Behavioral responses to two-part tariffs:
Evidence from the introduction of volumetric water pricing *

Praharsh M. Patel[1], Daniel A. Brent[1] and Casey J. Wichman[2,3]

[1]Pennsylvania State University
[2]Georgia Institute of Technology
[3]Resources for the Future

December 2024

Working paper
Preliminary, please do not circulate

Abstract

Economists often advocate for pricing mechanisms based on marginal prices to promote efficient water usage. However, consumers may find it difficult to interpret complex water pricing structures, such as nonlinear two-part tariffs that serve multiple policy objectives. In this study, we examine consumer behavior following the introduction of volumetric pricing in Sacramento, California. Our analysis reveals that volumetric pricing reduces average water consumption by 5%, although there is considerable variation among consumer groups. Low-consumption households, whose bills decreased under the new pricing structure, increased their water usage by 4% despite facing increases in marginal price. These results suggest that consumers are more responsive to changes in their total bills than to marginal prices. We interpret this finding as evidence that consumers respond to fixed bills when making consumption decisions, which has important implications for designing rate structures in systems using two-part tariffs.

Key Words: water demand, pricing, water conservation, behavioral response
JEL Codes: Q25, Q52, Q53, L95

*Corresponding author: Praharsh M. Patel (praharsh@psu.edu).

The authors thank the California Data Collaborative and City of Sacramento Public Utilities for providing data. For valuable comments and feedback, we also thank seminar and conference participants at Camp Resources, Association of Environmental and Resource Economics (AERE) summer conference 2024, Energy and Environmental Economics and Policy Initiative Seminar Series at Pennsylvania State University, and Social Cost of Water Pollution Workshop.

1 Introduction

Efficient water management remains a critical challenge today, with over half of the global population affected by water shortages for at least one month each year (Intergovernmental Panel on Climate Change, 2022). Climate change, urbanization, population growth, and inadequate water storage infrastructure are expected to exacerbate this issue, further decreasing per capita water availability and heightening the urgency for sound water management. In the Western U.S., the growing gap between water supply and demand is highlighted by the recent agreement for Arizona, California, and Nevada to conserve an additional 3 million acre-feet of water by 2026. Over the past two decades, the region has experienced its driest period in over 1,000 years (Williams et al., 2022). Addressing water scarcity involves both supply- and demand-side solutions. Although most cost-effective supply-side innovations have already been exhausted, there is significant potential for water conservation through demand-side management. This approach hinges on understanding how individual and institutional incentives can drive behavioral change. Economists advocate for using prices as an effective means to encourage water conservation and reflect the scarcity of water (Olmstead, 2010). However, political feasibility, social acceptance, and the lack of metering infrastructure to implement volumetric pricing pose significant challenges. Moreover, a critical question persists: how do consumers truly respond to water prices?

In this paper, we evaluate behavioral responses to water pricing by exploiting a wide-scale policy that installs water meters and switches consumers from fixed monthly bills to a two-part tariff containing a fixed monthly fee and the introduction of a volumetric price. Transitioning to volumetric pricing generates a large and salient increase in the volumetric price of water. Standard economic models presume that consumers respond to marginal price and therefore all consumers should weakly decrease consumption when marginal prices increase from zero to a positive level. Our results show that, on average, consumers do decrease consumption after the introduction of volumetric prices. However, the response is heterogeneous depending on how a household's total bill changes. The policy was designed to be revenue neutral, so a household consuming the average amount of water would see little change in their bill. Low-use consumers see their water bills decrease after switching to volumetric rates and these consumers increase their water use even though their marginal price increases. This finding provides evidence consistent with prior research that consumers are not aware of marginal prices (Ito, 2014; Brent and Ward, 2019) and may respond to changes in total bills as opposed to marginal prices (Ito and Zhang, forthcoming). This suggests consumers may respond to the fixed component of utility bills, which creates challenges for water managers that want to use marginal incentives to manage demand. It also questions recent recommendations for using progressive fixed costs to achieve equity goals in utility pricing (Levinson and Silva, 2022) that assume fixed costs do not affect demand.

Our empirical setting is the City of Sacramento. In 2004, California passed a law that required all urban water suppliers to meter and charge customers based on volumetric water consumption by 2025. The motivation for the policy rested on the assumption that metering and volumetric pricing were effective water conservation tools. Due to the cost and logistic requirements of installing water meters, the meter installation roll-out was staggered. Importantly for our research design, household consumption was metered for a minimum of 12 months after the meter was installed, in what the utility referred to as the "tracking period," before the household was switched to volumetric pricing. This design allows us to identify the effect of volumetric pricing by using households with installed meters, but still on the fixed bill, as the control group.

Our results find that the introduction of volumetric pricing reduces water consumption by about 4-8% on average. The range of estimates depends on how we deal with staggered adoption, which in our setting is extreme since households transition almost continuously over the sample period, with some large cohorts that transition contemporaneously. Additionally, the treatment effects are highly heterogeneous. When conditioning across the difference between "structural winners" (who are expected to see a decrease in their bills) vs. "structural losers" (who are expected to see an increase in their bills).[1] This finding suggests that consumers who see their total bill decrease, increase consumption even though the marginal price of water increases. The most logical explanation is that consumers adjust their behavior based on the change in their total bill, including the fixed component.

Responsiveness to fixed costs was theorized by Liebman and Zeckhauser (2004) as "schmeduling" when consumers respond to a simplified version of a complicated price schedule. Liebman

[1] We define a structural winner (loser) as a household whose bill goes down (up) absent any demand response.

and Zeckhauser (2004) describe one such simplification as "ironing" where consumers respond to average costs at a specific portion of the price schedule as opposed to marginal cost. Ito and Zhang (forthcoming) find evidence of this behavior in a similar setting in heating reform in China where consumers switched from flat to volumetric energy prices. Rees-Jones and Taubinsky (2020) find similar behavior in the context of progressive income taxes.

Our findings speak to the recent debate about whether consumers respond to average or marginal prices when facing nonlinear price schedules. While there is evidence that marginal prices matter for residential demand under nonlinear rate structure (i.e., increasing-block rates) (Nataraj and Hanemann, 2011; Elinder et al., 2024), other evidence suggests consumers respond to average price (Ito, 2014; Wichman, 2014). However, as noted in Cook and Brent (2021), the papers arguing that consumers respond to average price use average *volumetric* price, which excludes the fixed fee in a two-part tariff. Brent and Ward (2019) show that consumers do not have accurate information about the marginal price of water and are more informed about their total bill. This literature relates to work on rational inattention since it is costly to understand and define marginal price in two-part increasing block tariffs (Sallee, 2014). However, assuming that consumers respond to average volumetric price requires more information and calculations than the marginal price in many settings, and the most salient incentive is likely the total bill. Responding to the fixed fee is a larger deviation from the standard consumer optimization model than simply misperceiving the correct marginal price (Wichman, 2017), and is consistent with a model of mental accounting where water bills are largely ignored unless they deviate enough from an expected range Thaler (1985).

Our research also contributes to the literature on estimating water demand and how consumer behavior affects water use. Meta-analysis of the demand elasticities literature by Dalhuisen et al. (2003) indicates that the variation in demand elasticities depends on the underlying tariff system, and the estimates also depend on the income in the study area. Recent papers estimate the impact of prices on water demand using household microdata and variation in prices and rate structures (Nataraj and Hanemann, 2011; Wichman, 2014; Wichman et al., 2016; Wichman, 2017; Brent and Ward, 2019; Brent and Wichman, 2022; Elinder et al., 2024; Li and Jeuland, 2023). However, there is still no clear consensus on which price is relevant for modeling consumer demand when facing nonlinear prices. Nataraj and Hanemann (2011) exploited the introduction of an additional block rate for some households served by the City of Santa Cruz Water Department, which almost doubled their marginal price. The affected consumers reduced water consumption and the demand elasticity to marginal price to be -0.12, which suggests that consumers do respond to marginal prices if the deviations are large enough. In contrast, analysis of increasing block rates policy by Wichman (2014) indicates that the consumers respond to average volumetric price with average price demand elasticities to be in a range of -0.43 to -1.14. A recent study by Elinder et al. (2024) uses daily hot-water usage data from a sample of Swedish apartments and the introduction of individualized billing to estimate that consumers reduce water consumption by 18% by the introduction of a marginal price. Understanding how consumers respond to prices is important for forecasting demand and setting rates to encourage conservation. We contribute to this literature with a an empirical set up with city-wide exogenous shock leading to large — ranging from more than 65% decrease to more than 65% increase in monthly total bill amount, and salient change — with monthly billing information in water rates.

The paper is divided in six sections. Section 2 discusses the conceptual background and describes Sacramento's water metering and volumetric pricing policy. We provide details on the data in Section 3. We describe our empirical approach and methodology in Section 4. The main findings are discussed in Section 5. We provide our conceptual explanation to the demand response in 6 before offering concluding remarks in Section 7.

2 Conceptual Background & Policy Setting

Municipal water suppliers are often bound by regulations to maintain budget neutrality. Therefore, large, revenue neutral changes in water rates may generate heterogeneous impacts to consumers with different consumption profiles. Our policy setting shifts consumers from a fixed monthly fee that does not depend on consumption (we call it fixed bill) to a two-part tariff with a smaller fixed fee and a uniform marginal price. Figure 1 graphs the total bill as a function of consumption under both rate structures. It is clear that low users pay less under the two-part tariff whereas high users pay more.

The heterogeneous effect permits an analysis of how consumers respond to marginal prices versus total bills. The bottom two panels of Figure 2 show the effects of the rate change for two

Figure 1: Change in total bill due to two-part tariff

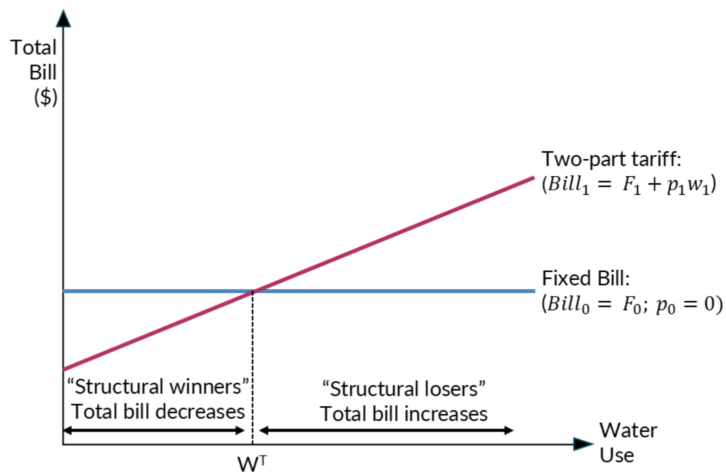

Note: This figure illustrates the rate structure change for Sacramento's household water metering under volumetric pricing. Prior to metering, consumers paid fixed bill based on number of rooms. With the new policy, a two-part (fixed fee + uniform marginal price) tariff was introduced. The fixed fee was reduced to about 30% of the fixed bill, and a uniform marginal prices were applied to water usage. As a result, consumers using less than W^T saw a drop in their total bill, while those consuming more than W^T faced an increase.

classes of consumers (Low and High).[2] The neoclassical model that assumes consumers respond to marginal price, shown in panel (b), predicts that both classes of consumers decrease consumption as the marginal price increases. An alternative model of consumer behavior, shown in panel (c), we assume the consumers respond to average price. Consumer response to either total bill (panel (a)) or average price (panel (c)) shall be different from responding to marginal price. Here, the consumer response is consistent with schmeduling (Liebman and Zeckhauser, 2004) — where Low users are "structural winners" whose total bill likely decrease will increase consumption. "Structural losers" are high users whose total bills increase along with their marginal price and they reduce their consumption.

The presence of heterogeneous responses based on pre-rate change consumption suggests that consumers are responding more to average prices or their perceived total cost rather than strictly to marginal prices. This behavior aligns with a more nuanced understanding of consumer decision-making, where changes in consumption patterns reflect an adaptation to total price rather than a straightforward response to marginal price signals.

[2]Low and High can broadly be defined as consumers previously consuming below or above W^T in panel (a) of Figure 2.

Figure 2: Theoretical basis of anticipated behavioral responses

(a) Consumers responding to total bill

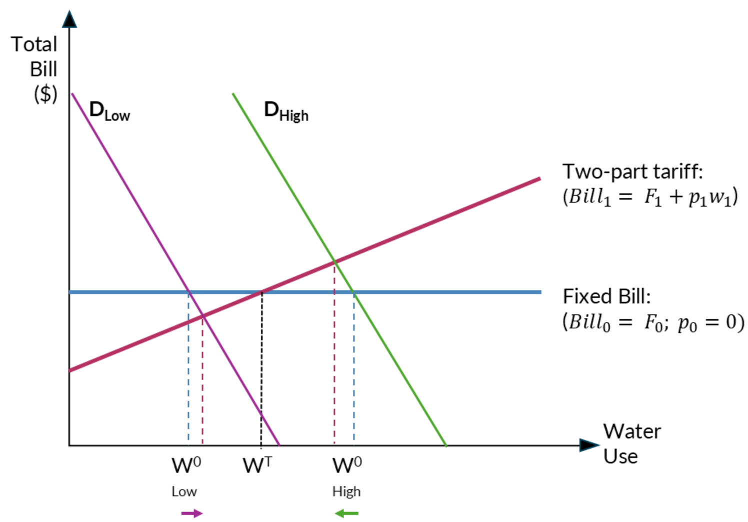

(b) Consumers responding to marginal price

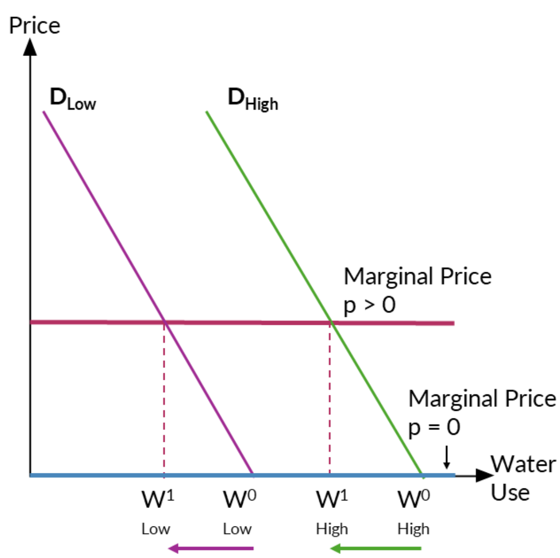

(c) Consumers responding to average price

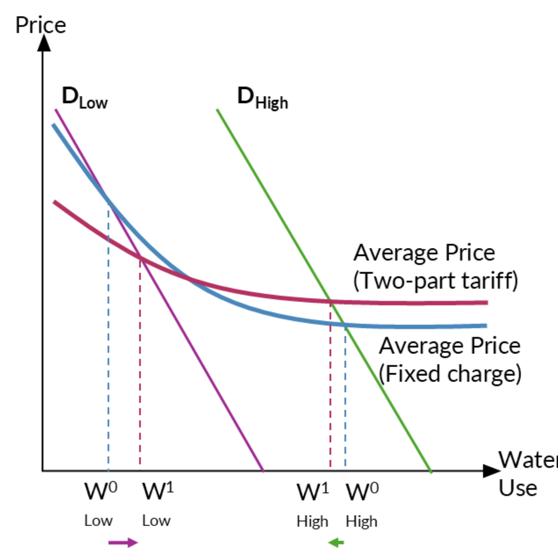

Note: Panel (a) illustrates the expected behavioral response to total bill change once the two-part tariff was introduced. The two part tariff has two component fixed fee which consumer pay irrespective of usage and the variable component based on the water usage and charged uniform marginal price. Consumers with higher water use experiences increased bill hence they shall reduce their consumption whereas the low water users' total bill decreases and shall be expected to increase their water consumption. Panel (b) shows the response to marginal prices, where demand drops across all consumption levels due to higher marginal prices. In Panel (c), consumers respond to average prices, leading to heterogeneous effects

2.1 Water metering and adoption of volumetric pricing

We exploit a large-scale policy in Sacramento, CA, where the City installed water meters and transitioned customers from a fixed bill billing system to a two-part tariff. In this new rate structure, the fixed component (referred as a fixed fee in this paper) of the total bill is lower than the previous fixed bill, and consumers were charged a positive, uniform volumetric rate. In 2004, California passed Assembly Bill 2572 to require all urban water suppliers to meter and charge customers for the actual amount of water they use by 2025. Starting in January 2010, Sacramento households that already had a meter for more than 12 months were transitioned to volumetric pricing. At the time, less than 20% of the City of Sacramento's water connections were metered.[3] As shown in Figure 4, the meters were installed gradually increased.

Meter installation stagnated during the drought period of 2014 to 2016, after which the city began a large-scale meter installation project in May 2017 called the Accelerated Water Meter Program (AWMP) with work scheduled to be completed in 2021. The cost of the AWMP was estimated to be $245 million and financed through increased water rates and municipal bonds. In June 2015, in the middle of a severe drought and before the metering project began, the City of Sacramento's residential customers averaged 208 gallons per capita per day (GPCD), compared to a statewide average of 98 GPCD (California State Water Resources Control Board, CASWRCB, 2024). This level was much greater than other large cities such as Los Angeles (83 GPCD), San Francisco (42 GPCD), and San Diego (71 GPCD). The water metering program had two goals. First, meters provide valuable information on each customer's water consumption as opposed to aggregated data, which could be used to detect inefficient usage or leaks. Second, meters allow the utility to charge customers based on how much water they use, which creates private incentives for conservation.

Before water metering was implemented, consumers were charged a fixed bill. The fixed bill was based on the number of rooms within the household. A room is defined as an area with a minimum of fifty square feet that is structurally or functionally distinct from other rooms or areas in a residential dwelling unit receiving domestic service. Room count was determined by the utilities department and/or the revenue division in accordance with the city's billing

[3]The City's information about metering is available here: https://www.cityofsacramento.org/Utilities/Water/Conservation/water-wise-tools/water-meters.

Figure 3: Sacramento City land Use and Water Metering Rollout

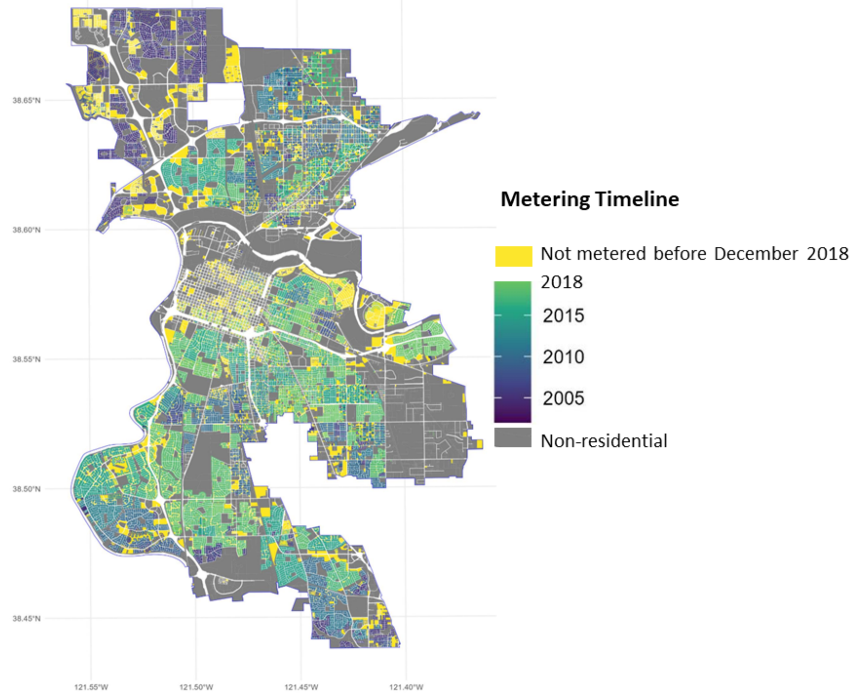

Note: Indicates the timeline for the meter installation in different single family residential properties as per the SWA staggered water metering policy. We have data from January 2008 to December 2018. Households not yet treated in our data are assumed to be treated in 2025.

criteria. There were no surcharges levied during the summer months. This meant that each household paid roughly the same amount every month, regardless of actual water usage. With the introduction of water metering, households were informed for the first time about how much water they consumed.

The metering program was structured so that households initially received a comparative bill, detailing their volumetric consumption and what their charges would be under the volumetric pricing system, although they continued to pay the fixed bill for 12 months. This 12-month period of receiving comparative bills is referred to as the "tracking period." After this period, households were automatically switched to a two-part tariff. The two-part tariff contains a fixed meter charge and a uniform volumetric rate charged on the water consumed. The fixed fee is based on the size of the meter installed. The volumetric rate are derived by the revenues department to recover the variable costs City of Sacramento incurs for maintaining the infrastructure that delivers water services to customers. The billing data indicates that the fixed fee dropped to an average of around 70% of the total bill once consumers switched to volumetric rates.

Figure 4: Rollout of water metering, volumetric pricing and mean consumption

(a) Number of households under tracking and volumetric billing

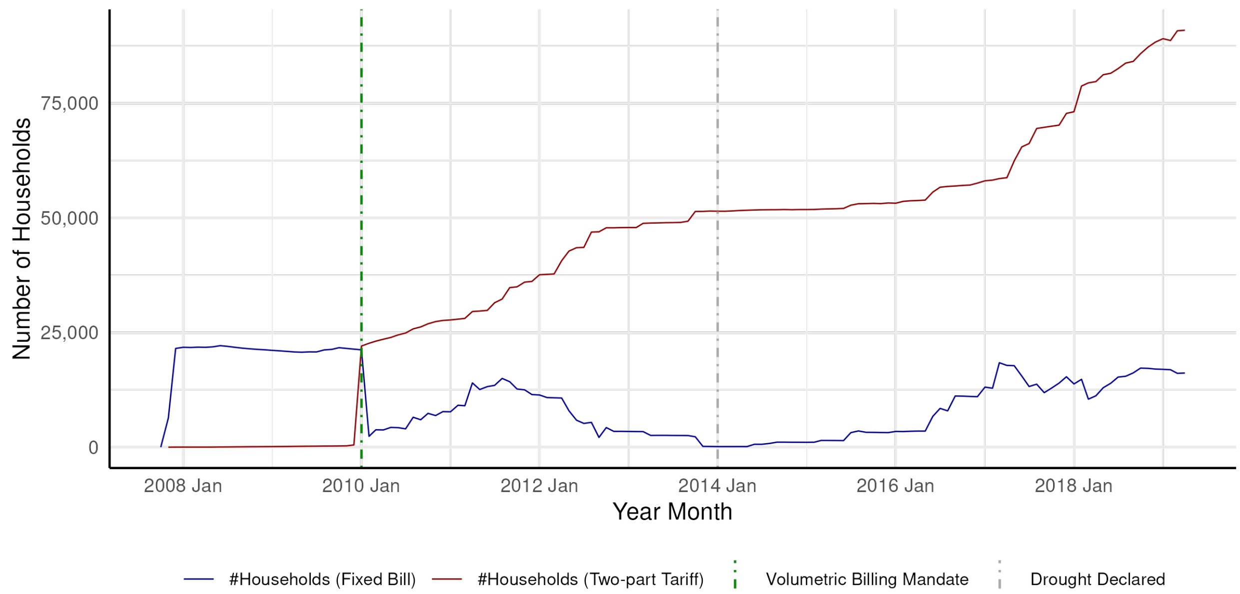

(b) Mean daily water consumption

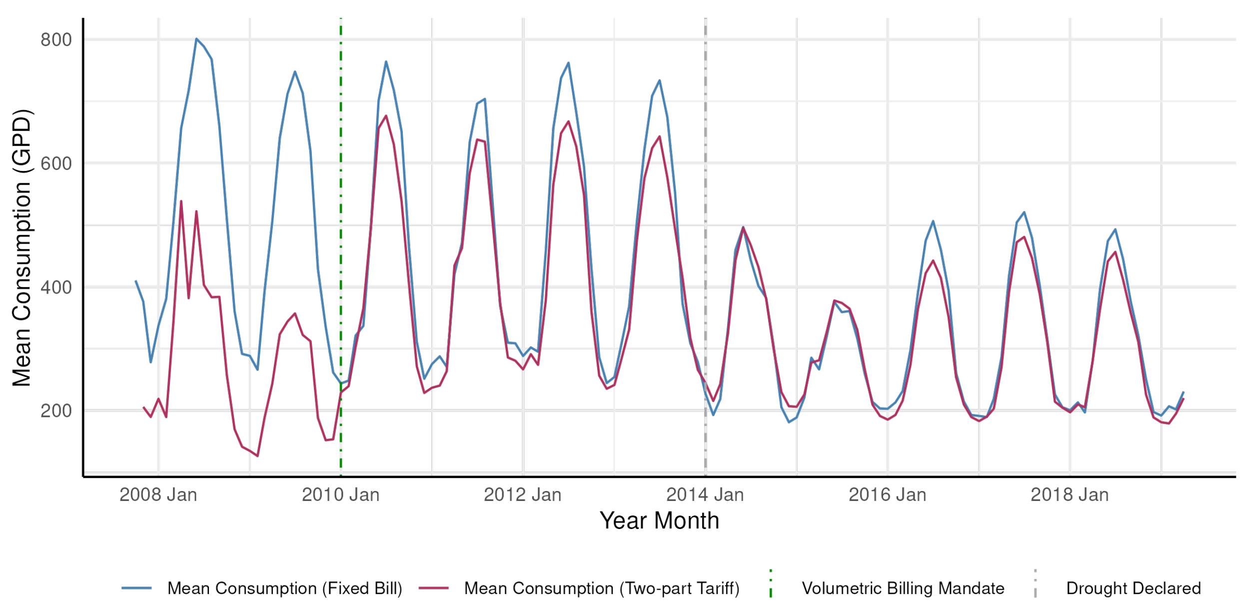

Note: The figure shows the number of households and their mean daily water consumption (gallons per day (GPD)) during the "tracking period" when consumers greceived information on volumetric consumption and a bill comparison between the flat-fee and the two-part tariff with volumetric pricing. Our data panel starts from Jan 2008 and ends in Dec 2018. Before volumetric pricing started in Jan 2010 (indicated by green line), around 20 percent of households were already metered but very few (<300) were charged a volumetric price.

2.2 Staggered roll-out of metering

The City of Sacramento's metering policy had three stages for an individual household. First, a meter is installed on a customer's property. Second, four months after installing the meter, during which the utility can monitor consumption and ensure the meter is functioning properly, the customer starts receiving the "comparative water bills" for approximately 12 months. During this period the customer still pays a flat (i.e., non-volumetric) monthly fee for water, but they observe how much they would pay once they switch to the uniform volumetric rate. This information can help a customer plan for the upcoming switch to volumetric rates. We call this period the "tracking period." We observe most that customers metered between January 2010 and January 2018 received 10 to 13 tracking period bills (see Figure A.3). Third, after the customer received comparative water bills for 12 months, they were permanently switched to the two-part tariff, which we call the "volumetric period." Customers can voluntarily switch to the volumetric rate at any time during the 12 months of comparative water bills. Figure 3 shows the spatial pattern of meter installations in Sacramento. Meters were installed based primarily on the engineering feasibility. Discussions with the utility manager suggested that meter installations were planned to coincide with scheduled maintenance for water distribution lines. For this reason, the timing of meter installations is not likely to be structurally correlated with water consumption. Additionally, our primary empirical strategy focuses on using households with meters installed within 4-8 months of treated households as a control group. Thus, eliminating potential endogeneity in metering roll-out.

The water metering and volumetric billing followed a staggered roll-out and it is unique in three ways: 1) there are many treatment cohorts, 2) there is limited pre-treatment data and thus limited overlapping time periods for counterfactual groups, 3) there is seasonality in water use.

New meters were being installed every month and households were transitioning into the tracking period and switching to volumetric billing continuously over time. Between January 2010 (when volumetric billing starts) and December 2018 (when our data ends), there are 108 treatment cohorts. Some of them have larger cohort sizes than others. Figure 4 shows the total number of households in tracking and volumetric billing periods at each point in time. Additionally, the staggered adoption does not exhibit uniform roll out, which means there are more than 2000 households treated in October 2013 but no treatment cohorts with more than 100 households treated between November 2013 and June 2015.

Household consumption is not measured until meters are installed. So we can only focus on the data available from the tracking and volumetric periods. This design has two implications: First, it greatly limits the available counterfactual groups for the treated groups. For example, a group of households treated in time t, called "Cohort t," will have $t-12$ to $t-1$ observations in the "tracking period." These observations are the pre-treatment period in the conventional difference-in-difference (DiD) design. Our data set ends in December 2018, so households in cohort t have $t+1$ through December 2018 as their post-treatment period. To construct a counterfactual for cohort t, we need tracking period data that overlaps the treatment period and overlapping pre-treatment data to test for parallel trends. This limits the set of potential counterfactual groups for cohorts treated up to 12 months after time t (i.e., cohorts $t+1$ through $t+12$). For example, using cohort $t+4$ as a counterfactual for cohort t results in four periods where the counterfactual cohort has pre-treatment data that overlaps with cohort t's post-treatment data ($t+1$ through $t+4$) and 8 periods with overlapping pre-treatment data to evaluate the parallel trends assumption. Ideally, we would want counterfactual data with perfect overlap with treatment timing. However, the policy design was conducted with an irregular roll-out, so there is varying degrees of comparison cohort observations before and after treatment.

Seasonality in water use also complicates our empirical design. As shown Figure 4, monthly water consumption is lowest during winter months. From March onwards, water use gradually increases to a summer peak in June and July, and it again starts gradually decreasing from August and repeats the cycle from December. Because we do not have a full 12-month overlap between treated and comparison groups, we need to rely on a counterfactual group with sufficient overlap. To do this we adopt a rule that ensures a minimum of four months of overlap between treatment and counterfactual consumption. Thus, we can use data from treatment groups $t+4$ to $t+8$ to ensure reasonable overlap in tracking period for cohort t's pre-treatment and post-treatment periods (see the illustration in Figure A.4).

The uniqueness of our design limits our ability to check for long-term treatment effects. First, we do not have enough counterfactual observation beyond 6 months, and secondly the seasonality aspects demands we compare the treated and control in a way that they atl east have sufficient overlap for pre-/post- treatment period. To make a fair comparison between treatment and control observations, we focus on short-term (within 6 months) treatment effects only.

3 Data

3.1 Household billing records

Our primary data are billing records obtained through a partnership with the California Data Collaborative (CaDC). The CaDC is a non-profit that works as an intermediary between researchers and water utilities and produces original analyses of California water data. Through this partnership, we obtained the universe of billing data for the City of Sacramento, including unique records of billed amounts and consumption from each account in the City's database. We use data for single-family residential properties in our analysis. The data distinguishes water accounts from the property. If property ownership is changed, the new resident will get a new unique utility billing id. This allows us to track consumption and billed amounts for each account. In our data, we observe 107,929 single-family properties and 150,823 unique accounts. The data include the total billed amount (whether on flat-rate or two-part tariff), water consumption during the tracking period and volumetric periods, and the volumetric price and fixed fee during the volumetric period. We also observe any changes to the bill due to utility assistance programs and vacancy adjustments during the flat period. Although the utility did monitor consumption for 2–4 months after installing the meters, but prior to sending comparative bills, we do not observe these records as they did not enter the billing system.

Our data include all billing records from January 2008 through December 2018. Figure 4 shows the transition over time to the different rates. The transition to volumetric rates began primarily in January 2010. After that time, the number of households on volumetric bill consistently increases. If accounts changed within a property, the new account-holder did not receive the tracking period information and they would start by paying volumetric charges only. Thus, in our sample, there are 93,606 such households with the tracking period consumption records available and we observe tracking as well as volumetric consumption data for 73,941 households.[4]

[4]We process the data for anomalous or potentially erroneous records. We remove all bills that are less than zero or over $1000. We also drop water consumption records greater than the 99.9th percentile of consumption. Lastly, we drop very short (<15 days) and very long billing (>90 days) periods; the typical billing period is one month,

Although the metering policy intended for each account to have 12 months of tracking data prior to transitioning billing, in practice the number varies. Consumers who were already metered before volumetric pricing started in January 2010 had longer than 12 months in the tracking period (around 25% of the sample). Around 5% of consumers opted into volumetric pricing at least two months earlier than 12 months leaving fewer tracking period observations.[5] The rest of the households have received 10 to 14 bills during the tracking period.

Customers are billed monthly but the exact number of days within each billing period may vary, thus, we use gallons per day (GPD) to calculate average daily consumption during the billing period. Figure 4 shows the cyclical pattern in water usage. Household water consumption peaks during the summer months typically between May to August and it is lowest during the winter period between November to February. The sharp rise during the summer months is due to outdoor use (such as for watering lawns or gardens). Summer consumption averages roughly three times winter consumption. Figure 4 shows the change of water use over time for households on the flat fee (during the tracking period) and on volumetric rates. Water use is typically higher for households in the tracking period, but it is important to note that the composition of households in the tracking and volumetric groups is also changing over time.

One contemporaneous shock to water use in our period was the severe drought in California that officially started in January 2014 when the governor declared the drought a state of emergency. Water consumption for both groups drops severely in 2014 in response to the drought. After the drought, the consumption gap between households billed on the flat fee vs. volumetric rate also decreased. This decrease is likely due to non-price demand-side management policies such as outdoor water restrictions and public messages about the need to conserve water (Buck et al., 2016; Newkirk, 2014; Brekke, 2014; Maddaus et al., 2020; California Department of Water Resources, 2021).

To highlight how the transition to volumetric billing changes households' total bills, we calculate the predicted bill that a household would pay had they been on volumetric rates during

or approximately 25-35 days (see Figure A.2). We also drop households with < 1 gpd average usage during the pre-treatment period, assuming these might be vacant households. The data cleaning drops fewer than 1% of the observations in our data.

[5]Our data does not reflect the early opt-in decisions but based on the length of tracking period we observe, 5% of total households who transitioned have tracking period less than 300 days and can be considered to have volunteered for the early switch.

the tracking period. The utility provides this information as part of their bill during the tracking period when they receive comparative bills. We only have access to the volumetric price so we use the household's average fixed fee during the volumetric period as an estimate of the "pre-treatment" fixed fee. However, we note that the fixed fees are rising over time in our data. The same households pay higher fixed fee as the metering charges increase over time. Thus the imputed fixed cost during the tracking period might be an overestimate, which would bias our estimate of the predicted bill upwards. In the early periods of the sample, almost all households would pay more under volumetric rates. However, as time progresses and the flat rate increases, more consumers would be better off under volumetric pricing absent any demand response. Typically, low users see smaller bills when switching to volumetric rates even if they do not change their behavior, while high users pay substantially more than they previously paid.

We use this variation between predicted and observed bills to test the heterogeneous response to an increase in marginal prices. We cluster the households within three categories based on the bill differences. As indicated in Figure 5, we take a 15 percentage-point buffer on each side of zero (bill difference = 0), assuming these households see essentially no change in their bill. The average difference in the predicted bill and actual bill during the tracking period for the "Neutral" group is between $-$5.73 and $4.89. Households with predicted bill differences greater than $4.89 threshold are called "Losers," and "Winners" are the households with bill decreases larger than $5.73.

3.2 Household characteristics

We also use CoreLogic's property characteristics data from 2020 (CoreLogic Inc, 2020). Given our focus on understanding how household characteristics explain household-level water consumption, we use three primary variables that explain distinct components of household water consumption. Total number of bathrooms (indoor use), potential lawn area defined as lot size minus ground floor area (potential irrigation water use), and the presence of a swimming pool (outdoor swimming pool use). We also use year built to check for any vintage effects that could affect water efficiency. The data is matched with the household billing data using the APN (Assessors' Parcel Number). Table 1 highlights the key household characteristics within sample subgroups defined based on heterogeneity in bill difference. Winner are low-use households and their houses are older, have relatively small lots and ground floor area, have fewer bathrooms, and are less likely to have a pool. Winner households experience an average 35% decrease in monthly bills, whereas losers are high water consuming households who see their bills increase by almost 50%.

Table 1: Household Characteristics by Heterogeneous Groups

		Bill Difference		
	Overall (1)	**Winners (2)**	**Neutral (3)**	**Losers (4)**
Sample	93,606	30,461	28,082	35,063
	Water Billing			
Water Consumption (Gallons Per Day)	387.2	211.2	333.0	583.7
	(241.7)	(102.3)	(125.5)	(257.7)
Mean Bill (Tracking) ($)	35.1	42.2	33.2	30.3
	(8.3)	(7.4)	(6.4)	(5.6)
Expected Bill Change ($)	0.4	-15.7	-0.1	14.8
	(14.9)	(8.8)	(3.0)	(9.2)
	Property Characteristics			
Assessed Value ('000 $)	253.5	229.4	243.3	282.7
	(162.5)	(155.7)	(153.1)	(171.0)
Year Built*	1971	1960	1972	1980
	(25.4)	(21.5)	(25.0)	(25.1)
Lot Area ('000 Sq ft)	6.9	6.7	6.6	7.3
	(4.9)	(3.7)	(4.8)	(5.7)
Ground Floor Area ('000 Sq ft)	1.3	1.2	1.2	1.4
	(0.4)	(0.4)	(0.4)	(0.5)
Total Bathroom	2.1	1.9	2.0	2.2
	(0.8)	(0.7)	(0.7)	(0.8)
Pool**	8%	7%	7%	8%

Notes: The table presents sample and subgroup means with standard deviation in parentheses. * The "Year Built" variable represents an average year of the house construction. ** Pool indicates percentage of properties with swimming pool.

4 Methodology

The empirical specification builds on a two-way fixed effects model (TWFE) with both account and year-month fixed effects to estimate the effect of transitioning to volumetric billing as shown in Equation 1.

$$ln(GPD)_{it} = \alpha_i + \tau_t + \delta\mathbb{1}[\text{Volumetric}_{it}] + \epsilon_{it} \tag{1}$$

Our dependent variable is the natural log of water consumption (in GPD) for household i in billing period t. $\mathbb{1}[\text{Volumetric}_{it}]$ is an indicator equal to one if a household faced volumetric prices at time t and zero otherwise. Within the difference-in-difference design, δ can be interpreted as the Average Treatment Effect on the Treated (ATT). We include household fixed effects (α_i) to account for any unobservables unique to the household and year-by-month fixed effects (τ_t) to control for any common time-varying factors at the utility level such as changes in weather, marginal prices, and non-price conservation policies over time. Our primary identifying assumption is that the tracking period observations serve as a valid counterfactual for consumption during the volumetric period, conditional on household and time fixed effects.

We also analyze heterogeneity by key factors that could affect how households respond to volumetric billing. First, we examine heterogeneity across the distribution of consumption and expected bill differentials. Many low users may see their bill decrease after transitioning to volumetric pricing, while higher users will experience larger bill increases. We estimate our primary specifications within quintiles of tracking period consumption and our definition of structural winners and losers based on the predicted bill differential. We also test whether the timing of the transition matters. Water use is much higher in the summer, and therefore volumetric bills will be relatively higher than flat rate bills during the summer. We run separate regressions for whether the transition occurred in the summer or winter. Lastly, the 2014 drought dramatically reduced water use due to public messaging campaigns and outdoor watering restrictions. We test whether volumetric pricing was more or less effective at reducing demand before and after the drought.

Our estimation framework is a difference-in-differences (DID) design with staggered adop-

Figure 5: Defining subgroups for heteroegeneity analysis

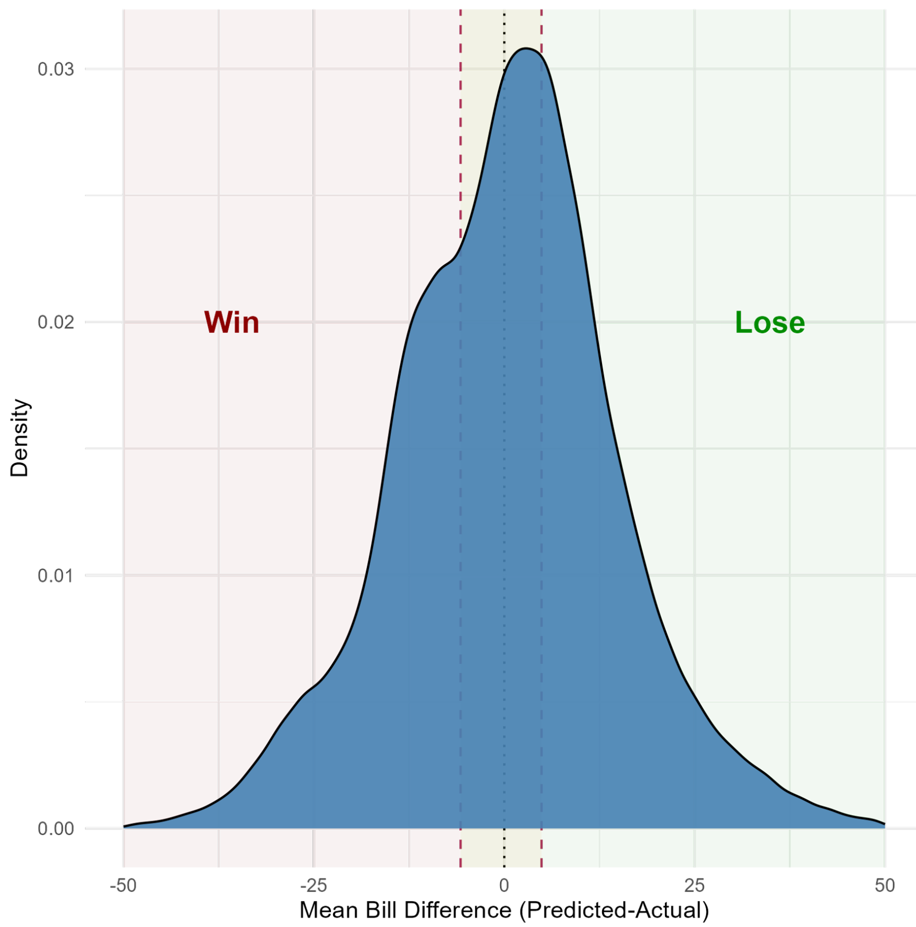

Note: Panel (a) explains the categorization for the treatment effect heterogeneity based on "Mean Bill Difference", defined as mean of the difference in predicted volumetric bill and the actual flat rate bill during tracking period. We define households whose mean bill decreases as winners and increases as losers. We leave 30 percent households who experience negligible change in bill difference as neutral (within dashed lines).

tion. Many recent papers show that the TWFE model is only appropriate when assuming homogeneous treatment effects (De Chaisemartin and d'Haultfoeuille, 2020; Goodman-Bacon, 2021; Callaway and Sant'Anna, 2021; Sun and Abraham, 2021; Dube et al., 2023), which may be unrealistic here. The staggered adoption in our setting is unusual, which prevents us from using the most common estimators proposed in the literature. In our setting, there are many treatment cohorts: every month new households are metered and transitioned to volumetric pricing and some cohorts have large number of treated households whereas most others have very few treated. This challenge makes applying off-the-shelf staggered DiD estimators computationally burdensome.

4.1 Custom matching-based treatment effect estimates

To address the staggered adoption we augment the stacked regression approach— motivated by Cengiz et al. (2019)— with matching to better defend the parallel trend assumption. We describe this for a representative cohort to make the estimation strategy transparent. First, we define a treated cohort as all households who transitioned to volumetric prices within the same month, have at least 4 consecutive bills in the tracking period, and at most 14 bills during the tracking period. We refer to a cohort that transitions in month t as cohort t. The corresponding eligible group is defined as households with at least 4 consecutive months of tracking period observations before and after cohort t transitions to treatment as explained in Section 2.2. For example, "cohort January 2010," a household is conserved an eligible control if and only if it has tracking period bills in the month of September 2009 ($t-4$) through April 2010 ($t+4$). The logic of this approach is identifying cohorts that transitioned at similar times. Overlap ($t-4$ through $t-1$) in the tracking period allows the most direct test for parallel trends that captures both the cyclical and acyclical time trends in water use (such as drought-induced demand response). Tracking-period observations in the control group during the treated cohort's treated period ($t+1$ through $t+4$) are required to estimate treatment effects. The cohort treated in time t together with the corresponding control group forms a "Group t." We estimate an average treatment effect for each such group using standard two-way fixed effects regression and refer to the estimand as the Group Average Treatment Effect (GATE), which is commonly used term in the staggered adoption literature.

We estimate GATEs for 105 groups between January 2008 and September 2018. We cannot estimate the GATE for the October 2018 to December 2018 groups due to fewer than 4 months of post-treatment data available. Since we drop households with more than 14 months or fewer than 4 months of tracking data, we are left with 48,254 households. We visualize the data availability for treated households and their respective counterfactual groups in Figure 6, panel (a). The figure shows the time series of number of eligible households available for each treated cohort, where we define eligibility based on having 4 month overlap during the pre-/post-treatment. We take a cross-section for the two sample cohorts, June 2011 and December 2017 to analyze availability of treated and control group observations. As shown in Figure 6, panels (b) and (c), there are not enough control group observations beyond six months of treatment during pre- or post-treatment periods. Thus, we restrict our estimation to within 6 months of treatment. Once we estimate the GATE for each group, we take a weighted-average of all groups to estimate the final treatment effects, as well as test for parallel trends using event-study plots. The assigned weights are defined based on the total number of treated households in each group divided by the sum of treated households across the groups. For instance, "Group June 2011" has 1,600 households treated out of a total 48,254 that remain after filtering households that have tracking period between 4 to 14 months. Hence, that cohort will receive a weight of 3.3% (= 1,600/48,254).

In each group-wise subset of the data, we estimate the GATE using three methods. First, we develop the group-wise stacked regression approach, estimating GATE by using two-way fixed effect models within each group. We then manually take a weighted average of GATE to calculate the ATT.

Second, we use matching on the pre-treatment outcome for control group within each treated cohort to strengthen the parallel trend assumption. While there is some skepticism about using matching on pre-treatment outcomes in DiD frameworks, Chabé-Ferret (2017) suggests either estimating DiD without matching (as done in the first case) or if matching on outcome variable is used, match on multiple (ideally 3 or more) pre-treatment outcomes. In our design, we are able to match on four pre-treatment observations of monthly water use. We perform nearest-neighbor matching based on these four months of overlapping data during the tracking period. Matching on more than four months reduces the availability of control group observations, especially post-treatment. To exclude households without close matches, we use a caliper of 0.2 standard deviations.

Third, we match on static household characteristics. We match households based on characteristics likely to influence water use, as listed in Table 1: year of construction, number of bathrooms (a proxy for indoor water use), presence of a pool (swimming pool use), and potential lawn area (calculated as the difference between lot size and ground floor area, representing potential irrigation use). For this matching, we use Mahalanobis distance, a preferred method for multivariate matching, and apply a caliper of 0.2 standard deviations to exclude households without close matches.

For these estimators, inference is drawn using standard errors derived through bootstrap

Figure 6: Treated and eligible control groups

(a) Cohort-wise treated and eligible control group

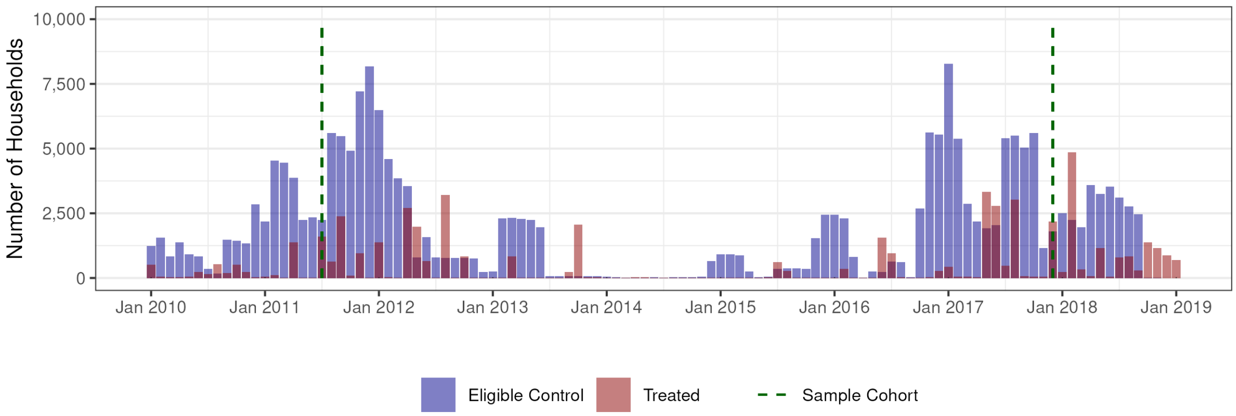

(b) July 2011 Group

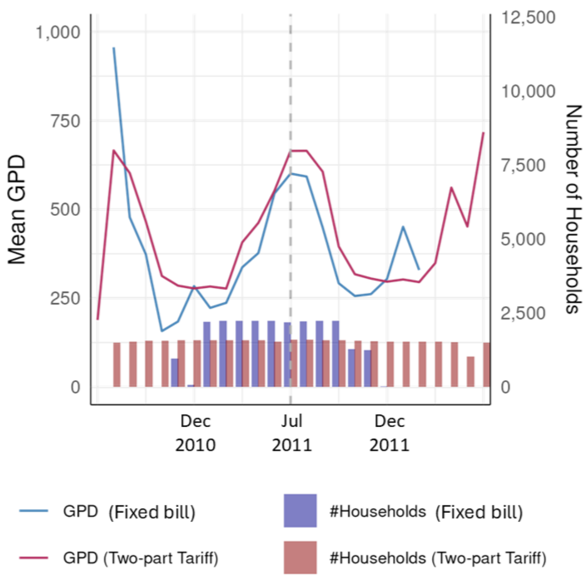

(c) December 2017 Group

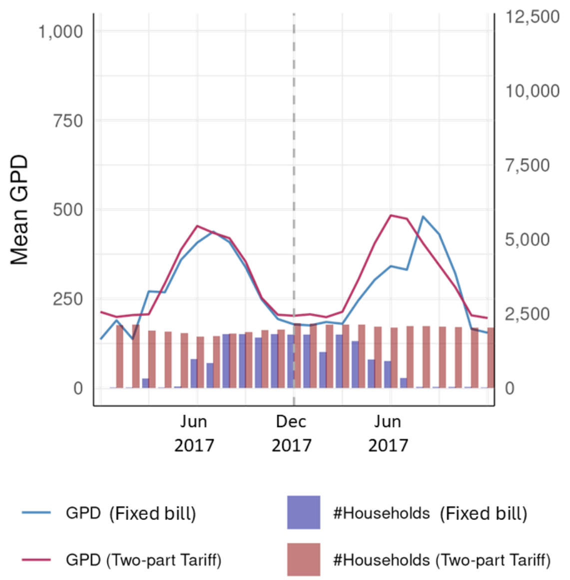

Note: This figure illustrates the staggered adoption of volumetric billing and our approach to constructing cohort comparison groups. Panel (a) shows the number of households treated in a specific cohort. Cohort "Jan 2010" refers to all households who switched from the tracking period to the volumetric period in Jan 2010. The corresponding "Eligible Control" group is defined such that the control group households have tracking period observations for at least 4 consecutive months before and at least 4 months after the switch for the treatment cohort. For example, for Cohort Jan 2010, a household is considered eligible control if and only if it has tracking period bills for Dec 2009 ($t-1$), Nov 2009 ($t-2$), Oct 2009 (t-3), Sep 2009 ($t-4$), and has 4 observations after Jan 2010 (t). Panel b & c show the number of treated and eligible control groups and their mean daily water consumption (gallons per day (GPD)) for early-treated cohorts "Jun 2011" and "Dec 2017", respectively. For the Jun 2011 cohort, the mean consumption for treated and control groups drops due to the seasonal nature of monthly water consumption, whereas the opposite is observed for the Dec 2017 cohort.

method. Bootstrapping includes four steps in our case. First, we use the group-level treated and control group data and a draw a random sample with replacement. Second, we estimate the GATEs for these randomly sampled group-level data and aggregate all GATEs into ATT for the iteration. Third, we iterate same process for 100 times.[6] Fourth, we take standard deviation of 100 ATTs from these iterations and use it as a standard error estimates.

For each of the three estimators, we restrict our sample to within 6 months of the treatment period. As shown in Figure 6b & 6c, we do not have enough valid counterfactual observations beyond 6 months of the treatment and we have no observations beyond 12 months, due to the 12-month tracking period. We check validity of the three new estimators using event study plots to defend the parallel trend assumption. We discuss the results in Section 5.

5 Results

5.1 Treatment Effect Estimates

Table 2 shows the base treatment effect of switching from flat-fee to volumetric pricing. The standard TWFE estimates with all household billing records are presented in column (1), which shows that consumers reduced their daily water consumption by approximately 10% upon switching to volumetric pricing, which is approximately 40 GPD. The TWFE estimates suffer from challenges due to staggered rollout such as dynamic and heterogeneous treatment effects (Dube et al., 2023; Goodman-Bacon, 2021; Callaway and Sant'Anna, 2021; Sun and Abraham, 2021). We try to address this within the standard TWFE estimator by dropping always treated households (column (2)) and limiting the data to within 6 months of treatment (column (3)). These procedures limit the problem of using previously treated households as controls for later cohorts. The magnitude of the ATT estimates decline to 7.6% and 2.8% respectively. The changes in the point estimates warrant using an estimation approach that is robust to staggered adoption.

Given the large data size and too many treated groups in our setup, the available methods to deal with staggered adoption such as Sun and Abraham (2021); Callaway and Sant'Anna (2021); Borusyak et al. (2024); Wooldridge (2022) failed to converge. On top of this, our panel data

[6] For sample iterations, we tried 200/500/1000 iterations as well but the standard deviations did not significantly change for additional iterations and the distributions for ATT estimates with 100 iterations were convincingly normal.

is unbalanced[7] which requires more scrutiny for applying the available methods even if they converge. The seasonality aspects water consumption also limits use of the available methods in our case. Thus, we adopt a novel design to estimate ATT based on weighted averages of the GATEs better suited to our setting as described in Section 4.1. The overall point GATE estimates, presented in columns (4) through (6) vary substantially.

Estimates based on the "No-Matching" and "Matching on Household Characteristics" are statistically insignificant, but matching on pre-treatment outcome estimate indicates overall 5.2% drop in water consumption after switching to volumetric billing. We use event study plots to check the validity of the parallel trend assumption during pre-treatment period. The event study plots for “No-Matching” and “Matching on Household Characteristics” (see Figure 7 (a) and (c)) suggest that the pre-treatment water use across treatment and comparison households are not convincing to assume parallel trends. For “Matching on Pre-treatment Outcomes,” we find a more convincing evidence for parallel trend assumption. It should also be noted that the parallel trends holds beyond the periods where we explicitly match ($t-5$ and $t-6$). Thus, we prefer the matching on pre-treatment outcome estimator for our further analysis and identification of overall GATEs, as well as the heterogeneous treatment effects. This estimator intuition is based on limiting the counterfactual households to households that had a meter installed during a similar time and had similar water use immediately prior to switching to volumetric rates.

[7] Data panel is considered unbalanced if unit of observation, in our case households,does not have observation for all time periods in the panel.

Table 2: Average treatment effects on the treated

	log(Gallons per day)					
	TWFE			**Aggregate-GATE**		
	All Records	No Always Treated	No Always Treated (Within 6 months)	No Matching	Matching on Pre-treatment Outcome	Matching on Household Characteristics
	(1)	(2)	(3)	(4)	(5)	(6)
Volumetric Pricing	-0.100*** (0.002)	-0.076*** (0.002)	-0.028*** (0.004)	-0.033 (0.026)	**-0.052***** **(0.006)**	-0.001 (0.01)
Accounts	150,823	93,606	93,550	93,606	**93,606**	93,606
Treated	73,941	73,941	73,885	-	-	-
Matched	-	-	-	48,254	**41,024**	48,254
Std. Errors	Clustered at account level	Clustered at account level	Clustered at account level	Bootstrapped	**Bootstrapped**	Bootstrapped
Dropped Accounts	-	No Always Treated	No Always Treated	No more than 14 months or less than 4 months tracking	**No more than 14 months or less than 4 months tracking**	No more than 14 months or less than 4 months tracking

Notes: This table compiles Average Treatment on Treated (ATT) estimates using different methods. In column (1), we use all available observations across all households with a standard two-way fixed effect (TWFE) estimator. For columns (2) and (3), we apply TWFE but drop always-treated households, as recommended by the staggered adoption DiD literature, and further limit observations to within 6 months of the switch to volumetric rates in column (3). For columns (4)-(6), we devise group-level estimators. We filter cohort-level treated households so that they were treated in the respective cohort and not tracked less than 4 months or longer than 14 months. The eligible control group consists of households with observations overlapping for at least 4 consecutive pre-treatment and 4 post-treatment months. The estimator in column (4) uses no matching and provides a cohort-weighted average of cohort-level TWFE estimates for treated and eligible control groups. Column (5) adds a 4-month dynamic matching of pre-treatment outcomes (gallons per day) for cohort-level estimates. In column (6), matching is based on observable household characteristics like the year built, number of bathrooms, potential lawn area (lot size-built area), and swimming pool availability. We prefer ATT estimates from column (5) as they effectively address staggered adoption concerns and offer a defensible event study analysis. $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

Figure 7: Event study plots for the overall estimates

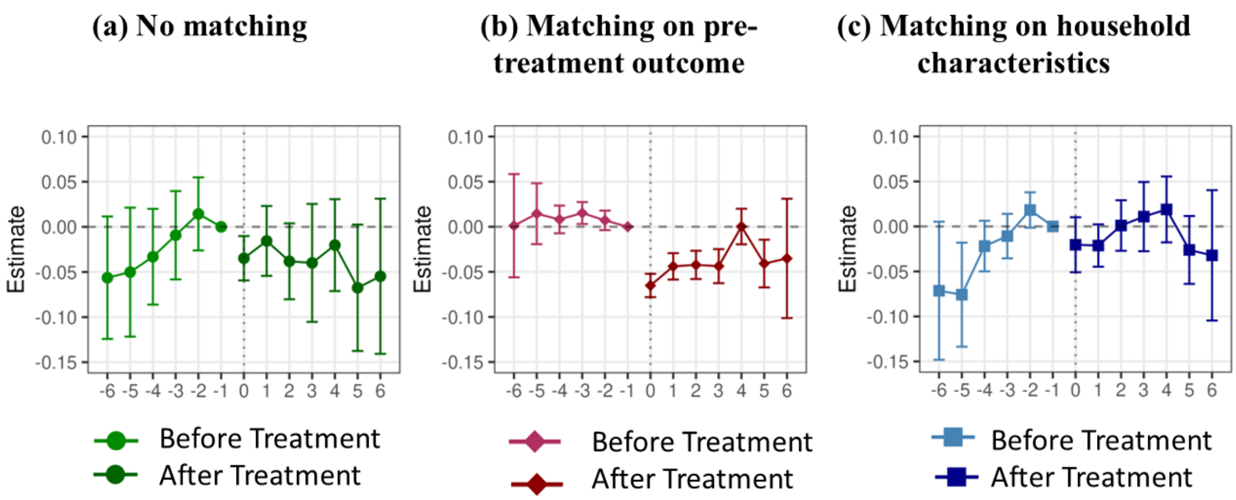

Note: In this figure, we use event study plots for within 6-months ATT based on three different estimation methods to identify the suitable overall treatment effect strategy. The method is suitable if there is no trend before the start of the treatment effect, hence defend the no parallel trend assumption for the difference-in-difference estimation. Parallel trend assumption is hard to defend for panel (a) and (c). Thus, matching on pre-treatment outcome-based estimator which suggests no pre-trend shall be a good fit for ATT estimates.

5.2 Treatment Effect Heterogeneity

The core of the paper attempts to test how different types of consumers, classified based on their consumption quintile and predicted bill differential respond to volumetric pricing. We prefer the ATT estimates derived from the matching on pre-treatment outcome estimator, but test the patterns of heterogeneity across the different methods robust to staggered adoption. If the consumers respond to the marginal prices consumption should go down in all subgroups. However, as shown in Figure 8, (b) and (c), the responses are heterogeneous and robust across all three estimators. The results indicate that the households experiencing a decrease in their total bills (Winners) increased their consumption by 4%, whereas consumers whose bills increased (Losers) decreased their consumption by almost 13%. The results are consistent for households classified based on the pre-treatment consumption categories within quintiles. The lowest quintile consumers, who were more likely to see their bills decrease, increased their consumption by 8%

and the highest consumers reduced their consumption by 15%. This finding indicates consumers are likely not responding to the change in marginal prices but they might be responding to the change in the total bill amount or average price.

Figure 8: ATT Estimates

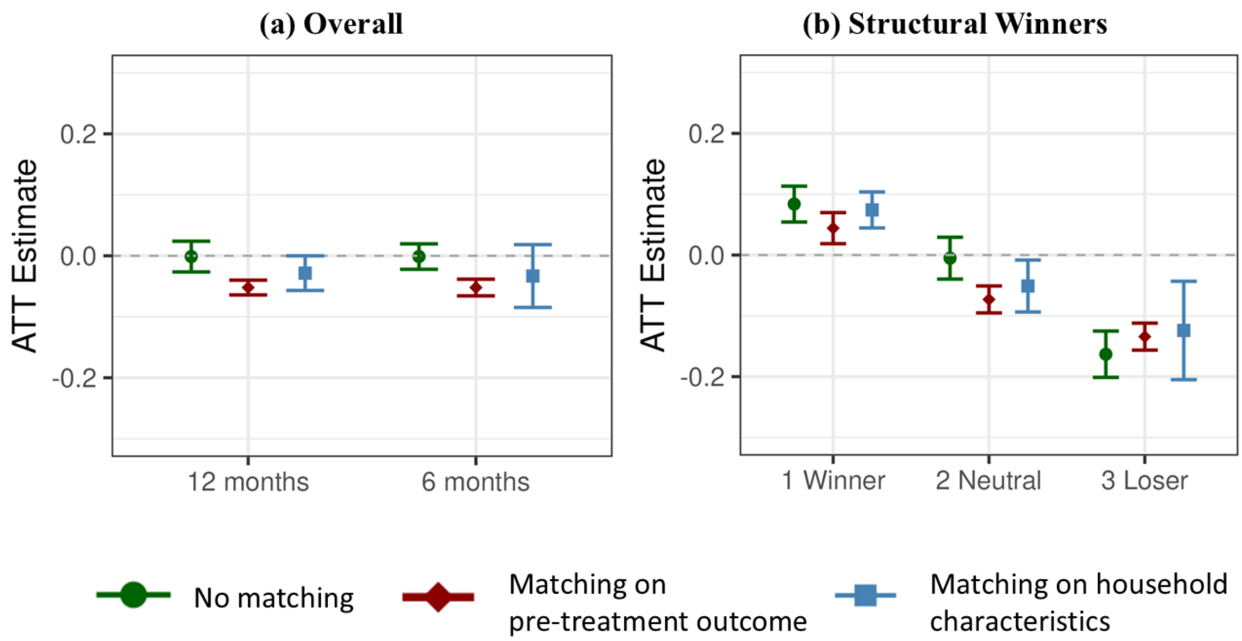

Note: The circles represent the point estimate and the bars are 95% confidence intervals generated from bootstrapped standard errors generated. Each parameter is based on a separate regression.

5.3 Treatment effects during drought

Lastly, we investigate treatment effect heterogeneity based on the timing of the 2014 drought. The drought generated substantial public pressure and demand side management that reduced water demand (California Department of Water Resources, 2021; Maddaus et al., 2020). Figure 4 shows a dramatic decrease in consumption for both volumetric and flat rate households after 2014. This decrease in consumption prior to switching to volumetric rates also makes households more likely to be winners. Figure 9 divides the sample based on 2014 drought. We split the individual group-level treatment effects to before and after January 2014.[8] There is a larger treatment effect prior to the drought, which could be due to non-price conservation crowding out the impact of

[8]The governor of California declared a state of emergency due to the drought on January 17, 2014, which makes this month reasonable starting period for the post-drought period. See the documentation at https://ca.water.usgs.gov/california-drought/california-drought-comparisons.html.

prices (Wichman et al., 2016). Alternatively, if consumers are responding to the change in total bills, the difference post drought will be smaller due to lower levels of water use from non-price conservation such as watering restrictions.

Although we see treatment effects aggregated over different treatment groups, we examine how the GATE changes across different groups. Figure A.6 plots GATE for each group with at least 0.5 percent of total records. The idea here is to understand how the treatment effects and heterogeneous groups vary over time. There are more structural winners after drought declaration since Jan 2014. Similarly, there is more weight placed on the later treated groups for Quintile 1 estimates. Alternatively, higher consumers are mostly concentrated in early treated groups. Since 2014, there have been many drought-related non-monetary water conservation programs, which may have already pushed the optimal usage. Restricting our sample to pre-drought can help with better identification of treatment effects without confounding drought-related non-monetary water conservation efforts. However, it will also leave us with a limited sample for structural winners and low consumption households.

Figure 9: Drought and change in ATT

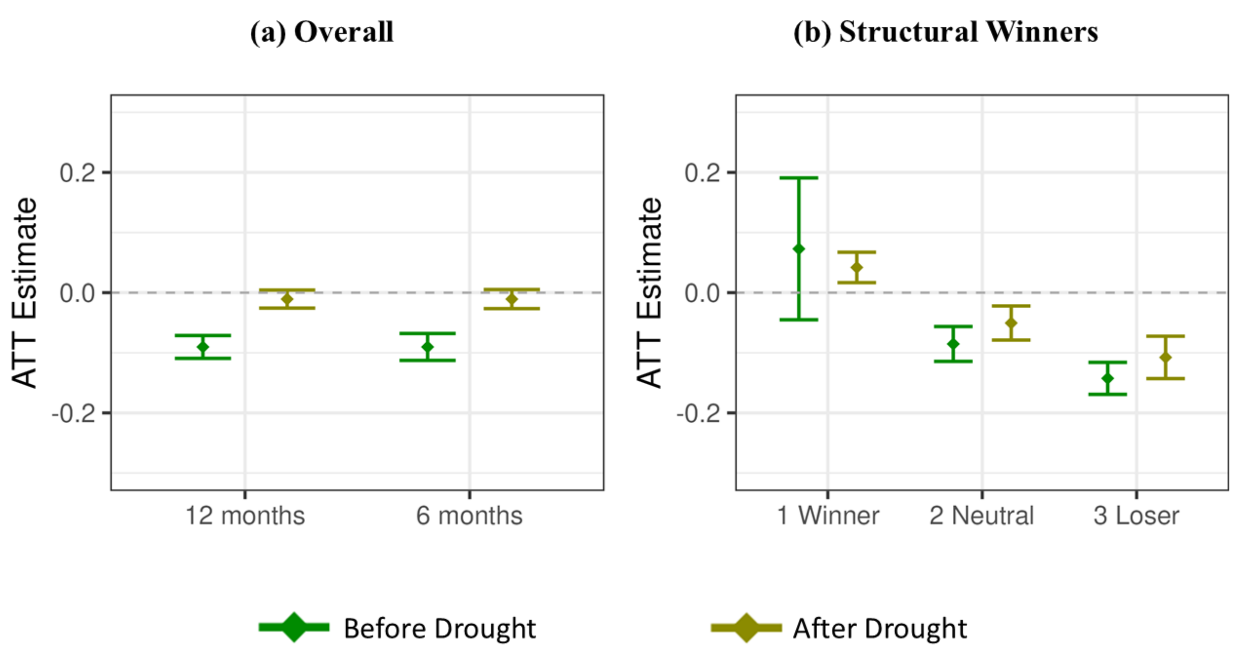

Note: In 2014, California governor declared drought and there were a lot of efforts not aligned with the pricing policy to curtail water consumption at household-level. Thus, we estimate treatment effects for pre-/post-droughts. Before droughts, the implementing volumetric pricing would yield significant reduction in overall consumption (10%) which turn negligible post drought years. There is also stronger heterogeneity in treatment effects post-drought years.

5.4 Placebo test

It is possible that the heterogeneous results are a product of mean reversion since low users in the pre-treatment period are more likely to be higher users in the post-treatment period. We believe our matching techniques alleviate this concern because we are comparing low users with other low users. We explicitly test for mean reversion using a placebo test. Here we drop all post-treatment data for two years which likely eliminates any dynamic treatment effects due to switching to volumetric rates. Then we randomly assign treatment cohorts for all households to mimic our original setup where households have 12 months of pre-treatment data. Focusing exclusively on the post-treatment period and eliminating the initial periods that likely exhibit dynamic treatment effects means that we are comparing treated households to other treated households. If our identification assumptions hold there will be a null placebo treatment effect. If our identification strategy suffers from mean revision we should see similar patterns of heterogeneity in the placebo estimates as we do in our actual estimates. The results, shown in Figure 10, show that the ATT estimates and all the heterogeneous effects are centered at zero and statistically insignificant. This placebo test alleviates concerns that the patterns of heterogeneity are primarily due to mean reversion.

Figure 10: Placebo test using post-treatment data

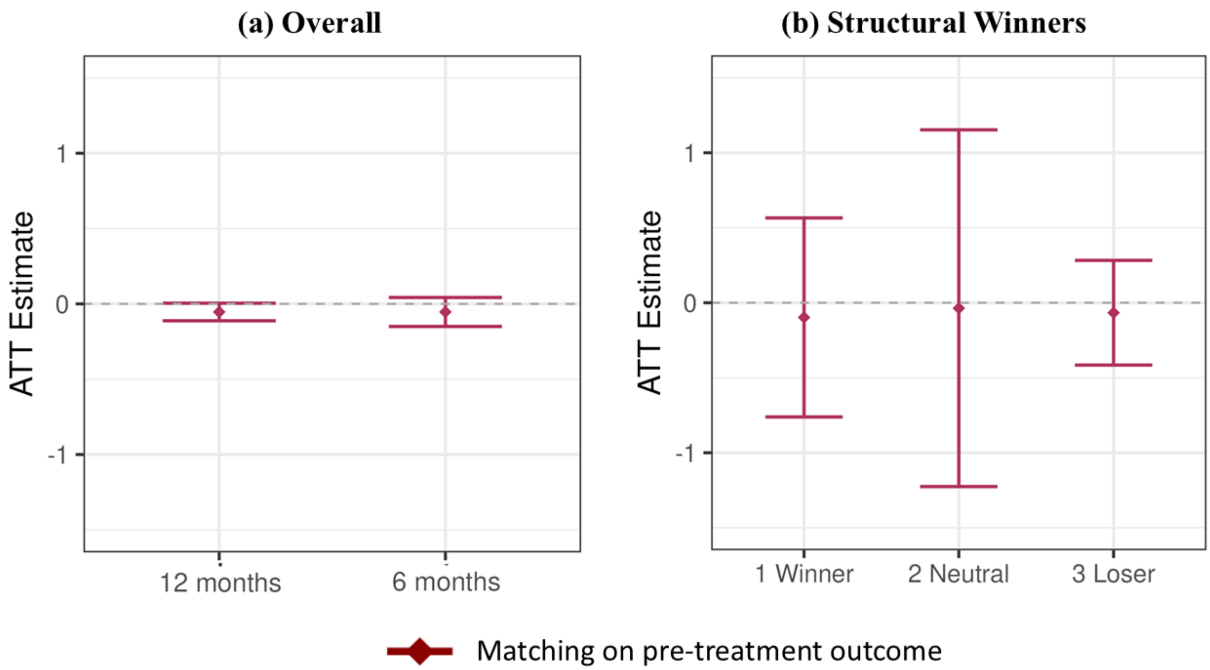

Note: These results show placebo estimates of the ATT with matching on pre-treatment water use. We restrict the sample to post-treatment data at least 2 years after transitioning to limit treatment dynamics. Comparing post-treatment households to other post-treatment households should generate null effects. We randomly assign treatment cohorts and then assign placebo tracking and volumetric periods to mimic the true data structure. The estimation in the panels mimics Figure 7 using matching on pre-treatment outcomes.

6 Demand response to total bill

The empirical estimates indicate that consumers do not respond to marginal prices and their behavioral response is more consistent with responding to average prices or total bills. Ito and Zhang (forthcoming) use a similar design for energy use in China conclude that consumers are responding to average price as suggested by Liebman and Zeckhauser (2004). They apply a non-parametric estimation method to test for schmeduling behavior. Their approach estimates treatment effects for deciles based on the change in average price. We follow a similar strategy by dividing the consumers into deciles based on two expected changes in water rates: the change in total bills (Figure 11). They show the point estimates and confidence intervals for each subgroup along with the change in total bills. The pattern, particularly for the lowest decile better fits the change in total bills. The highest decile consumers do not respond to the change in the rate structure: their consumption does not change significantly even when the total bill increases by approximately 65%.[9]

However, the relationship does not suggest a clear distinction of total bill over average price and it remains difficult to disentangle the specific price to which consumers are responding. Although the empirical evidence is not decisive, it is helpful to consider what is likely the most salient information for consumption decisions. Not responding to marginal price is consistent with models of rational inattention (Sallee, 2014). In these models, the difference in the behavioral response between costly and cheap information does not warrant the investment in acquiring costly information. In our setting, the cheapest and most salient information is the total bill. This information is written directly on the bill and is what they will see on their credit card or bank account. Calculating the average price takes significantly more effort and is not an intuitive metric for most consumers. Brent and Ward (2019) show that consumers do not have good knowledge of marginal or average prices, but do know their total bills. This is also consistent with a model of mental accounting where consumers do not pay much attention to their water bill unless it goes beyond some defined interval (Thaler, 1985).

A key element of responding to either average price[10] or total bill include the fixed price. Consumers would include the fixed price in their response accounting rather than neoclassical assumption of responding to variable cost only.

[9] We also check the demand response to average price in the Figure A.8

[10] We define average price as total bill divided by quantity as opposed to some research that uses average volumetric price interchangeably with average price (Cook and Brent, 2021)

Figure 11: Volumetric Pricing Induced Change in Water Consumption

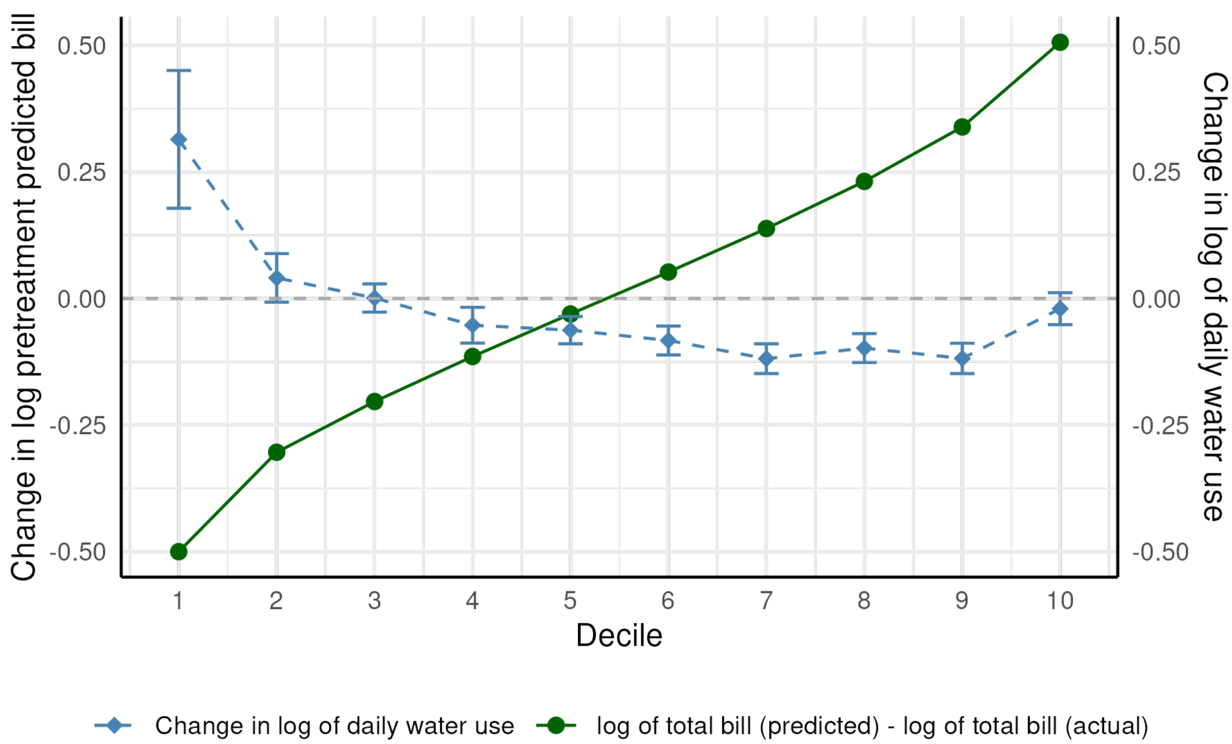

Note: This figure estimates demand based on response to the total bill. In Panel a, we divide the consumers by decile based on changes in their log transformed total bill ($log(TotalBill_{predicted}) - log(TotalBill_{actual})$). The total bill changes are indicated with a solid line. For each decile, we estimate the ATT of the volumetric pricing on log of daily water usage. We use matching on outcome-based estimator for estimating ATT for each decile. The error bars show 95 percent confidence intervals where standard errors are bootstrapped.

7 Conclusion

In this analysis, we estimate the impact of transitioning from fixed water rates that do not depend on consumption to volumetric rates in the City of Sacramento. The transition to volumetric water rates was enabled by a large-scale metering campaign, before which consumers did not know how much water they used. Moving to volumetric rates decreased water use by about 5% on average. However, the lowest users who are most likely see their total bill go down increased water use even though their marginal price increased. This is consistent with consumers using heuristics to respond to complicated rate schedules (e.g., Liebman and Zeckhauser, 2004) instead of optimizing for the correct marginal incentives. These results may be considered a lower bound because we use water in a "tracking period" as the control group. During the tracking period, consumers are shown how much they will pay once they transition to volumetric rates, so some behavioral changes may occur before consumers start paying volumetric rates.

Our results have implications for water rate design. First, changes in marginal incentives may have limited impacts on conservation without large changes in total bills, which is politically challenging. A recent trend towards budget-based rates that individualize the thresholds in an increasing block rate based on household characteristics insulates consumers from dramatic bill changes because high marginal prices typically only affect a small proportion of water use (Wietelman et al., 2024). One way to interpret that low-users increase their water use is that consumers respond to changes in their total bill, including the fixed fee. If the consumers incorporate the fixed fee in their response, there is a challenge to make utility bills more progressive without affecting the efficiency properties. Higher fixed bills will cause marginal price to deviate from marginal cost and hence less efficient recovery of bills (Levinson and Silva, 2022). If consumers respond to fixed costs, this strategy becomes more complicated since changes in the fixed cost will affect demand.

One positive element of consumers responding to fixed costs is revenue stability. Utilities often have high fixed costs but want to generate revenue with volumetric charges to encourage conservation. Our design does not allow us to estimate long-term effects. Still, if we assume the treatment effects persist, utilities can maintain revenue stability and conservation if consumers respond to total bills by collecting more revenue from fixed fees. However, if the goal of the utilities is to reduce household water demand, they should invest more time and effort in communicating bills to customers.

References

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting Event-Study Designs: Robust and Efficient Estimation," *The Review of Economic Studies*, 02 2024.

Brekke, Dan, "California Considering Fines for Wasting Water," July 9 2014. Accessed: 2024-10-15.

Brent, Daniel A. and Michael B. Ward, "Price Perceptions in Water Demand," *Journal of Environmental Economics and Management*, 2019, *98.*

Brent, Daniel and Casey J Wichman, "Do behavioral nudges interact with prevailing economic incentives? Pairing experimental and quasi-experimental evidence from water consumption," *Working Paper*, 2022.

Buck, Steven, Maximilian Auffhammer, Stephen Hamilton, and David Sunding, "Measuring welfare losses from urban water supply disruptions," *Journal of the Association of Environmental and Resource Economists*, 9 2016, *3*, 743–778.

California Department of Water Resources, "CNRA Drought Report," Technical Report, State of California 2021.

California State Water Resources Control Board, CASWRCB, "Conservation Reporting - Statewide Conservation," 2024. Accessed: 2024-09-01.

Callaway, Brantly and Pedro HC Sant'Anna, "Difference-in-differences with multiple time periods," *Journal of econometrics*, 2021, *225* (2), 200–230.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The effect of minimum wages on low-wage jobs," *The Quarterly Journal of Economics*, 2019, *134* (3), 1405–1454.

Chabé-Ferret, Sylvain, "Should We Combine Difference In Differences with Conditioning on Pre-Treatment Outcomes?," 2017.

Chaisemartin, Clément De and Xavier d'Haultfoeuille, "Two-way fixed effects estimators with heterogeneous treatment effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

Cook, Joseph and Daniel Brent, "Do households respond to the marginal or average price of piped water services?," in "Oxford Research Encyclopedia of Global Public Health" 2021.

CoreLogic Inc, "Real Estate Data, 2020," 2020. Accessed: 2024-06-30.

Dalhuisen, Jasper M., Raymond J.G.M. Florax, Henri L.F. De Groot, and Peter Nijkamp, "Price and income elasticities of residential water demand: a meta-analysis," *Land Economics*, 2003, *79* (2), 292–308.

Dube, Arindrajit, Daniele Girardi, Òscar Jordà, and Alan M. Taylor, "A local projection approach to difference-in-difference event studies," 2023.

Elinder, Mikael, Xiao Hu, Che-Yuan Liang, and Shane Minckley, "Mind the tap – how volumetric pricing affects residential hot water consumption," *Journal of Association of Environmental and Resource Economics*, 2024.

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

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

_ **and Shuang Zhang**, "Do Consumers Distinguish Fixed Cost from Variable Cost? "Schmeduling" in Two-Part Tariffs in Energy," *American Economic Journal: Economic Policy*, forthcoming.

Levinson, Arik and Emilson Silva, "The electric gini: income redistribution through energy prices," *American Economic Journal: Economic Policy*, 2022, *14* (2), 341–65.

Li, Li and Marc Jeuland, "Household water savings and response to dynamic incentives under nonlinear pricing," *Journal of Environmental Economics and Management*, 2023, *119.*

Liebman, Jeffrey B and Richard J Zeckhauser, "Schmeduling," 2004.

Maddaus, Lisa, Michelle Maddaus, Tess Kretschmann, Andrea Pacheco, Hannah Braun, Annikki Chamberlain, and Anil Bamezai, "Use and Effectiveness of Municipal Irrigation Restrictions During Drought Executive Summary," Technical Report, Alliance for Water Efficiency 1 2020.

Nataraj, Shanthi and W. Michael Hanemann, "Does marginal price matter? A regression discontinuity approach to estimating water demand," *Journal of Environmental Economics and Management*, 3 2011, *61* (2), 198–212.

Newkirk, Barrett, "California adopts mandatory water restrictions to tackle drought," July 15 2014. Accessed: 2024-10-15.

Olmstead, Sheila M., "The Economics of Managing Scarce Water Resources," *Review of Environmental Economics and Policy*, 2010, *4* (2), 179–198.

on Climate Change, IPCC Intergovernmental Panel, "Fact Sheet-Food and Water, Sixth Assessment Report, Working Group II - Impacts, Adaptation and Vulnerability," 2022. Accessed: 2024-09-05.

Rees-Jones, Alex and Dmitry Taubinsky, "Measuring "schmeduling"," *The Review of Economic Studies*, 2020, *87* (5), 2399–2438.

Sallee, James M., "Rational Inattention and Energy Efficiency," *Journal of Law and Economics*, 2014, *57* (3), 781–820.

Sun, Liyang and Sarah Abraham, "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

Thaler, Richard, "Mental Accounting and Consumer Choice," *Marketing Science*, 1985, *4* (3), 199–214.

Wichman, Casey J., "Perceived price in residential water demand: Evidence from a natural experiment," *Journal of Economic Behavior & Organization*, 2014, *107*, 308–323.

— , "Information provision and consumer behavior: A natural experiment in billing frequency," *Journal of Public Economics*, 2017, *152*, 13–33.

— **, Laura O. Taylor, and Roger H. von Haefen**, "Conservation policies: Who responds to price and who responds to prescription?," *Journal of Environmental Economics and Management*, 2016, *79*, 114 – 134.

Wietelman, Derek C, Casey J Wichman, and Daniel A Brent, "Conservation and Distributional Consequences of Pricing Scarce Water During Droughts," 2024.

Williams, A Park, Benjamin I Cook, and Jason E Smerdon, "Rapid intensification of the emerging southwestern North American megadrought in 2020–2021," *Nature Climate Change*, 2022, *12* (3), 232–234.

Wooldridge, Jeffrey M., "Simple Approaches to Nonlinear Difference-in-Differences with Panel Data," August 7 2022. Available at SSRN: https://ssrn.com/abstract=4183726 or http://dx.doi.org/10.2139/ssrn.4183726.

Figure A.1: Information Pamphlets Circulated by Sacramento Water Authority

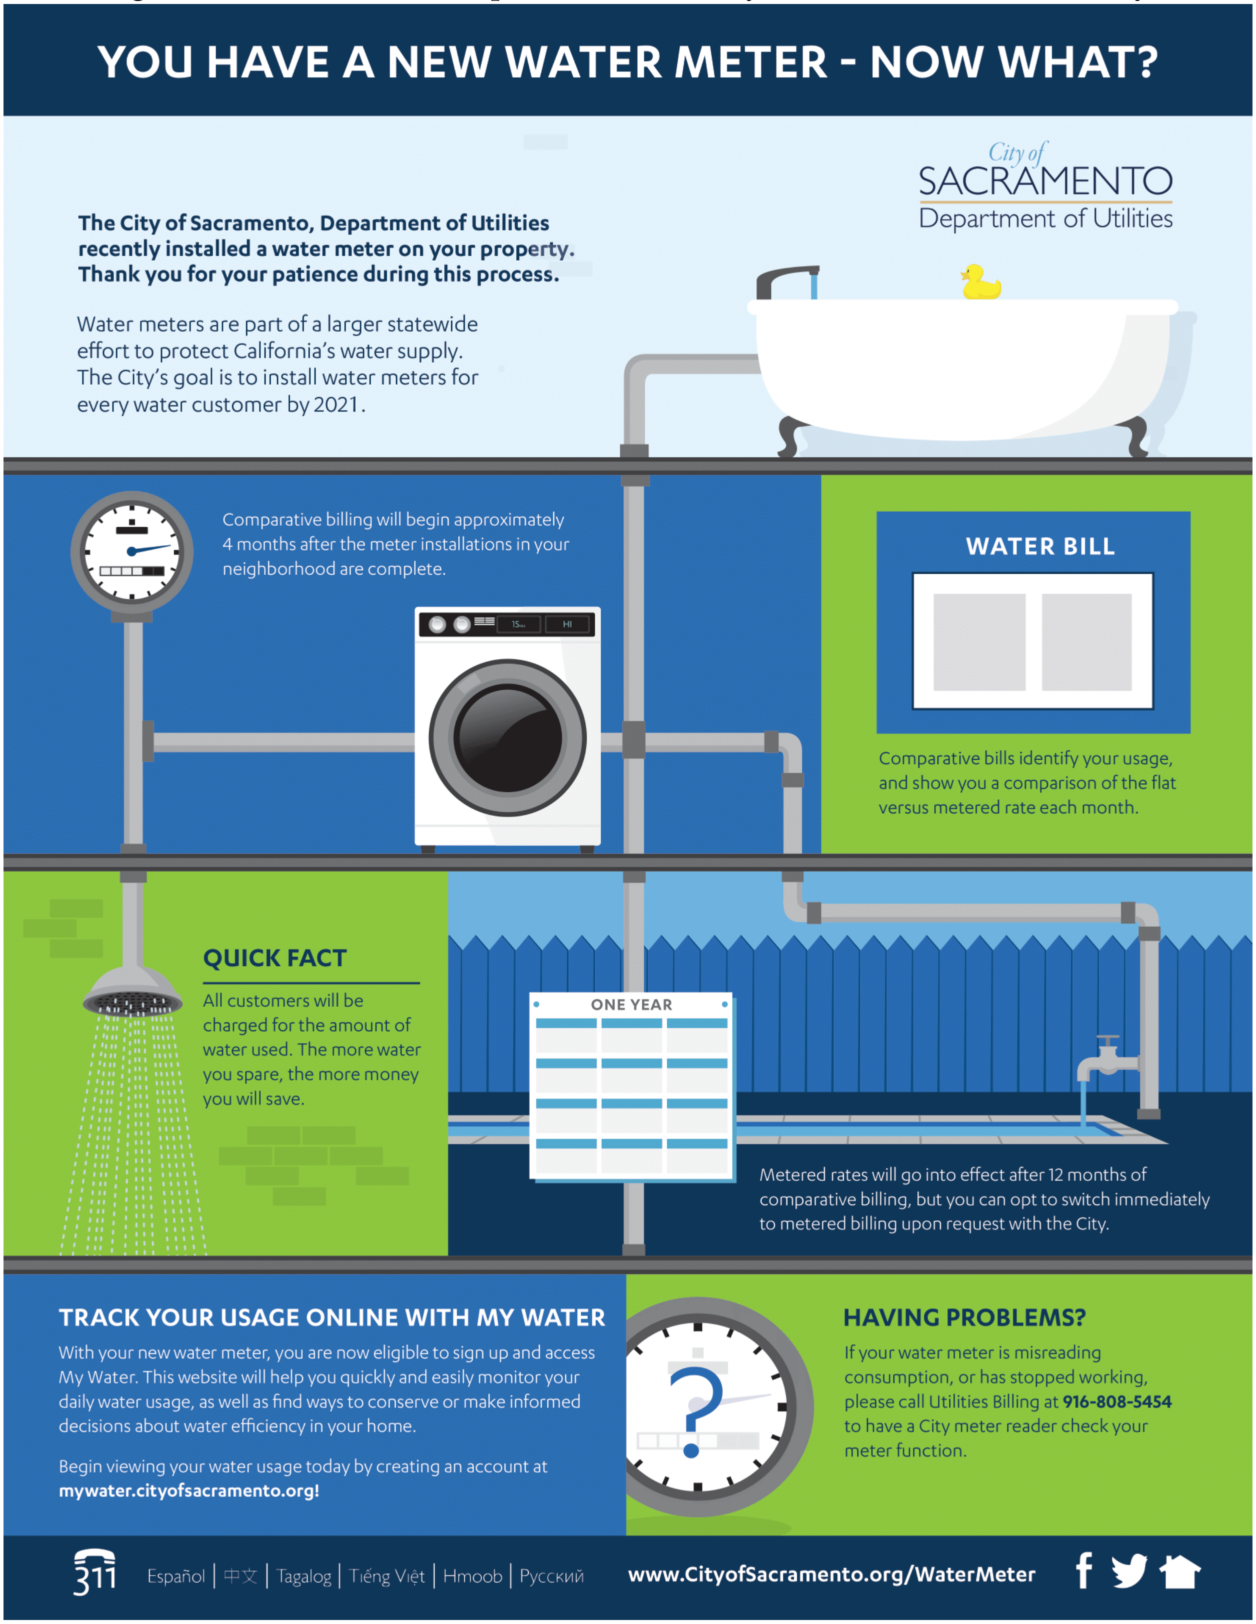

Front Side | *Continued on next page...*

HOW TO READ YOUR NEW WATER METER BILL

City of Sacramento
Utility Services Bill

Account Number:
Billing Date:

Consumption History

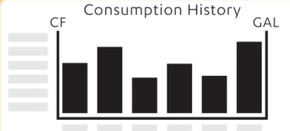

Monthly Service Charge	$ 21.34
2,259 cubic feet @ $0.0075 per cubic foot	$16.94
Subtotal	$38.29

Compared to the flat rate system, metered water bills are typically higher in the summer and lower in the winter. When reviewed over the course of the year, the majority of households on a metered rate find that metered bills are the same as – or less than – flat rate bills.

You will pay a monthly service fee on your new metered bill. The charge is based on the size of your water meter, as well as a charge for the amount of water you actually use.

	Current Meter Read		Previous Meter Read				Usage	

Keep in mind, your first metered bill will not be an accurate representation of your average monthly water usage. When your meter is installed, it is set to "0." The Automated Meter Infrastructure (AMI) system that wirelessly sends your meter reading to the City's billing center is not brought online until at least six months after installation. So, your first comparative bill will show use for multiple months, but your second will show a more accurate monthly use. For billing/usage questions, contact **916-808-5454.**

MOST WATER USE AND WASTE OCCURS OUTDOORS

We can make a big impact in the amount of water we use everyday by making simple changes. Take the following actions to see how much water you can spare each day.

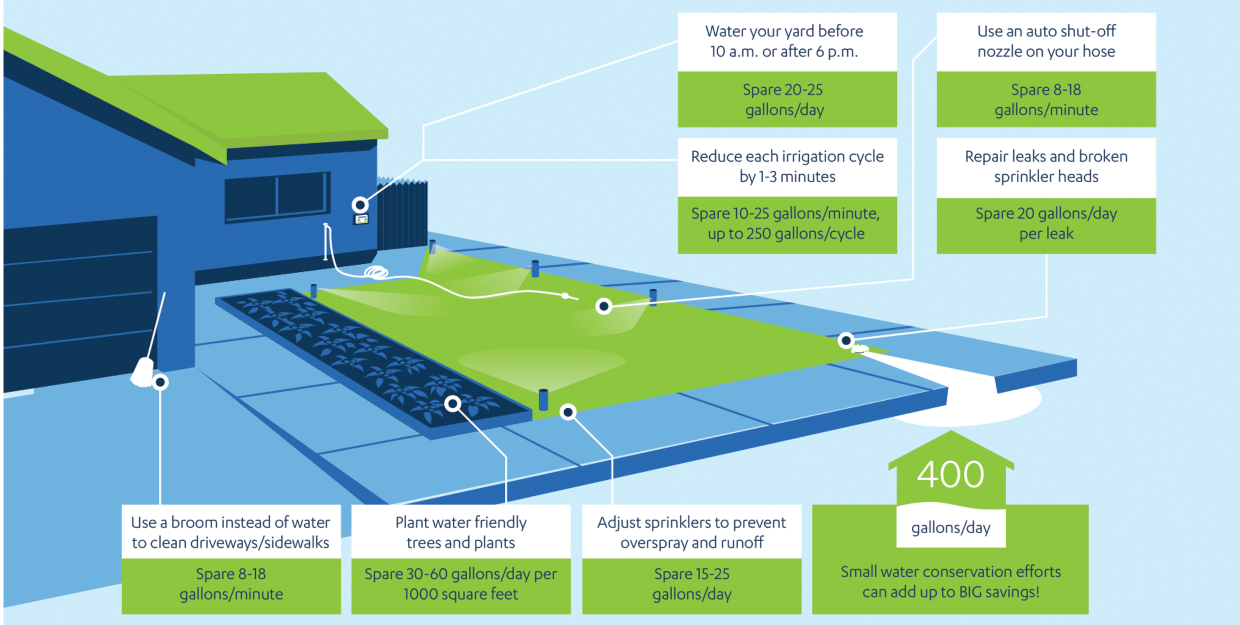

Back Side

Figure A.2: Number of days in a billing period

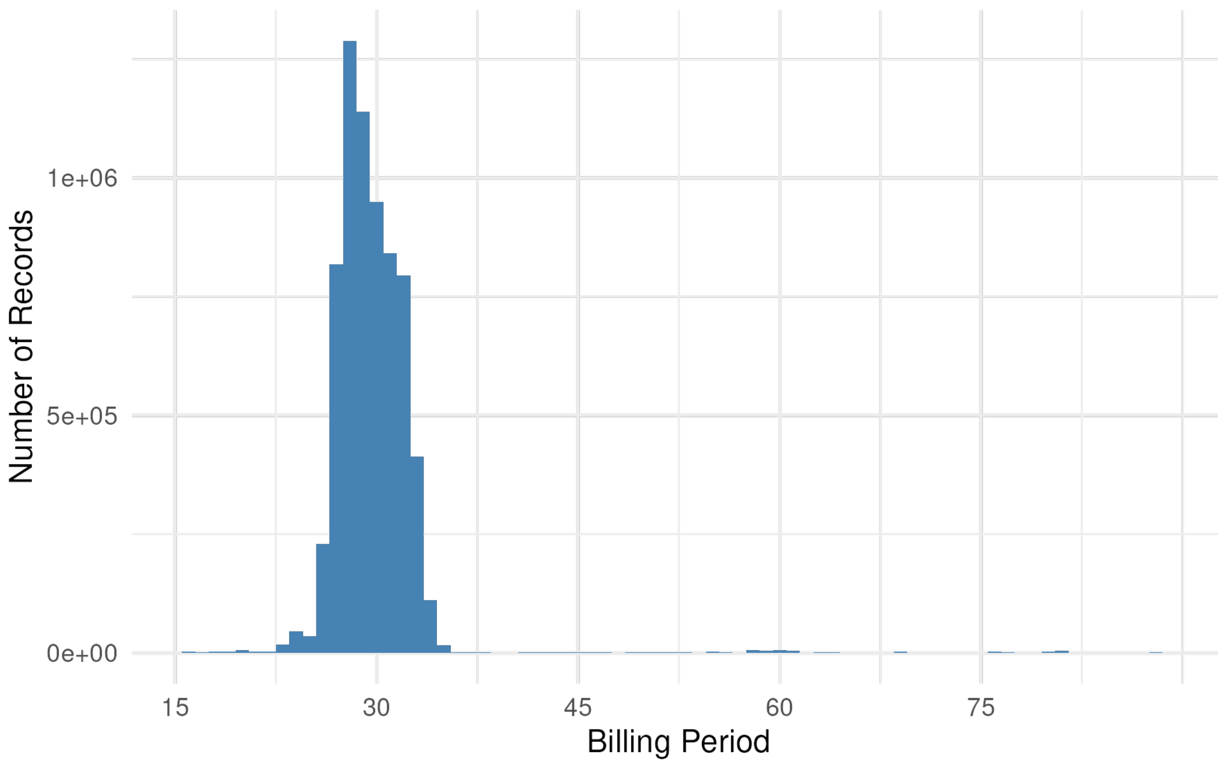

Note: Consumers are typically served monthly water bills. Thus most of the bill period lengths are 27 to 32 days. For some extraordinary circumstances a billing period is longer than 90 days or shorter than 15 days. We dropped the records with less than 16 days (0.5 % of records) and more than 89 days (0.57 % percent of records) from this figure and subsequently from our analysis analysis in the beginning itself as they do not represent a typical bill a consumer would receive.

Figure A.3: Number of Tracking Period Bills

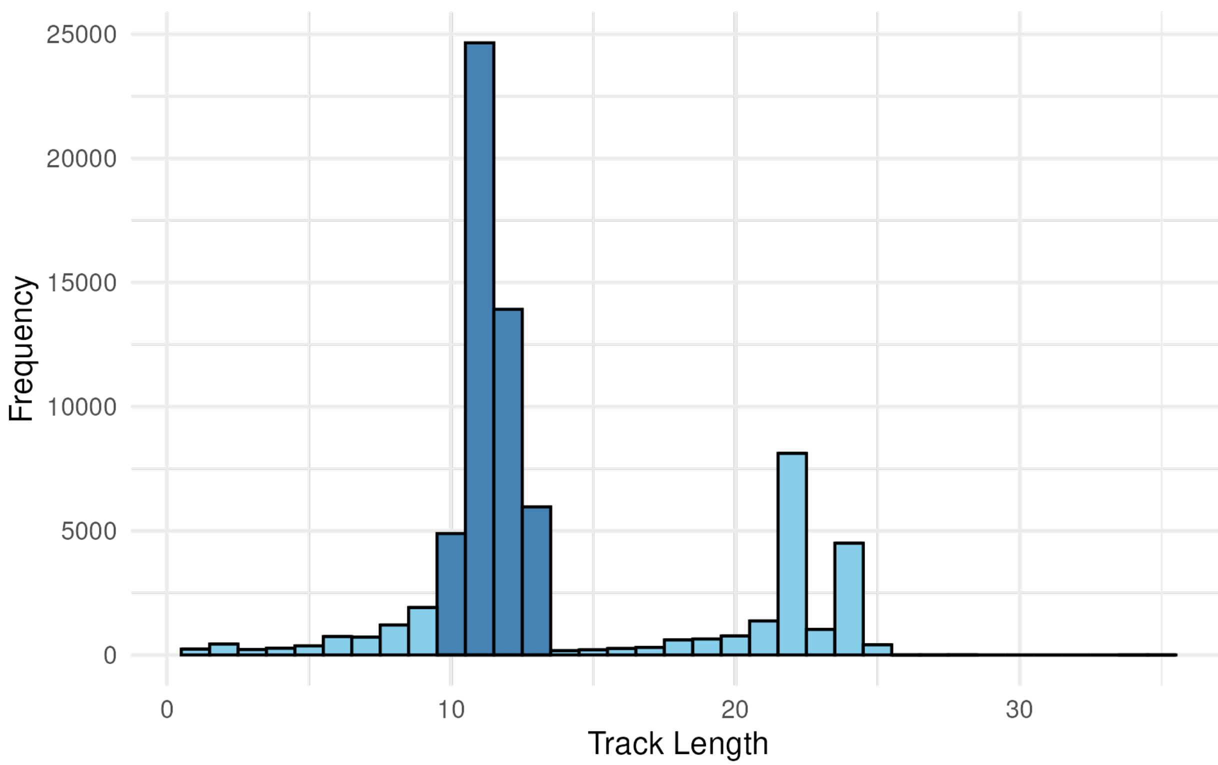

Note: Consumers are typically had 12 months of tracking period. Data suggests most of the consumers received 10-13 bills during their tracking period, highlighted with darker shade bars. There are consumers with more than 14 months of tracking, who were already metered before January 2008 but their volumetric billing did not start until January 2010.

Figure A.4: Staggered Adoption

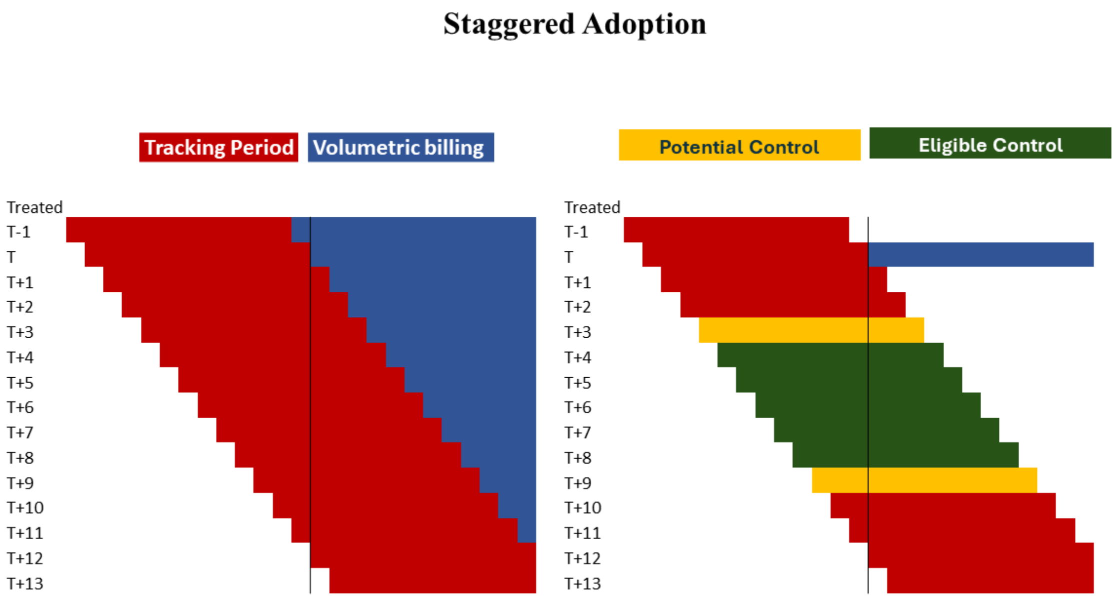

Note: The figure illustrates the nature of our set up. It is staggered adoption with imbalanced panel. We observe only 12 months pre-treatment data. So, if a cohort is treated in time T, Only groups having any pretreatment as well as posttreatment overlapping with the cohort T are cohorts T+1 to T+11. For T+12, there will be no overlap for pre-treatment observation and for T-1, there will be no post-treatment data overlap. Also, water consumption follows seasonality, water consumption changes based on the month of the year. To balance out the seasonality, we need sufficient control group spanning as many calendar months as possible. The ideal group for control shall be T+6. However, the staggered adoption does not follow uniformity across the months, there is strong possibility that there were no or very few households treated in T+6 cohort. Thus, we relax the 6 month overlap condition to 4-month overall condition. This will allow use to find eligible control groups from treatment cohorts from T+4 to T+8 (Indicated in green). 3-month overlap (indicated in yellow) might be too short, and we will not be able to find good match for the pre-treatment outcome-based matching to be used for ATT estimation.

Figure A.5: Event study plots | Bill difference heterogeneity

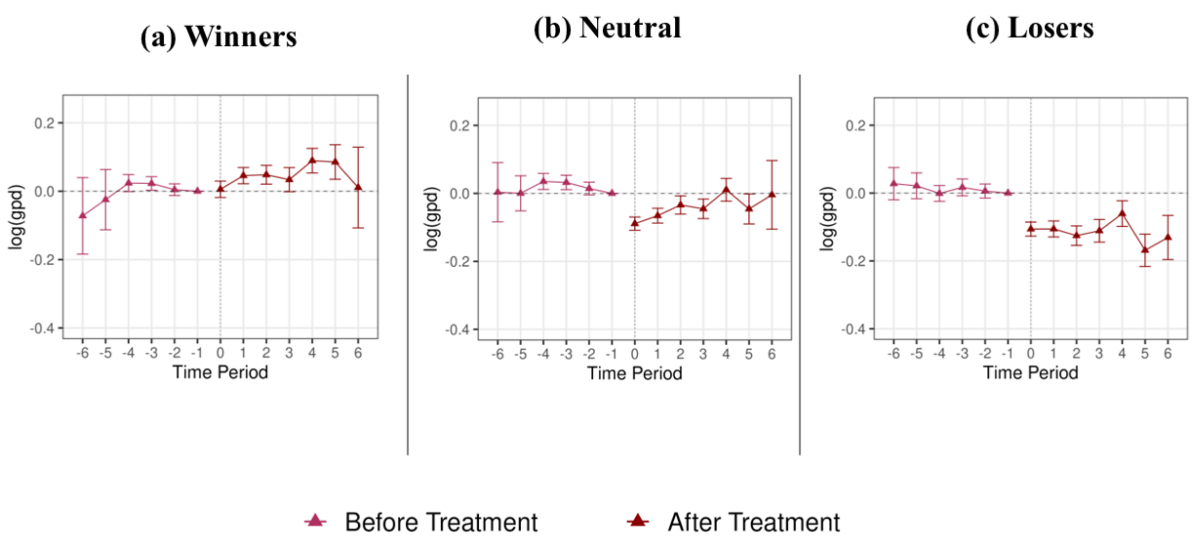

Note: The event study plots for the outcome matching-based ATT for the groups defined based on the bill difference (structural winners vs losers). The plots suggests that the parallel trend holds even beyond the matched observations (i.e. t-5 and t-6). Thus, te matching on outcome-based estimator holds critical assumptions for the difference-in-difference set up and we use it for our demand estimation.

Figure A.6: Cohort-level treatment estimates

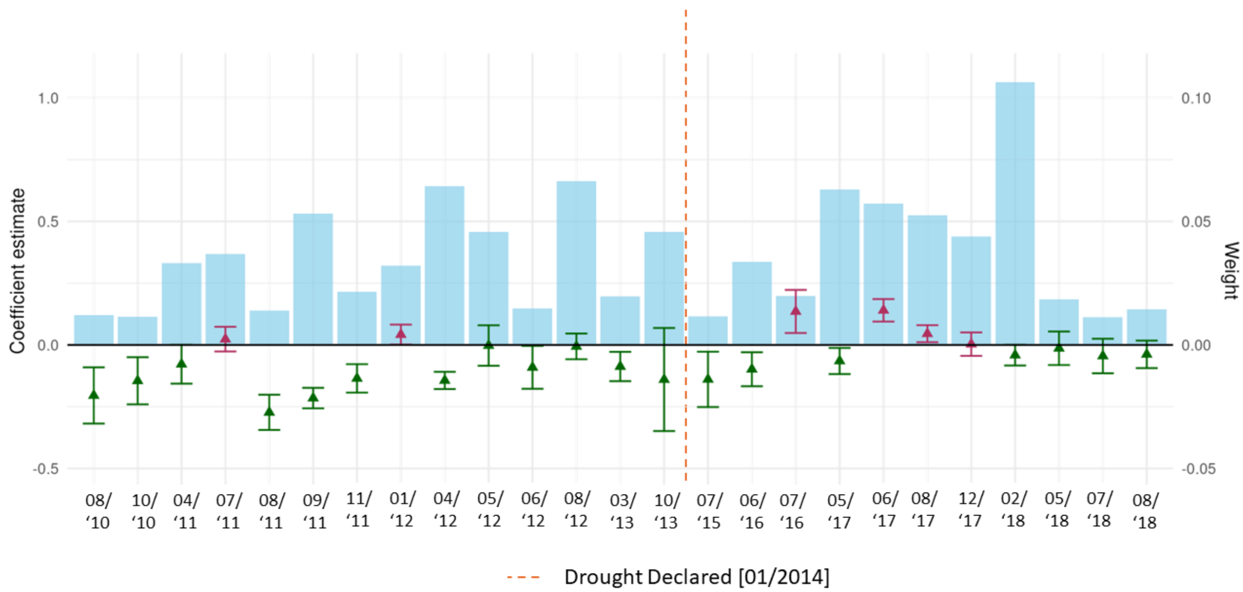

Note: This figure shows the group-wise heterogeneity in treatment effects for treatment groups having at least 1% weight in the overall ATT. Pointers and error bars in green indicated negative estimates and in maroon are for positive estimates.

Figure A.7: Cohort-level treatment estimates for winners and losers

(a) Winners

(b) Losers

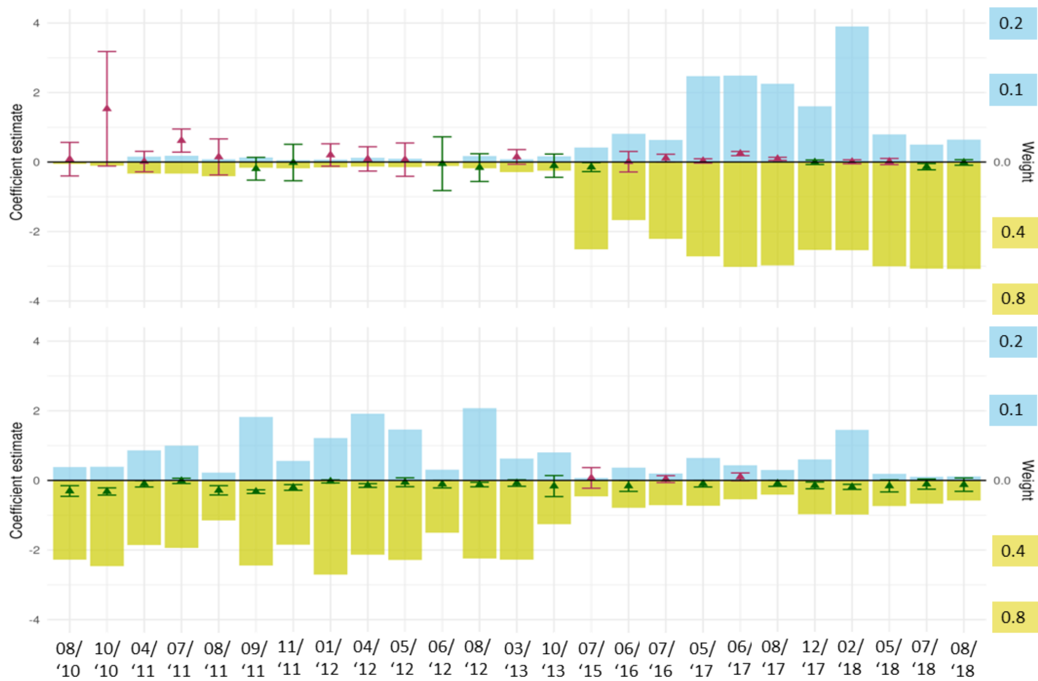

Note: This figure shows the group-wise heterogeneity in treatment effects for treatment groups having at least 1% weight in the overall ATT. Blue bards indicate relative weight of each group in the ATT (indicating the percent treated in the cohort as compared to total treated) and yellow bars indicate the category percentage within the same cohort. (For winners, percent of winner households in the respective cohort). Pointers and error bars in green indicated negative estimates and in maroon are for positive estimates.

Figure A.8: Consumers Responding to Average Price

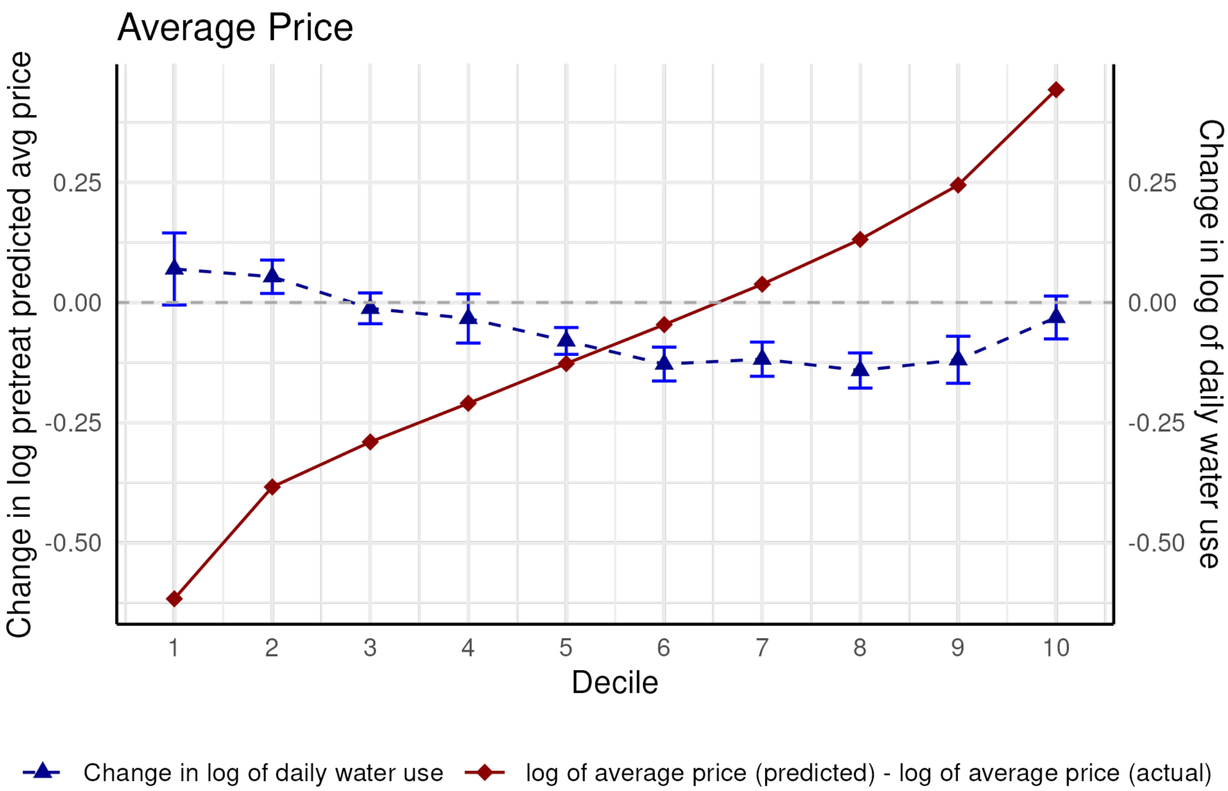

Note: Use the similar set up as Figure 11 except the deciles are defined based on the change in average price instead of total bill. We find that although having sharp drop (rise) in their average bill, Decile 1 (Decile 10) consumers are not changing their consumption. This indicates the consumers are not responding to the average price.

Figure A.9: Water Consumption vis-à-vis income groups

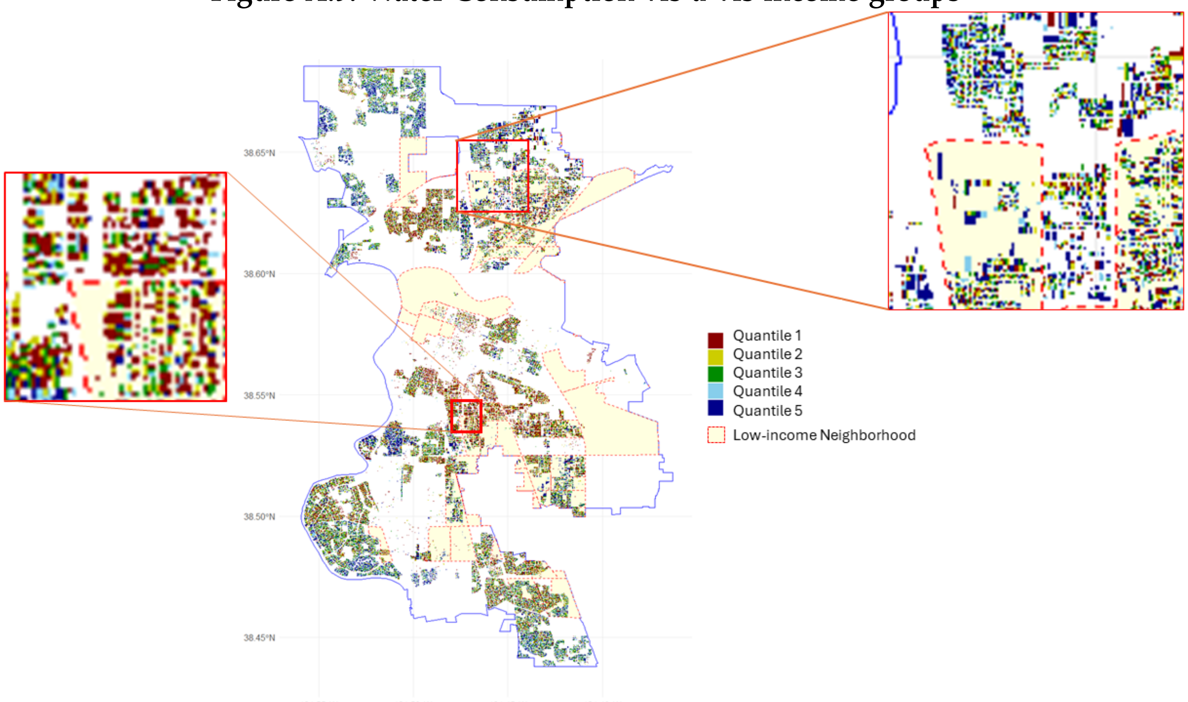

Note: Consumption quintile vis-à-vis layer of low-income (<$50,000) neighborhoods.