UNIVERSITY OF CALIFORNIA

SANTA CRUZ

**LESSONS IN DEMAND MANAGEMENT FROM THE WATER SECTOR**

A dissertation submitted in partial satisfaction
of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

ECONOMICS

by

**Bryan Edward Pratt**

June 2019

The Dissertation of Bryan Edward Pratt
is approved:

_______________________________

Professor Carlos Dobkin, Co-Chair

_______________________________

Professor Jeremy West, Co-Chair

_______________________________

Professor George Bulman

_______________________________

Lori Kletzer
Vice Provost and Dean of Graduate Studies

# Table of Contents

# List of Figures

# List of Tables

**Abstract**

Lessons in Demand Management from the Water Sector

by

Bryan Edward Pratt

This dissertation examines lessons in demand management through consumer behavior in the water sector. The first chapter identifies conservation behavior among households who were not targeted by a local drought policy and explains the mechanisms by which the policy appears to act on these households. Using billing records from two neighboring water districts, I implement a difference-in-differences approach among accounts near the arbitrary border between these districts and identify not only aggregate average treatment effects, but also effects across the distribution. I examine the marginal response to a fine-based policy instituted in only one of the two districts, controlling for contemporaneous common shocks across the two districts. While the policy was targeted only at the high end of the distribution, I establish a response even among households for which the fines were unlikely to bind. Using a simple structural estimation, I rule out a variety of price-based arguments and argue that the most consistent theory for the observed behavior is that the policy announcement temporarily shifted demand, reducing the marginal utility of consumption at all levels. Using both a randomized experiment and an arbitrary threshold in a concurrent natural policy experiment, the second chapter examine the effectiveness of a social comparison intervention across two states of the world with higher and lower pecuniary incentives. Contrary to some of the existing literature, we do not observe evidence of pecuniary incentives crowding out social pressure in this context, despite the fact that both policies were effective. The third chapter focuses on the external validity of drought policies and the population characteristics which might influence their effectiveness. I investigate whether renters are differentially sensitive to price-based and normative instruments. In answering this question, I find that responsiveness is primarily driven by the account characteristics that determine how likely a policy's effective

mechanism is to bind. In other words, the main source of heterogeneity in each policy's effectiveness appears to be driven by exposure to the policies themselves. Because this exposure is correlated with my proxy for renting, there are differential responses among homeowners and renters, but these are not due to landlord-tenant relationship.

To my family.

## Acknowledgements

Foremost, I would like to thank Tara Sinclair for encouraging me on my path to pursue a doctoral degree in economics and Brent Haddad for introducing me to the world of water utilities. Without them, none of this work would have happened. I wish to thank WaterSmart Software for implementing the randomized Home Water Reports and William Holleran in particular for assistance with the associated data and also to thank the California Data Collaborative and Patrick Atwater for data from additional utilities. I also sincerely appreciate the cooperation of the utilities directly providing data and materials for this work, without whom this research would not have been possible. In addition, I wish to gratefully acknowledge the financial and non-financial support of the Property and Environment Research Center. Furthermore, I am grateful for insightful feedback from Carlos Dobkin, Jeremy West, George Bulman, Dylan Brewer, John Creamer, Dahyeon Jeong, Justin Marion, Jon Robinson, Alan Spearot, Joshua Abbott, Hunt Allcott, Kelly Bishop, Daniel Brent, Lopamudra Chakraborti, Jesse Cunha, Lucas Davis, Daniel Friedman, Nicolai Kuminoff, Natalia Lazzati, Jonathan Meer, Adam Millard-Ball, Steven Puller, Jason Scorse, John Sorrentino, Forrest Williams, Kerry Smith, Ian Lange, David Simpson, Randy Rucker, and Wally Thurman. I also thank attendees at the UCSC Microeconomics Workshop, the AERE Water Session at SEA 2017, the 2018 meetings of the UEA at Columbia University, and numerous seminar participants at the Property and Environment Research Center. The panel regressions in this paper make use of the `reghdfe` package, created by Sergio Correia. Any errors, opinions and conclusions in this study are my own and do not necessarily represent the position of WaterSmart Software, the California Data Collaborative, the Property and Environment Research Center, nor any other parties.

# Chapter 1

# A Fine Is More Than a Price: Evidence from Drought Restrictions

An extensive literature examines financial penalties through the lens of fine avoidance. Becker [1993] describes the economic approach to illegal activity as examining the behavior of individuals by understanding the incentives that make misbehavior more or less profitable. Deterrence theory posits that fines cause potential violators to comply with policies to avoid the financial penalty and any related consequences, and this theory has been tested and reformulated in dozens of research applications [For a recent review, see: Pratt et al., 2009]. The mechanisms by which financial penalties or incentives cause fine avoidance behavior include price-based mechanisms, but they also include stigma, impacts on intrinsic motivation, and other non-price channels [Gneezy et al., 2011, Faure and Escresa, 2011]. Recent related work by Homonoff [2018] examines the extent to which even very small fines can cause large responses through behavioral channels.

In this paper, I consider a separate channel through which financial penalties affect behavior, which I term a 'severity signal.' This severity signal conveys to the entire population, not just likely violators, the importance of the problem or externality and causes an adjustment in behavior that is additional to any observed fine avoidance. This concept is most similar to the findings from Gneezy and Rustichini [2000] that the imposition of a financial penalty can cause individuals to update their beliefs about the importance of contributing to a public good. However, to the best of my knowledge, I am the first to examine the effect of fines on contributions to a public good among individuals who were subject to a policy but unlikely to face fines themselves.

To provide insight into how a fine policy affects individuals separately from fine avoidance behavior, I need to observe responses among individuals who face the information signal without the financial and other consequences of the fine itself. I provide this context by examining a quota-based fine policy intended to generate water conservation during exceptional drought conditions. The structure of the financial penalty, which was imposed on households consuming beyond a uniformly specified allotment, allows for examining the consequences of the policy on households who faced the 'severity signal' of the fine policy without the direct consequences of the fines themselves. I use a difference-in-differences methodology with quasi-random assignment to treatment across an arbitrary border to identify the policy response. In addition, I interact treatment assignment with measures of pre-period consumption, which are highly correlated with and predictive of a household's exposure to the direct component of the fine. This allows me to see the average policy response among households with a low risk of being fined and among households throughout the distribution of fine risk.

The reduced-form results from this exercise not only demonstrate that the policy was effective but also highlight the relative similarity of the response among households with very different probabilities of facing fines. Using the water bill prior to treatment, households at the 25th percentile and 75th percentile of pre-period consumption responded by 10.7 percent and 13.3 percent in the treated district, respectively, relative to the counterfactually. This is despite the fact that households at the 25th percentile faced only a six percent chance of being fined, while households at the 75th percentile faced a 30.1 percent chance of being fined.

This relatively minimal difference in policy responsiveness across different levels of exposure is consistent throughout the vast majority of the consumption distribution. In addition, the policy response is quantitatively similar for households that had never consumed more than 40 percent of their allotment and households that had previously consumed above the allotment at least once in the pre-period. The consequence is that a fine equivalent to increasing the marginal price for high consumption by 25 to 50 dollars per hundred cubic feet of water induced a larger aggregate response than would be predicted from a 100 percent increase in prices for all households, even for the most conservative estimation of the policy's effect.[1]

Given these findings, it is important to consider what factors might explain these behavioral changes. Potential explanations could include risk aversion, over-weighting of low-probability events, uncertainty, and incorrect beliefs about own consumption. The structure of billing for water consumption, primarily monthly in this setting, would suggest that households are not perfectly informed about their consumption. However, the institutional setting includes extensive effort by the water district to provide all necessary information, and the results are not consistent with imperfect information as the primary mechanism. Among other efforts, the water agency includes 24 months of historical consumption in the form of a vertical bar graph on each affected bill, with a horizontal red line indicating the threshold for penalties. Under the assumption that a lack of information or biased beliefs are acting in combination with fine avoidance, the predicted response would increase in magnitude with the size of expected fines much more strongly than is present in the data.

In addition, households may also respond to the probability of receiving a fine through a risk aversion channel. With a high degree of risk aversion, it would be possible to rationalize strong responses even among low-risk households. However, assuming such a high degree of risk aversion would lead again to much larger pre-

[1]The median total bill for a household in the month before the policy is $41.01, including fixed charges. The mean total bill for that month is $57.07. The median and mean amount of fixed charges for the sample in the same month are $22.20 and $22.85, respectively. The median average price of 100 cubic feet of water in that month is $8.97, while the mean average price is $10.23.

dicted responses among households higher in the distribution. Because the financial penalty grows linearly in consumption, there is no reason to expect a plateau in response magnitude if individuals are only responding to the financial consequences. Moreover, observed policy responses do not vary extensively by accounts' maximum historical consumption, neither unconditionally nor conditional on consumption in the bill before implementation, which is in direction opposition to the predictions of high risk aversion.

One hypothesis that is potentially more consistent with the results is that individuals over-weight low-probability events and under-weight high-probability events, as argued by Tversky and Kahneman [1992]. Under this prospect theory, the distribution of local average treatment effects would include a stronger than otherwise predicted response among low users and a weaker response among high users, which is qualitatively similar to my results. However, the findings presented here would imply a highly skewed weighting of event probabilities that is quantitatively inconsistent with their work. In order to match the response curve in this research to a prospect theory formulation, individuals would need to have weights that are substantially more skewed than in Tversky and Kahneman [1992]. In a related vein, Homonoff [2018] attributes a large response to a relatively small tax on disposable plastic bags to loss aversion, noting that a subsidy of similar size was substantially less effective. In my setting, however, an explanation of loss aversion would again predict a larger response among individuals with higher exposure, contrary to the results I find. Broadly, the evidence most strongly supports the explanation that the policy announcement directly reduced water consumption through the non-price channel of signalling the importance of conservation.

The implications of these results are important for optimal policy design, especially for policies targeting resource conservation and other contributions to public goods. The literature on optimal fines does not consider the potential aggregate effects that a policy can generate through a severity signal [See, for example: Becker, 1993, Polinsky and Shavell, 1979]. As a consequence, if that signal either increases or decreases the contributions to a public good, it will cause the traditional calculation of optimal fine magnitudes to be overstated or understated, respectively. Furthermore, policymakers choosing between imposing fines or employing a different policy mechanism should consider not only the fine avoidance behavior, but also the aggregate effect of fines on the population. This may cause a change in optimal policy choice, especially when the policymaker faces institutional constraints.

In addition, this paper contributes to the extensive literature on utility policy and demand management. Several authors have demonstrated that non-price channels are important for conservation policy in both the electricity [Allcott, 2011, Allcott and Rogers, 2014, Asensio and Delmas, 2015] and the water sectors [Brent et al., 2015, Ferraro and Price, 2013]. In addition, several studies have examined outdoor watering restrictions, a form of partial rationing [Grafton and Ward, 2008, Hensher et al., 2006, Fielding et al., 2012]. However, this is the first study to examine similar restrictions placed on total consumption by households. To the best of my knowledge, no study has yet examined restrictions on the total combined indoor and outdoor consumption by households, which is a relatively rare tool but increasingly under consideration in the most drought-prone areas of the world. From a policy perspective, this work expands on a growing literature documenting what conservation policies work for residential water consumption and why these policies are successful [See, for example: Bruvold, 1979, Wichman, 2017, Datta et al., 2015, Fielding et al., 2012]. This information is increasingly important for water managers, as well as policymakers in other areas of conservation, who are considering which of these policies to implement.

## 1.1 Data and Setting

The context considered in this research is the implementation of a fine-based policy for high water consumption during a severe drought in California. The study region, as with the vast majority of California, receives the nearly all of its annual rainfall during the months of October through May. After receiving just over 13 inches of rain - less than 50 percent of its historical average - from October 2013 through May 2014, the region was facing an 'exceptional' hydrological drought, according to the Palmer Drought Severity Index (PDSI).[2] Figure 1.1 plots monthly historical values for the PDSI for the Central Coast region of California, which includes Santa Cruz County, while Figure 1.2 illustrates the evolution of drought intensity across the state during the study period.[3]

Figure 1.1: Historical Drought Severity, California - Central Coast

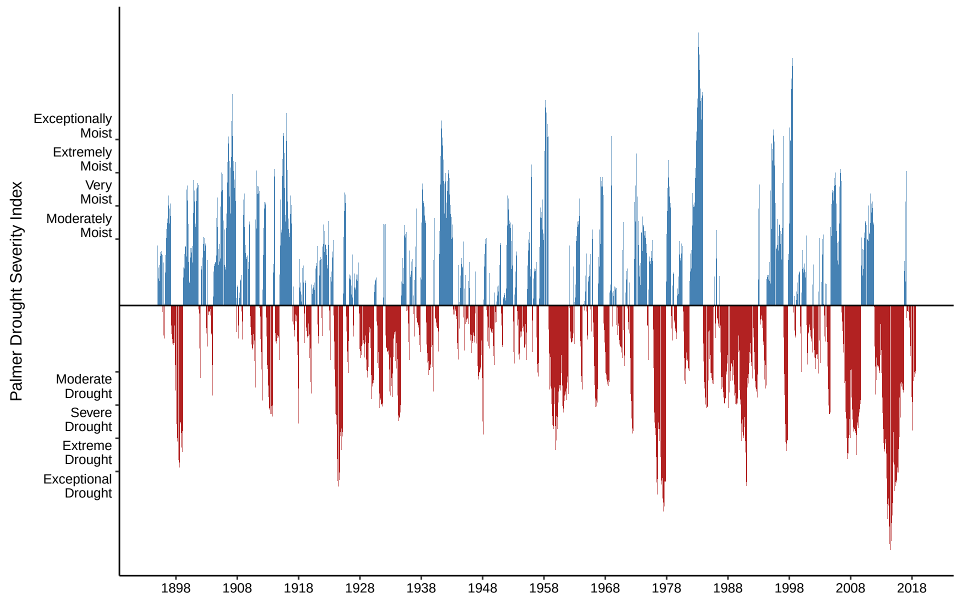


Source: NOAA

Despite the fact that the drought was formally considered exceptional, an 'exceptional' level of hydrological stress has actually occurred in this area during three separate periods in the past four decades.[4] Given the increasing prevalence of intense drought in this region and elsewhere, the findings from this research could have very immediate policy implications for water management. Notably, the policy studied here was ultimately covered by the New York Times and Santa Monica specifically cited Santa Cruz's example when imposing a similar policy in 2015, suggesting this policy and other similar policies are likely to be implemented in

[2]The PDSI is the standard measure of hydrological conditions used by the United States government, based on seminal work by Palmer [1965].

[3]Data on drought severity are available from the National Oceanic and Atmospheric Administration (NOAA). NOAA also provides an explanation of the drought indices. See [NOAA, 2015b].

[4]The drought of 1987-1992 is worth noting, as it is the last time drought restrictions were enforced in this area. Appendix A.1.6 compares the estimates for the most recent drought restrictions with an estimate for these previous restrictions and finds similar effects.

the future.[5]

Figure 1.2: California Drought Conditions, April 2013 - April 2016

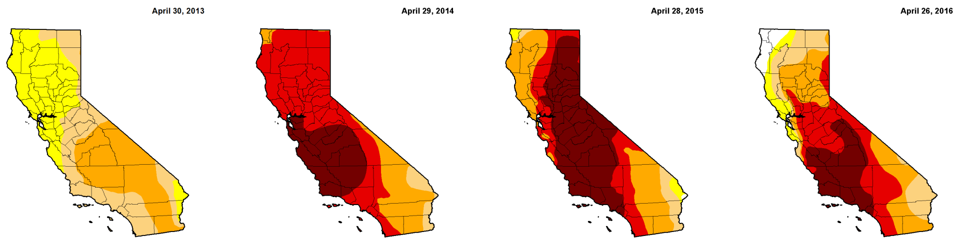


| Color | Yellow | Tan | Orange | Light Red | Dark Red |
|---|---|---|---|---|---|
| PDSI | -1 to -1.99 | -2 to -2.99 | -3 to -3.99 | -4 to -4.99 | -5 or lower |
| Classification | Abnormally Dry | Moderate Drought | Severe Drought | Extreme Drought | Exceptional Drought |

The U.S. Drought Monitor is jointly produced by the National Drought Mitigation Center at the University of Nebraska-Lincoln, the United States Department of Agriculture, and the National Oceanic and Atmospheric Administration. Maps courtesy of NDMC.

### 1.1.1 Policy Setting

In response to this exceptional hydrological stress, the two water districts in the greater Santa Cruz - Capitola metropolitan area chose two distinct responses to the drought, with these actions beginning in 2014, well before the onset of separate state mandates.[6] One chose a policy that imposed consumption quotas and fines for noncompliance, while the other raised marginal rates under an emergency rate schedule. The rate systems for both Santa Cruz Municipal Utilities ("West") and Soquel Creek Water District ("East") customers are shown in Figure 1.3.

Taking a quantity-oriented approach, the West district imposed volumetric restrictions in 2014 and 2015, but these restrictions were removed from November of 2014 through April of 2015 [SCMU, 2014, 2015a]. The stated policy was that these restrictions would be removed indefinitely, and their reinstatement was only an-

[5] [1] "In California, Spigots Start Draining Pockets." New York Times. https://www.nytimes.com/2014/05/09/us/a-thirsty-california-puts-a-premium-on-excess-water-use.html
[2] "City Council Report; City Council Meeting: January 13, 2015; Agenda Item: 8-B." City of Santa Monica. https://www.smgov.net/departments/council/agendas/2015/20150113/s2015011308-B.htm

[6] In 2015, the governor announced that the state would impose financial penalties on water districts if they failed to meet state-imposed conservation goals [Kostyrko, 2015].

Figure 1.3: Marginal Rate Comparison Over Time, 2014 Policy Period

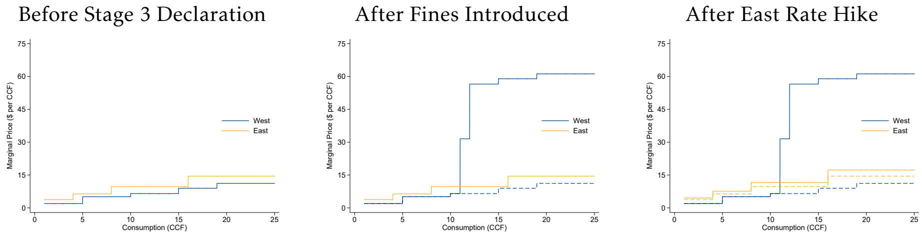


nounced in April of 2015. The restrictions were lifted as of November 1, 2015, with the shift announced on October 28, 2015 [Gomez, 2015]. The initial implementation covered bills starting in May of 2014. The policy was announced shortly before it was enforced, and households were notified through mail. Furthermore, the local press covered the imposition and removal of the policy.

While framed as rationing to the consumer, the penalty for using water beyond the allocation is a fine. The restrictions are as follows:

- Consumers are given **monthly water allotments**, which were the same in both the 2014 and 2015 rounds of restrictions.
- Single family residential accounts are assigned a monthly water allotment of 1,000 cubic feet of water (10 CCF)[7]
- Every CCF beyond the allotment is charged a fine in addition to the tiered pricing structure. These are termed **excessive water use penalties**:
  - For the first 10 percent beyond the allotment (one CCF for a single family residential account), the fine is $25 per CCF
  - For every CCF beyond the first 10 percent, the fine is $50 per CCF
- Instead of paying the fine, a first-time violator may elect to attend water school, a one-time class where violators learn about conserving water and why the restrictions are in place. Notably, this avoidance can only be done once.

[7] CCF is the standard billing unit of water, and it represents a "centum cubic foot" of water, or 100 cubic feet of water. One CCF is equivalent to 748 gallons of water.

It is important to note that the imposed restriction for the single-family household is 1,000 cubic feet of water (10 CCF), which is near the threshold between tiers 2 and 3 of the five-tiered system in the West service area. However, the fine-based marginal price increase is more than 20 times as large as the jump in marginal prices between tiers 2 and 3. Indeed, the fines can change the total cost of a monthly water bill by an order of magnitude for violators. Another important consideration is the fact that violators can choose to attend water school instead of paying a fine, if it is the first violation. According to Gomez [2015], "[a]lmost 400 people opted to attend water school rather than pay fines." While this may soften the financial consequences of the policy, the treatment is a clear punitive measure incurred by crossing an arbitrary threshold that does not coincide with a marginal price change in the tiered pricing structure. For perspective, 4,107 accounts with a single-family allotment received at least one fine, which is 13.7 percent of all single-family households in the district. For these 4,107 accounts, the median fine assessed is $75, and the median total value of fines billed during 2014 and 2015 is $175. An example bill, with a penalty, is provided in the Appendix, in Figure A.2. It is worth noting that bills generated during the policy period include a vertical bar chart with the last 24 months of consumption and a horizontal red line with the policy-imposed allotment. As a result, households are given all of the information necessary to understand the policy.

As shown in Figure 1.3, the East district raised rates during the West policy period, but this rate increase was modest by comparison. A quota system was proposed, but it was not enacted for idiosyncratic reasons related to the timing of planning cycles years before the study period.[8] Instead, the Stage 3 emergency declaration led to the district increasing rates by approximately 19 percent starting with billing periods beginning on or after July 1, 2014, through the end of the calendar year.

[8] [1] "Soquel Creek customers face rate increases and water rationing." KION. http://www.kionrightnow.com/news/local-news/soquel-creek-customers-face-rate-increases-and-water-rationing/26563910

[2] "Soquel Creek Water District water emergencies declared." Santa Cruz Sentinel. http://www.santacruzsentinel.com/general-news/20140617/soquel-creek-water-district-water-emergencies-declared

[3] "Water Rationing, Rate Increases Coming Soon." Aptos Community News. http://aptoscommunitynews.org/news/2014/03/19/water-rationing-rate-increases-coming-soon/

Both districts implemented small rate increases near the end of the year, which had been approved before the onset of the drought. In 2015, the East district again raised rates to address the worsening drought, with an increase of approximately 15 percent for billing periods beginning on or after June 1, 2015. These price changes most likely generated some reduction in consumption, but it is worth noting that the literature on price elasticities among residential water consumers has found estimates ranging from -0.1 to -0.2 [Ito, 2013]. My estimates of the policy's impact are therefore similar to the predicted consequences of a 100 percent increase in prices. However, it is worth noting that the cumulative increase in prices in the East district may be sufficient to seriously attenuate any estimates for the effect of the fine-based policy in 2015.[9] As a result, I will focus my attention on estimates for 2014 and present estimates for 2015 for completeness. The focus of this paper is on the effect of the financial penalties, using the arbitrary boundary between these districts to assess the mechanisms by which it altered consumption.

While one might be concerned that the choice of policy was endogenous to the severity of the drought or other characteristics of each district, publicly available information and private communications suggest that the policy divergence was largely due to idiosyncratic factors. Furthermore, while pricing changes were endogenously timed in the East district to reduce demand as conditions worsened, the financial penalty structure was implemented on a timeline independent of other pricing changes in the West district. All other rate structure changes during the study period in the West district were planned and voted into force before the beginning of the drought. Both districts acted independently, despite coordination on other issues of importance.[10] Furthermore, the policy was timed not around rate structures but around the seasonality of water consumption and rainfall. At the end of the 'wet' season of the water year, officials in the West district voted the financial penalty system into force to be implemented as soon as possible. While not designed to be temporary, the policy was implemented at the beginning of a

[9] Figure A.10 suggests that total expenditure stayed at a similar level and retained its historical seasonality in the control district while rates there rose.

[10] The two districts have open communication but have different legal structures for historical reasons and operate fully independently.

Figure 1.4: Water District Service Areas

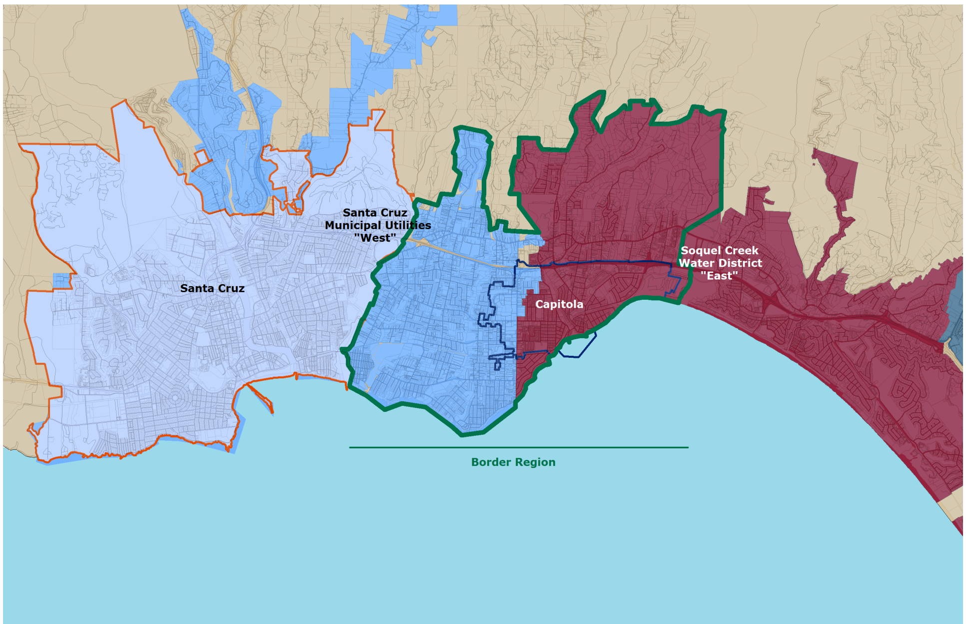


seasonal rise in consumption and removed with the seasonal fall in consumption in both years that it was in force. Moreover, while it is reasonable to posit that the divergence in policy was the result of divergent circumstances, the two districts share many things in common. Among other things, the water districts share their main aquifer, with each district also utilizing additional sources.

### 1.1.2 Geographic Setting and Identification

In order to identify the effects of this policy on households in the West district, I leverage the fact that these two districts share an arbitrary boundary. Figure 1.4 provides both an overview of the service area boundaries of these water districts and a visual representation of the border design I employ. Notably, the border runs through both commercial and residential areas and is independent of political boundaries. Both service areas have tiered pricing, and the primary difference between them is that the East district implemented higher marginal rates instead of a fine-based system.

The border design in this paper resembles Card and Krueger [1994] and Black [1999], using the assumption that households on either side of the water service

Table 1.1: Balance, within Border Sample

| | **West** | **East** | diff | p-value |
|---|---|---|---|---|
| *Pct Owner Occupied* | 48.04 | 48.61 | -0.57 | 0.96 |
| | (9.24) | (9.01) | | |
| *Pct Renter Occupied* | 42.05 | 38.62 | 3.43 | 0.81 |
| | (10.64) | (9.77) | | |
| *Pct Seasonal* | 5.51 | 6.60 | -1.08 | 0.88 |
| | (4.83) | (5.63) | | |
| *Median Income* | $64,940.42 | $71,138.07 | -6,197.64 | 0.79 |
| | (17,305.14) | (14,772.32) | | |
| *Per Capita Income* | $34,207.55 | $38,512.50 | -4,304.94 | 0.67 |
| | (5,253.05) | (8,589.57) | | |

Source: American Community Survey, 5-year estimates at the block group level, 2009-2014. Household-weighted averages of Census block group estimates, standard deviations in parentheses based on ACS confidence intervals.

area boundary are similar. Because the border does not coincide with borders for school attendance zones, tax purposes, or political representation, it is unlikely to be a prominent consideration for individuals choosing a new residence. While this boundary is completely independent of other administrative and legal boundaries, it is important to consider whether households sort on the basis of water district policy. This is highly unlikely, given that water and sewer charges comprise a small part of most households' budgets. However, Table 1.1 demonstrates that there appears to be balance running East-to-West across the boundary on key metrics, with these estimates representing population-weighted aggregations of the census block groups on either side of the boundary in the border area shown in Figure 1.4. These data are drawn from the American Community Survey, using the five-year estimates at the block group level for 2009-2014. I assign data from the block group level to accounts located within those block groups, and I correct the standard errors using the standard errors for the block group estimates.

While the border design strongly enhances the assumption of quasi-random assignment to treatment, this analysis uses a difference-in-differences methodology to utilize variation in treatment across both time and space. The exercise of using quasi-random variation from the housing market is more than sufficient to satisfy the parallel trends assumption, but it makes the assumption that the treatment group would have behaved similarly to the control group in the absence of the fine-based policy more credible. I present evidence of parallel trends in Figure 1.5, which provides consumption data for single-family residential households by district. The points in the graph are averages by week and district, weighted by the number of bills they represent. The West district generally bills less than one percent of customers in the last week of each month, leading to approximately one low and small dot for each month. Several key points are visible in this graph. First, consumers in the two water districts are strikingly similar before the application of the Stage 3 Water Emergency declaration. Second, during the drought period, water consumption decreased in both water districts. Third, water consumption is highly seasonal, and without a convincing counterfactual, any evidence of the policy's effect would be difficult to assess. In order to examine just the marginal impact of the fine-based policy, I use a difference-in-differences approach that removes variation driven by trends in conservation that are common across districts. Furthermore, I use reasonably frequent time-period fixed effects to account for seasonality and confounding variation in the time dimension.

Due to the form of geographic identification, I am not able to select accounts by exact distance to the border. Instead, I have geographic blocks which may be included or excluded based on each block's distance to the boundary. For my main specifications, I chose accounts within approximately two miles of the boundary in each direction. As shown in Appendix A.1.2, all key conclusions of the paper are robust to alternative distance specifications. The only exception to this is that the data and results are very noisy when restricting attention to less than half a mile from the border. However, these results are qualitatively similar despite the small sample.

Figure 1.5: Single-Family Residential Consumption by District, Border Sample Only

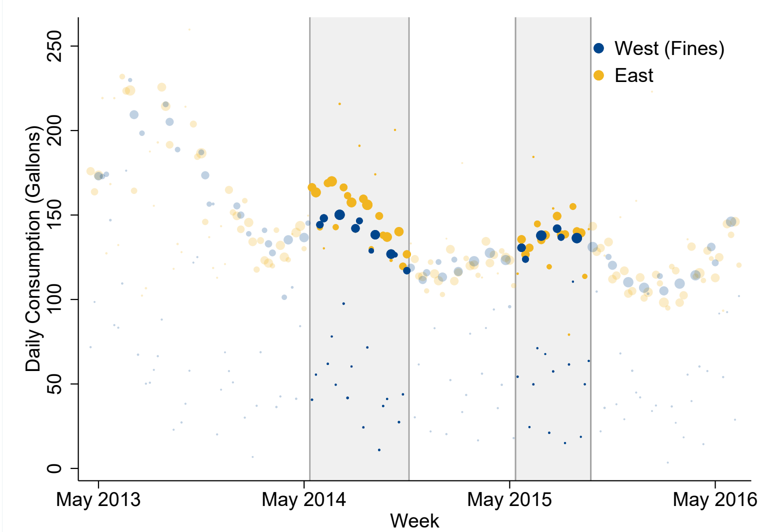


Average consumption collapsed by week of meter read and district. The size of each marker is proportional to the number of bills in each observation. Shaded areas represent periods when the West district was actively imposing fines.

### 1.1.3 Consumption Data

This paper leverages a combined dataset covering single-family residential accounts in both service areas. The core of this information is the universe of single-family residential billing records from January 1, 2013, through June 30, 2016. This includes 29,926 accounts in the West district and 17,955 accounts in the East district, with the primary border sample including 8,148 accounts in the West district and 5,705 accounts in the East district. Accounts are observed for an average of 26.75 billing periods in the border sample. Each billing record includes the date that the meter was read, as well as the recorded consumption and the account number.

Consumption is measured in discrete units of consumption, CCF, which are equivalent to 748.05 gallons. In order to compare accurately across billing cycles and for ease of interpretation, I convert consumption into daily consumption in gallons using the length of the billing cycle. The billing records for the East district include the date of the reading but not the number of days billed, and as a result, I must construct an algorithm to infer the number of days in the billing cycle. While the constructed daily consumption data are generally driven by the level of consumption, a small number of observations appear to be driven by an erroneous denominator. In order to ensure that the daily consumption metric accurately reflects consumption, I drop the highest one percent and lowest one percent of readings in terms of gallons consumed per day by an account for all estimation. This exercise does not substantially alter the results.

I also incorporate information on price schedules to examine the price elasticity of households and the extent to which households respond to the price component of the fine-based policy. Household data include a variety of characteristics, such as elevation and jurisdiction, which determine the monthly fixed and variable costs for the households. Each billing record is matched with both pricing details and the pricing regime in effect at the time.

## 1.2 Empirical Specifications

While the border design restricts the sample to similar households in the treatment and control group, the core method of this study is a difference-in-differences design. I apply this design first to measure the aggregate effect and then to document the channels through which the policy was successful in generating conservation. Specifically, I look at the effect by pre-period consumption to examine how households with different exposure to the policy responded. In order to cleanly examine the extent to which price plays a role in these findings, I also estimate a simple model of demand.

All of the results presented in this paper use two-way fixed effects and two-way clustering with the `reghdfe` package authored by Sergio Correia, as described in Correia [2016]. All regressions include fixed effects for the household and week of time, with standard errors clustered in both of these dimensions. Furthermore, the command drops singleton clusters to avoid overstating statistical significance, as noted in Correia [2015]. For comparison across specifications and to ensure appropriate inference, all regression results are restricted to accounts that are not missing a bill during the first bill of the fine-based policy. While additional accounts may serve as appropriate control or treated units, the key to inference in this panel setting is the behavior of individuals across time. Notably, the results are robust to a variety of specifications for the standard errors, border width, and controls.

### 1.2.1 Impacts on Aggregate Consumption

The first effect to consider is the impact on aggregate consumption among accounts in the West district. Since identification rests on the parallel trends assumption, the treatment effect is identified using equation 1.1, where consumption by household $i$ that is read on date $t$ covering billing period $b$ is measured in gallons per day.[11]

$$\ln(\text{Cons}_{i,t}) = \delta \text{Policy}_{i,b} \times \text{West}_i + \beta \text{X}_t + \gamma_i + \psi_w + \varepsilon_{i,t} \tag{1.1}$$

The measure of gallons per day, $\text{Cons}_{i,t}$, is calculated by converting the observed consumption in CCF into gallons and dividing by the number of days in the billing cycle. This ensures an accurate comparison across billing cycles with differing lengths. For ease of interpretation, the regressions will utilize the natural log of daily consumption to estimate a treatment effect, $\hat{\delta}$, that approximates the observed conservation as a percentage. The policy variable, $\text{Policy}_{i,b}$ is defined based on the water district's implementation guidelines, with bills having all consumption within the policy period being subject to fines. $\text{West}_i$ is a binary indicator of whether a household is in the West (fine) district. The control variables $\mathbf{X}_t$ are weather measures, with only one relevant weather station for all observations.[12] A precipitation variable controls for the sum of total precipitation during the previous 30 days. Evapotranspiration and maximum and minimum temperature variables refer to the average of the respective measure over the previous 30 days. I include a second-order polynomial for all weather variables. In addition, account

[11] The panel data is effectively missing all other dates other than dates that a household had its meter read, but this construction is meant to coerce the rolling billing cycles of true consumption into a form that allows for appropriate temporal controls. Mathematically, only one observation per household per billing cycle is included.

[12] Weather controls do not substantially impact the results. The choice of treatment and control households eliminates spatial heterogeneity in exposure to weather shocks.

and week-of-time fixed effects are included as $\gamma_i$ and $\psi_w$, respectively.[13] Standard errors are clustered by account and week of time.

In order to provide the most consistent estimates, I restrict the sample to accounts approximately within two miles of the border. While this restriction reduces power, it is an attempt to provide a homogeneous group of accounts. In Appendix **??**, I demonstrate that the effects hold for the broader sample. I estimate the effect for each year separately using samples that do not overlap in time. Each regression uses the five months before implementation and the five months after implementation, without the month of implementation. Results are robust to the inclusion of the month of implementation and qualitatively robust to a wide variety of alternate temporal samples. Data prior to the 2014 implementation sample contains a large number of bimonthly bills, which generate problems for the time-period fixed effects specification when mixed with monthly bills. After six months, the policy was temporarily removed for the winter. For symmetry, I include the same amount of pre-implementation and post-implementation months in 2014 and 2015.

### 1.2.2 Reduced Form Effects Across the Distribution

Facing this fine-based policy, a rational and informed consumer would anticipate the new price schedule and conserve to avoid paying the penalties. To understand the extent to which consumers are avoiding penalties and conserving as a function of their ex ante perceived risk, I examine a difference-in-difference effect of the response across all points in the consumption distribution to examine which households are responding to the policy. The specification is shown below in equa-

[13] The billing cycle is closed on a given date $t$, and I use the week of that date for the fixed effects. A finer level of fixed effects, such as the date, could lead to erroneously using the time fixed effect to control for the characteristics of the majority of accounts read on that date. Because the date a meter is read is correlated geographically, this can confound the intended purpose of controlling for weather and other similar effects. A coarser level of fixed effects, such as the month, would inaccurately group bills read weeks apart, which may have faced substantially different weather and other exogenous factors.

tion 1.2:

$$\ln\left(\text{Cons}_{i,t}\right) = \sum_{j=1}^{J} \delta_j \cdot \text{Policy}_{i,b} \times \text{West}_i \times 1\left\{\text{Cons}_{i,ref} = j\right\} + \sum_{j=1}^{J} \kappa_j \cdot \text{Post}_{i,t} \times 1\left\{\text{Cons}_{i,ref} = j\right\} + \beta \text{X}_t + \gamma_i + \psi_w + \varepsilon_{i,t} \quad (1.2)$$

The specification is similar to the aggregate effect specification in equation 1.1, except that this specification includes an array of difference-in-difference indicators for each level of pre-period consumption, $j$. I conduct this analysis using bins for consumption during the bill before implementation ("prior bill") and the maximum observed consumption by the household in the pre-period data ("maximum pre-period"). For the prior bill specification, $J = 15$, and all consumption above 15 CCF in the prior bill is counted as a single bin. For the maximum pre-period specification, $J = 25$, and all consumption above 25 CCF in the pre-period is counted as a single bin. The coefficients $\kappa_j$ address the mean reversion and persistence by pre-period consumption level that is common across the two water districts.

As before, the estimation is conducted only for the border sample and includes weather controls and household-specific and week-of-time-specific fixed effects. I present results only for the first year of the policy, not only because of the attenuated estimates for the second year, but also because the pre-period measures would either be endogenous to 2014 behavior if using 2014 and 2015 data or less meaningfully connected to whether the policy binds in 2015 if using data from before the initial intervention. Standard errors are again clustered by household and week of time, with the same temporal range as the 2014 regression for aggregate effects.

The results of this exercise provide reduced-form evidence to test which of the possible hypotheses can explain the observed data. Risk aversion and uncertainty would lead to a significant response among households with a low probability of being fined, as they consider the penalties of consumption higher in the distribution. However, these hypotheses would also predict a large increase in responsiveness when considering households with higher probabilities of being fined.

Because the fine scales with consumption, households higher in the distribution would then also be considering – and reacting to – the even larger fines still higher in the distribution. This would be true if individuals are strongly risk averse, but it would also be true if individuals have low certainty regarding their own consumption or inaccurately believe that their own consumption is higher than it is. A plateau in the relationship between response and exposure to fines would require a sophisticated correlation between inaccurate beliefs and levels of risk aversion which is not grounded in available theoretical or empirical evidence, with consumers in the low end of the consumption distribution having substantially different risk preferences or misbeliefs.

This reduced-form evidence can also test the possibility that the results stem from over-weighting low-probability events, such as preferences proposed by prospect theory. This set of preferences is consistent with a response that is stronger than predicted by price elasticity and risk aversion among low users and a response that is potentially lower than predicted among high users. Tversky and Kahneman [1992] estimate preferences of this sort in a lab setting, and I will discuss the extent to which my results would require much more skewed preferences than they estimate. Having addressed these explanations, I note that the evidence is consistent with a world in which all households with scope to conserve reduce consumption in response to the signal given by the policy. I explain the effects across the distribution by combining this signal-induced response with fine avoidance behavior and optimization with respect to the new price schedule, including fines.

### 1.2.3 Demand Estimation

In a fourth specification, I can adjust for price elasticity under a range of assumptions about prices, primarily that households respond to the average price. Because prices are determined by quantity, I instrument for prices with the price regime.

To understand the extent to which this perception of the policy or a lack of information might be driving the results, I estimate a simple model that incorporates uncertainty and an indicator for the policy's presence. Following the work of Ito (2013, 2014), I empirically test for consumer response to price signals other than marginal prices by assuming a simple demand function. As shown in equation 1.3, I assume that demand by household $i$ in period $t$ is the product of household-specific factors, time-specific factors, a drought emergency signal, and a function of the price of water.

$$\text{Cons}_{it} = \lambda_i \cdot \tau_t \cdot d_{it}^{\delta} \cdot [\tilde{p}_{it}(\text{Cons}_{it})]^{\varepsilon_p} \cdot \epsilon_{it} \tag{1.3}$$

Here, $\lambda_i$ are household fixed effects and $\tau_t$ are time fixed effects, each of which shifts demand either across households or across time. These effects account for heterogeneity in preferences across households and for seasonality and time trends in demand. I also add an indicator for whether the utility is in a drought emergency, $d_{it}$, with an exponent that scales the importance of this information for demand, $\delta$. Under the assumption that the drought declaration has no importance other than through a price regime shift, $\delta$ would be equal to zero. If the drought declaration does have a separation mechanism which encourages conservation, then $\delta < 0$.

When including price, I specifically denote the perceived price function as $\tilde{p}_{it}(\cdot)$ to capture the potential for households to misperceive the marginal price of water or to use heuristic methods, such as the average price. Notably, I use both the average price and the marginal price of water in various specifications, and I find that the average price of water more closely fits the data. While limited by the variation available in the data, I also attempt to estimate more flexible parameterizations of $\tilde{p}_{it}(\cdot)$. The $\epsilon_{it}$ term captures idiosyncratic demand shocks.

In practice, I estimate the demand equation after taking the natural log of both sides, as shown in equation 1.4.

$$\ln(\text{Cons}_{it}) = \eta_i + \omega_t + \delta \cdot \ln(d_{it}) + \varepsilon_p \cdot \ln[\tilde{p}_{it}(\text{Cons}_{it})] + u_{it} \tag{1.4}$$

Here, $\eta_i$ are household fixed effects in the log-consumption specification and $\omega_t$ are week-of-time fixed effects. For tractability reasons, I impose that $d_{it} \in \{1, e\}$ such that $\ln(d_{it}) \in \{0,1\}$. Note also that $\omega_t$ captures the true underlying state of drought across households in both districts over time. Therefore, $\hat{\delta}$ is only capturing effects derived from the differential signal in the fine policy district and not from the common hydrological conditions and stated drought emergency stage.

In order to estimate equation 1.4, however, I first need to handle the inherent endogeneity of estimating demand when the price is a function of the choice variable. Because I observe consumption under a variety of price schedules, I instrument for observed prices with the price regime in effect during each bill. This is a relatively weak set of instruments, because each change in the pricing regime only represents generally small changes in prices and the price changes are differentially impactful at different points in the consumption distribution. To overcome these challenges, I include the full range of time for which I have observations for both districts, January 2013 through June 2016. For comparison, I also present the reduced form results here for this increased temporal sample. Estimation is still restricted only to accounts in the border sample.

A challenge for this estimation is that, while the expanded temporal sample aids in the identification of price elasticity, the expanded sample creates new challenges to the identification of the true effects of the policy. By extending the beginning of the sample backward in time, I include more billing cycles that cover two months instead of just one. This reduces the effectiveness of the time fixed effects, since the seasonal demand shocks may be different for the two months leading up to a date, as opposed to just the one month prior. It would be useful to allocate the bimonthly bill across the two months, but this would require strict structural assumptions or far more data on monthly and bimonthly consumption in the pre-period.

Including additional time after the removal of the policy requires making assumptions about when the policy's effect is over. I have included an indicator for the policy that turns off with the policy's removal, but households may have made permanent changes to their consumption patterns. For example, if the policy induces investment in appliances or even persistent changes to irrigation or other water-intensive behavior, households will not quickly return to prior consumption patterns. This would attenuate my estimates under my current specification, but including all time after the policy's implementation as the post-period would also attenuate my estimates.

In addition to these concerns, the dynamic situation led the East district to make policy adjustments over time which I cannot fully capture. Under the assumption that pecuniary policies have an effect through other channels than just price elasticity, then using the price elasticity for both districts still will not fully capture the effects of the East district's sequence of actions. By repeatedly imposing 'emergency' rate increases, the East district may have generated reductions among its consumers beyond the price effects. My measure of price elasticity is an average effect across standard and emergency rate increases, partially for power purposes, and will not subsume this effect.

It is also important to control for the ways that households may react to unexpected fines. For this specification, I include both the effects of fines in the first months of the policy each year, as well as an estimate of the demand shift each year separately. In equation 1.5, $\hat{\delta}_{2014}$ and $\hat{\delta}_{2015}$ estimate the demand shift for each year of implementation, while $\hat{\epsilon}_p$ estimates the observed price elasticity. The price elasticity is assumed to be constant across households and time, while fixed effects for each household and week of time account for level shifts in demand.[14]

$$
\begin{aligned}
\ln\left(\text{Cons}_{i,t}\right) = & \phi_{2014} \cdot f\left(\text{fine}_{i,b=fm_{2014}}\right) + \eta_{2014} \cdot h\left(\text{Cons}_{i,b=fm_{2014}}\right) \\
& + \phi_{2015} \cdot f\left(\text{fine}_{i,b=fm_{2015}}\right) + \eta_{2015} \cdot h\left(\text{Cons}_{i,b=fm_{2015}}\right) \\
& + \delta_{2014} \cdot \text{Policy}_{i,b}^{2014} \times \text{West}_i + \delta_{2015} \cdot \text{Policy}_{i,b}^{2015} \times \text{West}_i \\
& + \epsilon_p \cdot \ln\left(\widehat{\text{Price}}_{i,t}\right) \\
& + \beta \cdot \boldsymbol{X}_t + \gamma_i + \psi_w + \varepsilon_{i,t}
\end{aligned} \tag{1.5}
$$

[14]The available variation in prices is insufficient to identify heterogeneous price elasticities.

## 1.3 Results

In this section, I first present reduced-form evidence of an aggregate effect from a difference-in-differences specification. I then examine the impacts on both targeted and non-targeted households. I find not only that some consumers act as if they were surprised to be fined in the first month of implementation, reacting by strongly curtailing their usage, but also that many consumers who are essentially inframarginal also respond to the penalties by reducing their usage. Specifically, I find statistically significant conservation among accounts with less than a 15 percent chance of being fined under business-as-usual. I present both reduced-form and demand estimation evidence that suggests a large component of the policy's effectiveness was driven by the policy changing demand by signalling the importance of conservation.

### 1.3.1 Impacts on Aggregate Consumption

Using single-family accounts in this restricted sample, I find that the fine-based policy reduced water consumption by an average of nearly 18 percent in 2014 and approximately three percent in 2015. Importantly, these results are measures of relative reduction in comparison to households in the East district, which I will discuss further below. Columns 1 and 2 in Table 1.2 present aggregate treatment effects for the first and second year of the policy, respectively, using the border sample with a full set of weather controls. Table A.1 in the Appendix provides evidence regarding the robustness of these estimates to the inclusion or exclusion of weather controls. Table 1.4 provides attenuated estimates from using a broader temporal sample with threats to identification which I discussed in the previous section. Even under conditions that are likely biased towards zero, I see statistically significant, meaningful, and qualitatively similar results.

It is noteworthy to recall that the East district imposed a mild rate increase for all tiers and accounts approximately two months into the fine-based policy. As Figure 1.6 demonstrates, the policy continued to elicit reductions throughout the

Table 1.2: Difference-in-Difference Results

| | Outcome Var: Log of Daily Consumption (Gallons) | | | |
|---|---|---|---|---|
| | 2014 | | 2015 | |
| | 1 | 2 | 3 | 4 |
| *Fine-Based Policy ATE* | -0.163*** | -0.170*** | -0.0241 | -0.0221 |
| | (0.0477) | (0.0563) | (0.0153) | (0.0152) |
| East Mean - Post | 140.76 | | 129.26 | |
| (Standard Error) | (86.92) | | (80.75) | |
| Controls | Precipitation, Max Temp, Min Temp, Evapotranspiration | | | |
| Fixed Effects | Household, Week of Time | | | |
| Cluster Variables | Household, Week of Time | | | |
| $R^2$ | 0.723 | 0.720 | 0.732 | 0.728 |
| Obs. | 96,048 | 86,587 | 100,752 | 91,299 |

Sample restricted to five months before and after the month of implementation (columns 2, 4), plus the month of implementation in columns 1 and 2. Lowest one percent and highest one percent of consumption observations dropped. Cluster-robust standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

summer months, but its effect was somewhat attenuated after the East district's policy change. Unfortunately, it is difficult to entirely disentangle whether this is a declining effect, attenuation from a response in the control group, or even statistical noise given the short window. Given the change in the East district, these estimates could be considered a lower bound.

Furthermore, the estimates presented here for the second year of the policy are attenuated by two additional rate increases enacted by the East district in 2015. As

Figure 1.6: Monthly Aggregate Treatment Estimates

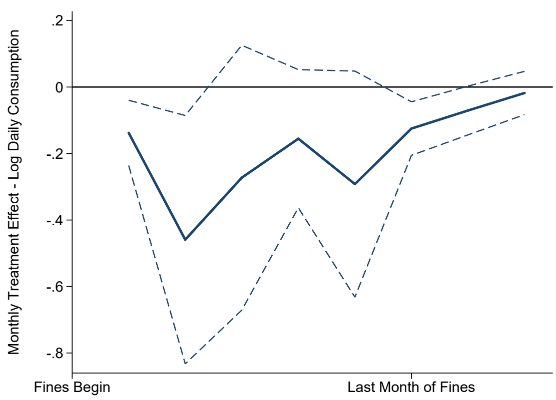

discussed previously, the East district faced local and state incentives to conserve beyond what East district consumers accomplished in 2014, and this led to the district becoming a less viable control group for the second year of the West district policy.

### 1.3.2 Mechanisms and Model Estimation

While the aggregate impact of the policy is notable, the only way to uncover the mechanisms by which the policy generated conservation is to examine the heterogeneous effect on households with different levels of exposure. Figure 1.7 plots raw collapsed consumption figures for households in the West (treated) and East (control) district plotted over time. There is clear evidence that the policy both caused reductions in the lower end of the distribution and was more impactful among households that were consuming more water and were, therefore, more exposed to the policy. There are two explanations for this behavior. The first is that many targeted households are forward-looking and avoid fines by proactively conserving before the policy is implemented. The second is that households who are unlikely to face fines are also responding to the policy.

In a neo-classical framework, the imposition of a fine-based policy like the one described here would be hypothesized to have almost exclusively a price effect. The water district raised the marginal price of consumption for water consumption levels above a certain quantity per month, and high users should be expected to conserve as they maximize their utility subject to a budget constraint. However, there are multiple potential avenues by which the policy would have other price and non-price effects. This section examines the potential roles of risk aversion, information gaps, and behavioral factors in determining the effects of this policy.

Panel A of Figure 1.8 plots the difference-in-difference coefficient estimates, and their 95 percent confidence intervals, for the effect of the policy by pre-policy consumption level. Confirming the evidence from the raw data, it is clear that there is both an aggregate response in the lower end of the distribution and a stronger re-

Figure 1.7: Responses to Fines across the Distribution, by District

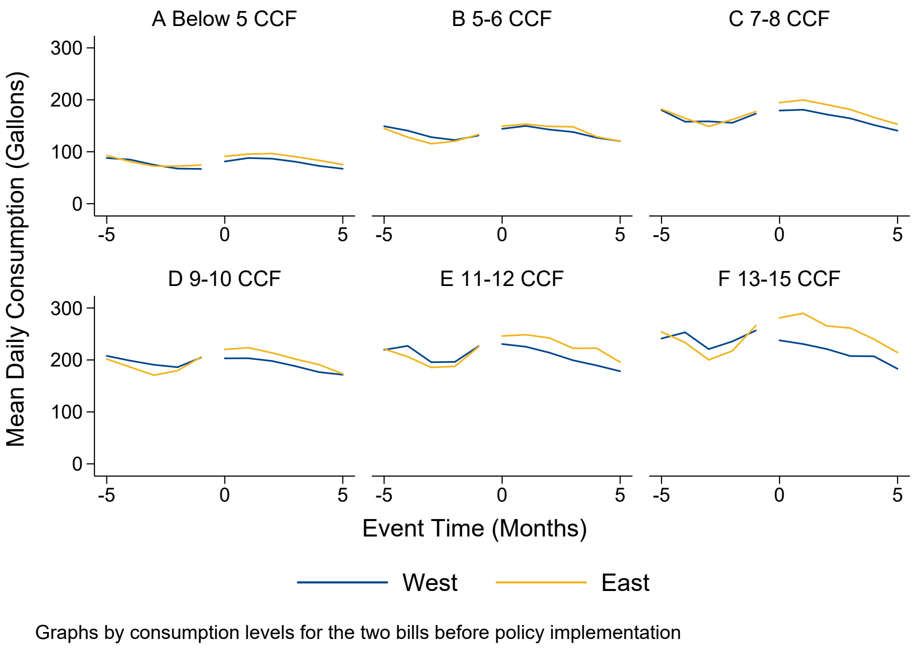


Figure 1.8: Effect of Fines by Pre-Implementation Consumption Level

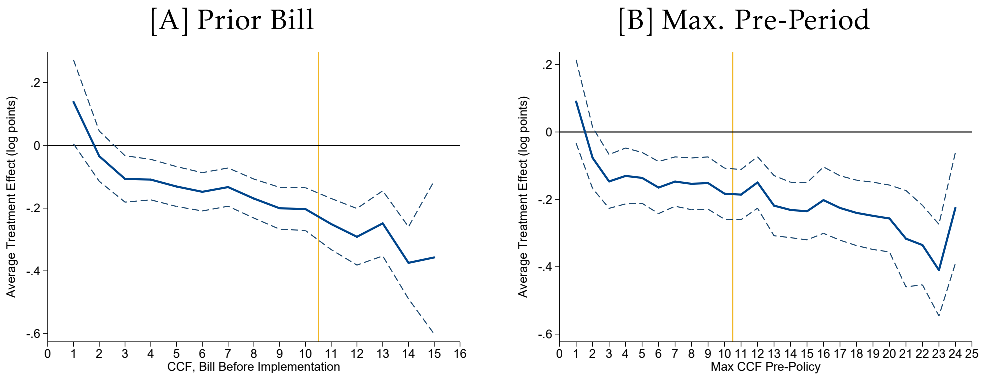


Data includes the five months before implementation, five months after implementation, and the month of implementation for the border sample during the 2014 policy period. Observations above 2,500 gallons per day omitted (<0.02% of observations).
Fixed effects and standard errors clustered by account and week. 95% Conf. Int. shown.

Table 1.3: Inframarginality of Pre-Policy Bins, Border Sample Only

| Prior | % Accounts Exceeding Allotment during Summer 2014 | | | | | |
|---|---|---|---|---|---|---|
| Consumption | **East** | **West** | diff. | *p-value* | $n_{\text{East}}$ | $n_{\text{West}}$ |
| *3* | 5.90 | 4.77 | -1.13 | 0.372 | 542 | 733 |
| *4* | 7.97 | 6.93 | -1.04 | 0.475 | 552 | 765 |
| *5* | 13.5 | 7.53 | -6.01*** | 0.000 | 502 | 770 |
| *6* | 20.3 | 16.6 | -3.68 | 0.133 | 409 | 626 |
| *7* | 30.1 | 25.0 | -5.12 | 0.110 | 322 | 476 |
| *8* | 47.1 | 29.8 | -17.3*** | 0.000 | 257 | 386 |
| *9* | 61.7 | 47.4 | -14.3*** | 0.006 | 141 | 253 |
| *10* | 65.9 | 64.2 | -1.74 | 0.759 | 129 | 159 |
| *11* | 77.3 | 67.5 | -9.85 | 0.139 | 75 | 123 |
| *12* | 87.0 | 83.7 | -3.34 | 0.589 | 54 | 92 |

Prior consumption is the billed 100 cubic feet in the bill before the policy was implemented. Data is collapsed at the account level and includes only accounts with a bill during the first month of the policy for variable construction purposes.

sponse among households with higher pre-policy consumption. To understand the extent to which this might be rationalized by risk aversion, I also present estimates of the probability of crossing above the fine threshold among households on either side of the border in Table 1.3, using the same bins of pre-policy consumption as presented in Panel A of Figure 1.8.

As shown in Table 1.3, the baseline (East) probability of exceeding the policy allotment is increasing in prior consumption but low even for accounts consuming five or six CCF in the bill before the policy was implemented. Still, Figure 1.8 and Table 1.3 suggest that treated households in the West district with this low expected exposure to the policy responded to the policy by conserving a meaningful amount of water. The treatment effect is evident for prior consumption as low as three or four CCF, but a risk aversion parameter strong enough to explain such behavior would imply far stronger reductions among higher users. By contrast, the treatment effect is not statistically different between prior consumption levels as low as three CCF and as high as 13 CCF. While the point estimates differ, this similarity in the response across households facing drastically different probabilities of facing fines or probability-weighted expected fines in dollar terms suggests that risk aversion could not simultaneously explain the responses observed across the distribution.

Figure 1.9: Response by Exposure and Variance

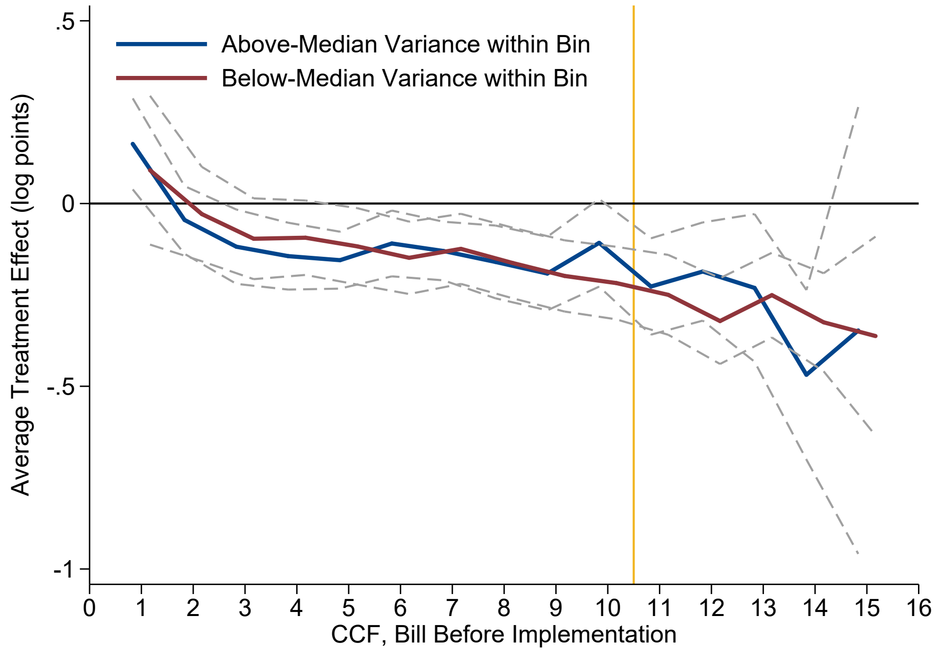


In addition to these estimates comparing all households across the policy and control districts within bin of prior consumption, I can disaggregate these estimates by observed pre-period variance in consumption. Specifically, Figure 1.9 plots the treatment effect by prior bill consumption separately for households that had above- or below-median variance in the pre-period, within their consumption bin. An above-median-variance household in the 3 CCF bin, for example, is a household with particularly high variation in consumption, and such a household would rationally expect a higher probability of being fined. While Figure 1.9 suggests that such low-consumption, high-variance households may have responded slightly more, a robustness check to including additional pre-period data actually confirms that this difference is only statistical noise (See Figure A.6). In other words, low-use households with historically high variance in consumption do not respond more strongly than low-use households with historically low variance in consumption, in contrast to what a risk-aversion explanation would predict.

An additional consideration might be that consumers are not well aware of their most recent consumption. Alternately, it might be worth considering how the threshold compares to the most that a household has ever consumed. Panel B of Figure 1.8 plots the difference-in-differences estimates (solid line) and their 95 percent confidence intervals, grouping households by their maximum pre-intervention consumption within the available data.[15] Consistent with the findings in Panel A and in Table 1.3, I observe a pattern of behavior that is highly inconsistent with a story of risk aversion. Households here are responding even if they never consumed above, for example, 40 percent of the threshold. This is remarkable, given that the policy bills include the previous 24 months of observed consumption, and these households are shown that they are in no danger of receiving a penalty. Indeed, the extent to which the response is very flat across the distribution strongly implies that the policy induced broad-based conservation that was relatively uncorrelated with pre-period consumption and separate from its price effects.

Panel A of Figure 1.10 reinforces this point by graphing the results from Panel A of Figure 1.8 by the predicted total dollar value of fines incurred during the 2014 policy period. I calculate predicted total fines using the policy-period consumption behavior of control households who share the same consumption as a given treated household for the bill prior to policy implementation. Using this metric, I find effects of 10 to 15 percent for households with a predicted exposure of between \$30 and \$50. Households with a predicted exposure of approximately \$700, by contrast, responded by approximately 30 percent.

Using Panel B of Figure 1.10, I can also assess the extent to which the observed behavior is the result of over-weighted low-probability events and under-weighting high probability events. Tversky and Kahneman [1992] provide laboratory evidence for this behavior as a product of prospect theory. Converting their estimates to predicted behavior within my framework, the observed data would require misperceptions of the probability at least approximately 50 percent more extreme, particularly among users below an approximately 20 percent probability of being fined. It is important to note, moreover, that this exercise is biased towards

[15] Since many accounts were on bimonthly billing at some point in the pre-period, I convert the data for bills with bimonthly billing cycles by dividing the total consumption by two and rounding up to the nearest whole number. The results are robust to alternate approaches.

Figure 1.10: Effects by Predicted Exposure, 2014

[A] Fine Exposure in Total Dollars

[B] Fine Exposure by Probability of Any Fine

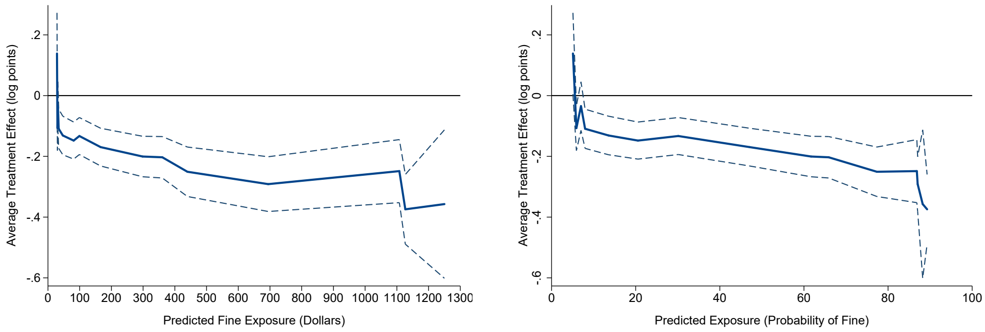


their estimates. This is because the (expected) value of the fine is increasing in the probability of being fined, while the estimates from Tversky and Kahneman [1992] are conducted within a fixed value for the negative outcome. As a consequence, the response among users with lower probabilities in my setting should be disproportionately low, while the response among users with higher probabilities in my setting should be disproportionately high. By contrast, I find that the laboratory estimates underestimate the response among users with lower probabilities.

In addition, two related hypotheses may explain the observed behavior. The first hypothesis is that households are trying to optimize consumption choices but lack the information necessary to do so. Each bill covers approximately one month of consumption, and it can be very difficult to map billing information to uses and choices throughout the billing period. Usage data is generally only available once per month (or once every two months when billing bimonthly), and this information provides a single total number for that period.[16] By contrast, usage choices are made daily. Under the assumption that households do not understand usage perfectly, households would optimize subject to uncertainty. A second hypothesis is that households do not know or understand the billing structure in detail. Instead, households use a heuristic, such as average prices, in their optimization. In practice, a reduced form empirical test for either or both of these hypotheses would test

[16]For the primary results of this paper, the sample has been constructed to minimize the inclusion of bimonthly billing data. However, bimonthly billing was prevalent in the summer of 2013 for these accounts, and results are robust to the inclusion of bimonthly bills.

the theoretical implication that consumers respond not to the true billing structure applied to their true consumption but to a set of price and quantity combinations that includes uncertainty (but not bias).[17] Mathematically, these hypotheses differ in where the uncertainty lies, but they can be reduced to an estimation of demand in which consumers optimize with respect to perceived prices which differ from actual prices. I therefore test for both of these explanations by addressing the potential for households to respond to perceived prices of water other than the marginal price. To the extent possible, I examine how prices both for lower tiers of consumption and for higher tiers of consumption affect water consumption choices.

I also propose a separate explanation that households perceive social pressure to conserve as a result of the policy. While the estimates control for social norm responses to the declaration of a Stage 3 Emergency and responses to statewide and local calls for conservation, the imposition of "water rationing" may have been interpreted as a signal of the severity of the water shortage or the importance of conservation locally. Such a mechanism would shift the demand curve for water inward, leading to lower consumption even when facing the same prices. As shown below, I posit that this mechanism is a primary driver of the observed results, accounting for the vast majority of the observed differences across districts during the policy period.

With these caveats in mind, Table 1.4 presents results for this temporally-expanded sample for both the reduced-form equations from the previous sub-sections and for equation 1.4. The first column presents estimates for the aggregate effect of the policy during each of the temporary implementations. Comparing these estimates

[17] If biased beliefs about consumption or prices are uncorrelated with observed consumption levels, this will have no bearing on the estimates I find. If consumers' beliefs about consumption or prices are biased depending on their observed consumption levels, this could confound my estimates. A priori, there are reasons to believe that consumers with lower consumption would understand their consumption more accurately in levels, since they may have a mechanically low variance. However, there are also reasons to believe that consumers with higher consumption and, therefore, higher bills would understand their consumption more accurately because it is rational for them to be more attentive as the significance of the water bill increases as a portion of a household's budget. Observed reduced form evidence would be consistent with a highly non-linear bias where only the middle section of the distribution over-estimates their consumption but this bias decays then levels out. While possible, this explanation seems less plausible by virtue of its complex assumptions without theoretical basis.

with those from Table 1.2, the attenuation effects for this sample are small but evident. However, I still find a statistically significant effect of around 13 percent for the first year and a statistically insignificant null effect for the second year. Notably, these estimates are for the effect on consumption during the policy period. In the second column, I estimate equation 1.4 using instrumented average prices. I can use alternate approaches, but no other perceived price function fits the data as well as average price, and the choice of perceived price functions does not qualitatively alter my findings. While statistically insignificant, I estimate a price elasticity of demand for water around -0.046, which is more price inelastic than the estimates from Ito [2013]. These are short-run estimates, however, and the 95 percent confidence interval for the estimate includes the point estimates from Ito [2013]. More notable, however, is that the estimate of the policy's non-price effect is similar to the policy's overall effect.

Table 1.4: Model Estimation, Average Price Elasticity

| | Reduced Form | Perceived Price Model |
|---|---|---|
| Policy-Induced Demand Shift | | |
| 2014 | -0.132*** | -0.135*** |
| | (0.0452) | (0.0452) |
| 2015 | -0.00334 | -0.00636 |
| | (0.0161) | (0.0149) |
| Average Price Elasticity | | -0.0455 |
| | | (0.0540) |
| Controls | Precipitation, Max Temp, Min Temp, Evapotranspiration | |
| Fixed Effects | Household ID, Week of Time | |
| Cluster Var. | Household ID, Week of Time | |
| $F$ (First Stage) | | 86.38 |
| Centered $R^2$ | 0.688 | 0.677 |
| Obs. | 299,443 | 299,443 |

Lowest one percent and highest one percent of consumption observations are removed. Households not observed in the first policy month of either year are also omitted. Cluster-robust standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Taking these estimates at face value, the initial implementation of the policy was equivalent to a more than 200 percent increase in average prices. Given that the policy increased average prices by more than 200 percent only for a very small fraction of consumers, there is compelling evidence that the policy had more than a price effect, including both reactions to being fined and avoidance behavior. As suggested by the reduced-form evidence presented previously, the policy appears to have had a direct impact on consumer demand.

The impact of the policy through a change in perception is similar to the observed behavior in Gneezy and Rustichini [2000]. While the title of their paper, "A Fine Is a Price," might seem to be the opposite of the findings presented here, the broader point that the authors make is that the imposition of a fine on a behavior changes the perceptions of those facing the fines. In the case of tardiness in picking up children from day care, the fine was small and lowered the parents' perceptions of how important it was for them to be timely. In this example, large fines on consumption in the higher end of the distribution increased households' perceptions of how important it was for them to conserve water.

## 1.4 Conclusion

Financial penalties form a fundamental piece of the public policy toolkit, invoked in essentially every sector. Neoclassical and even many behavioral economic predictions of how these policies impact people generally assume that individuals react almost exclusively to the financial content of the policy. However, this paper identifies a separate and important mechanism by which fine-based policies induce contributions to the public goods. The evidence provided here indicates that the enactment of fine-based policies conveys a signal of the importance of making these contributions to the public good, and this channel can account for a large component of the total policy impact.

Combining a policy experiment with quasi-random assignment to treatment across an arbitrary border, I examine this policy channel in the context of residential water consumption. With some households facing a minimal probability of facing the fines and bins of households with varying exposure to fines, I am able to identify similar conservation across extensively different direct exposure to the financial penalties. In addition to examining policy effects by consumption shortly before the policy's enactment, I examine policy responses by a household's maximum consumption in the observable pre-period. Both of these approaches indicate a response that increases in exposure but includes strong responses among low-exposure individuals and limited increase in response with exposure.

All of these results are relative to control households who also conserved in response to ambient media pressure and other common factors. These findings are also in a context where each bill includes historical consumption, current consumption, and a red line indicating the penalty threshold. In addition, the district imposing the policy conducted extensive direct outreach. The observed pattern of behavior is not consistent with or fully explained by risk aversion, prospect theory, or a lack of information. Instead, the most consistent explanation is a shift in demand induced by the policy itself through a signal of the importance of conservation. This mechanism is important in understanding — and predicting — the consequences of public policy. Under the assumption that this 'severity signal' channel is a substantial component of policy effectiveness in other sectors, it may change the expected relative performance of potential policy choices. Regardless of specific policy content, however, the biggest impact of a policy may be via the signal of importance provided simply by enacting the policy.

# Chapter 2

# Do Regulatory Policies Crowd Out Social Pressure for Resource Conservation?

**Coauthored with Jeremy West, Robert Fairlie, and Liam Rose.**

Society relies heavily on both regulatory policies and encouragement of voluntary actions to increase resource conservation and other prosocial behavior. Increasingly, policymakers combine the two tactics, strengthening both regulations and formalized normative interventions to internalize externalities and promote voluntary contributions towards public goods. For example, the U.S. Forest Service attempts to reduce forest fires using both extensive prohibitions and large-scale media campaigns of normative messaging such as "only YOU can prevent wildfires."[1] One outstanding question that is thus critical for the efficacy of such two-fold policy approaches is how stronger regulations interact with normative interventions targeting the same behavior.

The broader behavioral literature provides compelling support for the notion that stronger regulatory incentives may crowd out encouragement for voluntary re-

[1] c.f. https://firerestrictions.us and https://smokeybear.com.

source conservation. An extensive literature finds that pecuniary incentives to contribute to a public good often have psychological or behavioral aspects that displace intrinsic motivations to contribute to the same good. As characterized by Gneezy et al. [2011], "Monetary incentives have two kinds of effects: the standard direct price effect, which makes the incentivized behavior more attractive, and an indirect psychological effect [that] in some cases works in an opposite direction to the price effect and can crowd out the incentivized behavior." Such crowding out of prosocial behavior by financial incentives has been demonstrated for a wide range of outcomes, from blood donations and charitable giving to day care pick up timeliness and nature conservation [e.g. Titmuss, 1970, Gneezy and Rustichini, 2000, Bénabou and Tirole, 2006, Ariely et al., 2009, Bowles and Polanía-Reyes, 2012, Rode et al., 2015]. Given this evidence of crowd out of intrinsic motivations, it is natural to ask: does strengthening regulatory policy incentives also crowd out the incentives of normative interventions for engaging in prosocial behavior such as resource conservation?

In this study, we examine whether the effects of external social pressure – the influence on people by their peers – are crowded out by a sharp increase in regulatory incentives targeting the same behavior. The specific intervention we study involves peer comparisons, a form of social pressure used widely to encourage behaviors such as charitable giving, civic participation, resource conservation, and workplace productivity [Frey and Meier, 2004, Gerber et al., 2008, Mas and Moretti, 2009, Allcott, 2011, Ferraro and Price, 2013, Castillo et al., 2014, Smith et al., 2015]. The context we study is residential water use, and our empirical research question is whether strengthening the enforcement of regulations pertaining to landscape irrigation reduces the effectiveness of normative peer comparisons that encourage household water conservation. To our knowledge, our study is the first to directly test whether stronger regulatory policies crowd out social pressure interventions by comparing the efficacy of randomized social pressure before versus after an exogenous strengthening of regulatory incentives targeting the same behavior.

To empirically investigate this question, we use data from WaterSmart Software and a Southern Californian water utility that measures hourly residential water consumption using Advanced Metering Infrastructure (AMI). A WaterSmart field experiment substantially increased social pressure for randomly-selected households to conserve water by providing them with Home Water Reports (HWR) comparing their water consumption with that of their neighbors. In this baseline condition, HWR should be highly effective; prior studies testing similar treatments in Georgia and California generally find water conservation effects ranging from three to six percent [Ferraro et al., 2011, Ferraro and Price, 2013, Mitchell and Chesnutt, 2013, Bernedo et al., 2014, Brent et al., 2015, Jessoe et al., 2019, Bhanot, Forthcoming].[2] Moreover, the literature finds that this effectiveness is attributable primarily to the social pressure component, rather than to any technical or financial information included in HWR [Ferraro and Price, 2013, Mitchell and Chesnutt, 2013, Allcott and Rogers, 2014, Brent et al., 2015, 2017, Bhanot, Forthcoming]. Altogether, the existing research strongly supports the possibility for there to be significant behavioral crowding out of HWR from strengthening incentives through stronger regulations pertaining to water conservation.

Several months into the field experiment, having established baseline HWR effects, the utility significantly strengthened regulatory enforcement of financial incentives for households to conserve water. Leveraging the AMI data, the utility undertook an innovative approach to enforcing existing mandatory irrigation restriction policies by automating detection of irrigation violations and issuing notice to offenders. These violation notices, discussed in Section 2.1, warned offending households of fines for continuing to violate the pre-existing municipal irrigation policies and made it clear that the enforcement was now fully automated through computer algorithms. As shown in Figure 2.1, within one week this enforcement decision increased the share of households that had ever been warned from 4.6 to 39.2 percent. This sharp change in the policy regime and associated substantial increase in regulation-driven incentives for conserving water affords us a causal test of crowd

[2] These effects of peer comparisons for water conservation are somewhat stronger than those found in the literature on similar interventions related to energy use by Opower, which tend to range from one to three percent [e.g. Allcott, 2011, Ayres et al., 2013, Costa and Kahn, 2013, Allcott and Rogers, 2014].

out of social pressure from the randomized HWR.

Figure 2.1: Time series for issued irrigation violation notices

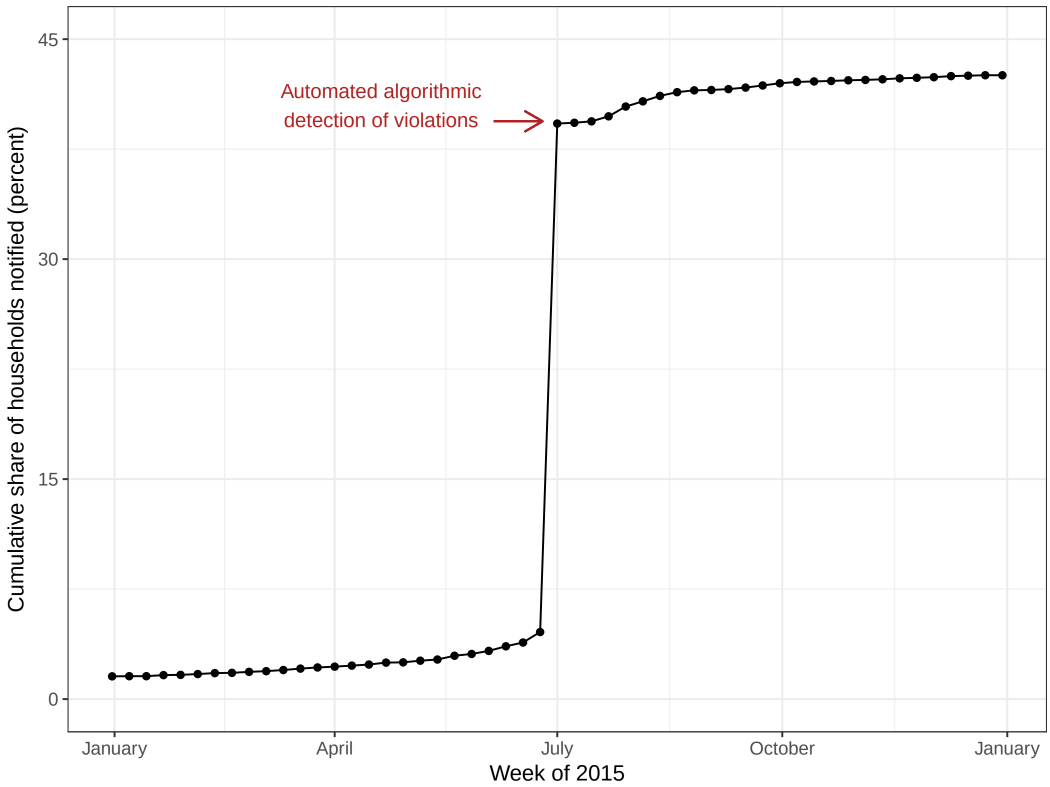


Notes: Figure 2.1 plots the cumulative share of in-sample households that had ever received an irrigation violation notice by week. Throughout this period, violations were determined when either a municipal employee or a neighbor of the offender reported unlawful irrigation to the city. As indicated by the annotation, the city also implemented an automated algorithmic detection of violations in early July.

Using the field experiment to identify intent-to-treat effects in administrative data on residential water consumption among the universe of utility customers, we find that the randomized HWR reduced average household water use by 76 gallons (3.2 percent) per week prior to the intensification of enforcement for the irrigation policy. After the automation of enforcement, we find HWR reduced average water consumption by 75 gallons per week – an economically and quantitatively identical treatment effect. Using difference in (randomized) differences tests to more formally quantify the change in treatment effects between neighboring weeks before and after the strengthened regulatory enforcement, we find no evidence that credibly stronger regulation crowds out the effectiveness of social pressure for water conservation.

We address several potential threats to inference. The key identification assumption in our test for crowd out is that the effect of randomized HWR would have held constant counterfactually, had the water utility not strengthened the regulatory incentives. Prior research shows that the effects of peer comparisons tend to gradually attenuate over time, particularly when the interventions are halted [Ferraro et al., 2011, Allcott and Rogers, 2014, Bernedo et al., 2014, Ito et al., 2018]. This concern seems minimal for our analysis, which focuses only on the initial few months of the intervention, entirely within a single summer water season, during which HWR continued to be sent (bi)monthly to treated households. Another concern is that the enhanced enforcement might crowd out social pressure simply due to limited scope for residential water conservation. Based on institutional details and prior research, it seems unlikely there would be significant crowd out for purely mechanical reasons, as we discuss in Section 2.1 [Castledine et al., 2014, Baker, 2017, Brelsford and Abbott, 2018]. A related consideration is that the HWR might have increased compliance with irrigation policies. All of the above considerations would generate attenuation bias, and if HWR treatment effects did attenuate for any aforementioned reasons, this would bias our estimates towards instead of against finding crowd out.

Given that we find that the effects of social pressure are invariant to the strengthened regulatory enforcement, the most serious threat to inference would be if the incentives from the regulations were non-binding. Akerlof and Dickens [1982] argue that changes to regulative enforcement cause individuals to switch their focus from "self-motivation to obey the law" towards the severity of the punishment. Although the irrigation violation notices explicitly assert they are a precursor to monetary penalties, recipients might have perceived the warnings to be cheap talk by the water utility.[3] To negate concerns that our finding of no crowd out is due

[3] In practice, irrigation restrictions are rarely enforced because of enforcement limitations. State reports show that during the drought we study, most water agencies that restricted irrigation to two days per week never issued a single penalty (California State Water Resources Control Board, www.waterboards.ca.gov).

to the “strengthened” incentive being weak and ineffective, we examine the direct effects of the automated violation notices on water conservation. Using a regression discontinuity design based on a somewhat arbitrary cutoff in the noncompliance algorithm, we find that violation notices cause significant reductions in water consumption: the local average treatment effect is a reduction of 550 gallons (29 percent) per week. In addition, the violation notices further increased policy compliance by shifting some residential water consumption within the week into irrigation-permitted time periods. Together, these findings demonstrate by revealed preference that the change in regulation-driven incentives is economically significant, supporting our interpretation of the experimental estimates that the efficacy of social pressure is not crowded out by strengthening regulatory incentives.

Our study makes several contributions to the literature. Most broadly, we provide insights regarding the interaction of pecuniary and non-pecuniary incentives for resource conservation. The three most closely-related studies are by Ito et al. [2018], who estimate separately and compare the effects of dynamic electricity price increases and those of peer comparisons, Pellerano et al. [2017], who find that layering on information about potentially relevant economic incentives may diminish the impact of normative appeals to conserve energy, and List et al. [2017], who find that adding a financial rewards program to a standard home energy report treatment provides additional conservation.[4] We build on these studies by directly testing whether (enforcement of) regulations crowd out or partially crowd out social pressure.

In addition, our study adds to a growing literature that examines social pressure as a form of policy. Governments are often limited in their use of first-best Pigouvian remedies, especially in energy and water conservation contexts. Various regulative and economic strategies for conserving water have been explored, including use restrictions and price increases; however, Grafton and Ward [2008] show that

[4]Within a broader study of various financial incentives for energy conservation, Gillan [2018] also tests a normative treatment. Given that the study finds insignificant effects of the normative message independent of the price incentive, it is unclear how to interpret the estimated interaction effect of the combined treatments.

mandatory restrictions can be welfare-reducing, and numerous studies find that water demand is fairly inelastic to price changes [Olmstead et al., 2007, Ito, 2013]. Moreover, utilities typically face substantial regulatory and political constraints to adjusting prices or restricting use. In response to these limitations, agencies are increasingly using interventions that operate outside the regulatory framework, and the use of social pressure – such as providing peer comparisons of consumption – is steadily rising among water utilities and in other sectors. Our study also provides some of the first empirical evidence on the impacts of enforceable and substantial financial penalties on water use. It highlights that households are responsive to the threat of large financial penalties, in contrast to a body of work highlighting a minimal response to water prices.

Interventions based around social pressure are particularly attractive due to low implementation costs for the policymaking agency. However, Allcott and Kessler [2019] find that the majority of resource conservation nudge recipients are unwilling to pay the marginal social cost of the nudge, indicating significant economic welfare considerations for social pressure policies. Our evidence of additive conservation benefits from social pressure and regulatory incentives is thus especially relevant for policymakers as they incorporate a broader variety of policies and optimize across an increasingly multidimensional set of interventions.

Finally, our focus on water conservation is particularly relevant as fresh water availability remains one of the most pressing environmental and economic concerns in many regions around the world. The United Nations forecasts that two-thirds of the world's population will live with water stressed conditions by 2025, and that the outlook will only worsen under existing climate change scenarios.[5] At a municipal level, there are also growing concerns about cities sinking as they over-extract local water resources.[6] As weather patterns are becoming more erratic, severe droughts such as the one experienced during the past several years in Cal-

[5] www.un.org/waterforlifedecade/scarcity.shtml.

[6] For example, www.nytimes.com/interactive/2017/02/17/world/americas/mexico-city-sinking.html and www.nytimes.com/interactive/2017/12/21/world/asia/jakarta-sinking-climate.html.

ifornia, other Southwestern states, and many other regions around the world are becoming increasingly common.

## 2.1 Empirical setting and research design

California has a history of extreme and persistent variation in precipitation, with the Coastal Southern California hydrological region suffering multiple periods of extended drought conditions during the past decade (see Appendix Figure A.12). The vast majority of annual precipitation falls during the winter months throughout the state and in Coastal Southern California in particular. With a dry 2011-2012 winter followed by an even drier 2012-2013 winter, authorities began implementing a wide range of conservation measures. This paper focuses on two of the most common interventions, which succinctly capture two major pillars of residential water policy: social pressure to encourage voluntary conservation and regulatory incentives to reduce consumption. The first intervention we study – normative peer comparisons as social pressure for voluntary conservation – is relatively contemporary in origin. The second intervention we evaluate – day-of-week and time-of-day outdoor water use restrictions (hereafter, DOWR) – has a long heritage in water conservation, although automating the associated enforcement is a first of its kind.

To study the effects of normative peer comparisons, we use data from WaterSmart Software and Burbank Water and Power (BWP) on a field experiment conducted.[7] Nearly 17,000 single-family households in Burbank, California are included in the randomized control experiment we study. Treated households were provided with WaterSmart HWR by mail or email, with the timing of initial treatment rolled out over the monthly water billing cycle during late April through mid May 2015. Notably, households in the "Experimental" treatment arm began to receive HWR at

[7] Operating in a parallel industry to Opower/Oracle Utilities for energy utilities, WaterSmart Software provides assistance to water utilities in California and around the world through analyzing and interpreting data on water use and through providing information to customers. The Burbank experiment was conducted partly in partnership with the Center for Water and Energy Efficiency at UC Davis [Jessoe et al., 2019].

least six weeks prior to the heightened enforcement of irrigation restrictions in early July. Very few treated households chose to opt-out, and nearly all continued to receive (bi)monthly HWR for at least several months after their initial treatment and throughout the time period we include in our empirical analyses.

A HWR has a few components (see Appendix Figure A.1 for an example report), with the core component being the normative comparison of the treated household's water consumption with that of a peer group formed from neighboring households with the same number of occupants and similar irrigable area. A household that used more water than an "average" comparison household would receive a frowning face on a red water drop and be informed that it "used more water than most of [its] neighbors." If a household used less water than an "efficient" household it would receive a smiling green water drop. Yellow drops were given to households using between the efficient level and the median level of water use.[8] In all three cases, households were provided with their water use in gallons per day (GPD), the median water use in GPD, and the efficient level of water use in GPD for comparison "neighbor" households.

The HWR includes some additional technical or financial information in text boxes on the report. For example, the HWR includes information on three water saving tips "buttons" at the bottom of the report (e.g. "Choose low water-use plants," "Upgrade to a low-flow toilet," "Install high-efficiency shower heads," "Fill bath 1/3 of the way," "Reduce shower to 5 minutes," and "Install faucet aerators"). Ferraro et al. [2011] experimentally test the effectiveness of this kind of information alone in HWR and find it largely ineffective in the absence of the normative peer comparison. Per Jessoe et al. [2019], some of the Burbank HWR also included a text box on the report with one of three randomly generated messages, some of which pertain

[8]The HWR does not clarify the specific thresholds, but the "efficient" neighbor benchmark is based on the 20th percentile of within-neighbor-group consumption, and the "average" neighbor benchmark is actually the 55th percentile of within-neighbor-group consumption, ensuring that most homes are considered better than "average."

more specifically to hot water use.[9] We verified that our estimated coefficients of interest are the same for the pooled three arms as for the subset receiving only the basic core HWR messaging. Another experiment with a different water utility and WaterSmart randomly varied the visual cue with the social comparison information on relative water use. Bhanot [Forthcoming] finds that most of the conservation effect of the HWR treatment comes from the basic social comparison without any visual cue, but that there is an additional effect from an injunctive message added into the visual cue (e.g. a frowning, neutral or smiling face icon inside the water droplet). We cannot identify the separate effects of each piece of information on the HWR in our data, and thus focus on the combined effect of the social comparison and other information in the HWR intervention, which is the relevant "treatment" from the policymaker's perspective.

The broader literature also supports that the normative social comparison is by far the most effective component of HWR, including research experimentally testing WaterSmart HWR nearly identical to the treatments examined in our study [Ferraro and Price, 2013, Mitchell and Chesnutt, 2013, Allcott and Rogers, 2014, Brent et al., 2015, 2017]. Moreover, as these "social comparisons impose a moral cost on consumption," we join the existing literature in interpreting the HWR treatment primarily as a form of social pressure for voluntary resource conservation [Brent et al., 2017].

In contrast to the contemporary novelty of HWR, DOWR are a common and long-standing conservation policy for water utilities. Although the benefits from water conservation do not vary across hours of the week, there are several institutional, horticultural, and behavioral reasons to impose restrictions by time of day and day of the week. First, DOWR generally prohibit irrigation between a few hours after

[9] The three messages are: i) "Surprised by your WaterScore? Your WaterScore compares your use to others in Burbank who also have X occupants and a similar yard size. Is your household different? Log on to tune your profile and see adjusted comparisons.", ii) "Reduce hot water use: Did you know that heating water is the second most energy intensive activity in your home? Log on for information and offers for the water, energy, and money saving actions below!", or iii) "Save hot water, win big! Reduce water use by 24 percent and gas use by 3 percent in the next 7 months and win one of: a) 25 high-efficiency Whirlpool clothes washers, b) 100 luxurious, efficient Evolve showerheads, or c) A hot water efficiency starter. conserveandwin.com, Use promo code "CONSERVE."'

sunrise and a few hours before sunset, which minimizes water lost to evaporation [Christiansen, 1942]. Furthermore, spacing out the days on which irrigation is allowed ensures water can be spread efficiently for the benefit of the plants. In addition, there are enforcement justifications for prohibiting outdoor uses on certain days. Before the introduction of AMI, "smart meters" which record high-frequency consumption data, the only method of detecting irrigation violators was visual inspection; in other words, either a water agency employee or an informant neighbor of the violator must visually observe the illegal irrigation. Prohibiting outdoor water use on specific days of the week – rather than, say, restricting total irrigation volume – facilitates enforcement, since an extensive margin of outdoor use on a given day is much more easily observed. Finally, outdoor water use comprises a large share of residential water use and provides the potential to conserve water without generating negative health or safety consequences.[10]

During the drought, BWP joined many other water utilities in implementing irrigation restrictions. On May 14, 2015, the Burbank City Council approved the implementation of stricter Stage 3 restrictions, which included limiting outdoor water use to Tuesdays and Saturdays. Notably, BWP initially enforced these DOWR using only the traditional method of visual inspection, resulting in only a very small number of households determined to be in violation. As seasonal dryness accumulated, the violation count increased slightly, but enforcement remained extremely low through the end of June. Then, after six weeks of the post-treatment period in our field experiment, BWP additionally undertook the unprecedented approach of using AMI data to algorithmically detect potential violators. BWP automatically sent more than one-third of single-family residential accounts in the district a violation notice during the first week of July 2015, as shown in Figure 2.1. These violation notices (shown in Appendix Figure A.13) warned that customers would be fined $100 if they continued to irrigate more than two days a week. Fines increased to $200 for a second violation and $500 for a third violation. The notices also clearly indicate that detection of violations was now algorithmic in nature.

[10]cf. California Department of Water Resources' *California Water Plan Update* (2013), Volume 3, Chapter 3: Urban Water Use Efficiency.

As the remainder of the content included in the violation notices was already being publicized widely throughout the community, these factors afford this treatment an interpretation as being a shock to household beliefs about detection and enforcement probabilities for a pecuniary threat. Because of potential spillovers across households, both directly and through media coverage of the heightened enforcement in outlets such as the *Los Angeles Times*, we interpret the treatment broadly as a substantial strengthening of regulation-driven incentives for conservation.[11]

Both of these policies were introduced into a landscape filled with media coverage and appeals from the State to conserve water, and BWP was facing threats of State-mandated pecuniary penalties tied to conservation goals.[12] Furthermore, recent evidence suggests that intense media coverage can directly influence residential water consumption [Quesnel and Ajami, 2017]. As such, one potential concern for our study is that the strengthened regulatory incentives could crowd out the effects of social pressure for purely mechanical reasons of limited scope to further reduce water consumption. We do not view this as a significant concern for several reasons. For one, the policies we study targeted primarily outdoor water use, for which there is ample scope for additional conservation [Castledine et al., 2014, Baker, 2017, Brelsford and Abbott, 2018]. In addition, the households treated by the automated violation notices are (unsurprisingly) almost exclusively included in the "above average" WaterSmart tier. As these tiers are defined within household size and irrigable area, above average consumption primarily reflects choices of the household rather than need. In any case, mechanical crowd out for reasons of scope would attenuate estimates of the HWR treatment effect and bias our estimates towards finding crowd out, whereas we find no evidence of crowd out.

Our research designs allow us to examine the impacts of social pressure and regulatory incentives within a given period across randomly and quasi-randomly as-

[11] For example, www.latimes.com/tn-blr-bwp-to-step-up-water-saving-efforts-20150705-story.html.

[12] California passed and ultimately enforced regulations providing for fines of $10,000 per day for water agencies that did not generate sufficient mandated conservation, as noted by local media such as www.mercurynews.com/2015/04/28/water-wasting-fines-of-10000-proposed-by-gov-jerry-brown/.

signed groups. One consideration for this particular setting is the external validity of conclusions drawn at the height of severe drought in a drought-prone region. However, this utility is representative of water utilities in Southern California, and studying policy during such an event is critical for understanding the effects of related policies in the contexts in which they are invoked.

## 2.2 Data and balance checks

Our study leverages the fairly recent introduction of AMI for residential water service, which records high-frequency data on water consumption. Unlike with smart meters for electricity, AMI measurement is relatively rare for household water use, although adoption of the technology exhibits a steep upward time trend. In 2015, only about seven million smart meters for water had been installed in the United States, compared to about 68 million smart electricity meters.[13] Burbank Water and Power (BWP), had installed smart water meters for eligible households throughout their service area approximately one year prior to the field experiment we evaluate.

Hourly AMI data offer several significant advantages over traditional monthly water consumption data. For the researcher, the availability of hourly consumption data avoids the measurement error that is typically present when trying to map metered water use to the actual timing of consumption. In addition, AMI data enable us to analyze not only broad consumption patterns but also the distribution of water consumption within a week; a portion of our empirical study takes advantage of this hourly disaggregation to identify patterns of within-week intertemporal substitution. For the utility, one benefit of AMI is the possibility for algorithmic detection of water leaks or, less commonly, irrigation. As single-family residential accounts typically do not have separate meters for landscape irrigation, utilities are generally unable to identify irrigation disaggregated from total household water use. However, the flow rate of irrigation controllers is so high that consumption

[13] cf. www.westmonroepartners.com/Insights/White-Papers/State-of-Advanced-Metering-Infrastructure

during an hour with irrigation far exceeds regular household consumption during any hour of the week without irrigation.[14] Thus, AMI technology allows for automating the enforcement of irrigation policies, and, as water utilities increasingly install smart meters throughout their jurisdictions, the scope for applying AMI technology is steadily growing.

In conducting the field experiment, WaterSmart excluded commercial, industrial, and multi-family residential accounts, leaving an experimental sample of 17,289 single-family residential water accounts. For our analyses, we drop an additional 653 households with missing water consumption data during the pre-treatment period or during the analysis period of May-October 2015.[15] Our final analysis sample includes 16,636 households in Burbank, California. WaterSmart intentionally weights their experiments to have significantly more treated participants, so 13,717 (82 percent of the total) of these households are assigned to the "Experimental" treatment arm, while 2,919 households form the untreated control.

Summary statistics and experimental balance tests for these households are presented in Table 2.1. The top panel shows that provision of the HWR is virtually 100 percent among households assigned to treatment; no households in the control group were treated by WaterSmart. Furthermore, the use of administrative data on water consumption for all households rules out concerns about differential attrition or loss in follow-up.

[14] cf. www.wsscwater.com/customer-service/rates/water-usage.html.

[15] We verified that these sample restrictions are orthogonal to the randomization we use for identification.

Table 2.1: Summary statistics and randomization balance checks

| | Group means | | t-tests | |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| Covariate | Control | Experimental | Difference | p-value |
| Number of households | 2919 | 13,717 | | |
| Sent WaterSmart HWR | 0 | 0.9972 | | |
| Prior water violation | 0.0483 | 0.0452 | -0.0031 | 0.48 |
| Lot size (SqFt) | 7347 | 7320 | -27 | 0.71 |
| Irrigable area (SqFt) | 3829 | 3794 | -35 | 0.45 |
| House size (SqFt) | 1619 | 1620 | 1 | 0.95 |
| Year built | 1945 | 1945 | 0 | 0.82 |
| Number of floors | 1.062 | 1.067 | 0.005 | 0.28 |
| Number of bedrooms | 2.912 | 2.919 | 0.007 | 0.69 |
| Number of bathrooms | 1.93 | 1.939 | 0.009 | 0.61 |
| Weekly water gallons | 2693 | 2678 | -15 | 0.66 |

Notes: Table 2.1 shows statistics by WaterSmart Home Water Reports (HWR) treatment arm for household-level covariates. The first two columns show means by treatment arm for all households in the randomization sample, Column (3) shows the difference in means, and Column (4) shows the p-values for t-tests of whether the difference in group means is significantly different from zero. By design, WaterSmart weighted the randomization to have about 82 percent of the sample in the “Experimental” treatment arm. Initial HWR were sent to treated households during the billing cycle ending in mid May 2015. All outcomes in the lower panel are determined prior to the randomization and prior to the automated pecuniary treatment. For pre-treatment weekly water consumption, we use each household’s average weekly gallons consumed during May 2014 through April 2015, spanning one full year prior to both treatments. Figure 2.2 shows the density distributions of households’ average weekly pre-treatment consumption by treatment arm.

The bottom panel of Table 2.1 shows averages and balance t-tests for time-invariant and pre-treatment household covariates. The administrative data used in our study include several residential attributes, showing overall a high degree of balance across groups in the characteristics of their homes and landscapes. Households across groups are also equally likely to have received prior irrigation violation notices. Most importantly, the groups are balanced in their water consumption during the year preceding our study (May 2014 - April 2015). Figure 2.2 confirms that the distribution – not just the average level – of pre-treatment water consumption is very similar across the two groups. Overall, as we should expect given randomization among a large set of households, the two arms are very balanced.

Figure 2.2: Balance test of distributions of pre-treatment water consumption

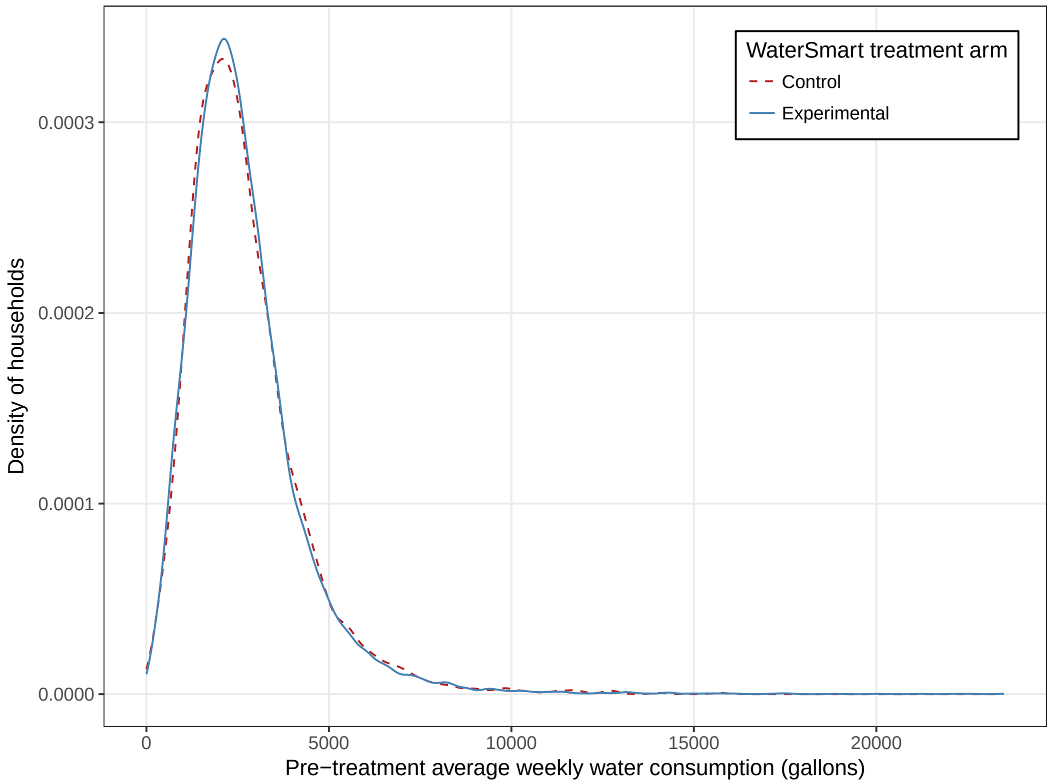


Notes: Figure 2.2 plots the distributions of average weekly water consumption for in-sample households during May 2014 through April 2015, the full year prior to both the social comparison and automated enforcement treatments. The solid line shows the distribution for households assigned to the WaterSmart Home Water Reports (HWR) treatment arm, and the dashed line includes the untreated control group.

## 2.3 Results

This section presents our empirical findings. First, we assess the impact of social pressure on water conservation using the randomized WaterSmart field experiment. We estimate treatment effects of the normative peer comparisons over the full sample period and then evaluate crowd out by comparing estimates between a baseline period and a period with strengthened enforcement of irrigation restrictions and associated automated notice of violations. Second, to validate that the heightened regulatory incentives are binding, we estimate the direct effects of the violation notices using a regression discontinuity design based on the computer algorithm used to automate enforcement.

### 2.3.1 Estimated effects of the randomized social pressure

We start by examining whether social pressure affects water conservation in both the absence and presence of stronger regulatory incentives. Our identification uses a field experiment in which randomly-selected households were provided HWR including normative social comparisons of water use.

Figure 2.3 displays average post-treatment water consumption by week for both the treatment group that received social comparisons and the control group that did not. The time range shown in the figure spans from late May through December 2015; treated households each had been sent exactly one HWR at the start of this time period and then continued to receive (bi)monthly reports throughout. The AMI data enable us to see treatment effects immediately following the initial HWR, and weekly water use is lower for the treatment group than the control group for every week over the experimental period.[16] The magnitude of treatment-control differences also appears to be fairly consistent across weeks. Visually, the evidence strongly supports that providing social comparisons to households causes substan-

[16]The immediacy of the treatment effect is consistent with Reiss and White [2008], who find that consumers in the electricity sector respond promptly to both price changes and normative appeals to conserve. These findings support our use of May-June consumption as a counterfactual for the magnitude of the effect of WaterSmart HWR in subsequent months.

tial water conservation.

Figure 2.3: Post-treatment water consumption by treatment arm

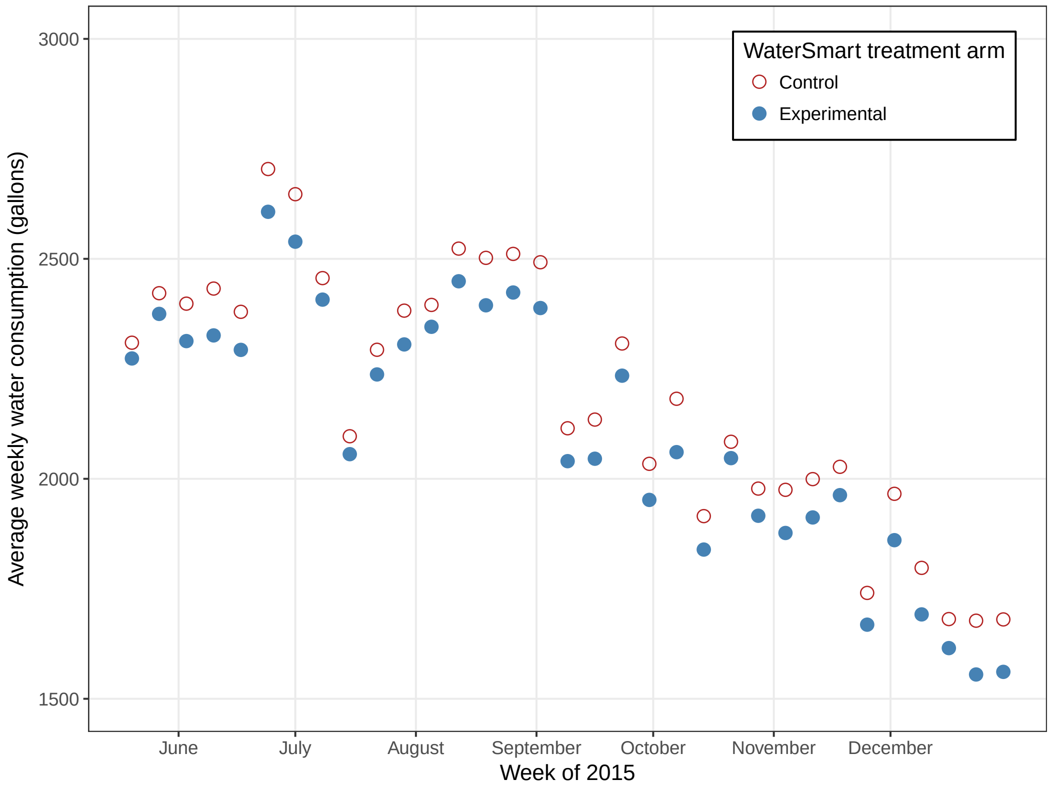


Notes: Figure 2.3 plots average water consumption by week for each WaterSmart treatment arm during late May through December 2015. "Experimental" households each had been sent one Home Water Report (HWR) as of the start of this time period and (bi)monthly reports continued to be sent to treated households throughout this time period.

Of importance for exploring the crowd out hypothesis, there is no discernible break over time in the magnitude of treatment-control differences in average weekly water use, despite the substantial increase in irrigation-restriction enforcement discussed earlier and shown in Figure 2.1. The post-treatment sample period can be divided into three ranges. During the six weeks from late May through June, the statutory summer watering season under irrigation restrictions was in effect, but it was prior to automation of the associated regulatory enforcement. Then, during July through October the statutory summer watering season under irrigation restrictions remained in effect, and automated violation notices were issued. Finally, Figure 2.3 displays water use through November and December, part of the statutory and technical winter water season, when precipitation seasonally increases and legal irrigation was further restricted to only on Saturdays.[17] Visually examining the treatment-control differences across these three periods, there does not appear to be any clear break in the magnitude of the treatment effect. In fact, there does not appear to be a break in the magnitude of treatment-control differences at any point over the entire sample period.

We investigate these patterns of water use more formally by estimating several regression specifications. Because of the significant change in seasonal rainfall patterns in California beginning in November (i.e. the winter rainy season), we focus our quantitative analyses on the summer watering season from late May to October. The starting regression equation is straightforward in the context of the random experiment:

$$\text{consumption}_{it} = \beta' X_i + \gamma \cdot \mathrm{I}\{\mathrm{HWR}\}_i + u_{it} \tag{2.1}$$

In Equation (2.1), $\text{consumption}_{it}$ is weekly water use for household $i$, $X_i$ is a vector of baseline controls, $\mathrm{I}\{\mathrm{HWR}\}_i$ is the randomly assigned WaterSmart Home Water Reports treatment indicator, and $u_{it}$ is the econometric error term. The baseline controls include administrative data on residential lot size, irrigable area, and the home's square footage, year of construction, number of floors, number of bedrooms, and number of bathrooms. The effect of receiving a HWR – the "intent-to-treat" estimate of social pressure – is captured by $\gamma$, which can be estimated separately under different regimes of regulatory incentives. All specifications are estimated using OLS and robust standard errors are reported, two-way clustered by household and week.[18]

Table 2.2 reports regression estimates for the effects of randomized HWR on post-treatment water consumption. Each column presents an estimate of the average intent-to-treat effect of the HWR for weekly water consumption during the months

[17] Following plans to the tariff structure set years in advance, water tariffs changed once near the beginning of our study period on June 2, 2015. The price change was a relatively small increase of 5.2 cents per hundred cubic feet (about 748 gallons) for the first consumption tier, with slightly larger increases on higher tiers of consumption. The median May water bill of 8,550 gallons would have been increased by only \$1.59, inclusive of a \$1.00 increase to the fixed service charge. There were no additional changes to tariffs until July 2016, over a year later and after the end of our study period.

[18] As our study period includes a fairly small number of weeks, we verified that standard errors are very similar when clustering only by household.

of 2015 indicated by the column titles. We first discuss the results for the late May-June period, which is covered by summer irrigation restrictions but prior to automated irrigation restriction enforcement. Columns (1) and (2) report estimates without controls and with controls, respectively. As expected with randomization of households into treatment, the inclusion of controls has no effect on the point estimate but does slightly improve precision of the estimate. Prior to heightened regulatory incentives, we find an intent-to-treat estimate from the social pressure treatment of 76 gallons per week reduction in water use per treated household. The reduction in water consumption represents 3.2 percent of average weekly water consumption. These findings closely align with those in the existing literature on HWR, both for WaterSmart in Burbank specifically [Jessoe et al., 2019] and in other jurisdictions more generally [Ferraro et al., 2011, Ferraro and Price, 2013, Mitchell and Chesnutt, 2013, Bernedo et al., 2014, Brent et al., 2015, 2017, Bhanot, Forthcoming].

Table 2.2: Estimated effects of randomized WaterSmart Home Water Reports

| | Weekly water consumption (gallons) | | | | | |
|---|---|---|---|---|---|---|
| | Late May - June | | July - October | | Late May - October | |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| I{HWR} | $-76.36^{**}$ | $-76.24^{***}$ | $-75.51^{***}$ | $-74.95^{***}$ | $-76.14^{***}$ | |
| | (31.23) | (29.00) | (28.25) | (25.55) | (28.71) | |
| I{Post July} | | | | | $-160.61^{**}$ | $-164.47^{**}$ |
| | | | | | (73.35) | (76.16) |
| I{HWR} X I{Post July} | | | | | 1.15 | 3.11 |
| | | | | | (14.42) | (19.88) |
| HH controls | No | Yes | No | Yes | Yes | — |
| HH FEs | No | No | No | No | No | Yes |
| Num. of HHs | 16,636 | 16,636 | 16,636 | 16,636 | 16,636 | 16,636 |
| Observations | 99,575 | 99,575 | 292,542 | 292,542 | 392,117 | 392,117 |

$^{*}$p<0.1; $^{**}$p<0.05; $^{***}$p<0.01 Columns (1) - (4) present estimates of the average intent-to-treat effect of the randomized WaterSmart HWR for weekly water consumption during 2015 for the month ranges indicated by the column titles. "Experimental" households each had been sent one HWR as of the start of this time period, and monthly reports continued to be sent throughout this time period. Automated irrigation violation notices were sent during the first week of July. The legal and technical local summer water season runs through the end of October. Columns (5) and (6) present difference in differences estimates testing whether the effect of the randomized HWR differed before versus after violation notices were sent. Standard errors in parentheses are two-way clustered by household and week. The household control terms include residential lot size, irrigable area, and the home's square footage, year of construction, number of floors, number of bedrooms, and number of bathrooms. Figure 2.3 shows average weekly water consumption by treatment group separately for each post-treatment week spanning from late May through December 2015.

Of much more novelty, we now turn to the results for the July-October period, which represents the regime when incentives for reduced irrigation were heightened and the remainder of the statutory summer watering season under irrigation restrictions.[19] From what we gather based on newspaper articles and institutional sources, there is evidence of abundant general public awareness of the monetary fines and associated automated enforcement for wrong day-of-week irrigation. Columns (3) and (4) report estimates without controls and with controls, respectively. Again, the inclusion of controls has no effect on the point estimate but does slightly improve precision. During this period, we find an estimate from the social pressure treatment of a 75 gallons per week reduction in water use per treated household.

The comparison of treatment estimates from before to after heightened regulatory enforcement reveals that they are virtually identical, suggesting no crowd out of treatment effects by the regulatory incentives. To formally investigate this comparison, we build on Equation (2.1) by including the full late May-October sample period and estimating the following equation:

$$\text{consumption}_{it} = \beta' X_i + \gamma \cdot \text{I}\{\text{HWR}\}_i + \lambda \cdot \text{I}\{\text{Post July}\}_t + \delta \cdot \text{I}\{\text{HWR}\}_i \cdot \text{I}\{\text{Post July}\}_t + u_{it} \tag{2.2}$$

In Equation (2.2), $\text{I}\{\text{Post July}\}_t$ is an indicator variable that equals zero during May-June and one in the post-automated enforcement period (i.e. July-October). In this specification, $\delta$ captures the difference between the two periods' average treatment effects and is the coefficient of interest because it provides a direct estimate of the crowd out effect. This empirical test is in the spirit of a difference in differences estimator, with one of the differences determined by a randomly assigned treatment. The other difference is a comparison across closely neighboring weeks. The identifying assumption for this test is thus much less stringent than that for typical difference in differences strategies based on heterogeneous policy changes across treatment units and time periods (often years).

[19] As discussed above, the irrigation policy itself did not change at this time, only the enforcement of the pre-existing restrictions.

Column (5) of Table 2.2 reports estimates of Equation (2.2). As expected given the previous results, we find a point estimate of crowd out that is essentially zero. The standard error is small enough to also rule out even moderate amounts of crowd out: under the null hypothesis of positive crowd out, at conventional significance levels we can rule out anything larger than 25 gallons per week, one third of the average treatment effect of the social pressure. Finally, in Column (6) of Table 2.2 we estimate a specification that includes household fixed effects. The inclusion of these fixed effects – which control for all unobserved time-invariant differences across households – does not change the results. The estimate of $\delta$ is very similar to that reported in Column (5). From this model that absorbs even more heterogeneity, we continue to find no evidence of crowd out of social pressure by the heightened regulatory incentives.

The interpretation of these findings is made stronger given that several factors potentially bias our estimates towards finding crowd out. As discussed previously, prior research shows that the effects of peer comparisons tend to attenuate over time, although generally over a much longer time span than that included in our analysis. As discussed in Section 2.1, there is also some minor possibility for mechanical crowd out due to limited scope for additional conservation. Finally, the intensity of the direct incentive from stronger regulatory enforcement might be slightly stronger for the HWR control group. Initial HWR were sent more than a month prior to the automated violation notices and they induced conservation among treated households. Because a household's total water consumption factors into the algorithm for detecting irrigation violations, HWR-treated households are about two percentage points less likely to have been sent an automated irrigation violation notice than those in the control group (36.9 vs. 39.0 percent). This minor difference in relative treatment intensity across groups does not affect causal inference about the randomized HWR treatment, but it does potentially bias the difference in differences estimates towards showing crowd out.[20] In light of these factors, the weight of empirical evidence is strengthened against there being crowd out of social pressure.

We conduct several additional robustness checks of these findings in Table 2.3, which presents estimates for six variants of Equation (2.2). Columns (1) and (2) reproduce the difference in differences tests from Columns (5) and (6) of Table 2.2. In Columns (3) and (4) we add fixed effects for the week of the sample, more flexibly controlling for any heterogeneity in consumption across time periods. Again, we find no evidence of crowd out. In Columns (5) and (6) of Table 2.3, we narrow the analysis window to include only late May through August, more tightly surrounding the July policy change. Using this shorter post-treatment time period also yields no evidence of crowd out.

[20]The degree of bias from this heterogeneity in intensity of regulatory incentives across HWR arms is well within our estimated confidence intervals; if anything, at face value this implies our estimated null for $\delta$ in Equation (2.2) supports there is some crowd-*in*. A more nuanced consideration would be if the HWR treatment effect were concentrated among a very small share of households, but this conjecture runs counter to the entire body of evidence in the literature on home energy reports and HWR, including our own findings.

Table 2.3: Robustness checks of difference in differences test for crowd out

| | Weekly water consumption (gallons) | | | | | |
|---|---|---|---|---|---|---|
| | Late May - October | | | | Late May - August | |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| I{HWR} | $-76.14^{***}$ | | $-76.15^{***}$ | | $-76.26^{***}$ | |
| | (28.71) | | (28.73) | | (28.76) | |
| I{Post July} | $-160.61^{**}$ | $-164.47^{**}$ | | | | |
| | (73.35) | (76.16) | | | | |
| I{HWR} X | 1.15 | 3.11 | 0.68 | 2.28 | 4.44 | 6.34 |
| I{Post July} | (14.42) | (19.88) | (14.49) | (19.95) | (12.94) | (20.58) |
| HH controls | Yes | — | Yes | — | Yes | — |
| HH FEs | No | Yes | No | Yes | No | Yes |
| Week FEs | No | No | Yes | Yes | Yes | Yes |
| Num. of HHs | 16,636 | 16,636 | 16,636 | 16,636 | 16,636 | 16,636 |
| Observations | 392,117 | 392,117 | 392,117 | 392,117 | 247,463 | 247,463 |

$^{*}$p<0.1; $^{**}$p<0.05; $^{***}$p<0.01 All columns present difference in differences estimates testing whether the effect of the randomized Home Water Reports (HWR) differed before versus after violation notices were sent. Standard errors in parentheses are two-way clustered by household and week. Columns (1) and (2) reproduce the estimates from Columns (5) and (6) of Table 2.2. Columns (3) - (6) add in fixed effects for the week of sample. Columns (5) and (6) further restrict the time period to May - August, 2015. The household control terms include residential lot size, irrigable area, and the home's square footage, year of construction, number of floors, number of bedrooms, and number of bathrooms. Figure 2.3 shows average weekly water consumption by treatment group separately for each post-treatment week spanning from late May through December 2015.

In Table 2.4, we explore heterogeneity in the treatment effect and assess any crowd out. Again, Columns (1) and (2) reproduce the difference in differences tests from Columns (5) and (6) of Table 2.2. Columns (3) and (4) and Columns (5) and (6) estimate these specifications for the subsets of households whose pre-treatment consumption was, respectively, below or above the median. Perhaps not surprisingly, we find that low-volume water users show smaller and statistically insignificant effects of HWR, whereas large-volume consumers exhibit large and significant conservation effects. Neither subset shows any evidence of crowding out between the two interventions. Especially when considering that nearly two-thirds of the high-volume group was treated with violation notices, the estimates in Table 2.4 serve as a further strong robustness check of the primary finding of our study. In sum, these estimates show consistent water conservation effects from social pressure. The effectiveness of social pressure does not change with respect to the strength of regulatory incentives targeting the same behaviors.

Table 2.4: Heterogeneity in difference in differences test for crowd out

| | Dependent variable: weekly water consumption (gallons) | | | | | |
|---|---|---|---|---|---|---|
| | Full sample of households | | Low volume consumers | | High volume consumers | |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| I{HWR} | $-76.14^{***}$ | | $-33.72$ | | $-116.89^{***}$ | |
| | (28.71) | | (21.11) | | (42.03) | |
| I{Post July} | $-160.61^{**}$ | $-164.47^{**}$ | $-42.84$ | $-44.36$ | $-275.52^{***}$ | $-284.16^{***}$ |
| | (73.35) | (76.16) | (53.58) | (55.58) | (95.25) | (99.62) |
| I{HWR} X | 1.15 | 3.11 | 5.13 | 5.13 | $-5.12$ | 0.49 |
| I{Post July} | (14.42) | (19.88) | (14.06) | (17.18) | (26.58) | (33.01) |
| HH controls | Yes | — | Yes | — | Yes | — |
| HH FEs | No | Yes | No | Yes | No | Yes |
| Num. of HHs | 16,636 | 16,636 | 8,317 | 8,317 | 8,319 | 8,319 |
| Sent violation | 37.29% | 37.29% | 12.77% | 12.77% | 61.80% | 61.80% |
| Observations | 392,117 | 392,117 | 196,103 | 196,103 | 196,014 | 196,014 |

$^{*}$p<0.1; $^{**}$p<0.05; $^{***}$p<0.01 All columns present difference in differences estimates testing whether the effect of the randomized Home Water Reports (HWR) differed before versus after violation notices were sent. Standard errors in parentheses are two-way clustered by household and week. Columns (1) and (2) reproduce the estimates from Columns (5) and (6) of Table 2.2, including all in-sample households. Columns (3) and (4) include only households with below-median pre-treatment water consumption. Columns (5) and (6) include only households with above-median pre-treatment water consumption. The household control terms include residential lot size, irrigable area, and the home's square footage, year of construction, number of floors, number of bedrooms, and number of bathrooms.

### 2.3.2 Regression discontinuity estimates for the regulation

Given that we find that the effects of social comparisons are invariant to the layering of stronger regulations, the most serious potential threat to inference would be if the "stronger" regulatory enforcement were non-binding. To negate concerns that our finding of no crowd out is attributable to the heightened enforcement being weak and ineffective, we next examine the direct revealed preference effects of the automated violation notices on water conservation. We find substantial conservation effects of the enhanced enforcement but do not attempt to disentangle the mechanism for the effect of the automated enforcement, such as whether it is the pecuniary incentive of threat of fines or the increased perception of regulatory oversight by the utility and city. We can only test the reduced-form net effects of a policy as it was employed and as it could readily be employed by other jurisdictions.

In doing so, we employ a regression discontinuity design based on a cutoff in the noncompliance algorithm the water agency initially used to determine irrigation violations (discussed above in Section 2.1). We include all households in this analysis, both those assigned to the HWR treatment and those assigned to the HWR control group.[21] The RDD is facilitated as follows. During a single week of June 2015, BWP conservatively estimated the number of days per week that each household was irrigating by counting the number of days on which the household consumed more than 125 gallons in any individual hour of the day. This arbitrary cutoff of 125 gallons for the daily peak consumption hour forms the basis of our regression discontinuity design. For example, a household would be assigned to be sent an automated notice if their seven daily peak hours were $\{200, 200, 125, 100, 100, 100, 100\}$. A household would *not* be assigned to be sent an automated notice if their seven daily peak hours were $\{200, 200, 124, 100, 100, 100, 100\}$. Thus, the third highest daily peak hour during that specific week of June forms the running variable in the

[21] In principle, we could test for crowd out by evaluating the difference in discontinuities across the HWR Treatment and Control arms. In practice, such an exercise is statically underpowered as there are too few households in the neighborhood of the RD cutoff (particularly among the much smaller Control group).

RDD.

For internal institutional reasons, in automating the violation notices the agency ultimately decided to allow for comparatively more detected irrigation days per week for (both HWR Treatment and Control) accounts with "average" or "efficient" consumption, per the WaterSmart tier categorizations discussed above in Section 2.1. Household HWR arm assignment is not a factor in the algorithm and households in both the HWR and Control arm were sent violation notices. However, because we imperfectly observe households' historical tier statuses (particularly among the WaterSmart control group), we assign all households to the running variable based on the third-highest peak consumption hour of the day that would have determined a violation under the strictest allowance. For this reason, the RDD is "fuzzy" and treated households nearly-exclusively fall into the "above average" water consumption tier.[22]

Before turning to the estimates, we conduct some standard exercises to support the validity of our RDD. First, we test for manipulation along the running variable, which measures the distance to the irrigation violation cutoff. Given that the automation of detecting irrigation violations was unprecedented and unannounced, a priori there should be no concern. As shown in Figure A.14, there is some measurement lumpiness from the underlying meter technology, but there is no evidence of any sorting of households around the threshold, which visually confirms the results of our statistical implementation of McCrary's (2008) test for manipulation. Further supporting the identification strategy, Figure A.15 demonstrates that there is also smoothness across the threshold in pre-treatment water consumption along the running variable. Finally, in Appendix Table A.3 we replicate the balance tests from Table 2.1 for only the subset of residences within an 80 gallon bandwidth of the RD cutoff to show that the RDD is not disproportionately capturing behavior of only one of the randomly-assigned HWR treatment arms. On the whole, the RDD appears to be very strongly supported and well-positioned to provide credi-

[22]While the vast majority of households detected using this scheme were likely violating the city's restrictions on outdoor water use, a small fraction of residences that received notices claimed to be using the water for other purposes or hand-watering, which was permitted under the policy.

ble causal inference.

Next, we check compliance. Figure 2.4 displays the share of households receiving initial automated violation notices along the running variable. For visual clarity, the running variable uses bins of 10 gallons in Figures 2.4-2.5, and the size of the markers corresponds to the number of households included in the local average. As evidenced graphically, zero households below the threshold received a violation notice. At the treatment threshold, we find a discontinuous jump of roughly 25 percentage points for receipt of an automated violation notice in the first week of July, a strong first stage.

Figure 2.4: First stage for automated violation notices in the regression discontinuity design

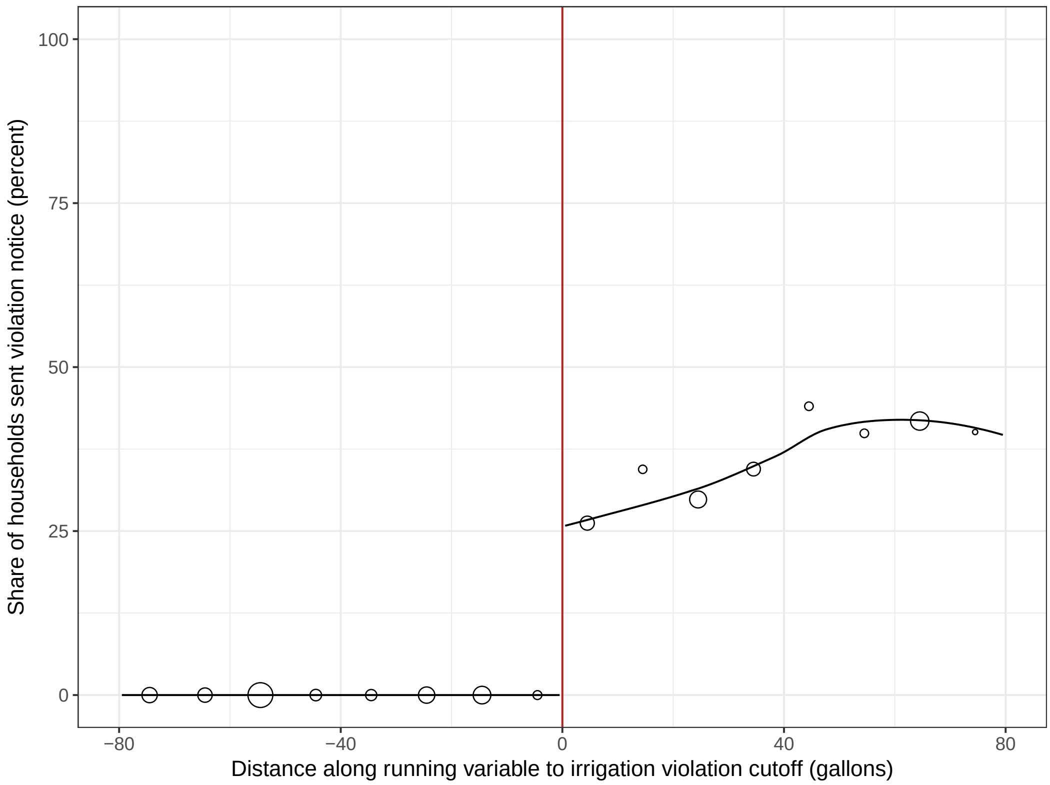


Notes: Figure 2.4 plots local averages for the first stage outcome of whether a household received an automated irrigation violation notice during the first week of July 2015. For clarity, the running variable uses 10 gallon bins. The size of the markers corresponds to the number of households included in the local averages. The LOESS curves shown are fit to the underlying microdata separately on each side of the threshold. Because the running variable represents a necessary but not sufficient condition for a household to be sent an automated violation notice, the first stage supports a "fuzzy" regression discontinuity design.

In other words, households with peak hourly water consumption of amounts just above the 125 gallon threshold are 25 percentage points more likely to have received a violation notice from the water utility than households just below the threshold. In addition, because households that have higher peak water consumption are more likely to be "above average" and thus treated using the stricter threshold, the treatment propensity increases with the running variable. At the right end of the range displayed in Figure 2.4, we find that 40 percent of households received a notice, an increasing slope that continues past the displayed range. Because of the heterogeneous treatment across the different consumption tiers, as discussed earlier, the computer algorithm used by BWP resulted in perfect compliance below the threshold but not above the cutoff. Thus, because the running variable represents a necessary but not sufficient condition for a household to be sent an automated violation notice, the first stage supports a fuzzy regression discontinuity design.

Having established the validity of our first-stage, we examine the effects of the automated violation notices on water use. We start by presenting figures plotting local averages of post-treatment water consumption measures against the running variable. Figure 2.5(a) plots the reduced-form relationship for water consumption during targeted periods of the week, when irrigation was not allowed (all hours excluding Tuesday and Saturday before 9:00 a.m. and after 6:00 p.m.). Figure 2.5(b) shows the reduced-form relationship for water consumption during the entire week. Average water consumption in these figures is pooled over July-October 2015, the four-month period immediately following the automated violation notices treatment and continuing through the end of the statutory local summer water season. This period corresponds with the "Post July" period used earlier in our RCT analysis.

We find a discontinuous drop in water consumption at the threshold for water use during irrigation-restricted periods of the week. The discontinuity in reduced-form is roughly 200 gallons per household per week, or about 7.5 percent of the pre-treatment average weekly water consumption. For water consumption during the entire week, in Figure 2.5(b) we also find a large drop at the threshold, although comparatively smaller, consistent with possible intertemporal substitution in response to the enhanced enforcement of an asymmetric restriction. Overall, water consumption both during the targeted hours of the week and across the entire week decreases across the violation notice threshold.

We investigate these patterns more formally by estimating nonparametric local linear regressions [Imbens and Kalyanaraman, 2012]. Table 2.5 reports RD estimates of the effects of irrigation violation notices. Panel [A] presents reduced-form estimates and Panel [B] presents the local average treatment effects, which essentially rescale the reduced-form estimates by the estimated magnitude of the first stage. Each cell in the table presents a nonparametric RD estimate at the cutoff for automated violation notices. Following Lee and Card [2008], standard errors are clustered along the running variable, which is discrete in gallons. The bandwidth for each specification is selected independently and nonparametrically using a triangular kernel [Lee and Lemieux, 2010].[23]

In Panel [A], we present the reduced-form estimates that correspond to Figure 2.5. Column (2) reports estimates for water consumption during irrigation-restricted times of the week. We find a statistically significant drop of 194.4 gallons per week at the threshold. Some of this decrease in water consumption, however, is offset by an increase in water consumption during non-restricted days and hours of the week. Column (3) shows that water consumption increased during the irrigation-allowed hours on Tuesdays and Saturdays before 9:00 a.m. or after 6:00 p.m.: specifically, water use during these periods increases at the threshold by 47.7 gallons per week. Thus, the violation notices partly shifted water consumption from irrigation-restricted times to irrigation-allowed times, showing intertemporal substitution in response to a policy with (intentionally) partial coverage. Focusing on total weekly water conservation in Column (4), we also find significant reduced-form effects at the threshold, with total water consumption decreasing by 139.7 gallons per week.

[23]Estimates are quantitatively and qualitatively similar when forcing a common bandwidth, such as 40 gallons, to be used across all outcomes.

As displayed in Figure 2.4, household receipt of violation notices increases substantially at the threshold. Confirming this visible discontinuous jump, nonparametric RD estimates indicate an increase of 0.2555 at the cutoff for automated violation notices. The estimate is reported in the first column of Panel [A] in Table 2.5. Given that only one out of four "barely-eligible" households received the violation notice, it is useful to rescale the RD estimates so that they can be interpreted as the effect of receiving a violation notice instead of as a reduced-form estimate of the effects of crossing the arbitrary threshold.

Table 2.5: Regression discontinuity estimates of effects of irrigation violation notice

| | | Weekly water consumption: July-Oct (gallons) | | |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| | First-stage | Non-irrigation hours | Irrig. hours | All hours |
| **Panel [A]: Reduced-form estimates** | | | | |
| Discontinuity | 0.2555*** | -194.4*** | 47.66** | -139.7** |
| | (0.0154) | (56.14) | (21.5) | (66.41) |
| **Panel [B]: Local average treatment effects** | | | | |
| Discontinuity | | -765.6*** | 189.8** | -550.5** |
| | | (227.9) | (84.86) | (264.3) |
| Bandwidth (gal) | 45.17 | 39.16 | 68.84 | 36.37 |
| Observations | 57,259 | 49,316 | 97,861 | 48,509 |

*p<0.1; **p<0.05; ***p<0.01 Each cell presents a nonparametric regression discontinuity estimate at the cutoff for automated violation notices. Following Lee and Card [2008], standard errors in parentheses are clustered along the running variable, which is discrete in gallons. The bandwidth for each specification is selected nonparametrically using a triangular kernel [Lee and Lemieux, 2010]. Column (1) provides the estimated first-stage for automated violation notices, corresponding to Figure 2.4. These notices were sent to households during the first week of July 2015. Columns (2) - (4) present estimates for average weekly water consumption during July through October 2015, the remainder of the legal and technical local summer water season following the violation notices. Panel [A] shows the reduced-form estimates and Panel [B] shows the estimated local average treatment effects. Column (2) includes consumption only during hours of the week when irrigation was not legally allowed, and Column (3) includes consumption only during hours irrigation was legally allowed: Tuesdays and Saturdays before 9:00 a.m. or after 6:00 p.m. Column (4) includes water consumption pooled across all hours. Figure 2.5 graphs the reduced-form local averages corresponding to Columns (2) and (4).

Panel [B] of Table 2.5 reports RD estimates for local average treatment effects of receiving a violation notice. As expected, the LATE estimates are roughly four times larger than the reduced-form estimates. The effect of receiving a violation notice is to reduce post-notice water consumption by 765.6 gallons per week during irrigation-restricted times of the week. In contrast, water use during irrigation-allowed portions of the week increases by 189.8 gallons per week. Finally, total weekly water use decreases by 550.5 gallons per week on average for households sent a violation notice. To place this into perspective, average household water use per week during the pre-treatment period is about 2700 gallons (Table 2.1), implying that these are economically significant effects on the order of more than 20 percent of total water consumption – and nearly 30 percent relative to the barely-untreated side of the RD treatment cutoff.[24]

We draw three main inferences from these estimates. First, strengthened regulatory incentives directly cause significant reductions in water consumption, in addition to any spillovers or other indirect conservation effects from enhanced enforcement. Second, the overall effectiveness of automating violation notices indicates that we can view the "Post July" period as one with now-much-stronger regulatory incentives to conserve water. Third, the results from our difference in (randomized) differences tests in Section 2.3.1 are not simply due to the heightened enforcement of irrigation restrictions being weak and ineffective, supporting the hypothesis that there is no crowd out from interacting the two policy instruments.

[24]In percentage terms, our estimated conservation effects for automating landscape irrigation violation notices are on par with those from subsidizing 'Water Smart Landscape' conversion [e.g. Baker, 2017, Brelsford and Abbott, 2018]. Although estimating the effects of enforcing irrigation restrictions is not the primary focus of our study, this is an interesting and policy-relevant finding in its own right.

Figure 2.5: Reduced-form local averages for post-treatment water consumption

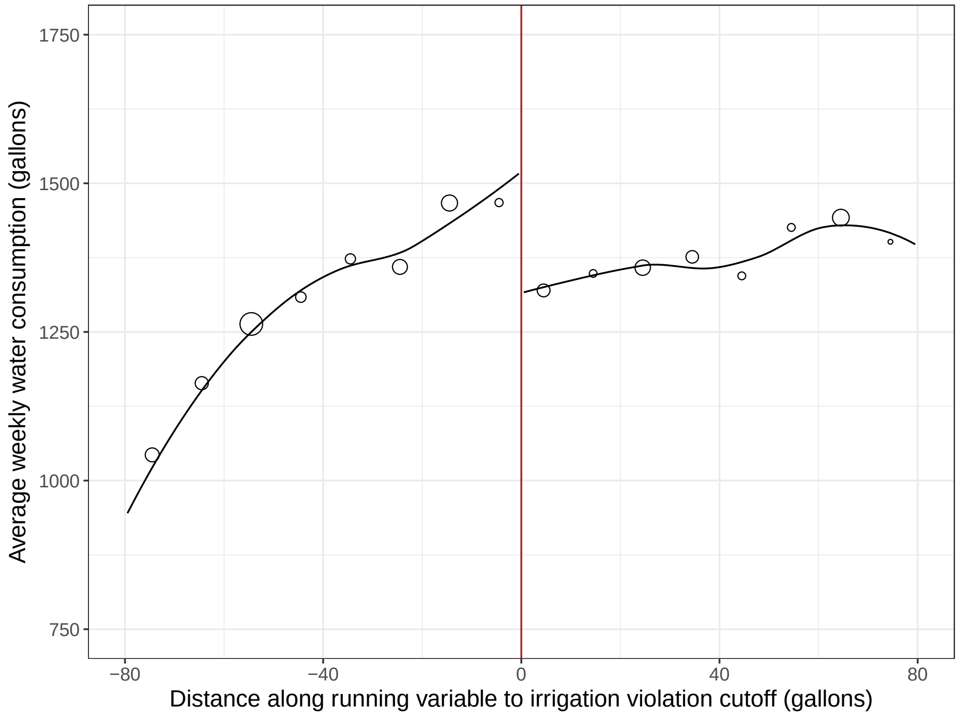


(a) Consumption during hours of the week when irrigation is not allowed

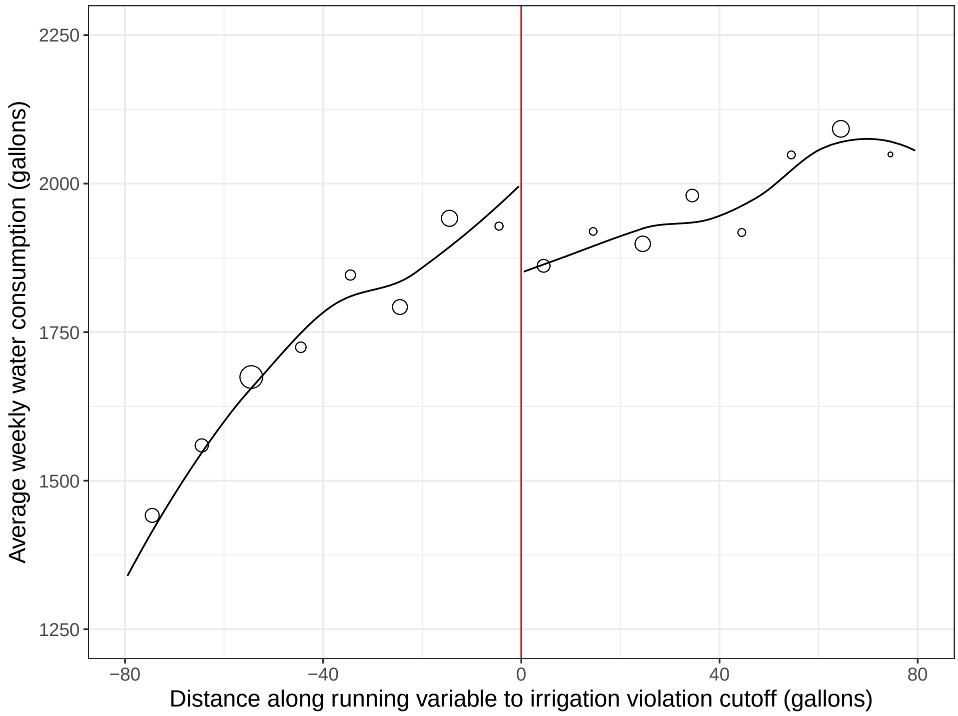


(b) Total weekly water consumption

Notes: Figure 2.5 plots local averages for weekly water consumption during July-October 2015, the period following the automated violation notices treatment. For clarity, the running variable uses 10 gallon bins. The size of the markers corresponds to the number of households included in the local averages. The LOESS curves shown are fit to the underlying microdata separately on each side of the threshold.

## 2.4 Conclusions

The research literature includes ample evidence that pecuniary incentives can crowd out intrinsic motivations for voluntary contributions to public goods across a wide range of settings, but these studies do not address whether monetary incentives also reduce the efficacy of external social pressure interventions to behave prosocially. Focusing on residential water conservation, we examine whether the effects of social pressure are crowded out by strengthened enforcement of landscape irrigation regulations. An important feature of the setting that we analyze is that the heightened regulatory policy is well-enforced and automated because of a new technology allowing for precise monitoring of hourly water use.

We test whether social pressure yields economically significant conservation in the face of binding, technology-enforced regulatory incentives for water conservation. Evaluating a large field experiment among residential water customers, we find that randomly-provided normative peer comparisons cause substantial water conservation over time periods in which strengthened enforcement of landscape irrigation restrictions was in place as well as when it was not in place. The evidence is compelling: in every post-treatment week we find lower water use among treated households that received normative social comparisons. More formally quantifying crowd out using difference in differences tests, in which one difference is randomly generated through the experiment and the other difference compares across neighboring weeks, we find no evidence that the effects of social pressure are crowded out by heightened regulatory incentives; rather, the evidence supports that the conservation benefits could be fully additive.

We also rule out a serious potential threat to the inference of no crowd out, which is that the "strengthened" pecuniary incentives might have been weak and nonbinding. If this were the case, it would not be surprising to find that the effects of the experimentally generated social pressure to conserve water did not change across neighboring weeks. To rule out this explanation, we employ a regression discontinuity design based on a discrete cutoff in the noncompliance algorithm used by the water utility. These estimates indicate that the heightened regulatory incentives also cause significant water conservation, confirming that we are comparing a baseline period to a period with substantially stronger regulation-driven incentives in our test of the crowd out hypothesis.

Governments rely on three major types of policies to encourage resource conservation and public goods provision: regulations, economic incentives, and social pressure. Because of limited ability to implement first-best corrective policies, often due to legal or political constraints, policymakers are increasingly drawing on normative appeals through social pressure as a means to change behavior. Our findings provide evidence that these relatively new policies based around non-pecuniary incentives can strengthen other-regarding preferences to serve as an effective complement to binding regulations. Demonstrating that social pressure remains effective even when interacted with enforcement of regulations is especially relevant as government and non-government institutions increasingly utilize combinations of normative interventions and technology-driven enforcement of policies.

# Chapter 3

# Property Tenure and Determinants of Sensitivity to Price and Non-Price Conservation Instruments

In the United States, the vast majority of conservation interventions in the utility sector rely on either messaging or price incentives for customers. However, 36.6 percent of households in the United States rent their residence, and rental contracts vary in whether the landlord or tenant pays such utilities as water or electricity [Cilluffo et al., 2017].[1] As a result, it is important to understand the effects of utility policy on renters, who generally do not face the same incentives that owners face, both from the utility and for long-term maintenance of a property. This project presents evidence to understand how the effects of price and non-price conservation interventions vary by property tenure status, with a specific focus on drought and residential water policy.

[1] It is difficult to identify the extent to which renters are responsible for water. A recent nationwide survey suggests that up to 94 percent of apartments have some form of recovery, but this includes a 'flat fee,' which means the marginal price of water in these communities is still zero [Munger and Yoon, 2018]. However, coverage of water in California suggests that many renters do not pay a variable cost for their water consumption [de Sousa Mills, 2016].

The challenge of using price-based instruments to induce conservation among tenants has been empirically studied, especially in the electricity sector. Levinson and Niemann [2004] demonstrate that tenants use more energy to heat their homes when heat is included in their rental contracts, but the magnitude of this difference is small. By contrast, Elinder et al. [2017] find very large reductions in electricity consumption – around 25 percent – when tenants change from having electricity included in rent to paying for electricity. In the commercial sector, Jessoe et al. [2018] also find that commercial end-users who pay for electricity consume up to 14 percent less electricity during summer months than commercial end-users who have electricity included in their property lease. In equilibrium, Brewer [2019] predicts residential electricity consumption among renters to fall by 25 percent with the imposition of tenant-pay contracts on all apartments.

Landlord-pay arrangements, when landlords pay the variable cost of water consumption, can thwart incentives both for short-term behavioral change and for long-term efficiency investments. Davis [2012] finds that similar renters and homeowners in similar contexts own different appliances, with homeowners more likely to buy home energy-efficient appliances. Given that the life of the appliance will most likely exceed the length of the rental contract, renters will not fully capture the benefits of appliance upgrades before leaving, even if they pay the utility bill. Renters may be unable to take these upgraded appliances with them when moving out, or they may not be valuable in future locations. Furthermore, renters may act as if they have high discount rates due to credit constraints or other factors, as suggested by findings from Hausman [1979]. If the owner pays the utilities, he or she has an incentive to increase efficiency, but the abuse and depreciation from renter use may make costly investment unprofitable. As a result, renter-occupied housing may be less likely to see water-conserving investments, such as efficient appliances, synthetic turf, or efficient irrigation systems.

In addition to appliance investment decisions, Myers (forthcoming) finds evidence that renters who do not pay for energy are uninformed about the cost of energy. Given long-term efficiency investments, it is also important to consider how renters respond differently with short-term consumption decisions. If raising prices is to be effective as a tool for conservation, the consumers must observe and understand the price of the water or energy that a utility is seeking to conserve.

Furthermore, one key reason to examine the relationship between tenure status and sensitivity to conservation interventions is to understand the external validity of estimated effects. While the share of households renting versus owning their property is only one dimension along which customers differ across service areas, understanding how the effects of price and non-price interventions are mediated by these characteristics is critical to forecasting consumer responses in areas considering similar policies. In order to understand both the behavior of renters and the policy implications of property tenure, I will examine both the differences in response between renter-occupied and owner-occupied housing and the extent to which these differences are due to incentives.

I provide evidence in this paper that the disparity between homeowners and renters in responsiveness to both price and non-price interventions is substantial. In fact, the estimates imply that homeowners are responsible for the majority of the gains from the price and normative interventions studied in this research. However, I also provide evidence that this differential response is driven not by differences in the incentives faced by homeowners and renters but by differences in pre-period consumption and other household characteristics. This appears to be primarily driven by differential exposure of homeowners and renters to the policies, based on the structure of the interventions and the pre-intervention differences in consumption levels between homeowners and renters. Conditional on these factors, there are no statistically significant differences in the response to either intervention by property tenure. I find this lack of differences despite the effects of the underlying interventions being statistically and meaningfully significant. The non-price intervention induced a conservation effect around 2.5 percent, while the price-based intervention generated a conservation effect of nearly 10 percent.

Based on the estimates in this paper, renters may be implicitly or explicitly facing similar incentives when controlling for other factors. A key finding of my study is that the estimated difference in responsiveness between renters and homeowners appears to be a product of the differential extent to which these policies bind for the two groups, on average. Notably, renters are more likely to be lower users, which makes both of the policies studied here less binding. For the price intervention, the price increase only affects users with monthly consumption beyond a utility-defined budget, causing it to affect homeowners more than renters. In the normative intervention, the communications provided negative feedback only to higher users, conditional on characteristics I observe. I observe that renters are less likely to receive negative feedback unconditionally, but there is no statistical difference when conditioning on covariates. In other words, renters are less exposed to negative feedback as a consequence of their baseline characteristics and the characteristics of their residences.

These results have important implications for both conservation policy and economic theory. As a matter of policymaking, these policies have important differential impacts on homeowners and renters. These political consequences are relevant for the policy process, regardless of the underlying mechanism. From a theoretical perspective, it is notable that renters behave as if landlords are able to perfectly pass through conservation incentives.

## 3.1 Setting

The majority of urban California relies on water that falls as rain or snow between October and May, while water demand is highest in the Summer months. With successive dry winters generating the most severe drought conditions on record for the State of California, the Governor declared a drought state of emergency in 2014 and imposed mandatory conservation measures in 2015 (Office of the Governor, 2015). Under the 2015 conservation measures, utilities throughout the state would face financial penalties of $10,000 per day if they did not reduce consumption by 25 percent relative to 2013 consumption [Rogers, 2015]. However, municipal utilities face political constraints on what policies they can implement and what rates

Figure 3.1: Historical Palmer Drought Severity Index, CA

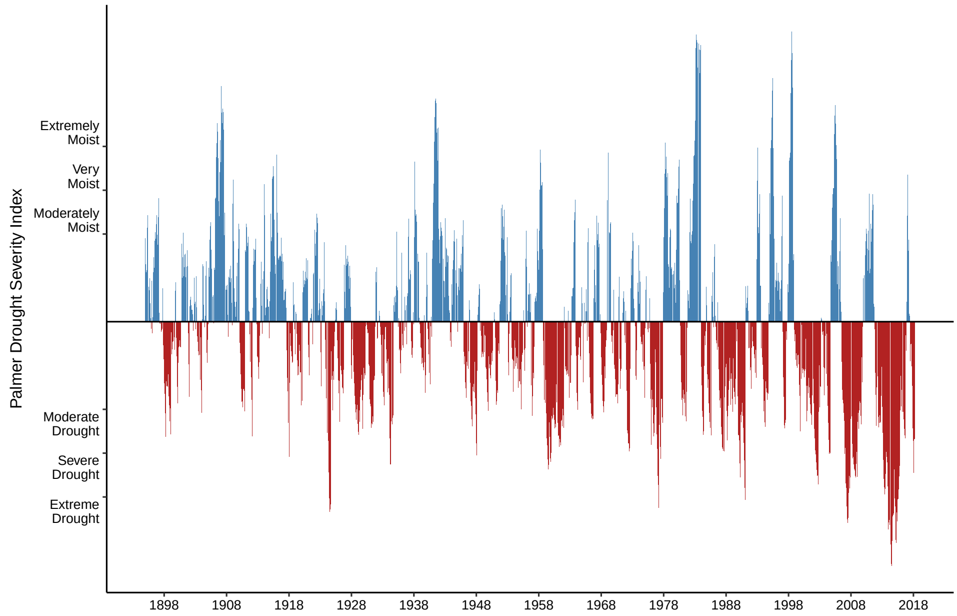


they can charge, while privately-owned utilities are subject to direct government regulation by the California Public Utilities Commission (CPUC).

Drought and water shortages are becoming an increasingly common challenge both globally and in my study area of California. The latest report from the Intergovernmental Panel on Climate Change (IPCC) forecasts that droughts and water shortages will be a major issue facing urban regions in the near-term and into the future [IPCC, 2014]. More specifically, California has historically experienced prolonged periods of drought, but drought has become increasingly common and pervasive. Figure 3.1 shows a surface-area-weighted average of the Palmer Drought Severity Index (PDSI) for the State of California since records began. The PDSI measures soil moisture conditions and is designed to reflect medium-term drought conditions [Palmer, 1965]. The drought conditions during the period I study both reached the most severe deficit of moisture state-wide and extended for years in severe condition. As a result, utilities implemented a diverse array of strategies designed to generate reductions in consumption.

For the purpose of this project, I consider two policies representative of the two primary types of interventions implemented during the drought – price interventions and normative interventions. Broadly speaking, the existing literature has found that residential customers are notably insensitive to price adjustments both specifically in the water sector [Ito, 2013] and more broadly in the utility space [Auffhammer and Rubin, 2018, Faruqui and Sergici, 2011, Gneezy et al., 2011, Ito, 2014, Ito et al., 2018, Puller and West, 2013], with generally low price elasticities of demand. This could be attributable to the small amount of money at stake for households. When larger price changes are implemented, evidence suggests more substantial effects, though this may also relate to non-price channels. For example, Chapter 1 finds evidence of large responses to quota-and-fine restrictions, but these effects appear to extend beyond a price response. In addition to a price effect, this policy appears to have had substantial, though temporary, effects on preferences, along the lines of Gneezy et al. [2011]. In other words, the policy caused not only a movement along the demand curve, but also a shift in the demand curve itself. Furthermore, Wichman et al. [2016] find that responses to both price and non-price interventions vary by certain key factors, such as income for price interventions and the use of an irrigation system.

I examine two neighboring districts, one of which implemented a price increase for part of its rate structure in 2015 and the other of which implemented a normative intervention in late 2016. For clarity, I will refer to the district which implemented the price increase as the 'price intervention district' and the district which implemented the normative intervention the 'normative intervention district.'[2]

### 3.1.1 Price Intervention

The price intervention in this paper involves the temporary increase of the marginal price of water for the middle of a five-tier increasing block rate schedule. For each household, the price intervention district calculates an indoor and an outdoor water budget for each billing cycle using such factors as the number of occupants and the extent of irrigable area, as described below. The marginal rates for the first two tiers were not changed under the 'emergency' rate change I study, and the district

[2] For confidentiality reasons, I am not able to disclose the names of the districts.

had already been implementing a five-tier rate structure with personalized budgets for years before the implementation of this emergency rate adjustment.

The structure of the five tiers is as follows. The first marginal rate tier applies to consumption between zero and the indoor water budget. The second rate tier is for consumption greater than the indoor budget up to the combined total of the indoor and outdoor budgets. Marginal rate tier 3 applies to water consumption greater than the total budget up to 1.25 times the total budget. Marginal rate tier 4 applies to consumption greater than 1.25 times the total budget up to 1.5 times the total budget. Marginal rate tier 5 applies above 1.5 times the total budget. These tiers also featured normative labels, with tiers 3, 4, and 5 labeled "Inefficient Use," "Wasteful Use," and "Unsustainable Use," respectively. Prices are given in dollars per 100 cubic feet, also known as 'centum cubic feet' (CCF).

The rate structure over time is shown in Table 3.1. These marginal rates do not include additional volumetric charges for pumping but do include a charge of $0.42 per CCF that the utility charges for the stated reason of reliability.[3] In response to drought conditions, the price intervention district 'removed' tiers 3 and 4 starting on July 1, 2015. Effectively, this means that tiers 3 and 4 are charged the tier 5 rate. As the water shortage became less urgent, tier 3 was 'reinstated' on July 1, 2016, with tier 4 still charged at the higher tier 5 rate. The emergency rates were completely removed as of March 1, 2017. The underlying rate structure did not otherwise change from January 1, 2015 until January 1, 2018.

Table 3.1: Treatment Rate Structure, Price Intervention District, Primary Service Area

| | 2014 | 2015, 2016, 2017 | | |
|---|---|---|---|---|
| | | Base Rates | Emergency 15-16 | Emergency 16-17 |
| Tier 1 [Indoor Budget] | $2.357 | $2.398 | $2.398 | $2.398 |
| Tier 2 [Outdoor Budget] | $2.637 | $2.726 | $2.726 | $2.726 |
| Tier 3 [1 – 1.25×Budget] | $3.240 | **$3.269** | **$5.734** | **$3.269** |
| Tier 4 [1.25 – 1.5×Budget] | $4.801 | **$4.844** | **$5.734** | **$5.734** |
| Tier 5 [> 1.5×Budget] | $5.691 | $5.734 | $5.734 | $5.734 |

[3]Pumping charges vary by location due to elevation, but the 'reliability' charge is the same for all households across the district.

The budgets are calculated using a limited amount of data, all of which are available in my data. Historical consumption does not enter the calculation for the indoor and outdoor budgets. The indoor budget calculation depends only on the number of occupants and the number of days in the billing cycle. The indoor budget for household $i$ with its meter read on date $t$ is calculated as follows:

$$\text{Indoor}_{it} = (60 \text{ gals./day}) \times \text{Occupants}_i \times \text{Days in the billing cycle}_{it} / (748 \text{ gals./CCF})$$

The outdoor budget depends on localized weather data, irrigated area, and a landscape factor. The evapotranspiration rate reflects the rate at which plants and irrigated areas, in general, lose water to the atmosphere. The landscape factor is based on the types of plants in customer yards, with a default value of 0.8 for accounts that opened before 2012 and a default value of 0.7 for accounts that opened during or after 2012. The outdoor budget for household $i$ with its meter read on date $t$ is calculated as:

$$\text{Outdoor}_{it} = \text{Irrigated acreage}_i \times \text{Evapotranspiration}_{it} \times \text{Landscape factor}_i \times .62$$

As shown by the scripting, both the indoor and outdoor budgets may over time. The indoor budget will only vary slightly as the length of the billing cycle changes, with 95 percent of billing periods falling between 26 and 36 days and a median of 29. There is more variation in the outdoor budget, driven by an evapotranspiration rate that is calculated during the billing cycle, with 95 percent of billing cycle evapotranspiration rates falling between 2.27 and 7.96 and a median of 5.28. Evapotranspiration rates are highly correlated with seasonal changes in the weather, with higher rates of evapotranspiration in the summer.

Because I use a sample from the normative intervention district as the control group for the price intervention, I will discuss its pricing structure here and the normative policy in the following subsection. In contrast to the rate structure in the price intervention district, the normative intervention district implemented a rate structure with four tiers and seasonal rates, as shown in Table 3.2, but it did not change rates during the period in question. For a small number of accounts, the normative intervention district charges 1.5 times the rates shown in Table 3.2 on the basis of location. All of these costs are only the marginal cost of water in each district. Households also face a variety of fixed costs per billing cycle based the size of their pipes, their elevation, and other factors.

Table 3.2: Control Rate Structure, Normative Intervention District, Primary Service Area

| | April 2014 - June 2018 | |
|---|---|---|
| | Summer | Winter |
| Tier 1 [≤ 15 CCF] | \$1.14 | \$1.13 |
| Tier 2 [16 – 35 CCF] | \$1.83 | \$1.64 |
| Tier 3 [36 – 60 CCF] | \$2.85 | \$2.26 |
| Tier 4 [> 60 CCF] | \$4.10 | \$2.75 |

Summer rates are in effect for service rendered during June, July, August, September, and October.

### 3.1.2 Non-Price Intervention

Non-price channels are popular in the utility sector for encouraging conservation, where price adjustments are not only potentially ineffective but also legally constrained. Experimental evidence has repeatedly demonstrated that comparing a household's usage to its peers is effective in generating reductions of varying magnitudes [Allcott, 2011, Allcott and Rogers, 2014, Ferraro and Price, 2013]. Similar to my study, Brent et al. [2015] find significant impacts from WaterSmart reports, which are very low-cost for the utility. Notably, however, it is unclear whether renters ever see these comparisons or if the entire effect is driven by owner-occupied homes.

My study also assesses the impact of experimentally-assigned WaterSmart Home Water Reports (HWRs). In the normative intervention district, the vast majority of single-family residential accounts were randomly assigned to either receive an HWR or to serve as control households.[4] An example HWR is shown in Figure A.1.

[4] Randomization was conducted and implemented by WaterSmart in conjunction with the utility.

Treatment involves monthly reports comparing a household's consumption to other households with similar characteristics, such as lot size and number of occupants. Households consuming less than the 20th percentile of their comparison group are considered 'efficient' and shown a smiling water drop on a green background, while households consuming more than the 55th percentile of their comparison group are told that they "used more water than most of [their] neighbors" and shown a frowning water drop on a red background. A neutral label and water drop on a yellow background is given to those consuming between 20th percentile 'efficient' level and 55th percentile 'average' level, both of which are labeled on a bar graph with the household's consumption.[5] While these reports include suggestions and other relevant information, the existing literature has found that the normative social comparison is by far the most effective component [Ferraro and Price, 2013].

Control households in the normative intervention district received no intervention, and certain additional households were not assigned to treatment or control, on the basis of insufficient pre-period data. While randomization was not stratified geographically, random sampling should ensure no correlation between geography and treatment status. WaterSmart has implemented the normative intervention I study in many water districts across California and in other states, beginning around 2012 and continuing through the present. Each participating district initiated the HWR treatment at a different point in time for both endogenous, such as worsening local water shortage or insufficient conservation, and exogenous reasons, such as bureaucratic processes. Existing evidence suggests that the program has been effective, with statistically significant conservation in the range of three to six percent Ferraro et al. [2011], Ferraro and Price [2013], Bernedo et al. [2014], Brent et al. [2015], Jessoe et al. [2019].

[5] HWRs report the 55th percentile as 'average,' ensuring that fewer than half of accounts are told that they are using more water than most of their neighbors.

## 3.2 Data

This project takes advantage of data on consumption and certain household characteristics for single-family residential accounts in two Southern California water districts. The data include information on irrigable area, number of occupants, and geographic identification. This geographic identification allows for linking the accounts to Census block groups. In addition, precisely identified households in one district can be linked to property data, including value and features of the house.

Using the 2016 five-year American Community Survey, I can include a range of demographics.[6] Most directly relevant to my study, I am able to ascertain the share of households in a Census block group that is owner- versus renter-occupied. Because other important characteristics are correlated with the probability that a housing unit is rented, I also include data on median home value of owner-occupied housing, in addition to the data on irrigable area and number of occupants. The value of homes in a neighborhood can be expected to be correlated with income, which I am not able to include at the same level. Irrigable area is defined as the surface area of a property which might reasonably be expected to be watered or irrigated. The districts infer this information from parcel assessments, satellite data, and engineering assumptions, and households have the ability to revise this estimate in communication with a water district. This measure is imperfect, as it may include pools and other areas of pavement, but it is highly correlated with the actual irrigation consumption, conditional on landscaping choices. The number of occupants is estimated by the utility using Census data and other information, but accounts are encouraged to update this information if incorrect. There may be asymmetric incentives for correcting this information, which I will address further in Section 3.3.

The billing data includes monthly consumption, with the initial and final date of each billing cycle. Combining this information with time fixed effects precludes the need to include weather data, because there is limited spatial variation within

[6] My study takes advantage of the publicly available summary file, found here: https://www.census.gov/programs-surveys/acs/data/summary-file.html

a given time period. These billing data are provided by WaterSmart and the California Data Collaborative, both of whom maintain technical support infrastructure for water districts throughout the state.

For the normative intervention, WaterSmart only included in the experiment single-family residential accounts without missing pre-period data. Among these households, accounts were randomly assigned to either a treatment or control group. Households who were not eligible for the experiment are excluded both from the normative intervention sample and from the pool of potential matches for households in the price intervention sample. All single-family residential accounts are included in the price intervention group, but the matching algorithm, which I describe in further detail in the next section, also requires data in the pre-period. As a result, households without at least one bill in 2014 were excluded from both samples.

## 3.3 Identification

This study includes two analyses with different identification strategies. To study the effects of my primary non-price instrument, WaterSmart HWR social comparisons, I am able to leverage random assignment to experimental treatment. For precision and to reduce confounding variation, I use a difference-in-differences framework with both household and time fixed effects. I also interact the difference-in-difference estimator with a measure of the share of renter-occupied housing in the household's Census block group, as well as with other controls.

With the price-based instrument in my sample, I must leverage a natural experiment. For this, I compare households across neighboring districts, where the price intervention district increased prices for part of the pricing schedule. The control district is the normative intervention district, but the price instrument time period includes only bills prior to the implementation of the normative intervention. For additional precision, I use a matched sample of control households, with replacement, but the results are consistent using the full sample.

In order to examine the empirical relationship between property tenancy and treatment effects for water conservation interventions, I need to use only variation in tenancy that is orthogonal to other determinants of a household's responsiveness to water policies. To that end, I use property-specific covariates and Census block group level data to control for confounding variation. Specifically, I include the irrigable area of a property, the number of occupants, and the median home value of owner-occupied homes in the Census block group.[7]

For each of the samples, I implement a difference-in-differences approach. For each intervention, I estimate three specifications. The base specification is shown in equation 3.1 and demonstrates the reduced-form effectiveness of each intervention in the total relevant sample. The outcome in the base specification is the natural logarithm of daily water consumption in gallons for household $i$ and the billing cycle ending on date $t$.[8] Here, I include a binary indicator of whether a household was assigned to treatment, $1\,(\text{Treated})_{it}$, in addition to fixed effects for both households, $\eta_i$, and week of consumption reading, $\tau_t$.[9] These fixed effects control for common weather and external shocks that vary across time but not across households.

$$\log\left(\text{Cons}\right)_{it} = \delta \times 1\left(\text{Treated}\right)_i + \eta_i + \tau_t + \epsilon_{it} \tag{3.1}$$

The second specification, shown in equation 3.2, is an unconditional approach to estimating the differential response to the policy by a proxy for whether the household is renter- or owner-occupied. The variable I use, Renter Share, is the percent

[7]The Census block group level data both contain underlying uncertainty and are only a proxy for the actual household characteristics. As a result, there is classical measurement error in the right-hand-side variables, potentially leading to attenuation bias. Given the estimates presented in Table 3.5, the attenuation bias and standard error inflation does not appear to be severe.

[8]The panel data is effectively missing all other dates other than dates that a household had its meter read, but this construction is meant to coerce the rolling billing cycles of true consumption into a form that allows for appropriate temporal controls. Mathematically, only one observation per household per billing cycle is included.

[9]The billing cycle is closed on a given date $t$, and I use the week of that date for the fixed effects. A finer level of fixed effects, such as the date, could lead to erroneously using the time fixed effect to control for the characteristics of the majority of accounts read on that date. Because the date a meter is read is correlated geographically, this can confound the intended purpose of controlling for weather and other similar effects. A coarser level of fixed effects, such as the month, would inaccurately group bills read weeks apart, which may have faced substantially different weather and other exogenous factors.

of households at the Census block group level that are renter-occupied. I include an interaction between the treatment indicator and Renter Share, which should identify the differential response of interest. I do not include the variable itself, as Renter Share does not vary over time, and the differences in consumption between households with different values here are captured in the household fixed effects. Under the assumption of no omitted variables, this specification would provide a direct estimate of the differential response between renters and homeowners, $\rho$.

$$\log(\text{Cons})_{it} = \delta \times 1(\text{Treated})_{it} + \rho \times 1(\text{Treated})_{it} \times \text{Renter Share}_i + \eta_i + \tau_t + \epsilon_{it} \quad (3.2)$$

With my third specification, I examine how this estimate changes when controlling for potentially confounding variation. In equation 3.3, I include interactions between the effect of the intervention and other covariates, $\mathbf{X}_i$, namely the irrigable area and number of occupants of the household, the median home value of the block group, and the account's historical consumption, taken as the average daily consumption billed for 2014. I take the natural log of both the irrigable area and home values, based on goodness of fit and visual inspection of the relationship between these covariates and the outcome. I use bins of historical consumption, with each bin having a width of 200 gallons. For example, a household consuming 341 gallons per day in 2014, on average, would be in the bin for (200, 400].

$$\begin{aligned}\log(\text{Cons})_{it} =& \delta \times 1(\text{Treated})_{it} \\ &+ \rho \times 1(\text{Treated})_{it} \times \text{Renter Share}_i \\ &+ \boldsymbol{B} \times 1(\text{Treated})_{it} \times \mathbf{X}_i \\ &+ \eta_i + \tau_t + \epsilon_{it}\end{aligned} \quad (3.3)$$

### 3.3.1 Normative Intervention

For the normative intervention, I use the entire randomized sample in the normative intervention district. While the assignment to treatment was random, there is a small difference in the number of occupants per household across groups. The size of this difference is ultimately very small, and the differences across treatment and control groups in other variables of interest are statistically insignificant and of inconsequential magnitudes.

Table 3.3 compares the means of household characteristics across the treatment and control groups within the District. Standard errors are shown for group means, and the p-value of the difference is shown below the magnitude of the difference.[10] For renter share and median home value, these variables are known at the Census block group level but assigned to individual households. As a consequence, there should be no bias of the group means in expectation, but the standard errors are artificially deflated by a lack of within-block-group variation that does, in fact, exist.[11] This does not appear to be an issue for this intervention, because the only meaningful difference in Table 3.3 is for the number of occupants. Notably, treated households might have an incentive to manipulate reported irrigable area and number of occupants, because these data dictate the likelihood of receiving negative feedback on the HWR, conditional on consumption. There is evidence of changes to occupancy, which is most likely the cause of the statistical difference between the treatment and control groups, but there is no evidence of any household changing its irrigable area. Unfortunately, I cannot use pre-period occupancy, but I can remove those households who changed their listed number of occupants from the analysis. Removing these households does not alter the results in any meaningful way. For the specifications presented in this paper, I drop the bottom one percent and top one percent of consumption readings within the broadest definition of the sample. Some of these observations are for negative consumption or physically incredible quantities of water. The results are, however, generally robust to outliers and other methods of controlling for them.

[10]The balance tables in this paper rely on the `balancetable` package for Stata.

[11]To understand why this is the case, consider the fact that the standard deviation of $\{3,3,3,5,5,5\}$ is less than the standard deviation of $\{2,3,4,4,5,6\}$, even though both sets are the merger of two subsets with the same median and mean. Because the true underlying data look like the second set but the data in my sample look like the first set, the standard errors of the means presented in Table 3.3 for the block group characteristics are articially smaller than they would be with the true underlying data. It is possible to simulate the underlying distribution to estimate the standard errors more accurately, but that would not change the conclusions I present and would still suffer from uncertainty.

Table 3.3: Balance Table, Normative Experiment

| | Control (SE) | Treated (SE) | Diff (p-value) |
|---|---|---|---|
| Renter Share | 0.383 | 0.380 | -0.003 |
| | (0.236) | (0.233) | (0.223) |
| Irrigable Area | 7,100.668 | 7,113.895 | 13.227 |
| | (19,320.705) | (23,462.727) | (0.948) |
| Median Home Value | \$287,835.19 | \$288,720.53 | \$885.35 |
| | (105,373.773) | (105,792.070) | (0.402) |
| Number of Occupants | 2.955 | 3.116 | 0.161 |
| | (1.272) | (1.386) | (0.000)*** |
| Observations | 38,524 | 14,020 | 52,544 |

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

To test the effects of the randomized HWRs, I use a difference-in-differences specification. While I have random variation in the assignment to treatment, I use pre-period data and household fixed effects to increase precision and reduce confounding variation that is correlated with the demographic characteristics of interest. Only data from households who were in the randomization sample are included, as described in Section 3.2. Moreover, for comparison across specifications, observations with missing data are excluded in all specifications. In order to avoid data from a prior pricing regime shift, I use two years of pre-period data starting in October of 2014, as well as one year of post-treatment data, through September 2017. The first HWRs were sent out on October 2, 2016, with nearly all accounts in the treatment group receiving an initial HWR before the end of October 2016. The last HWRs were sent in September of 2017. Not all accounts received HWRs throughout this period, either by choice or due to changes to the account, but these cases are rare. While the persistence of the intervention after its conclusion may be of interest, my analysis is focused on heterogeneity in how different households responded during the intervention. Evaluating behavior over a longer time horizon brings in different hydrological conditions and would only assess persistence within this specific context.

### 3.3.2 Price Intervention

For the price intervention, I again use a difference-in-differences estimator, but without perfectly random assignment to treatment. Instead of random assignment, my identification rests on the assumption that similar households in each district would have remained on parallel consumption trends in the absence of the price intervention. While the two districts are different, they serve overlapping populations, which forms the basis of my matching approach, and the timing and details of the intervention are arguably exogenous and related to idiosyncratic institutional and bureaucratic factors.

In order to execute this strategy, I utilize pre-period data to match households in the price intervention district with households in the normative intervention district, which serves as the control for this analysis. Specifically, I match all single-family residential households in the price intervention district to single-family residential households in the experimental sample in the normative intervention district using block-group renter share and median home value, as well as household-level irrigable area, occupancy, and pre-period consumption. To capture pre-period consumption, I take the average daily consumption in gallons across all billing cycles concluded prior to the price intervention. Only accounts with all of this information were used in the matching procedure.

I use propensity scores from a full Mahalanobis matching algorithm and match one-to-one with replacement.[12] As seen in Table 3.4, the average number of treated observations per control observation is 2.274 (median of 2). I present unweighted specifications, but all results are robust to weighting observations by the number of matches.[13] As in the normative intervention, all regressions remove the top and bottom one percent of observations, which are likely errors. Contrary to the district implementing the normative intervention, the minimum reported occupancy in

[12] I take advantage of the `psmatch2` Stata package to enerate propensity scores and identify matches for each treated household.

[13] I conduct this weighting such that an observation of a control account that is matched to three treated accounts is given three times the weight of an observation of a control account that is only matched to one treated account.

this district is three people. This stems from a combination of the price intervention district's default inference and the fact that households have no incentive to revise their occupancy lower, because it would result in higher prices.

To support the validity of the difference-in-differences estimator, I show evidence of parallel trends among matched households. Figure 3.2 graphs collapsed weekly average consumption in gallons per household per day by district. Each district-by-week observation is shown in size proportional to the number of bills observed in that district during that week, with a local polynomial fit for each district and vertical line showing the beginning of the price intervention. The districts exhibit similar pre-trends, but after the price increase was enacted, the separation between treatment and control households shows a modest expansion. I test this treatment effect empirically later in this subsection.

Figure 3.2: Parallel Trends Assumption, Matched Sample

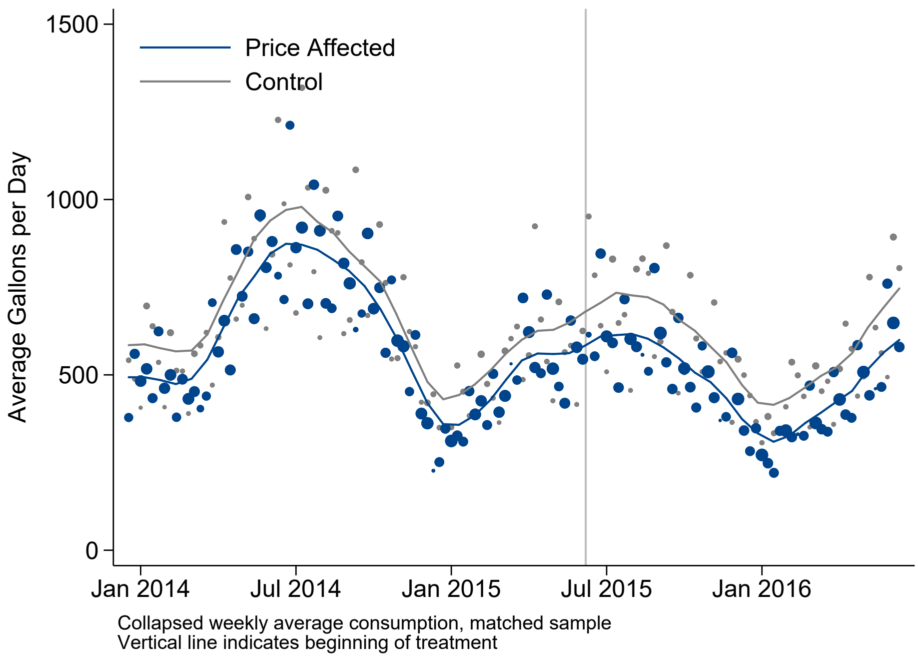


As shown in Table 3.4, household characteristics are statistically different across the districts. However, the magnitudes of the differences are fairly modest, and the comparability between treated and control households is substantially better than when using the full sample of households. For renter share and median home value, these variables are again known at the Census block group level but assigned to individual households. Because not all households in a given block group are included, especially for the control group, these estimates may be slightly inaccurate, and the standard errors will also be artificially deflated by a lack of within-block-group variation in the data that does exist in reality. Given the limited differences across groups, if the standard error deflation is more important than any bias, the groups are more similar than they would appear in Table 3.4. Given the strong evidence for parallel tends in Figure 3.2, these modest differences should not pose a challenge to identification.

Table 3.4: Balance Table, Matched Sample

| | Control | Treated | Diff |
|---|---|---|---|
| | (SE) | (SE) | (p-value) |
| Renter Share | 0.211 | 0.180 | -0.031 |
| | (0.171) | (0.145) | (0.000)*** |
| Irrigable Area | 9,695.297 | 9,909.885 | 214.588 |
| | (29,242.762) | (68,008.266) | (0.742) |
| Median Home Value | $375,939.66 | $396,564.16 | $20,624.51 |
| | (130,642.453) | (68,699.977) | (0.000)*** |
| Number of Occupants | 3.787 | 4.088 | 0.301 |
| | (1.224) | (2.256) | (0.000)*** |
| Observations | 11,798 | 26,829 | 38,627 |

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

I do not include price or conduct a full demand estimation. Despite the fact that this policy uses prices as a lever, related evidence suggests that even price-oriented policies have non-price effects and impact non-targeted households (See Chapter 1). Given this evidence, estimating the effect of the policy only through the price elasticity of demand would yield biased estimates of that elasticity. Furthermore, the household fixed effects remove heterogeneous preferences and baseline differences in household demand curves. As a result, I am looking only at the reduced-form effect of the policy net of baseline demand, including both the effects stemming from the price elasticity of demand and from non-price channels.

## 3.4 Results

Estimating the specifications from Section 3.3, I find suggestive evidence of a differential response by property tenure to the normative intervention and statistically significant evidence of a differential response for the price intervention. In both cases, this differential response is driven primarily by differences in the characteristics of renter- and owner-occupied homes.

### 3.4.1 Normative Intervention

The first column of Table 3.5 presents the results from estimating equation 3.1 for the HWR intervention. Because I cannot determine whether any individual household opened a HWR that was mailed to them, the estimated value of $\delta$, $\hat{\delta}$, is an intent-to-treat (ITT) estimate. This implies that the average treatment effect (ATE) of actually reading the report is likely larger in magnitude. I find that the HWRs were successful in generating reductions of around 2.4 percent in the treatment group, which is consistent with estimates from the existing literature on HWRs [Brent et al., 2015, Ferraro and Price, 2013]. The effectiveness of the intervention overall allows us to examine the extent to which there was a heterogeneous response along various characteristics, particularly property tenure.

The second column of Table 3.5 estimates equation 3.2 for the HWR intervention. As shown in Table 3.3, the Renter Share variable is balanced across treatment and control groups and not correlated with assignment to treatment. In the second column of Table 3.5, I show that my estimate of $\rho$, $\hat{\rho}$, is statistically insignificant but positive. It is important to note that the magnitude of the point estimate would imply that homeowners are responsible for nearly all of the aggregate effect. Given both the size of the estimate and the noise with which I proxy for property tenure, it is worth identifying this effect more precisely.

The third column of Table 3.5 provides estimates of Equation 3.3 for the HWR intervention. Here we see that controlling for these covariates essentially eliminates the differential response between homeowners and renters.[14] While not shown here, these findings are consistent with the fact that renters are unconditionally less likely to receive an 'inefficient' rating on their first HWR, but there is no difference in the likelihood of receiving such a score conditional on the covariates included in column 3.

Table 3.5: Sensitivity to Normative Intervention by Property Tenure

| | 1 | 2 | 3 |
|---|---|---|---|
| HWR | -0.0243*** | -0.0460*** | - |
| | (0.00310) | (0.0141) | |
| **HWR** | | **0.0594** | **-0.000247** |
| **×Renter Share** | | **(0.0390)** | **(0.0191)** |
| HWR | | | -0.00290 |
| ×log(Irr. Area) | | | (0.00401) |
| HWR | | | -0.0228 |
| ×log(Home Value) | | | (0.0274) |
| Occupancy Bin Interactions | No | No | Yes |
| 2014 Cons. Bin Interactions | No | No | Yes |
| FE | Household ID, Week of Time | | |
| Cluster Var. | Household ID, Week of Time | | |
| $R^2$ | 0.751 | 0.751 | 0.752 |
| Obs. | 1,067,021 | 1,067,021 | 1,067,021 |

Highest one percent and lowest one percent of observations dropped.
Cluster-robust standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 3.3 plots regression coefficients for equations akin to columns 2 and 3, but with separate coefficients by each quartile of renter-share. Neighborhoods with the lowest share of renter-occupied housing is the omitted category. This further reinforces the extent to which the intervention primarily generates savings from homeowners, but it also confirms the findings from Table 3.5 that this differential response is completely due to differences in baseline characteristics of households in each group. For the unconditional specification, the findings are noisy, but the estimated differences between the quartiles shown and the omitted quartile is two-thirds the estimated effect for the high-homeowner neighborhoods. These sugges-

[14]This finding is driven more by pre-period consumption than any other control, but each factor is relevant.

Figure 3.3: Sensitivity to HWRs by Renter-Share Quartile

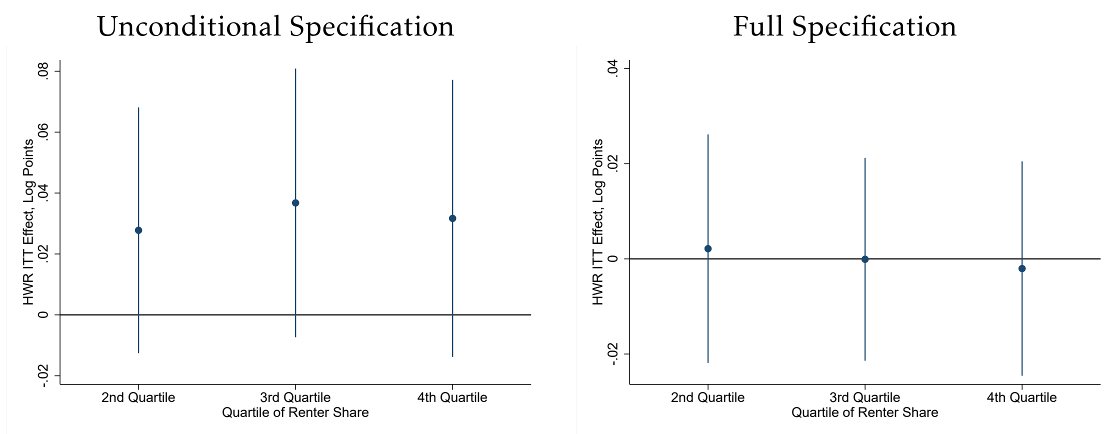


tive differences do not persist in the full specification.

It is worth noting that irrigable area and median home value are both relatively more important than Renter Share, though statistically insignificant, in determining the effectiveness of the policy. In order to interepret the coefficients in relative terms, one must know the mean and, particularly, standard deviation of the covariates. For the sample, the mean and standard deviation of each variable is: Renter Share, 0.365, 0.221; log(Irr. Area), 8.51, 0.676; log(Median Home Value), 12.5, 0.360. As a result, the effect of a one standard deviation increase in Renter Share is only approximately 2.5 percent as large as a one standard deviation decrease in the natural log of irrigable area. Furthermore, the effect of a one standard deviation increase in Renter Share is about one percent the effect of a one standard deviation decrease in the natural log of the median home value. The prominence of home value is not surprising, given that it is correlated with income, which Wichman et al. [2016] found to be the primary driver of differential response to water policy in a separate setting.

More important than either of these effects, however, is the difference in response by pre-period consumption. Figure 3.4 plots the coefficients (solid line) and 95 percent confidence intervals (dashed lines) on indicators for pre-period daily consumption in bins of 200 gallons per day. The omitted reference group are households consuming less than 200 gallons per day in 2014, so the coefficient estimates should be interpreted as the difference in response to HWRs between households consuming in a given bin in 2014 and otherwise similar households consuming less than 200 gallons per day in 2014. Because consumption is persistent and because higher consumption conditional on household characteristics (the x axis in the Figure) is a direct determinant of the Water Score that a household receives, it is unsurprising that accounts with higher conditional consumption respond more strongly. In other words, Figure 3.4 provides evidence that getting more negative normative feedback is correlated with a stronger response to the policy.

Figure 3.4: HWR Effects by Pre-Period Consumption

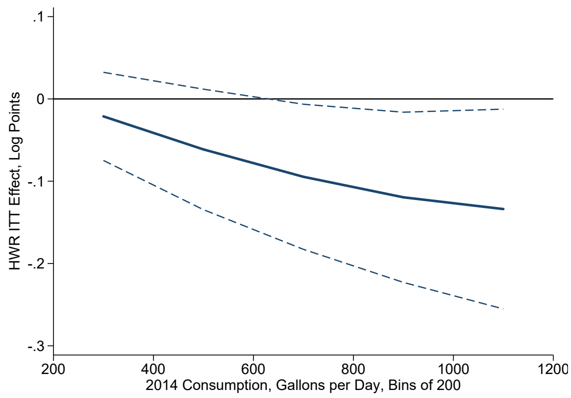


Given that one key to understanding the difference in response appears to be the underlying differences in baseline characteristics between owner-occupied and renter-occupied units, it is sensible to consider whether these findings are primarily about the types of residences inhabited by each group. Within my data for this experiment, I observe whether an account is listed as single-family housing or a condominium. Notably, condominiums are located in neighborhoods with a higher share of renter-occupied housing and more likely to by renter-occupied. As shown in Figure 3.5, the reduction in the effect is visible only for single-family housing units, but the estimates for condominiums are noisy due to a relatively small number of accounts classified as condominiums. It is possible that these point estimates reflect differences in the pass-through of incentives between renters in condominiums or apartments and renters living in single-family housing, but the noise in this sample is too large to make a well-supported conclusion.

Figure 3.5: HWR Effects by Renter Share, Single-Family Houses versus Condominiums

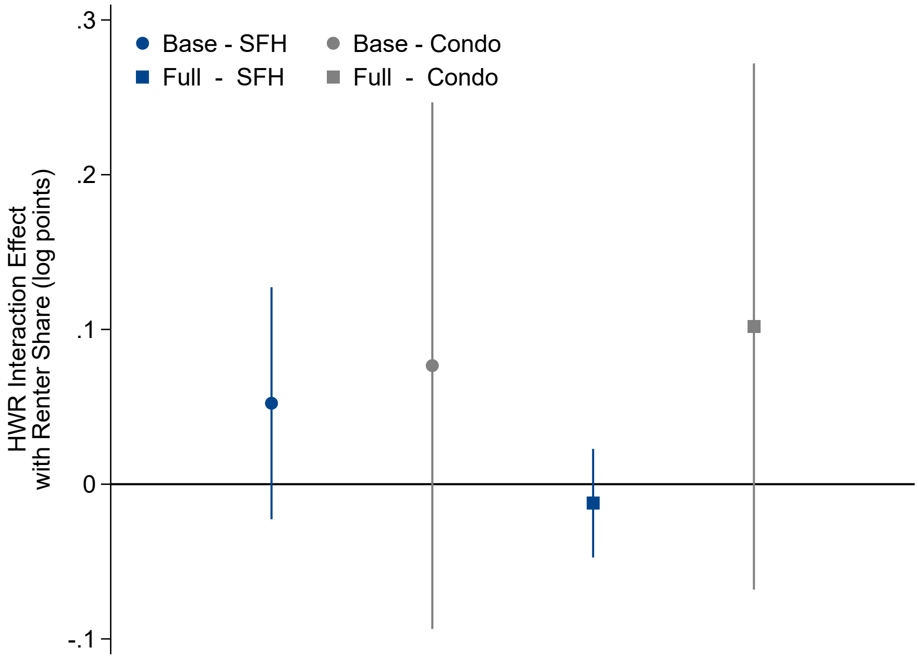


### 3.4.2 Price Intervention

Column 1 of Table 3.6 presents estimates of $\hat{\delta}$, which is the average treatment effect (ATE) of the policy for all households. The policy was effective, with an estimated nearly nine percent conservation in the post-period relative to the neighboring district. Column 2 of Table 3.6 provides evidence that households in high-renter-share neighborhoods were substantially less responsive. Indeed, the point estimates would suggest that homeowners were responsible for nearly all of the effect. This finding is relevant for the policy conversation, as it identifies which households shoulder the burden of the policy. However, this is not necessarily due to differences in the incentives facing renters and homeowners, nor to differences in responses to those incentives.

As shown in column 3 of Table 3.6, controlling for how response to the policy varies by these characteristics eliminates any variation in response between homeowners and renters. Notably, the key control variable is 2014 average daily consumption,

Table 3.6: Sensitivity to Price Intervention

| | 1 | 2 | 3 |
|---|---|---|---|
| Price Increase | -0.0880*** | -0.112*** | |
| | (0.0105) | (0.0149) | |
| **Price×Renter Share** | | **0.137***** | **0.00139** |
| | | **(0.0510)** | **(0.0534)** |
| Price×log(Irr. Area) | | | 0.0325*** |
| | | | (0.00551) |
| Price×log(Med. Home Value) | | | 0.0209 |
| | | | (0.0243) |
| [-.7em] | | | |
| Occupant Bin Interactions | No | No | Yes |
| 2014 Cons. Bin Interactions | No | No | Yes |
| FE | Household ID, Week of Time | | |
| Cl. Var. | Household ID, Week of Time | | |
| $R^2$ | 0.772 | 0.772 | 0.773 |
| Obs. | 949,972 | 949,972 | 949,972 |

Highest one percent and lowest one percent of observations dropped.
Cluster-robust standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

which is not surprising.[15] Consumption is seasonal and prone to noisy variation, but historical consumption is one of the best predictors of current consumption. As a consequence, households with higher pre-period consumption were more likely to be affected by the price increase, which was only in place for consumption beyond the indoor and outdoor usage budget. For the first month of the price increase, the median household's total budget was equivalent to approximately 830 gallons per day, and consumption below that threshold was not subject to this large price increase. Calculating forecasted reductions, only households consuming in the 600 to 800 gallons per day bin or higher in 2014 would be predicted to have reduced consumption in response to the price increase. Given that summer consumption is generally substantially higher than average consumption, this is the first consumption bin where households would have reasonably faced higher prices, in expectation.

[15]Specifically, a regression specification which expands on column 2 only by including controls for 2014 average daily consumption generates a small, negative, statistically insignificant effect of Price×Renter Share.

Figure 3.6 plots the point estimates (solid line) and 95 percent confidence intervals (dashed lines) of the marginal additional treatment effect by pre-period consumption bin. Each point estimate should be considered as a comparison in the average treatment effect among households consuming at that level in the pre-period with otherwise similar households consuming less than 200 gallons per day in the pre-period. I will subsequently discuss the ways in which renters and homeowners differ historically and how this differentially exposed them to this policy.

Figure 3.6: Price Effects by Pre-Period Consumption

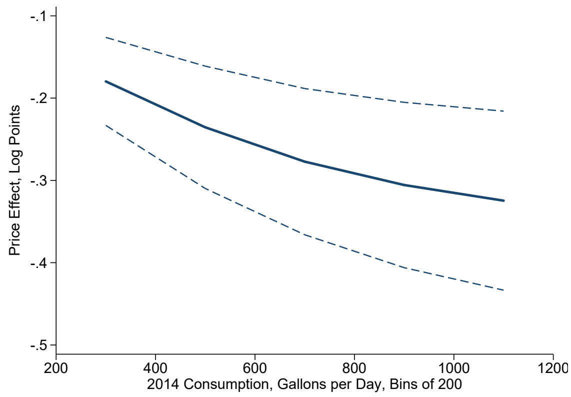


As with the normative intervention, there is evidence to confirm the findings from columns 2 and 3 from Table 3.6 with a more non-parametric approach. Figure 3.7 plots the differential effect by quartile of renter share, with the omitted category being neighborhoods with the lowest share of renter-occupied housing. While the unconditional approach might suggest complex variation in sensitivity to the price instrument, any differential impacts by renter share are eliminated with the inclusion of relevant controls.

### 3.4.3 Discussion

Given the estimates I find, it is natural to examine differences in water consumption between neighborhoods with above- and below-median shares of renter-occupied housing. As demonstrated in Figure 3.8, neighborhoods with lower shares of renter-occupied housing have a higher consumption distribution. While not shown, this relationship is also true conditional on household characteristics, such as irrigable

Figure 3.7: Sensitivity to Price Intervention by Renter-Share Quartile

Unconditional Approach　　Full Specification

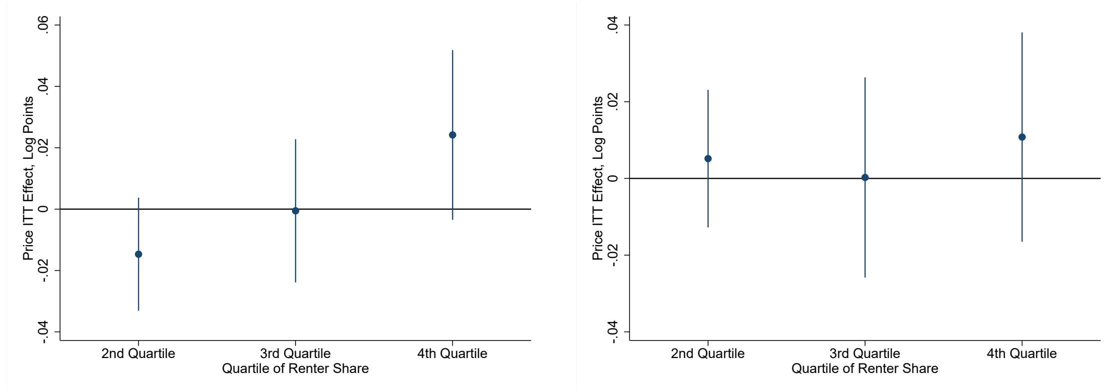


area, home values, and the number of occupants. One potential explanation for this fact is that the consumption patterns for a family of four are likely different from the consumption patterns for a house with four adult housemates. Renters may also have less interest in maintaining a green lawn, and the income necessary to live in a neighborhood with a certain set of home values may be lower for renters.

Figure 3.8: Differences in Consumption, Above- and Below-Median Renter Share

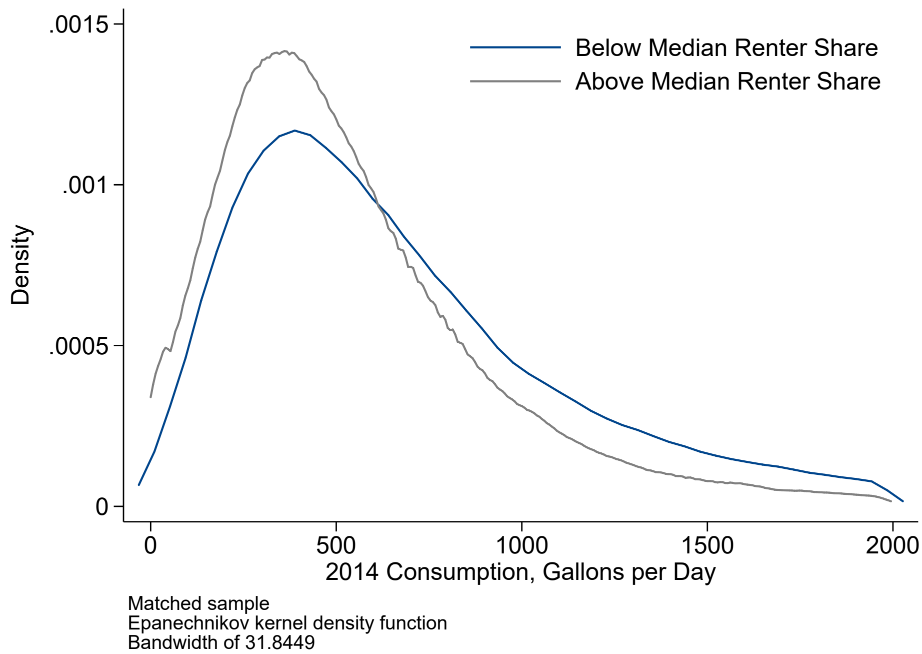


Despite this discrepancy, Figure 3.9 demonstrates that high- and low-renter-share neighborhoods do exhibit parallel trends, especially before the price and non-price interventions. The vertical lines in the figure represent the dates of initial implementation for each policy, with the price intervention occurring earlier and the non-price intervention occurring later. These estimates are pooled across both districts and collapsed by above- or below-median renter share and week of time. The size of the dots representing each observation are proportional to the number of bills observed in each district and week. Given that these estimates are not conditional on other characteristics, the evidence of a more limited response by renters is visible in the graph. Conditional on covariates, this evidence disappears.

Figure 3.9: Consumption Trends, Above- and Below-Median Renter Share

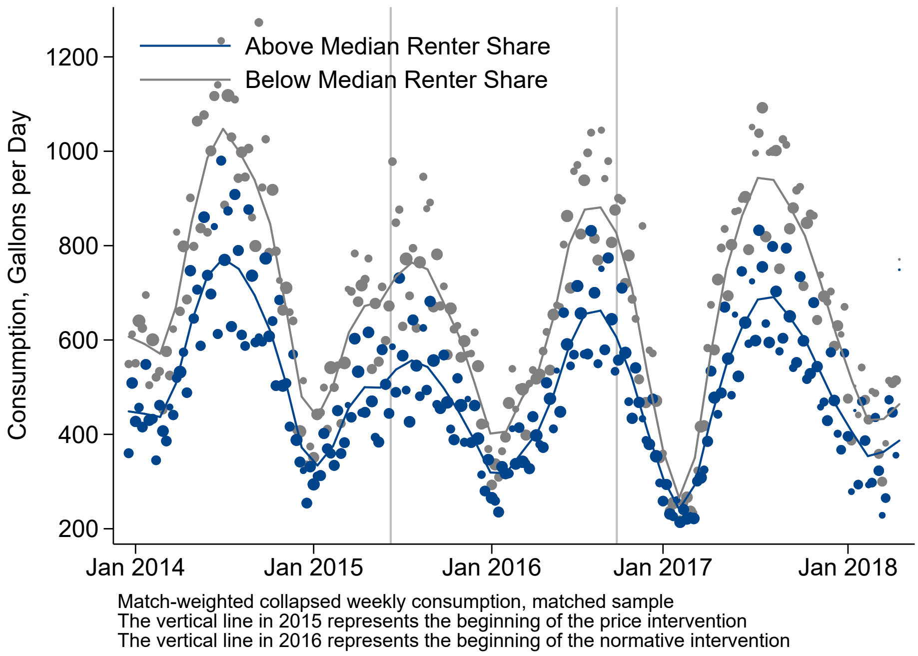


Given the similar response among households in low- and high-renter-share neighborhoods, it is also important to understand whether landlords pass costs and information on to renters. If renters pay the variable cost of water, they will have similar incentives to respond to a price intervention, even if they may be restricted in the ways they can respond. For example, renters may not be able to upgrade their appliances, but they will have incentives to reduce consumption through channels they control, such as clothes- and dishwashing. A brief investigation of rental advertisements suggests that many renters do not pay for water, but at least an important minority of rental situations do require tenants to pay for their water bill. Even when utilities are covered, renters may fear that high water use could be used as a reason to raise rental prices. It is outside the scope of this paper to assess whether landlords raised rents in response to high water bills, but it would also be difficult to estimate, given that renters ultimately appear to have conserved similarly. As a consequence, threatened rent increases for water costs were likely not implemented.

For the non-price intervention, the more salient concern is whether renters receive the HWR. While it is impossible to validate the identity of HWR recipients, 57.73 percent of reports were mailed to an address on file with the District. These reports were mailed to the accountholder, not the service address, and as a consequence, one might expect renters to be less likely to receive a physical copy of the HWRs that pertain to them. For the remaining HWRs sent via email, 70.93 percent of accounts receiving emailed reports did open at least one HWR email. To the extent that the results in Table 3.5 suggest a slightly weaker response, this difference in the intent-to-treat estimates could be driven by differences in whether a household receives a HWR. If a tenant is not the named accountholder, receiving an HWR would only occur if the landlord forwards the email or passes along the physical HWR. To the extent that I can measure this impact, accountholders with service addresses in high-renter-share neighborhoods are not less likely to open HWR emails. In probit regressions not shown here, Renter Share is not a statistically significant or meaningful predictor of whether an account opens an email containing an HWR. However, if renters are less likely to see a report, this would only bias my findings toward a more muted response by renters. I see only statistically insignificant evidence of this, which implies that the true difference in responsiveness is likely to be even smaller than estimated. In fact, it is notable that I can rule out large differences in responsiveness between homeowners and renters, given these substantial obstacles to passing through the interventions' incentives.

## 3.5 Conclusion

This research finds that renters and homeowners have differential responses to price and non-price interventions. However, these differential responses appear to be driven by pre-period consumption and other household characteristics, such as irrigable area, home values, and the number of occupants. Notably, there is no evidence that differences in property tenure drive this differential response, as my proxy for renting is both statistically insignificant and small in magnitude when controlling for these characteristics. Moreover, there is no evidence that accountholders differentially engage with HWRs based on whether the housing is renter- or owner-occupied.

The factors that drive the differential responses among households are not particularly suprising, given that these factors are highly correlated with the extent to which these policies bind. For the price intervention, renter-occupied households were much less likely to consume enough water to face the increased prices, because fewer renter-occupied houses were consuming beyond their combined indoor and outdoor budget. Notably, for the normative intervention, treated renter-occupied households were more likely to have lower water scores unconditional on house characteristics. Conditional on household characteristics, however, there is no discernible difference between renter-occupied and owner-occupied households in the likelihood of having a high (bad) water score, among treated households.

It is worth noting that additional characteristics of the household are also only weakly driving the effectiveness of these interventions. For the price intervention, high consumers conserved more than low consumers, while prices were less effective with wealthier neighborhoods. Notably, the price intervention was less effective for households with larger irrigable areas, conditional on consumption and other factors. This may be due to the construction of personalized water budgets and the fact that households with larger yards were allowed to consume more water at the lower tier 2 price. For the normative intervention, the effectiveness of the HWRs does not appear to be driven in a meaningful way by irrigable area or home values. However, pre-period consumption is a primary driver, with high consumers conserving more than low consumers. This effect is driven in part by the normative pressure being placed primarily on high consumers, especialy conditional on household characteristics.

These results are very promising for external validity and the promise of exporting effective policies to other districts. While my study does not consider a welfare analysis, both of the policies studied here proved effective at achieving their public policy goals. Moreover, this effectiveness appears to have limited dependence on demographic characteristics, which suggests that the same policies could potentially be similarly effective in other districts. This promising result should, however, be tempered by the caveat that my study intentionally sought to control other potentially important factors for external validity. Continued investigation may yield other important factors which are relatively constant across these districts, such as climate, landscaping, and cultural or social factors. Notably, these districts were chosen specifically for their similarity to one another, and these policies may be more or less effective for other water utilities which differ from both of the districts used in this study. Future work can address the extent to which these policies are externally valid across more diverse contexts.

# Appendix

Figure A.1: Example of a WaterSmart Home Water Report

## A.1 Supplemental Material for Chapter 1

### A.1.1 Policy Details

Figure A.2 provides an example bill presented during the Water School that violators could attend [SCMU 2015c]. The bill is exclusively for presentation purposes. It is important to note, however, that the red line in the usage plot remained even while fines were suspended. In other words, the allotment remained even when it was not enforced.

Figure A.2: West Bill

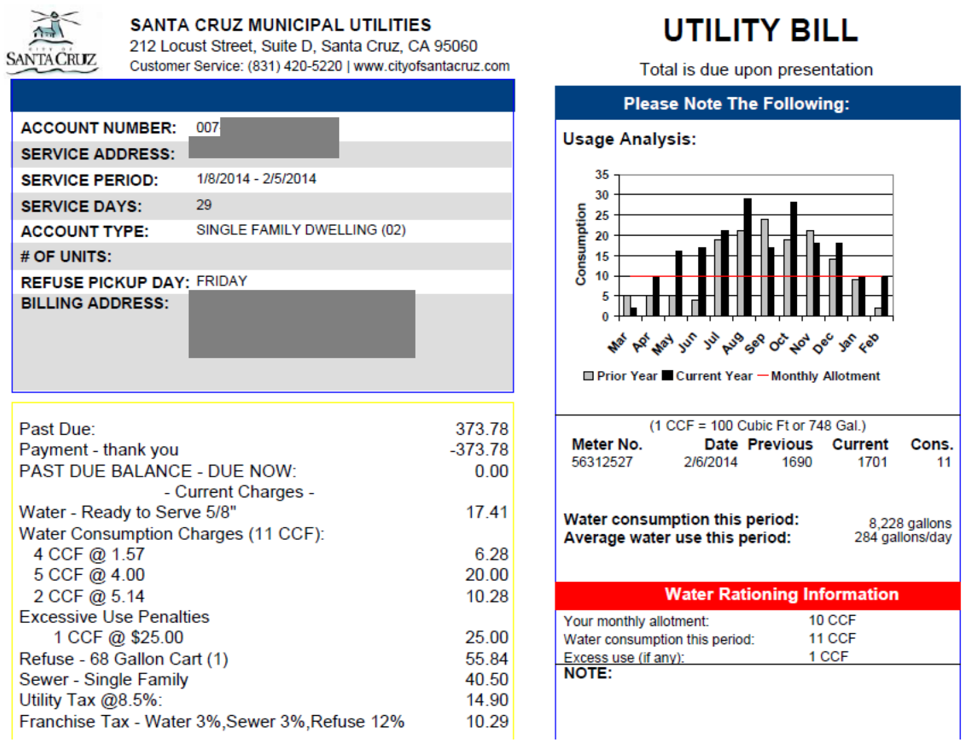

SANTA CRUZ MUNICIPAL UTILITIES
212 Locust Street, Suite D, Santa Cruz, CA 95060
Customer Service: (831) 420-5220 | www.cityofsantacruz.com

UTILITY BILL
Total is due upon presentation

| | |
|---|---|
| ACCOUNT NUMBER: | 007[redacted] |
| SERVICE ADDRESS: | [redacted] |
| SERVICE PERIOD: | 1/8/2014 - 2/5/2014 |
| SERVICE DAYS: | 29 |
| ACCOUNT TYPE: | SINGLE FAMILY DWELLING (02) |
| # OF UNITS: | |
| REFUSE PICKUP DAY: | FRIDAY |
| BILLING ADDRESS: | [redacted] |

| | |
|---|---|
| Past Due: | 373.78 |
| Payment - thank you | -373.78 |
| PAST DUE BALANCE - DUE NOW: | 0.00 |
| - Current Charges - | |
| Water - Ready to Serve 5/8" | 17.41 |
| Water Consumption Charges (11 CCF): | |
| 4 CCF @ 1.57 | 6.28 |
| 5 CCF @ 4.00 | 20.00 |
| 2 CCF @ 5.14 | 10.28 |
| Excessive Use Penalties | |
| 1 CCF @ $25.00 | 25.00 |
| Refuse - 68 Gallon Cart (1) | 55.84 |
| Sewer - Single Family | 40.50 |
| Utility Tax @8.5%: | 14.90 |
| Franchise Tax - Water 3%,Sewer 3%,Refuse 12% | 10.29 |

Please Note The Following:

Usage Analysis:



(1 CCF = 100 Cubic Ft or 748 Gal.)

| Meter No. | Date | Previous | Current | Cons. |
|---|---|---|---|---|
| 56312527 | 2/6/2014 | 1690 | 1701 | 11 |

| | |
|---|---|
| Water consumption this period: | 8,228 gallons |
| Average water use this period: | 284 gallons/day |

Water Rationing Information

| | |
|---|---|
| Your monthly allotment: | 10 CCF |
| Water consumption this period: | 11 CCF |
| Excess use (if any): | 1 CCF |

NOTE:

### A.1.2 Weather Controls

Table A.1 demonstrate robustness of the aggregate estimates and fine estimates to the inclusion and exclusion of weather controls. For precision, all estimates shown in the main sections of the paper include weather controls. Precipitation controls for the sum of total precipitation during the previous 30 days. Evapotranspiration and maximum and minimum temperature refer to the average of the respective measure over the previous 30 days.

Table A.1: Difference-in-Difference Robustness Tables, 2014

| | Outcome Var: Log of Daily Consumption (Gallons) | | | |
|---|---|---|---|---|
| | No Controls | Rain | Temperature | ET |
| Fine-Based Policy (2014) | -0.147*** | -0.145*** | -0.146*** | -0.154*** |
| | (0.0501) | (0.0447) | (0.0443) | (0.0427) |
| Precipitation | | 0.0731 | 0.0694 | 0.0772 |
| | | (0.0464) | (0.0498) | (0.0465) |
| Precipitation Squared | | -0.0640*** | -0.0826*** | -0.0981*** |
| | | (0.0152) | (0.0203) | (0.0204) |
| Maximum Temperature | | | -0.238** | -0.239** |
| | | | (0.0983) | (0.0931) |
| Max. Temp. Squared | | | 0.00164** | 0.00172*** |
| | | | (0.000664) | (0.000624) |
| Minimum Temperature | | | 0.0529 | 0.0585 |
| | | | (0.105) | (0.0978) |
| Min. Temp Squared | | | -0.000448 | -0.000471 |
| | | | (0.00104) | (0.000970) |
| Evapotranspiration | | | | -7.251*** |
| | | | | (2.724) |
| Evapotrans. Squared | | | | 15.68* |
| | | | | (8.355) |
| Fixed Effects | Household, Week of Time | | | |
| Clustered SE | Household, Week of Time | | | |
| $R^2$ | 0.714 | 0.715 | 0.715 | 0.715 |
| Obs. | 123,197 | 123,197 | 123,197 | 123,197 |

Estimation based on the 12 months before and five months after implementation the month of implementation and the month of implementation. Cluster-robust standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

### A.1.3 Sensitivity of Effects across the Distribution

Figure A.3 demonstrates that the finding from Panel A of Figure 1.8 holds whether I examine only accounts within approximately 1 mile of the border or the entirety of each district. Figure A.4 similarly demonstrates the robustness of Panel B of Figure 1.8 to various border samples. Samples closer to the boundary appear to suffer from challenges brought on by the small sample size.

In addition, Figure A.5 illustrates that Panels A and B of Figure 1.8 are robust to the inclusion of an additional six months of pre-period data, despite the inclusion of bimonthly billing data. Figure A.6 demonstrates that low-use households with a higher historical variance do not respond more strongly than low-use households with a lower historical variance, using 11 months of data prior to the month of implementation.

Figure A.3: Treatment Effect by Prior Bill Consumption, Sample Robustness

Border Sample (Within 2 Miles)

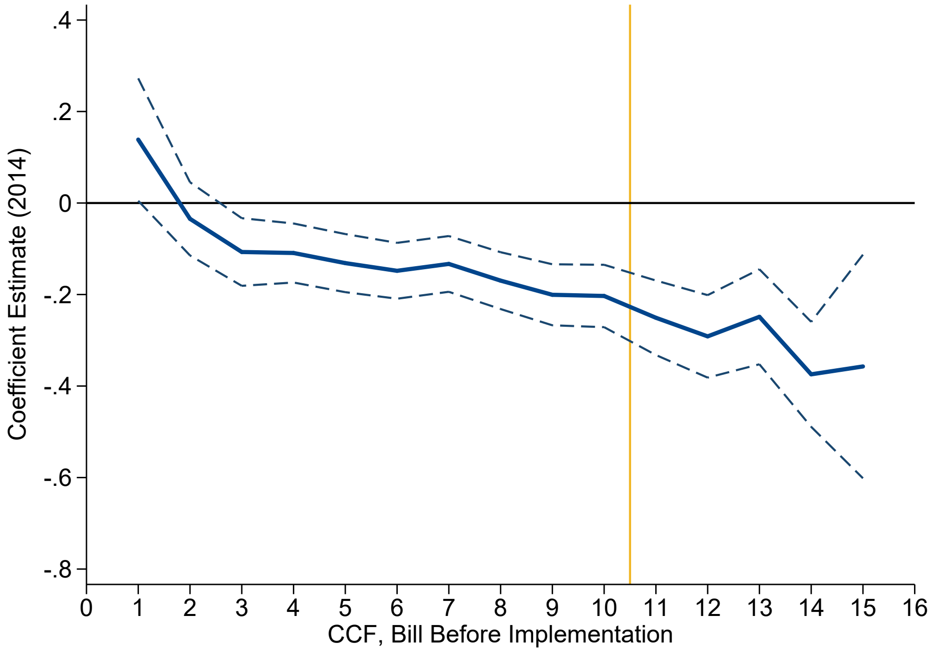


Within 0.5 Miles

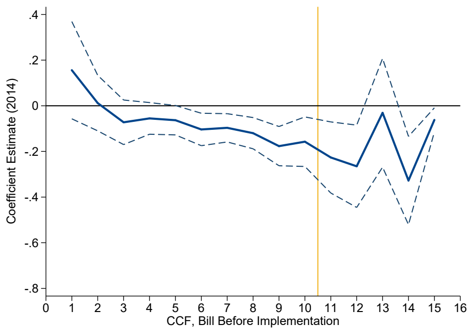


Within 1 Mile

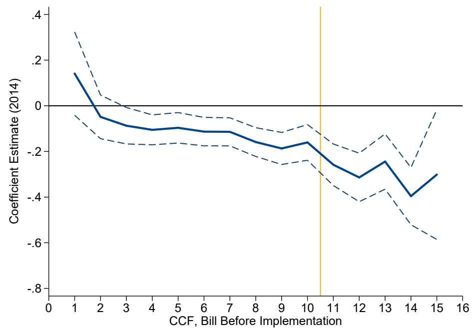


Within 1.5 Miles

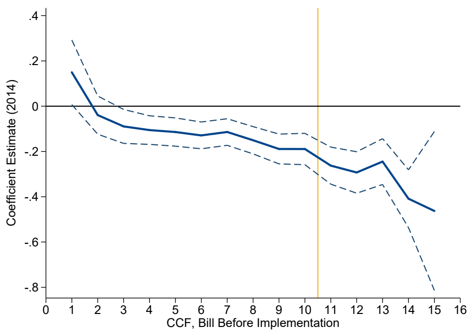


Whole Sample

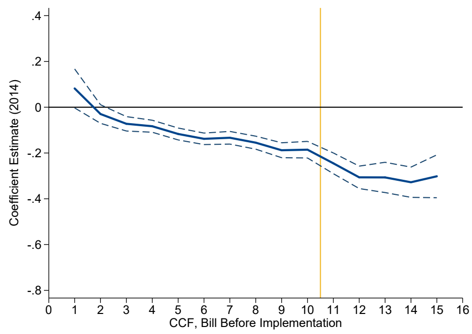

Figure A.4: Treatment Effect by Pre-Policy Maximum Consumption, Sample Robustness

Border Sample (Within 2 Miles)

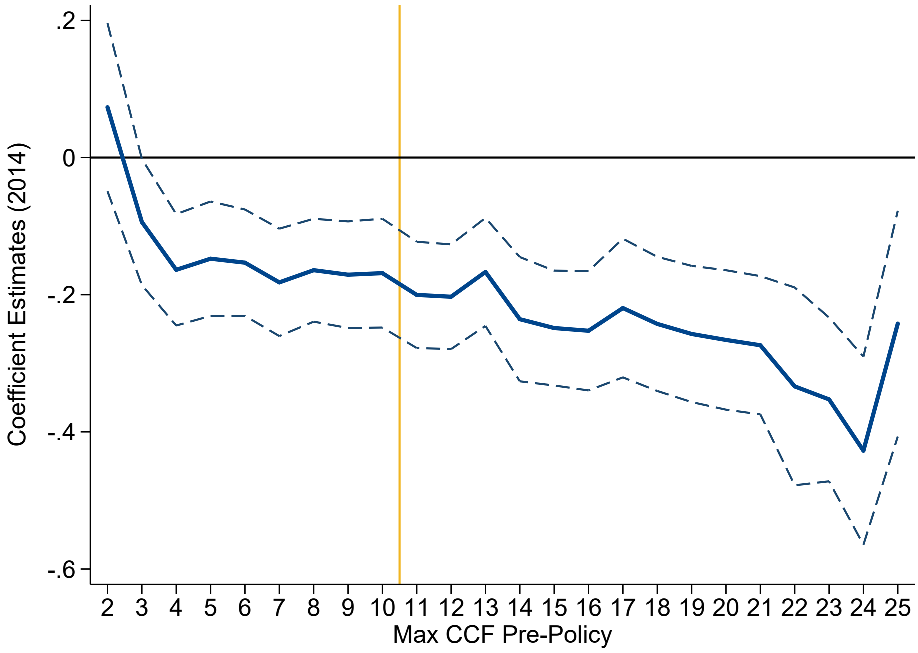


Within 0.5 Miles

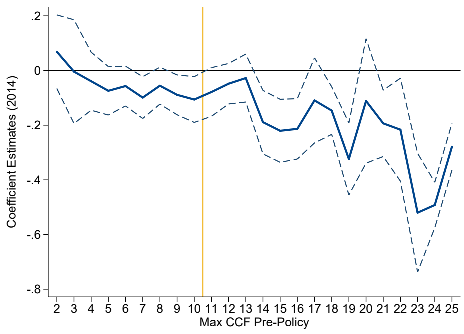


Within 1 Mile

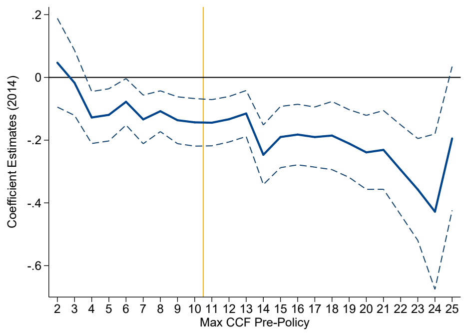


Within 1.5 Miles

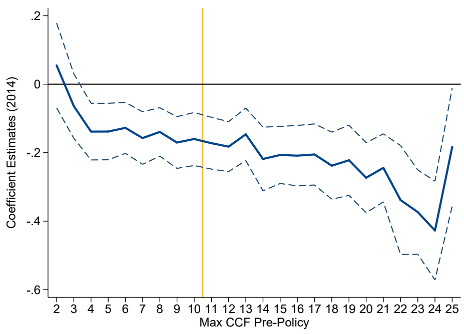


Whole Sample

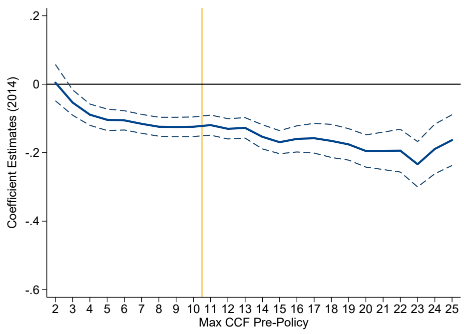

Figure A.5: Effect of Fines by Pre-Implementation Consumption Level, Robustness to Additional Pre-Period Data

Five months before implementation

[A] Prior Bill

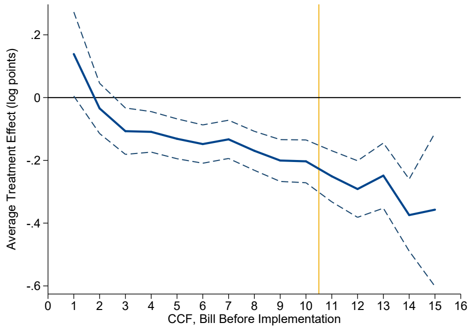


[B] Max. Pre-Period

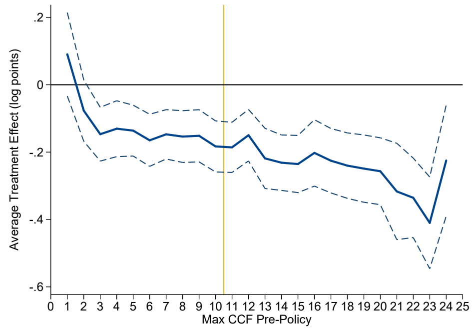


11 months before implementation

[C] Prior Bill

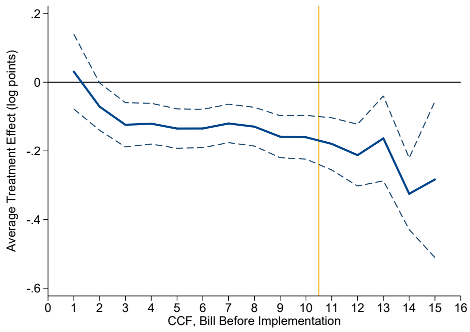


[D] Max. Pre-Period

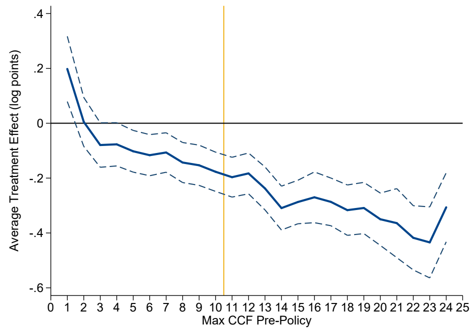


### A.1.4 Full District Estimates

As discussed in Section 1.1, this paper relies on the assumption that sorting across an arbitrary district boundary is as good as random assignment. However, we could also consider a difference-in-difference approach using the entirety of each district. Figure A.7 demonstrates that the parallel trends assumption appears to hold not only across the border, but also more generally between the two districts. Consequently, one might wish to understand whether the effects estimated using the border sample might be replicated within the broader sample. As shown in Table A.2, these estimates largely hold. The effects actually seem slightly higher using the full sample, but this may reflect differential responses among segments of the West district population or bias from unobservables.

Figure A.6: Response by Exposure and Variance, Robustness to Additional Pre-Period Data

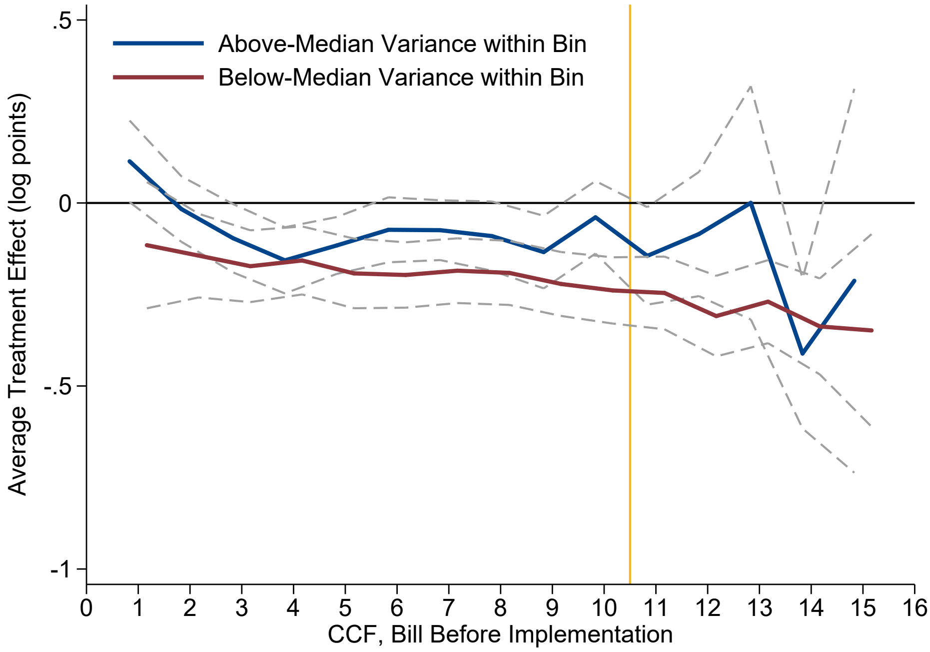


Figure A.7: Single-Family Residential Consumption by District

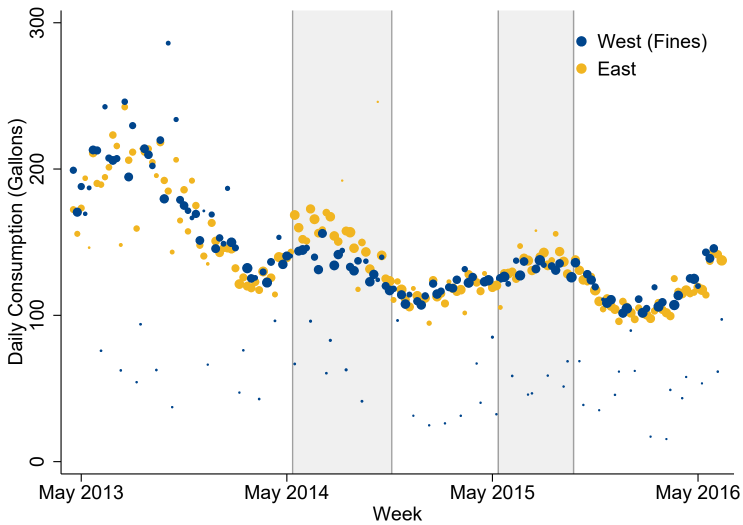

Table A.2: Full District Estimates

| | Outcome Var: Log of Daily Consumption (Gallons) | | | |
|---|---|---|---|---|
| | 2014 | | 2015 | |
| | 1 | 2 | 3 | 4 |
| *Fine-Based Policy ATE* | -0.126*** | -0.123*** | -0.0317*** | -0.0306** |
| | (0.0184) | (0.0190) | (0.0106) | (0.0117) |
| East Mean - Post | 140.28 | | 126.81 | |
| (Standard Error) | (91.36) | | (82.63) | |
| Controls | Precipitation, Max Temp, Min Temp, Evapotranspiration | | | |
| Fixed Effects | Household, Week of Time | | | |
| Cluster Variables | Household, Week of Time | | | |
| $R^2$ | 0.711 | 0.708 | 0.718 | 0.713 |
| Obs. | 318,093 | 287,115 | 322,678 | 292,017 |

Sample restricted to five months before and after the month of implementation (columns 2, 4), plus the month of implementation in columns 1 and 2. Lowest one percent and highest one percent of consumption observations dropped. Cluster-robust standard errors in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

### A.1.5 Effects across the Full Distribution

Figure A.8 illustrates the effect of the policy on the full distribution. The cumulative density supports the evidence of conservation along the majority of the distribution.

Figure A.8: Impacts on the Full Distribution

Summer 2013 v Summer 2014, Full Sample

Density

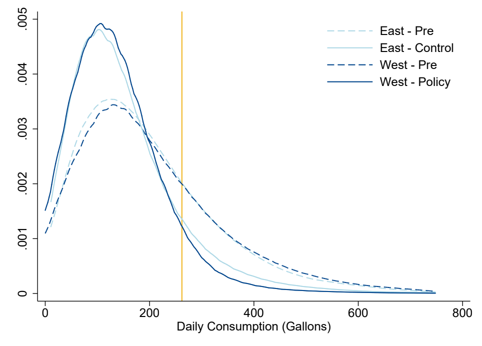


Cumulative Density

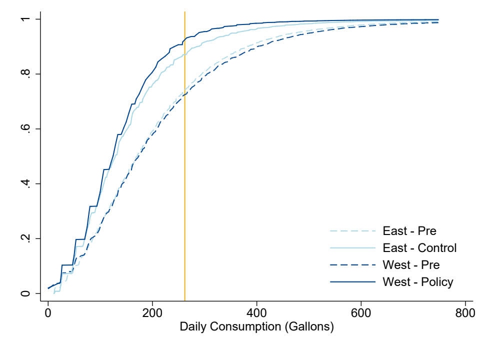


Five Months Before and After Policy Implementation, Full Sample

Density

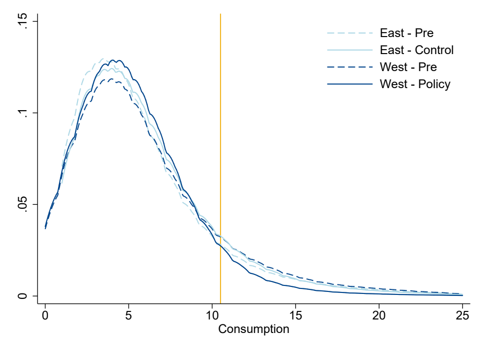


Cumulative Density

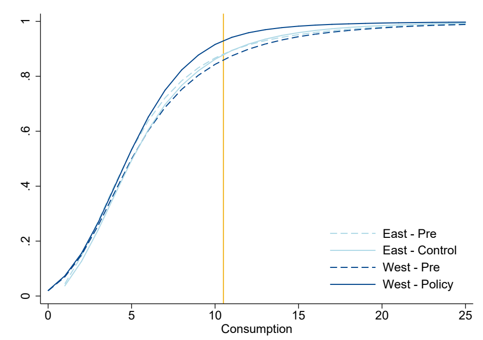


### A.1.6 Historical Comparison

In 1990, the West district applied a similar water rationing program. Using monthly aggregate data for 30 years and business consumption as a control, I find similar estimates for 1990 and 2014, as shown in Figure A.9. Notably, these estimates are from regressions controlling for the Palmer Drought Severity Index of the area, by month. However, the drought index becomes statistically insignificant when including indicator variables for these three policies, as well as an indicator for high-use penalties in 1991 and 1992. Given the less precise counterfactual of using business consumption, the effects are both imprecise and slightly overstated for 2014 and 2015. As a result, one might expect the true reductions achieved in 1990 to similarly be slightly smaller.

Figure A.9: Aggregate Effects of Mandatory Reductions, 1990 vs 2014/2015

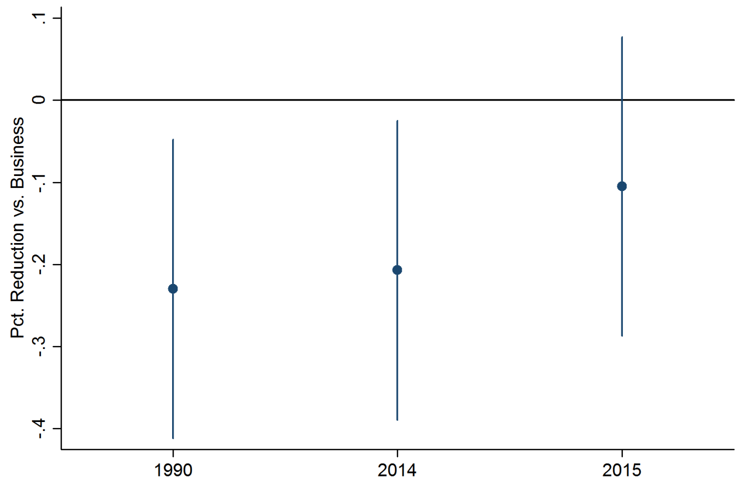


In order to compare estimates to 1990, I use a dataset of monthly aggregate consumption by single-family residential, multi-family residential, business, industrial, and other users, recorded over approximately 33 years. I first define three indicator variables, one for each period of mandatory drought restrictions - 1990, 2014, and 2015 - during the six peak months of the year - May through October - and an indicator variable for excessive use penalties in 1991 and 1992. In addition to fixed effects for year and month-of-year, I control for the monthly Palmer Drought Index for the climate division of California that includes the West service area. The dependent variable is the log of consumption per single-family residential account minus the log of consumption per business account, as below:

$$
\begin{aligned}
\text{log-diff Cons per Account}_{m,y} =& \delta_{1990} 1 \left(m \in \{5, \ldots, 10\}, y = 1990\right) \\
& + \delta_{2014} 1 \left(m \in \{5, \ldots, 10\}, y = 2014\right) \\
& + \delta_{2015} 1 \left(m \in \{5, \ldots, 10\}, y = 2015\right) \\
& + \pi_{pricing} 1 \left(m \in \{5, \ldots, 10\}, y \in \{1991, 1992\}\right) \\
& + \mu_m + \gamma_y + \varepsilon_{m,y}
\end{aligned}
$$

Where the dependent variable is calculated as follows:

$$
\begin{aligned}
\text{log-diff Cons per Account}_{m,y} =& \log\left(\text{Cons}_{SFR,m,y}/\text{No. of Accts}_{SFR,m,y}\right) \\
& - \log\left(\text{Cons}_{Bus,m,y}/\text{No. of Accts}_{Bus,m,y}\right)
\end{aligned}
$$

and $m$ denotes the month of year and $y$ denotes the year. The $\pi_{pricing}$ estimate is of tangential interest, but is not statistically different from zero. Figure A.9 provides graphical estimates of $\hat{\delta}_{1990}$, $\hat{\delta}_{2014}$, and $\hat{\delta}_{2015}$.

### A.1.7 Expenditures and Bunching

It is worth noting that, while the fine policy eliminated the seasonality in consumption of water in the West district, it did not eliminate the seasonality in expenditure. Figure A.10 plots average expenditure by households in each district, including all fixed and variable charges as the pricing structure evolves over time. Since many fines were forgiven, the blue fitted line in the figure is an over-estimate of actual expenditure, but it captures the fact that this policy had a potent pricing channel, especially for high-consumers. With inelastic demand, as prices rise, expenditure rises, so the behavior is not unexpected. Taken together with the response across the distribution, Figure A.10 suggests that high consumers did respond to the price component of the fine by lowering their consumption.

Figure A.10: Effects on Expenditure

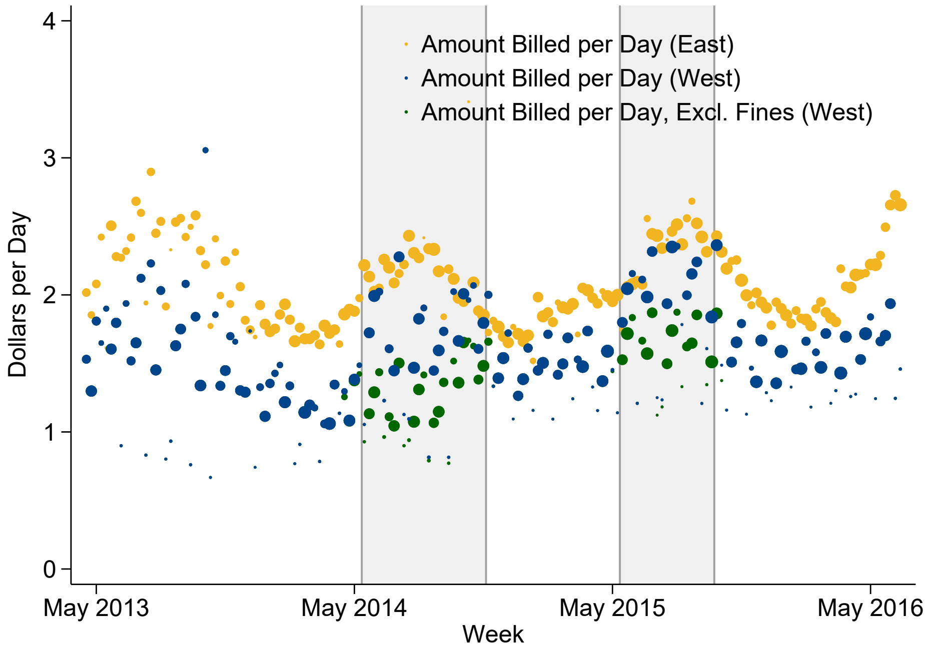


On the other hand, there is no evidence that reductions were made on the basis of a marginal price response. As in Ito [2013], I find no evidence of bunching at the marginal price discontinuity, even though the discontinuity is even larger in both absolute and percentage terms. Figure A.11 demonstrates that the distribution of consumption was very smooth across the penalty-induced marginal price jump throughout the first year of the policy.

Figure A.11: No Evidence of Bunching During the Policy Period, 2014

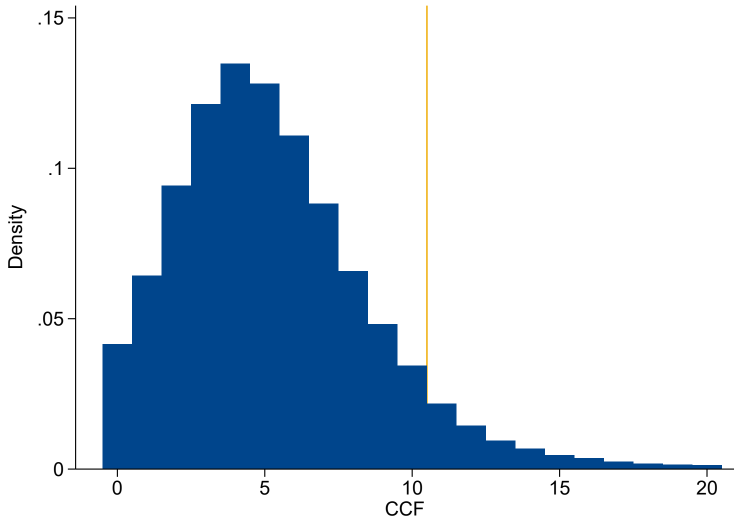

## A.2 Supplemental Materials for Chapter 2

Figure A.12: Historical perspective on severity of drought in southern California

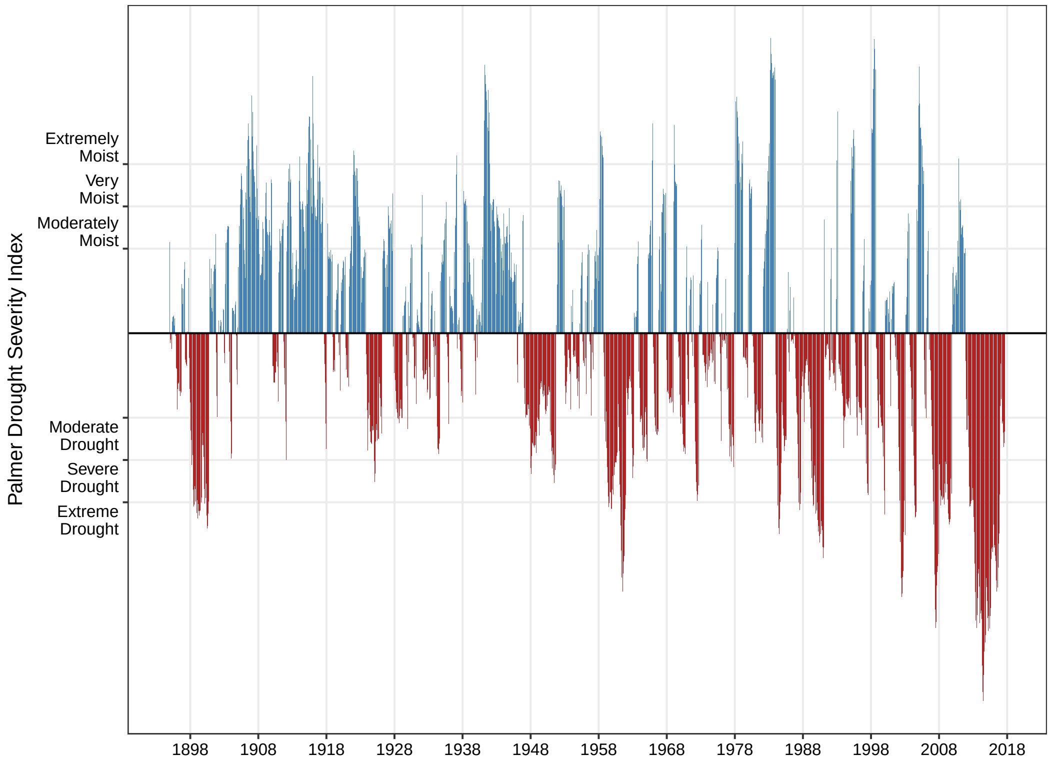


Notes: Figure A.12 plots historical monthly observed drought severity on the Palmer Drought Severity Index for the hydrological region of Coastal Southern California. Our study period during 2015 lies within the most severe drought on record for the region, but lengthy periods of drought are common.

Figure A.13: Example of an automated irrigation violation notice

CITY OF BURBANK
Water and Power
Since 1913
*Always There for You!*

July 7, 2015

**RE: Violation of Irrigation Requirements at St**

Dear Customer,

Our system has detected that you are watering your landscape more than two days per week, for longer than 15 minutes per irrigation station, or on incorrect days, all of which are in violation of Burbank's Sustainable Water Use Ordinance.

Burbank is in Stage III of the Ordinance and landscape watering during April through October is limited to **Tuesdays and Saturdays, before 9am or after 6pm,** and **no more than 15 minutes per watering station**. Starting November 1, watering will be limited to Saturdays.

Immediately, please take the necessary steps to comply with the drought requirements, including appropriately programming your irrigation controller if you have one. If you use a gardener or landscape professional, make sure that they are aware that only Tuesdays and Saturdays are allowed for watering in Burbank, either before 9am or after 6pm.

There is no need to contact BWP. Simply correct your irrigation schedule as quickly as possible. Provided your changes are made within the week, our system will detect your corrections and discontinue sending you notices. However, if you continue watering your landscape more than two days per week, you will receive a second notice. If you remain in violation of the Ordinance thereafter you will be fined $100. Any additional fines will be $200 for the second violation and $500 for the third or more violations.

Governor Brown has mandated that Burbank reduce water usage by one billion gallons by February 2016 or face fines of $10,000 per day. We thank you in advance for doing your part to meet this mandate and preserve the water supply to get us through this historic drought.

Sincerely,

Burbank Water and Power

---

Burbank Water and Power
164 West Magnolia Blvd., P.O. Box 631, Burbank, CA 91503-0631

Figure A.14: Distribution of households along the running variable

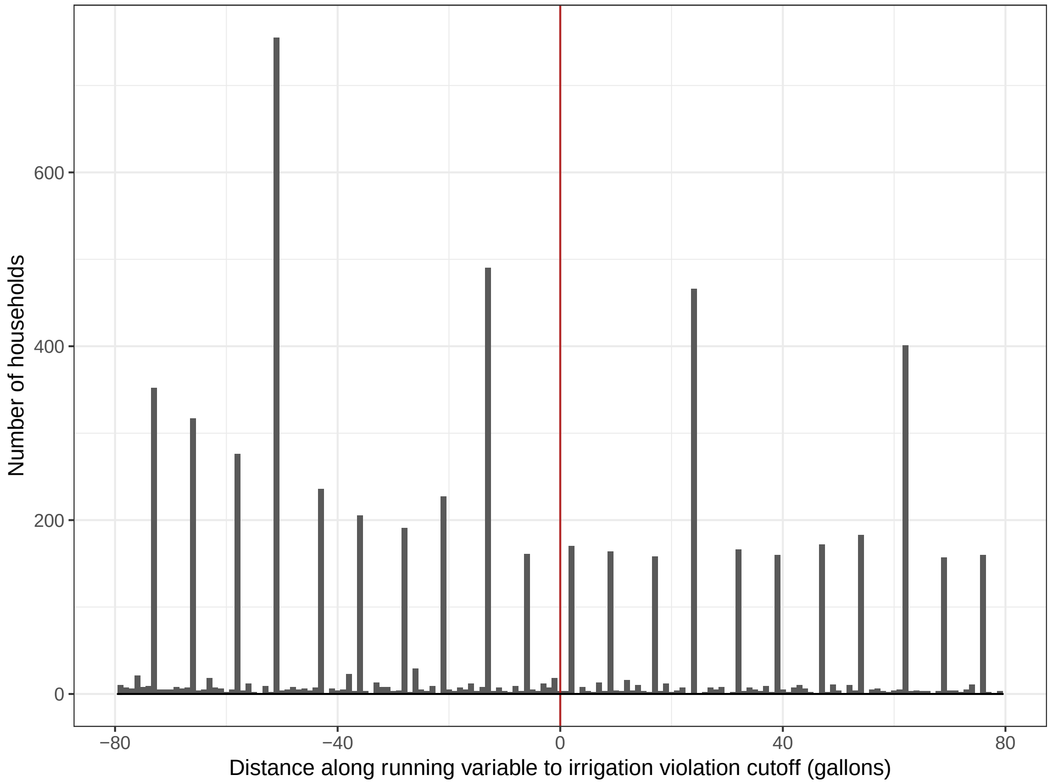


Notes: Figure A.14 plots the distribution of households along the running variable used in our regression discontinuity design, providing a graphical version of the McCrary [2008] bunching test for manipulation with respect to treatment assignment. Due to heterogeneity in the granularity of measurement for included water meters, there is significantly more mass at cubic foot (7.48 gallons) and five cubic feet increments. Importantly, there is no evidence of any excess distributional mass in the region surrounding the cutoff used for determining automated irrigation violation notices.

Figure A.15: Identification check for pre-treatment consumption along the running variable

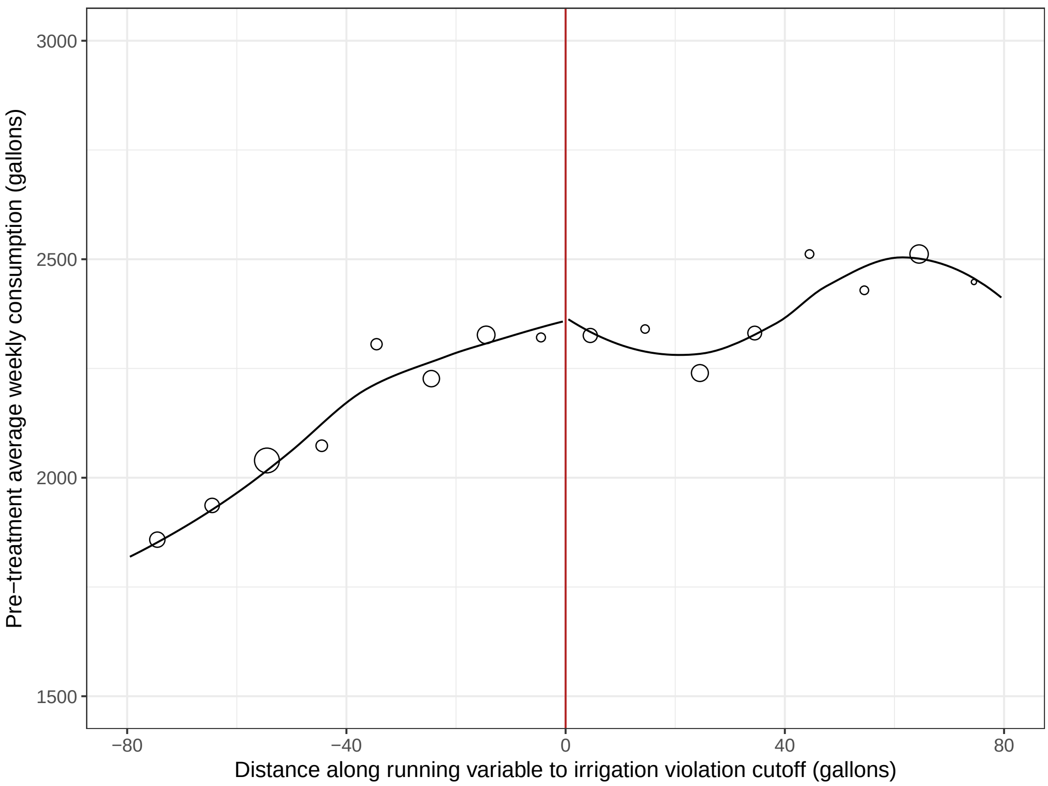


Notes: Figure A.15 plots local averages for weekly water consumption during May 2014 through April 2015, the full year prior to both the automated enforcement and social comparison treatments. For clarity, the running variable uses 10 gallon bins. The size of the markers corresponds to the number of households included in the local averages. The LOESS curves shown are fit to the underlying microdata separately on each side of the threshold. This identification check shows that pre-treatment water consumption is smooth across the cutoff used for determining automated irrigation violation notices.

Table A.3: Randomization balance checks in 80 gallon bandwidth around RD cutoff

| | Group means | | t-tests | |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| Covariate | Control | Experimental | Difference | p-value |
| Number of households | 1104 | 5222 | | |
| Sent WaterSmart HWR | 0 | 0.9992 | | |
| Algorithmic violation | 0.1395 | 0.152 | 0.0125 | 0.28 |
| Prior water violation | 0.03261 | 0.03658 | 0.00397 | 0.5 |
| Lot size (SqFt) | 6861 | 6906 | 45 | 0.53 |
| Irrigable area (SqFt) | 3594 | 3576 | -18 | 0.68 |
| House size (SqFt) | 1532 | 1555 | 23 | 0.27 |
| Year built | 1944 | 1944 | 0 | 0.87 |
| Number of floors | 1.057 | 1.055 | -0.002 | 0.75 |
| Number of bedrooms | 2.854 | 2.892 | 0.038 | 0.18 |
| Number of bathrooms | 1.846 | 1.871 | 0.025 | 0.36 |
| Weekly water gallons | 2191 | 2236 | 45 | 0.25 |

Notes: Table A.3 shows statistics by WaterSmart Home Water Reports (HWR) treatment arm for household-level covariates for the subset of households near to the RD cutoff. See notes for Table 2.1 for additional details.

## A.3 Supplemental Materials for Chapter 3

### A.3.1 Full Sample - Price Intervention

Figure A.16 shows the strength of the parallel trends assumption for the full sample. While reasonable, the matched sample provides a far better fit. Moreover, Table A.4 demonstrates that the disparity between treated and control households

Figure A.16: Parallel Trends Assumption, Full Sample - Price Intervention

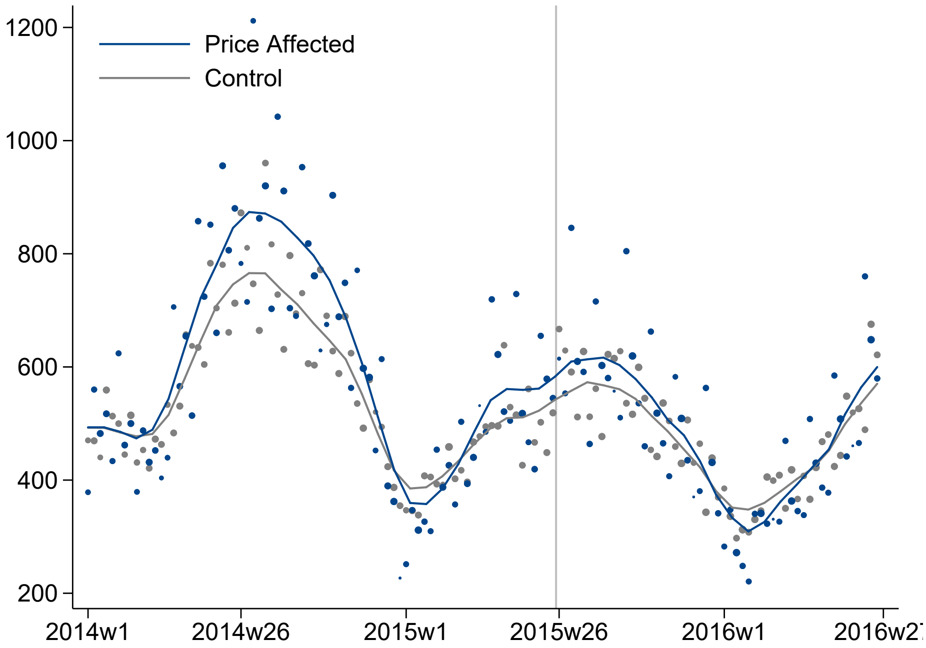


is larger and more significant among the full sample of both districts. Nevertheless, estimates in both Table A.6 and A.5 reflect similar trends to those in the matched sample. There is evidence of differential response, but this differential response appears to again be driven by differences in pre-period consumption, which is highly correlated with the extent to which the price increase binds and affects a given household.

Table A.4: Balance Table, Full Sample - Price Intervention

| | Control | Treated | Diff |
|---|---|---|---|
| | (SE) | (SE) | (p-value) |
| Renter Share | 0.384 | 0.181 | -0.203 |
| | (0.236) | (0.146) | (0.000)*** |
| Irrigable Area | 7,057.519 | 9,581.682 | 2,524.163 |
| | (20,080.273) | (14,866.735) | (0.000)*** |
| Median Home Value | $287,577.31 | $396,534.47 | $108,957.16 |
| | (104,785.891) | (68,643.789) | (0.000)*** |
| Number of Occupants | 2.986 | 4.075 | 1.090 |
| | (1.301) | (1.523) | (0.000)*** |
| Observations | 55,573 | 24,294 | 79,867 |

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.5: Sensitivity to Price Intervention, Logs, Full Sample

| | (1) | (2) | (3) | (4) | (5) |
|---|---|---|---|---|---|
| Price Increase | -0.0540*** | -0.0765*** | | 0.173*** | |
| | (0.0131) | (0.0182) | | (0.0289) | |
| **Price×Renter Share** | | 0.128*** | 0.127** | -0.0248 | -0.00388 |
| | | (0.0468) | (0.0500) | (0.0433) | (0.0511) |
| Price×log (Irr. Area) | | | 0.00162 | | 0.0333*** |
| | | | (0.0101) | | (0.00519) |
| Price×log (Med. Home Value) | | | -0.0133 | | 0.0184 |
| | | | (0.0260) | | (0.0229) |
| Occupant Bin Interactions | - | - | Yes | - | Yes |
| 2014 Cons. Bin Interactions | - | - | - | Yes | Yes |
| FE | Household ID, Week of Time | | | | |
| Cl. Var. | Household ID, Week of Time | | | | |
| $R^2$ | 0.771 | 0.771 | 0.771 | 0.772 | 0.772 |
| Obs. | 2,186,607 | 2,176,309 | 2,103,012 | 2,097,435 | 2,027,166 |

Highest one percent and lowest one percent of observations dropped.

Cluster-robust standard errors in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.6: Sensitivity to Price Intervention, Levels, Full Sample

| | (1) | (2) | (3) | (4) | (5) |
|---|---|---|---|---|---|
| Price Increase | -39.37*** | -66.09*** | | 97.57*** | |
| | (10.99) | (15.80) | | (21.08) | |
| **Price×Renter Share** | | 146.8*** | 101.3** | -14.74 | -12.82 |
| | | (39.38) | (38.97) | (25.98) | (34.35) |
| Price×log(Irr. Area) | | | -47.78** | | 19.14*** |
| | | | (18.44) | | (4.752) |
| Price×log(Med. Home Value) | | | -46.11 | | 4.884 |
| | | | (31.09) | | (25.57) |
| Occupant Bin Interactions | - | - | Yes | - | Yes |
| 2014 Cons. Bin Interactions | - | - | - | Yes | Yes |
| FE | Household ID, Week of Time | | | | |
| Cl. Var. | Household ID, Week of Time | | | | |
| $R^2$ | 0.748 | 0.748 | 0.749 | 0.754 | 0.754 |
| Obs. | 2,211,961 | 2,201,599 | 2,127,417 | 2,113,391 | 2,042,628 |

Highest one percent and lowest one percent of observations dropped.

Cluster-robust standard errors in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

# Bibliography


George A. Akerlof and William T. Dickens. The economic consequences of cognitive dissonance. *The American Economic Review*, 72(3):307–319, June 1982.

Hunt Allcott. Social norms and energy conservation. *Journal of Public Economics*, 95(9-10):1082–1095, October 2011.

Hunt Allcott and Judd B. Kessler. The welfare effects of nudges: A case study of energy use social comparisons. *American Economic Journal: Applied Economics*, 11(1):236–276, January 2019.

Hunt Allcott and Todd Rogers. The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation. *The American Economic Review*, 104(10):3003–3037, October 2014.

Dan Ariely, Anat Bracha, and Stephan Meier. Doing good or doing well? Image motivation and monetary incentives in behaving prosocially. *The American Economic Review*, 99(1):544–555, March 2009.

Omar I. Asensio and Magali A. Delmas. Nonprice incentives and energy conservation. *Proceedings of the National Academy of Sciences*, 112(6):E510–E515, 2015. doi: 10.1073/pnas.1401880112. URL http://www.pnas.org/content/112/6/E510.abstract.

Maximillian Auffhammer and Edward Rubin. Natural gas price elasticities and optimal cost recovery under consumer heterogeneity: Evidence from 300 million natural gas bills. EI @ Haas WP 287, 2018.

Ian Ayres, Sophie Raseman, and Alice Shih. Evidence from two large field experiments that peer comparison feedback can reduce residential energy usage. *The Journal of Law, Economics, and Organization*, 29(5):992–1022, October 2013.

Jonathan E. Baker. Subsidies for succulents: Evaluating the Las Vegas Cash-for-Grass rebate program. Harvard Environmental Economics Program Discussion Paper 17-74, June 2017.

Gary S. Becker. Nobel lecture: The economic way of looking at behavior. *Journal of Political Economy*, 101(3):385–409, June 1993.

Roland Bénabou and Jean Tirole. Incentives and prosocial behavior. *The American Economic Review*, 96(5):1652–1678, December 2006.

Maria Bernedo, Paul J. Ferraro, and Michael K. Price. The persistent impacts of norm-based messaging and their implications for water conservation. *Journal of Consumer Policy*, 37(3):437–452, September 2014.

Syon P. Bhanot. Isolating the effect of injunctive norms on conservation behavior: New evidence from a field experiment in California. *Organizational Behavior and Human Decision Processes*, Forthcoming.

Sandra E. Black. Do better schools matter? parental valuation of elementary education. *The Quarterly Journal of Economics*, 114(2):577–599, May 1999.

Samuel Bowles and Sandra Polanía-Reyes. Economic incentives and social preferences: Substitutes or complements? *Journal of Economic Literature*, 50(2):368–425, June 2012.

Christa Brelsford and Joshua K. Abbott. How smart are "Water Smart Landscapes"? arXiv.org Working Paper, March 2018.

Daniel A. Brent, Joseph H. Cook, and Skylar Olsen. Social comparisons, household water use, and participation in utility conservation programs: Evidence from three randomized trials. *Journal of the Association of Environmental and Resource Economists*, 2(4):597–627, December 2015.

Daniel A. Brent, Corey Lott, Michael Taylor, Joseph H. Cook, Kim Rollins, and Shawn Stoddard. Are normative appeals moral taxes? Evidence from a field experiment on water conservation. LSU Working Paper 2017-07, June 2017.

Dylan Brewer. Equilibrium sorting and moral hazard in residential energy contracts. Working Paper, 2019.

Wiliam H. Bruvold. Residential response to urban drought in central California. *Water Resources Research*, 15(6):1297–1304, December 1979.

David Card and Alan B. Krueger. Minimum wages and employment: A case study of the fast-food industry in new jersey and pennsylvania. *The American Economic Review*, 84(4):772–793, September 1994.

Marco Castillo, Ragan Petrie, and Clarence Wardell. Fundraising through online social networks: A field experiment on peer-to-peer solicitation. *Journal of Public Economics*, 114:29–35, June 2014.

Anita Castledine, Klaus Moeltner, Michael K. Price, and Shawn Stoddard. Free to choose: Promoting conservation by relaxing outdoor watering restrictions. *Journal of Economic Behavior and Organization*, 107(A):324–343, November 2014.

Jerald Emmet Christiansen. *Irrigation by Sprinkling*, volume 670. University of California, College of Agriculture: Agricultural Experiment Station, Berkeley, California, 1942.

Anthony Cilluffo, Abigail Geiger, and Richard Fry. More u.s. households are renting than at any point in 50 years. Facttank, Pew Research Center, July 2017.

Sergio Correia. Singletons, Cluster-Robust Standard Errors and Fixed Effects: A Bad Mix. *Working Paper*, 2015.

Sergio Correia. A Feasible Estimator for Linear Models with Multi-Way Fixed Effects. *Working Paper*, 2016.

Dora L. Costa and Matthew E. Kahn. Energy conservation “nudges” and environmentalist ideology: Evidence from a randomized residential electricity field experiment. *Journal of the European Economic Association*, 11(3):680–702, June 2013.

Suagato Datta, Matthew Darling, Karine Lorenzana, Oscar Calvo Gonzalez, Juan Jose Miranda, and Laura de Castro Zoratto. A behavioral approach to water conservation: Evidence from a randomized evaluation in costa rica. ideas42 and the World Bank, 2015.

Lucas W. Davis. *The Design and Implementation of U.S. Climate Policy*, chapter Evaluating the Slow Adoption of Energy Efficient Investments: Are Renters Less Likely to Have Energy Efficient Appliances?, pages 301– 316. University of Chicago Press, 2012.

Paula C.P. de Sousa Mills. California adopts measures targeting multifamily residential buildings for water conservation, October 2016.

Mikael Elinder, Sebastian Escobar, and Ingel Petré. Consequences of a price incentive on free riding and electric energy consumption. *Proceedings of the National Academy of Sciences of the United States of America*, 2017. ISSN 0027-8424.

Ahmad Faruqui and Sanem Sergici. Dynamic pricing of electricity in the mid-Atlantic region: Econometric results from the Baltimore gas and electric company experiment. *Journal of Regulatory Economics*, 40(1):82–109, 2011. ISSN 0922680X. doi: 10.1007/s11149-011-9152-5.

Michael Faure and Laarni Escresa. *Production of Legal Rules*, chapter Social Stigma, pages 205–227. Maastricht Faculty of Law Working Paper, Cheltenham, 2011.

Paul J. Ferraro and Michael K. Price. Using nonpecuniary strategies to influence behavior: Evidence from a large-scale field experiment. *The Review of Economics and Statistics*, 95(1):64–73, March 2013.

Paul J. Ferraro, Juan José Miranda, and Michael K. Price. The persistence of treatment effects with norm-based policy instruments: Evidence from a randomized environmental policy experiment. *The American Economic Review: Papers & Proceedings*, 101(3):318–322, May 2011.

Kelly Fielding, Anneliese Spinks, Sally Russell, Aditi Mankad, Rod McCrea, and John Gardner. Water demand management: Interventions to reduce household water use. techreport 94, Urban Water Security Research Alliance, 2012.

Bruno S. Frey and Stephan Meier. Social comparisons and pro-social behavior: Testing "conditional cooperation" in a field experiment. *The American Economic Review*, 94(5):1717–1722, December 2004.

Alan S. Gerber, Donald P. Green, and Christopher W. Larimer. Social pressure and voter turnout: Evidence from a large-scale field experiment. *American Political Science Review*, 102(1):33–48, February 2008.

James M. Gillan. Dynamic pricing, attention, and automation: Evidence from a field experiment in electricity consumption. Energy Institute at Haas WP284, March 2018.

Uri Gneezy and Aldo Rustichini. A fine is a price. *The Journal of Legal Studies*, 29 (1):1–17, January 2000.

Uri Gneezy, Stephan Meier, and Pedro Rey-Biel. When and why incentives (don't) work to modify behavior. *Journal of Economic Perspectives*, 25(4):191–210, Fall 2011.

Phil Gomez. Water restrictions in Santa Cruz lifted, 2015. KSBW News. http://www.ksbw.com/news/watrer-restrictions-in-santa-cruz-lifted/36096802.

R. Quentin Grafton and Michael B. Ward. Prices versus rationing: Marshallian surplus and mandatory water restrictions. *Economic Record*, 84(S):S57–S65, September 2008.

Jerry A. Hausman. Individual discount rates and the purchase and utilization of energy-using durables. *The Bell Journal of Economics*, 10(1):33–54, 1979.

David Hensher, Nina Shore, and Kenneth Train. Water Supply Security and Willingness to Pay to Avoid Drought Restrictions. *Economic Record*, 82(256):56–

66, 2006. ISSN 0013-0249. doi: 10.1111/j.1475-4932.2006.00293.x. URL http://doi.wiley.com/10.1111/j.1475-4932.2006.00293.x.

Tatiana A. Homonoff. Can small incentives have large effects? the impact of taxes versus bonuses on disposable bag use. *American Economic Journal: Economic Policy*, 10(4):177–210, November 2018.

Guido Imbens and Karthik Kalyanaraman. Optimal bandwidth choice for the regression discontinuity estimator. *The Review of Economic Studies*, 79(3):933–959, July 2012.

IPCC. Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change, 2014. URL https://www.ipcc.ch/report/ar5/syr/. R.K. Pachauri and L.A. Meyer (eds.), Geneva, Switzerland.

Koichiro Ito. How do consumers respond to nonlinear pricing? Evidence from household water demand. Mimeo, April 2013.

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

Koichiro Ito, Takanori Ida, and Makoto Tanaka. Moral suasion and economic incentives: Field experimental evidence from energy demand. *American Economic Journal: Economic Policy*, 10(1):240–267, February 2018.

Katrina Jessoe, Maya Papineau, and David Rapson. Utilities included: Split incentives in commercial electricity contracts. Technical report, E2e Working Paper 029, 2018.

Katrina Jessoe, Gabriel E. Lade, Frank Loge, and Edward Spang. Spillovers from behavioral interventions: Experimental evidence from water and energy use. UC Davis Working Paper, January 2019.

George Kostyrko. Top Story: California's Cumulative Water Savings Continue to Meet Governor's Ongoing Conservation Mandate, 2015. State of California. http://ca.gov/drought/topstory/top-story-51.html.

David S. Lee and David Card. Regression discontinuity inference with specification error. *Journal of Econometrics*, 142(2):655–674, February 2008.

David S. Lee and Thomas Lemieux. Regression discontinuity designs in economics. *Journal of Economic Literature*, 48(2):281–355, June 2010.

Arik Levinson and Scott Niemann. Energy use by apartment tenants when landlords pay for utilities. *Resource and Energy Economics*, 26:51–75, 2004.

John A. List, Robert D. Metcalfe, Michael K. Price, and Florian Rundhammer. Harnessing policy complementarities to conserve energy: Evidence from a natural field experiment. NBER Working Paper 23355, April 2017.

Alexandre Mas and Enrico Moretti. Peers at work. *The American Economic Review*, 99(1):112–145, March 2009.

Justin McCrary. Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, 142(2):698–714, February 2008.

David L. Mitchell and Thomas W. Chesnutt. Evaluation of East Bay Municipal Utility District's pilot of WaterSmart Home Water Reports. Commissioned report, for the California Water Foundation and the East Bay Municipal Utility District, December 2013.

Paula Munger and Paul Yoon. 2018 naa survey of operating income & expenses in rental apartment communities. Technical report, National Apartment Association, 2018.

Erica Myers. Asymmetric information in residential rental markets: Implications for the energy efficiency gap. *American Economic Journal: Economic Policy*, forthcoming.

n.a. Water Rationing Begins May 1, 2014. SCMU Review, Santa Cruz Municipal Utilities. http://www.cityofsantacruz.com/home/showdocument?id=44394.

n.a. Stage 3 Water Shortage Emergency: 2015 Update for Santa Cruz Municipal Utility Customers, 2015a. City of Santa Cruz. http://cityofsantacruz.com/home/showdocument?id=36799.

n.a. Historical Palmer Drought Indices, 2015b. National Oceanic and Atmospheric Administration. http://www.ncdc.noaa.gov/temp-and-precip/drought/historical-palmers/psi/199512-201511.

n.a. Santa cruz water school presentation, 2015c. City of Santa Cruz. http://www.cityofsantacruz.com/departments/water/drought/water-school.

Office of Governor Edmund G. Brown Jr. Governor brown directs first ever statewide mandatory reductions. Press Release, April 2015.

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

Wayne Palmer. Meteorological drought. Research Paper 45, US Department of Commerce, US Weather Bureau, Office of Climatology, 1965.

José A. Pellerano, Michael K. Price, Steven L. Puller, and Gonzalo E. Sánchez. Do extrinsic incentives undermine social norms? Evidence from a field experiment in energy conservation. *Environmental and Resource Economics*, 67(3):413–428, November 2017.

A. Mitchell Polinsky and Steven Shavell. The optimal tradeoff between the probability and magnitude of fines. *The American Economic Review*, 69(5):880–891, December 1979.

Travis C. Pratt, Francis T. Cullen, Kristie R. Blevins, Leah E. Daigle, and Tamara D. Madensen. The empirical status of deterrence theory: A meta-analysis. *Advances in Criminilogical Theory*, 15:367–395, 2009.

Steven L Puller and Jeremy West. Efficient Retail Pricing in Electricity and Natural Gas Markets. *American Economic Review*, 103(3):350–355, 2013. ISSN 0002-8282. doi: 10.1257/aer.103.3.350. URL http://pubs.aeaweb.org/doi/abs/10.1257/aer.103.3.350.

Kimberly J. Quesnel and Newsha K. Ajami. Changes in water consumption linked to heavy news media coverage of extreme climatic events. *Science Advances*, 3 (10):1–9, October 2017.

Peter C. Reiss and Matthew W. White. What changes energy consumption? Prices and public pressures. *RAND Journal of Economics*, 39(3):636–663, Autumn 2008.

Julian Rode, Erik Gómez-Baggethun, and Torsten Krause. Motivation crowding by economic incentives in conservation policy: A review of the empirical evidence. *Ecological Economics*, 117:270–282, September 2015.

Paul Rogers. California drought: Gov. jerry brown says $10,000-a-day fines for water-wasting cities are not 'bluster'. The Mercury News, June 2015.

Sarah Smith, Frank Windmeijer, and Edmund Wright. Peer effects in charitable giving: Evidence from the (running) field. *The Economic Journal*, 125(585):1053–1071, June 2015.

Richard Titmuss. *The Gift Relationship: From Human Blood to Social Policy*. Allen and Unwin, 1970.

Amos Tversky and Daniel Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5:297–323, 1992.

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

Casey J. Wichman, 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*, 79:114–134, 2016.