentially greater adoption among higher income households (Borenstein, 2017). I estimate elasticities for households with different income levels, finding that low-income consumers are less responsive to changes in prices in the short run, but that in the long-run, this trend reverses and low income consumers are actually more responsive than higher-income consumers.

To understand the policy implications of the long-run elasticities estimated here, I conduct a back-of-the-envelope calculation of the impacts of a $50 per ton carbon tax on residential electricity consumption and emissions, comparing my elasticity estimates with the assumptions of climate and electricity simulation models. Under a long-run elasticity of -2.25 and a series of strong assumptions, a $50 per ton carbon tax would lead to an annual reduction of nearly 220 million tons of greenhouse gas emissions across the US from residential electricity consumption alone, or about a 13.6% reduction in total power sector emissions. This reduction is more than double the reduction under the assumptions of existing climate and electricity simulation models, demonstrating that the impact of price-based energy policies may be far greater than previously thought.

The primary contribution of this paper is that it is one of the first causally identified empirical estimates of the difference between long- and short-run price elasticities of demand, across any field. There are numerous papers across a variety of fields that use aggregated data and structured dynamic panels to estimate long-run elasticities (Kamerschen and Porter, 2004; Storchmann, 2005; Dergiades and Tsoulfidis, 2008). However, these studies rely on strong assumptions about the form of serial correlation on the error term. Many fewer studies use household- or individual-level data to estimate price responsiveness in the short- and long-run: Deryugina, MacKay and Reif (2019) estimates the dynamics of price elasticity for residential electricity consumers up to two years, finding a price elasticity of -0.09 in the first six months and -0.28 after two-and-a-half years.

Karlan and Zinman (2019) estimate that, in the context of borrowing, the elasticity with respect to the interest rate is -1.1 in the first year and rises to -2.9 in the third year. Feehan (2018) is perhaps more closely related to this work, where the author leverages a natural experiment in Canada to estimate 20-year price elasticities for residential electricity customers, finding a long-run elasticity of -1.2, but does not compare with the short-run. Here, I build on this work by estimating a price elasticity of 30 to 40 years, estimating the magnitude of difference to the short run, and exploring mechanisms.

Second, this paper contributes to the literature on durable goods investment, particularly in response to input prices. Here, I estimate how temperature responsiveness, adoption of solar and energy efficiency measures, and property characteristics varies in response to electricity prices. Chesser et al. (2018) and Crago and Chernyakhovskiy (2017) explore the impact of electricity prices on investment in rooftop solar, using aggregated data to show that electricity prices are an important drivers of residential solar adoption. Bushnell, Muehlegger and Rapson (2021) finds that electric vehicle adoption is impacted by both gasoline prices and electricity prices. Several papers have shown the importance of electricity price in appliance adoption, including Davis (2020) in the case of heating type and Rapson (2014) in the case of air conditioning. In contrast, with this literature, I find that price variation does not drive adoption of solar and energy efficiency. However, I find evidence that prices drive differences in temperature responsiveness and consumption for new builds, suggesting likely differences in investments of unobserved durable goods. This paper is the first to attribute differences in overall long-run electricity consumption to durable goods mechanisms.

Finally, this paper contributes to the literature exploring how household income impacts price responsiveness among residential electricity customers. Evidence in this literature is somewhat conflicting – Alberini, Gans and Velez-Lopez (2011) and Reiss and White (2005) find that price elasticities of demand are highest among the poorest households and monotonically decrease as income grows, while Brolinson (2019) and Schulte and Heindl (2017) find that wealthier households are more responsive to prices. Recent work by Cong et al. (2022) finds that low-income households wait until higher summer temperatures to turn on their air conditioning. This paper is the first to estimate causally identified short- and long-run price elasticities by income. I show that while higher income households are much more responsive to prices in the short- and medium-run, this trend reverses in the long run. This suggests that bill salience is particularly relevant among low-income households and that investment in durable goods may play a significant role, even among a low-income population which may have capital constraints.

The paper proceeds as follows: Section 2 discusses the institutional background; Section 3 presents the data and empirical strategy; Section 4 estimates elasticities in both the long and short run; Section 5 explores the mechanisms driving these responses to price; Section 6 examines heterogeneity; Section 7 discusses the policy implications of this work; and Section 8 concludes.

2 Background

2.1 Increasing block pricing and baseline territories

The setting for this study is Pacific Gas & Electric (PG&E), a large investor owned utility company in Northern California. PG&E uses a non-linear price schedule called "Increasing Block Pricing" to set prices for electricity. This pricing mechanism charges a higher marginal prices for higher levels of electricity usage. A customer faces the lowest Tier 1 price until they reach their baseline allowance (henceforth referred to as the "baseline"), after which they face a higher Tier 2 price. The number of pricing tiers varies over time, with PG&E using as few as two pricing tiers and as many as five.

Because of the vast variation in climate within PG&E's service territory, there is spatial variation in the quantity of electricity that is consumed before a higher tier price kicks in. For instance, PG&E's service territory includes both Fresno, with an average June high temperature of 92 degrees, and San Francisco, with average June high temperatures of 60 degrees. Because of the gap in climatic conditions, the level of electricity demand necessary to meet basic heating and cooling needs across the utility's service territory is not equal. As such, PG&E divide their customers into different territories that determine the baseline (henceforth referred to as "baseline territories") – in other words, the level of electricity that can be used before the higher marginal price takes effect. Furthermore, baselines are different in summer and winter, as well as for customers with electric versus gas heat.

PG&E divides its service territory into ten different baseline territories,[5] as shown in Figure 1. These baseline territories were established in 1982[6] by the California legislature, and adopted by by the California Public Utility Commission in 1983. Between 1983 and 1990, the CPUC continued to make small changes to where the baseline territory boundaries lay. From 1990 to 2019, the baseline territory map stayed the same, with an adjustment to one community in 2019. This community has been dropped from the sample, meaning that the baseline territory boundaries are constant over the sample period of this study. More generally, baselines are determined based on three potential factors: geopolitical demarcations (e.g. zip code/city/county boundaries, roads), elevation discontinuities, and lines drawn between two points on a map. For example, Trinity Country is divided into Territories X, Y, and Z, where residents of Trinity County below 2,000 feet of altitude are in Territory X, residents between 2,001 feet and 4,500 feet are in Territory Y, and residents above 4,500 feet are in Territory Z. Meanwhile, Contra Costa County is divided into Territories S, T, and X according to a mix of geopolitical demarcations (e.g. city boundaries, county line) and lines drawn on a map (e.g. "a line running from the most easterly point of the city limits of the City of Pinole as of July 1, 1980 to the point on the county line common to Contra Costa and Alameda Counties about one-fourth of a mile east of the intersection of Grizzly Peak

[5]Note that PG&E's baseline territories are different boundaries than the California Electricity Commission's (CEC's) "climate zones," which are used to determine building codes. Baseline territory boundaries are nearly universally separate from CEC climate zone boundaries, with a very small number of exceptions.

[6]A precursor to baseline territories was established in 1976, called "climate bands," though there were only four climate bands based purely on heating degree-days.

Figure 1: PGE baseline territories (PGE, 2020)

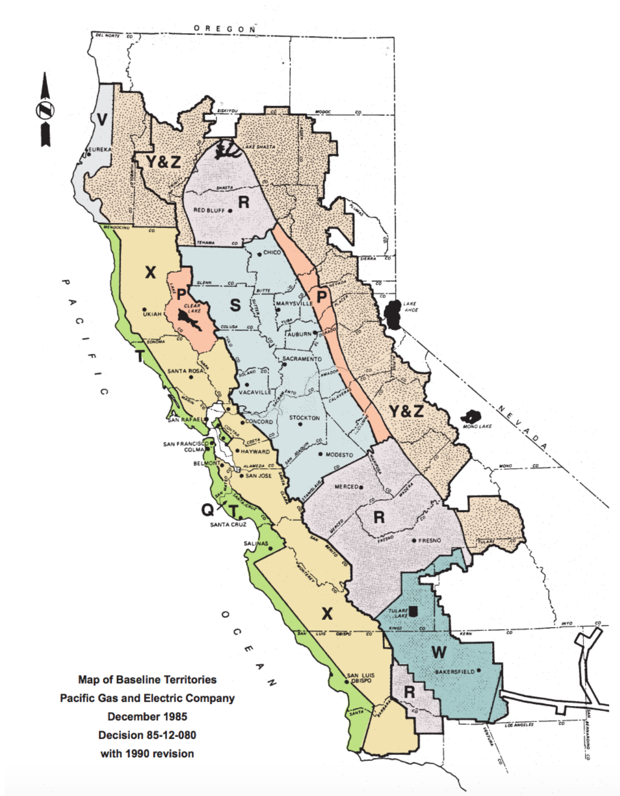

Note: This map shows PG&E's baseline territories. These borders did not change from 1990 to 2019. Source: PG&E

Boulevard and Wildcat Canyon Road...”). A full list of baseline territory boundaries defined by elevation and codes is provided in Appendix A.5.[7]

As an illustrative example, Figure 2 shows the city of Hercules, CA, a city within Contra Costa County, by baseline territory and with Census Block Group divides shown in dark red. Notably, there are households who live within the same Census Block Group, neighborhood, and even across the street from each other who are divided into different baseline territories.

To determine the level of each baseline, PG&E sets quantities so that 50 to 60 percent of expected residential electricity consumption in each climate zone is set as baseline consumption (equivalently, so that 50 to 60 percent of consumption is in Tier 1)[8]. As a result, there is a discontinuous change in the level of the baseline at the boundary of a baseline territory. Households directly on either side of this border, by virtue of their baseline territory assignment, face different price schedules. As an illustrative example, Figure 3 exhibits the price schedule faced by customers on the most common border in my sample, the border between Territories X and T, during the

[7] A full list of baseline territory definitions including those defined by non-elevation definitions can be found at https://www.pge.com/tariffs/assets/pdf/tariffbook/ELEC_PRELIM_A.pdf.

[8] One might be concerned that this could lead to endogeneity, where the actions of a household impact the baseline allowance in future periods. I assume that individual households do not exhibit market power, an assumption supported by the fact that each baseline territory contains at least 6,000 households.

Figure 2: Hercules, CA by baseline territory

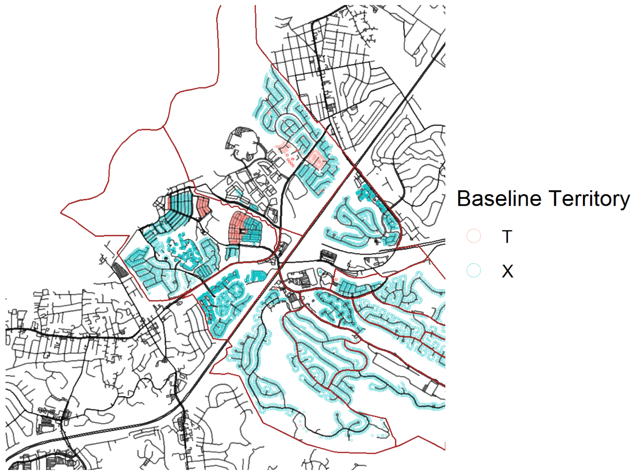

Note: This map shows the city of Hercules, CA by baseline territory. Census Block Group outlines are shown in dark red.

summer months of 2017, with the distribution of monthly consumption for Territories X and T shown below the figure. The figure shows differences in both marginal price and average price across the consumption distribution – differences across territories in average price are present throughout the consumption distribution, while differences in marginal price are larger but only present for certain pockets of consumption. Note that throughout the paper, I follow Ito (2014) and assume that consumers respond to average price, though I show that my results are similar under marginal prices as well.

Note that the existing energy economics literature, including Ito (2014), Auffhammer and Rubin (2018), and Shaffer (2020), often uses utility service territory boundaries as a source of exogenous spatial price variation. Utility service territory boundaries, however, are vulnerable to confounding non-price factors along the utility border, such as utility-specific programs and potential household selection effects. Because baseline territories boundaries are within a single utility's service territory, they are not subject to the same confounding effects to prices. Furthermore, utility service boundaries are often limited to only a narrow geographic range. PG&E's baseline territories cover a much broader spatial area, allowing for a more diverse set of households that may be more representative of the broader population.

In addition to spatial variation in prices driven by these baseline territory discontinuities, PG&E customers face price variation over time as well. Appendix Figure A1 shows the evolution of each

Figure 3: Price variation in Territory X versus Territory T

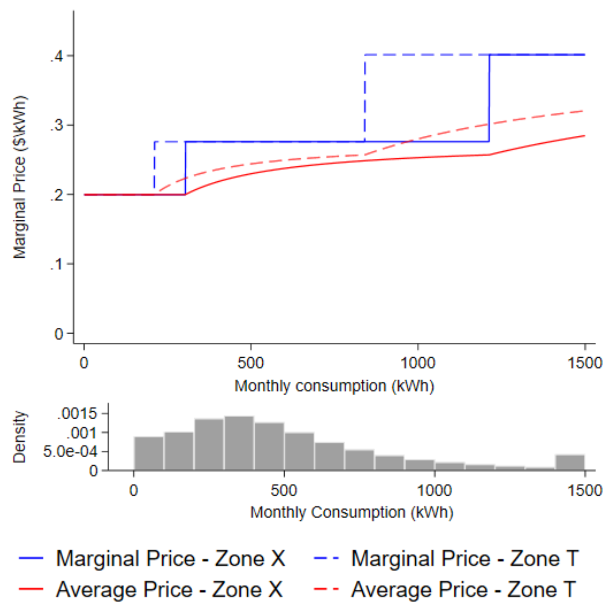

Note: This top panel of this figure shows marginal and average prices for baseline territory X compared with baseline territory T in June 2017 for customers with gas heat. Territories X and T are directly adjacent to one another. The bottom panel shows the portion of the sample with monthly consumption in each range.

price over time for the standard residential tariff, E-1. Note that there is price variation within each tier as well as changes to the number of tiers over time. Further, there is variation in levels of baselines themselves as well. Appendix Figure A2 shows the daily baselines from 1995 to 2020 for each baseline territory within PG&E. Baseline variation exists across baseline territories and even within territories from summer to winter and between customers with electric versus gas heat. Over the course the sample for this study (2008 to 2020), baseline quantities change four times.

Figure 4 incorporates all these sources of price variation to exhibit how average prices vary over time, once again using the T/X border as an illustrative example. This figure shows the average price over time for customers with gas heat and median seasonal consumption levels.[9] There are persistent price differences between the two baseline territories across the full sample period, typically between one and three cents per kWh. Note that these price differences are magnified at higher levels of monthly consumption, as shown in Appendix Figure A3.

3 Data

To estimate how consumers respond to price variation and the mechanisms driving these responses, I combine several datasets form a variety of sources. First, I leverage administrative utility billing

[9]Median summer usage is 394 kWh per month and median winter usage is 423 kWh per month for gas heat customers across the sample period.

Figure 4: Average prices over time at median consumption levels by baseline territory

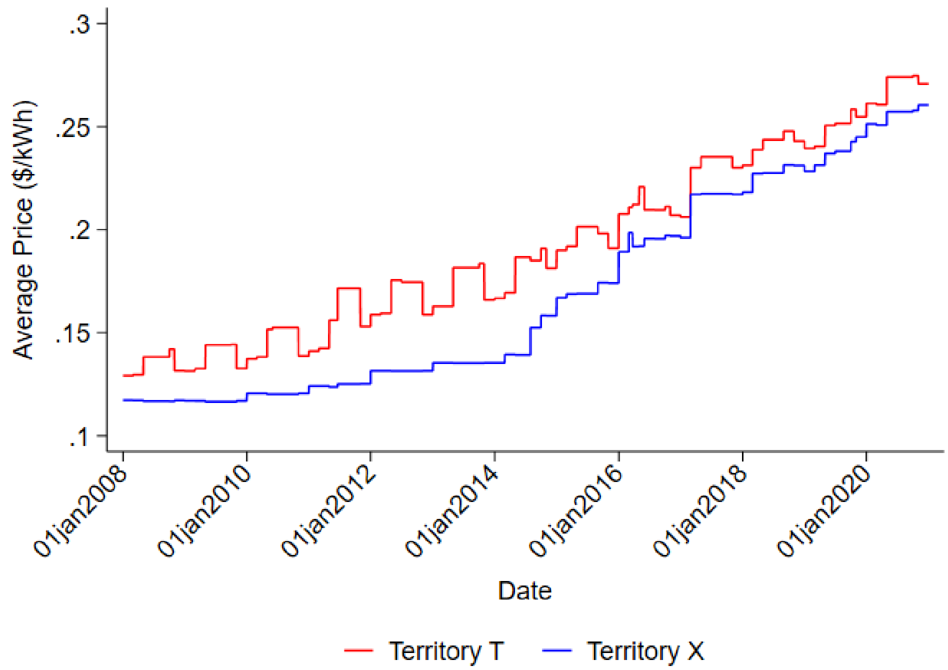

Note: This figure shows average volumetric prices over time for two households with identical consumption levels, one living in Territory T and the other in Territory X. Consumption is assumed to be the median level of consumption by season – 394 kWh per month in summer and 423 kWh per month in winter.

data at the account level for a subset of PG&E electricity customers from 2008 to 2020.[10] I observe data at the monthly level on electricity usage,[11] billing, electricity tariff, and adoption of durable goods (e.g. solar panels, electric heat, energy efficiency), as well as address and some limited demographic information. Second, I use census data from the 2017 5-Year American Community Survey (ACS) at the Census Block Group (CBG) level to obtain demographic information.[12] Third, I use daily weather data from the National Oceanic and Atmospheric Administration. The dataset, called Global Historical Climatology Network, reports daily temperatures for land surface stations across the globe. I use PG&E address data to determine the closest weather station to each household in the sample. For each billing period from 2008 to 2020, I determine the number of Heating and Cooling Degree Days[13] at the nearest weather stations for a particular household. Finally, I scrape and address-match parcel data for a five-county subsample. I obtain this data

[10] PG&E has granted me access to this data under a confidentiality agreement

[11] The electricity usage data that I use throughout this paper is *net* monthly electricity consumption. For solar customers who both generate and consume electricity, their net consumption is the difference between their gross monthly consumption and their gross monthly generation.

[12] In my sample, there are an average of 588 households in each CBG.

[13] Heating degree days and cooling degree days are common measures of how hot or cold are region is across a period of time. In particular they provide a measure of the distance from a standard neutral indoor temperature (65 degree Fahrenheit). Cooling degree days are calculated according to the following formula: $\sum_{M_i}(\text{Daily Mean Temperature}_i - 65)$ where M_i refers to the monthly billing period for each household i., while heating degree days as calculated as $\sum_{M_i}(65 - \text{Daily Mean Temperature}_i)$.

from county assessor websites for five of the largest counties in my sample[14] and it includes data on housing values, year built, square footage, and number of bedrooms. All details of the data scraping, cleaning, and merging processes are provided in Appendix A.1.

In the empirical analysis, I restrict this sample in several ways: first, I omit households with non-standard baselines such as medical baselines. Second, for consistency across bills, I include only bills ranging between 28 and 33 days. Third, because baselines change over the course of the sample, there are some geographic areas which have higher baselines than their neighbors at some point in the sample and lower baselines than their neighbors at other points. I drop these observations, ensuring that I only include households who are consistently in the "high" or "low" price region throughout the sample. After making these sampling restrictions, I retain 219,000 premises and 524,000 accounts.[15] A map with the counts of premises per CBG and the mean difference in average price by CBG is shown in Figure 10.

Figure 5: Census Block Group maps

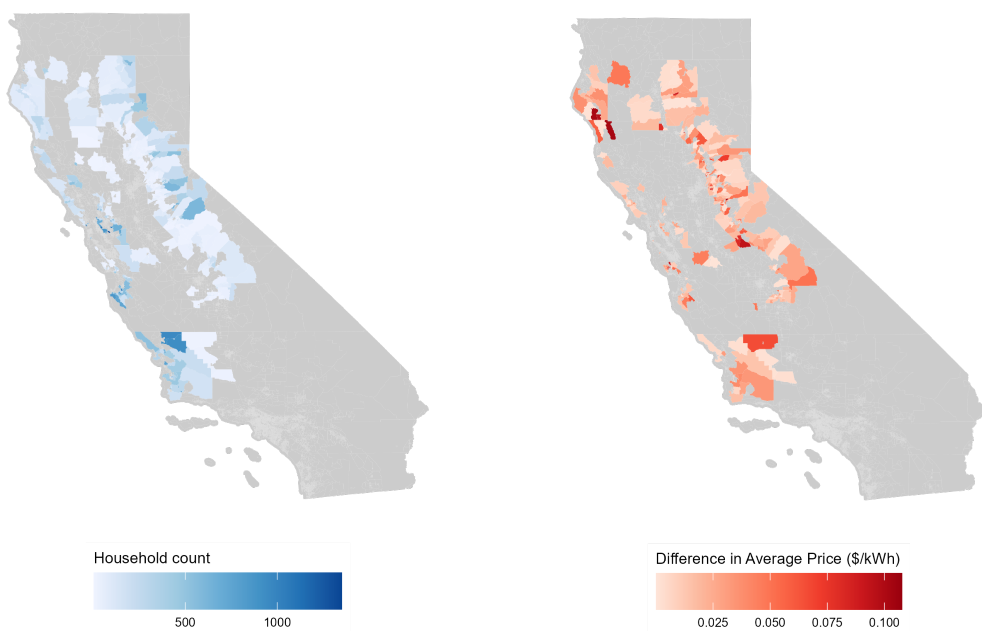

(a) Customer count by CBG

(b) Average price differences by CBG

Note: This figure shows two maps by Census Block Group. Panel A shows customers counts under the above sampling restrictions for each CBG in the sample. Panel B shows the mean differences in average price ($/kWh) across the baseline territory boundaries within each CBG in the sample.

Table 1 presents summary statistics, showing the means and standard deviations for my sample in comparison with the full set of PG&E residential customers (4.8 million premises and 21 million accounts) from 2008 to 2020. The sample in this study seems to be fairly representative of the

[14] These counties are Contra Costa, El Dorado, Monterey, Santa Barbara, and San Mateo.

[15] Note that many more accounts exist than premises. This is because most premises have multiple different account-holders over the course of the 13-year sample, due to customers moving to new premises.

Table 1: Summary statistics

	In Sample		*All PG&E*	
	Mean	Std.Dev.	Mean	Std.Dev.
Monthly baseline (kWh)	386.39	490.33	348.26	426.12
Monthly consumption (kWh)	473.63	802.62	393.37	639.80
Average Price ($/kWh)	0.22	0.49	0.20	0.50
Percent Electric Heat	0.23	0.42	0.21	0.40
Percent Solar	0.04	0.20	0.03	0.18
Percent CARE	0.22	0.42	0.25	0.43
Income per capita	40,097.17	20,843.70	39,474.15	24,393.10
N	523,927		20,922,019	

broader PG&E population in the share of customers adopting electric heat and solar and enrolling in CARE. However, the average customer in this sample faces baselines that are 11% higher than in the broader PG&E population and consumes 20% more electricity. These differences are driven by the locations of the baseline territory borders – for instance, no border runs through San Francisco, which has a very large population and low baselines and average consumption, driven by its mild climate.

4 Research Design and Results

While many papers have estimated causally identified short-run price elasticities of demand for residential electricity customers with household data, very few have done the same in the long run. In the short run, I follow the existing literature, relying on price variation over time and across space using panel methods. In the long run, however, I employ a new approach, leveraging a novel source of persistent and long-lasting cross-sectional price variation. In this section, I describe my empirical approach and results in the long run. I then anchor my results to the existing literature by using standard methods to estimate short-run price elasticities of demand. In Appendix A.7, I explore the dynamics of consumers' price responses by extending the short-run approach to the medium-run.

4.1 Long run empirical strategy and results

To estimate price elasticities of demand, economists typically leverage price variation over time and estimate how consumers react to dynamic changes in the price schedule. However, this type of analysis is limited in that it only captures certain margins of response among certain customers. First, it only captures customers who are continuously present in the sample over a period of time. However, electricity usage and price responses may vary substantially across housing age and tenure. Furthermore, as recent work by Davis (2020) and Costa and Kahn (2011) suggests, important durable good decisions such as whether a home is heated by gas or electricity may often be decided when a home is built, with substantial switching costs that lead to low incidence of switching behavior. Alternatively, investment decisions may be made when a utility account switches due to a new owner or tenant moving in, or when durable goods break and need to be replaced. The typical dynamic methods used in the literature to estimate elasticities can fail to capture the variation in recently-built homes, recently-renovated homes, or homes making investment decisions in response to non-price factors. In fact, any approach that leverages only price shocks over time to estimate consumption shocks will fail to capture these critical margins of response.

The challenge, then, is determining how to estimate differences in consumption without a "pre-period". Instead of leveraging price changes over time, I leverage cross-sectional differences in prices due to baseline territory divisions to observe long-run differences in consumption. Because the borders in my sample were drawn in 1982, finalized in 1990, and have not changed since, the differences in price driven by these borders are a long-standing, persistent source of price variation.

The baseline territory boundaries divide PG&E's service territory into a series of high price and low price regions. Because these boundaries are often drawn according to criteria distinct from other administrative boundaries – such as elevation discontinuities, or lines drawn between two points on a map – households are similar in expectation on either side of the border except for the difference in baseline (and therefore price) that they face. Importantly, I restrict my sample to borders where the ordering of baselines has remained consistent across the boundary since the start of my sample in 2008. To estimate a long-run elasticity, I leverage this cross-sectional price variation with a regression discontinuity approach.

With variation in prices and baselines across both space and time, it's important to consider exactly what variation in baselines comes from each source of variation. In Appendix Table A2, I decompose the variation in baselines according to space and time, demonstrating that the vast majority of variation in baselines comes from baseline territories, electric heat, and season.[16] Hence, to isolate the price variation driven by spatial differences in baseline territories, I include electric-heat-by-season fixed effects in all specifications.

Throughout the main body of the paper, I assume that households response to average volumetric prices[17] rather than marginal prices. Past studies from Ito (2014) and Shaffer (2020), have found that households are unlikely to respond to marginal prices in the short run in increasing block pricing settings. In particular, Ito (2014) finds that households are more responsive to average prices than to marginal prices. As such, I follow the existing literature and use average prices, though all results are qualitatively similar under marginal prices as shown in several robustness checks.

[16]Note that "season" here refers only to the difference between summer and winter.

[17]I use average volumetric prices rather than average prices due to nuances surrounding solar billing. Solar net metering customers are billed for the balance of their energy consumption once per year, rather than on a monthly basis, leading to negative bills in most months and potentially a large positive bill in a single month. Average prices for these customers do not reflect their incentives and have the potential to create bias in the sample. Rather, I use average volumetric prices, which are identical to average prices for the vast majority of customers but reflect the true incentives for solar customers.

Figure 6: Monthly baseline (kWh) by distance to border

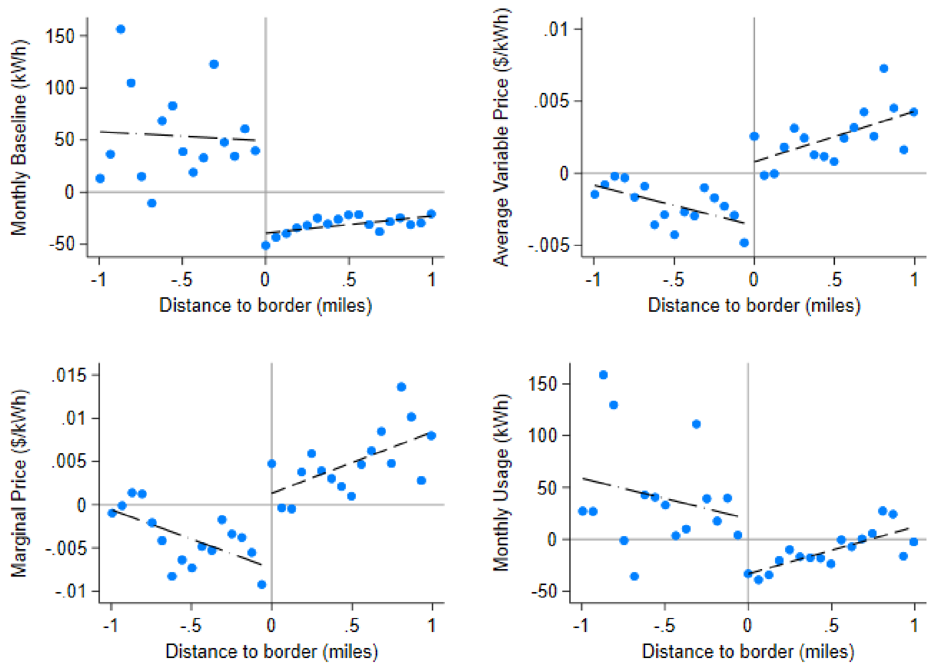

Note: In this figure, each dot represents all households within a 100-meter bin. The vertical axis shows residuals after regressing monthly baselines on CBG-by-month-of-sample and electric-heat-by-season fixed effects and taking the mean of the residual within each bin. The dotted lines are lines of best fit across all bins within a 1-mile bandwidth on either side of the baseline discontinuity.

I begin by estimating how several important variables change across the baseline territory boundary. To implement this approach, I restrict my sample to households within one mile of the baseline territory boundary. For each variable of interest, I regress on CBG-by-month-of-sample and electric-heat-by-season fixed effects, taking the means of the residuals across 100-meter bins. I then estimate linear regressions across all bins on each side of the border, plotting the line of best fit. The results are shown in Figure 6.

First, to establish that baselines vary across the border as expected, Panel A plots the magnitude of daily baselines against the distance to the border. There is a clear discontinuity at the border with magnitude of approximately 90 kWh, demonstrating that baselines are significantly impacted by the border discontinuity. This difference in baselines impacts prices, as shown in Panels B and C. While there is more noise in these regressions since marginal and average variable prices are endogenous to consumption, there is a clear impact at the border, where households with lower baselines face higher marginal prices and average variable prices, by an average of 2.8 and 1.6 cents per kWh respectively. At median levels of electricity usage, this difference in price would imply a bill difference of about $10.80 per month. Panel D shows differences in electricity consumption, where monthly electricity consumption on the "high price" side of the border is about 200 kWh lower than electricity consumption on the "low price" side of the border.

Given the persistent differences in prices and consumption, this is a natural setting to estimate a long-run price elasticity of demand. However, an important consideration when estimating elasticities with non-linear price schedules is that prices are endogenous to consumption. As customers use more electricity, the marginal price of electricity increases. As such, the marginal price of electricity is correlated with, and even partly determined by, consumption. To solve this issue, I use an instrumental variables approach, instrumenting for price with a regression discontinuity design that isolates the exogenous variation in prices driven by spatial differences in baselines. Within a narrow bandwidth of the baseline discontinuity of one mile, the only mechanism through which the baseline impacts electricity consumption is through prices.

Leveraging this approach combining a regression discontinuity and instrumental variables, I run the following regressions:

$$\text{First stage: } ln(p_{it}) = \alpha_0 + \alpha_1 \mathrm{Hi}_i + \alpha_2 \mathrm{d}_i + \alpha_3 \mathrm{d}_i \mathrm{Hi}_i + \gamma_{ct} + \eta_{es} + \epsilon_{it} \tag{1}$$

$$\text{Second stage: } ln(z_{it}) = \beta_0 + \beta_1 \widehat{ln(p_{it})} + \beta_2 \mathrm{d}_i + \beta_3 \mathrm{d}_i \mathrm{Hi}_i + \gamma_{ct} + \eta_{es} + \epsilon_{it} \tag{2}$$

where c identifies Census Block Groups, e is a dummy variable indicating if a customer has electric heat, s denotes whether the bill is in the summer or winter, z_{it} represents electricity consumption for household i in month t, p_{it} denotes the contemporaneous average variable price, d_t denotes the running variable in the regression discontinuity – the distance (in meters) between the household and the baseline territory boundary, and ϵ denotes an idiosyncratic error term. CBG-by-month-of-sample fixed effects are included to control for variation in demographic traits that may change over time. Electric-heat-by-summer fixed effects are included since baselines vary according to heating type and season. Standard errors are clustered according to baseline territory, since prices are assigned at the baseline territory level, and month-of-sample, to account for unobserved correlation in variances across seasons and over time. The exclusion restriction in this instrumental variables specification is that the baseline territory border can only impact consumption through electricity prices, conditional on the fixed effects. Similarly, the identifying assumption under this regression is that customers living close the baseline territory border with the same type of heating systems would consume similar amounts of electricity absent the differences in prices driven by baseline territory divisions.

Because an indicator for electric heat is included in the fixed effects, this specification makes a parametric correction for heating type, as the length of a baseline is partly determined by heating type. This is a necessary fixed effect to prevent an endogenous heat type choice to bias the estimates. However, the inclusion of this fixed effect eliminates heating choice as a potential mechanism to impact consumption. Therefore, the long-run elasticity estimated here is a lower bound, as a theoretical specification that allowed a heating type margin to impact consumption would only increase the estimated elasticity.

The results of these regressions are shown in Table 2. My preferred specification, shown in Column (1), restricts the sample to households within one mile of the border and uses CBG-by-month-of-sample and electric-heat-by-summer fixed effects. In this specification, I find that consumers are highly responsive to price variation in the long run, with an estimated elasticity of

Table 2: Long-run IV estimate of elasticity

	(1)	(2)	(3)	(4)
$ln(AVP)$	-2.249***	-2.386***	-1.949***	
	(0.561)	(0.563)	(0.573)	
	[0.003]	[0.002]	[0.008]	
$ln(MP)$				-1.500***
				(0.435)
				[0.007]
Distance to Border (100 meters)	-0.012***	-0.014***	-0.003	-0.014***
	(0.003)	(0.002)	(0.002)	(0.004)
	[0.006]	[0.000]	[0.220]	[0.007]
Distance x Hi	0.026***	0.037***	0.001	0.029***
	(0.005)	(0.007)	(0.004)	(0.006)
	[0.001]	[0.001]	[0.840]	[0.001]
F Stat	62.285	57.403	68.560	108.294
Bandwidth	1 Mile	1/2 Mile	2 Miles	1 Mile
CBG x MoS FE	Yes	Yes	Yes	Yes
Electric Heat x Summer FE	Yes	Yes	Yes	Yes
Observations	7,327,914	5,597,526	8,932,262	7,327,914

Note: This table shows the results from four different regressions. In all regressions, the outcome variable is logged monthly electricity consumption. Columns (1), (2), and (3) show the results of the preferred specification at bandwidths of one mile, one-half mile, and two miles respectively. Column (4) shows a regression with a one-mile bandwidth under marginal prices rather than average prices. Column (5) shows a regression analogous to Column (1) with only observations matched in the county assessor dataset. Fixed effects include CBG-by-month and electric-heat-by-season. Standard errors are clustered by baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

-2.25. Columns (2) and (3) show that estimates with half and double the bandwidth are similar in magnitude, demonstrating the robustness of these estimates to the bandwidth . Column (4) shows the same specification with respect to marginal prices rather than average prices. Consumers are somewhat less responsive to marginal prices than average variable prices, though still quite responsive with a price elasticity of -1.5. Appendix Tables A9, A10 show the first stage and reduced form estimates associated with my preferred specification. These estimates imply an average price difference of 1.7 cents per kWh, implying a price difference of $8.64 per month. Meanwhile, the reduced form estimates show that for the same change in monthly baseline, consumers respond by decreasing their consumption by about 229 kWh per month, or over 37% of the mean monthly usage.

An elasticity of -2.25 implies that that customers are highly responsive to price changes in the long run. While this estimate is substantially larger than the existing literature, there are several reasons that one should expect a larger estimate in this setting and under this methodology: first, the existing literature tends to use panel methods that compares consumption for a customer before and after a price change. This type of estimation misses important margins of response, as discussed earlier in this section. By leveraging a persistent source of cross-sectional price variation, the specification here captures the investment margin, including in new and recently transacted homes, both of which are often missed by studies that rely primarily on price variation over time, as opposed to across space.

Second, there are very few existing quasi-experimental estimates of long-run elasticities in the literature. Most estimates rely on strong structural assumptions made by researchers. One of the only quasi-experimental long-run elasticity estimate to date, Deryugina, MacKay and Reif (2019), looks only at a time horizon up to three years, and estimates elasticities using the panel methods described above, which are likely to miss important margins of response. The estimates in that paper are more directly comparable to the medium-run results shown in the Appendix, not these long-run estimates, because of the parallels in both time horizon and in the identifying variation. The other quasi-experimental long-run elasticity estimate to date, Feehan (2018), finds a long-run elasticity of -1.2 in Newfoundland and Labrador, Canada. Critically, the setting for this paper is in a different climate. Consumers in Newfoundland and Labrador face lower temperatures than California year-round, leading to less flexibility in decisions around heating and cooling. Furthermore, solar irradiance and air conditioner adoption is substantially lower, diminishing the value of some of the most important margins of long-run response observed in California. With additional margins of response and more flexible heating and cooling loads, one would expect consumers to be more responsive to prices in this setting.

4.2 Threats to identification

The identifying variation of this long-run estimation strategy is driven by cross-sectional variation in baseline territory assignment. A natural concern arises that the borders cannot be taken as quasi-random because of selection concerns, as demonstrated in Bayer, Ferreira and McMillan (2007). In particular, one might be concerned that there is non-random selection that occurred after the baseline territories were assigned, such as higher energy users self-selecting to the low price side of the border. While I cannot fully rule out this concern, I take steps to explore whether there is evidence of selection.

Before presenting such evidence, it's important to note that it's difficult for consumers to know precisely where the baseline territory divide falls. The definitions, which are themselves a mix of obtuse geopolitical definitions, elevation discontinuities, and lines drawn between two points on a map, can only be found in the Appendices of rate filings. The map that is easily accessible on PG&E's website is that shown in Figure 1 and only shows the territories in vague terms. PG&E does not, for example, provide a tool where a customer can enter an address and see the associated baseline territory. The simplest way for a customer to see their baseline territory is to look at their bill, but knowledge of where the baseline territory boundary falls would require observing the bills of customers in the surrounding area. That said, home buyers have proven to be sophisticated in what factors are capitalized into housing prices (Myers, 2019; Linden and Rockoff, 2008; Gibbons, 2004). As such, I use data on property values and characteristics to explore whether there is any evidence of changes in home values and property characteristics.

I scrape county assessor data for five of the largest counties in my sample. I then address-match this data with utility billing data, as described in more detail in Appendix A.1. The result is a dataset of over 170,000 properties that includes housing values, housing structure values, housing size, and year built. Notably, housing size and year built are not fully populated in the assessor datasets, leading to some incompleteness in those variables.

If there has been non-random selection where the counterfactual electricity consumption of individuals differ on either side of baseline territory boundaries, one might expect to see these differences reflected in housing values and property characteristics. Past research has demonstrated a strong correlation between housing value, age, and size and electricity consumption, as shown by Bao and Li (2020); Costa and Kahn (2011); and Bruegge, Deryugina and Myers (2019). Therefore, I estimate how each of these property characteristics vary across the baseline territory boundary with a regression discontinuity design:

$$\mathrm{Y}_i = \beta_0 + \beta_1 \mathrm{Hi}_i + \beta_2 \mathrm{d}_i + \beta_3 \mathrm{d}_i \mathrm{Hi}_i + \gamma_c + \eta_e + \epsilon_i \tag{3}$$

where Y_i denotes the housing value, housing structure value, square footage, and year built, and all other variables are as previously defined. CBG and electric heat fixed effects are included, as in the long-run elasticity estimation, to absorb variation in neighborhood and heating type. Standard errors are again clustered according to baseline territory, as prices are assigned at the baseline territory level.

Figure 7: Property characteristics regression discontinuities

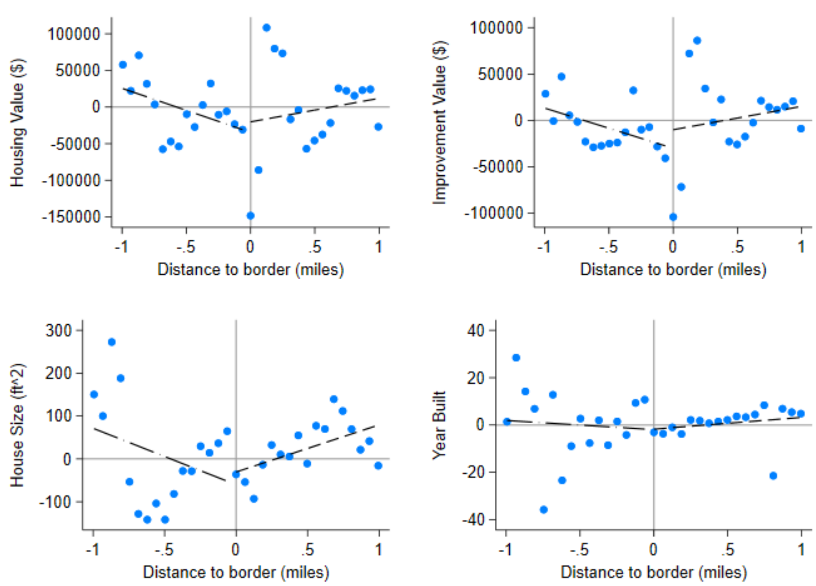

Note: In this figure, each dot represents all households within a 100-meter bin. The vertical axis shows residuals after regressing the outcomes of interest on CBG and electric-heat fixed effects and taking the mean of the residual within each bin. The dotted lines are lines of best fit across all bins within a 1-mile bandwidth on either side of the baseline discontinuity.

The results of the regressions discontinuities are shown in Figure 7, with the associated regression tables are shown in Appendix Table A11. While standard errors are large and these estimates are somewhat noisy, I do not find evidence of changes in any property characteristics crossing the border. That said, because of the size of the standard errors, I cannot rule out differences on either side of the border.

To further explore the potential impacts of any endogenous selection, I directly control for housing value in the long-run elasticity estimation, as shown in Table 3. Column (1) estimates a long-run elasticity as described in Equations (1) and (2), but restricting the sample to only households where assessor data is observed. Restricting to these households, the long-run elasticity is somewhat smaller at -1.5 but still large and highly significant. Columns (2) and (3) estimate the same specifications, flexibly controlling for housing value and structure value respectively with a quartile fixed effect. While effects are slightly attenuated, controlling for housing and structure value makes little qualitative difference, indicating that the differences in elasticities are not primarily driven by differences in housing values or characteristics.

Table 3: Long-run elasticity controlling for housing value

	(1)	(2)	(3)
$ln(AVP)$	-1.538**	-1.499**	-1.358*
	(0.598)	(0.588)	(0.630)
Distance to Border (100 meters)	-0.010***	-0.011***	-0.015***
	(0.002)	(0.002)	(0.002)
Distance x Hi	0.024***	0.024***	0.025***
	(0.002)	(0.002)	(0.002)
F Stat	25.846	33.304	52.091
Bandwidth	1 Mile	1 Mile	1 Mile
CBG x MoS FE	Yes	Yes	Yes
Electric Heat x Summer FE	Yes	Yes	Yes
Sample	Assessor	Assessor	Assessor
Controls	None	Housing Value Quartile	Structure Value Quartile
Observations	4,623,959	4,623,959	3,729,307

Note: This table shows the results from three different regressions. In all regressions, the outcome variable is logged monthly electricity consumption. Column (1) shows an estimated long-run elasticity under Equations (1) and (2) when restricting to the assessor data sample. Column (2) shows the estimated elasticity controlling for housing value quartile. Column (3) shows the estimated elasticity controlling for structure value quartile. Fixed effects include CBG-by-month and electric-heat-by-season. Standard errors are clustered by baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

Of course, there is the possibility of selection on individual preferences that aren't are not observed. For instance, households that are more attentive to their bills and the rate schedule may select to areas with lower prices and adjust their usage behaviors accordingly. This remains a weakness of the empirical approach outlined thus far in the paper – although I am able to estimate a causally-identified long-run elasticity spanning over 30 years and rule out that property characteristics are driving the magnitude of the results, the approach relies on the unverifiable assumption that households do not select on unobserved preferences that are uncorrelated with property characteristics, Census Block Group, or heating type.

4.3 Robustness checks

In Appendix A.3, I include several robustness checks. First, I test several different levels of fixed effects, finding that my long-run elasticity estimates are robust to other fixed effects strategies. Second, one might be concerned that California Electricity Climate Zones used for building energy codes are drawn along similar lines as baseline territory boundaries. This is not the case, and elasticities are similar even directly controlling for building climate zones.

Third, there is a reasonable question about what the right price variation is in this setting. Costa and Kahn (2011) show that electricity prices at the time of home construction drive future electricity consumption, suggesting that choices made at the time of home construction may drive some portion of the elasticity observed here. In that case, consumers would respond to the price at the time of home construction rather than today. I construct counterfactual prices under the 1982 and 1990 price schedules to estimate long-run elasticities under different price schedules. Because the dispersion in the price schedule is higher during my sample period than in earlier periods, using earlier price schedules actually results in higher elasticity estimates.

In addition, one might be concerned about endogeneity due to the adoption of energy-saving durable goods. For example, when a customer adopts a durable good such as solar, their net electricity usage decreases dramatically, often putting them into a different pricing tier and decreasing both their marginal and average prices. Because I use the contemporaneous average variable price as the variable of interest, there is a concern that endogenous adoption of durable goods may decrease the price difference on either side of the border, thereby biasing upwards the estimated of elasticity. In Appendix A.3.1, I test alternative definitions of price, where prices are determined by a fixed level of consumption in a baseline year, finding similar long-run elasticity estimates.

Finally, one may also be concerned that this result is driven by the presence of outliers. To rule out this possibility, I separately estimate coefficients for every Census Block Group in the sample. As shown in Appendix Figure A4, I find that 52% of CBGs exhibit elasticities between 0 and -10, and that the few outliers that do exist are not the primary factor driving the results. In Appendix Figure A5, I explore the distribution of elasticity estimates across space, finding no demographic trends that are predictive of the elasticity magnitude.

4.4 Short run estimation

In this section, I anchor these results within the existing literature by estimating short-run elasticities, which are much more commonly estimated. in contrast with my long-run approach, in the short run, I follow standard methods including Ito (2014). I rely on three primary sources of identifying variation: (1) spatial discontinuities in the baseline and therefore price that a customer faces; (2) temporal variation in prices; and (3) temporal variation in baselines.[18] In combination,

[18] Because month-of-sample fixed effects are included in all specifications, temporal variation in baselines is limited to policy changes and does not include seasonal variation.

these three sources of variation lead to prices that vary both in time and across space.

Let c_{it} denote consumption for customer i in month t and P_{it} denote the price that customer i faces in month t. For expositional purposes, P_{it} will refer to average variable price, and I assume that all customers respond to average variable price, though I will test this assumption later in this section. Typically, one could consider the following first differences estimating equation:

$$\Delta\ln(c_{it}) = \beta_1\Delta\ln(P_{it}) + \gamma_{ct} + \lambda es + \eta_{it} \tag{4}$$

where $\Delta\ln(c_{it}) = \ln(c_{it}) - \ln(c_{i,t-12})$ is the difference between log consumption today and the same month one year prior, $\Delta\ln(P_{it}) = \ln(P_{it}) - \ln(P_{i,t-12})$ is the difference between log marginal price today and the same month one year prior, γ_{ct} denotes CBG-by-time fixed effects, λes denotes electric-heat-by-season fixed effects, and $\eta_{it} = \epsilon_{it} - \epsilon_{i,t-12}$ is an idiosyncratic error term. Using this first differences estimator removes household-by-month-of-year variation. However, the structure of electricity rates in California raises issues for this estimation.

As described in the background section, electricity providers in California employ increasing-block pricing. Hence, as customers use more electricity, the marginal price of electricity increases. The marginal price of electricity is therefore correlated with consumption, meaning that in the Equation (1), the marginal price is correlated with the unobserved error term η_{it}.

To solve this issue, I follow Ito (2014). Ito instruments for price using the policy-induced price change. The instrument, called a simulated instrument in the tax literature, is

$$\Delta\ln(P_{it})^I = \ln(P_t(c_{i,t-6})) - \ln(P_{t-12}(c_{i,t-6})) \tag{5}$$

This instrument isolates the change in price induced by exogenous policy change at a specific consumption level. For it to be valid, $c_{i,t-6}$ must be uncorrelated with the unobserved error η_{it}. Some past studies have used the base year consumption, $c_{i,t-12}$, here. However, as Ito points out, mean reversion presents a challenge in this setting, as transitory shocks to consumption in month $t-12$ will cause mean reversion in consumption that will be correlated $\epsilon_{i,t-12}$ and therefore η_{it}. Blomquist and Selin (2010) and Saez, Slemrod and Giertz (2012) suggest that in an income tax setting, using consumption in a period midway between t and $t-12$ can be used to address this mean reversion problem.

This instrument might still be correlated with η_{it} if specific types of electricity users (e.g. high- and low-usage customers) have different consumption paths over time. This is where I make use of the border discontinuity that results from baseline territories. Ito uses the border discontinuity between utility regions. Here, I build on his approach by leveraging within-utility price variation driven by baseline territories. In different utility regions, there are often different incentives and marketing strategies for energy durable goods, such as solar and energy efficiency, that go beyond the price that customers face. Leveraging price variation across baseline territories allows me to isolate the price variation, without concern for these confounding factors. Furthermore, baseline territory borders are not limited to one concentrated geographic area as utility borders are, leading to a more representative sample.

Table 4: Short-run price elasticity

	MP	AP
	(1)	(2)
$\Delta\ln(MP_{it})$	-0.11***	
	(0.0043)	
$\Delta\ln(AP_{it})$		-0.14***
		(0.0048)
Observations	5284051	5272212
F	17167.8	30484.5

Note: Across all columns, the dependent variable is $\Delta\ln(c_{it})$. Fixed effects include CBG-by-month, electric-heat-by-season and 6-month-lagged consumption deciles. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

To ensure that households across the baseline territory boundary are comparable, I restrict my sample to census block groups that have at least 50 service accounts in multiple different climate zones. The resulting identifying assumption is that customers in the same census block groups on either side of the climate zone boundary would consume the same amount of energy absent the price variation that results from the climate zones.

With this instrument, I estimate a two-stage least squares regression of consumption on average volumetric price, instrumenting for average variable price with the simulated instrument described above:

$$\text{First stage: } \Delta\ln(P_{it}) = \alpha_1\ \Delta\ln(P_{it})^I + f_t(c_{i,t-6}) + \gamma_{ct} + \lambda_{es} + \eta_{it} \tag{6}$$

$$\text{Second stage: } \Delta\ln(c_{it}) = \beta_1\ \widehat{\Delta\ln(P_{it})} + f_t(c_{i,t-6}) + \gamma_{ct} + \lambda_{es} + \eta_{it} \tag{7}$$

where $f_t(c_{i,t-6})$ is a set of dummy variables determined by the decile of consumption in period $t-6$. Formally, for percentile j, $f_{j,t} = 1\{c_{j,t-6} < c_{i,t-6} \leq c_{j+1,t-6}\}$. Standard errors are clustered according to household by month-of-sample.

In this specification, β_1 represents a short-run elasticity to average volumetric price – it estimates how any exogenous price change over the previous year leads to a difference in consumption within that period. Note that, as in the long-run specification, I also include an alternative specification with marginal prices instead of average volumetric prices.

As shown in Table 4, I find results that are consistent with the existing literature (Zhu et al., 2018). Elasticities are approximately -0.11 and -0.14 for marginal and average variable prices respectively. These results primarily serve to anchor my results within the existing literature. Much of the literature on price elasticities in the residential electricity sector have focused on the short-run, and has typically found similar results to those that I present here – short-run responses to prices are relatively inelastic.

In Appendix A.7 I explore the dynamics of households' responses to price changes by estimating price elasticities in the medium run. While not the primary focus of this paper, these results provide some insight into how household behavior changes over time in response to prices.

While the magnitude of these short-run estimates (-0.14) is largely consistent with the existing literature, it draws a sharp contrast with the long-run results of -2.25. However, it is important to consider the potential mechanisms that may be captured by each specification. A short-run elasticity estimation is designed to capture consumption changes in response to intertemporal price shocks. Causal inference is simple here in a two-way fixed effects framework, where there is plausibly quasi-random variation due to fluctuations in the price schedule that impact households differently because of cross-section differences. However, this is only a narrow slice of price variation and does not capture numerous potential margins of response. In particular, price variation when a home is built may drive different choices over building materials, property characteristics, housing size, appliance choices, and more, as suggested by Costa and Kahn (2011). Similarly, price variation when energy-intensive durable goods are purchased (such as furnaces, refrigerators, clothes dryers, air conditioners, etc.) may drive consumers to make different choices. Finally, there may be learning over time that strengthens a consumption response, as shown by Deryugina, MacKay and Reif (2019). These margins of response are likely to have substantial impacts on electricity consumption, emphasizing the importance of estimating a long-run elasticity that captures these mechanisms. In the next section, I empirically estimate how customers respond in their adoption of durable goods to better understand the specific mechanisms driving the observed long-run response.

5 Mechanisms

To this point, I have shown that residential electricity customers are highly responsive to electricity prices in the long run. This contrasts with the short run, where electricity consumption is relatively inelastic, as shown in Section 4.1 and in much of the literature. This begs the question of which mechanisms are driving this response. In this subsection, I explore four such mechanisms: temperature responsiveness, solar adoption, energy efficiency adoption, and learning over time. In particular, first, I test for differences in how electricity consumption responds to temperature across the baseline territory border, and explore heterogeneity. I then directly test for differences in adoption for the two durable goods that I directly observe, solar and energy efficient appliances that are supported by PG&E incentive programs. Finally, I explore the consumption path of electricity consumption after a new home is built.

5.1 Temperature response

Past research has shown that one of the primary drivers of electricity consumption is outdoor temperature (Pardo, Meneu and Valor, 2002; Santamouris et al., 2015). Temperature responsiveness may be driven by both extensive margin durable goods investments, such as air conditioning, [19] and

[19] Air conditioning adoption is especially relevant in California, where air conditioning adoption rates are among the lowest in the country at just 72%

by intensive margin behaviors, where consumers in different price regimes may behave differently in their appliance usage and where they choose to set their thermostats. While my data do not allow for direct testing of these mechanisms, in this subsection, I explore how consumers respond to temperature.

Specifically, I estimate the relationship between outdoor heat and electricity consumption, following the "degree day method" used by Thorpe (2013) and Fowlie, Greenstone and Wolfram (2018), among households in different price regimes. I obtain temperature data for all California weather stations from NOAA. For each household bill in my sample, I determine the closest weather station with complete temperature data at the time of the billing period to that household, and merge that temperature data into the billing data. I then use that temperature data to calculate the average daily heating degree days (HDDs) and average daily cooling degree days (CDDs) in order to normalize across billing periods of different lengths. To ensure comparability across price regimes, I restrict my sample to households within a 1-mile bandwidth of the baseline territory border.

To compare households' temperature responses, I estimate the following regression specifications. Limiting my sample to summer months (May to October), I estimate how electricity consumption responds to cooling degree days across the electricity pricing border. Similarly, I estimate how electricity consumption responds to heating degrees days across the electricity pricing border, restricting my sample to winter months (November to April).

$$\mathrm{kWh}_{it} = \beta_0 + f(\mathrm{CDD}_{it})\mathrm{Hi}_i + f(\mathrm{CDD}_{it})\mathrm{Low}_i + \gamma_{ct} + \eta_{es} + \epsilon_{it} \text{ for months May to October} \tag{8}$$
$$\mathrm{kWh}_{it} = \beta_0 + f(\mathrm{HDD}_{it})\mathrm{Hi}_i + f(\mathrm{HDD}_{it})\mathrm{Low}_i + \gamma_{ct} + \eta_{es} + \epsilon_{it} \text{ for months November to April} \tag{9}$$

In both specifications, kWh_{it} denotes electricity consumption for household i in month t, CDD_{it} and HDD_{it} refer to average daily cooling and heating degree days respectively, and Hi_i is an indicator for whether the household is on the "high price" side of the pricing border. $f(CDDit)$ is a flexible nonparametric function of cooling degree day bins. Specifically, $f(CDD_{it})$ is set of indicator variables determined by the number of cooling degree days for household i in period t. Standard errors are clustered by baseline territory and month-of-sample.

Results of these regressions are shown in Figures 8 and 9. First, I find that households facing higher prices have significantly flatter responses to both hotter weather in the summer and colder weather in the winter. In the summer, the differences in prices lead to consumption differences that generally rise at higher levels of cooling degree days. In the winter, the difference is especially prevalent for bills with between 15 and 25 heating degree days per day, before consumption decreases in the low price region.[20]

The differences in electricity consumption between the high- and low-price regions can be decomposed into two parts: temperature invariant differences and temperature dependent differences.

[20] This effect at higher levels of HDD is likely driven by differences in electric heat adoption among those households in the coldest regions of California.

Figure 8: Summer monthly consumption (kWh) by CDDs

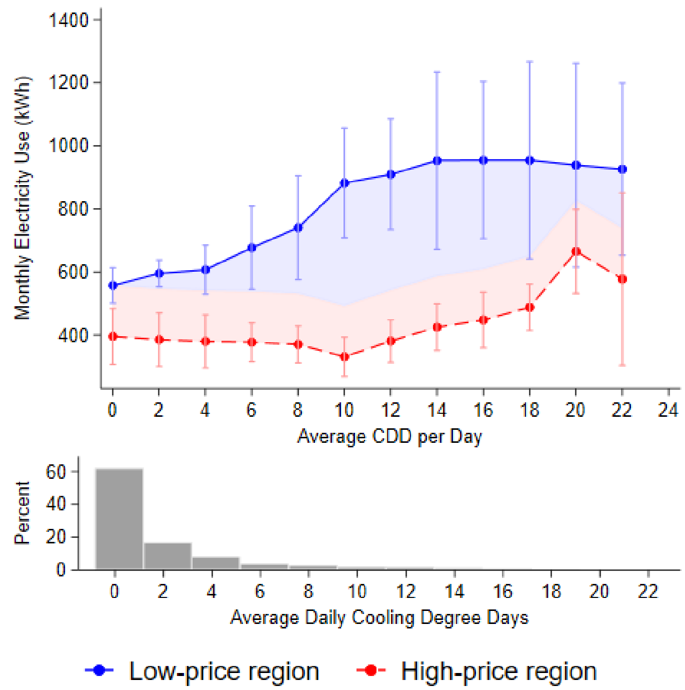

Note: This figure shows the results of regressing monthly electricity use on indicator terms for twelve different cooling degree day bins during summer, interacted with an indicator for whether a household is on the high or low price side of the border. Fixed effects include CBG-by-month and electric-heat-by-season and standard errors are clustered by baseline territory and month of sample.

Temperature invariant differences are the consumption differences that would occur in the absence of temperature changes, and can be quantified as the portion of the wedge between the low- and high-price temperature curves that holds fixed the consumption difference of the lowest CDD bin, weighted by the number of observations in each CDD bin. Temperature dependent differences are those driven by changes in the gradient of the temperature response curve, and can be quantified as the portion of the same wedge driven by differences in the slope of the temperature response curve, again weighted by the number of observations in each CDD bin. This decomposition is shown Figures 8 and 9, with the temperature invariant differences in red and the temperature dependent differences in blue.

Under this decomposition, I find that 44% of consumption differences in summer and 28% of consumption differences in winter can be explained by these differences in temperature responsiveness. In total, 34% of the observed consumption differences driving the long-run elasticity estimate can be explained by this mechanism alone.

These results follow a long literature estimating the temperature response of electricity consumption (Auffhammer and Mansur, 2014; Davis and Gertler, 2015; Kumar et al., 2020; Fazeli, Ruth and Davidsdottir, 2016; Auffhammer and Aroonruengsawat, 2012). The results here are broadly consistent with this literature – residential electricity consumption is U-shaped with respect to

Figure 9: Winter monthly consumption (kWh) by HDDs

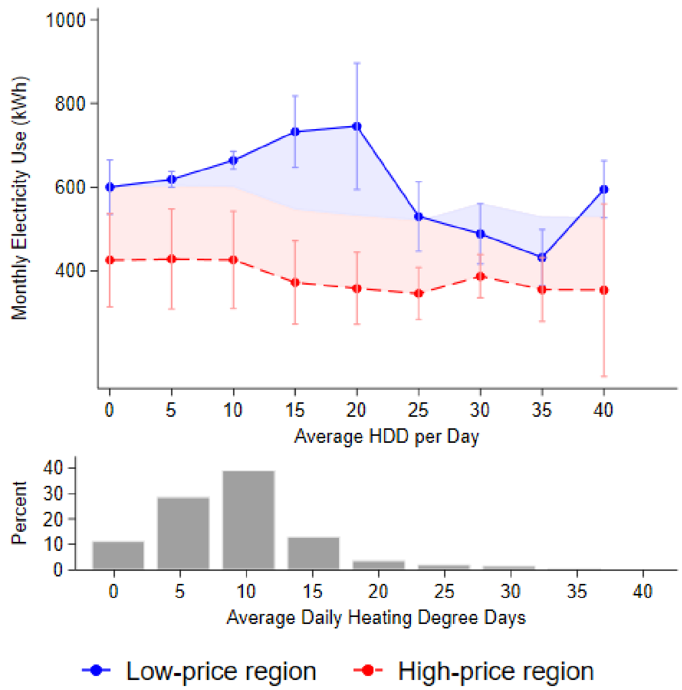

Note: This figure shows the results of regressing monthly electricity use on indicator terms for ten different heating degree day bins during winter, interacted with an indicator for whether a household is on the high or low price side of the border. This specification uses Huber-White robust standard errors.

temperature, where very low temperatures and very high temperatures both lead to increases in electricity consumption. I build on this literature by showing compelling evidence that electricity prices are particularly impactful in how households respond to temperature variation.

5.2 Solar and energy efficiency programs

To understand how customers respond to prices with durable good investment in the long run, I estimate the following model, using a standard regression discontinuity approach:

$$\text{Adoption}_i = \beta_0 + \beta_1 \text{Hi}_i + \beta_2 \text{d}_i + \beta_3 \text{d}_i \text{Hi}_i + \gamma_c + \epsilon_i$$

where Adoption_i is a binary variable indicating whether a customer ever adopts the durable good over the course of the sample; Hi_i is a binary variable indicating whether customer i lives in the "high price" baseline territory within a CBG, d_i is a measure of distance from the baseline territory boundary, c denotes CBG, and ϵ_i denotes an idiosyncratic error term.

Similar to the long run price elasticities estimated in Section 4.1, the identifying variation in this specification is cross-sectional variation in prices driven by the baseline territory discontinuity. I show the graphs resulting from this regression specification.

In Table 5, I show the results for the two durable goods of interest: residential solar adoption and utility energy efficiency programs. There is very little evidence of statistically significant changes in either solar adoption or energy efficiency adoption across the baseline territory border.[21] In fact, the 95% confidence interval in the solar specification rules out a solar adoption effect of greater than 0.5 percentage points and an energy efficiency program enrollment effect of greater than 0.3 percentage points. These results effectively rule out solar and energy efficiency programs as mechanisms driving the observed long-run elasticities.

Table 5: Adoption of energy durable goods

	(1) Solar	(2) Energy Efficiency
High Price	-0.013	-0.003
	(0.009)	(0.003)
Distance to Border (100 meters)	-0.001	-0.000
	(0.001)	(0.000)
Distance x High	0.002	0.000
	(0.001)	(0.000)
Observations	272,358	272,358

Notes: In Column (1), the outcome variable is an indicator for whether a household ever adopts solar over the course of the sample. In Column (2), the outcome variable is an indicator for whether a household ever enrolls in an energy efficiency program over the course of the sample. Fixed effects include CBG and electric heat. Standard errors are clustered by baseline territory. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

5.3 Consumption trajectories over time for new builds

There are several points in time at which households make decisions that will have significant ramifications over future energy consumption. In particular, when a home is built, there are numerous decisions made over durable goods, such as building materials, appliances, and housing size, which drive future energy use. Although many of these margins are unobserved in this setting, in this subsection, I explore how household usage differs at the time in which it is built, compared with how consumption evolves over time after a home is built.

Although housing built dates are not observed for most of the sample, I do observe when a new premises is established in the sample after 2008 and therefore when a new home is built.[22] I restrict my sample to premises established in 2009 or later to ensure that all premises appearing in the data are truly new builds. I then estimate how electricity consumption evolves over time for newly built houses over the course of the first eight years. Formally, I estimate the following

[21] Note that the only observed energy efficiency measures are PG&E programs (e.g. utility-run subsidies for energy efficiency appliances and energy audits), which are only a portion of energy efficiency measures that households adopt in practice.

[22] In particular, I observe the period in which a bill first appears.

specification:

$$\text{kWh}_{it} = \beta_0 + \sum_{k=1}^{96}\Big[\beta_k \mathbb{1}\{\text{Bill Number}_{it} = k\}\text{Hi}_i\Big] + \sum_{k=1}^{96}\Big[\beta_k \mathbb{1}\{\text{Bill Number}_{it} = k\}\text{Low}_i\Big] + \gamma_{ct} + \eta_{es} + \epsilon_{it}$$

where Bill Number_{it} refers to the numbers of bills since house i was built.

The results from this specification are shown in Figure 10, where Panel A shows the consumption trajectories over time for both the high and low price sides of the border, while Panel B shows the difference in consumption between the two pricing regions. Panel A demonstrates that there is a ramp-up period of about one year where consumption continues to grow as tenants might still be moving in or haven't fully purchased all of their appliances. After that initial period, there are immediate and persistent differences in electricity usage for new builds across the pricing border, suggesting that there are differences in housing traits or the initial set of appliances in the home. These differences grow over time, with a difference in usage of 162 kWh after 12 months that ramps up to a peak of 356 kWh after 96 months. While standard errors are large, these point estimates do provide suggestive evidence that consumers change their behaviors and/or investments over time in a way that reflects these differences in price. Furthermore, this eight-year time horizon does not capture the period over which many energy-intensive durable goods are expected to break and require replacement investments.[23] While not directly observed here, to the extent that the decisions over the replacement of these durable goods reflect price incentives, it's likely that these differences would continue to grow over time.

Figure 10: Electricity consumption over time for new builds

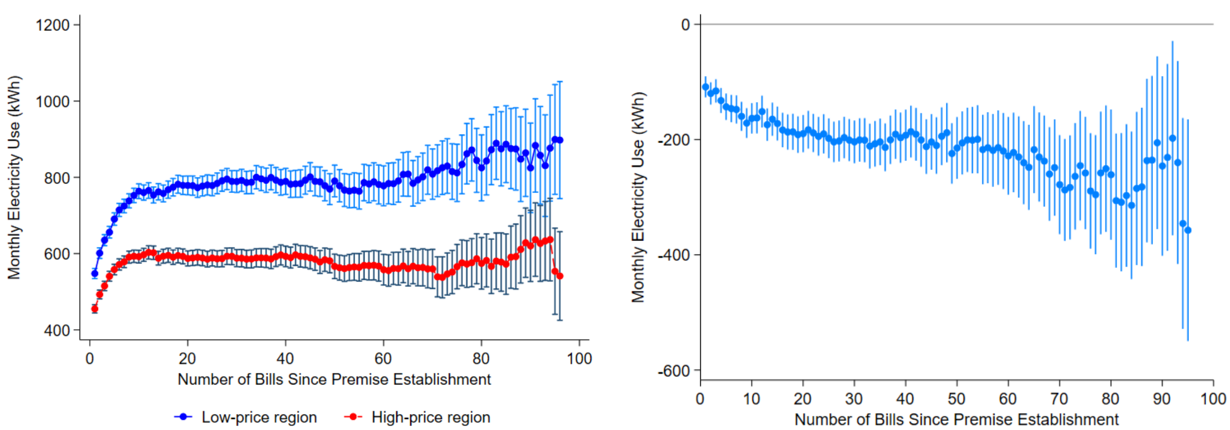

(a) Electricity consumption over time

(b) Differences in consumption

[23]For instance, air conditioners typically last 10-15 years, furnaces and boilers typically last 15-20 years, water heaters last about 10 years, refrigerators last 10-15 years, and clothes driers last 10-15 years.

6 Heterogeneity in price responses

To better understand which consumers are most responsive to prices, I explore heterogeneity in price responses, both in electricity consumption and in some of the mechanisms described in the previous section. Many papers, including Shaffer (2020) and Alberini, Gans and Velez-Lopez (2011) show there are significant heterogeneities in how customers respond to prices in their energy choices, driven by factors including information, salience, access to capital, and more. Different responses across customer groups induces heterogeneity in welfare changes. While in theory, transfers could be used to equitably redistribute any gains (or losses) from a policy, work by Sallee (2019) emphasizes the challenge that targeting presents, especially in the context of energy policy. In a context with limited transfers, understanding these mechanisms and heterogeneities is highly important for designing and evaluating policy, especially when equity is a policy objective.

6.1 Measures of income

In this setting, the primary demographic variable of interest is income. Because adoption of durable goods requires access to capital, we might expect that higher income customers are more likely to invest in durable goods that impact long-run price responsiveness. On the other hand, past work (Alberini, Gans and Velez-Lopez, 2011; Reiss and White, 2005) seems to indicate that low-income consumers tend to be more aware of their bills and may therefore may be more responsive to price fluctuations, especially in the short-run. Furthermore, Cong et al. (2022) show that low-income households wait until higher temperatures before turning on their air conditioning, by up to 7.5 degrees Fahrenheit.

While I do not directly observe income at a customer level, there are two primary ways that I explore demographic heterogeneity. First, I use CBG-level data on income from the 2017 5-year American Communities Survey to compare high-income CBGs with lower-income CBGs. Census data has numerous measures of income; my preferred measure in this work is average per-capita income.

Second, I use a proxy for income that is observed at the account-level: participation in the California Alternative Rates for Energy (CARE) program, following Auffhammer and Rubin (2018) among others. CARE is a program that is available to all energy customers in the state of California with incomes below 200% of the federal poverty level (FPL). Customers enroll directly through PG&E, who conducts income verification checks at various time intervals to ensure that customers are compliant with the income requirements. PG&E estimates that 95% of eligible customers are enrolled in CARE. Note that there is some endogeneity in which customers are enrolled in CARE that may be correlated with information and bill attention. In Appendix Figure A8, I compare how CARE participation correlates with CBG-level income deciles. While there is strong correlation between CARE enrollment and the CBG-level per-capita income, there is substantial heterogeneity in income levels within each CBG. Because of the potential endogeneity concerns with CARE, CBG-level per-capita income will be used as the primary proxy for income, while similar specifications that use CARE as a proxy for income will be shown in the Appendix.

6.2 Estimation of heterogeneity

Table 6 shows heterogeneity in elasticities by CBG-level income. To estimate this specification, I use an indicator variable to denote whether CBG-level per-capita income is above or below the sample median. I then estimated Equations 1 and 2 separately for CBG with above-median income, and compare with households with below-median income.

For robustness, I also include a specification with an indicator for whether a household has enrolled in CARE at any point during the sample in Appendix A.6.

Table 6: Long run elasticities by income

	Low income	High income
Log Average Price	-3.34***	-1.02**
	(0.59)	(0.39)
Distance to border (100m)	-0.0047	-0.019***
	(0.0046)	(0.0015)
Hi x Distance	0.019***	0.031***
	(0.0055)	(0.0048)
Observations	4,080,008	3,247,906
F	35	60

Note: This table shows an instrumental variable regressions by CBG-level per-capita income. Fixed effects include CBG-by-month and electric-heat-by-season. Standard errors are clustered by baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

Table 7 estimates a similar specification in the short-run. I estimate the same specification separately for CBGs with income levels below versus above the median in my sample. While there are some papers that find similar results (Brolinson, 2019; Schulte and Heindl, 2017), this result is in contrast with the majority of the literature, which finds that price elasticities of demand are higher among the poorer households (Alberini, Gans and Velez-Lopez, 2011; Reiss and White, 2005).

Table 7: Short-run price elasticity by income

	Low Income (1)	High Income (2)
$\Delta\ln(AP_{it})$	-0.14***	-0.20***
	(0.0047)	(0.0076)
Observations	4,801,871	4,944,301
F	34478	25215

Note: Across all columns, the dependent variable is $\Delta\ln(c_{it})$. Fixed effects include CBG-by-month and 6-month-lagged consumption deciles. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

While in the short-run specification, higher-income consumers were more responsive, in the long run, lower-income consumers consumers are significantly more responsive. This is a somewhat surprising result – higher income consumers have more access to capital with which they can invest in durable goods that impact their consumption. High levels of price responsiveness among low-income households in the long-run indicates that there may be less capital intensive margins of response that low-income households are able to leverage. In Appendix A.6 I show that this result holds using CBG-level per-capita income as a proxy for CARE.

In addition, while not the focus of this paper, in A.6 I show how consumption responds to price across several other demographic variables of interest at the CBG level, including homeownership, age of housing stock, race, and income inequality.

7 Policy Implications

The elasticities estimated in this paper have critical implications for policy, both in forecasting future electricity demand for stakeholders in the energy industry and in understanding the implications for priced-based energy policy. With respect to the former, a variety of stakeholders rely on long-term demand forecasts to make planning and investment decisions, including generation utilities, transmission and distribution utilities, grid planners, and more. These planning decisions determine when and where transmission upgrades occur and generation is built, decisions that are especially critical to reliability in the context of potential widespread electrification and other future energy transitions.

Furthermore, the estimates presented here have vast implications for price-based energy policies. Carbon taxes, electricity rates, and other price-based policies can provide strong incentives for households to engage in energy-saving behaviors. However, past research focused on short-run responses has cast doubt on how much consumers respond to electricity prices and therefore how effective price-based electricity policies can be. Many simulation models used to assess the long-run impacts of policy assume relatively low price elasticities of demand. For instance, the IMF-World Bank Climate Policy Assessment Tool (CPAT) assumes a price elasticity of demand of -0.42 for residential electricity, while the Engineering, Economic, and Environmental Electricity Simulation Tool (E4ST)[24] assumes elasticities ranging from -0.7 to -1.0, all substantially lower than the estimates in this paper of -2.25.

To demonstrate the importance of these differences, I explore how a carbon tax would lead to differences in consumption under different levels of price elasticities. Consider a carbon tax of \$50/ton, reflecting EPA's current Social Cost of Carbon (SCC).[25] Under this level of carbon tax, I calculate the associated change in residential electricity consumption and emissions under each of three different long-run elasticities: the assumptions under CPAT, E4ST, and the estimates presented in this paper. The change in residential electricity consumption is given by the following:

[24] E4ST is the result of a collaboration between Cornell University, Arizona State University, and Resources for the Future.

[25] EPA has recently proposed raising the SCC to \$190/ton. I don't make any claims as to the true cost of carbon. The thought experiment presented here is intended to demonstrate the impacts of a politically feasible policy under different elasticity assumptions, as opposed to a policy that truly reflects the social cost of carbon.

$$\Delta C = \frac{\Delta P}{P} \times C \times e_{LR}$$ [26]

I use averages across the United States for each of these parameters: the change in price ΔP under a \$50/ton carbon tax is calculated to be the marginal emissions rate for power generation in 2021 (eGRID, 2023) multiplied by the carbon tax;[27] the baseline electricity price P is the average retail US electricity price in 2021 (EIA, 2023); baseline consumption C is total residential electricity sales in the US in 2021 (EIA, 2023); and the long-run price elasticity of demand takes on one of three values, as discussed above.

Table 8 displays the results of this back-of-the-envelope analysis. In total, the emissions reductions from a carbon tax of \$50/ton under an elasticity of -2.25 represent a 35% reduction in total emissions from residential electricity consumption, or 14% of all power sector emissions, more than double of that under counterfactual elasticity assumptions.[28]

Table 8: Impacts of a \$50/ton carbon tax back-of-envelope calculation

	This paper Elasticity = -2.25	E4ST Elasticity = -1.0	CPAT Elasticity = -0.42
Price increase (cents/kWh)	2.14	2.14	2.14
Consumption reduction (%)	-35.3%	-15.7%	-6.6%
Change in residential electricity consumption (million MWh)	-519	-231	-97
Change in GHG emissions (million metric tons)	-222	-99	-42
Change in GHG emissions (% of residential power sector emissions)	-35.3%	-15.7%	-6.6%
Change in GHG emissions (% of total power sector emissions)	-13.6%	-6.1%	-2.5%

Note: This table shows the implications of a \$50/ton carbon tax under three different price elasticities of demand. Column (1) shows estimates under the CPAT elasticity of -0.4; Column (2) shows estimates under the E4ST elasticity of -1.0; and Column (3) shown estimates under the long-run elasticity estimated here of -2.25.

While the specific magnitude back-of-the-envelope calculation requires a series of strong assumptions, the qualitative differences across elasticity estimates demonstrate the scale to which long-run elasticity estimates make a difference when estimating the impacts of policy. While this section explores the impact of a carbon tax, the long-run elasticities estimates here apply to any policy that has permanent price impacts, including utility rate reform, rate subsidies, infrastructure investments that are passed on to ratepayers, and more. As policymakers consider changes to the energy industry and transitions to a more sustainable systems, the long-run elasticities estimated here suggest that price-base policies may drive long-term impacts in demand-side investments and behaviors much more than previously thought.

[26] Note that this is a just a rearrangement of a standard formula for elasticity: $e = \frac{\Delta Q}{Q} / \frac{\Delta P}{P}$.

[27] I assume that the carbon tax is fully passed through from wholesale electricity prices to retail rates.

[28] These estimates hold under a series of strong assumptions, namely that (1) long-run elasticities are similar throughout the rest of the country; (2) the cost of a carbon tax is fully passed through from wholesale to retail electricity prices; and (3) the existing power generation mix holds true to 2021 levels.

8 Conclusions

In this paper, I leverage a novel source of cross-sectional price variation for one of the first household-level causally-identified estimates of how consumers respond differently to prices in the short and long run, across any field. The magnitude of this estimated long-run elasticity is more than an order of magnitude times larger than in the short run. Furthermore, the long run estimate is larger than the existing residential electricity literature, potentially due to methodological differences that allow me to capture additional margins of response. Typical quasi-experimental methods rely on tracking the same consumers before and after a price change, missing investment choices made at the time a home is built or before new tenants move in. In this paper, estimation of long-run elasticities relies on cross-sectional price variation, allowing me to capture additional margins of response in the comparison of similar households facing different price regimes.

To explore the difference in magnitudes between short- and long-run elasticities, I explore several mechanisms that may drive the observed elasticity. First, I find that temperature responsiveness varies widely by price regime, and differences in temperature responsiveness can explain 34% of the observed long-run elasticities. Second, I directly observe adoption behaviors of two durable goods that might be used in response to price changes – rooftop solar and utility energy efficiency programs. Consumers are largely unresponsive to prices in their adoption of solar or energy efficiency programs. Third, I explore how electricity consumption evolves for new builds, finding that while there are immediate and persistent differences in consumption across price regimes, the differences grow over time, suggesting a role of learning or behavioral adjustment over time.

In addition, I explore the impact that income has on price responsiveness. I find that low-income consumers are less responsive to price changes in the short run, but that low-income consumers are more responsive than higher-income consumers in the long run. These findings highlight that higher income consumers may have more margins to adjust their usage in the short- and medium-run (e.g. more appliances that they are able to turn off in response to price changes), but that prices may play a larger role for low-income household in making investment choices surrounding energy-intensive goods.

These results have important implications for policy. I show that under $50/ton carbon tax, residential electricity consumption would decrease by over 35%, leading to a 13.6% reduction in total power sector emissions from the residential sector alone. This back-of-the-envelope calculation demonstrates the important role that price-based policies can play in decarbonization efforts. When consumers respond to prices to the extent estimated here, price-based policies are an appealing option to internalize emissions externalities. It also, however, emphasizes the important of getting prices right. Electricity prices that depart from the social marginal cost may lead to inefficient consumption levels, potentially leading to losses in social welfare.

References

Alberini, Anna, Will Gans, and Daniel Velez-Lopez. 2011. "Residential consumption of gas and electricity in the US: The role of prices and income." *Energy Economics*, 33(5): 870–881.

Auffhammer, Maximilian, and Anin Aroonruengsawat. 2012. "Hotspots of climate-driven increases in residential electricity demand: A simulation exercise based on household level billing data for California." *Public Interest Energy Research Program.*

Auffhammer, Maximilian, and Edward Rubin. 2018. "Natural gas price elasticities and optimal cost recovery under consumer heterogeneity: Evidence from 300 million natural gas bills." *National Bureau of Economic Research*, , (w24295).

Auffhammer, Maximilian, and Erin Mansur. 2014. "Measuring climatic impacts on energy consumption: A review of the empirical literature." *Energy Economics*, 46: 522–530.

Bao, Helen XH, and Steven Haotong Li. 2020. "Housing wealth and residential energy consumption." *Energy Policy*, 143(111581).

Bayer, Patrick, Fernando Ferreira, and Robert McMillan. 2007. "A unified framework for measuring preferences for schools and neighborhoods." *Journal of political economy*, 115(4): 588–638.

Blomquist, Sören, and Håkan Selin. 2010. "Hourly wage rate and taxable labor income responsiveness to changes in marginal tax rates." *Journal of Public Economics*, 94(11-12): 878–889.

Borenstein, Severin. 2017. "Private net benefits of residential solar PV: The role of electricity tariffs, tax incentives, and rebates." *Journal of Public Economics*, 4(S1): S85–S122.

Borenstein, Severin, and James Bushnell. 2015. "The US Electricity Industry After 20 Years of Restructuring." *Annual Review of Economics*, 7: 437–463.

Brolinson, Becka. 2019. "Does Increasing Block Pricing Decrease Energy Use?: Evidence from the Residential Electricity Market." *Georgetown University, Department of Economics.*

Bruegge, Chris, Tatyana Deryugina, and Erica Myers. 2019. "The distributional effects of building energy codes." *Journal of the Association of Environmental and Resource Economists*, 6(S1): S95–S127.

Bushnell, James, Erich Muehlegger, and David Rapson. 2021. "Do Electricity Prices Affect Electric Vehicle Adoption?" *Georgetown University, Department of Economics.*

Busse, Meghan R, Christopher R Knittel, and Florian Zettelmeyer. 2013. "Are consumers myopic? Evidence from new and used car purchases." *American Economic Review*, 103(1): 220–256.

Chesser, Michael, Jim Hanly, Damien Cassells, and Nicholas Apergis. 2018. "The positive feedback cycle in the electricity market: Residential solar PV adoption, electricity demand and prices." *Energy Policy*, 122: 36–44.

Cong, Shuchen, Destenie Nock, Yueming Lucy Qiu, and Bo Xing. 2022. "Unveiling hidden energy poverty using the energy equity gap." *Nature Communications*, 13(2456).

Costa, Dora L, and Matthew E Kahn. 2011. "Electricity consumption and durable housing: understanding cohort effects." *American Economic Review*, 101(3): 88–92.

Crago, Christine Lasco, and Ilya Chernyakhovskiy. 2017. "Are policy incentives for solar power effective? Evidence from residential installations in the Northeast." *Journal of Environmental Economics and Management*, 81: 132–151.

Davis, Lucas. 2020. "What Matters for Electrification? Evidence from 70 Years of US Home Heating Choices." *National Bureau of Economic Research*, , (w28324).

Davis, Lucas, and Catherine Hausman. 2016. "Market Impacts of a Nuclear Power Plant Closure." *American Economic Journal: Applied Economics*, 8(2): 92–122.

Davis, Lucas, and Paul Gertler. 2015. "Contribution of air conditioning adoption to future energy use under global warming." *Proceedings of the National Academy of Sciences*, 112(19): 5962–5967.

Dergiades, Theologos, and Lefteris Tsoulfidis. 2008. "Estimating residential demand for electricity in the United States, 1965–2006." *Energy Economics*, 30(5): 2722–2730.

Deryugina, Tatyana, Alexander MacKay, and Julian Reif. 2019. "The long-run dynamics of electricity demand: Evidence from municipal aggregation." *American Economic Journal: Applied Economics*, 12(1): 86–114.

EPA. 2023. "Inventory of U.S. Greenhouse Gas Emissions and Sinks: 1990-2021." *U.S. Environmental Protection Agency*, , (EPA 430-R-23-002).

Fazeli, Reza, Matthias Ruth, and Brynhildur Davidsdottir. 2016. "Temperature response functions for residential energy demand – A review of models." *Urban Climate*, 15: 45–59.

Feehan, James. 2018. "The long-run price elasticity of residential demand for electricity: Results from a natural experiment." *Utility Policy*, 51: 12–17.

Fowlie, Meredith, Michael Greenstone, and Catherine Wolfram. 2018. "Do energy efficiency investments deliver? Evidence from the weatherization assistance program." *Quarterly Journal of Economics*, 133(3): 1597–1644.

Gibbons, Steve. 2004. "The costs of urban property crime." *The Economic Journal*, 114(499): F441–F463.

Ito, Koichiro. 2014. "Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing." *American Economic Review*, 104(2): 537–63.

Kamerschen, David, and David Porter. 2004. "The demand for residential, industrial and total electricity, 1973–1998." *Energy Economics*, 26(1): 87–100.

Karlan, Dean, and Jonathan Zinman. 2019. "Long-run price elasticities of demand for credit: evidence from a countrywide field experiment in Mexico." *The Review of Economic Studies*, 86(4): 1704–1746.

Kumar, Rohini, Benjamin Rachunok, Debora Maia-Silva, and Roshanak Nateghi. 2020. "Asymmetrical response of California electricity demand to summer-time temperature variation." *Scientific Reports*, 10(1): 1–9.

Kyritsis, Evangelos, Jonas Andersson, and Apostolos Serletis. 2017. "Electricity prices, large-scale renewable integration, and policy implications." *Energy Policy*, 101: 550–560.

Linden, Leigh, and Jonah E Rockoff. 2008. "Estimates of the impact of crime risk on property values from Megan's laws." *American Economic Review*, 98(3): 1103–1127.

Myers, Erica. 2019. "Are home buyers inattentive? Evidence from capitalization of energy costs." *American Economic Journal: Economic Policy*, 11(2): 165–188.

Pardo, Angel, Vicente Meneu, and Enric Valor. 2002. "Temperature and seasonality influences on Spanish electricity load." *Energy Economics*, 24(1): 55–70.

Rapson, David. 2014. "Durable goods and long-run electricity demand: Evidence from air conditioner purchase behavior." *Journal of Environmental Economics and Management*, 68(1): 141–160.

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

Saez, Emmanuel, Joel Slemrod, and Seth Giertz. 2012. "The elasticity of taxable income with respect to marginal tax rates: A critical review." *Journal of Economic Literature*, 50(1): 3–50.

Sallee, James M, Sarah E West, and Wei Fan. 2016. "Do consumers recognize the value of fuel economy? Evidence from used car prices and gasoline price fluctuations." *Journal of Public Economics*, 135: 61–73.

Sallee, Jim. 2019. "Pigou creates losers: On the implausibility of achieving pareto improvements from efficiency-enhancing policies." *National Bureau of Economic Research*, , (w25831).

Santamouris, Matheos, Constantinos Cartalis, Afroditi Synnefa, and Dania Kolokotsa. 2015. "On the impact of urban heat island and global warming on the power demand and electricity consumption of buildings—A review." *Energy and buildings*, 98: 119–124.

Schulte, Isabella, and Peter Heindl. 2017. "Price and income elasticities of residential energy demand in Germany." *Energy Policy*, 102: 512–528.

Shaffer, Blake. 2020. "Misunderstanding nonlinear prices: Evidence from a natural experiment on residential electricity demand." *American Economic Journal: Economic Policy*, 12(3): 433–61.

Storchmann, Karl. 2005. "Long-run gasoline demand for passenger cars: the role of income distribution." *Energy Economics*, 27(1): 25–58.

Thorpe, David. 2013. *Energy Management in Buildings: The Earthscan Expert Guide.* Routledge.

Zhu, Xing, Lanlan Li, Kaile Zhou, Xiaoling Zhang, and Shanlin Yang. 2018. "A meta-analysis on the price elasticity and income elasticity of residential electricity demand." *Journal of Cleaner Production*, 201: 169–177.

A Appendix

A.1 Data Appendix

The primary dataset used throughout this analysis is an administrative dataset of monthly electricity billing data for PG&E. I pair this data with weather data from NOAA, demographic data from the US Census, and property characteristics data for five counties in my sample. In this section, I discuss the cleaning and merging process for each of these datasets.

A.1.1 PG&E Data

PG&E data is divided into six different datasets: First, billing data is observed at the monthly level from 2008 to 2020 for all residential PG&E customers (about 5 million) and includes customer and premise identifiers; bill start and end dates; tariff code; baseline territory; electricity consumption (kWh); bill amount ($), monthly baseline, and indicators for solar, electric heat, CARE, FERA, balanced payment plans, and solar true-up. Second, customer data is a cross section of all PG&E service accounts that are open at any point during the sample period and includes customer and premise identifiers; Census Block Group; CEC Climate Zone; elevation estimate; indicator for e-Billing; an indicator for master-metered homes; and premise start dates.[29] Third, energy efficiency data is observed at the service account ID level and includes information on enrollment in PG&E energy efficiency programs, including installation date and EE measure description and codes. Fourth, PG&E provided addresses for a subset of the sample. In particular, PG&E limited the number of addresses that they were willing to provide to 500,000. As such, I asked only for and received address data only customers from Census Tracts that had at least 15 customers on either side of a baseline territory boundary, leaving 463,292 premises. This dataset included premise identifiers and street address including Zip-9 for most customers. Fifth, PG&E provides public information on historic electricity rates and baselines. In private email communications with PG&E's rate team, they shared rates data going back further all the to way to 1977. From all this data, I manually created a daily panel of rates and baseline data for all periods from 1982 to 2020.

To clean PG&E's billing data, I began by taking a number of simple steps: I drop bills where billing dates, electricity consumption, and bill size are identical. I also drop any bills longer than 33 days or shorter than 27 days. There are some cases where bills are divided into two periods in the data, where the first day of a bill is shown as a different observation – for these instances, I combine that first day with the rest of the bill. At this point, I merge the billing, customer, energy efficiency, and address datasets together on customer and premise IDs.

The next step is to determine the distance from a household to the nearest baseline territory border. This process presents a challenge, as baseline territory boundaries are not cleanly defined continuous lines, as shown in Figure 2. To overcome this challenge, for household i, I define the running variable for the regression discontinuity as the geographic distance from household i to the nearest household in a different baseline territory:

$$RV_i = \min_{b \in B} \{dist(coord_i, coord_b)\} \tag{10}$$

where RV_i denotes the final RD running variable for household i; B denotes all premises in baseline

[29] Note that while premise start dates are observed, communications with PG&E revealed that theses data are imperfectly collected and should not be trusted, particularly for households built prior to 2005. As such, this field is largely ignored in the analysis.

territory B; $coord_i$ denotes the latitude and longitude of household i; and $dist$ refers to a function to calculate geographic distance between two coordinates.[30] I calculate these distances from household i to all baseline territory boundaries of interest within the same Census Block Group.

The next challenge is defining which side of the baseline territory boundary is the "high price" side of the border. Note that there are occasionally three or more baselines within a single CBG, all with different baselines. I define "high price" according to a comparison with the closest neighboring baseline territory to facilitate a direct comparison in the RD. Because there is variation in baselines over time, it's possible that one side of the border would have a higher baseline in one period, but a lower baseline in a later period. To ensure that order of baselines does not change throughout the sample for a given comparison group, I drop all borders where this is the case.

For there comes the most challenging portion of the data cleaning process – bringing in PG&E rates and baseline data. Rates and baselines change over time in two ways – first, when there are updates to the rate schedule; and second, baselines and certain tariffs vary from season to season. Furthermore, PG&E doesn't use a standard billing cycle across all households, instead bill start and end dates vary throughout a month. Because these changes over time typically do not line up with billing cycles, rates often change within a single billing period. In these instances, rates and baselines must be prorated according to the portion of the billing cycle to which they apply.

To prorate bills, I convert rates, baselines, and bills to daily panels. The three datasets are then merged together at the daily level, then collapsed back to the bill cycle level where the price schedules and monthly baselines are now weighted averages based on the portion of the bill cycle during which that rate schedule and baseline applied. I then calculate marginal and average volumetric prices. Marginal prices are defined to be the price under the maximum level of usage within a bill cycle. Average volumetric prices are defined to be:

$$AVP = \frac{c_{0:100}p_{0:100} + c_{100:130}p_{100:130} + c_{130:200}p_{130:200} + c_{200:300}p_{200:300} + c_{300:400}p_{300:400} + c_{>400}p_{>400}}{c} \tag{11}$$

where AVP denotes average volumetric price; $c_{x:y}$ denotes consumption from x percent of the baseline to y percent of the baseline; $p_{x:y}$ denotes price for the same; and c denotes total monthly electricity consumption.

A.1.2 Weather data

I use weather data from the Global Historical Climatology Network from NOAA. I use the R package "rnoaa" to directly download daily average, minimum, and maximum temperatures for the 720 land surface stations in California from 2008 to 2020. I convert temperatures to degrees Faranheit, then calculate daily Heating and Cooling Degree Days (HDD and CDD) according to the following formulas:

$$HDD = 65 - \frac{T_{max} - T_{min}}{2}$$
$$CDD = \frac{T_{max} - T_{min}}{2} - 65$$

[30] I use the R function "pointDistance" from the raster package to calculate distances.

The data are then restricted to stations with data for the full period of the sample, leaving 181 total weather stations in the state.

From there, I follow another nearest neighbor matching procedure, where I again use the "pointDistance" function to calculate the geographic distance from each household to the nearest weather station. I then merge HDD and CDD information into the billing data.

A.1.3 Demographic data

I download demographic data from the 2017 5-year American Communities Survey using the US Census API for each CBG in my sample. The primary demographic variable of interest in this study is per-capita income, but I also download data on housing vintage, ownership status, racial demographics, and measures of income inequality. These data are saved and merged into the billing data at the CBG level.

A.1.4 Property characteristics

Finally, I scrape county parcel assessor data for five of the largest counties in my sample - Contra Costa, El Dorado, Monterey, Santa Barbara, and San Mateo. Each county assessor website and available information differs, leading me to use a mix of manual and automated scraping procedures to download all available data, including addresses, housing values, and any available information on property characteristics.[31]

Matching addresses to PG&E addresses presented an additional challenge. Addresses often do not have a standard form across datasets with slightly different formats, styles, and representations for the same address.[32] To allow for matching across the two datasets, I standardize addresses in each dataset. I parse addresses into their component parts, before run a matching algorithm designed to find matches on several specific components.

The matching algorithm works as follows: for a given household in the PG&E address dataset, the algorithm looks for houses that are matches on street number, street name, either city or zip code, and unit number. Notably, there are cases where individual apartments are separately metered in PG&E's data but are represented by a single unit in the county assessor data – therefore, I allow for matches where a unit number is present in the PG&E data but missing in the assessor data. Under this algorithm, the match rate and total number of matched households is shown in Table A1.

It is worth highlighting that there other methods of address-matching. In particular, fuzzy string matching algorithms are commonly used in similar applications, where the algorithm calculates a distance between two strings and only keeps matches with a distance below a threshold. By opting for my own algorithm, I make a decision to be overly conservative and to minimize the false positive rate, potentially at the expense of the overall match rate. The goal here is to only include matches that I have a very high degree of confidence are true matches across the two datasets.

[31] Note that under the NDA signed with PG&E, I am not permitted to enter PG&E addresses into an external web browser. All scraping procedures were therefore done to download all addresses within a county, with merging occurring only locally after the data was fully downloaded.

[32] A few examples include different street suffixes ("Street" versus "St"); different names for apartments ("Unit" versus "Apartment" versus "Apt"); and multiple different spelling for the same street name ("Saint" versus "St" or "1st Street" versus "First Street").

Table A1: Assessor data matches by county

County	Match Rate	Household Count
Contra Costa	85%	32,066
El Dorado	89%	52,579
Monterey	91%	15,405
Santa Barbara	87%	24,633
San Mateo	85%	45,775

A.2 Supplementary background tables and figures

This section of the Appendix displays supplementary figures and tables relevant to the institutional background and price variation. Figures A1 and A2 show variation in the price schedule and levels of the baselines over time for all PG&E customers.

Figure A1: PG&E price evolution over time

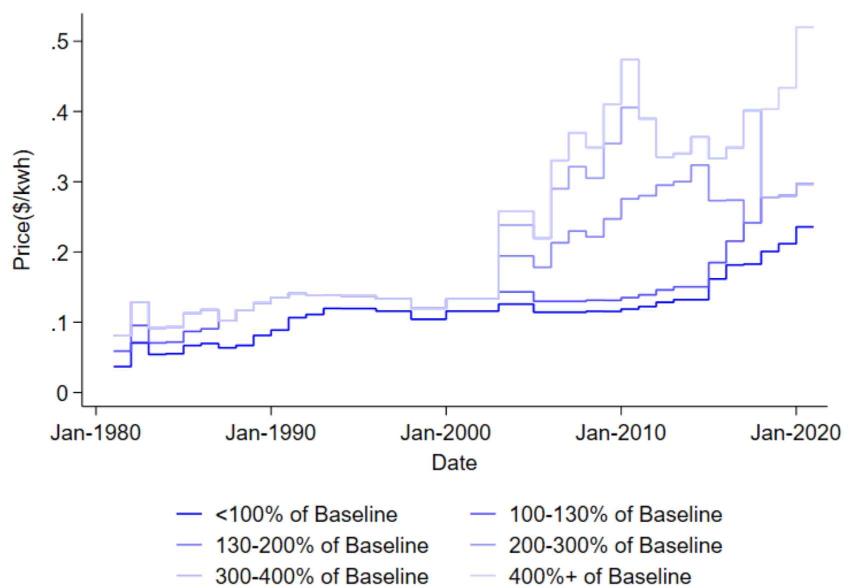

Note: This figure shows the price schedule over time for PG&E's standard default non-time varying tariff (Tariff E-1) from 1982 to 2020. The darkest line shows marginal prices for the lowest level of usage over time (under 100% of the baseline), while lighter shades show marginal prices for higher levels of usage.

Table A2 decomposes the variation in baselines, by regressing baselines on baseline territory and month of sample fixed effects, along with controls for the other determinants of baselines – electric versus gas heating and summer versus winter. Column 1 shows how much of the variation in baselines can be explained by controls alone, while Columns 2 through 4 add spatial and time series fixed effects sequentially to demonstrate the extent to which each type of fixed effect explains variation in the baseline. The vast majority of variation in baselines not accounted for by the controls can be explained by spatial fixed effects, with a small amount explained by temporal fixed effects, suggesting that cross-sectional variation plays a major role in creating price differences. It

Figure A2: Baselines over time

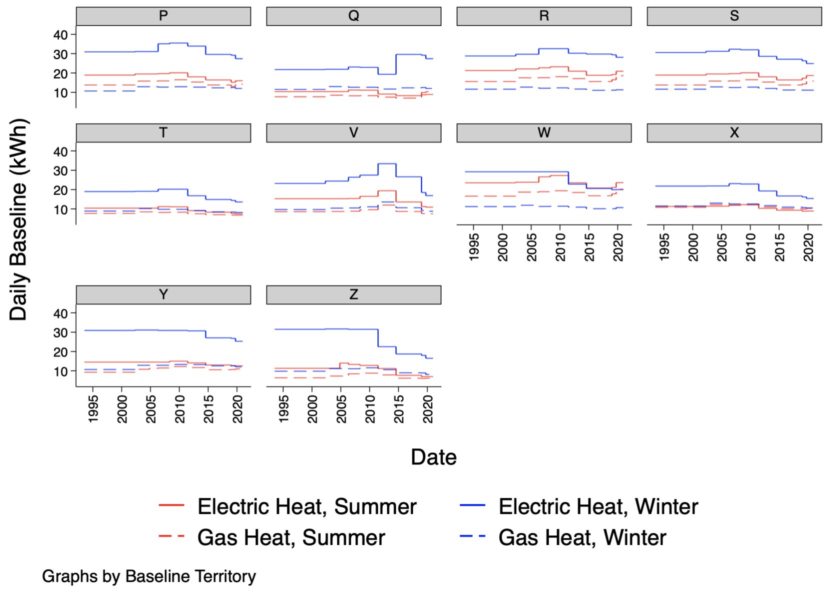

Note: This figure shows the level of baselines over time for each baseline territory and for households with electric heat in summer and winter. Data are shown from 1995 to 2020, as this is the period for which public data is available.

is also worth noting that, as expected, almost all variation in baselines (99%) can be explained by spatial and temporal fixed effects in combination with controls for the other determinants of baselines.

Table A2: Baseline variation decomposition

baseline_{it}	(1)	(2)	(3)	(4)
R^2	0.70	0.93	0.75	0.99
ElectricxSummer FE	Yes	No	No	No
ElectricxSummerxBT FE	No	Yes	No	No
ElectricxSummerxMofS FE	No	No	Yes	No
ElectricxSummerxBTxMofS FE	No	No	No	Yes

Notes: This table shows the results of four regressions, all with the length of baseline as the dependent variable. The level of observation is a customer account by month of sample.

Figure A3: Average prices over time at high consumption levels by baseline territory

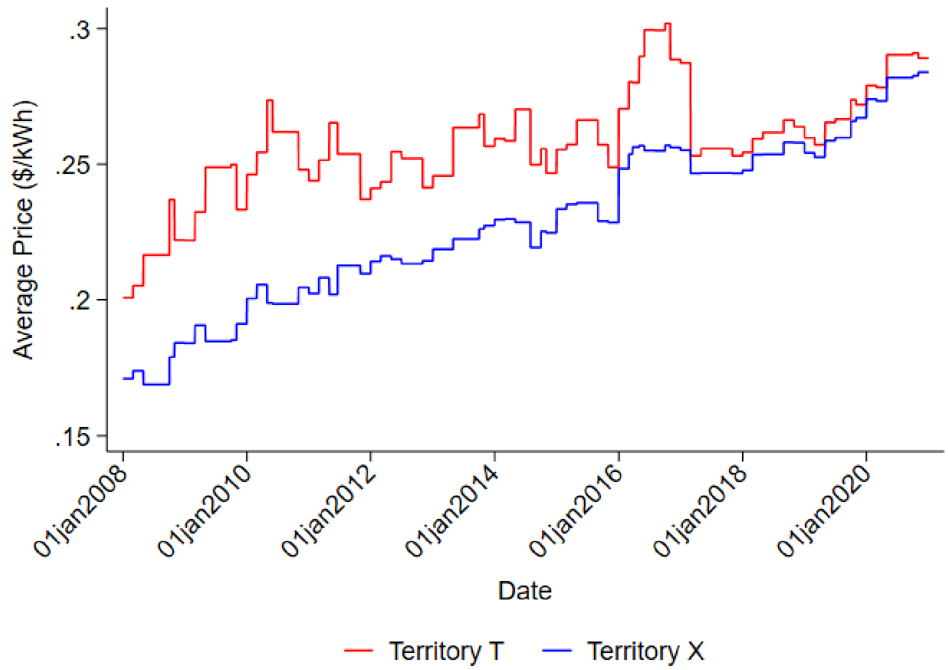

Note: This figure shows average volumetric prices over time for two households with identical consumption levels, one living in Territory T and the other in Territory X. Consumption is assumed to be the double the median level of consumption by season – 788 kWh per month in summer and 846 kWh per month in winter.

A.3 Robustness checks

This section shows results for the robustness checks outlined in the main text. To begin with, I show that these estimates are not driven by a few outlier Census Block Groups, as demonstrated with a histogram and map of CBG-level long-run elasticity estimates in Figures A4 and A5.

Figure A4: Histogram of long-run elasticity distribution across Census Block Groups

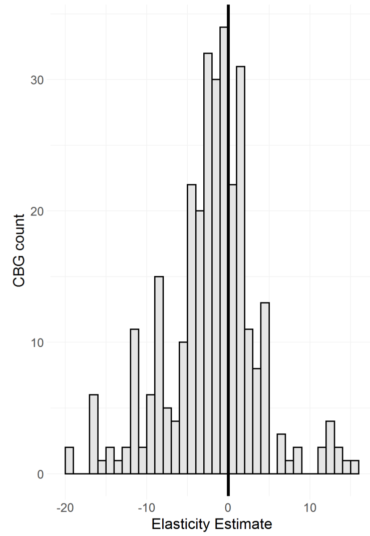

Note: This figure shows the results of estimating long-run elasticities within each Census Block Group. The histogram shows the distribution of CBG-specific elasticity estimates. CBGs with elasticities below -20 or above 20 have been omitted for scaling purposes.

Figure A5: Map of long-run elasticities by Census Block Group

Note: This map shows the results of estimating long-run elasticities within each Census Block Group. The color of each CBG on the map indicates the long-run elasticity estimate for that CBG. CBGs with elasticities below -20 or above 20 have been omitted for scaling purposes.

Next, I test several different levels of fixed effects, finding that the long-run elasticity estimates are robust to alternative fixed effects strategies, as shown in Table A3.

Table A3: Long run elasticities alternative fixed effects

	(1)	(2)	(3)	(4)
Log Average Price	-2.42**	-2.54**	-2.48***	-2.25***
	(0.96)	(0.85)	(0.61)	(0.56)
Distance to border (100m)	-0.015	-0.014	-0.013	-0.012***
	(0.0092)	(0.0100)	(0.010)	(0.0035)
Hi x Distance	0.049***	0.049**	0.049***	0.026***
	(0.015)	(0.016)	(0.015)	(0.0051)
F Stat	14.898	19.545	21.888	62.285
Electric Heat x Summer FE	No	Yes	Yes	Yes
MofS FE	No	No	Yes	Yes
CBG x MofS FE	No	No	No	Yes
Observations	7,333,627	7,333,627	7,333,627	7,327,914

*Note: Standard errors are clustered by baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

Second, I control for the California Electricity Climate Zones which are used for building energy codes, as shown in Table A4.

Table A4: Controlling for building climate zones

	IV
Log Average Price	-2.24***
	(0.57)
Distance to border (100m)	-0.013***
	(0.0037)
Hi x Distance	0.027***
	(0.0053)
F Stat	65.803
Observations	7,242,829

*Note: Fixed effects include CBG-by-month and electric-heat-by-month-of-sample. Standard errors are clustered by baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

A.3.1 Alternative price definitions

Third, I explore how long-run elasticity estimates vary under different definitions of price. There are two concerns that I check here: first, differences in the price schedule when baseline territories were drawn, as opposed to in the contemporaneous period of the sample; second, the endogeneity of prices later in the sample given that consumption may change as a result of durable goods investment.

Potential endogeneity of prices In the main body specifications, one might be concerned about endogeneity, where the adoption of an energy-saving technology might cause a dramatic change in usage and push a household into a lower pricing tier. I test this concern with two alternative definitions of price: first, I calculate what monthly average variable price would have been under monthly consumption levels from 2008 (the first year of data in my sample) and under the present-period price schedule[33]. Second, to confirm that prices aren't dependent on that single year of data, I repeat the same exercise with 2009 consumption levels. As shown in Tables A5 and A6, I find similar estimates, demonstrating that this potential endogeneity is not driving the observed results.

Table A5: Long run elasticities with prices according to 2008 consumption

	1 mile
Log Average Price under 2008 Consumption	-1.99***
	(0.48)
Distance to border	-0.00012***
	(0.000035)
Hi x Distance	0.00025***
	(0.000048)
Observations	7270569
F	68.8

*Note: Fixed effects include CBG-by-month and electric-heat-by-season. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

Table A6: Long run elasticities with prices according to 2009 consumption

	1 mile
Log Average Price under 2009 Consumption	-2.01***
	(0.47)
Distance to border	-0.00012***
	(0.000035)
Hi x Distance	0.00025***
	(0.000048)
Observations	7276524
F	61.8

*Note: Fixed effects include CBG-by-month and electric-heat-by-season. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

Changes to the rate schedule over time In Figure A1, it's clear that there is variation in rates over time. In fact, PG&E began with only two tiers, then expanded the rate schedule to as many as five tiers, before finally returning to two tiers by the end of the sample. Because many energy-intensive choices are made at the time a home is built, as demonstrated by Costa and Kahn (2011), household may respond to prices at the time a home is built. If price differences across baseline territories are substantially different than they are in the contemporaneous period, this could lead to bias in the first stage estimates.

[33]Specifically, I match according to the month. For instance, average variable price in February 2011 would be determined from consumption levels in February 2008 and the price schedule in February 2011.

I construct counterfactual prices under the 1982 and 1990 price schedules to estimate long-run elasticities under these different price schedules. The results of this exercise are shown in Tables A7 and A8. In these periods, the dispersion in the price schedule was actually lower than in the contemporaneous period. Thus, under these previous price schedules, first stage estimates of the variation in price across the border are actually smaller, leading to even larger elasticity estimates. In this sense, the estimates presented in the main body can be thought of as a lower bound.

Table A7: Long-run elasticity estimates under a 1982 price schedule

	1 mile
Log Average Price under 1982 Price Schedule	-3.53**
	(1.26)
Distance to border	-0.014***
	(0.0039)
Hi x Distance	0.028***
	(0.0061)
Observations	7327914
F	242.2

*Note: Fixed effects include CBG-by-month and electric-heat-by-month-of-sample. Standard errors are clustered by baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

Table A8: Long-run elasticity estimates under a 1990 price schedule

	1 mile
Log Average Price under 1982 Price Schedule	-4.76**
	(1.71)
Distance to border	-0.014***
	(0.0041)
Hi x Distance	0.028***
	(0.0063)
Observations	7327914
F	241.8

*Note: Fixed effects include CBG-by-month and electric-heat-by-month-of-sample. Standard errors are clustered by baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

A.4 Supplementary RD graphs and tables

In the main body text, I show a mix of graphs and tables for each regression discontinuity specification, with reference to additional supplementary tables. These tables are provided in this section.

Table A9: Long-run first stage regression on average variable price

	1 mile	1/2 mile	2 miles
hi	.017***	.016***	.019***
	(.0018)	(.0017)	(.002)
Distance to border (100m)	-2.2e-06*	-1.0e-06	1.2e-08
	(1.1e-06)	(1.0e-06)	(7.8e-07)
Hi x Distance	7.6e-06***	1.0e-05***	-1.8e-07
	(1.5e-06)	(1.1e-06)	(1.6e-06)
Observations	7.4e+06	5.6e+06	9.0e+06

Note: This table shows the results from three different regressions. In all regressions, the outcome variable is logged monthly electricity consumption. Fixed effects include CBG-by-month and electric-heat-by-season. Standard errors are clustered by baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

Table A10: Long-run reduced form regression on monthly electricity consumption

	1 mile	1/2 mile	2 miles
hi	-229**	-208**	-227**
	(82)	(82)	(80)
Distance to border (100m)	-.025	-.066	6.1e-04
	(.042)	(.067)	(.025)
Hi x Distance	.031	.074	-.022
	(.052)	(.092)	(.025)
Observations	7.5e+06	5.7e+06	9.1e+06

Note: This table shows the results from three different regressions. In all regressions, the outcome variable is logged monthly electricity consumption. Fixed effects include CBG-by-month and electric-heat-by-season. Standard errors are clustered by baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

Figure A6: Solar adoption

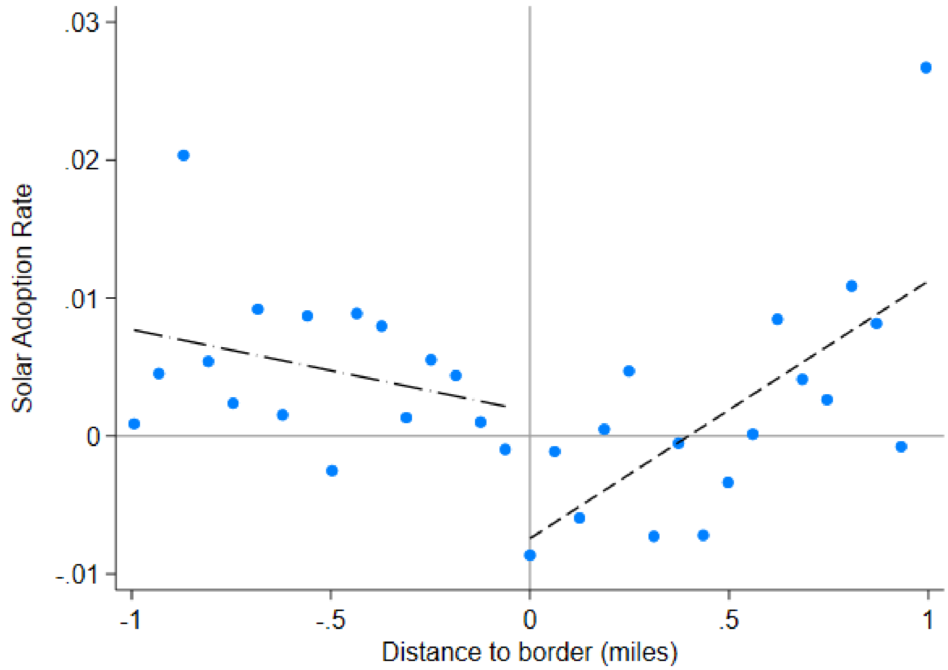

Figure A7: Energy efficiency adoption

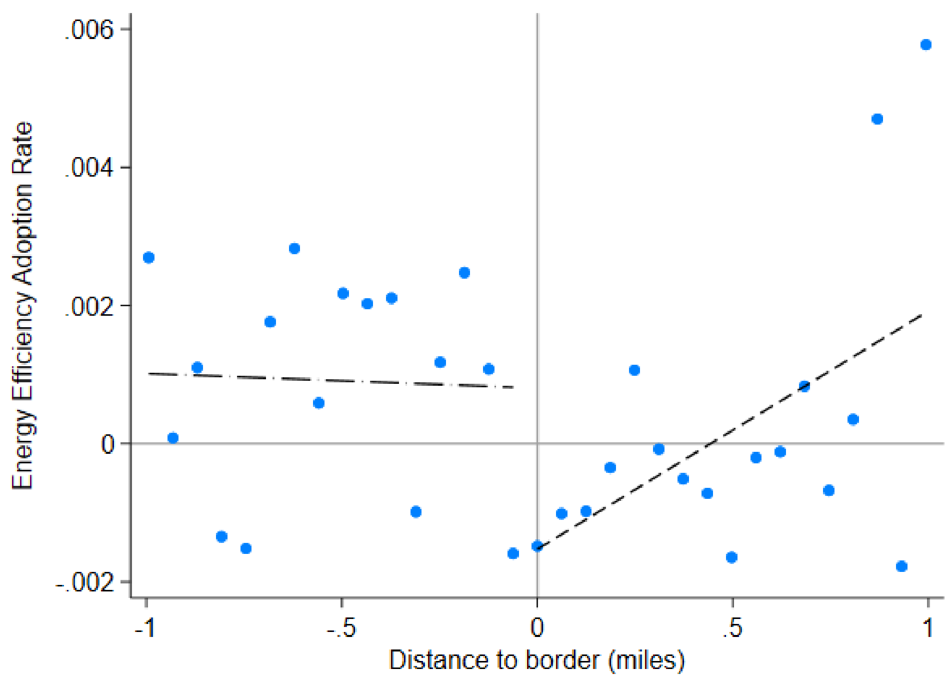

Table A11: Parcel characteristics regression discontinuity

	(1) Housing Value	(2) Land Value	(3) Improvement Value	(4) House Size (Square Feet)	(5) Number of Bedrooms
High Price Zone	-6.58e+04	8762.731	65497.303	-102.546	-7.244
	(1.38e+05)	(21228.800)	(71476.084)	(117.043)	(6.081)
	[0.650]	[0.697]	[0.402]	[0.421]	[0.287]
Distance to Border (100 meters)	-7691.875	-7276.812	-1.48e+04	-18.017	0.749
	(17728.932)	(4604.923)	(11127.753)	(11.938)	(0.867)
	[0.680]	[0.175]	[0.242]	[0.192]	[0.427]
Distance x Hi	18670.403	13402.566	27843.708	48.828*	-0.372
	(34728.229)	(10156.357)	(21743.958)	(20.369)	(0.852)
	[0.610]	[0.244]	[0.257]	[0.062]	[0.681]
_cons	6.23e+05***	1.80e+05***	2.80e+05***	1854.288***	1984.000***
	(1.25e+05)	(20704.774)	(40172.391)	(89.362)	(5.461)
	[0.002]	[0.000]	[0.001]	[0.000]	[0.000]
Bandwidth	1 Mile	1 Mile	1 Mile	1 Mile	1 Mile
CBG x MoS FE	Yes	Yes	Yes	Yes	Yes
Electric Heat x Summer FE	Yes	Yes	Yes	Yes	Yes
Observations	51066	43006	43006	41380	29957

Note: Each column shows the result of a regression discontinuity, where the outcome variable is exhibited at the top of the column. Fixed effects include CBG and electric heat. Standard errors are clustered by baseline territory. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.

A.5 Definition of baseline territories

Pacific Gas and Electric Company		Revised	*Cal. P.U.C. Sheet No.* 34601-E
U 39 *San Francisco, California*	*Cancelling*	Revised	*Cal. P.U.C. Sheet No.* 12081-E

ELECTRIC PRELIMINARY STATEMENT PART A Sheet 1
DESCRIPTION OF SERVICE AREA & GENERAL REQUIREMENTS

A. DESCRIPTION OF SERVICE AREA AND GENERAL REQUIREMENTS

1. TERRITORY SERVED BY PG&E

   a. The Pacific Gas and Electric Company (PG&E) supplies electric service in all or portions of 47 counties in the northern and central part of the State of California. A map of counties and associated zip codes that PG&E provides service to can be found on PG&E's website at http://www.pge.com/tariffs/ under electric maps. (N) I (N)

   b. The baseline territories used in the residential rate schedules are shown below for each county. Baseline territories correspond with elevation lines, unless specific boundaries were drawn to avoid dividing communities or neighborhoods as described in Section A.1.c. (T) (T)

County	Locations, Elevation Range or Description at c. Below	Baseline Territory Code
ALAMEDA	c.(1)(S)	S
	c.(1)(T)	T
	All Other	X
ALPINE*	All	Z
AMADOR	Under 1,500'	S
	1,500'-3,000'	P
	3,001'-6,000'	Y
	Over 6,000'	Z
BUTTE	Under 1,500'	S
	1,500'-3,000'	P
	3,001'-4,800'	Y
	Over 4,800'	Z
CALAVERAS	Under 1,500'	S
	1,500'-3,000'	P
	3,001'-6,000'	Y
	Over 6,000'	Z
COLUSA	All	S
CONTRA COSTA	c.(2)(S)	S
	c.(2)(T)	T
	All Other	X
EL DORADO*	Under 1,500'	S
	1,500'-3,000'	P
	3,001'-6,000'	Y
	Over 6,000'	Z
FRESNO*	Under 3,500'	R
	3,501'-6,500'	Y
	Over 6,500'	Z
GLENN	Under 3,000'	S
	Over 3,000'	Y
HUMBOLDT	c.(3)(V)	V
	All Other	Y
KERN*	Under 1,000'	W
	Over 1,000'	R
KINGS*	All	W
LAKE*	All	P
LASSEN*	Under 4,800'	Y
	Over 4,800'	Z

*Pertains to PG&E electric service area only.

(Continued)

Advice	4535-E-A	*Issued by*	*Date Filed*	December 15, 2014
Decision		***Steven Malnight***	*Effective*	December 17, 2014
		Senior Vice President	*Resolution*	
		Regulatory Affairs		

A. DESCRIPTION OF SERVICE AREA AND GENERAL REQUIREMENTS (Cont'd.)

1. TERRITORY SERVED BY PG&E (Cont'd.)

County	Locations, Elevation Range or Description at c. Below	Baseline Territory Code	
MADERA*	Under 4,000'	R	
	4,001'-6,500'	Y	
	Over 6,500'	Z	
MARIN	c.(4)(T)	T	
	All Other	X	
MARIPOSA	Under 3,500'	R	
	3,501'-6,000'	Y	
	Over 6,000'	Z	
MENDOCINO	c.(5)(T)	T	
	All Other	X	
MERCED	All	R	
MONTEREY	c.(6)(T)	T	
	All Other	X	
NAPA	All	X	
NEVADA	Under 1,500'	S	
	1,500'-3,000'	P	
	3,001'-5,500'	Y	
	Over 5,500'	Z	
PLACER*	Under 1,500'	S	
	1,500'-3,000'	P	
	3,001'-5,500'	Y	
	Over 5,500'	Z	
PLUMAS*	Under 4,800'	Y	
	Over 4,800'	Z	
SACRAMENTO	All	S	
SAN BENITO	c.(7)(T)	T	
	All Other	X	
SAN FRANCISCO	All	T	
SAN JOAQUIN	All	S	
SAN LUIS OBISPO	c.(8)(R)	R	
	c.(8)(T)	T	
	All Other	X	
SAN MATEO	c.(9)(T)	T	
	c.(9)(Q)	Q	
	All Other	X	
SANTA BARBARA*	c.(10)(R)	R	
	c.(10)(T)	T	
	All Other	X	
SANTA CLARA	c.(11)(Q)	Q	
	All Other	X	
SANTA CRUZ	Under 1,500'	T	
	1,500' & Over	Q**	(T)

* Pertains to PG&E electric service area only.
** Territory Q also includes customers in the following locations (zip codes) within Santa Cruz County at elevations less than 1,500 feet: Ben Lomond (95005), Boulder Creek (95006), Brookdale (95007), Felton (95018), Mount Hermon (95041) and unincorporated areas (95033). (N)

(Continued)

ELECTRIC PRELIMINARY STATEMENT PART A Sheet 3
DESCRIPTION OF SERVICE AREA & GENERAL REQUIREMENTS

A. DESCRIPTION OF SERVICE AREA AND GENERAL REQUIREMENTS (Cont'd.) (T)

1. TERRITORY SERVED BY PG&E (Cont'd.) (T)

County	Locations, Elevation Range or Description at c. Below	Baseline Territory Code	
			(L)
SHASTA	Under 2,000'	R	|
	2,001'-4,500'	Y	|
	Over 4,500'	Z	|
SIERRA	Under 5,500'	Y	|
	Over 5,500'	Z	|
SISKIYOU*	Under 4,500'	Y	|
	Over 4,500'	Z	|
SOLANO	c.(12)(X)	X	|
	All Other	S	|
SONOMA	c.(13)(T)	T	|
	All Other	X	|
STANISLAUS	All	S	|
SUTTER	All	S	|
TEHAMA	Under 2,500'	R	|
	2,501'-4,800'	Y	|
	Over 4,800'	Z	|
TRINITY	Under 2,000'	X	|
	2,001'-4,500'	Y	|
	Over 4,500'	Z	|
TULARE*	Under 1,000'	W	|
	1,001'-3,500'	R	|
	3,501'-6,500'	Y	|
	Over 6,500'	Z	|
TUOLUMNE*	Under 1,500'	S	|
	1,500'-3,500'	P	|
	3,501'-6,000'	Y	|
	Over 6,000'	Z	|
YOLO	All	S	|
YUBA	Under 1,500'	S	|
	1,500' & Over	P	(L)

* Pertains to PG&E electric service area only. (D)

(Continued)

A.6 Supplementary heterogeneity figures and tables

This section of the Appendix shows figures and tables that supplements the heterogeneity estimates. Figure A8 shows the correlation between CARE enrollment and CBG-level per-capita income.

Figure A8: CARE enrollment versus CBG-level per-capita income decile

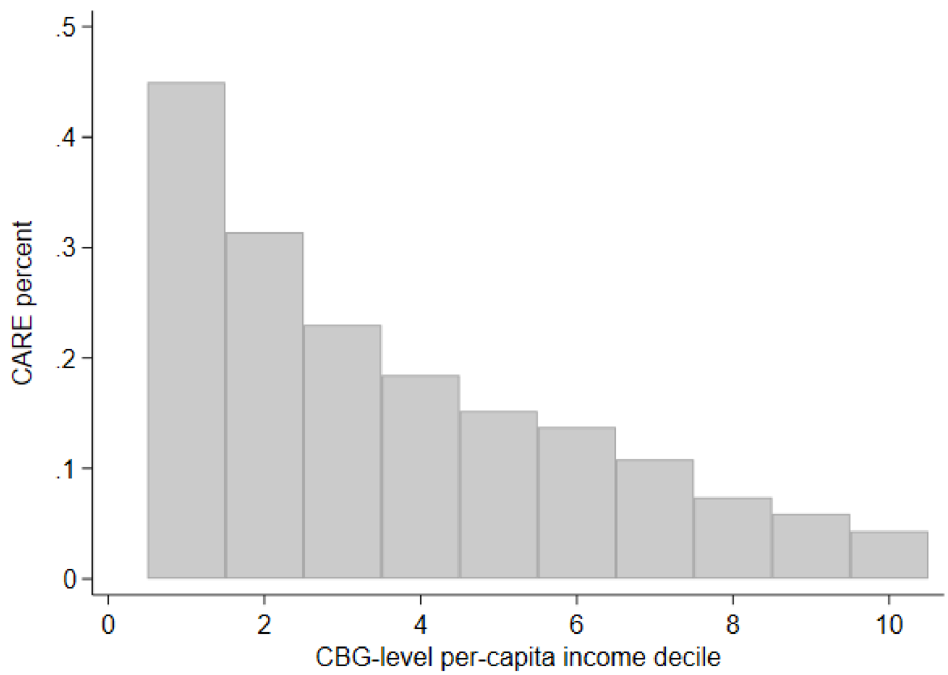

Note: In this figure, the horizontal axis is a measure of CBG-level per-capita income, while the vertical axis represents the proportion of households enrolled in CARE within a CBG.

Next, I present heterogeneity results under an alternative measure of income. To estimate heterogeneity in long-run elasticities, I first create an indicator for whether a household has ever enrolled in CARE over the full sample period as a proxy for income. I then interact this indicator with the specification shown in Equations 2, comparing price responsiveness across the two groups, as shown below:

$$\begin{aligned} ln(z_{it}) = \beta_0 + \sum_{j=0}^{1} \Big[\beta_{1j} \widehat{ln(p_{it}) \mathbb{1}\{LI_i = j\}} + \beta_{2j} \mathbb{1}\{LI_i = j\} \\ + \beta_{3j} \mathrm{d}_i \mathbb{1}\{LI_i = j\} + \beta_{4j} \mathrm{d}_i \mathrm{Hi}_i \mathbb{1}\{LI_i = j\} \Big] + \gamma_{ct} + \eta_{es} + \epsilon_{it} \end{aligned}$$

where $ln(p_{it})\mathbb{1}\{LI_i = j\}$ is instrumented for with $\mathrm{Hi}_i \mathbb{1}\{LI_i = j\}$ and $\mathbb{1}\{LI_i = j\}$ is an indicator function for whether a household is low-income or not (as defined by enrolling in CARE at any point during the sample period).

Results of this specification are shown in Table A12.

Table A12: Long run elasticities by CARE

	CARE	non-CARE
Log Average Price	-3.31***	-2.02**
	(0.80)	(0.84)
Distance to border (100m)	-0.0081	-0.015***
	(0.0065)	(0.0030)
Hi x Distance	0.019**	0.030***
	(0.0060)	(0.0065)
Observations	2,160,961	5,155,241
F	67	139

*Note: This table shows an instrumental variable regressions by income of log consumption. Income is proxied by whether a customer has ever enrolled in CARE. Fixed effects include CBG-by-month and electric-heat-by-season. Standard errors are clustered by baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

I then do the same in the short run, where I interact Equation 6 with an indicator for CARE:

$$\Delta\ln(c_{it}) = \beta_0 + \sum_{j=0}^{1}\left[\beta_{1j}\widehat{\Delta\ln(P_{it})\mathbb{1}\{LI_i = j\}} + \beta_{2j}\mathbb{1}\{LI_i = j\}\right]$$
$$+ f_t(c_{i,t-6}) + \gamma_{ct} + \lambda_{es} + \eta_{it}$$

where $\Delta\ln(P_{it})\mathbb{1}\{LI_i = j\}$ is instrumented for with $\Delta\ln(P_{it})^I\mathbb{1}\{LI_i = j\}$. The results of this specification are shown in Table A13.

Table A13: Short-run price elasticity by CARE

	CARE	nonCARE
	(1)	(2)
$\Delta\ln(AP_{it})$	-0.080***	-0.24***
	(0.0038)	(0.0081)
Observations	1,702,801	8,027,799
F	37131	21781

*Note: Across all columns, the dependent variable is $\Delta ln(c_{it})$. Fixed effects include CBG-by-month and 6-month-lagged consumption deciles. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

Finally, I show estimates for per-capita income, home ownership, age of housing stock, race, and within-CBG income inequality (as measured by a Gini coefficient).

Figure A9: Long-run estimates by CBG-level per-capita income

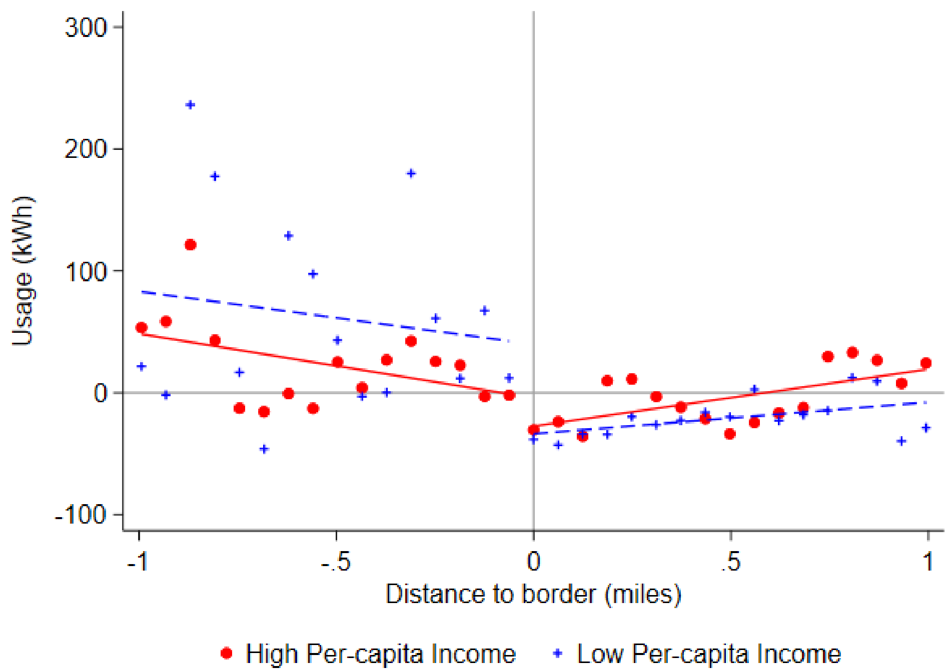

Figure A10: Long-run estimates by CBG-level home ownership rate

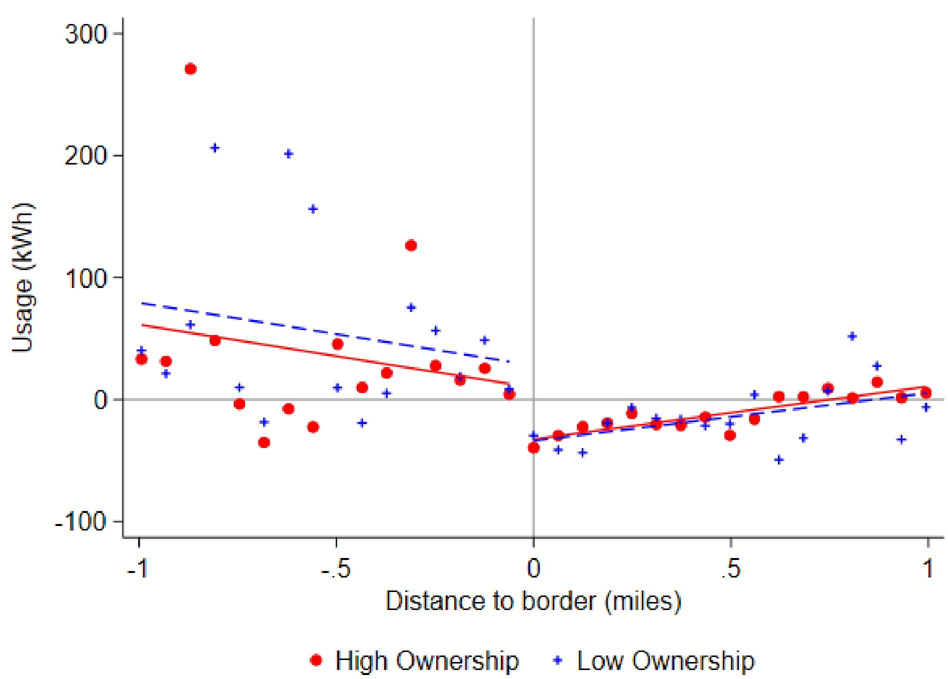

Figure A11: Long-run estimates by CBG-level age of housing stock

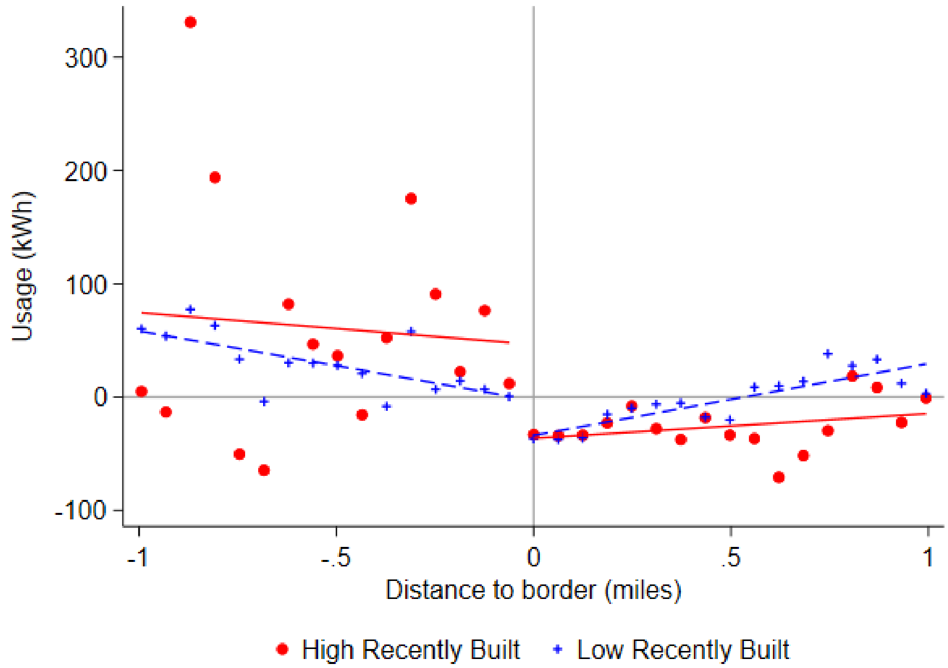

Figure A12: Long-run estimates by CBG-level race

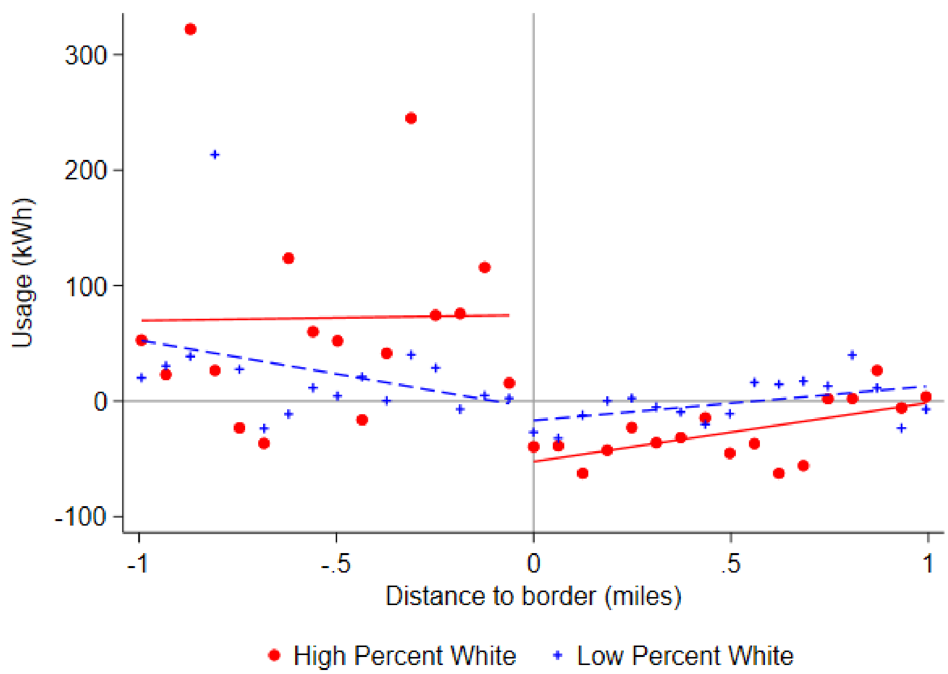

Figure A13: Long-run estimates by CBG-level Gini coefficient

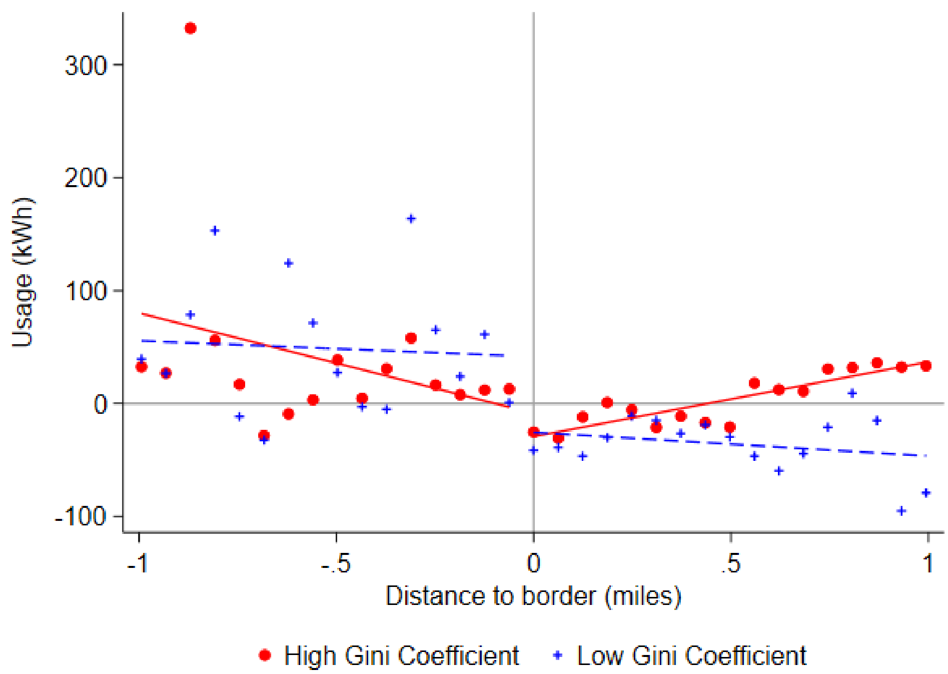

A.7 Medium-run responses to prices

Here, I build on the short-run specifications to estimate elasticities in the medium run. In my primary medium-run specification, I refer to the medium run as a time horizon of four years. Currently, there is little existing work in the literature on medium-run elasticities, especially in a quasi-experimental setting. Deryugina, MacKay and Reif (2019) find that customers are more responsive in the medium-run than in the short-run. This is is consistent with a number of studies using aggregated state-level data that similarly find that consumption responses to price build over time. In this section, I estimate how four-year consumption differences can be attributed to price changes that occur within that four-year period.

There are several different channels through which customers might respond to prices. After observing a change in price, consumers might respond in the short run by reducing their consumption of certain appliances – for example, a consumer might turn off their lights more frequently. If this short-run behavior becomes a habit for the consumer, we might continue to see this response carry through to the medium-run. However, customers may also respond by changing their investment of durable goods, such as energy efficient appliances, electric heat, or solar panels. We should expect that durable good adoption will impact consumption in both the short-run and the medium-run. These two channels – habit formation of conservation behaviors and durable good adoption – are the primary channels through which past prices can impact current consumption.

In order to estimate medium-run elasticities in this setting, I extend Ito's approach using time lags. Now, the dependent variable is the difference in consumption between the contemporaneous period and four years prior for a given household. I include the full price path as right-hand side variables, with a series of annual price differences within a four-year window as the primary covariates of interest. I used two-stage least squares, with four endogenous variables and four instruments:

$$\text{First stage: } \Delta\ln(MP_{i,t,l}) = \ln(MP_{i,t-12l}) - \ln(MP_{i,t-12(l+1)}) \text{ for each } l \in 0,1,2,3 \tag{12}$$

$$\text{Second stage: } \Delta\ln(c_{i,t,t-48}) = \sum_{l=0}^{3} \beta_l \, \Delta\ln(\widehat{MP}_{i,t,l}) + f_t(c_{i,t-60}) + \gamma_{ct} + \eta_{it} \tag{13}$$

where $\Delta\ln(c_{i,t,t-48}) = \ln(c_{i,t}) - \ln(c_{i,t-48})$ and $f_t(c_{i,t-60})$ is a set of dummy variables determined by the percentile of consumption in period $t-60$.

As in the short-run specifications, each price difference is endogenous to consumption due to the nature of increasing block pricing. Again, I use simulated instruments to solve this issue. For each endogenous price covariate, an associated simulated instrument is included.

In the short-run specifications, consumption levels from period $t-6$ were used in the instrument. Here, however, consumption in period $t-6$ is endogenous to the price differences included as covariates. Instead, in the medium-run specifications, consumption levels from period $t-60$ (one year prior to the first included price period) are used. This ensures that the instrument isolates exogenous changes in the price schedule and eliminates all endogenous price variation driven by consumption changes. Note that this specification includes only utility accounts continuously present in the sample over the course of five years (months ranging from t to $t-60$). Any customers who move over the course of this period are dropped from the sample. Hence, the external validity of these medium-run estimates is limited to consumers who are fairly stable and live in a single location for an extended period of time.

Table A14: Dynamic medium-run average variable price elasticities

	kWh
	(1)
$\Delta\ln(AP\_it)$	-0.18***
	(0.044)
$\Delta\ln(AP\_i,t,1)$	-0.19***
	(0.042)
$\Delta\ln(AP\_i,t,2)$	-0.13***
	(0.030)
$\Delta\ln(AP\_i,t,3)$	-0.12***
	(0.034)
Observations	2606624
F	194.9

*Note: Fixed effects include CBG-by-month and consumption deciles from the period twelve months prior to the initial price period. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

It's important to note that the empirical setting in this paper is quite different than in past work, including Deryugina, MacKay and Reif (2019). Deryugina et al. leveraged on a one-time change in prices and followed customers' demand levels over time. Here, price schedule fluctuations frequently occur and impact customers differently depending on their baseline territories and underlying consumption levels. As such, the interpretation of estimates is different in this setting: while elasticity estimates in Deryugina et al. should be interpreted as a consumption response to a single permanent change in price, the estimates in this paper tell us how to attribute changes in consumption to changes in price over the relevant period. When consumption changes over a four year period, how much of that change should be attributed to price changes in each year? Examination of each coefficient in the regression demonstrates how elasticities evolve over time.

Results of these medium-run regressions over a four-year period are shown in Table A14. These results demonstrate that responses to prices last over the course of several years, indicating that habits and/or durable good adoption play a vital role. When including past price periods, customers are similarly responsive to short-run fluctuations in price, with a price elasticity of -0.18. This elasticity stays close to -0.2 through two years, before fading towards -0.1 by the fourth year.

In the Table A15, I also estimate medium-run elasticities over an eight-year period. Note that this sample is even more highly selected to include only customers who do not move over a nine year period within my twelve year sample. Again, the external validity of these estimates is restricted only to consumers who live in a single location for an extended period of time – in this case, nine years.

These results are consistent with a combination of a short-run transient behavioral responses and significant durable good investment. After price fluctuations, consumers respond by changing their consumption. However, customers may also respond to price changes by investing in durable goods, which last for the duration of the sample. As a result, they still demonstrate responsiveness to price changes that occurred in more distant periods – in this case, three to four years prior to the contemporaneous period.

Table A15: Dynamic medium-run price elasticities - all average variable prices (8 year stable sample)

	12
	(1)
$\Delta\ln(AP\_it)$	-0.072
	(0.059)
$\Delta\ln(AP\_i,t,1)$	-0.14**
	(0.058)
$\Delta\ln(AP\_i,t,2)$	-0.030
	(0.053)
$\Delta\ln(AP\_i,t,3)$	0.048
	(0.056)
$\Delta\ln(AP\_i,t,4)$	0.027
	(0.068)
$\Delta\ln(AP\_i,t,5)$	-0.0088
	(0.056)
$\Delta\ln(AP\_i,t,6)$	-0.058
	(0.061)
$\Delta\ln(AP\_i,t,7)$	-0.21**
	(0.088)
Observations	955837
F	62.6

*Note: Fixed effects include CBG-by-month and consumption deciles from the period twelve months prior to the initial price period. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*

In addition, I estimate heterogeneity in medium-run elasticities across income, again using CARE enrollment as a proxy for income, as shown in Table A16. Consistent with the short-run results, non-CARE (higher-income) consumers are more responsive to changes in their electricity prices in all periods.Once again, this suggests that investment in durable goods may play a substantial role in how consumers respond to energy prices.

Furthermore, these results suggest that consumers' responses to price changes may accumulate over time as consumers continue to respond to prices from four years prior. However, the type of dynamic two-way fixed effects panel regressions shown to this point only allow for evaluation up to the length of the observed sample, and may miss important mechanisms, such as durable goods investments in new homes. These are better captured by long-run elasticities, as described in Section 4.1.

Table A16: Dynamic medium-run average variable price elasticities by CARE

	CARE	nonCARE
	(1)	(2)
$\Delta\ln(AP\_it)$	-0.11**	-0.21***
	(0.053)	(0.050)
$\Delta\ln(AP\_i,t,1)$	-0.12**	-0.23***
	(0.052)	(0.049)
$\Delta\ln(AP\_i,t,2)$	0.0038	-0.18***
	(0.043)	(0.034)
$\Delta\ln(AP\_i,t,3)$	-0.056	-0.13***
	(0.052)	(0.037)
Observations	413612	2192841
F	166.4	200.7

*Note: Fixed effects include CBG-by-month and consumption deciles from the period halfway between the present period and the lagged consumption period. Standard errors are clustered by CBG-baseline territory and by month of sample. ***, **, * indicate significance at the 1% and 5% and 10% level, respectively.*