Quantifying Heterogeneity in the Price Elasticity of Residential Natural Gas

Edward Rubin Maximilian Auffhammer*

November 2, 2022

**Abstract**

We exploit a spatial discontinuity in two natural gas utilities' service territory—combined with variation in their block-rate pricing structure and a difference in how prima facie determined wholesale prices are deferentially passed though to consumers—to identify average, seasonal, and income-specific own-price elasticities of residential natural gas demand. We estimate an average elasticity ranging from 0.15–0.19 depending on the measure of price used. We further estimate that this elasticity varies substantially across seasons, income groups, and their interaction. We find no significant difference in consumers' responses to average versus marginal prices.

**JEL Codes:** Q41, Q48, D12, H23, Q31
**Keywords:** Heterogeneity, Elasticity of Demand, Distribution, Natural Gas, Big Data

*We thank the Pacific Gas and Electric Company and the Southern California Gas Company for making their billing data available. We thank Eric Zou, Reed Walker, Glen Waddell, Sofia Berto Villa-Boas, Sarahelen Thompson, James Sallee, William Pizer, Karen Palmer, Nathan Miller, Solomon Hsiang, Meredith Fowlie, Lucas Davis, Jonathan Davis, Richard Carson, Severin Borenstein, Peter Berck, Abhay Aneja, and seminar/workshop participants at UC Berkeley ARE, Cornell University, the Luleå Department of Economics, the University of British Columbia, UC Davis, UC San Diego, the University of Oregon, Yale University, the AERE Annual Summer Conference, Camp Resources, Duke's Energy Data Analytics Symposium, the Heartland Environmental and Resource Economics Workshop, the Federal Reserve Bank of Dallas, Renmin University, and the WCERE for valuable feedback. Auffhammer gratefully acknowledges support from the Julian Simon Fellowship at Property and Environment Research Center (PERC). All remaining errors are the authors'. Rubin: Assistant Professor of Economics, University of Oregon (edwardr@uoregon.edu) Auffhammer: George Pardee Professor of Sustainable Development, UC Berkeley. (auffhammer@berkeley.edu).

# 1 Introduction

Identifying and estimating structural parameters in demand and supply systems has greatly progressed since Ezekiel (1925) and Working (1927) laid out the fundamental framework of simultaneity (Wright 1928; Koopmans 1945; Girshick and Haavelmo 1947; Angrist and Krueger 2001; Stock and Trebbi 2003). In the field of resource economics, much of the focus has been on identifying a price elasticity of demand in the presence of block-rate pricing (*e.g.*, Hanemann (1984), Olmstead, Hanemann, and Stavins (2007), Ito (2014)), with some recent work focusing on heterogeneity in price response (*e.g.*, El-Khattabi et al. (2021), Hahn and Metcalfe (2021)).

Estimates of the price elasticity of demand are of first-order importance—not only to consumers and firms but also to regulators—both in the short and long run. First, in market settings where externalities exist, the welfare consequences of a policy intervention (*e.g.*, a carbon tax) depend on the long-run price elasticity. Second, in settings with significant fixed costs—as is the case in the energy and natural-resource sectors—regulators choose between recovering these fixed costs through fixed charges or through volumetric fees. Once again, the welfare implications of this decision hinge upon the elasticity of demand. The short-run elasticity of natural gas demand is both important for price-regulated utilities' procurement planning, but also from a short-run public policy perspective. For example, the optimal response to the supply crisis in Europe's natural gas supplies due to the war in Ukraine hinges critically on the magnitude of the short-run elasticity.

Further, customers' responsiveness to prices can vary both across customers' characteristics (*e.g.*, income) and time. Because the welfare effects of policy, regulation, and taxation depend upon customers' price responsiveness, estimating and understanding these heterogeneities is important, especially as policy making in the energy space is increasingly concerned with the distributional and justice consequences of policy choices.

Economists have long studied time-varying pricing. Examples of settings where such time-varying pricing may actually improve welfare abound— especially in the areas of electricity demand and transportation. Designing optimal and fair time-varying pricing structures requires causal estimates of group- and time-specific price elasticities—a task that is often difficult due to the data and exogenous variation required to identify temporally and cross-sectionally heterogeneous consumption responses.

This paper is an effort to fill this gap for the main source of residential energy in North America: residential natural gas. In 2016, natural gas surpassed coal as the main source of energy for electricity generation in the United States. Furthermore, half of US residences use natural gas as their main heating fuel (U.S. Energy Information Administration 2016b). Depending on the severity of the winter, US residential consumers spend 50-80 billion dollars per year on natural gas. The average household spends about as much money on natural gas as it spends on water (BLS 2017; U.S. Energy Information Administration 2016b).

We exploit access to the universe of natural gas bills for the two biggest investor-owned utilities in California to econometrically identify the own-price elasticity of demand by zooming in on areas in which natural gas prices discontinuously change due to arbitrary, historical network boundaries. Our approach uses the same spatial discontinuity approach as Ito (2014). However, the natural gas market setting differs in two major ways from the market for residential electricity. First, natural gas is priced on a two-tier block-rate structure. In the case of electricity, the bottom tier is fixed and variation arrives in the form of up to *seven*-tier changes over time from utility-specific rate cases filed with the regulator. Second, utilities are allowed to shift both tiers month by month in order to pass though variation in lagged procurement costs driven by changes in the wholesale cost of gas. Electricity prices on the other hand are infrequently updated as a consequence of rate cases. Finally, utilities rely on different sources of natural gas for their procurement.

This setting requires that we build on Ito (2014) and augment his identification strategy, as the two-tier pricing structure provides very limited variation as both the width of the blocks as well as the price difference between them vary infrequently. In addition to the geographic boundary, we construct a price instrument by interacting this spatial discontinuity with predetermined exogenous variation in the upstream spot market (Henry Hub) for natural gas—in order to introduce exogenous time-series variation in prices. We use the fact, and confirm in a first stage, that there is differential pass through of this exogenous to the consumer lagged wholesale price, which provides additional variation on the two sides of the utility territories. Taken together, the spatial discontinuity, differential pass through and standard variation in bill beginning and end dates allow us to provide precise and causal estimates of the own-price elasticity of demand.

In what follows, we provide a well-identified medium-run (two month) "average price elasticity", which we then decompose into the price elasticity of demand by season, by income, and by their interaction. Our results demonstrate that consumers are much more price sensitive in cool months (when households heat), and this difference is economically and statistically significant. We show suggestive evidence that lower-income households are more price sensitive than their more-affluent counterparts with sizable differences in the point estimates. However, the coefficients are not statistically different across these groups. Finally, we show that the largest point estimates for the price elasticity of demand are for lower-income households in cool months. On their own, these heterogeneous elasticities are key to understanding an important component of household behavior—they paint a very different picture than that of a single, pooled elasticity. For instance, pooling across households and seasons, our estimate of the marginal price elasticity of demand for residential natural gas is approximately −0.17, but our season-by-income estimates range from unresponsive (affluent households in warm, non-heating months) to −0.46 (lower-income households in cool, heating months). Our estimates are robust to an extensive barrage of specification tests and different applications and sources of identifying variation. We also provide estimates for different measures of price. We find no evidence that the response to marginal price is different from that to average price, which is likely due to the relatively small size of the "step" increase in the natural gas tier.

While quantifying heterogeneity is important on its own, in this setting our results open a new avenue for a welfare-enhancing and progressive policy change, which can easily be implemented by state utility regulators. Specifically, if utilities and governments shift small program fees currently charged on a volumetric basis throughout the year to more inelastic periods of the year (non-heating months in our context), consumers would suffer less deadweight loss—and the reduction in deadweight loss is potentially higher for poorer households. Such a switch is especially plausible in the context of natural gas, as many utilities and governments already apply seasonal changes in the pricing structure of residential natural gas to respond to seasonal changes in the demand for heating.

## 2 Institutional setting

A basic understanding of the institutional and physical setup of the natural gas sector in the United States is helpful for understanding our identification strategy. This market is commonly divided into four segments: (1) production and processing, (2) transportation, (3) storage, and (4) local distribution companies (LDCs). Figure 1 illustrates the basic institutional organization of the natural gas industry.[1] The four segments we discuss below roughly follow Figure 1 except that they exclude end users (those users who only consume natural gas) and the liquid natural gas import/export-based segments of the market. While we focus on the behavior of residential natural gas consumers, part of our identification strategy relies upon a basic understanding of the wider industry—specifically in understanding which factors may shift supply without affecting demand.

### 2.1 Market segments

**Production and processing** Natural gas enters the market at the wellhead, where it is produced and first sold (Brown and Yücel 1993). Some wells produce only natural gas, while other wells produce natural gas in addition to crude oil. The raw product then moves from wellheads to processors. Processors remove impurities and separate the raw product into multiple commodities (Levine, Carpenter, and Thapa 2014).

**Transportation** High-pressure pipelines transport processed natural gas from production and processing areas to both intermediate users (storage facilities, processors, LDCs) and

1. We include liquid natural gas (LNG) in the figure for completeness, but liquid natural gas does not play a large role in the natural gas market in the United States: LNG imports currently account for less than one percent of natural gas imports and accounted for three percent of imports at their peak in 2007 (Levine, Carpenter, and Thapa 2014). For this reason, we omit LNG for the rest of the paper.

final users (electricity generators, industrial users, commercial users, and residential users).[2] Extensive spot markets and futures markets sit at the major hubs along this pipeline. Notably, Louisiana's Henry Hub connects to 13 intrastate and interstate pipelines. The Henry Hub is the designated delivery point for the New York Mercantile Exchange's natural gas futures contracts, and the Henry Hub price is generally regarded as one nationally relevant price (Levine, Carpenter, and Thapa 2014). Figure 2 depicts the Henry Hub spot price from 1997 through 2016. In general, transportation costs represent a substantial percentage of natural gas prices.[3] Thus, the natural gas transportation network creates a nationally integrated market and simultaneously contributes to a sizable portion of the prices paid by natural gas end users.

**Storage** Storage plays a major role in several parts of the natural gas market, but all parties store mainly for the same reason: volatility within the market. Due to its major role in heating and electricity production, natural gas demand is strongly driven by weather and can be unpredictable in the short run. To combat price volatility and to be able to meet peak demand, both local distribution companies and large natural gas consumers store gas underground. Producers utilize storage to smooth production.

**Local distribution companies** Local distribution companies' primarily distribute natural gas to their contracted end users—industrial, residential, and commercial consumers of natural gas. To accomplish this task, LDCs purchase natural gas through both spot markets and longer-term contracts. In addition, LDCs own and operate their own pipeline and storage networks. To cover the fixed costs involved in their pipelines, storage, and administration, LDCs often utilize a combination of two-part tariffs and multi-tiered pricing regimes—though some utilities fold all of their costs into their volumetric pricing. State utility commissions (*e.g.*, the California Public Utilities Commission) regulate LDCs' price regimes, allowing the LDCs to earn a regulated rate of return (Brown and Yücel 1993; Davis and Muehlegger 2010; Levine, Carpenter, and Thapa 2014).

## 2.2 Natural gas pricing in California

The California Public Utilities Commission (CPUC) regulates the two utilities from which we draw data: Pacific Gas and Electric Company (PG&E) and Southern California Gas Company (SoCalGas). Because we analyze residential natural gas consumers' responses to natural-gas

2. Figure A1 maps this pipeline network for the continental United States. Private companies own and operate segments of the pipeline; these pipeline companies' rates are regulated at the state level and the national level (Levine, Carpenter, and Thapa 2014).

3. According to Levine, Carpenter, and Thapa, in 2011–2012, 72 percent of consumers' average heating costs originated in "transmission and distribution charges". They also note that in 2007–2008 "transmission and distribution charges" accounted for 41 percent of consumers' average heating costs.

retail prices, the most relevant regulations facing PG&E and SoCalGas are the CPUC's price and quantity regulations.

For households served by PG&E or SoCalGas, a bill depends upon five elements:

1. The **two-tiered price schedule** set by the utility
2. The **total volume of natural gas consumed** during the billing period
3. The **season** (*summer* or *winter*) in which the bill occurs
4. The **climate zone** into which the household's physical location falls, as defined by the California Energy Commission (CEC) (see Figure A3),
5. The household's **CARE (California Alternate Rates for Energy) status**, which provides a 20% discount to households whose income falls below pre-specified income thresholds or are enrolled in certain public assistance programs.

Figure A4 provides an example of a typical residential natural gas bill from PG&E.

Both PG&E and SoCalGas utilize two-tiered pricing regimes. The California Energy Commission divides California into 16 climate zones in which households' needs for heating should be relatively homogeneous (California Energy Commission 2015, 2017; Pacific Gas and Electric Company 2016). The utilities also divide the year into heating (winter) and non-heating (summer) seasons. Based upon a household's climate zone (determined by the household's location) and the season, the CPUC determines a volume of natural gas that should be adequate for heating during the course of one day. This volume of natural gas is called the household's *daily allowance*. Multiplying the household's daily allowance by the number of days in the billing period gives the household's *total allowance* for the bill. For each unit (*therm*[4]) of natural gas up to the bill's total allowance, the household pays the first tier's per-unit price (baseline price). For each unit of gas above the household's total allowance, the household pays the second tier's per-unit price (excess price). Figure 3a illustrates an example of the two-tier block-pricing regime used by PG&E and SoCalGas. Figure 3b depicts how residential consumers' (daily) tier-one allowances vary through time within a given climate zone (PG&E's climate zone *R* and SoCalGas's climate zone *1*).

In terms of variation employed later in the paper, it is important to note that consumers' billing periods do not perfectly align with calendar months. The beginning date of the bill is fully dependent on the day of the month the account was opened, which for most households does not coincide with the first of the month. We have checked for bunching of bill beginning days and find a pretty uniform distribution across days of the month. However, PG&E's and SoCalGas's price changes do align with calendar months. The two utilities deal with this misalignment of billing periods and price regimes slightly differently. PG&E calculates individual bills for each calendar month under the assumption that consumption is constant throughout the billing period. SoCalGas calculates a single bill using time-weighted average

4. The utilities work in units of volume called *therms*. One therm is equal to 100,000 Btu (U.S. Energy Information Administration 2016c).

prices (averaging across the different price regimes). These methods are equivalent under a single linear price but differ under the actual multi-tiered price regimes. Please see the *Calculating bills* section in the appendix for more detail.

If retail prices were not updated month to month, as is the case for electricity for most of California's customers, we could not write this paper. We use the fact that each month, the natural gas utilities are permitted by regulators to update their block price schedules. The absolute difference between the first-tier price and the second-tier price also varies but tends to remain constant for several months.[5] These monthly price changes allow the utilities to charge customers at rates that reflect the lagged prevailing price of natural gas in the wholesale markets they procure gas from. In fact, the utilities tie their price updates to their costs—thus linking residential rates to *lagged* spot market prices.[6] If the utilities wish to change the formula with which their prices are tied to market prices and other costs, they must receive authorization following a review process with CPUC. Figure 5a illustrates these monthly price-regime changes and the fairly fixed step between the two tiers. Figure 5b depicts the correlation between the utilities' baseline (first-tier) prices and the spot market price of natural gas at the Henry Hub, which is significant.

At the insistence of a reviewer, we thankfully engaged in several conversations with natural gas experts at the CPUC, who reaffirmed that price variation month to month is significant and that what gets passed through (*e.g.*, how much of Henry Hub price variation) may differ for the two utilities. Perhaps the biggest factor is the regional difference in prices for the gas they buy, as the two utilities procurement portfolios differ (PG&E procures a lot of supply from Canadian basins; SoCalGas procures supply from basins located in the Southwest of the US). Secondly, each investor-owned utility also engages in market practices to decrease core commodity costs, which include purchasing different kinds of gas contracts, park and loan programs, price arbitrage, and selling excess gas into the market. Many of their contracts are based on the first-of-month Henry Hub index price, but not all of them. The Henry Hub price hence plays a significant role for both utilities' procurement and ultimate pricing strategy, yet the degree to which Henry Hub matters varies across the two suppliers. In fact, for our identification strategy, the fact that Henry Hub price is passed through differentially is a key source of variation we use for identification of the price elasticity of demand.

Finally, a household's participation in the CARE (California Alternate Rates for Energy) program affects the prices that the household faces. Customers that are enrolled in the CARE program receive a 30-35 percent discount on their electric bill and a 20-percent discount on their natural gas bill. Households qualify for CARE by either meeting low-income qualifications or receiving benefits from one of several state or federal assistance programs (*e.g.*, Medi-Cal or the National School Lunch Program) (Southern California Gas Company 2016). CARE prices are 80 percent of standard prices at both tiers. In a recent paper, Hahn and Met-

5. The utilities differ in the frequencies at which they change this absolute difference: PG&E adjusts the distance between the two tiers' price much more frequently than SoCalGas.

6. The utilities report their weighted average costs of gas to the CPUC.

calfe (2021) exploit the fact that some household's enrollment lapses and use a randomized encouragement design to estimate the price elasticity of demand for these customers, with their central estimate being −0.21. Our paper trades in the excellent experimental design for an ex post observational design in order to identify a residential price elasticity for CARE and non-CARE customers. Throughout the paper, we will often refer to the elasticity of the CARE customers as "low-income", which implies that either the income threshold of the program is met or the household is enrolled in one of the eligible public assistance programs. One could be concerned about whether households that are eligible are actually signing up for the CARE program. In their appendix table A2, (Hahn and Metcalfe 2021) estimate the share of eligible customers that are signed up for CARE. For the two utilities in our paper they estimate that 84% of eligible customers do indeed sign up for CARE. This is encouraging. To be clear: when we use the phrase "low income", it explicitly refers to households who have enrolled in CARE.

# 3 Data

## 3.1 Natural gas billing data

Our billing data come from two major utilities in California: Pacific Gas and Electric Company (PG&E) and Southern California Gas Company (SoCalGas). The PG&E data cover all residential natural gas bills in PG&E's territory from January 2003 through December 2014. The SoCalGas data cover residential natural gas bills from May 2010 through September 2015. Thus, the two utilities' data overlap from May 2010 through December 2014. After excluding zip codes with fewer than 50 households, PG&E's service area covers 597 5-digit zip codes, with a total of 5,888,276 households and 180,663,705 bills. After excluding zip codes with fewer than 50 households, SoCalGas's service area covers 611 5-digit zip codes, with a total of 2,526,503 households and 95,335,393 bills.

The left panel of Figure 4 depicts PG&E's and SoCalGas's service areas at the 5-digit zip code level. Table 2 provides a brief summary of the billing data with regard to the numbers of bills, households, zip codes, and monetary values of the bills. Tables 2 and 3 summarize prices, quantities, and other variables of interest—pooling across all observations and also splitting the data by season or CARE status. Both tables summarize the full dataset—all zip codes across both utilities– as well as the subset of the data used in our preferred estimation which limits the sample to all 5-digit zip codes served by both utilities. We describe this subset in further detail below.

The utilities' billing data are at the household-bill level: a single row of the dataset represents a single billing period for a given household. Table A17 describes the variables (columns) in this dataset. We follow the natural gas utilities' convention in defining a household (or customer) as the interaction between a unique utility account and a unique physical location identifier. It is important to note that there is variation in the beginning and end dates (and hence bill lengths) across households, which depends on the day of the month the account was opened.

We also scraped historical data on pricing from the two utilities off their websites as well as CPUC documentation. As described above, these pricing data include (1) each utility's monthly two-tier pricing regime and (2) the daily allowance for each climate zone during each season. After joining these pricing data to the households' billing data, we are able to determine both the marginal price and average price (and average marginal price) for each bill received by each household. We checked whether we could replicate the bill totals for each bill using our scraped data and declare success.

### 3.2 Wholesale Prices

Daily data on wholesale gas prices come from (U.S. Energy Information Administration 2016d) and are reported in dollars per million BTU. Figure 2 displays the time series of prices over the time of our sample. There is significant variation day to day, which we exploit in the first stage of our analysis that follows.

### 3.3 Weather data

Data on daily weather observations originate from the PRISM project at Oregon State University (PRISM Climate Group 2004). We match these local, daily weather data to the household consumption data at the day by 5-digit-zip-code level. The PRISM dataset contains daily gridded maximum and minimum temperatures for the continental United States at a grid cell resolution of roughly 2.5 miles (4 km). Figure A5 maps a single day of average temperature from the PRISM data for the continental United States.

We observe these daily data for California from 1980–2015. In order to match the weather grids to zip codes, we obtained a GIS layer of zip codes from ESRI (Esri 2017), which is based on US Postal Service delivery routes for 2013. For small zip codes not identified by the shape file, we purchased the centroid latitude and longitude of these zip codes from a private vendor (zip-codes.com). We matched the PRISM grids to the zip code shapes and averaged the daily temperature data across multiple grids within each zip code for each day. For zip codes identified as a point, we use the daily weather observation in the grid at that point. This exercise results in a complete daily record of minimum and maximum temperatures—as well as precipitation—at the zip-code level from 1980–2015.

## 4 Empirical strategy

We now describe our empirical strategy to identify the price elasticity of demand for residential natural gas consumers. First, we present the basic estimating equation that motivates the paper's results. Next, we discuss the inherent challenges to identification in this setting. We then discuss potential solutions to these challenges and detail which of these solutions are feasible in our specific setting.

## 4.1 Estimating equation

The relationship at the heart of our elasticity estimates is

$$\log(q_{i,t}) = \eta \log(p_{i,t}) + \lambda_{i,t} + \varepsilon_{i,t} \tag{1}$$

where $i$ and $t$ index household and time, $q$ denotes quantity demanded, and $p$ denotes price. Rather than choosing a specific type of price (*e.g.*, marginal), we present results for five variants of price. These five types of price include the price that classical economic theory deems relevant—the marginal price—in addition to average price, average marginal, baseline (first-tier) price, and *simulated* marginal price (defined and discussed below). We define *average marginal price* as the quantity-weighted marginal price paid by a customer during her billing period. A single billing period may contain multiple marginal prices, as the block rates are updated on the first of each month and most bills do not begin and end on the first and last day of each month. *Average marginal* price does not include fixed charges, while *average* price does. In the results section, we also discuss which lag of price is most salient to consumers. The term $\lambda_{i,t}$ represents household fixed effects, time-based fixed effects, and/or household-by-time fixed effects—depending on the specification. Our main specification uses household fixed effects and city by month-of-sample fixed effects (*e.g.*, Fresno in January 2010; also called city by year by month). A causally identified estimate of $\eta$ yields the own-price elasticity of demand.

## 4.2 Challenges to identification

In settings where researchers must rely on highly aggregated data on prices and quantity there is endogeneity resulting from the simultaneous determination of price and quantity that resolves the supply and demand equilibrium (*e.g.*, Woolridge 2009), which we overcome by using household level consumption data and by showing that households respond to lagged prices. However, assuming that the price households face in a given billing period is exogenous is wrong, since the two-tiered price schedule within California's natural gas market makes the marginal price a weakly increasing, monotonic function of quantity. For the same reason, average price is also a function of quantity. Thus, when a household consumes more, its marginal and average prices mechanically increase. In terms of identifying the price elasticity of demand, this price variation is bad variation: the marginal price that a household faces is endogenous because it is correlated with unobserved demand shocks (Ito 2014). This bias resulting from marginal and average prices being a function of quantity results in upwardly biased estimates of demand elasticities, possibly yielding estimates that suggest upward-sloping demand curves.

Table 5 demonstrates the consequences of failing to address these challenges to identification by estimating the price elasticity of demand—$\eta$ in equation 1 via ordinary least squares (OLS) using marginal price (columns 1–3). Each column varies the set of controls. Column (1) applies the simplest set of controls to a full 5% sample of all bills in the sample. The "identification strategy" presented in Table 5 makes no attempt to correct for endogeneity outside of a fairly rich set of fixed effects—household fixed effects and city by month-of-sample fixed effects. All regressions control for within-bill heating degree days (HDDs) during the billing period.[7] Columns 2 and 3 subset the data to the ZIP codes containing our border-discontinuity, in which we keep all zip codes where the zip code receives natural gas from both PG&E and SoCalGas (discussed in detail below; also see Figure 4).

The fact that all marginal-price based elasticity estimates are positive (implying upward-sloping demand curves) suggest that the fact that price is a function of quantity is a first order concern here, which is not surprising.

### 4.3 Identification Strategy

Having confirmed that and OLS regression of quantity on price at the household level does not cleanly identify the own-price elasticity of demand in this setting, we now discuss our route for identifying the causal effect of price on quantity in our setting.

#### 4.3.1 Discontinuities

A common path toward identification in applied microeconomics involves finding relatively small geographic units that receive different exogenous treatments within the same time period. The assumption is that observable and unobservable characteristics and, more importantly, households' price responsiveness do not differ across the border, yet they are exposed to different supply side driven price changes, allowing for econometric identification. Arbitrary administrative boundaries that determine policies' catchment areas provide a popular tool in this context, *e.g.*, Dell (2010), Chen et al. (2013), and Ito (2014).

We exploit the boundary between PG&E and SoCalGas, as it offers arbitrary within-city (and within-zip code) variation in prices during a month. Specifically, the boundary between PG&E's and SoCalGas's natural gas service areas bisects eleven cities—in three clusters—in southern California: Arvin, Bakersfield, Fellows, Fresno, Del Ray, Fowler, Paso Robles, Selma, Taft, Tehachapi, and Templeton. The left panel of Figure 4 displays the two utilities' service areas throughout California (with sufficient coverage in our billing data). The right panel of Figure 4 zooms in on the eleven cities (39 zip codes) that PG&E and SoCalGas both serve. Within these eleven cities, PG&E serves all 39 zip codes, while SoCalGas serves 18 of the zip

7. The number of heating degrees in a day is equal to the difference between the day's average temperature and 65. Days with average temperatures above 65°F receive zero heating degrees. More formally, we calculate the number of heating degrees for day $t$ with mean temperature $\bar{T}_t$ (in °F) as $\text{HDD}_t = \mathbb{1}\left\{\bar{T}_t < 65\right\} \times \left(65 - \bar{T}_t\right)$. The HDDs variable above is thus $\text{HDDS} = \sum_t \text{HDD}_t/1000$.

codes. It is important to note that we know *that* a household is located within a given ZIP code, but we do not know *where* exactly, which prevents us from "zooming in" at the street level and directly comparing neighbors.

This identification strategy rests upon the assumption that households on one side of the utilities' border provide a valid control group for households on the other side of the border. In other words, households on both sides of the border are identical and experience the same demand shocks (*e.g.*, temperature or income shocks). Because the boundary mainly represents the extent of each utilities' underground distribution network and is unlikely to enter into households' preferences, the exogeneity of the boundary to household characteristics is likely to be valid, as extensively shown by (Ito 2014). The main threat to this identification strategy, as he discusses, is that utilities' networks correlate with geographic or neighborhood characteristics over which individuals have preferences'. However, we use household fixed effects, which absorb time invariant observable and unobservable differences across households.

Figure 5b suggests the generating distributions for the utilities' prices are quite similar (the standard deviations of the price series are 0.0940 and 0.1053 for PG&E and SoCalGas, respectively). In addition, Table 4 provides some limited evidence of balance across the utility border, comparing PG&E and SoCalGas households within season (summer or winter) and within income group. We remind the reader that we use income status synonymously with CARE status. Figure A7 provides evidence of the relationship between CARE status and income. Within a season-income group, the utilities' customers appear to consume similar volumes of natural gas, receive similar numbers of days per bill, receive similar allowances on the first tier, and face similar numbers of heating degree days. SoCalGas customers tend to receive slightly lower bills, but the difference is less than half of one standard deviation of total bill amount. Figure A2 illustrates the relevant natural-gas and electricity service areas. Notably, the border between PG&E's *natural-gas* service area and SoCalGas's *natural-gas* service area is covered entirely by PG&E's *electricity* service area. This overlap allows for neighboring households to receive natural gas from two different utilities, while only receiving electricity from PG&E.

Ito (2014) employs a similar strategy within the context of electricity consumption. However, there is at least one significant difference between the electricity and natural gas contexts that prevents us from entirely adopting Ito's identification strategy, which relies on discontinuities within electricity utilities' *seven-tier* pricing regime. By law, the electricity utilities in Ito's study cannot change the prices of their first two tiers—they must recover changes in their costs by moving tiers three through seven. In addition, electric utilities in California generally do not change consumer's prices each month—and prices do not change across all utilities at the same time. Thus, marginal prices in Ito's setting move differently depending upon a household's tier and utility after a rate case has passed. Ito argues that the residual variation—combining the spatial discontinuity with this pricing discontinuity and spatiotemporal fixed effects—is plausibly exogenous from demand shocks. In our natural gas setting we only have the benefit of two tiers and we lack significant variation in the absolute difference between the two tiers which powers Ito's results. Consequently, we need to supplement this utility-border-based discontinuity by exploiting an additional source of exogenous variation to power our estimates, which we discuss next.

### 4.3.2 Instrumental variables

Due to the way natural gas utilities are regulated in California, they are allowed to pass through procurement costs (*e.g.*, wholesale prices) from the *last* week of the *previous* month to consumers for the coming month. Hence variation in wholesale prices provides a significant source of variation in the prices that consumers face on both sides of the utility border. This would not be useful, if both utilities passed wholesale prices through using the same formula, but fortunately for us, they do not. We exploit the fact that wholesale prices from the previous month are (1) strongly correlated with the prices that the natural gas utilities charge their customers (the first stage), and (2) uncorrelated with residual shocks affecting consumers' demand as they are temporally disconnected (Angrist and Pischke 2009). Our specific instrument, which injects differential variation on each side of the border, is based upon the Henry Hub spot price for natural gas, which is one important source of cost for both utilities as we discussed above.

**Henry Hub spot price** Specifically, we instrument the prices that consumers face (*e.g.*, marginal price, average price, baseline price) with the average spot price at Louisiana's Henry Hub in the week preceding the change in prices, which matches the rule the utilities have to adopt. While the pass-through formula is not solely based on the Henry Hub price—it is the price of a variety of sources of gas—we interact the Henry Hub spot price with utility to allow the utilities to differentially pass through this wholesale price change. The Henry Hub spot price represents a nationally prevailing price for short-term natural gas contracts as the hub sits at the intersection of thirteen intrastate and interstate pipelines (U.S. Energy Information Administration 2016a).

We argue that the necessary exclusion restriction for this spot-price based instrument is plausibly valid. First, we interact the spot price instrument with utility. This interaction, conditional on city by month-of-sample fixed effects, means that the identifying variation in our instruments comes from the difference in how the two utilities' pass monthly Henry Hub spot-price shocks on to their consumers. As discussed in section 2, this is due to how the two utilities procure gas and the relative importance of the Henry Hub price in their procurement portfolio. The CPUC-approved pass-through further is assumed to not correlate with demand shocks within a city (or zip code) because of the disconnect in timing. Additional plausibly exogenous variation in residential prices across households and time within a given city and utility comes from the different days on which households' bills begin.[8]

It is important to stress how timing shuts down the much-feared classical simultaneity channel. As the spot price is temporally disconnected from the billing period the utilities' costs (and approved prices) are based upon spot prices that *precede* the billing period by several weeks. Thus, shocks that affect the Henry Hub spot price are distinct in time from shocks that affect natural gas demand—our fixed effects will absorb any of these shocks, so long as they do not differ across the utilities' border within a month.

The next important point that needs to be clarified is which lag of price is the salient one. We show that the most salient lag of price is the *second* lag of price, further disconnecting contemporaneous local demand shocks from market-level supply shocks two months prior.[9] We also control for the number of heating degree days (HDDs) in the household's zip code during the households' billing period. Because residential consumers primarily use natural gas in heating applications, controlling for HDDs further reduces the opportunity for local demand shock to affect national prices.

One final exclusion-restriction concern is that price variation in the spot market for natural gas may affect both residential natural-gas prices and residential electricity prices. In this scenario, we would not be able to separate the effect of an electricity price shock from a natural-gas price shock. However, in addition to the institutional setup that rules out this channel, Figure 7 supports our argument that (1) residential natural gas and electricity prices are uncorrelated in both levels and differences (across the utilities' border within a month of sample), and (2) variation in the residential price for electricity is uncorrelated with variation in the Henry Hub spot price of natural gas.[10] Therefore, we argue that the exclusion restriction is plausibly valid for our spot-price instrument.

Our identification strategy thus interacts the spatial discontinuity between PG&E's and SoCalGas's service areas with the Henry Hub spot price. Specifically, the identifying variation stems from the two utilities' divergent pass-through of the spot market price—differentially projecting variation in the natural gas spot market across a tenably arbitrary border between the two utilities. The border discontinuity ensures that we compare similar (neighboring) households. The spot-market pass-through instrument generates arbitrarily different exogenous prices for these neighboring households. Variation in the day-of-month on which a households bills begin provides additional plausibly exogenous variation at the household level—though our results do not depend upon this additional source of variation. Finally, the temporal separation between utilities' price changes (reflecting the upstream market) and consumers' demand decisions lends further credibility to this design's ability to separate supply shocks from demand shocks.

8. Our results are robust to eliminating this additional source of variation, *i.e.*, treating all households within a city-utility to the same price within a month of sample.

9. See Tables 7, A3–A6 for the second-stage results comparing consumers' responses to various lags in price.

10. This observation also draws upon Figure 5b.

**Spot price IV, first stage** Panel A of Table 6 provides the first-stage estimates for the two-stage least squares (2SLS) equations

$$\log(p_{i,t}) = \pi_{1a} p^{\text{spot}}_{i,t} + \pi_{1b} p^{\text{spot}}_{i,t} \times \text{SCG}_i + \pi_2 \text{HDD}^{\text{bill}}_{i,t} + \text{HH}_i + \text{City}_{i,t} + u_{i,t} \tag{2}$$

$$\log(q_{i,t}) = \eta_1 \widehat{\log(p_{i,t})} + \eta_2 \text{HDD}^{\text{bill}}_{i,t} + \text{HH}_i + \text{City}_{i,t} + v_{i,t} \tag{3}$$

where $\text{HH}_i$ is a household fixed effect, $\text{City}_{i,t}$ is a city by month-of-sample fixed effect, $\text{SCG}_i$ is an indicator for whether the household's retail utility is SoCalGas, and $\text{HDD}^{\text{bill}}_{i,t}$ is the number of heating degree days for household $i$ during its billing period that began in month $t$.

While the utilities update their prices $p_{i,t}$ for each calendar month, households' bills often straddle two calendar months—and consequently two pricing regimes—which gives the household a time-weighted average of prices for the months its bill straddles (see Figure 6). We construct the spot-price instrument $p^{\text{spot}}_{i,t}$ to correspond to this pricing policy—calculating the weighted average of the relevant months' spot prices.[11] Consequently, the spot-price instrument $p^{\text{spot}}_{i,t}$ is not collinear with the city-by-month-of-sample fixed effects: within a given city-utility-month, there is residual variation in prices (spot and retail) due to staggered bill start dates.[12]

Figure 5b provides visual evidence of the first stage—illustrating (1) the strong link between the two utilities' prices and the Henry Hub spot price and (2) the utilities' differential responses to the spot price. Notably, SoCalGas's price series appears to track the spot-market price for natural gas more closely than PG&E's price series. Throughout the rest of the paper, we define the Henry Hub spot price as the average spot price for natural gas at the Henry Hub during the seven days preceding the utility's change in pricing.

Panel A of Table 6 displays the first-stage results for equation 2 using five different prices that may be relevant to households: marginal price, average price, average marginal price, baseline price, and simulated marginal price[13] (using the log of each price). Each price is the second lag of the contemporaneous price.[14] Table 7 and Tables A3–A6 compare consumers' varying responses to different lags of price.

Both Figure 5b and Panel A of Table 6 demonstrate that the spot-price based instruments are quite strong: the $F$ statistics testing the joint significance of the instruments range from 369.9 to 1,333.2. This significance is unsurprising because the utilities purchase gas on the spot market and incorporate these costs directly into their price regimes. The significance

11. Weighting the months' spot prices by their temporal share of the bill. See the appendix section Calculating bills for further discussion of bills spanning multiple months.

12. We can drop the within city-utility-month variation induced by the staggered bill start dates by assigning bills the price (a) at the start of the billing period or (b) for the month in which the bill spends the majority of its time. These changes do not notably change the results. We include this variation because it better matches the actual mechanisms generating variation in the first stage.

13. *Simulated marginal price* refers to a simulated instrument for marginal price. We discuss this measure of price in the next section.

14. The current bill is lag zero, the prior bill contains the first lag of price, and the bill preceding the prior bill contains the second lag of price.

of the interaction between spot price and utility (SoCalGas) in the second row of Panel A in Table 6 confirms that the utilities differ appreciably when incorporating spot-market costs into their pricing regimes: PG&E's pricing regime appears to be much less responsive to the contemporaneous spot price than that of SoCalGas, matching the observation from Figure 5b above.[15] Though the city-by-month-of-sample fixed effect should control for most local demand shocks, bills do not perfectly match months. The within-bill HDDs variable $\text{HDD}_{i,t}^{\text{bill}}$ in equation 2 controls for any remaining weather-based demand shocks. The results in Table A9 demonstrate robustness to excluding (odd-numbered columns) or including (even-numbered columns) within-bill heating degree days, which suggests that the instrument is exogenous to local weather shocks, one of the key local-demand drivers in natural gas (Davis and Muehlegger 2010; Levine, Carpenter, and Thapa 2014; Hausman and Kellogg 2015).

#### 4.3.3 Instrumented prices and simulated instruments

In the preceding sections, we discussed how we interact a spot-price–pass-through-based instrument with a spatial discontinuity in utilities' service area. We now discuss one additional level of our identification strategy that directly deals with the price-is-a-function-of-quantity endogeneity present in multi-tiered pricing contexts.[16] We present three separate options for breaking this endogenous link between price and quantity, but, in the end, the options yield very similar results.

**Option 1: Instrumented prices** One option is simply to instrument the household's price with a variable that is aggregated at a unit above household. Consider the IV strategy discussed above: instrumenting a household's price with the Henry Hub spot-price interacted with utility. Because this instrument only varies at the level of billing-period by utility, when we regress a household's endogenous price on this instrument (and our set of fixed effects) in the first stage, the variation captured by the predicted prices is only the variation that correlates with the spot price, which is determined weeks, if not months, before the household's consumption decision.

**Option 2: Baseline price** In a similar manner, the baseline price provides a valid instrument that breaks the price-is-a-function-of-quantity endogeneity. Because a household's baseline price is not a function of its quantity consumed, baseline price does not suffer from the same endogeneity. Baseline price is also strongly predictive of marginal (or average) price.[17] Thus, in application, one could either replace marginal (or average) price with baseline price or

15. One difference between the utilities' pricing regimes is that PG&E does not have a fixed charge, while SoCalGas does. Thus, PG&E recovers both fixed and volumetric costs through volumetric charges to its customers.

16. This endogeneity is present in marginal price, average price, and average marginal price—all three prices are functions of the individual household's quantity consumed.

17. The correlation between baseline price and marginal price is approximately 0.79; the correlation between baseline price and average price is approximately 0.94. See Table A1 for all bivariate correlations between our five measures of price.

instrument one of the endogenous prices with baseline price. There is at least one drawback to this approach: baseline price, by definition, fails to capture the higher price that a household faces when the household exceeds its total monthly allowance.

**Option 3: Simulated instrument** Simulated instruments (sometimes called *policy-induced instruments*) provide a third option for breaking the price-is-a-function-of-quantity flavor of endogeneity. The simulated-instrument approach follows a methodology suggested in Ito (2014). Specifically, this approach creates an instrument (or proxy) for marginal (or average) price by plugging a lagged level of consumption into the current price regime, *i.e.*,

$$z_{i,t} = p_{i,t}\left(q_{t-k}\right) \tag{4}$$

The main idea for this instrument is using a household's consumption history to predict whether a household will face the baseline or excess price in the current period. As with any instrument, we want to accomplish this prediction in a way that is strongly predictive of the true outcome (the first stage) and that is uncorrelated with any recent shocks to the household (the exclusion restriction) (Angrist and Pischke 2009). For these reasons, we modify equation 4 slightly. First, we use the households' lagged consumption levels (from lagged bills 10 through 14 months prior) to calculate the share of lagged periods that exceed this billing period's baseline allowance, *i.e.*,

$$s_{i,t} = \frac{1}{5}\sum_{k=10}^{14} \mathbb{1}\{q_{i,t-k} > \bar{A}_{i,t}\} \tag{5}$$

where $\bar{A}_{i,t}$ is household $i$'s baseline allowance in time (bill) $t$. We then calculate the simulated instrument for marginal price, $z_{i,t}$, as

$$z_{i,t} = \mathbb{1}\{s_{i,t} \leq 0.5\} \times p_{i,t}^{\text{base}} + \mathbb{1}\{s_{i,t} > 0.5\} \times p_{i,t}^{\text{excess}} \tag{6}$$

Summarizing equations 5 and 6: this simulated instrument for marginal price predicts that a household will exceed its allowance when the majority of the household's past bills (using lagged months 10 through 14) exceed the current bill's allowance.[18]

Table A2 provides evidence that this simulated-instrument approach significantly predicts households' marginal prices. Specifically, Table A2 provides the estimate and standard error for $\beta$ in the equation

$$p_{i,t}^{\text{mrg}} = \beta p_{i,t}^{\text{sim}} + \text{HH}_i + \text{City}_{i,t} + w_{i,t} \tag{7}$$

where $p_{i,t}^{\text{mrg}}$ is household $i$'s marginal price in time $t$ and $p_{i,t}^{\text{sim}}$ is our simulated instrument for

18. This simulated instrument is robust to the choice of months 10 through 14. The goal is to keep the instrument in the same season as the current bill (maintaining a strong first stage), while allowing some temporal distance between the lags and the current period (the exclusion restriction: preventing short- and medium-run shocks from affecting both periods).

household $i$'s marginal price in time $t$ $\left(\text{i.e., } p_{i,t}^{\text{mrg}}\right)$. The estimates for $\beta$ in Table A2 confirm the strong "first stage" for this simulated instrument. Marginal price and simulated marginal price are strongly and significantly correlated—both $t$ statistics are approximately 148. The two columns in Table A2 also provide evidence of the robustness of the simulated instrument to the choice of lags: the estimates using lags 10–14 or 11–13 are virtually indistinguishable. In addition, the bottom row of Table A1 demonstrates that this simulated instrument is strongly correlated with marginal price ($r \approx 0.85$) in addition to the other four measures of price.

Column 5 of Table 6 (Panel A) provides the first-stage results consistent with equation 2 but with the simulated instrument of marginal price substituted (proxying) for actual marginal price (and still instrumenting with spot price interacted with utility across the utilities' border).[19] The first stage is again quite strong in this specification, and the results are qualitatively similar to the results in columns 1–4 of Table 6, Panel A. Henceforth we will refer to the simulated instrument for marginal price as simulated marginal price.

# 5 Results

## 5.1 Pooled price elasticity of demand for natural gas

Panel B of Table 6 displays the elasticity results from the second-stage regression specified in equation 3. These results instrument log price with the Henry Hub spot price (interacted with utility), exploit the spatial discontinuity, and use the log of daily average consumption (in therms) as the outcome. The five columns each estimate the elasticity using the log of a different type of price: marginal price, average price, average marginal price, baseline price, and simulated marginal price. As discussed above, each price is the second lag of price, as opposed to the contemporaneous price. The estimates for the price elasticity of demand range from −0.19 (average price) to −0.15 (average marginal price).

Comparing the estimated elasticities across the different types of price shows little difference. Most notably, the estimated elasticity with respect to marginal price (−0.174) is slightly lower than that with respect to average price (−0.192). Both differ from zero but do differ from each other at the 10% level. This result is reasonable in our setting, as the step in natural gas prices across tiers is not as significant as it is in the case of electricity.

Table A10 demonstrates that the estimated elasticity is also robust to the inclusion/exclusion of heating degree days and to the levels of fixed effects—ranging from city by month-of-sample fixed effects to zip-code by week-of-sample fixed effects (while still including household fixed effects). The robustness to type of price also demonstrates robustness to how we control for the price-is-a-function-of-quantity endogeneity discussed above. Tables A10–A14 demonstrate the robustness of the estimated elasticity to excluding within-bill heating degree days and

19. It is worth noting that in our paper, any result using the simulated instrument will have fewer observations than other results, as the simulated instrument is greedier for data—for an observation to remain in the dataset, its 14th lag must also be in the dataset. Our other price measures are not as greedy.

varying the spatiotemporal fixed effects. Finally, Table A15 contains marginal-price based elasticity estimates as we incrementally extend the study-area. Beginning with the study area (*Common Zips*), we add the zip codes neighboring (bordering) the study area (*Neighbors 1*); we then add the neighbors of the neighbors (*Neighbors 2*); last, we add a third band of neighbors (*Neighbors 3*). Figure A6 illustrates these groups of neighboring zip codes. The estimated elasticity from the first group of neighbors (−0.15 (0.05) in column 2 of Table A15) is quite close to the elasticity previously discussed (−0.17, (0.06) in column 1); the elasticities that include the second and third peripheral neighbors diminish in magnitude (−0.08 and −0.06) but still differ significantly from zero.

Compared to their OLS-based counterparts in Table 5, the marginal-price based 2SLS estimates for the elasticity of demand now have opposite—and theoretically correct—signs. The magnitudes of our preferred estimate of the elasticity (approximately −0.17) is precisely estimated and within the range of previous findings. Furthermore, these estimates are identified and hence carry a causal interpretation.

As discussed above, the results discussed up to this point—*e.g.*, the results in Table 6—estimate the price elasticity of demand for residential natural gas using the second lag of the various prices. In order for a household to know the price of its contemporaneous bill, the household would need to closely follow the approved advice letters published online by the utility or the California Public Utilities Commission *and* pay careful attention to their daily gas consumption, which is much harder to do than for electricity for technical metering reasons as gas meters are less smart. Otherwise, the household will learn about prices from past bills—hence the use of lagged prices. Because a household will not receive the bill for the previous billing period for several days into its current billing period—and because the household may not view the previous bill until it pays the bill (or its credit card bill, if the household uses automatic bill payment) weeks later—the household may not know the prices from its immediately previous bill until the current period is nearly over. For these reasons, it is plausible that the second lag of price is the most salient price to many households. Figure 6 illustrates an example of the timing for bill delivery, bill payment, and the relevant lags of prices.

Table 7 replicates the second-stage results for marginal price and average price but varies the lag/lead of price: beginning with the first lead of price, followed by contemporaneous price, the first lag of price, and finally the second lag of price. Tables A3–A6 provide further detail, varying the lead/lag of each of the five prices—ranging from the first lead of price to the third lag of price. Across the five types of measures of price, none of the first leads of price, contemporaneous prices, or first lags of price differ significantly from zero. For each type of price, both the second and third lags of price differ significantly from zero. For each price, the second-lag based elasticity slightly exceeds the third-lag based elasticity in magnitude, but the difference does not exceed the standard error. These results are consistent with households responding to two-to-three lags of price—as opposed to contemporaneous price—suggesting some degree of inattention by the household to the true price, akin to previous work on inattention and salience, *e.g.*, Chetty, Looney, and Kroft (2009), Sallee (2013), and Allcott and Taubinsky (2015).

At the suggestion of a referee, we use an alternate instrument instead of Henry Hub price: lagged heating degree days (HDDs) on the East Coast in various forms (polynomials and bins). While we maintain the Henry Hub prices are exogenous to a customer's bill, this serves as a robustness check. Table A16 displays the estimated elasticities for marginal price using six different functional forms of population-weighted, lagged HDDS east of the Mississippi River. The estimated elasticities range from −0.20 to −0.14—consistent with our previous results that ranged from −0.19 to −0.15. The most flexible polynomial results in an estimate of −0.172, which is almost identical to the estimate using the Henry Hub prices and statistically different form zero at the 5% level.

## 5.2 Heterogeneity

We now examine evidence of heterogeneity in the price elasticity of demand for natural gas. Our institutional setting motivates two relevant dimensions of heterogeneity—income level and season—as the CPUC and utilities already apply different price regimes to households depending upon (1) the season of year (summer *vs.* winter) and (2) the household's income level (as proxied by CARE status). If heterogeneity exists, then the regressions in the preceding section pool across the heterogeneous effects. This pooled parameter estimate may still be relevant for policy applications—particularly for policies that cannot differentiate between seasons and/or income groups. However, because OLS weights heterogeneous treatment effects by their shares of the residual variation in the variable of interest—which is itself a function of (1) the numbers of observations in the heterogeneous groups and (2) the (residual) within-group variance in the variable of interest (Solon, Haider, and Wooldridge 2015)—one might wonder whether the pooled estimator always provides a policy-relevant estimate.

The income-based heterogeneities that we estimate should be taken as descriptive for the given group, rather than causal effects of income level or season. In other words, while we estimate heterogeneous elasticities with respect to income level/CARE status, this heterogeneity may have nothing to do with income and could instead result from some other (omitted) variable that correlates with income/CARE status, *e.g.*, the age and size of the physical home. However, identification of the sources of heterogeneity is not our goal; instead, we aim to identify the elasticity of demand and demonstrate dimensions of heterogeneity. We leave it for future papers to identify the mechanisms behind this heterogeneity. For seasonal heterogeneity, we split the calendar into winter months (October through March) and summer months (April through September).[20]

20. This definition reflects Southern California's two seasons: warm and slightly less warm.

#### 5.2.1 Seasonal heterogeneity

To estimate seasonal heterogeneity in the price elasticity of demand for residential natural gas, we estimate the two-stage least squares equations 2 and 3 separately for winter months and for summer months. Columns (1) and (2) of Table 8 supply the second-stage results from these regressions, providing estimates of the elasticity of demand by season.

The results in columns (1) and (2) of Table 8 indicate a stark and significant difference between price elasticities in summer and winter months. Using marginal price, we estimate that the price elasticity of demand for natural gas in summer months is approximately 0.03 (0.03), which does not differ from zero statistically at the 10% level. The estimated elasticity for winter months is approximately −0.37 (0.11) and differs significantly from zero at the 1 percent level. The two point estimates are different form each other at the 1% level. These results provide strong evidence that households' consumptive and price-response behaviors vary considerably by season—the winter-based elasticity is nearly twice the "pooled" elasticity.[21]

#### 5.2.2 Income heterogeneity

To examine income-based heterogeneity in the price elasticity of demand for natural gas, we estimate the two-stage least squares equations 2 and 3 separately for CARE households and non-CARE households. Columns (3) and (4) of Table 8 supply the second-stage results from these regressions, providing estimates of the elasticity of demand by income level (CARE status).

The results in columns (3) and (4) of Table 8 suggest that the elasticity results in the previous section may in fact pool across heterogenous elasticities; we estimate that the price elasticity for CARE (lower-income) households is approximately twice that of non-CARE (higher-income) households. Specifically, using the marginal price, we estimate an elasticity of approximately −0.20 (0.062) for CARE (lower income) households and −0.13 (0.062) for non-CARE households. The "pooled" estimate corresponding to these results is −0.174 (0.065) (column (1) of Panel B in Table 6)—slightly higher than the midpoint between the CARE estimate and the non-CARE estimate. We note that the two point estimates are not statistically different from each other at the 10% level, which in part may be due to a conservative (but we believe correct) clustering strategy.

#### 5.2.3 Income-by-season heterogeneity

To estimate seasonal-by-income heterogeneity in the own-price elasticity of demand for residential natural gas, we estimate the two-stage least squares equations 2 and 3 separately for the four potentially heterogeneous subgroups: CARE households in the summer, non-CARE

21. Table A7 reproduces these heterogeneity results using average price—rather than marginal price—with very similar results.

households in the summer, CARE households in the winter, and non-CARE households in the winter. Table 9 displays the second-stage results from these regressions, providing estimates of the elasticity of demand by season and CARE status.

The results in Table 9 are consistent with heterogeneous elasticities that depend upon the interaction between household income (CARE status) and season. In other words, the difference between a household's winter and summer price elasticities varies by the household's income level. Specifically, the results in Table 9 indicate that both income groups are essentially inelastic to prices in summer months; we estimate a "summertime" price elasticity of 0.025 (0.035) for CARE households and 0.062 (0.033) for non-CARE households. Both elasticities are positive, but only one is significantly different from zero at the 10% level and small. In winter months, both sets of consumers are significantly and substantially more responsive to price, but the point estimates for CARE households indicate a higher price responsiveness. We estimate the "wintertime" price elasticity of demand for natural gas is −0.459 (0.110) for CARE households and −0.365 (0.110) for non-CARE households. For both types of households the summer elasticity is statistically different from the winter elasticity. We cannot, however, reject equality of the winter elasticities across both types of households given our clustering strategy.

### 5.3 Policy implications

The heterogeneity discussed above suggests that households are much more price sensitive during their high-consumption months—the winter. These high-consumption winter months also correspond to the time of year in which consumers use natural gas in its most salient form: heating. When we break down the price elasticity across users and seasons, we show that subsidized (lower income) consumers display the largest price sensitivity during the winter (−0.46). Neither type of customer displays a substantial price response in the summer. These results suggest that, if one wanted to pass through the small program and administrative fees charged volumetrically by regulators to consumers, summertime may be best—both for efficiency and possibly for progressivity.[22] This reasoning suggests an implementable dimension for regulators charging small administrative and program fees that we have not seen recommended in the literature or applied in practice in the utility space. For a much more thoughtful discussion of these issues, we refer the reader to Hausman (2019) and Borenstein, Fowlie, and Sallee (2021).

22. This point hinges critically on the assumption that external costs from natural gas combustion are properly priced. For global pollutants, this is the case in California because the natural gas sector is part of California's cap and trade system.

## 6 Discussion and conclusion

We use a large sample of household natural gas bills with a multi-part identification strategy to provide micro-data based causal estimates of the own-price elasticity of demand for residential natural gas across seasons and income groups for California. Using cross-border price variation between California's two largest natural gas utilities—resulting from the utilities' differential pass-through of wholesale price variation—we isolate plausibly exogenous variation in residential natural gas prices. We estimate an elasticity of −0.17 with respect to marginal price. This estimate is robust to specification choices that include within-bill weather, several price instruments, and the definition/type of price. The point estimates for the own-price elasticity range from −0.19 to −0.15 across five considered measures of price. Notably, we do not find a differential consumer response to average versus marginal price. Because households respond significantly to price changes two to four months prior to the period of consumption—and following Ito (2014)—we interpret these estimated elasticities as medium-run elasticities. We agree with Ito (2014) who notes that the medium-run elasticity is often the most policy-relevant elasticity.

As a second important finding, we estimate that the own-price elasticity of demand varies across seasons and possibly customer types. We show that the price elasticity varies greatly across seasons. If we average across types of households, the summer price elasticity is close to, and only marginally different from, zero. The winter price elasticity is −0.37. Households are most responsive to natural-gas prices when they are heating. We show suggestive evidence that households on a popular low-income program, which subsidizes households' natural gas and electricity, appear to be more elastic in their response to price as households who are not part of the program. We further decompose the estimated price elasticity across users and seasons, and show that lower-income consumers display the largest price sensitivity during the winter (−0.46), although this estimate is not statistically different from that for non-CARE (higher income) customers.

Finally, based upon the heterogeneity underlying the own-price elasticity of demand, we suggest that small volumetric administrative program fees could be billed during inelastic periods (here: warm, non-heating months). In the setting of natural gas, this policy idea has the potential to slightly enhance efficiency and possibly equity.

## References

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

Allcott, Hunt, and Dmitry Taubinsky. 2015. "Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market." *American Economic Review* 105 (8): 2501–2538.

Angrist, Joshua D, and Alan B Krueger. 2001. "Instrumental Variables and the Search for Identification: From Supply and Demand to Natural Experiments." *Journal of Economic Perspectives* 15 (4): 69–85.

Angrist, Joshua D., and Jörn-Steffen Pischke. 2009. *Mostly harmless econometrics: an empiricist's companion.* Princeton University Press.

Borenstein, Severin, Meredith Fowlie, and James Sallee. 2021. "Designing electricity rates for an equitable energy transition." *Energy Institute at Haas working paper* 314.

Brown, SPA, and MK Yücel. 1993. "The Pricing of Natural Gas in U.S. Markets." *Economic Review-Federal Reserve Bank of Dallas,* no. April, 41–52.

California Energy Commission. 2015. "California Building Climate Zone Areas." Accessed March 01, 2017. Accessed February 1, 2015. http://www.energy.ca.gov/maps/renewable/building_climate_zones.html.

———. 2017. "Title 24: Building Energy Efficiency Program." Accessed March 01, 2017. Accessed March 1, 2017. http://www.energy.ca.gov/title24/.

Chen, Yuyu, Avraham Ebenstein, Michael Greenstone, and Hongbin Li. 2013. "Evidence on the Impact of Sustained Exposure to Air Pollution on Life Expectancy from China 's Huai River Policy." *Proceedings of the National Academy of Sciences of the United States of America* 110 (32): 12936–12941.

Chetty, Raj, Adam Looney, and Kory Kroft. 2009. "Salience and taxation: Theory and evidence." *American Economic Review* 99 (4): 1145–1177.

Davis, Lucas W., and Erich Muehlegger. 2010. "Do Americans consume too little natural gas? An empirical test of marginal cost pricing." *RAND Journal of Economics.*

Dell, Melissa. 2010. "The Persistent Effects of Peru's Mining Mita." *Econometrica* 78 (6): 1863–1903.

Esri. 2017. "USA ZIP Code Areas." Accessed January 11, 2017. https://www.arcgis.com/home/item.html?id=8d2012a2016e484dafaac0451f9aea24.

Ezekiel, Mordecai. 1925. "The Assumptions Implied in the Multiple Regression Equation." *Journal of the American Statistical Association* 20 (151): 405–408.

Garcia-Cerrutti, L. Miguel. 2000. "Estimating elasticities of residential energy demand from panel county data using dynamic random variables models with heteroskedastic and correlated error terms." *Resource and Energy Economics* 22 (4): 355–366.

Girshick, M. A., and Trygve Haavelmo. 1947. "Statistical Analysis of the Demand for Food: Examples of Simultaneous Estimation of Structural Equations." *Econometrica* 15 (2): 79–110.

Hahn, Robert W., and Robert D. Metcalfe. 2021. "Efficiency and Equity Impacts of Energy Subsidies." *American Economic Review* 111, no. 5 (May): 1658–1688.

Hanemann, W. Michael. 1984. "Discrete/continuous models of consumer demand." *Econometrica: Journal of the Econometric Society,* 541–561.

Hausman, Catherine. 2019. "Shock Value: Bill Smoothing and Energy Price Pass-Through." *The Journal of Industrial Economics* 67 (2): 242–278.

Hausman, Catherine, and Ryan Kellogg. 2015. "Welfare and Distributional Implications of Shale Gas." *Brookings Papers on Economic Activity,* no. Spring, 71–125.

Herbert, John H., and Erik Kreil. 1989. "Specifying and evaluating aggregate monthly natural-gas demand by households." *Applied Economics* 21 (10): 1369–1382.

Houthakker, H. S., and Lester D. Taylor. 1970. *Consumer Demand in the United States: Analyses and Projections.* 2nd ed. Cambridge, Mass.: Harvard University Press.

Ito, Koichiro. 2014. "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing." *American Economic Review* 104 (2): 537–563.

El-Khattabi, Ahmed Rachid, Shadi Eskaf, Julien P Isnard, Laurence Lin, Brian McManus, and Andrew J Yates. 2021. "Heterogeneous responses to price: Evidence from residential water consumers." *Journal of Environmental Economics and Management* 107:102430.

Koopmans, Tjalling. 1945. "Statistical Estimation of Simultaneous Economic Relations." *Journal of the American Statistical Association* 40 (232): 448–466.

Levine, Steven, Paul Carpenter, and Anul Thapa. 2014. *Understanding natural gas markets.* Technical report. American Petroleum Institute.

Maddala, G. S., Robert P. Trost, Hongyi Li, and Frederick Joutz. 1997. "Estimation of short-run and long-run elasticities of energy demand from panel data using shrinkage estimators." *Journal of Business and Economic Statistics* 15 (1): 90–100.

Meier, Helena, and Katrin Rehdanz. 2010. "Determinants of residential space heating expenditures in Great Britain." *Energy Economics* 32 (5): 949–959.

Metcalf, Gilbert E., and Kevin A. Hassett. 1997. "Measuring the energy savings from home improvement investment." *The Review of Economics and Statistics* 81 (3): 516–528.

Olmstead, Sheila M, W Michael Hanemann, and Robert N Stavins. 2007. "Water demand under alternative price structures." *Journal of Environmental economics and management* 54 (2): 181–198.

Pacific Gas and Electric Company. 2016. "Guide to California Climate Zones." Accessed August 13, 2016. Accessed June 14, 2016. http://www.pge.com/myhome/edusafety/workshopstraining/pec/toolbox/arch/climate/index.shtml.

PRISM Climate Group. 2004. "PRISM Gridded Climate Data." Oregon State University. Accessed February 4, 2004. http://prism.oregonstate.edu.

Rehdanz, Katrin. 2007. "Determinants of residential space heating expenditures in Germany." *Energy Economics* 29 (2): 167–182.

Sallee, James M. 2013. "Rational Inattention and Energy Efficiency." *Journal of Law and Economics* 57 (3): 781–820.

Solon, Garry, Steven J. Haider, and Jeffrey Wooldridge. 2015. "What are we weighting for?" *Journal of Human Resources* 50 (2): 301–316.

Southern California Gas Company. 2016. "California Alternate Rates for Energy (CARE)." Accessed August 08, 2016. Accessed August 8, 2016. https://www.socalgas.com/save-money-and-energy/assistance-programs/california-alternate-rates-for-energy.

Stock, James H, and Francesco Trebbi. 2003. "Retrospectives: Who Invented Instrumental Variable Regression?" *Journal of Economic Perspectives* 17 (3): 177–194.

U.S. Department of Labor, Bureau of Labor Statistics. 2017. *Consumer Expenditure Survey.*

U.S. Energy Information Administration. 2016a. "Average annual natural gas spot price in 2015 was at lowest level since 1999." Accessed November 04, 2017. Accessed January 5, 2016. https://www.eia.gov/todayinenergy/detail.php?id=24412#.

———. 2016b. "Electric Power Monthly - Net Generation by Energy Source: Total (All Sectors), 2006-May 2016." Accessed August 05, 2016. Accessed August 5, 2016. https://www.eia.gov/electricity/monthly/epm_table_grapher.cfm?t=epmt_1_1.

———. 2016c. "Frequently Asked Questions." Accessed March 01, 2017. Accessed April 6, 2016. https://www.eia.gov/tools/faqs/faq.cfm?id=45&t=8.

———. 2016d. "Henry Hub Natural Gas Spot Price." Accessed August 10, 2016. Accessed August 10, 2016. https://www.eia.gov/dnav/ng/hist/rngwhhdD.htm.

U.S. Energy Information Administration. 2017. "Maps: Layer Information for Interactive State Maps - Natural Gas Interstate and Intrastate Pipelines." Accessed October 23, 2017. Accessed March 1, 2017. https://www.eia.gov/maps/layer_info-m.php.

Woolridge, Jeffrey. 2009. *Introductory Econometrics - A Modern Approach.* 4e. Mason, OH: South-Western Cengage Learning.

Working, E. J. 1927. "What Do Statistical "Demand Curves" Show?" *The Quarterly Journal of Economics* 41 (2): 212–235.

Wright, Phillip G. 1928. *The Tariff on Animal and Vegetable Oils.* 1st ed. New York: MacMillan.

# 7 Figures

Figure 1: **U.S. natural gas institutional organization**

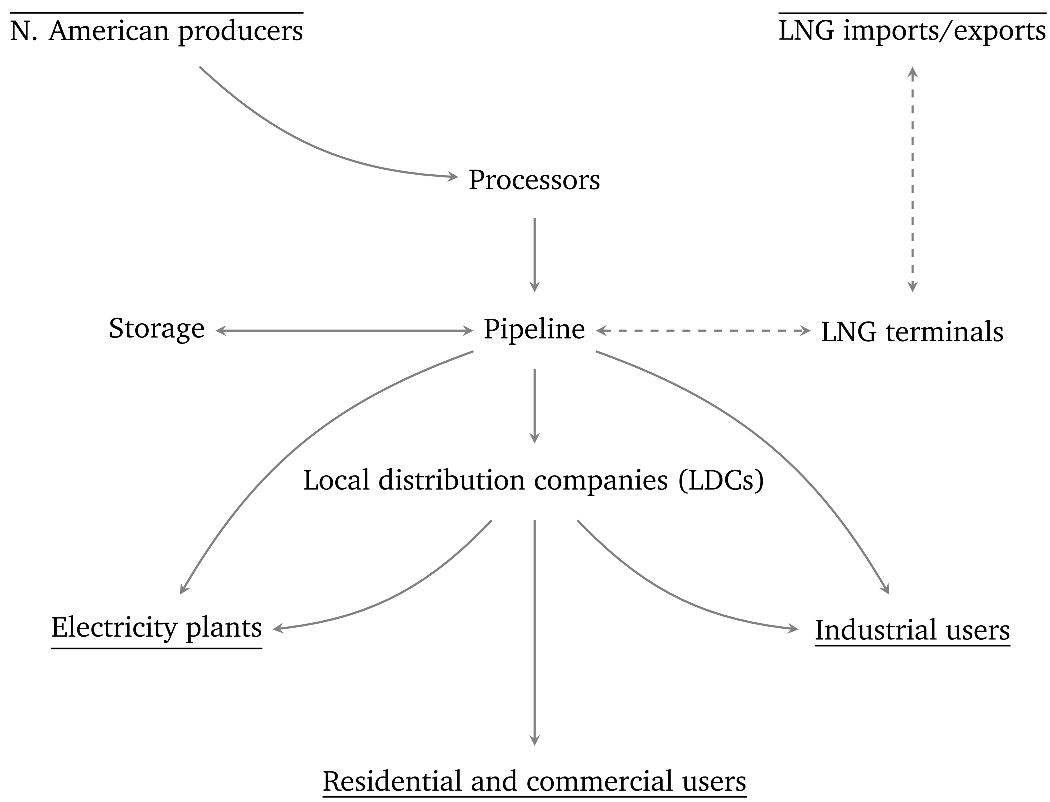


**Notes:** Overbars represent points of entry into the U.S. natural gas market; underbars represent end points in the market; all other labels represent intermediaries. Arrow directions correspond to the direction of the flow of natural gas. The acronym *LNG* abbreviates *liquid natural gas*. This figure is based on Levine, Carpenter, and Thapa (2014) with modification following Brown and Yücel (1993).

Figure 2: **Henry Hub natural gas spot price:** Daily, 1997–2016

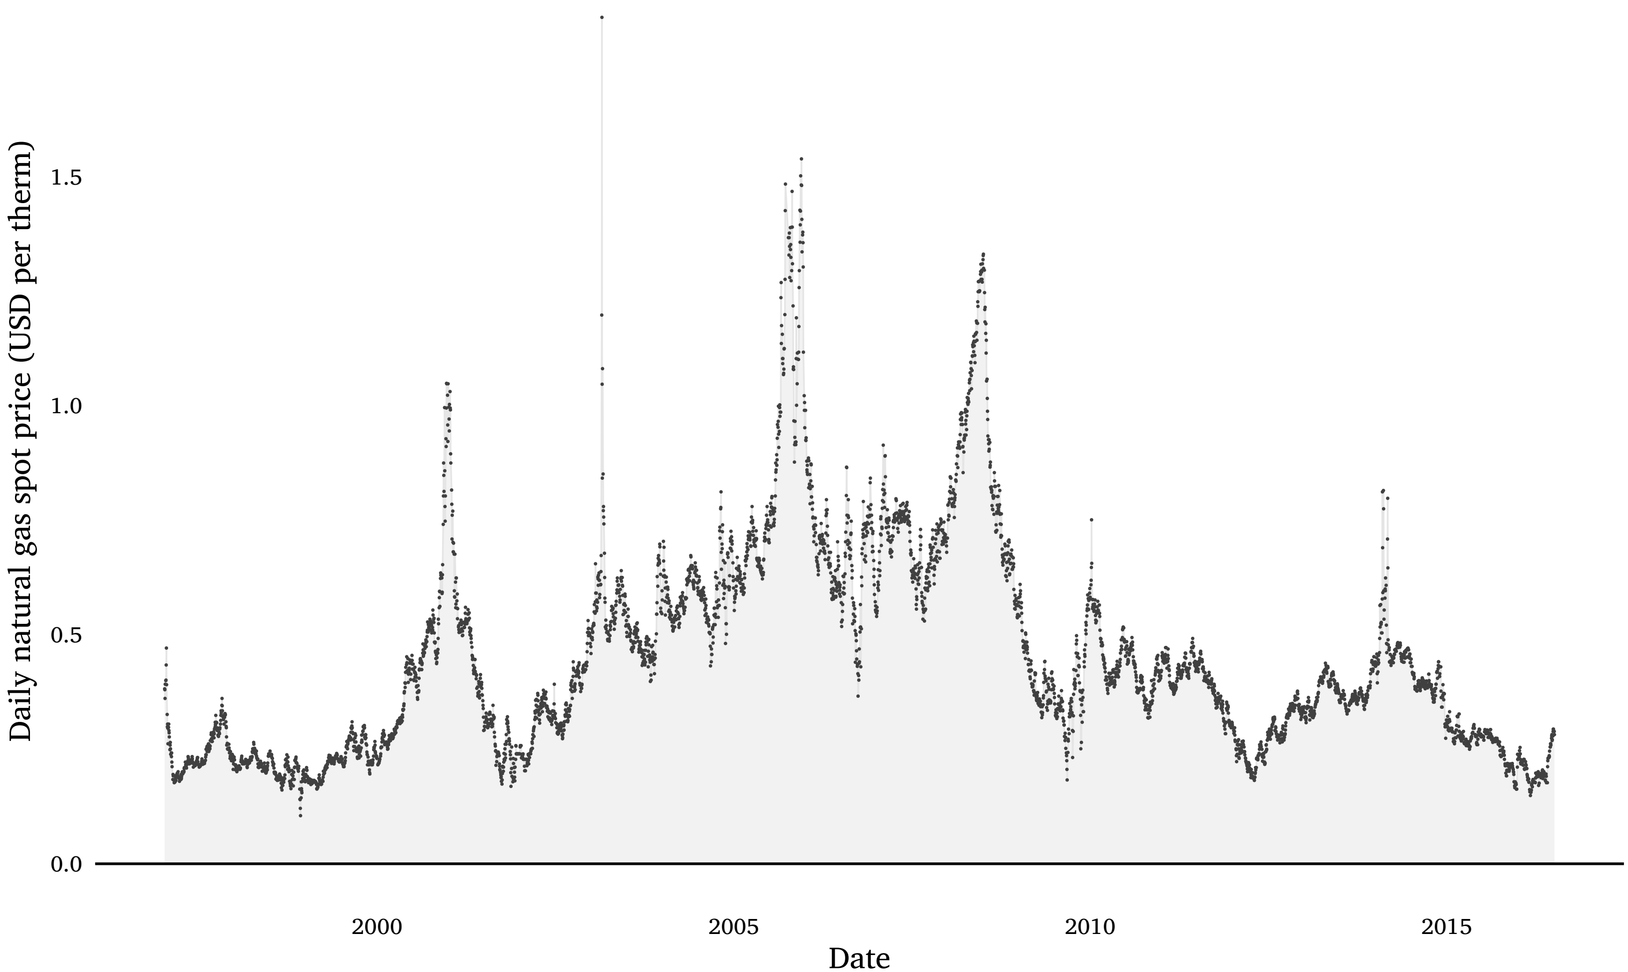


*Source:* U.S. Energy Information Administration

Figure 3: **Households' allowances and prices**

(a) **Allowance and marginal *vs.* average price example:** PG&E, January 2009, climate zone R

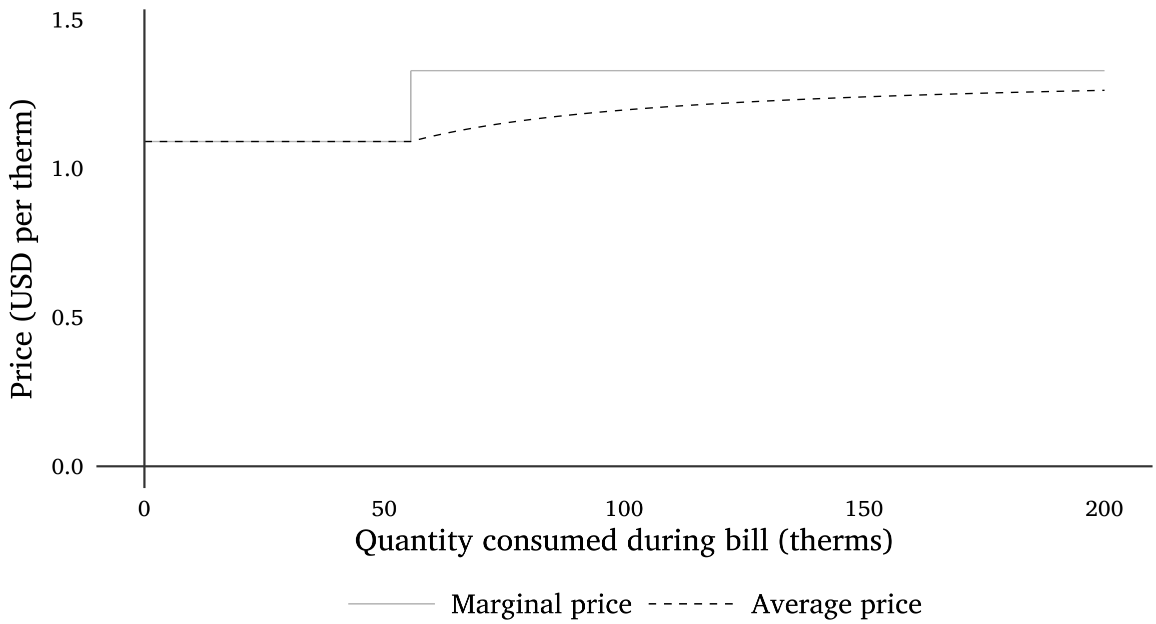


(b) **Tier-one daily allowances over time:** PG&E (zone R) and SoCalGas (zone 1), 2009–2015

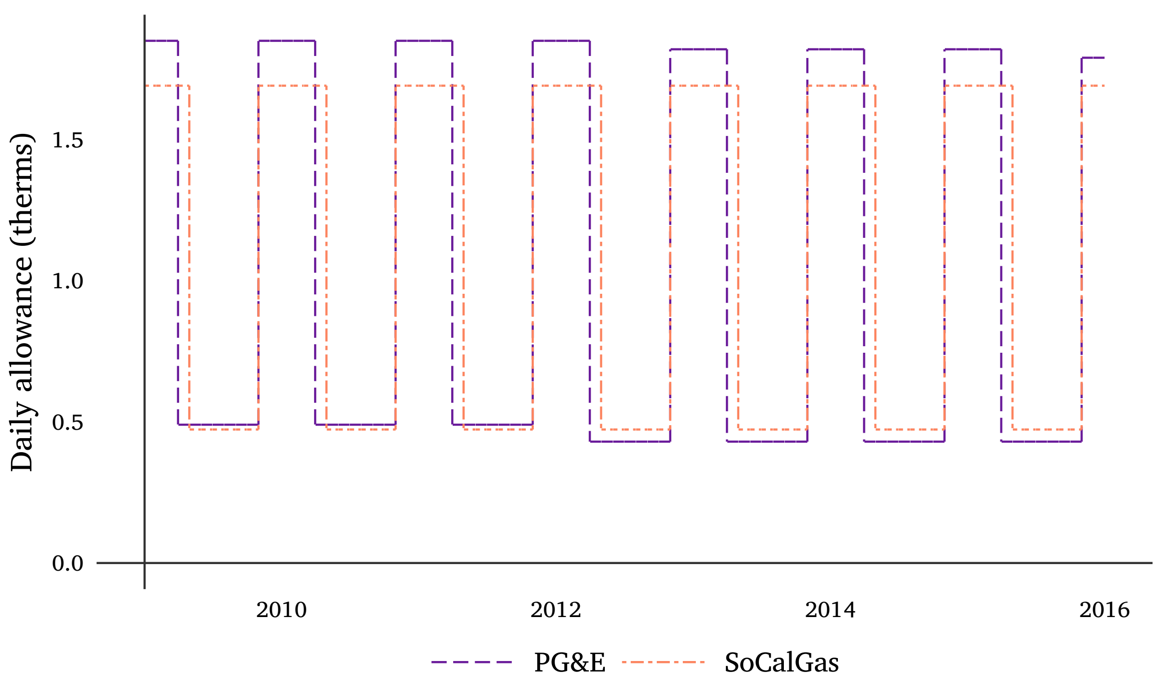


**Notes:** Households receive daily allowances for baseline (first-tier) consumption as a function of location and season (*e.g.*, climate zone *R*, January 2009). The household pays the second-tier price on all units that exceed its allowance—comparing total consumption (during the billing period) to total allowance (daily allowance summed across the bills' days).

Figure 4: **Natural gas service areas and the study-area discontinuity**

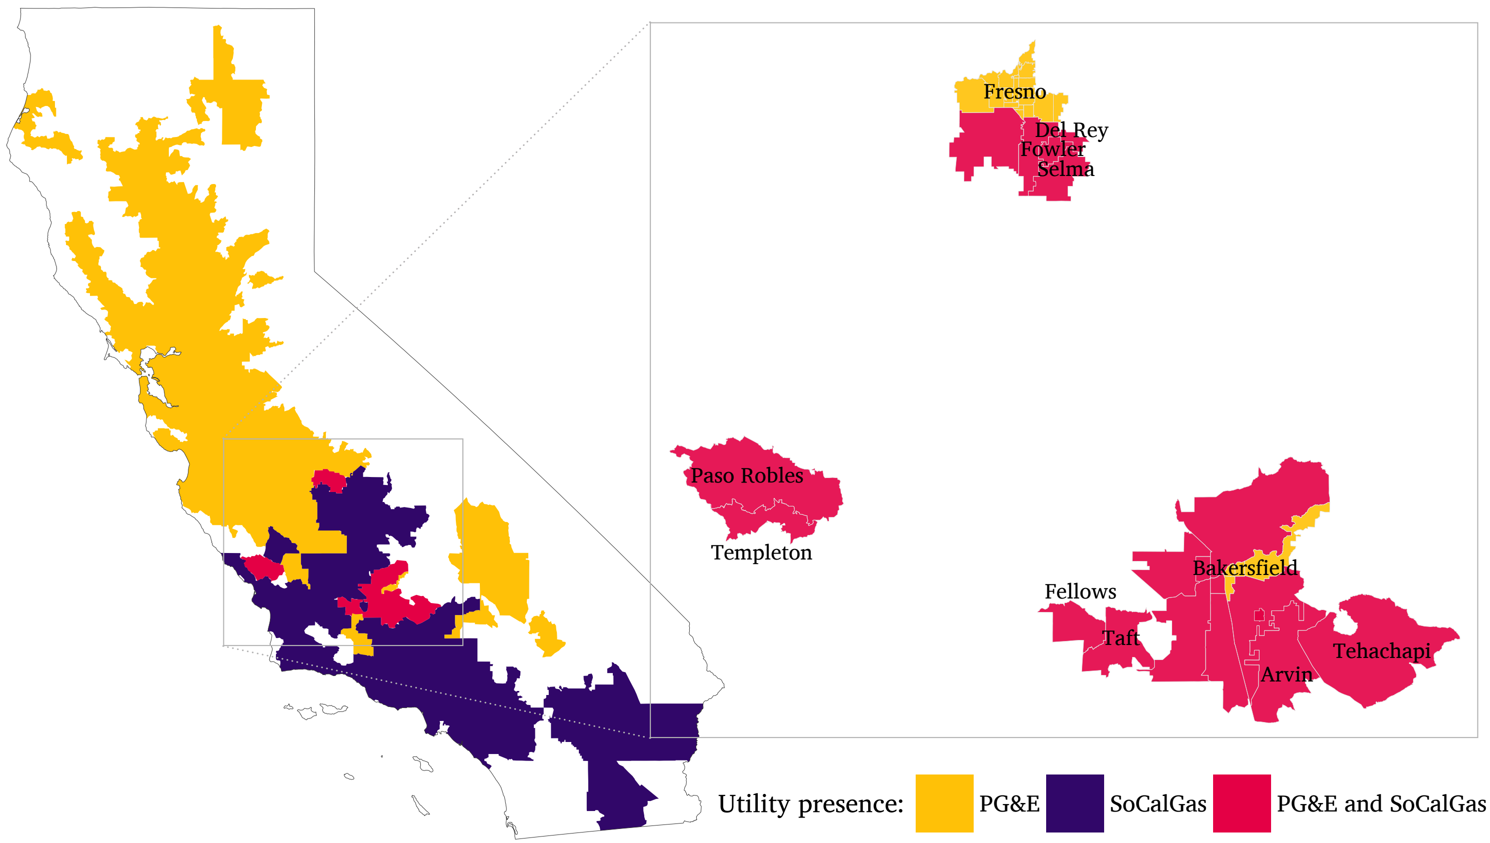


**Notes:** The left side of the figure displays PG&E's and SoCalGas's services areas (by 5-digit zip code). The right side of the figure zooms in on three clusters of cities that receive service from both utilities. These three clusters of cities encompass 39 zip codes; 18 of these (5-digit) zip codes receive service from both PG&E and SoCalGas. These 18 zip codes represent the main study area for the paper.

Figure 5: **Prices across utilities, tiers, and in the spot market**, 2009–2015

(a) **Price regimes over time:** PG&E and SoCalGas, 2009–2015

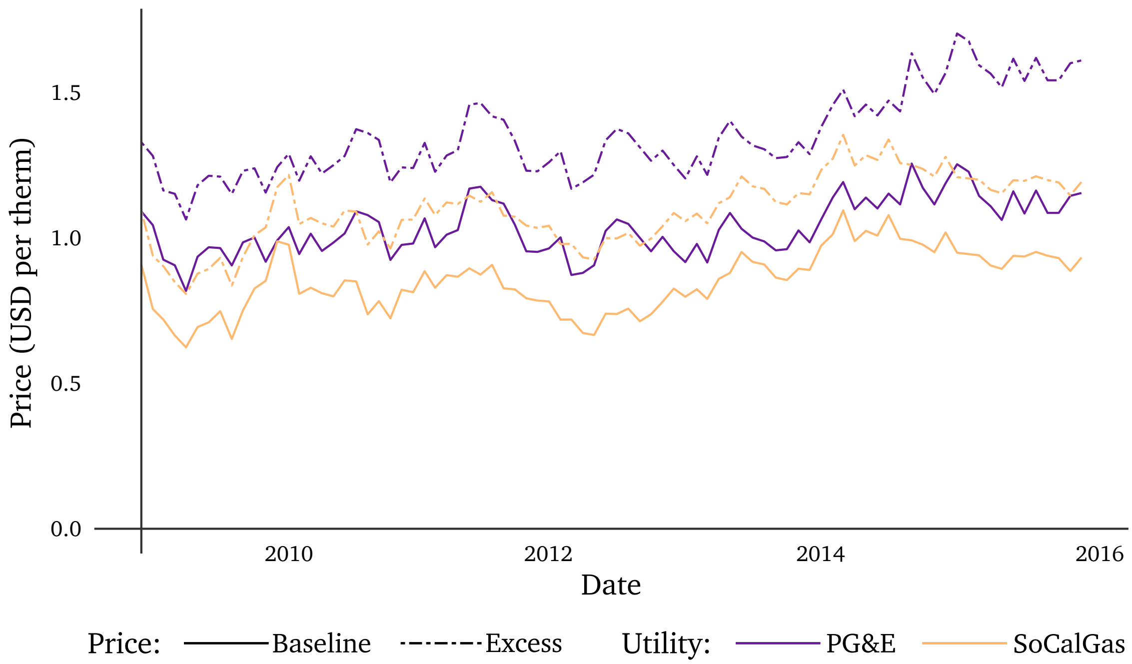


(b) **Correlation across prices** Three relevant natural gas price series, 2009–2015

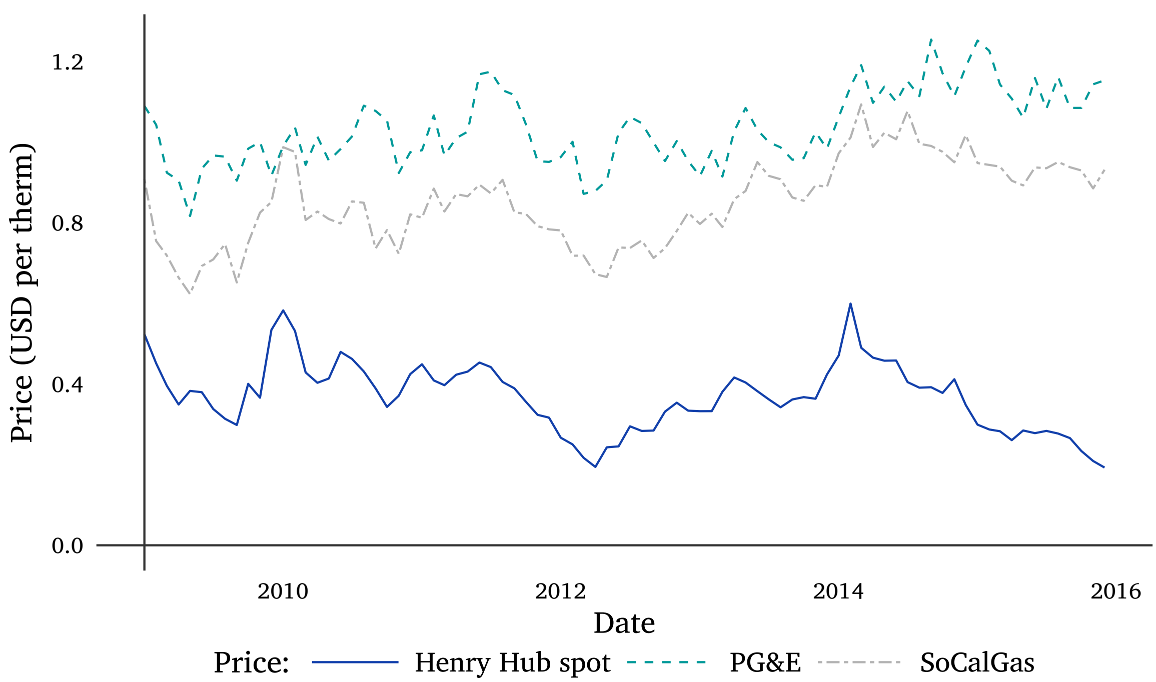


**Notes:** *Baseline* refers to first-tier price, *i.e.*, the price a household pays for its first therm of natural gas. *Excess* refers to the second-tier price, *i.e.*, the price a household pays for each therm that exceeds its first-tier allowance (see Figure 3). The Henry Hub spot price is generally recognized as a national benchmark (U.S. Energy Information Administration 2016a; Levine, Carpenter, and Thapa 2014).

Figure 6: **Calendar months and billing periods:** Four 30-day bills and five months

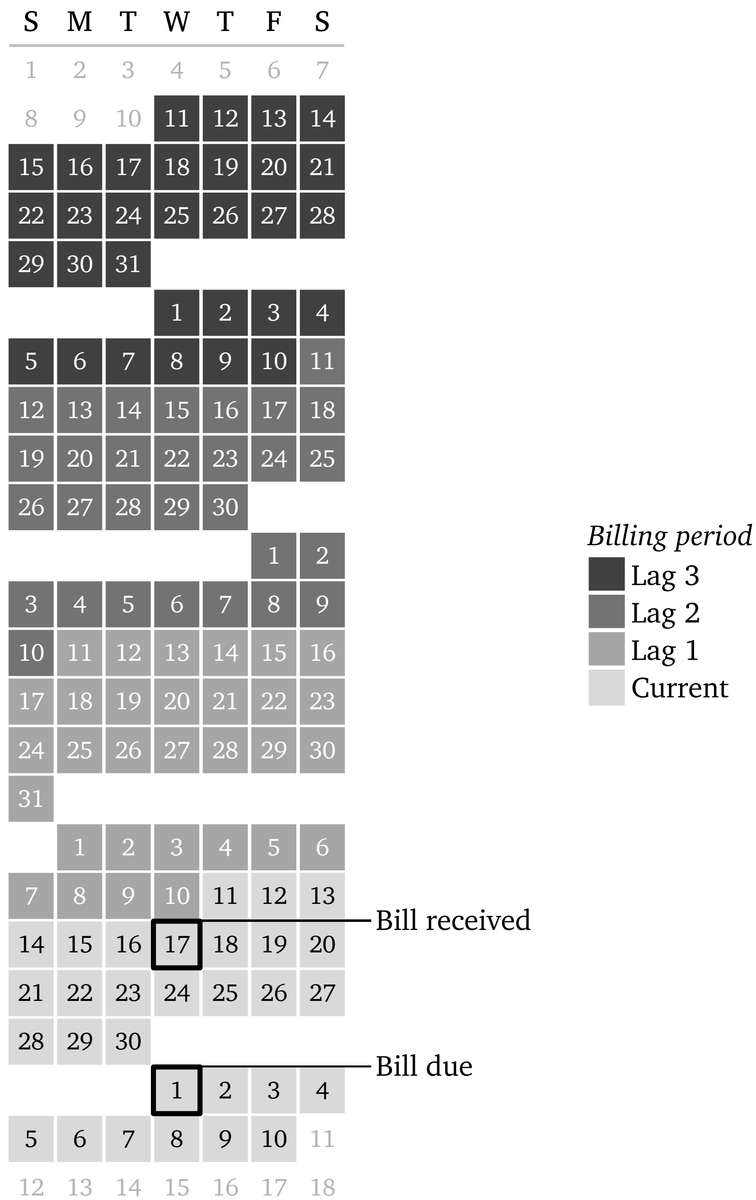


**Notes:** The household receives its bill from the *Lag 1* period on the fifth business day of its *current billing period* (the 17th). Payment for the *Lag 1* bill is due two weeks later (on the 1st). Now consider the question "Which lag of price is relevant?" **Current:** For the household to know the price structure for its current billing period, it must pay attention to the approval status of its utility's advice-letters correspondence with the CPUC. **Lag 1:** Again, unless the household pays attention to the utility's CPUC-approved advice letters, it will not know the prices in the *Lag 1* billing period until it receives and opens the bill. The bill arrives several days into the new period, and if the household does not see the bill until payment, it may not learn about the prices of the *Lag 1* bill until the current billing period is nearly complete. Autopay may extend this moment of salience even further into the future. **Lag 2:** Throughout the entirety of the *Current* billing period, the household will know the prices from its *Lag 2* bill, and for a non-zero amount of time, the *Lag 2* bill is likely to be the most recent set of prices the household knows. **Lag 3:** Same level of knowledge as *Lag 2* but less recent.

Figure 7: **Natural gas and electricity prices:** Comparing utilities and industries

(a) **Comparing trends in levels**, 2010–2014

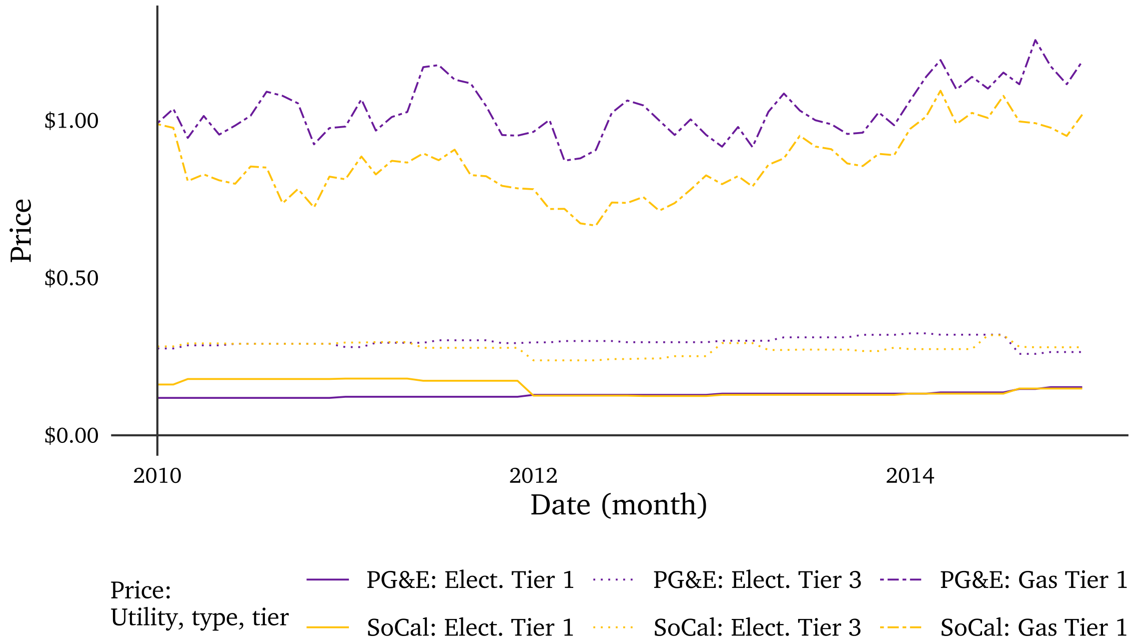


(b) **Comparing trends in differences** across utilities, 2010–2014

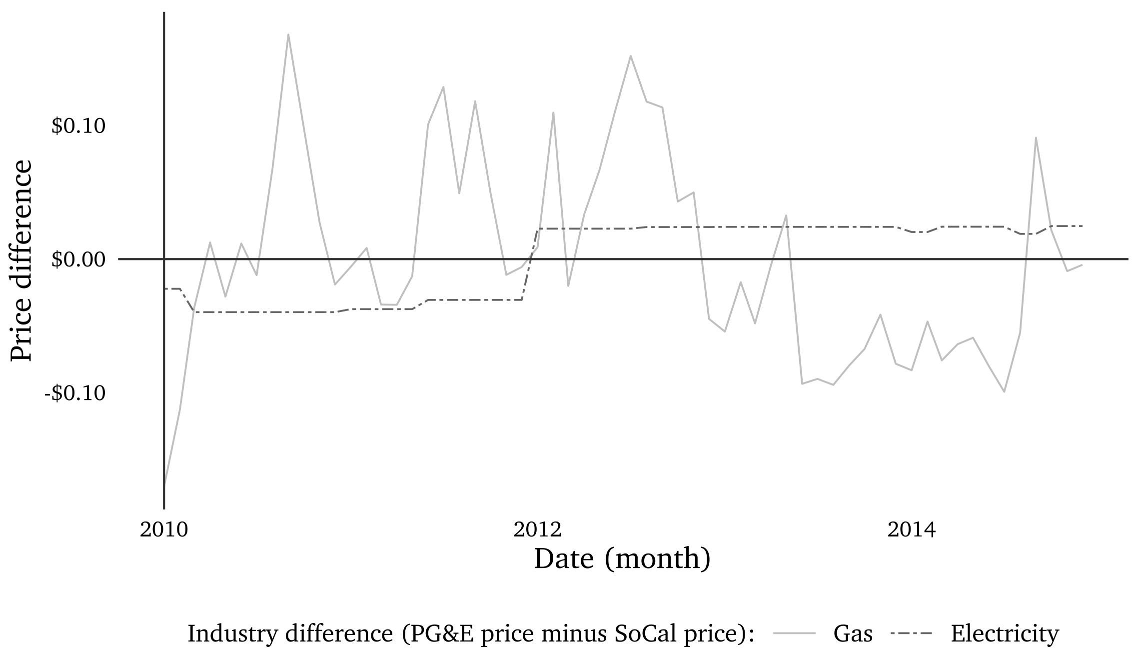


**Notes:** For the consumers in this paper, natural gas prices do not significantly correlate with electricity prices—neither in levels (Panel A), nor in differences (Panel B). *Differences* constitute PG&E minus SoCal within the same calendar month. We demean the time series of differences for each industry (natural gas and electricity). *SoCal* denotes the *Southern California Gas Company* for natural gas and *Southern California Edison* for electricity. The underlying data come from publicly available CPUC letters for the relevant utilities.

Figure 8: **Increasing tax efficiency** using seasonal heterogeneity

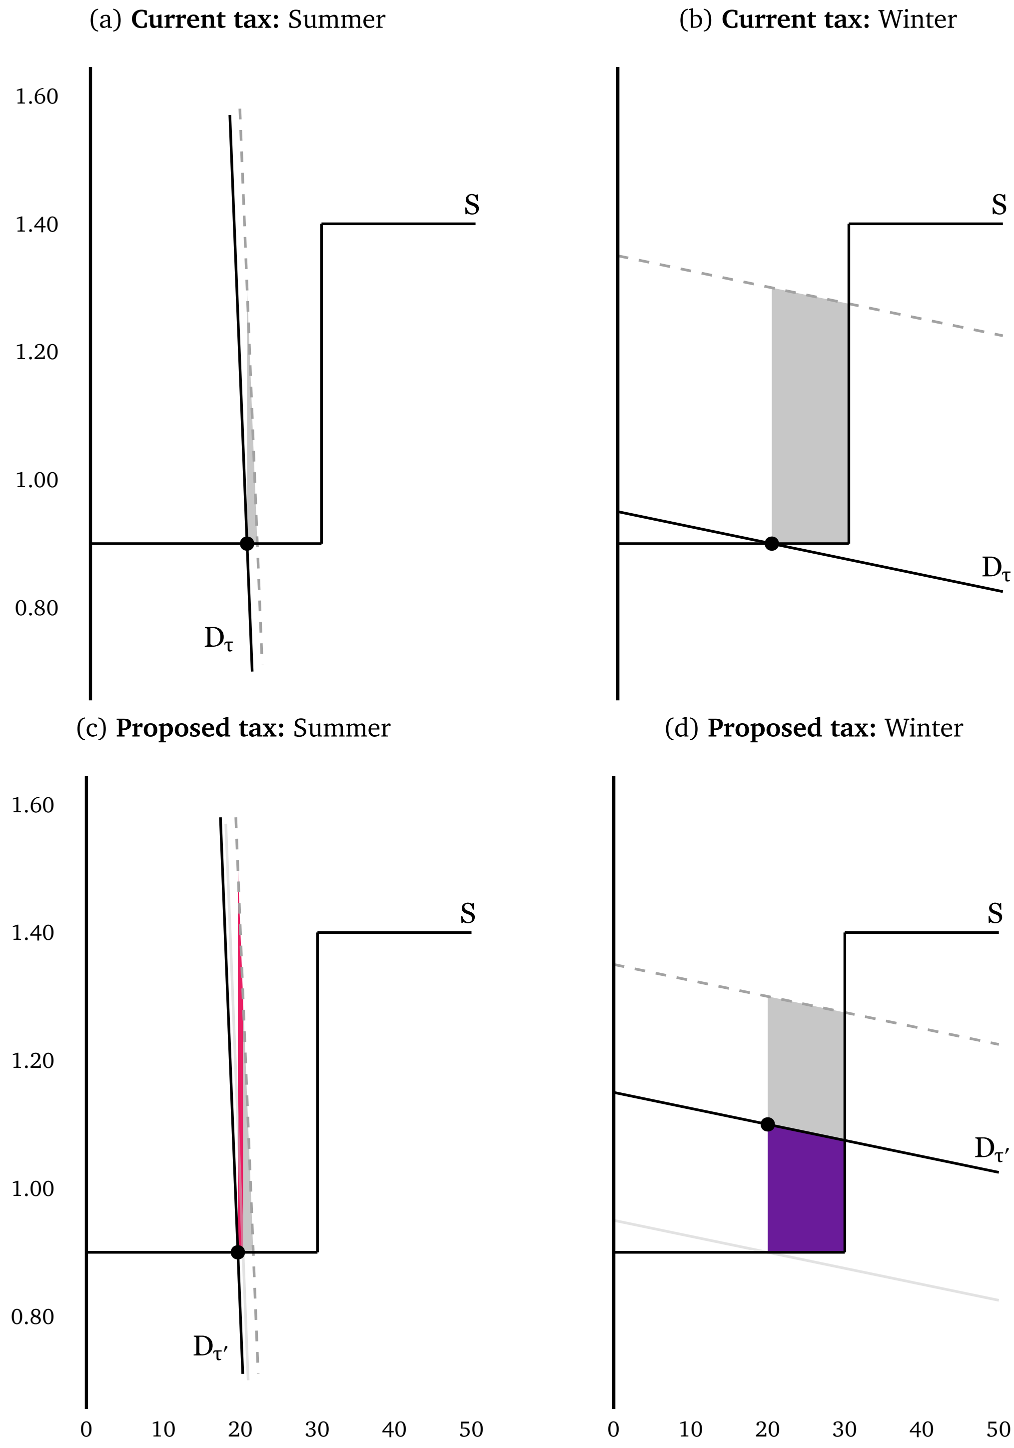


**Notes:** Each figure presents the combination of a tax (current *vs.* proposed) and a season; the $x$ and $y$ axes are quantity and price, respectively. The top row illustrates the two seasons under the **current tax**, where households pay the same tax per therm in both seasons. The shaded gray area gives the deadweight loss (DWL) under this tax. **Proposed tax:** The bottom row doubles the tax in the summer—increasing DWL by the narrow pink region—and halves the tax in the winter—reducing DWL by the purple region. Overall DWL decreases.

# 8 Tables

Table 1: **Prior point estimates:** The price elasticity of demand for residential natural gas

| **Paper** | **Data** | **Estimate** |
|---|---|---|
| Davis and Muehlegger (2010) | US state panel | −0.278 |
| Maddala et al. (1997) | US state panel | −0.09 to −0.18 |
| Garcia-Cerrutti (2000) | Calif. county panel | −0.11 |
| Hausman and Kellogg (2015) | US state panel | −0.11 |
| Herbert and Kreil (1989) | Monthly time series | −0.36 |
| Houthakker and Taylor (1970) | Time series | −0.15 |
| Metcalf and Hassett (1997) | RECS HH panel | −0.08 to −0.71 |
| Meier and Rehdanz (2010) | UK HH panel | −0.34 to −0.56 |
| Rehdanz (2007) | Germany HH panel | −0.44 to −0.63 |
| Hahn and Metcalfe (2021) | California HH panel (CARE) | −0.35 |

*Sources:* Authors and Alberini, Gans, and Velez-Lopez (2011)

Table 2: **Billing data summaries**

| | **Full dataset** | | **Border-area dataset** | |
|---|---|---|---|---|
| | **PG&E** | **SoCalGas** | **PG&E** | **SoCalGas** |
| N. 5-digit zip codes | 597 | 611 | 18 | 18 |
| N. unique households | 5,888,276 | 2,526,503 | 152,418 | 68,407 |
| N. bills | 180,663,705 | 95,335,393 | 3,401,947 | 2,352,141 |
| Approx. value (USD) | $5.71B | $3.28B | $120M | $70.5M |

**Notes:** *Full dataset* refers to all of the PG&E and SoCalGas bills in the data. *Border-area (discontinuity) dataset* refers to the subset of the *full dataset* for households located in the 18 5-digit zip codes served by both utilities during 2010–2014.

Table 3: **Numerical summaries:** Prices, quantities, and other variables of interest

| | **5% Sample of California** | | | **Border-discontinuity sample** | | | | |
|---|---|---|---|---|---|---|---|---|
| | | **Split by utility** | | | **Split by season** | | **Split by CARE** | |
| **Variable** | **Overall** | PG&E | SoCalGas | **Overall** | Winter | Summer | CARE | Non-Care |
| Baseline price | 0.8901 | 0.9823 | 0.7432 | 0.9026 | 0.8836 | 0.9204 | 0.8080 | 0.9811 |
| | [0.1686] | [0.1206] | [0.1242] | [0.1419] | [0.1361] | [0.1448] | [0.0854] | [0.1311] |
| Average price | 1.0138 | 1.1053 | 0.8680 | 1.0211 | 1.0008 | 1.0402 | 0.9086 | 1.1147 |
| | [0.1845] | [0.1439] | [0.1439] | [0.1621] | [0.1583] | [0.1633] | [0.1004] | [0.1430] |
| Marginal price | 1.0206 | 1.1277 | 0.8500 | 1.0387 | 1.0121 | 1.0637 | 0.9338 | 1.1259 |
| | [0.2260] | [0.186] | [0.173] | [0.1983] | [0.1905] | [0.2021] | [0.1448] | [0.1944] |
| Therms | 35.4626 | 37.7541 | 31.8135 | 33.8273 | 50.9544 | 17.7311 | 33.1136 | 34.4204 |
| | [33.7995] | [36.0107] | [29.5791] | [30.7697] | [35.2487] | [11.5803] | [28.7629] | [32.3306] |
| Days | 30.3992 | 30.4282 | 30.3530 | 30.3994 | 30.5876 | 30.2225 | 30.4040 | 30.3955 |
| | [1.4275] | [1.2667] | [1.6505] | [1.3038] | [1.3843] | [1.1966] | [1.2761] | [1.3263] |
| Therms per day | 1.1592 | 1.2355 | 1.0378 | 1.1063 | 1.6588 | 0.5871 | 1.0840 | 1.1249 |
| | [1.0921] | [1.1698] | [0.9426] | [0.9936] | [1.1354] | [0.3838] | [0.9304] | [1.0429] |
| Total bill | 36.8703 | 42.3938 | 28.0747 | 34.9508 | 52.0750 | 18.8573 | 30.3135 | 38.8040 |
| | [39.5758] | [44.0564] | [29.0445] | [33.8812] | [39.8973] | [14.0069] | [27.2567] | [38.1017] |
| (Percent) CARE | 27.43% | 26.35% | 29.15% | 45.38% | 45.00% | 45.74% | 100% | 0% |

**Notes:** Unbracketed values provide the variables' means; bracketed values denote the variables' standard deviations. The *5% sample of California* is based upon 5% of PG&E's and SoCalGas's natural gas bills from 2010–2014, sampling at the 5-digit zip code. The *border-discontinuity sample* represents all bills from PG&E and SoCalGas for the 18 5-digit zip codes served by both utilities from 2010–2014.

Table 4: **Balance on observables:** Comparing utilities' customers across the discontinuity

| | Non-CARE | | | CARE | | |
|---|---|---|---|---|---|---|
| **Variable** | PG&E | SoCalGas | **Diff.** | PG&E | SoCalGas | **Diff.** |
| **Panel A: Summer** | | | | | | |
| Therms consumed | 17.61 | 17.29 | 0.32 | 19.35 | 18.00 | 1.34 |
| | [10.8] | [11.7] | [11.3] | [11.3] | [11.3] | [11.3] |
| Days in bill | 30.31 | 29.97 | 0.34 | 30.29 | 29.96 | 0.33 |
| | [1.16] | [1.36] | [1.28] | [1.16] | [1.36] | [1.22] |
| Allowance | 14.17 | 17.22 | −3.05 | 14.14 | 17.11 | −2.96 |
| | [0.805] | [8.05] | [6.14] | [0.851] | [8.17] | [4.33] |
| Total bill | 21.58 | 16.45 | 5.14 | 19.03 | 13.52 | 5.51 |
| | [14.8] | [12.4] | [13.8] | [12.4] | [9.35] | [11.9] |
| HDDs | 0.16 | 0.25 | −0.08 | 0.14 | 0.26 | −0.12 |
| *(thousands)* | [0.309] | [0.407] | [0.367] | [0.267] | [0.418] | [0.315] |
| *N* | 810,949 | 961,824 | 1,772,773 | 973,063 | 320,082 | 1,293,145 |
| **Panel B: Winter** | | | | | | |
| Therms consumed | 51.40 | 54.07 | −2.67 | 49.60 | 49.94 | −0.34 |
| | [33.8] | [35.7] | [34.8] | [31.1] | [33.1] | [31.6] |
| Days in bill | 30.55 | 30.78 | −0.24 | 30.57 | 30.83 | −0.26 |
| | [1.31] | [1.8] | [1.59] | [1.31] | [1.81] | [1.45] |
| Allowance | 46.70 | 49.07 | −2.37 | 47.16 | 49.68 | −2.52 |
| | [12.8] | [10.7] | [11.8] | [12.4] | [10.4] | [12] |
| Total bill | 59.79 | 50.60 | 9.19 | 45.35 | 36.51 | 8.84 |
| | [41.8] | [36.4] | [39.4] | [30.3] | [26.5] | [29.7] |
| HDDs | 1.69 | 1.73 | −0.04 | 1.70 | 1.75 | −0.05 |
| *(thousands)* | [0.467] | [0.437] | [0.452] | [0.439] | [0.422] | [0.435] |
| *N* | 746,140 | 800,037 | 1,546,177 | 871,795 | 270,198 | 1,141,993 |

**Notes:** Unbracketed values provide the variables' means; bracketed values denote the variables' standard deviations. The standard deviations below the difference column (*Diff.*) are pooled across utilities. The difference column denotes the difference in means across utilities for the given cross-section of data. For example, the rightmost **Diff.** column in **Panel A** gives the difference between the PG&E mean and the SoCalGas mean for CARE households in summer months, $\bar{X}_{\text{PGE}} - \bar{X}_{\text{SCG}}$. *CARE* households participate in the California Alternative Rates for Energy (CARE) program. CARE targets low-income households and provides a 20 percent discount on natural gas bills. Heating degree days (HDDs) are in thousands. We calculate the number of heating degrees for day $t$ with mean temperature $\bar{T}_t$ (in °F) as $\text{HDD}_t = \mathbb{1}\left\{\bar{T}_t < 65\right\} \times \left(65 - \bar{T}_t\right)$. The HDDs variable above is thus $\text{HDDS} = \sum_t \text{HDD}_t/1000$.

Table 5: **OLS Results**: Estimating elasticities, varying the dataset, and fixed effects

| | **Dependent variable:** Log(Consumption, daily avg.) | | |
|---|---|---|---|
| | (1) | (2) | (3) |
| Log(Marginal price) | $0.4698^{***}$ | $0.4346^{***}$ | $0.4276^{***}$ |
| | (0.0106) | (0.0136) | (0.0134) |
| Bill HDDs | T | T | T |
| Household FE | T | T | T |
| Month-of-sample FE | T | T | F |
| City by month-of-sample FE | F | F | T |
| Sample | 5% CA | Border | Border |
| *N* | 12,855,910 | 5,754,088 | 5,754,088 |

**Notes:** Each column denotes a separate regression. Errors are two-way clustered within (1) household and (2) utility by climate-zone by billing-cycle (the level at which price varies). All regressions include heating degree days (HDDs) within the households' billing cycle. Price in the table is the second lag of price, *i.e.*, the prices from two bills prior to the current bill. *Significance levels:* *10%, **5%, ***1%.

Table 6: **First- and second-stage results**:
Instrumenting consumers' prices with the Henry Hub spot price

**Dependent variable:** Log(Consumption, daily avg.)

| | (1) Marginal | (2) Average | (3) Avg. Mrg. | (4) Baseline | (5) Sim. Mrg. |
|---|---|---|---|---|---|
| **Panel A:** First-stage results | | | | | |
| Spot price | 0.4377*** | 0.3698*** | 0.3264*** | 0.4365*** | 0.4117*** |
| | (0.0839) | (0.0548) | (0.0586) | (0.0439) | (0.1068) |
| Spot price × SoCalGas | 0.7833*** | 0.7161*** | 0.9377*** | 0.8231*** | 0.8126*** |
| | (0.0299) | (0.0182) | (0.0194) | (0.0172) | (0.0310) |
| **Panel B:** Second-stage results | | | | | |
| Log(Price) *(instrumented)* | −0.1743*** | −0.1919*** | −0.1500*** | −0.1666*** | −0.1749*** |
| | (0.0645) | (0.0700) | (0.0543) | (0.0600) | (0.0673) |
| First-stage F stat. | 152,617.4 | 353,807.8 | 540,457.5 | 718,718.5 | 116,118.5 . |
| Bill HDDs | T | T | T | T | T |
| Household FE | T | T | T | T | T |
| City mo.-of-sample FE | T | T | T | T | T |
| *N* | 5,782,724 | 5,782,724 | 5,782,724 | 5,782,724 | 4,658,087 |

**Notes:** Each column denotes a separate regression. Errors are two-way clustered within (1) household and (2) utility by climate-zone by billing-cycle (the level at which price varies). All regressions include heating degree days (HDDs) within the houesholds' bill. *(HH) Spot price* refers to the weekly average spot price for natural gas at Louisiana's Henry Hub in the week preceding the utility's price change. Each price in the table is the second lag of price, *i.e.*, the prices from two bills prior to the current bill. *Avg.* or *average* price is the total bill divided by quantity. *Avg. Mrg.* or *average marginal* price denotes the quantity-weighted average of the household's marginal price. *Sim. Mrg.* or *simulated marginal* price is the household's marginal price (using the relevant pricing regime) as a function of the household's historical consumption patterns (lagged bills 10 through 14). As discussed in the empirical strategy section, the numbers of observations differ due to the lags required to calculate the *simulated instrument* for marginal price. *Significance levels:* *10%, **5%, ***1%.

Table 7: **Comparing lags, second-stage results**: Marginal and average prices with HH spot price IV

**Dependent variable:** Log(Consumption, daily avg.)

| | **Marginal Price** | | | | **Average Price** | | | |
|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| | 1 Lead | No lag | 1 Lag | 2 Lags | 1 Lead | No lag | 1 Lag | 2 Lags |
| Log(Price) | 0.0320 | −0.1197* | −0.0211 | −0.1743*** | 0.0366 | −0.1400* | −0.0180 | −0.1919*** |
| *instrumented* | (0.0892) | (0.0715) | (0.0630) | (0.0645) | (0.0971) | (0.0772) | (0.0692) | (0.0700) |
| First-stage F stat. | 113,462.1 | 124,673.8 | 139,532.6 | 152,617.4 | 248,826.0 | 308,230.1 | 329,840.9 | 353,807.8 |
| Bill HDDs | T | T | T | T | T | T | T | T |
| Household FE | T | T | T | T | T | T | T | T |
| City month-of-sample FE | T | T | T | T | T | T | T | T |
| *N* | 5,711,976 | 5,782,724 | 5,782,724 | 5,782,724 | 5,704,716 | 5,782,724 | 5,782,724 | 5,782,724 |

**Notes:** With regard to lags: *No lag* refers to the price for the household's contemporaneous bill; *1 Lag* refers to the price on the household's previous bill; *etc. Avg.* or *average* price is the total bill divided by quantity. *(HH) Spot price* refers to the weekly average spot price for natural gas at Louisiana's Henry Hub in the week preceding the utility's price change. Each column denotes a separate regression. Errors are two-way clustered within (1) household and (2) utility by climate-zone by billing-cycle (the level at which price varies). All regressions include heating degree days (HDDs) within the households' billing period. *Significance levels:* *10%, **5%, ***1%.

Table 8: **Heterogeneity by season or income**:
Second-stage results, instrumenting marginal price with HH spot price

**Dependent variable:** Log(Consumption, daily avg.)

| | Marginal Price | | | |
|---|---|---|---|---|
| | Split by **Season** | | Split by **CARE (Income)** | |
| | (1) Summer | (2) Winter | (3) CARE | (4) Non-CARE |
| Log(Price) | 0.0342 | −0.3662*** | −0.2033*** | −0.1255** |
| *instrumented* | (0.0296) | (0.1071) | (0.0618) | (0.0618) |
| First-stage F stat. | 104,026.5 | 55,379.5 | 57,217.1 | 86,181.0 |
| Bill HDDs | T | T | T | T |
| Household FE | T | T | T | T |
| City month-of-sample FE | T | T | T | T |
| *N obs.* | 5,782,724 | 5,782,724 | 5,782,724 | 5,782,724 |

**Notes:** Columns (1)–(2) result from a single regression; columns (3)–(4) come from a single regression. Errors are two-way clustered within (1) household and (2) utility by climate-zone by billing-cycle (the level at which price varies). All regressions include heating degree days (HDDs) within the households' billing period. We interact fixed effects and the HDD control with indicators for the dimension of heterogeneity (season or CARE). Each price in the table is the second lag of price, *i.e.*, the prices from two bills prior to the current bill. *Summer* includes April through September. *Winter* includes October through March. *CARE* households participate in the California Alternative Rates for Energy (CARE) program. CARE targets low-income households and provides a 20 percent discount on natural gas bills. We estimate the heterogeneity results by splitting the sample along the dimension(s) of heterogeneity and then individually estimating the models. *Significance levels:* *10%, **5%, ***1%.

Table 9: **Heterogeneity by season and income**:
Second-stage results, instrumenting marginal price with HH spot price

**Dependent variable:** Log(Consumption, daily avg.)

| | Marginal Price | | | |
|---|---|---|---|---|
| | (1) Summer CARE | (2) Summer Non-CARE | (3) Winter CARE | (4) Winter Non-CARE |
| Log(Price) | 0.0250 | 0.0625* | −0.4588*** | −0.3647*** |
| *instrumented* | (0.0350) | (0.0330) | (0.1100) | (0.1098) |
| First-stage F stat. | 37,705.5 | 58,925.5 | 22,604.3 | 30,227.7 |
| Bill HDDs | T | T | T | T |
| Household FE | T | T | T | T |
| City month-of-sample FE | T | T | T | T |
| *N* | 5,782,724 | 5,782,724 | 5,782,724 | 5,782,724 |
| *N* in subset | 1,370,158 | 1,841,678 | 1,087,563 | 1,483,325 |

**Notes:** The four columns result from a single regression. Errors are two-way clustered within (1) household and (2) utility by climate-zone by billing-cycle (the level at which price varies). All regressions include heating degree days (HDDs) within the households' billing period. We interact the fixed effects and the bill HDDs control with both dimensions of heterogeneity (season and CARE status). Each price in the table is the second lag of price, *i.e.*, the prices from two bills prior to the current bill. *Summer* includes April through September. *Winter* includes October through March. *CARE* households participate in the California Alternative Rates for Energy (CARE) program. CARE targets low-income households and provides a 20 percent discount on natural gas bills. We estimate the heterogeneity results by splitting the sample along the dimension(s) of heterogeneity and then individually estimating the models. Each price in the table is the second lag of price, *i.e.*, the prices from two bills prior to the current bill. *Significance levels:* *10%, **5%, ***1%.

**For Online Publication**

# A Appendix

## A.1 Appendix figures

Figure A1: **U.S. natural gas pipeline network**

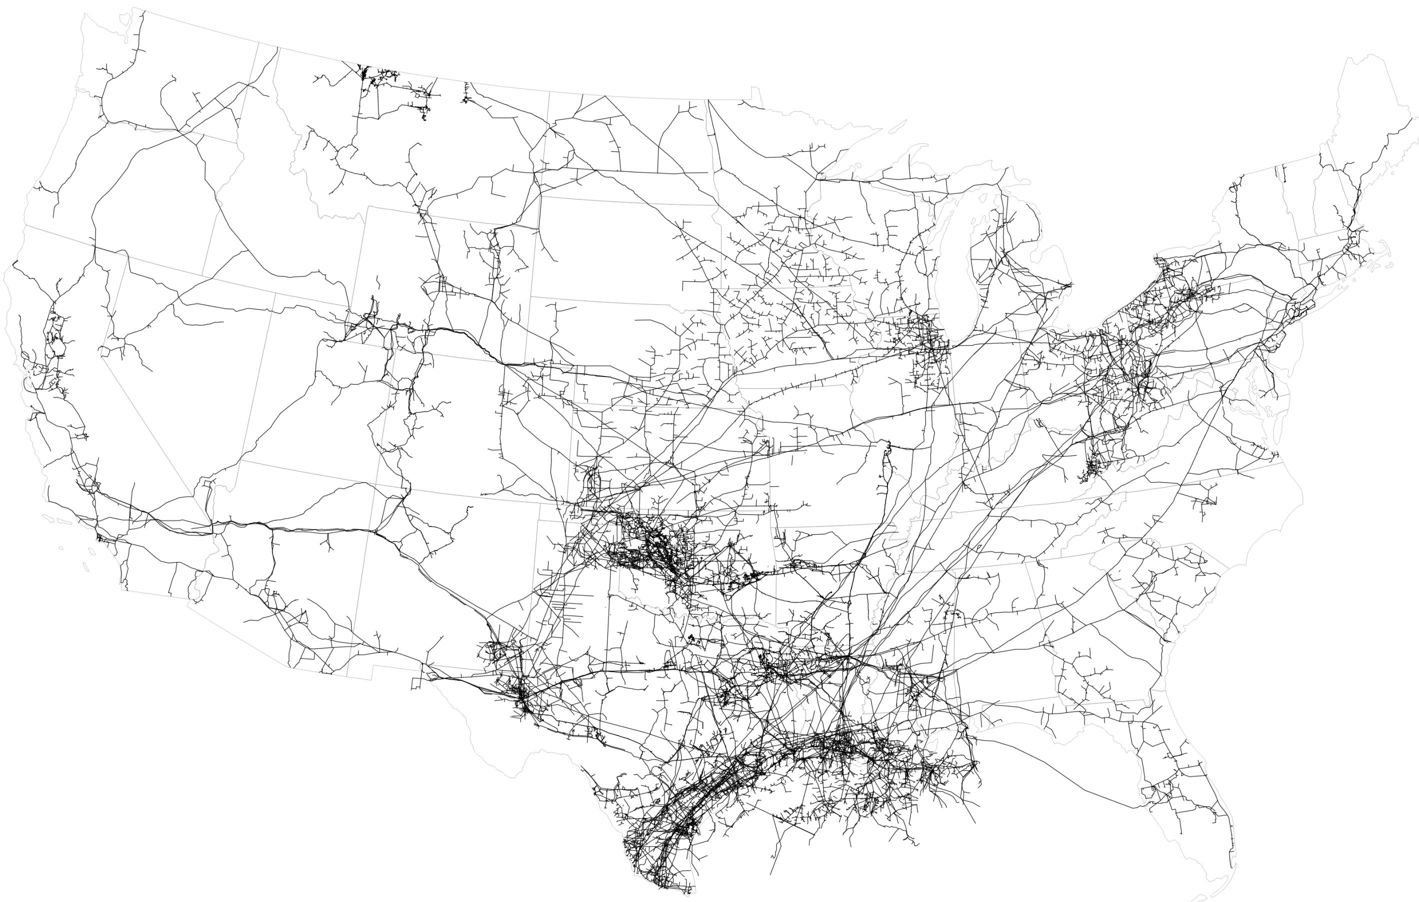

**Notes:** This map depicts the intra- and inter-state natural gas pipeline network for the (continental) United States (in black) overlayed on a map of the (continental) U.S. (light gray). *Source:* U.S. Energy Information Administration

Figure A2: **Natural gas service areas and PG&E's electricity service area**

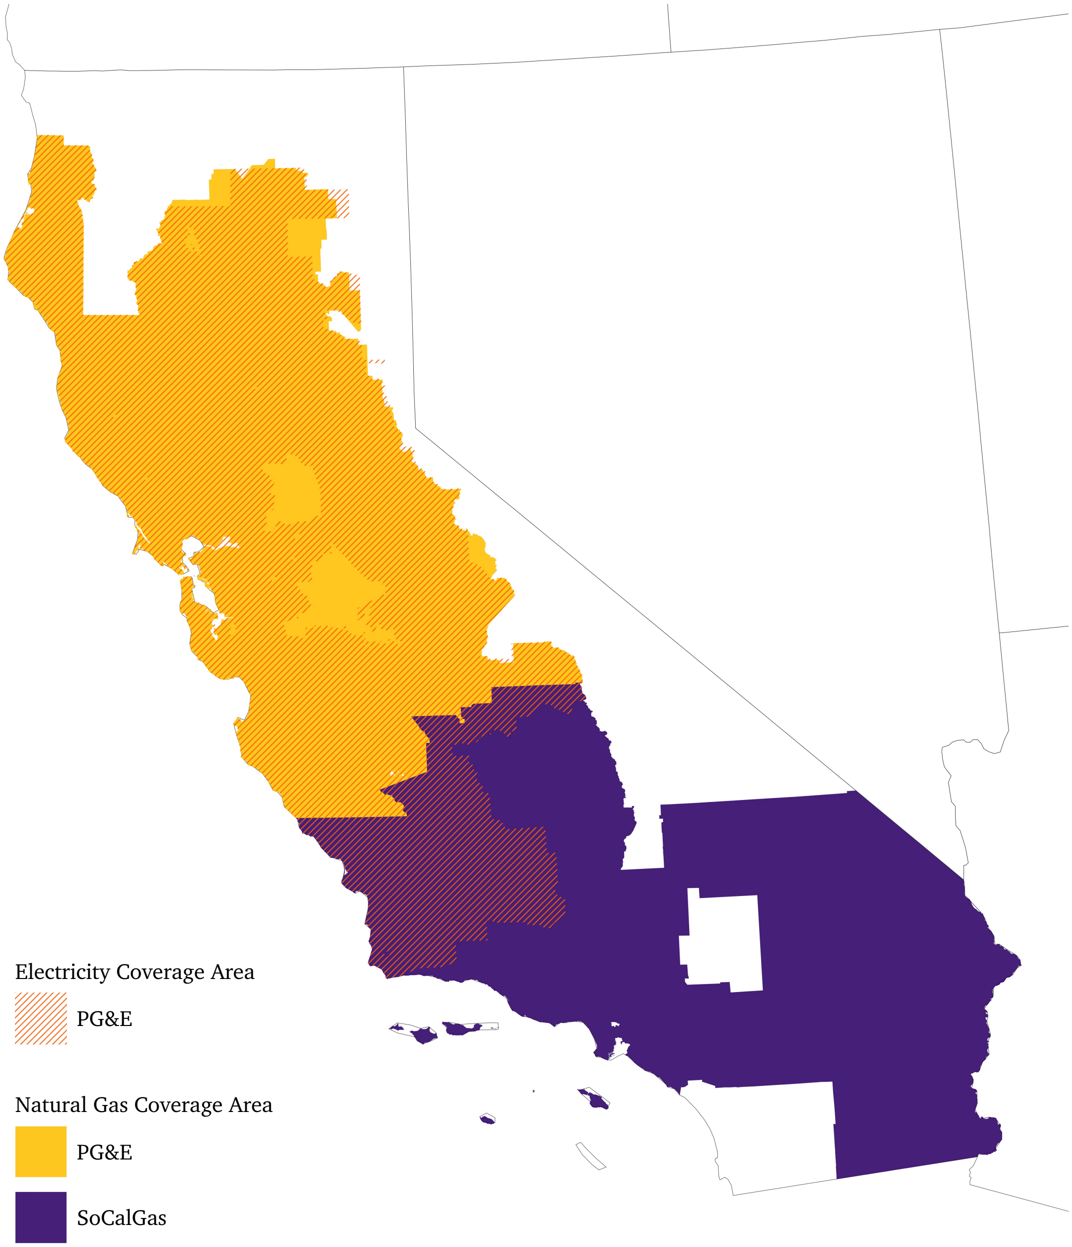


**Notes:** This map illustrates the natural gas service areas covered by PG&E (darker purple) and SoCalGas (lighter gold), as well as PG&E's service area for electricity (diagonally hashed pink). PG&E's electricity service covers both sides of the natural-gas service area. This overlap allows two neighboring households to receive natural gas from two different utilities, while they both receive electricity from PG&E. Note that this map differs slightly from Figure 4 because this map uses the official boundaries from California Energy Commission, whereas Figure 4 maps zip codes contained in our billing data.

Figure A3: **California's 16 CEC climate zones** determine daily allowance within season

**Notes:** The shapefile underlying this map comes from the California Energy Commission (CEC). This map constitutes the CEC's climate-based building zones, which affect a number of energy policies, including households' baseline allowances. (California Energy Commission 2015, 2017)

Figure A4: **Example bill:** PG&E residential natural gas bill

PG&E

## ENERGY STATEMENT

www.pge.com/MyEnergy

| | |
|---|---|
| Account No: | [redacted] |
| Statement Date: | 12/25/2016 |
| **Due Date:** | **01/17/2017** |

### Details of Gas Charges

11/24/2016 - 12/23/2016 (30 billing days)

Service For: [redacted]
Service Agreement ID: [redacted]
Rate Schedule: G1 T Residential Service

| **11/24/2016 – 11/30/2016** | **Your Tier Usage** | **1** | 2 |
|---|---|---|---|

| | | | |
|---|---|---|---|
| Tier 1 Allowance | 12.53 Therms | (7 days x 1.79 Therms/day) | |
| Tier 1 Usage | 2.333330 Therms @ $1.42692 | | $3.33 |
| Gas PPP Surcharge ($0.10197 /Therm) | | | 0.24 |
| Berkeley Utility Users' Tax (7.500%) | | | 0.25 |

| **12/01/2016 – 12/23/2016** | **Your Tier Usage** | **1** | 2 |
|---|---|---|---|

| | | | |
|---|---|---|---|
| Tier 1 Allowance | 41.17 Therms | (23 days x 1.79 Therms/day) | |
| Tier 1 Usage | 7.666670 Therms @ $1.36245 | | $10.45 |
| Gas PPP Surcharge ($0.10197 /Therm) | | | 0.78 |
| Berkeley Utility Users' Tax (7.500%) | | | 0.78 |

**Total Gas Charges** **$15.83**

**Service Information**

| | |
|---|---|
| Meter # | [redacted] |
| Current Meter Reading | 3,036 |
| Prior Meter Reading | 3,026 |
| Difference | 10 |
| Multiplier | 1.046971 |
| Total Usage | 10.000000 Therms |
| Baseline Territory | T |
| Serial | C |

**Gas Procurement Costs ($/Therm)**

| | |
|---|---|
| 11/24/2016 - 11/30/2016 | $0.45875 |
| 12/01/2016 - 12/23/2016 | $0.39428 |

**Notes:** This 30-day bill for a PG&E customer (one of the authors) overlaps two calendar months in 2016: 7 days in November (24–30) and 23 days in December (01–23). Because PG&E's prices vary with the calendar month, PG&E needs to split consumption by calendar month. To achieve this task, PG&E assumes the customer consumed evenly across all days in the bill. Specifically, PG&E calculates that the customer consumed 10 therms and assigns the same amount of consumption to each day during the 30-day period. Thus, PG&E assigns $10 \times 7/30 \approx 2.33$ to November (the consumer spent 7 days in November in this 30-day bill) and $10 \times 23/30 \approx 7.67$ to December (the consumer spent 23 days in November in this 30-day bill).

Figure A5: **PRISM:** Mean temperature raster for 15 June 2010

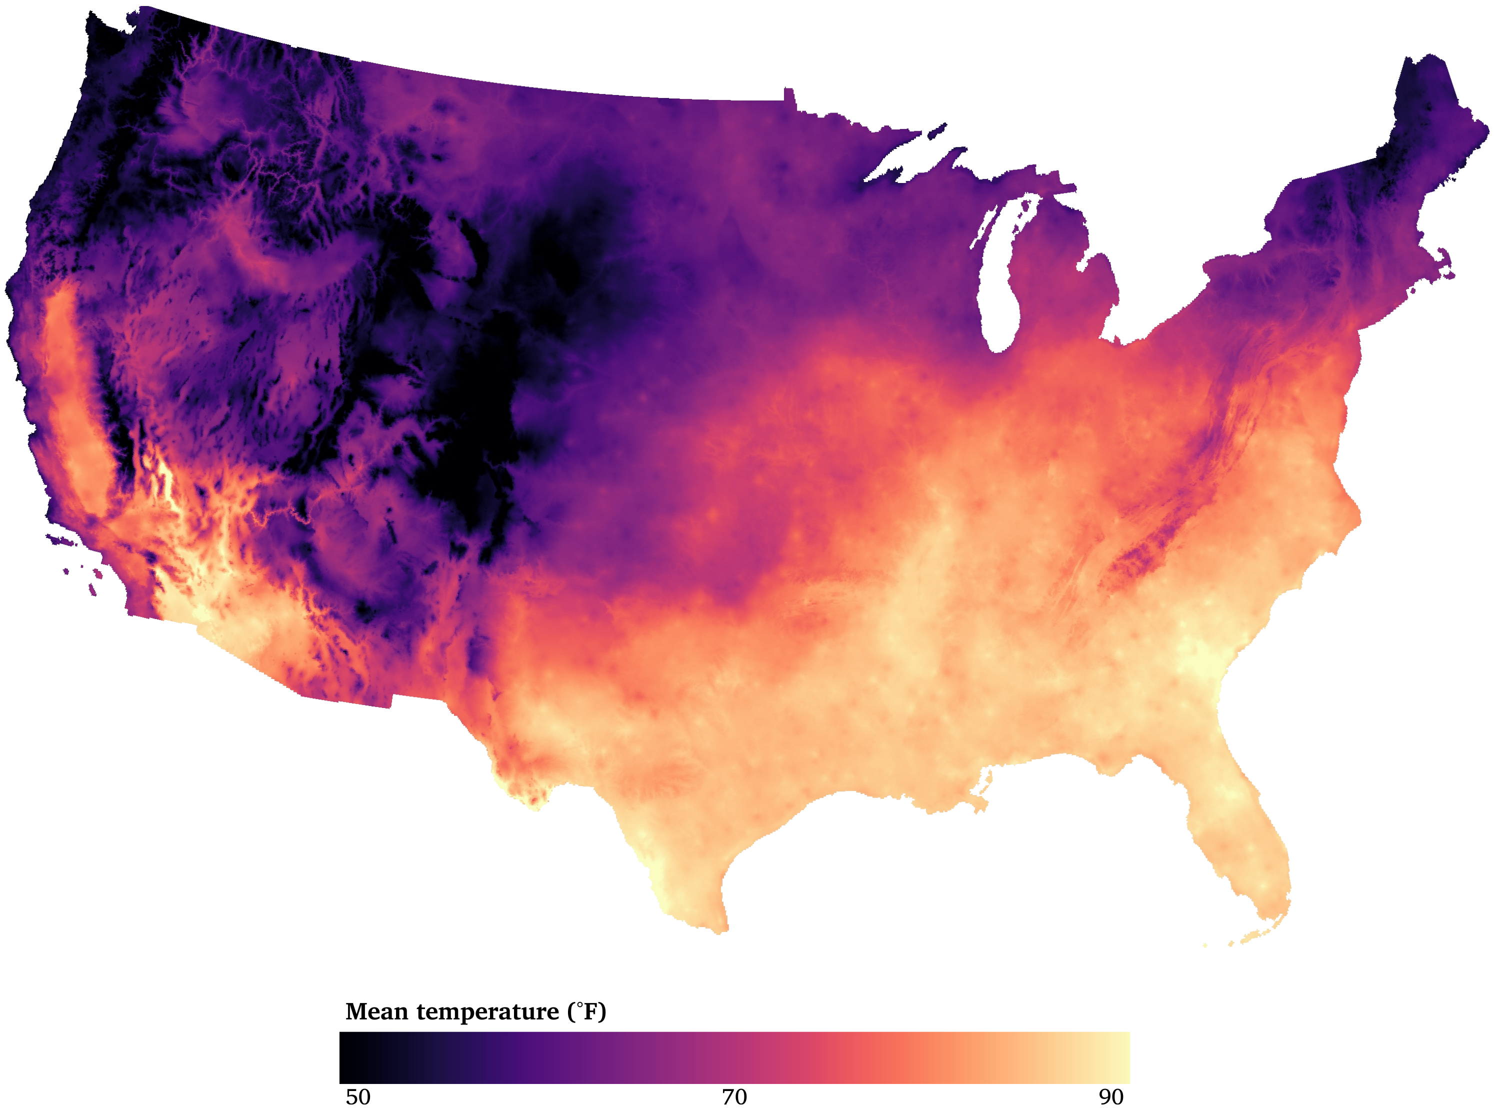

Figure A6: **Expanding the study area:** Zip codes neighboring the study's zip codes

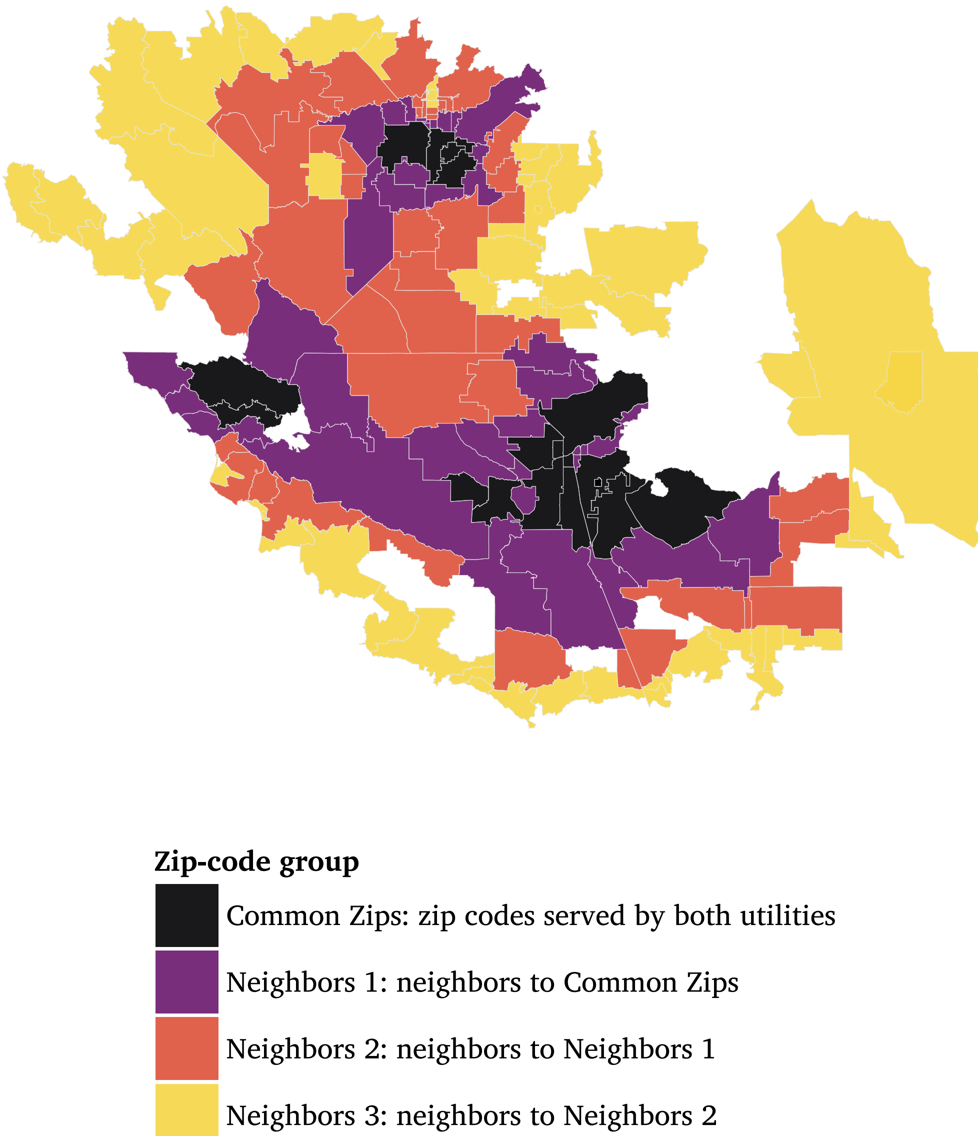


**Notes:** This figure illustrates the four groups of zip codes referenced in Table A15. The groups begin with *Common Zips*—the group in which each zip code receives natural gas service from both PG&E and SoCalGas—and expands by adding each group's immediately proximate neighbors. *E.g.*, *Neighbors 2* consists of all zip codes that neighbor a zip code in *Neighbors 1* (excluding those zip codes already included in another group).

Figure A7: **CARE status and income:** Zip-code income *per capita* *vs.* share of bills with CARE

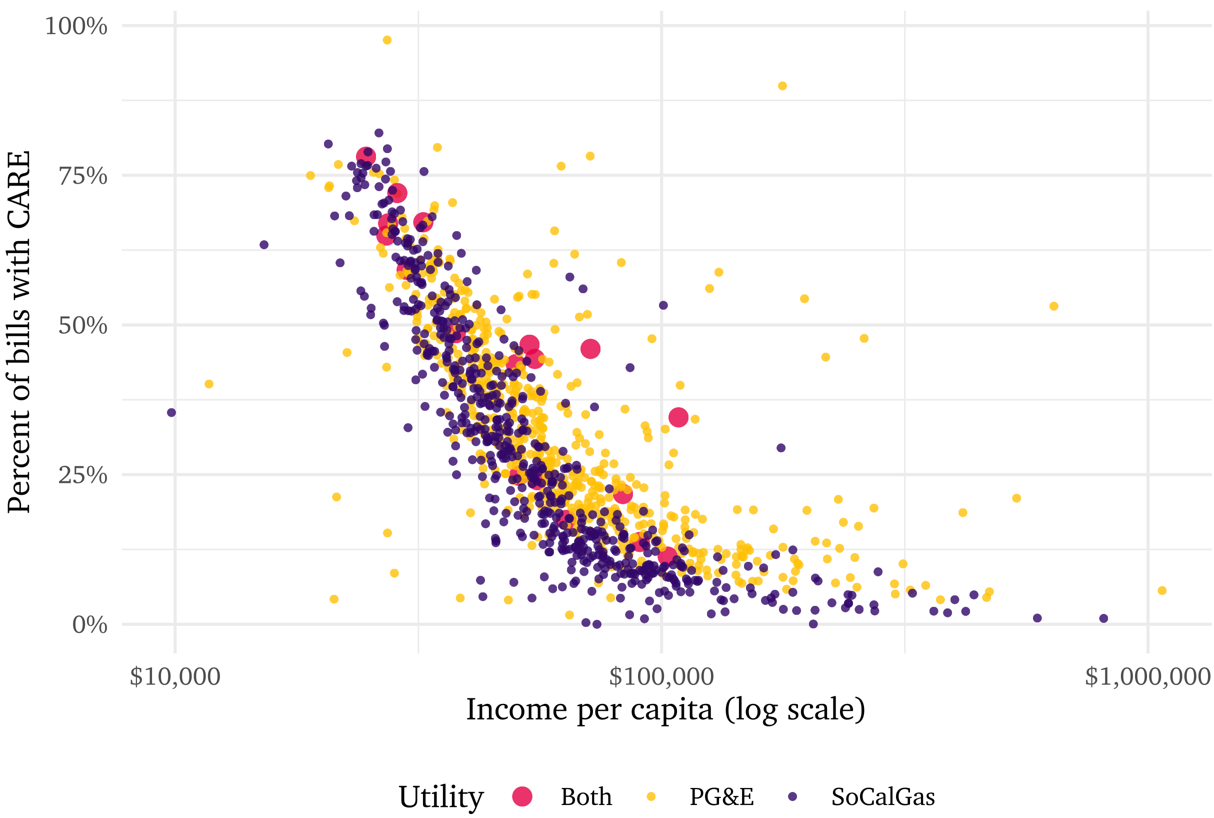


**Notes:** CARE/utility data come from the respective utilities; income data come from California Franchise Tax Board.

## A.2 Appendix tables

Table A1: **Price correlation**: Bivariate correlations between types of prices

| | Type of Price | | | | |
|---|---|---|---|---|---|
| | Marginal | Average | Avg. Mrg. | Baseline | Sim. mrg. |
| Marginal | 1 | | | | |
| Average | 0.8898 | 1 | | | |
| Avg. Mrg. | 0.8628 | 0.9421 | 1 | | |
| Baseline | 0.7901 | 0.942 | 0.9202 | 1 | |
| Sim. mrg. | 0.8503 | 0.849 | 0.8174 | 0.781 | 1 |

**Notes:** *Avg.* or *average* price is the total bill divided by quantity. *Avg. Mrg.* or *average marginal* price denotes the quantity-weighted average of the household's marginal price. *Base* or *baseline* price refers to the price the household pays for its first unit (*therm*) of natural gas. *Sim. Mrg.* or *simulated marginal* price is the household's marginal price (using the relevant pricing regime) as a function of the household's historical consumption patterns (lagged bills 10 through 14).

Table A2: **Testing the simulated instrument**:
Regressing marginal price on *simulated* marginal price

**Dependent variable:** Marginal price

| | (1) | (2) |
|---|---|---|
| Simulated marginal price | 0.1812*** | 0.1667*** |
| | (0.0056) | (0.0054) |
| Bill HDDs | T | T |
| Household FE | T | T |
| City month-of-sample FE | T | T |
| Lags used for sim. inst. | 10–14 | 11–13 |
| *N* | 4,658,087 | 4,567,773 |

**Notes:** Each column denotes a separate regression. Errors are two-way clustered within (1) household and (2) utility by climate-zone by billing-cycle (the level at which price varies). All regressions include heating degree days (HDDs) within the households' billing period. *Sim. Mrg.* or *simulated marginal* price is the household's marginal price (using the relevant pricing regime) as a function of the household's historical consumption patterns (lagged bills 10 through 14 or 11 through 13). As discussed in the empirical strategy section, the numbers of observations differ due to the lags required to calculate the *simulated instrument* for marginal price. *Significance levels:* *10%, **5%, ***1%.

Table A3: **Comparing lags, second-stage results**: Sim. marginal price with HH spot price IV

**Dependent variable:** Log(Consumption, daily avg.)

| | Lag of Simulated Marginal Price | | | | |
|---|---|---|---|---|---|
| | (1) 1 Lead | (2) No lag | (3) 1 Lag | (4) 2 Lags | (5) 3 Lags |
| Log(Price) | 0.0042 | −0.0691 | 0.0148 | −0.1749** | −0.1715** |
| *instrumented* | (0.0882) | (0.0669) | (0.0599) | (0.0673) | (0.0686) |
| First-stage F stat. | 112,639.7 | 121,567.3 | 120,863.7 | 116,118.5 | 111,067.5 |
| Bill HDDs | T | T | T | T | T |
| Household FE | T | T | T | T | T |
| City month-of-sample FE | T | T | T | T | T |
| *N* | 4,912,363 | 4,864,359 | 4,758,259 | 4,658,087 | 4,567,768 |

**Notes:** With regard to lags: *No lag* refers to the price for the household's contemporaneous bill; *1 Lag* refers to the price on the household's previous bill; *etc. Sim. Mrg.* or *simulated marginal* price is the household's marginal price (using the relevant pricing regime) as a function of the household's historical consumption patterns (lagged bills 10 through 14). *(HH) Spot price* refers to the weekly average spot price for natural gas at Louisiana's Henry Hub in the week preceding the utility's price change. Each column denotes a separate regression. Errors are two-way clustered within (1) household and (2) utility by climate-zone by billing-cycle (the level at which price varies). All regressions include heating degree days (HDDs) within the households' billing period. *Significance levels:* *10%, **5%, ***1%.

Table A4: **Comparing lags, second-stage results**: Avg. marginal price with HH spot price IV

**Dependent variable:** Log(Consumption, daily avg.)

| | Lag of Average Marginal Price | | | | |
|---|---|---|---|---|---|
| | (1) 1 Lead | (2) No lag | (3) 1 Lag | (4) 2 Lags | (5) 3 Lags |
| Log(Price) | 0.0315 | −0.0908 | −0.0354 | −0.1500*** | −0.1417*** |
| *instrumented* | (0.0734) | (0.059) | (0.0538) | (0.0543) | (0.0559) |
| First-stage F stat. | 435,512.6 | 477,158.8 | 511,605.5 | 540,457.0 | 505,347.9 |
| Bill HDDs | T | T | T | T | T |
| Household FE | T | T | T | T | T |
| City month-of-sample FE | T | T | T | T | T |
| *N* | 5,704,716 | 5,782,724 | 5,782,724 | 5,782,724 | 5,714,679 |

**Notes:** With regard to lags: *No lag* refers to the price for the household's contemporaneous bill; *1 Lag* refers to the price on the household's previous bill; *etc. Sim. Mrg.* or *simulated marginal* price is the household's marginal price (using the relevant pricing regime) as a function of the household's historical consumption patterns (lagged bills 10 through 14). *(HH) Spot price* refers to the weekly average spot price for natural gas at Louisiana's Henry Hub in the week preceding the utility's price change. Each column denotes a separate regression. Errors are two-way clustered within (1) household and (2) utility by climate-zone by billing-cycle (the level at which price varies). All regressions include heating degree days (HDDs) within the households' billing period. *Significance levels:* *10%, **5%, ***1%.

Table A5: **Comparing lags, second-stage results**: Average price with HH spot price IV

**Dependent variable:** Log(Consumption, daily avg.)

| | Lag of Average Price | | | | |
|---|---|---|---|---|---|
| | (1) 1 Lead | (2) No lag | (3) 1 Lag | (4) 2 Lags | (5) 3 Lags |
| Log(Price) | 0.0366 | −0.1400* | −0.0185 | −0.1919*** | −0.1797** |
| *instrumented* | (0.0971) | (0.0771) | (0.0692) | (0.0700) | (0.0729) |
| First-stage F stat. | 248,826.0 | 308,320.1 | 329,840.9 | 353,807.0 | 295,600.1 |
| Bill HDDs | T | T | T | T | T |
| Household FE | T | T | T | T | T |
| City month-of-sample FE | T | T | T | T | T |
| *N* | 5,704,716 | 5,782,724 | 5,782,724 | 5,782,724 | 5,714,679 |

**Notes:** With regard to lags: *No lag* refers to the price for the household's contemporaneous bill; *1 Lag* refers to the price on the household's previous bill; *etc. Avg.* or *average* price is the total bill divided by quantity. *(HH) Spot price* refers to the weekly average spot price for natural gas at Louisiana's Henry Hub in the week preceding the utility's price change. Each column denotes a separate regression. Errors are two-way clustered within (1) household and (2) utility by climate-zone by billing-cycle (the level at which price varies). All regressions include heating degree days (HDDs) within the households' billing period. *Significance levels:* *10%, **5%, ***1%.

Table A6: **Comparing lags, second-stage results**: Baseline price with HH spot price IV

**Dependent variable:** Log(Consumption, daily avg.)

| | Lag of Baseline Price | | | | |
|---|---|---|---|---|---|
| | (1) 1 Lead | (2) No lag | (3) 1 Lag | (4) 2 Lags | (5) 3 Lags |
| Log(Price) | 0.0279 | −0.1336** | −0.0063 | −0.1666*** | −0.1452** |
| *instrumented* | (0.0819) | (0.0665) | (0.0607) | (0.0600) | (0.0604) |
| First-stage F stat. | 631,126.3 | 649,648.6 | 659,016.2 | 718,718.0 | 715,115.2 |
| Bill HDDs | T | T | T | T | T |
| Household FE | T | T | T | T | T |
| City month-of-sample FE | T | T | T | T | T |
| *N* | 5,711,976 | 5,782,724 | 5,782,724 | 5,782,724 | 5,721,093 |

**Notes:** With regard to lags: *No lag* refers to the price for the household's contemporaneous bill; *1 Lag* refers to the price on the household's previous bill; *etc. Base* or *baseline* price refers to the price the household pays for its first unit (*therm*) of natural gas. *(HH) Spot price* refers to the weekly average spot price for natural gas at Louisiana's Henry Hub in the week preceding the utility's price change. Each column denotes a separate regression. Errors are two-way clustered within (1) household and (2) utility by climate-zone by billing-cycle (the level at which price varies). All regressions include heating degree days (HDDs) within the households' billing period. *Significance levels:* *10%, **5%, ***1%.

Table A7: **Heterogeneity by season or income**:
Second-stage results, instrumenting average price with HH spot price

**Dependent variable:** Log(Consumption, daily avg.)

| | Average Price | | | |
|---|---|---|---|---|
| | Split by **Season** | | Split by **CARE (Income)** | |
| | (1) Summer | (2) Winter | (3) CARE | (4) Non-CARE |
| Log(Price) | 0.0375 | −0.4401*** | −0.2173*** | −0.1377** |
| *instrumented* | (0.0326) | (0.1240) | (0.0761) | (0.0674) |
| First-stage F stat. | 244,242.9 | 113,729.8 | 160,233.1 | 295,231.7 |
| Bill HDDs | T | T | T | T |
| Household FE | T | T | T | T |
| City month-of-sample FE | T | T | T | T |
| *N* | 5,782,724 | 5,782,724 | 5,782,724 | 5,782,724 |

**Notes:** Columns (1)–(2) come from a single regression; columns (3)–(4) come from another regression. Errors are two-way clustered within (1) household and (2) utility by climate-zone by billing-cycle (the level at which price varies). All regressions include heating degree days (HDDs) within the households' billing period. *Summer* includes April through September. *Winter* includes October through March. *CARE* households participate in the California Alternative Rates for Energy (CARE) program. CARE targets low-income households and provides a 20 percent discount on natural gas bills. We estimate the heterogeneity results by splitting the sample along the dimension(s) of heterogeneity and then individually estimating the models. *Avg.* or *average* price is the total bill divided by quantity. Each price in the table is the second lag of price, *i.e.*, the prices from two bills prior to the current bill. *Significance levels:* *10%, **5%, ***1%.

Table A8: **Heterogeneity by season and income**:
Second-stage results, instrumenting average price with HH spot price

**Dependent variable:** Log(Consumption, daily avg.)

| | Average Price | | | |
|---|---|---|---|---|
| | (1) Summer CARE | (2) Summer Non-CARE | (3) Winter CARE | (4) Winter Non-CARE |
| Log(Price) | 0.0270 | 0.0688* | –0.5251*** | –0.4353*** |
| *instrumented* | (0.0378) | (0.0365) | (0.1241) | (0.1263) |
| First-stage F stat. | 106,853.3 | 216,106.7 | 60,255.1 | 90,857.6 |
| Bill HDDs | T | T | T | T |
| Household FE | T | T | T | T |
| City month-of-sample FE | T | T | T | T |
| *N* | 5,782,724 | 5,782,724 | 5,782,724 | 5,782,724 |
| *N* in subset | 1,370,158 | 1,841,678 | 1,087,563 | 1,483,325 |

**Notes:** The four columns result from a single regression. Errors are two-way clustered within (1) household and (2) utility by climate-zone by billing-cycle (the level at which price varies). All regressions include heating degree days (HDDs) within the households' billing period. We interact the fixed effects and the bill HDDs control with both dimensions of heterogeneity (season and CARE status). Each price in the table is the second lag of price, *i.e.*, the prices from two bills prior to the current bill. *Summer* includes April through September. *Winter* includes October through March. *CARE* households participate in the California Alternative Rates for Energy (CARE) program. CARE targets low-income households and provides a 20 percent discount on natural gas bills. We estimate the heterogeneity results by splitting the sample along the dimension(s) of heterogeneity and then individually estimating the models. *Avg.* or *average* price is the total bill divided by quantity. Each price in the table is the second lag of price, *i.e.*, the prices from two bills prior to the current bill. *Significance levels:* *10%, **5%, ***1%.

Table A9: **First-stage results**:
Robustness to specification: Marginal price instrumented with spot price

**Dependent variable:** Log(Marginal price)

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| Spot price | 0.4080*** | 0.4377*** | 0.3631*** | 0.4666*** |
| | (0.0815) | (0.0839) | (0.0649) | (0.0588) |
| Spot price × SoCalGas | 0.7828*** | 0.7833*** | 0.7729*** | 0.7337*** |
| | (0.0299) | (0.0299) | (0.0303) | (0.0372) |
| First-stage F stat. | 151,153.1 | 152,617.4 | 122,411.2 | 34,337.6 |
| Bill HDDs | F | T | T | T |
| Household FE | T | T | T | T |
| City by month-of-sample FE | T | T | F | F |
| City by week-of-sample FE | F | F | T | F |
| Zip by week-of-sample FE | F | F | F | T |
| *N* | 5,782,724 | 5,782,724 | 5,782,724 | 5,782,724 |

**Notes:** Each column denotes a separate regression. Errors are two-way clustered within (1) household and (2) utility by climate-zone by billing-cycle (the level at which price varies). All regressions include heating degree days (HDDs) within the households' billing period. Each price in the table is the second lag of price, *i.e.*, the prices from two bills prior to the current bill. *(HH) Spot price* refers to the weekly average spot price for natural gas at Louisiana's Henry Hub in the week preceding the utility's price change. *Significance levels:* *10%, **5%, ***1%.

Table A10: **Second-stage results**:
Robustness to specification: Marginal price instrumented with spot price

**Dependent variable:** Log(Consumption, daily avg.)

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| Log(Marginal price) | −0.3285*** | −0.1743*** | −0.1434** | −0.1448*** |
| *instrumented* | (0.0798) | (0.0645) | (0.0578) | (0.063) |
| First-stage F stat. | 151,153.1 | 152,617.4 | 122,511.2 | 34,337.6 |
| Bill HDDs | F | T | T | T |
| Household FE | T | T | T | T |
| City by month-of-sample FE | T | T | F | F |
| City by week-of-sample FE | F | F | T | F |
| Zip by week-of-sample FE | F | F | F | T |
| *N* | 5,782,724 | 5,782,724 | 5,782,724 | 5,782,724 |

**Notes:** Each column denotes a separate regression. Errors are two-way clustered within (1) household and (2) utility by climate-zone by billing-cycle (the level at which price varies). All regressions include heating degree days (HDDs) within the households' billing period. Each price in the table is the second lag of price, *i.e.*, the prices from two bills prior to the current bill. *Significance levels:* *10%, **5%, ***1%.

Table A11: **Second-stage results**:
Robustness to specification: Average price instrumented with spot price

| **Dependent variable:** Log(Consumption, daily avg.) | | | | |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| Log(Average price) | $-0.3585^{***}$ | $-0.1919^{***}$ | $-0.1662^{***}$ | $-0.1551^{***}$ |
| *instrumented* | (0.0857) | (0.0700) | (0.0626) | (0.0571) |
| First-stage F stat. | 354,196.9 | 353,807.8 | 289,495.4 | 91,329.2 |
| Bill HDDs | F | T | T | T |
| Household FE | T | T | T | T |
| City by month-of-sample FE | T | T | F | F |
| City by week-of-sample FE | F | F | T | F |
| Zip by week-of-sample FE | F | F | F | T |
| *N* | 5,782,724 | 5,782,724 | 5,782,724 | 5,782,724 |

**Notes:** Each column denotes a separate regression. Errors are two-way clustered within (1) household and (2) utility by climate-zone by billing-cycle (the level at which price varies). All regressions include heating degree days (HDDs) within the households' billing period. Each price in the table is the second lag of price, *i.e.*, the prices from two bills prior to the current bill. *Avg.* or *average* price is the total bill divided by quantity. *Significance levels:* *10%, **5%, ***1%.

Table A12: **Second-stage results**:
Robustness to specification: Avg. mrg. price instrumented with spot price

| **Dependent variable:** Log(Consumption, daily avg.) | | | | |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| Log(Avg. marginal price) | $-0.2646^{***}$ | $-0.1500^{***}$ | $-0.1357^{***}$ | $-0.1323^{**}$ |
| *instrumented* | (0.0661) | (0.0543) | (0.0483) | (0.0483) |
| First-stage F stat. | 541,171.7 | 540,457.5 | 451,337.9 | 121,389.8 |
| Bill HDDs | F | T | T | T |
| Household FE | T | T | T | T |
| City by month-of-sample FE | T | T | F | F |
| City by week-of-sample FE | F | F | T | F |
| Zip by week-of-sample FE | F | F | F | T |
| *N* | 5,782,724 | 5,782,724 | 5,782,724 | 5,782,724 |

**Notes:** Each column denotes a separate regression. Errors are two-way clustered within (1) household and (2) utility by climate-zone by billing-cycle (the level at which price varies). All regressions include heating degree days (HDDs) within the households' billing period. Each price in the table is the second lag of price, *i.e.*, the prices from two bills prior to the current bill. *Avg. Mrg.* or *average marginal* price denotes the quantity-weighted average of the household's marginal price. *Significance levels:* *10%, **5%, ***1%.

Table A13: **Second-stage results**:
Robustness to specification: Simulated marginal price instrumented with spot price

**Dependent variable:** Log(Consumption, daily avg.)

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| Log(Simulated mrg. price) | $-0.2641^{***}$ | $-0.1719^{**}$ | $-0.1637^{***}$ | $-0.1354^{*}$ |
| *instrumented* | (0.0913) | (0.0743) | (0.0618) | (0.0767) |
| First-stage F stat. | 97,898.2 | 96,831.5 | 91,646.3 | 15,809.6 |
| Bill HDDs | F | T | T | T |
| Household FE | T | T | T | T |
| City by month-of-sample FE | T | T | F | F |
| City by week-of-sample FE | F | F | T | F |
| Zip by week-of-sample FE | F | F | F | T |
| *N* | 4,864,359 | 4,864,359 | 4,864,359 | 4,864,359 |

**Notes:** Each column denotes a separate regression. Errors are two-way clustered within (1) household and (2) utility by climate-zone by billing-cycle (the level at which price varies). All regressions include heating degree days (HDDs) within the households' billing period. Each price in the table is the second lag of price, *i.e.*, the prices from two bills prior to the current bill. *Sim. Mrg.* or *simulated marginal* price is the household's marginal price (using the relevant pricing regime) as a function of the household's historical consumption patterns (lagged bills 10 through 14). *Significance levels:* *10%, **5%, ***1%.

Table A14: **Second-stage results**:
Robustness to specification: Baseline price instrumented with spot price

**Dependent variable:** Log(Consumption, daily avg.)

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| Log(Baseline price) | $-0.3136^{***}$ | $-0.1667^{***}$ | $-0.1429^{***}$ | $-0.1358^{***}$ |
| *instrumented* | (0.0726) | (0.0600) | (0.0533) | (0.0497) |
| First-stage F stat. | 724,660,0 | 718,718.5 | 598,286.8 | 190,744.7 |
| Bill HDDs | F | T | T | T |
| Household FE | T | T | T | T |
| City by month-of-sample FE | T | T | F | F |
| City by week-of-sample FE | F | F | T | F |
| Zip by week-of-sample FE | F | F | F | T |
| *N* | 5,782,724 | 5,782,724 | 5,782,724 | 5,782,724 |

**Notes:** Each column denotes a separate regression. Errors are two-way clustered within (1) household and (2) utility by climate-zone by billing-cycle (the level at which price varies). All regressions include heating degree days (HDDs) within the households' billing period. Each price in the table is the second lag of price, *i.e.*, the prices from two bills prior to the current bill. *Base* or *baseline* price refers to the price the household pays for its first unit (*therm*) of natural gas. *Significance levels:* *10%, **5%, ***1%.

Table A15: **Second-stage results**: Extending the set of zip codes to neighboring zip codes

**Dependent variable:** Log(Consumption, daily avg.)

| | Marginal Price | | | |
|---|---|---|---|---|
| | (1) Common Zips | (2) Neighbors 1 | (3) Neighbors 2 | (4) Neighbors 3 |
| Log(Marginal price) | $-0.1743^{***}$ | $-0.1543^{***}$ | $-0.0806^{**}$ | $-0.0616^{*}$ |
| *instrumented* | (0.0645) | (0.0479) | (0.0402) | (0.0352) |
| First-stage F stat. | 152,617.4 | 240,064.9 | 413,472.4 | 604,498.9 |
| Bill HDDs | T | T | T | T |
| Household FE | T | T | T | T |
| City by month-of-sample FE | T | T | T | T |
| Levels of neighboring zip codes | 0 | 1 | 2 | 3 |
| *N* | 5,782,724 | 12,648,457 | 21,305,253 | 30,606,756 |

**Notes:** *Common zips* refers the set of zip codes in which each zip code receives natural gas from both PG&E and SoCalGas. *Neighbors 1* includes the *common zips* and the zip codes that immediately neighbor the common zips. *Neighbors 2* adds the neighbors of these neighbors (adding the neighbors of *Neighbors 1*). *Neighbors 3* adds the neighbors of *Neighbors 2*. Figure A6 depicts these sets of zip codes. Each column denotes a separate regression. Errors are two-way clustered within (1) household and (2) utility by climate-zone by billing-cycle (the level at which price varies). All regressions include heating degree days (HDDs) within the households' billing period. Each price in the table is the second lag of price, *i.e.*, the prices from two bills prior to the current bill.
*Significance levels:* *10%, **5%, ***1%.

Table A16: **Second-stage results**:
Alternative instruments: Marginal price instrumented with lagged eastern HDDs

**Dependent variable:** Log(Consumption, daily avg.)

| | Polynomials in HDDs | | | | | *Binned* |
|---|---|---|---|---|---|---|
| | (1) 2nd deg. | (2) 3rd deg. | (3) 4th deg. | (4) 5th deg. | (5) 6th deg. | (6) 5 bins |
| Log(Marginal price) | −0.1990* | −0.2009** | −0.1872** | −0.1436* | −0.1724** | −0.1500** |
| *instrumented* | (0.1028) | (0.0888) | (0.0838) | (0.0761) | (0.0729) | (0.0763) |
| Bill HDDs | T | T | T | T | T | T |
| Household FE | T | T | T | T | T | T |
| City by month-of-sample FE | T | T | T | T | T | T |
| *N* | 5,782,724 | 5,782,724 | 5,782,724 | 5,782,724 | 5,782,724 | 5,782,724 |

**Notes:** The instrument for marginal price *in this table* is the number of population-weighted heating-degree days (HDDs) in states east of the Mississippi River, six months prior to the household's relevant marginal price. Each column denotes a separate 2SLS regression with the same second stage but varying the functional form of the HDDs in the first stage. Columns (1)–(5) use polynomials varying in degree from second through sixth; column (6) uses a semi-parametric binned specification where bins are based upon HDD quintiles. Errors are two-way clustered within (1) household and (2) utility by climate-zone by billing-cycle (the level at which price varies). All regressions include heating degree days (HDDs) within the households' billing period. Each price in the table is the second lag of marginal price, *i.e.*, the prices from two bills prior to the current bill. *Significance levels:* *10%, **5%, ***1%.

Table A17: **Billing data description:** Columns within the billing data

| **Feature name** | **Description** |
|---|---|
| Account ID | Unique identifier for household account with the utility |
| Premise ID | Unique physical-location based identifier |
| Prior read date | Effectively the start date of the bill |
| Current read date | Effectively the end date of the bill |
| Gas rate schedule | Classifies type of customer (and the customer's price regime) |
| Gas usage | Volume of gas consumed during billing period (in therms) |
| Bill revenue | Total bill charged to household for the current billing period |
| Climate band | California Public Utility Commission-based climate region |
| Service address 9-digit zip | Household's 9-digit zip code |
| Service start date | Date on which the household began service |
| Service stop date | Date on which the household ended service |

## A.3 Calculating bills

As discussed in the body of the paper, the majority of bills do not line up with calendar months. Consequently, households' billing periods do not line up with utilities' monthly changes in price (or with changes in daily allowances resulting from changes in seasons). Thus, a single bill will typically span multiple price regimes. The two utilities deal with change in price in subtly different ways. This "problem" results from the fact that neither utility knows households' *daily* consumption.

**PG&E** When a PG&E customer's bill spans multiple calendar months (price regimes), PG&E calculates individual bills for each month. However, because PG&E does not know the daily consumption levels, they assume constant daily consumption throughout the billing period.

**SoCalGas** In the case that a SoCalGas customer's bill spans multiple calendar months (price regimes), SoCalGas computes time-weighted average prices (and allowances) by aggregating the prices and allowances from the calendar months by the number of days the bill spent in each month.

## A.4 Data work

In this section, we describe the exclusion and cleaning choices that we made while preparing the data for analysis. Our R scripts are available upon request, though the data themselves cannot be shared due to agreements with the utilities and the IRB.

**Exclusions**:

- We omitted SoCalGas price data from advice letters 3644, 3680, 3695, 3807, 4053, and 4061, as they were updated by letters 3660, 3697, 3697, 3810, 4055, and 4069, respectively.
- We dropped pre-2008 data (PG&E and prices/allowances), as SoCalGas did not share billing data for pre-2009 bills.
- We trimmed the shortest 2.5% and longest 2.5% bills (resulted in keeping bills of length between 28-34 days). We did this to omit the first or last bills for a household and bills

that were irregular for any other reasons. We applied this requirement of 28–34 days to the current bill and the first through the third lagged bills, because we consider the effect of lagged prices on contemporaneous consumption.

- We dropped bills missing any critical information: number of therms (quantity), revenue, *etc.*
- We dropped bills outside the central 99% of data (*i.e.*, the bill's revenue or volume fall in the bottom 0.5% or in the top 0.5%). Our main results apply this rule for the contemporaneous and the first three lagged bills.
- We dropped bills whose total revenue we could not predict within five percent (using known prices, quantities, and discounts).
- We dropped bills for exactly zero therms.

**CARE status** While the datasets presumably denoted CARE (California Alternate Rates for Energy) households, we found many households not denoted as CARE households whose charges were consistent with CARE pricing (*i.e.*, charges were 80 percent of the standard tariffs). We classified these households as CARE households.