

Automated Enforcement of Irrigation Regulations and Social Pressure for Water Conservation

Jeremy West, Robert W. Fairlie, Bryan Pratt, Liam Rose

Abstract: This study evaluates two interventions for residential water conservation. Comparing households across an enforcement algorithm's cutoff using a regression discontinuity design, we find that automated irrigation violation warnings cause substantial water conservation but also shift some consumption from regulated to unregulated hours within the week. In contrast, we show using data from a randomized experiment with the same customers that normative home water reports reduce water use by a much smaller amount, but that this social pressure is effective during all hours both before and after automating irrigation policy enforcement. Our findings highlight the merits of implementing multidimensional conservation programs.

JEL Codes: D04, L98, Q25, R22

Keywords: resource conservation, program evaluation, smart meter technology

FRESH WATER AVAILABILITY REMAINS ONE of the most pressing environmental and economic challenges 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

Jeremy West (corresponding author) is at the University of California, Santa Cruz, and the E2e Project (westj@ucsc.edu). Robert W. Fairlie is at the University of California, Santa Cruz, and the National Bureau of Economic Research (rfairlie@ucsc.edu). Bryan Pratt is at the US Department of Agriculture Economic Research Service (bryan.pratt@usda.gov). Liam Rose is at the US Department of Veterans Affairs and Stanford University School of Medicine (liamrose@stanford.edu). We thank WaterSmart Software and William Holleran in particular for sharing data. This work benefited from comments from editor Mary Evans and several anonymous reviewers, as well as Joshua Abbott, Hunt Allcott, Kelly Bishop, Daniel Brent, Lopamudra Chakraborti, Jesse Cunha, Lucas Davis, Carlos Dobkin, Zoe Fairlie, Daniel Friedman, Brent Haddad, Katrina Jessoe, Nicolai Kuminoff, Natalia Lazzati, Justin Marion, Jonathan Meer, Adam Millard-Ball, Steven Puller, *Dataverse data*: <https://doi.org/10.7910/DVN/BZWANC>

Received June 15, 2020; Accepted May 16, 2021; Published online September 15, 2021.

Journal of the Association of Environmental and Resource Economists, volume 8, number 6, November 2021.
© 2021 The Association of Environmental and Resource Economists. All rights reserved. Published by The University of Chicago Press for The Association of Environmental and Resource Economists.
<https://doi.org/10.1086/715472>

and that this outlook will only worsen with climate change.¹ At more local levels, concerns are also growing about cities sinking from overextracting underground water resources.² And, as weather patterns become more erratic and severe droughts more frequent, this will continue to increase the prevalence of prolonged water shortages such as those recently experienced in California, other southwestern states, and in many regions around the world.

In conventional market settings, the response to a severe supply shortage is that prices will increase and demand will adjust accordingly; however, prices in water markets are rarely set competitively, and governments are typically constrained in implementing Pigouvian remedies to curtail water consumption (Olmstead et al. 2007). Instead, policy makers rely heavily on water use regulations such as irrigation restrictions and social pressure—the influence on people by their peers—to encourage voluntary water conservation.

The efficacy of social pressure is often limited. For instance, many water utilities provide residential customers with home water reports (HWR) that compare each household's water consumption with that of their neighbors. Evaluations of randomized HWR typically find conservation effects of around 2%–5%—certainly not trivial but far from adequate to address the magnitude of the shortages (e.g., Ferraro et al. 2011; Ferraro and Price 2013; Mitchell and Chesnutt 2013; Bernedo et al. 2014; Brent et al. 2015; Jessoe et al. 2019, 2021; Bhanot 2021). Moreover, Allcott and Kessler (2019) find that the majority of conservation nudge recipients are unwilling to pay the marginal social cost of the nudge, which questions the economic efficiency of such policies.

With regulations, the challenge is enforcement. For example, a common form of water rationing is to restrict outdoor water use to certain days of the week, but the reliance on violators' neighbors to be informants means that in practice these restrictions are rarely enforced.³ However, for water conservation—as in many contexts—technology is advancing rapidly to enable automated detection of violations and enforcement

Jonathan Robinson, Jason Scorse, Kerry Smith, John Sorrentino, Alan Spearot, Forrest Williams, and seminar participants at Arizona State University, the Middlebury Institute, University of California, Santa Cruz, and the 2018 AERE summer conference. This research was supported in part by the US Department of Agriculture, Economic Research Service. The findings and conclusions in this work are ours and should not be construed to represent any official USDA or US Government determination or policy. Any errors and conclusions are our own and do not necessarily represent the views of WaterSmart Software nor any other parties.

1. The United Nations' forecasts are available at www.un.org/waterforlifedecade/scarcity.shtml.

2. Some examples of sinking cities are 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.

3. For instance, California state reports show that during the recent drought we study, most water agencies that restricted irrigation never issued a single penalty (CA State Water Board, www.waterboards.ca.gov). More generally, the poor quality of reporting of violations hinders the scope of regulatory enforcement in many settings of environmental compliance monitoring (e.g., Evans et al. 2009; Gilpatric et al. 2011).

of regulations. Thus, given the aforementioned limitations, a promising second-best policy to potentially address water shortages is technology-enforced restrictions on water use.

We examine such a policy in this study, providing the first evaluation to our knowledge of the effects of automating the enforcement of water conservation regulations. In addition, we reevaluate an HWR social pressure intervention in the same setting and with the same customers (Jessee et al. 2019, 2021), providing for a direct comparison between these two types of conservation policies as well as an evaluation of their layered effects. Our findings indicate that automating the enforcement of preexisting residential irrigation restrictions has large effects, inducing treated households to curtail their summer water consumption by about 31%. In comparison, we find—consonant with prior WaterSmart studies in the literature—that HWR reduce household water consumption by about 3%. That is, the effects of technology-enforced regulations in this context are an order of magnitude larger than those of the social pressure. Facilitated by novel high-frequency data, we additionally show that automating enforcement causes households to shift some water consumption into unregulated time periods, a form of environmental leakage; in contrast, HWR yield water conservation across the hours of the week.

To empirically arrive at these findings, we use data from WaterSmart Software and a Southern Californian water utility that measures hourly residential water consumption using advanced metering infrastructure. In July 2015, the utility tested an innovative approach to enforcing existing mandatory restrictions by leveraging the smart meter data to automate detection of irrigation violations and notify offending households. These notices stipulated fines for continued violation of the preexisting irrigation policies and made it clear that violations had been detected by computer algorithms. This pilot program dramatically increased the scope of enforcement: within 1 week, the one-time application of automation increased the share of households that had ever been warned from 4.6% to 39.2%.⁴ We exploit the algorithmic nature of this strengthened enforcement in a regression discontinuity design that compares barely treated households to those whose targeted water use fell just below the essentially arbitrary cutoff that designated an irrigation violation. Throughout the postautomation summer water season, we estimate local average treatment effects (LATEs) of 566 gallons saved weekly per household—31% of the mean—and that this overall effect is composed of gross conservation of 756 gallons during irrigation-prohibited time periods and a gross increase in consumption of 190 gallons when irrigation was allowed.

To directly compare these water conservation effects to those of social pressure, we utilize a field experiment conducted by WaterSmart Software in partnership with Jessee

4. The notices never state that this was a one-time pilot program testing the algorithm, but to the extent that some households viewed the warnings as cheap talk, this implies that our estimates are lower bounds of the conservation treatment effects of a more persistent automated enforcement of irrigation restrictions.

et al. (2019, 2021). The experiment sent HWR to randomly selected households in the same city during the months spanning the utility's automated irrigation enforcement. Using this randomization to identify intent-to-treat effects, we find—replicating their results—that HWR reduce average household water consumption by about 78 gallons per week during the summer water season, roughly 3% of the control group mean and well within the range of estimates in the sizable literature on HWR. Extrapolating the regression discontinuity (RD) estimates for the automated enforcement to apply to all treated households (37% of the sample population), a comparison of the two interventions on a per-treated-household basis shows that, while the overall effects of HWR are economically significant, the total conservation induced is about a tenth of the magnitude of that from automated enforcement of landscape irrigation regulations.

We then build upon this work to explore the layering of the two interventions. By estimating treatment effects in the weeks just before versus just after the automated violation notices, we find average conservation effects from HWR of 79 gallons per week prior to the automated enforcement (in May and June) and effects of 78 gallons just after (in July through October). That is, our evidence is suggestive that the effects of the randomized HWR are completely invariant to the regime change in irrigation enforcement. We additionally show that following this heightened enforcement, HWR continue to cause water conservation both during time periods when irrigation is allowed and during irrigation-prohibited days, including among households that are high-volume water consumers (the vast majority of whom were sent violation notices). This pattern starkly contrasts the increased water use during irrigation-allowed periods caused by the automated enforcement. An actionable policy implication of these collective findings is that, while technology-enforced regulations induce much greater behavioral change, social pressure serves as an effective simultaneous intervention to obtain additional conservation and mitigate some of the environmental leakage from asymmetric regulations.

Our study makes three primary contributions to the research literature. Most directly, we provide the first empirical evidence on automating the enforcement of regulations that target household resource conservation. The small literature on technology-enforced regulations has focused on automated traffic enforcement for speeding and red light violations, generally finding mixed evidence of social benefits (Retting et al. 2008; Hu and McCartt 2016; Gallagher and Fisher 2020). The evidence from our study supports the potential for large conservation benefits by the households targeted by automated enforcement for water policies, especially given the technological facilitation provided by the rapid growth in smart water meter deployment.⁵

5. Forecasts from ABI Research and IHS Markit predict that there will be a global base of 400 million smart water meters by 2026, one-third of new meter installations (see media coverage of these forecasts at [businesswire.com/news/home/20190801005064/en/400-Million-Smart-Water-Meters-Installed-Worldwide](https://www.businesswire.com/news/home/20190801005064/en/400-Million-Smart-Water-Meters-Installed-Worldwide) and [technology.informa.com/610487/water-meters-market-to-be-one-third-smart-by-2025](https://www.technology.informa.com/610487/water-meters-market-to-be-one-third-smart-by-2025)).

Second, we join a growing literature on the multifaceted approach of simultaneously increasing both social pressure and financial or regulatory incentives to change behavior. Policy makers regularly employ tactics such as layering strengthened regulations onto heightened normative interventions to internalize externalities and promote voluntary contributions toward public goods (Browne et al. 2020). For example, the US Forest Service attempts to reduce forest fires using both extensive prohibitions of risky behaviors—backed by statutory financial penalties—and large-scale media campaigns of normative messaging such as “only YOU can prevent wildfires.”⁶ Prior studies focus on how financial incentives might interact with social pressure for resource conservation (List et al. 2017; Pellerano et al. 2017; Gillan 2018; Ito et al. 2018; Holladay et al. 2019). We build on this literature by evaluating a context in which strengthened enforcement of existing regulations is implemented concurrently with a widely used form of social pressure. We also show promising new results regarding their use in combination.

Finally, our findings serve as novel evidence that a credible threat of fines can substantially influence people’s behavior. The threat of financial penalty is used as a primary policy instrument to back regulations in many settings, particularly so for environmental contexts such as littering, emissions control, and trespassing in protected areas. Whereas statutory penalties can readily be changed by policy makers, the expected value of a fine also depends greatly on the level of enforcement (Becker 1968). In the setting we study, official restrictions and penalties pertaining to irrigation remained unchanged while enforcement increased from virtually none to essentially complete enforcement. We provide causal evidence that households are very responsive to this newly credible threat of financial penalties, in contrast to a body of work highlighting relatively minimal sensitivity of water consumption to changing prices (e.g., Olmstead et al. 2007; Browne et al. 2020). Technological advances are likely to further enable utilities and governments to automate the enforcement of policies not only for water conservation, but for many other resources and behaviors as well.

1. STUDY SETTING AND RESEARCH DESIGN

As in many regions of the world, Southern California has a history of extreme and persistent variation in precipitation, including regular periods of extended drought conditions (see fig. A1; figs. A1–A6 are available online). During the most recent drought (2011–17), water utilities across California explored a wide variety of conventional and novel approaches to managing water demand in the face of rapidly shrinking supply. Our paper evaluates two simultaneous policy interventions for residential water conservation that were implemented in Burbank (Los Angeles County) by the local utility, Burbank Water and Power (BWP). Methodologically, we utilize a regression discontinuity design to identify the effects of automating enforcement of preexisting landscape irrigation restrictions,

6. See US Forest Service messaging at firerestrictions.us and smokeybear.com.

and we use a randomized field experiment to identify the effects of social pressure via home water reports (HWR).

Our primary research focus is on the automated enforcement of day-of-week and time-of-day outdoor water use restrictions (hereafter DOWR). Irrigation restrictions have a long heritage in water conservation and are widely used across California during droughts.⁷ Although the social benefits from water conservation do not vary across hours of the week, there are several institutional and horticultural reasons to impose DOWR.⁸ Perhaps more importantly, restricting outdoor irrigation to specific days facilitates enforcement. Prior to the introduction of “smart” meters that record high-frequency water consumption data using automated metering infrastructure (AMI), the only method of detecting irrigation violations was visual inspection by a utility employee or an informant neighbor of the violator. While DOWR allow informants to focus only on the extensive margin of “wrong day” water use, enforcement is believed to be quite low—as we provide empirical support for below. BWP’s use of a computer algorithm to enforce DOWR is highly novel, and ours is the first study to our knowledge to evaluate automated enforcement of residential water policies.

On May 14, 2015, the Burbank City Council approved the implementation of tighter restrictions, which included limiting outdoor water use to only on Tuesdays and Saturdays before 9:00 a.m. and after 6:00 p.m.⁹ Notably, BWP initially enforced these DOWR using only the conventional method of visual inspection. Then, during the first week of July, the utility conducted a pilot test of a computer algorithm that uses AMI data to automatically detect DOWR violations. Within 1 week, the share of single-family residences that had ever been found in violation jumped from 4.6% to 39.5%—that is, more than one-third of Burbank households were sent their first water violation notice in early July, as shown in figure 1. These notices (shown in fig. A2) clearly indicate that the violations were detected by a computer algorithm, which had not been previously announced. In addition, the notices reminded customers of the existing fines of \$100 for irrigating more often than twice per week, with fines increasing to \$200 and then to \$500 for subsequent violations. The notices do not state whether automated enforcement would continue, and BWP ultimately used the algorithm only once, so this treatment is best interpreted as being

7. See California Water Resources Control Board’s Water Conservation Portal: Conservation Reporting. www.waterboards.ca.gov/water_issues/programs/conservation_portal/conservation_reporting.html.

8. DOWR generally prohibit irrigation between a few hours after 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 that water can be spread efficiently for the benefit of plants. Finally, outdoor water use constitutes a large share of municipal water use and provides the potential to conserve water with minimal health and safety consequences (California Department of Water Resources 2013).

9. These revised DOWR replaced existing outdoor water use limits in Burbank of 3 days per week.

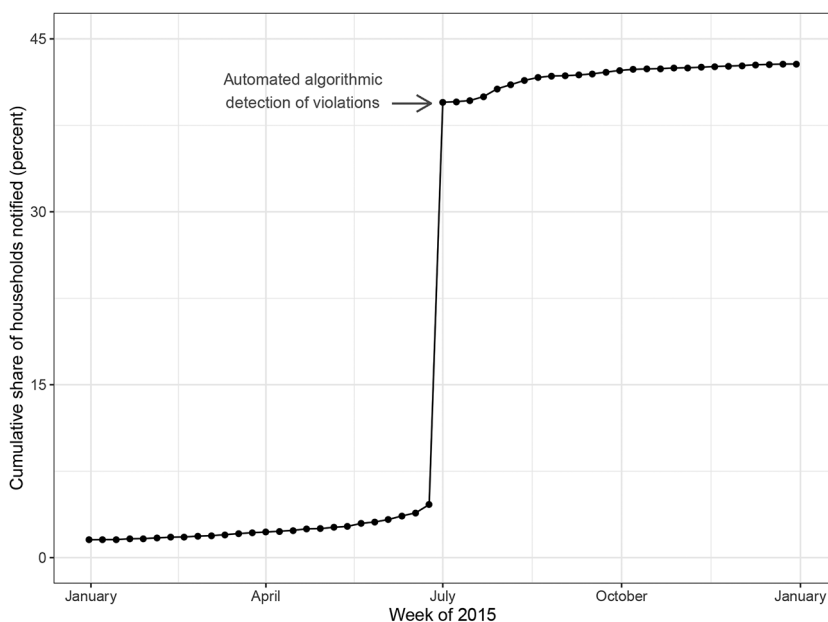


Figure 1. Time series for issued residential irrigation violation notices. This figure plots the cumulative share of in-sample households that had ever received an irrigation violation notice by week. Throughout this period, violations were determined when either a municipal employee or a neighbor of 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.

a shock to household beliefs about detection and enforcement probabilities pertaining to existing regulations and the associated pecuniary penalties.

To identify the effects of this novel enforcement, we use a regression discontinuity design (RDD) based on a cutoff in the algorithm that the BWP employed to determine violations. Using data from a single week in late June 2015, the noncompliance algorithm estimated the number of days per week that each household was irrigating, defined by whether the water meter had flow of more than 125 gallons during any individual hour of the day. For each household's third-highest daily peak consumption hour, this arbitrary cutoff of 125 gallons thus 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}, and a household would not be assigned to treatment if their seven daily peak hours were {200, 200, 124, 100, 100, 100, 100}. Although HWR assignment is not a factor in the algorithm, BWP 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 below. Because we imperfectly observe households' historical tiers

(particularly among the HWR control group), we assign all households to the running variable based on the third-highest daily peak consumption hour, which would determine violations under the strictest allowance. For this reason, we evaluate the automated enforcement using a “fuzzy” RDD.

We additionally evaluate the effects of HWR in Burbank during the same time period, which provides points of comparison for both overall conservation effects as well as more nuanced behavioral changes. Acting in partnership with BWP and the Center for Water and Energy Efficiency at University of California, Davis (Jesso et al. 2019, 2021), WaterSmart Software included nearly 17,000 single-family households in the randomized control study, with the timing of initial HWR treatment (by mail or email) rolled out over the monthly water billing cycle during late April through mid-May 2015. Notably, treated households began to receive HWR at least 6 weeks prior to the automated enforcement of irrigation restrictions in early July and continued to receive HWR for at least several months—throughout the October end of the 2015 summer water season that we include in our empirical analyses.

HWR include several components (see fig. A3). The primary component is a normative comparison of the treated household’s water consumption with that of a peer group of neighboring households with the same number of occupants and similar irrigable area.¹⁰ The reports also provide treated households with some potentially water-saving suggestions (e.g., “Upgrade to a low-flow toilet” or “Reduce showers to 5 minutes”). Some Burbank HWR also included one of three randomly generated messages, two of which pertain to hot water use.¹¹ Jesso et al. (2019) find “that households did not respond to the messaging or recommendations,” and we also verified that our estimated coefficients of interest are the same for the pooled three arms as for the random subset receiving only the core HWR messaging. Broadly, the literature has determined that the social pressure is by far the most effective component of HWR (and home energy reports) and that the other components are largely ineffective in the absence of normative peer comparisons (Ferraro et al. 2011; Ferraro and Price 2013; Mitchell and Chesnutt 2013; Allcott and Rogers 2014; Brent et al. 2015, 2020; Jesso et al. 2019, 2021; Bhanot 2021). We

10. HWR do not state the specific thresholds for the qualitative bins, but the “efficient” neighbor benchmark is based on the 20th percentile of peer group consumption and the “average” benchmark is the 55th percentile. The reports also provide the gallons per day values corresponding to each of these thresholds.

11. These 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% and gas use by 3% in the next 7 months and win one of: a) 25 high-efficiency Whirlpool clothes washers, b) 100 luxurious, efficient Evolve shower heads, or c) A hot water efficiency starter.”

cannot separately identify effects of each included piece of information, and we join prior literature in viewing the overall HWR intervention as the relevant “treatment” from a policy maker’s perspective.

Our research designs allow us to examine the impacts of automated enforcement and social pressure within a given period across quasi-randomly and randomly assigned groups. One consideration for this particular setting is the external validity of conclusions drawn at the height of a severe drought in a drought-prone region. Both of these interventions were introduced into a landscape filled with media coverage and other policies encouraging water-saving behavior, and BWP was facing threats of state-mandated penalties tied to conservation targets.¹² Although these factors might bias our estimates toward finding smaller effects, prior research supports that there is ample scope for residential water conservation, especially for outdoor water use (Castledine et al. 2014; Brelsford and Abbott 2018; Pratt 2019; Browne et al. 2020; Baker 2021). Moreover, as periods of extended drought become increasingly common, it is important to study policies during such an event in order to better understand the effects of related policies within the contexts in which they will be invoked.

2. DATA

Our study primarily uses data sourced from automated metering infrastructure residential water meters, which provide high-frequency records of household water consumption. 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. For municipal utilities, one widely recognized benefit of AMI is the ability to implement algorithmic detection of water leaks. In addition, the technology facilitates automated detection of landscape irrigation, as we study. Single-family residential accounts typically do not have separate meters for irrigation, so utilities (and researchers) are generally unable to identify irrigation disaggregated from total household water use; however, the flow rate of irrigation controllers is so large that consumption during an hour with irrigation far exceeds regular household consumption during any other hour of the week.¹³ Unlike with smart meters for electricity, AMI adoption has been relatively rare for household water use historically, although deployment of the technology exhibits a steep upward time trend.¹⁴ Thus, as water utilities increasingly install smart meters throughout their jurisdictions, the scope for applying AMI technology to enforce water policies will continue to grow steadily.

12. California enacted fines of \$10,000 per day for water agencies that did not meet mandated conservation targets (www.mercurynews.com/2015/04/28/water-wasting-fines-of-10000-proposed-by-gov-jerry-brown).

13. See an example comparison at www.wsscwater.com/customer-service/rates/water-usage.html.

14. In 2015, only about 7 million smart meters for water had been installed in the United States, compared to about 68 million smart electricity meters (see www.westmonroepartners.com/Insights/White-Papers/State-of-Advanced-Metering-Infrastructure).

Burbank Water and Power had installed AMI throughout their service area approximately 1 year prior to implementing the interventions we study. WaterSmart Software provided us with data on all single-family residential accounts within BWP's service territory in Burbank, California. The city of Burbank has a population of about 100,000 people, but the Home Water Report program focused on single-family homes with some irrigable landscape area. We observe hourly water consumption for nearly 17,000 single-family households during April 2014 through October 2015. We collapse these hourly data to the household-weekly level for our analysis, although we utilize the hourly disaggregation to identify patterns of within-week intertemporal substitution and to provide a deeper understanding of how and when households respond to the two water conservation policies that we evaluate.

In addition to data on water consumption, BWP and WaterSmart Software provided us with data on all home water reports and on violations of irrigation restrictions. We also incorporate several household-level covariates into our analyses. These control terms were compiled by WaterSmart Software through property records and inferred when necessary from other house characteristics; the covariates included are the size and irrigable area of each residential lot, and each house's size in square feet, year of construction, number of floors, number of bedrooms, and number of bathrooms.

Table 1 presents summary statistics for these variables, as well as demonstrating balance between the HWR treatment and control groups. By design and as expected with randomization across a large sample, implementation of the HWR treatment shows virtually complete compliance and the two arms are highly balanced. Prior to the automated enforcement pilot, roughly 5% of sample households had been sent a (nonautomated) violation notice. During the year prior to the interventions (April 2014 through March 2015), average household water use was about 2,700 gallons per week. The full distribution of this water consumption is plotted by HWR arm in figure A4. As is typical of household resource consumption, the distribution is skewed and has a long right tail; the standard deviation is 1,645 gallons and the Pearson's moment coefficient of skewness is 2.55. For comparison, during our primary analysis period of late May through October 2015, average water use by the control group was 2,323 gallons ($SD = 1,776$).

3. EMPIRICAL SPECIFICATIONS AND RESULTS

This section presents our empirical findings. First, we evaluate the impacts of the violation notices on water usage using a regression discontinuity design based on the computer algorithm used to automate enforcement. We estimate treatment effects both during hours when watering was allowed and during hours when irrigation was prohibited, as well as for total household water consumption. Second, we assess the impact of social pressure using the randomized home water reports field experiment. We estimate these treatment effects both for total consumption and separately for watering-allowed/prohibited hours. Finally, we examine the layering of the two policies by comparing estimates

Table 1. Summary Statistics and Randomization Balance Checks

Covariate	Group Means		<i>t</i> -Tests	
	Control (1)	HWR Treated (2)	Difference (3)	<i>p</i> -Value (4)
Number of households	2,920	13,703		
Sent WaterSmart HWR	0	.9972		
Prior water violation	.0486	.045	-.0036	.41
Lot size (sq ft)	7,346	7,322	-24	.74
Irrigable area (sq ft)	3,829	3,796	-33	.47
House size (sq ft)	1,619	1,620	1	.93
Year built	1,945	1,945	0	.86
Number of floors	1.061	1.067	.006	.27
Number of bedrooms	2.912	2.919	.007	.69
Number of bathrooms	1.929	1.939	.01	.59
Weekly water gallons	2,734	2,721	-13	.68

Note. This table 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, col. 3 shows the difference in means, and col. 4 shows the *p*-values for *t*-tests of whether the difference in group means is significantly different from zero. Initial HWR were sent to treated households during the billing cycle spanning from mid-April through mid-May 2015. All outcomes in the lower panel are determined prior to the randomization and prior to the automated irrigation restrictions enforcement. For pretreatment weekly water consumption, we use each household's average weekly gallons consumed during April 2014 through March 2015, spanning a full year prior to both treatments.

for the HWR intervention during weeks just before and just after the automated enforcement of irrigation restrictions.

3.1. Regression Discontinuity Estimates for Automated Enforcement

To examine the effects of the automated violation notices on household water consumption, we use a regression discontinuity design based on an arbitrary cutoff in the algorithm used to determine irrigation violations (discussed in sec. 1). We do not attempt to disentangle the mechanism(s) for these conservation effects, such as whether it is the pecuniary incentive from threat of fines or the increased perception of regulatory oversight by the utility and city. We can only test the reduced-form impact of this policy as it was implemented in Burbank—and as it could readily be deployed by other jurisdictions.

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 is no reason for concern. As shown in figure A5, there is some measurement lumpiness from the underlying meter technology, but there is no evidence of any asymmetric sorting of households around the threshold, which visually confirms the results of our statistical implementation




Figure 2. Regression discontinuity design: first stage for automated violation notices. This figure 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.

of McCrary’s (2008) test for manipulation. Further supporting the identification strategy, figure A6 demonstrates that there is also smoothness across the threshold in pretreatment water consumption along the running variable.

Next, we evaluate treatment compliance. Figure 2 displays the share of households receiving automated violation notices along the running variable. For visual clarity, figures 2–4 use bins of 10 gallons for the running variable. The size of the markers corresponds to the number of households included in each local average. Consistent with the algorithm, zero households below the threshold received a violation notice; at the threshold, there is a clear discontinuous jump for receipt of an automated violation notice in the first week of July. Because of the heterogeneous intensity of 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.

Quantitatively, we estimate that households with peak hourly water consumption just above the 125 gallon threshold are 25.5 percentage points more likely to have received a violation notice from the water utility relative to households just below the threshold. In

addition, because households that have higher peak water consumption are more likely to consume “above average” amounts of water per WaterSmart designation—and are thus more likely to be assigned to treatment using the stricter allowance—the treatment propensity increases with the running variable.¹⁵ At the right end of the range displayed in figure 2, we find that 40% of households received a notice, an increasing slope that continues past the displayed range. Thus, the running variable represents a necessary but not sufficient condition for a household to be sent an automated violation notice, and this first stage supports a fuzzy RDD.

Having established the validity of our first stage, we examine the effects of the automated enforcement on water use. We start by plotting local averages of posttreatment water consumption against the running variable. Figure 3 shows total average weekly water consumption in gallons; figure 4*a* shows water consumption during targeted periods of the week, when irrigation was not allowed; and figure 4*b* shows water consumption during hours of the week when irrigation was allowed (Tuesday and Saturday before 9:00 a.m. and after 6:00 p.m.). Average water consumption in these figures is pooled over July–October 2015, the 4-month period immediately following the automated violation notices treatment and continuing through the end of the statutory local summer water season.

Figure 4*a* shows a substantial discontinuous drop at the threshold in water consumption during irrigation-restricted periods of the week. In reduced form, the discontinuity is roughly 200 gallons per household per week, about 14.5% of the respective sample mean. For water consumption during the entire week, in figure 3 we also find a large drop at the threshold of about 8% in reduced form. This overall conservation is comparatively smaller, consistent with possible intertemporal substitution in response to the enhanced enforcement of an asymmetric restriction. Figure 4*b* reinforces evidence for such substitution, showing that water consumption discontinuously increased during hours of the week when irrigation was allowed.

We investigate these patterns more formally by estimating nonparametric local linear regressions of the following form:

$$\text{WaterUse}_{it} = \beta_0 + \tau \text{AboveCutoff}_i + f(\text{Peak Water Consumption}_i) + \varepsilon_{it}. \quad (1)$$

In equation (1), WaterUse_{it} is the quantity of water consumed by household i in week t . This outcome variable is either the household’s weekly total water consumption or the total consumption during subsets of hours of the week, such as consumption during irrigation-allowed hours only. In all regressions, each observation is one household for 1 week. AboveCutoff_i is an indicator for whether the household’s water use is above the cutoff in the algorithm used to determine irrigation violations (discussed in sec. 1). The term $f(\text{Peak Water Consumption}_i)$ is a nonparametric function of the household’s

15. As discussed in sec. 1, the algorithm uses “above average” households’ third-highest daily peak hour during a specific week in late June 2015 but allows more leniency for lower-volume consumers.

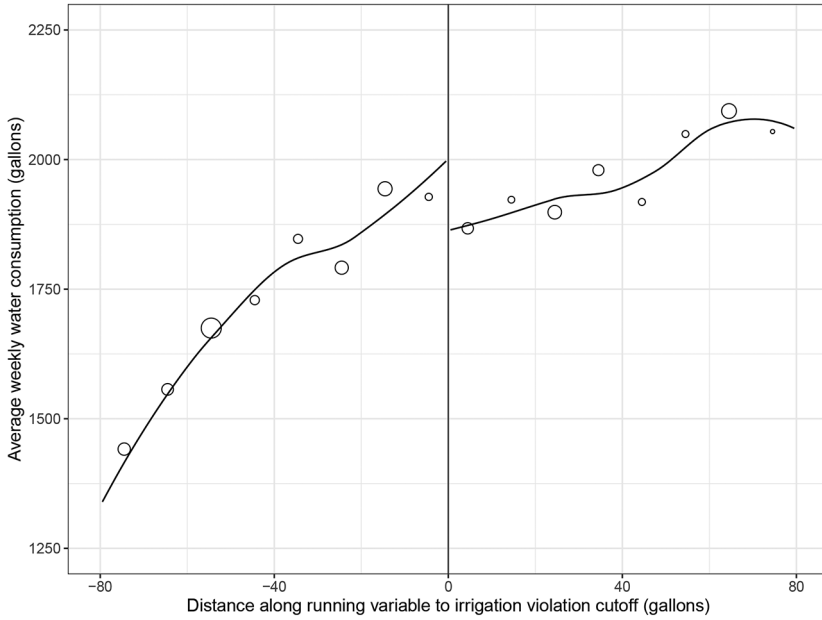


Figure 3. Reduced-form local averages for posttreatment weekly water consumption. This figure plots local averages for weekly water consumption during July–October 2015, the period following the automated violation notices treatment. For clarity, the running variable uses 10-gallon bins. The size of 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.

water use that entered into the algorithm and that is represented as the RDD running variable. The term ϵ_{it} is a mean-zero idiosyncratic error term.

In addition to presenting reduced-form results using equation (1), we also estimate local average treatment effects using nonparametric local linear regressions of the following form:

$$\text{WaterUse}_{it} = \beta_0 + \tau \text{AutomatedViolation}_i + f(\text{Peak Water Consumption}_i) + \epsilon_{it}. \quad (2)$$

Here, $\text{AutomatedViolation}_i$ is an indicator for whether household i was sent an automated violation notice due to the enforcement algorithm. Leveraging the RDD framework, we instrument for $\text{AutomatedViolation}_i$ using AboveCutoff_i to obtain the LATE estimates. Essentially, this rescales the estimates from equation (1) by the magnitude of the first-stage discontinuity. We estimate all RDD specifications using techniques from Calonico et al. (2014). With their provided statistical software package, we use local linear regressions with a triangular kernel.

Table 2 reports RD estimates of the effects of irrigation violation notices. Panel A presents reduced-form estimates, and panel B presents the local average treatment

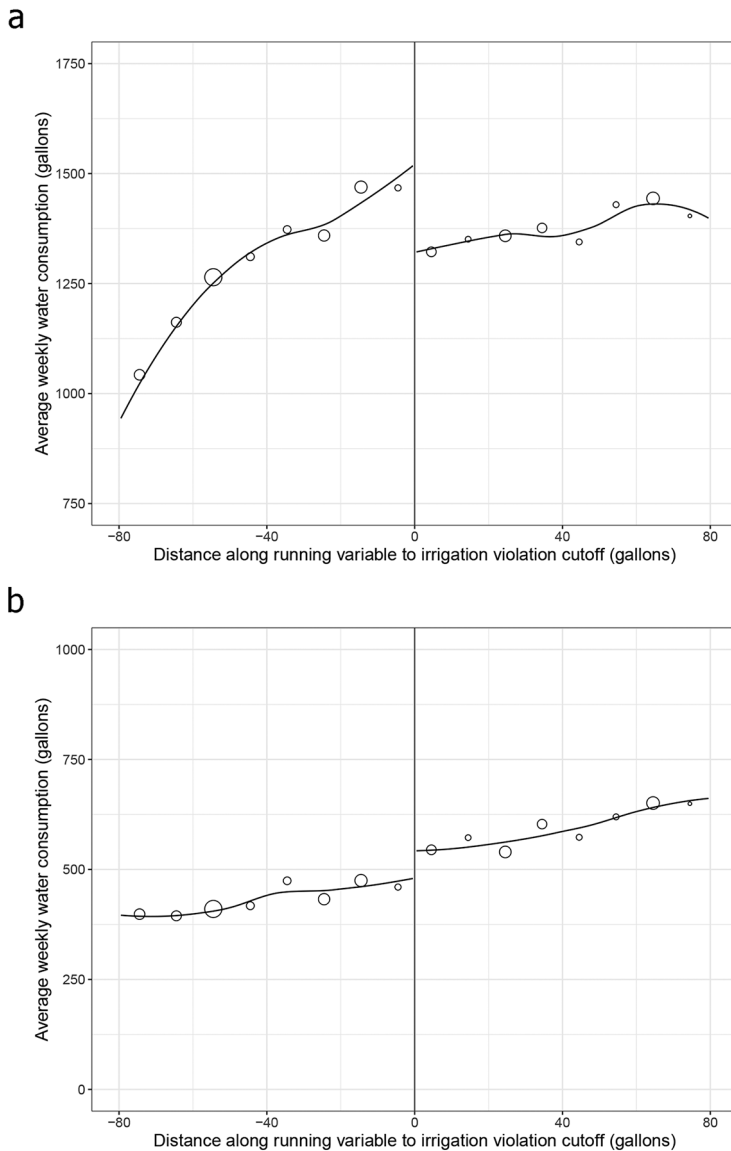


Figure 4. Reduced-form for posttreatment weekly water consumption by hours of the week. *a*, Consumption during hours of the week when irrigation is not allowed. *b*, Consumption during hours of the week when irrigation is allowed. This figure plots local averages for weekly water consumption during 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.

Table 2. Regression Discontinuity Estimates of Effects of Irrigation Violation Notice

	First Stage (1)	Weekly Water Consumption: July–October 2015 (Gallons)		
		All Hours (2)	Nonirrig. Hours (3)	Irrig. Hours. (4)
A. Reduced-form estimates:				
Discontinuity	.2549*** (.0068)	−144.2*** (26.1)	−192.6*** (21.42)	48.4*** (10.67)
B. Local average treatment effects:				
Discontinuity		−565.9*** (106.5)	−755.8*** (89.53)	189.9*** (41.29)
Sample mean	.15	1,833	1,334	499
Bandwidth (gal)	80	80	80	80
Observations	111,211	111,211	111,211	111,211

Note. Each cell presents a nonparametric regression discontinuity estimate at the cutoff for automated violation notices. All regressions use the “rdrobust” software package developed and provided by Calonico et al. (2014). Heteroskedasticity-robust bias-corrected standard errors are estimated using the same package. Column 1 provides the estimated first stage for automated violation notices. These notices were sent to households during the first week of July 2015. Columns 2–4 present estimates for household 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 water consumption pooled across all hours of the week. Column 3 includes consumption only during hours of the week when irrigation was not legally allowed. Column 4 includes consumption only during hours irrigation was legally allowed: Tuesdays and Saturdays before 9:00 a.m. or after 6:00 p.m.

* $p < .1$.

** $p < .05$.

*** $p < .01$.

effects, which essentially rescale the reduced-form estimates by the estimated magnitude of the first stage. Each cell in the table presents an RD estimate at the cutoff for automated violation notices. We use a consistent bandwidth of 80 gallons and compute heteroskedasticity-robust bias-corrected standard errors using the approach provided by Calonico et al. (2014).¹⁶

In panel A, we present the reduced-form estimates that correspond to figures 3 and 4. Column 3 shows estimates for water consumption during irrigation-prohibited hours of the week. We find a statistically significant drop of 192.6 gallons per week at the threshold. Some of this decrease in water consumption, however, is offset by an increase in water consumption during nonrestricted periods within the week. Column 4 shows that water

16. Estimates are quantitatively and qualitatively similar when using data-driven “optimal” bandwidths.

consumption increased discontinuously at the treatment threshold on average by 48.4 gallons per week during the irrigation-allowed hours on Tuesdays and Saturdays before 9:00 a.m. or after 6:00 p.m. 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 2, we also find significant effects at the threshold, with total water consumption decreasing in reduced form by 144.2 gallons per week.

As displayed in figure 2, household receipt of violation notices increases substantially at the threshold. Confirming this visible discontinuous jump, nonparametric RD estimates indicate an increase of 25.5% at the cutoff for automated violation notices. The estimate is reported in the first column of panel A in table 2. 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.

Panel B of table 2 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 755.8 gallons per week during irrigation-restricted times of the week. In contrast, water use during irrigation-allowed portions of the week increases by 189.9 gallons per week. Finally, total weekly water use decreases by 565.9 gallons per week on average for households sent a violation notice. To place this into perspective, these estimates imply that automated enforcement decreases household water use by about 31% on average, a very economically significant effect.¹⁷

In table 3, we show that these conservation effects persist throughout (at least) the 2015 summer water season. The table presents reduced-form and local average treatment effect estimates for household weekly water use by month during July, August, September, and October. While the estimates fluctuate somewhat—as is expected of data on a highly variable outcome—there is compelling evidence that estimated impacts of the violation notices persist for months after the one-time use of automated enforcement.¹⁸ While the largest point estimate is for July, a LATE of 641.3 gallons per week, the estimated LATE remains at 526.2 gallons in October. Moreover, the estimates for July

17. Our estimates for automated enforcement of irrigation policies are substantially larger than those for conventionally enforced DOWR (e.g., Hayden and Tsvetanov 2019). In percentage terms, we find water conservation on par with that from subsidizing “Water Smart Landscape” conversion (e.g., Brelsford and Abbott 2018; Baker 2021) or increasing marginal water prices by about 150% (Browne et al. 2020).

18. While we have no data on specific household behavioral changes, these findings are consistent with an induced change in irrigation controller settings. Given that (counterfactual) households clearly were not making similar changes absent the notifications, this potential explanation is interesting in its own right.

Table 3. RD Estimates of Effects of Irrigation Violation Notice by Month

Time Period	Weekly Water Consumption in 2015 (Gallons)						Sample Mean
	Reduced Form		LATE		Observations	Bandwidth	
	Coefficient	SE	Coefficient	SE			
July	-163.0***	(48.99)	-641.3***	(202.4)	31,385	80	1,890
August	-111.0*	(57.58)	-437.5*	(233.3)	24,735	80	2,006
September	-159.0***	(49.99)	-625.2***	(205.0)	30,695	80	1,793
October	-135.7***	(51.07)	-526.2**	(206.4)	24,396	80	1,635

Note. Each row presents nonparametric regression discontinuity estimates at the cutoff for automated violation notices. The outcome variable is total weekly water consumption. The month for each week is defined based on which month the first day of the week falls within. All regressions use the "rdrobust" software package developed and provided by Calonico et al. (2014). Heteroskedasticity-robust bias-corrected standard errors are estimated using the same package.

* $p < .1$.

** $p < .05$.

*** $p < .01$.

through October form a rather tight range between 437.5 and 641.3 gallons per week. As shown in table A1 (tables A1–A4 are available online), the treatment effects appear to persist even through the summer of 2016.

On the whole, these results for the effects of automated enforcement of irrigation restrictions are both academically interesting and directly policy relevant. Following a one-time application of automated enforcement, it is clear that households respond by significantly altering consumption patterns. While the aggregate impact on consumption is the most consequential for addressing water shortages, it is also valuable to understand how these reductions are achieved, including the evidence of within-week intertemporal substitution.

3.2. Estimated Effects of the Randomized Home Water Reports

We next examine how social pressure affects water conservation. Our identification strategy uses a field experiment in which randomly selected households were provided HWR including normative social comparisons of water use. Building on previous work by Jessoe et al. (2019, 2021) that explores the overall effect of the intervention, we primarily focus on time periods corresponding to features of the automated enforcement of irrigation policy. Of novel interest, we explore the sensitivity of HWR to the sharp change in enforcement by estimating treatment effects in the weeks just before versus after the automated irrigation violation notices.

Figure 5 displays average household weekly water consumption by month for both the treatment group that received HWR and the control group that did not. The time range shown in the figure spans one full year from November 2014 to October 2015, consisting of three distinct policy regimes. From November 2014 through March 2015

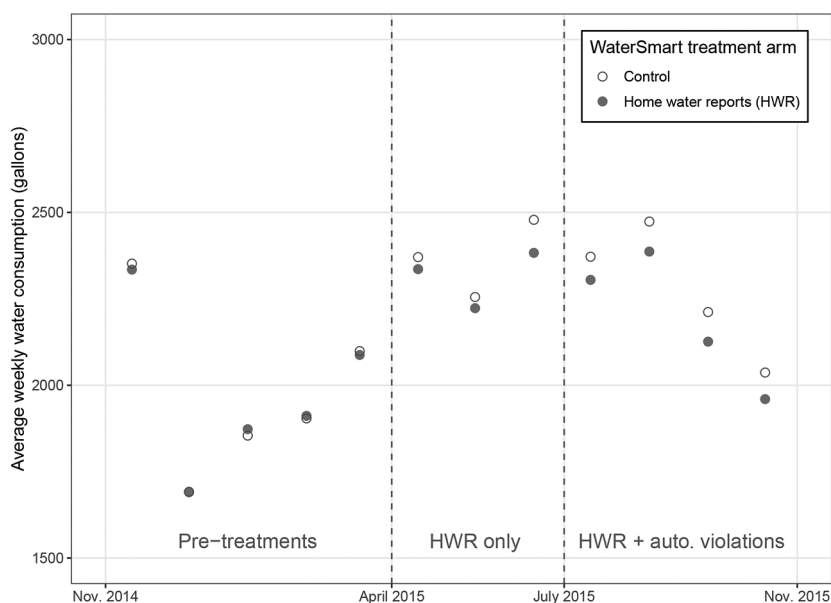


Figure 5. Average weekly water consumption by WaterSmart arm across the three regimes. This figure plots average weekly residential water consumption by month for each WaterSmart home water reports treatment arm during November 2014 through October 2015. The dashed vertical lines indicate the three regimes of policy interventions. HWR were sent to households starting with the April–May 2015 billing cycle and throughout the end of 2015 (and later). Algorithmic automated notices for violations of irrigation restrictions were sent to households during the first week of July 2015.

is a “pretreatments” regime when neither HWR nor automated enforcement were in effect. HWR-treated households were sent initial social comparisons during April through May 2015, and the “HWR only” regime runs from April through June. Then, the automated enforcement was conducted in the first week of July; the “HWR plus automated violations” regime runs from July through October, concluding the 2015 summer water season. In the pretreatments period, there is clearly no difference between the treatment and control groups, which were defined randomly. As initial HWR were sent in April and May, the AMI data enable us to see treatment effects immediately: average water use is visibly lower for HWR-treated households compared to the control group for every month over the experimental period.¹⁹ Finally, the magnitude of HWR treatment-control

19. The immediacy of the treatment effect is consistent with Reiss and White (2008), who find that electricity consumers respond promptly to both price changes and normative appeals. These findings support our use of May–June as a counterfactual for HWR effects during the “HWR plus automated violations” regime.

differences appears to remain strikingly similar in the third regime following the automated enforcement of irrigation restrictions.

We investigate the effects of HWR on water use more formally by estimating regression specifications of the following form:

$$\text{WaterUse}_{it} = \beta_0 + \gamma \text{HWR}_{it} + \mu_i + \omega_t + \delta X'_i + \varepsilon_{it}. \tag{3}$$

As in section 3.1, WaterUse_{it} is the quantity of water consumed by household i in week t . The term HWR_{it} indicates whether the household had been randomly assigned to be sent a home water report(s). The term X_t is a vector of household control terms, including 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 terms μ_i and ω_t are fixed effects for each household and week of the sample, respectively. The term ε_{it} is a mean-zero idiosyncratic error term.

Column 1 of table 4 uses a univariate ordinary least squares (OLS) regression of household weekly water consumption on assignment to treatment during the posttreatment period from late May through October. In column 2, we augment this specification with a vector of baseline controls for residential lot size, irrigable area, and the home’s

Table 4. Estimated Effects of Randomized WaterSmart Home Water Reports

	Weekly Water Consumption in 2015 (Gallons)		
	Late May–October		January–October
	(1)	(2)	(3)
I{HWR}	-77.88*** (28.38)	-77.67*** (25.62)	-68.76*** (17.44)
Household controls	No	Yes	...
Household fixed effects	No	No	Yes
Week of sample fixed effects	No	No	Yes
Control group mean	2,323	2,323	2,217
Number of households	16,623	16,623	16,623
Observations	391,552	391,552	706,882

Note. This table presents 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. Home water report–treated households each had been sent one HWR as of late May 2015, and monthly reports continued to be sent throughout (and following) October, the end of the legal and technical local summer water season. The household control terms in col. 2 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. Standard errors in parentheses are two-way clustered by household and week.

* $p < .1$.
 ** $p < .05$.
 *** $p < .01$.

square footage, year of construction, number of floors, number of bedrooms, and number of bathrooms. We also estimate a difference-in-differences specification in column 3 by adding preperiod data from January through late May of 2015, including household fixed effects, week of sample fixed effects, and interacting the assignment to treatment term with a binary indicator for post-HWR implementation. For each regression, standard errors in parentheses are two-way clustered by household and week.²⁰

Across the specifications, we find an average intent-to-treat effect of HWR of about 70–80 gallons per household per week. Using the control group sample averages (in the lower portion of table 4), the point estimates correspond to an average reduction in water consumption of 3.1%–3.4%, with a 95% confidence interval spanning 1.0%–5.7%.²¹ This finding essentially replicates that of Jessoe et al. (2019, 2021), using the same field experiment but slightly different measures. Our results also closely align with evidence on HWR in other jurisdictions (Ferraro et al. 2011; Ferraro and Price 2013; Mitchell and Chesnutt 2013; Bernedo et al. 2014; Brent et al. 2015, 2020; Bhanot 2021). Whereas the conservation benefits of HWR are qualitatively well documented, our study provides a novel point of comparison: we show that automated enforcement of irrigation restrictions has effects on water consumption that are about 10 times as large as those of HWR (31% vs. 3%).

Building on these existing results of the social pressure field experiment, we next explore how simultaneously using both HWR and automated enforcement policies affects household water conservation. We explore the layering of the policies in two ways. First, we examine how water conservation changes as automated enforcement is layered onto the social comparison policy over time. Second, we focus on how layering the policies changes water conservation by days of the week, that is, when outdoor irrigation is allowed or not allowed.²²

As discussed just above in the context of figure 5, the HWR posttreatment period consists of two regimes. During the 6 weeks from late May through June of 2015, 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, but automated violation notices had been abruptly issued to more than one-third

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

21. If, instead of calculating the percentage treatment effect relative to the randomly assigned control group, we compute it using pretreatment consumption for the treated (from table 1), these estimates represent a change of 2.9%.

22. In principle, we could test for policy interactions 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 near the RD cutoff, particularly among the much smaller control group. We present results in table A2 for the regression discontinuity estimates using only the sample of HWR-treated households, showing a very similar pattern as that shown for all households in table 2.

of households (as shown in fig. 1). We show above that this automated enforcement reduced average household water consumption by about 31% per week. Thus, there was a sharp change in the enforcement regime while other aspects of the setting and HWR experiment remained unchanged.²³ Empirically, we explore potential policy interactions and behavioral changes by estimating the effects of randomized HWR on post-treatment water consumption separately for different periods of time.

In table 5, we report these regression estimates. Column 1 includes our full HWR time period of late May through October, directly repeating table 4's column 2 for convenience. For this full posttreatment period, we find an intent-to-treat estimate from social pressure of 78 gallons per week reduction in water use per treated household. The remaining columns of table 5 maintain this same specification but vary the included time periods. In column 2, we estimate an average reduction of 79 gallons per household-week during late May through June, prior to the automated enforcement. In column 3, we estimate a conservation effect of 78 gallons per week during July through October, following the use of automated enforcement.²⁴ The comparison of these treatment estimates reveals that they are virtually identical, providing suggestive evidence that the policies may be fully additive.²⁵ The finding also provides some suggestive evidence on the interactions between mechanisms. Two mechanisms through which HWR theoretically might affect water conservation are by increasing responsiveness to intrinsic incentives (e.g., moral costs) and by reducing responsiveness to extrinsic incentives (e.g., financial costs) via substitution effects. Concerns about the behavioral response to intrinsic incentives being sensitive to the strength of extrinsic incentives have been conjectured in the broader economics literature (e.g., Bénabou and Tirole 2006; Gneezy et al. 2011; Pellerano et al. 2017), but we do not find evidence of a strong interaction between the two interventions in this setting.

The remaining two columns of table 5 explore the effects of HWR for different times within the week during the postautomation regime. Specifically, column 4 includes only

23. Following plans to the tariff structure announced years in advance, water prices 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. The median May water bill of 8,550 gallons would have increased by only \$1.59, inclusive of a \$1.00 increase to the fixed service charge. There were no additional changes during our study period.

24. As discussed above, the irrigation policy did not change at this time—only the approach to enforcement.

25. In table A3, we show that these results remain unchanged when adding additional controls for weather (temperature and precipitation) and to adding week-of-sample fixed effects. A formal test of the difference across time periods yields a point estimate of essentially zero. The standard error is small enough to rule out interactions of larger than a 25 gallons per week reduction or increase, about one-third of the average treatment effect of the HWR.

Table 5. Estimated Effects of Randomized WaterSmart Home Water Reports by Time Period

	Weekly Water Consumption in 2015 (Gallons)				
	Late May–October	Late May–June	July–October		
	All Hours (1)	All Hours (2)	All Hours (3)	Nonirrig. (4)	Irrig. (5)
I{HWR}	-77.67*** (25.62)	-78.80*** (28.96)	-77.51*** (25.57)	-39.71** (18.31)	-37.79** (15.28)
Household controls	Yes	Yes	Yes	Yes	Yes
Control group mean	2,323	2,442	2,283	1,414	869
Number of households	16,623	16,623	16,623	16,623	16,623
Observations	391,552	99,477	292,075	292,075	292,075

Note. This table presents estimates of the average intent-to-treat effect of the randomized WaterSmart HWR for weekly water consumption during 2015 for the time periods indicated by the column titles. Specifically, col. 1 includes the full 2015 summer treatment period from late May through October; col. 2 includes only late May through June, before the automated violation notices were sent; and cols. 3–5 includes July through October, after the automated violation notices were sent. Column 4 includes consumption only during hours of the week when irrigation was not legally allowed. Column 5 includes consumption only during hours irrigation was legally allowed: Tuesdays and Saturdays before 9:00 a.m. or after 6:00 p.m. Home water report–treated households each had been sent one HWR as of late May 2015, and monthly reports continued to be sent throughout (and following) October, the end of the legal and technical local summer water season. 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. Standard errors in parentheses are two-way clustered by household and week.

* $p < .1$.

** $p < .05$.

*** $p < .01$.

hours of the week when irrigation was prohibited, and column 5 includes only hours when outdoor irrigation was allowed. We find that HWR reduce average water use by 40 gallons per week during irrigation-prohibited time periods and by 38 gallons per week during irrigation-allowed time periods. This ubiquity of HWR effects with respect to irrigation permission starkly contrasts the intertemporal substitution shown earlier for the automated enforcement. Even though automated enforcement of irrigation regulations has large conservation effects during irrigation-prohibited periods, it also shifts some water use to nonrestricted days of the week; in contrast, social pressure has a blanketed effect of reducing water use regardless of irrigation restrictions. These results provide further evidence suggesting that social pressure works independently of automated enforcement policies for water conservation.

We bolster this evidence by focusing on high-volume water consumers, of which the vast majority (70%) were sent an automated violation notice. Table 6 reports estimates for the same outcomes as in table 5 but uses the subset of households who had pretreatment

Table 6. Estimated Effects of HWR by Time Period for High-Volume Consumers

	Weekly Water Consumption in 2015 (Gallons)				
	Late May–October	Late May–June	July–October		
	All Hours (1)	All Hours (2)	All Hours (3)	Nonirrig. (4)	Irrig. (5)
I{HWR}	-163.97*** (48.93)	-162.45*** (54.97)	-164.79*** (49.93)	-79.68** (40.20)	-85.11*** (32.15)
Household controls	Yes	Yes	Yes	Yes	Yes
Control group mean	3,572	3,832	3,483	2,060	1,423
Number of households	5,536	5,536	5,536	5,536	5,536
Observations	130,414	33,122	97,292	97,292	97,292

Note. This table replicates the specifications in table 5 for the subsample of households that had pretreatment water consumption in the top tercile of weekly volume. Of these households, 70% were sent an automated violation notice in early July 2015. Specifically, col. 1 includes the full 2015 summer treatment period from late May through October; col. 2 includes only late May through June, before the automated violation notices were sent; and cols. 3–5 includes July through October, after the automated violation notices were sent. Column 4 includes consumption only during hours of the week when irrigation was not legally allowed. Column 5 includes consumption only during hours irrigation was legally allowed: Tuesdays and Saturdays before 9:00 a.m. or after 6:00 p.m. 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. Standard errors in parentheses are two-way clustered by household and week.

* $p < .1$.

** $p < .05$.

*** $p < .01$.

water consumption in the top tercile of weekly volume. Of note, 70% of these households were sent an automated violation notice during the first week of July. For these high-volume consumers, we find even larger and more significant conservation effects from HWR. In column 1, we estimate an average intent-to-treat reduction of 164 gallons per week for the full time period (4.8% of the mean). Focusing on time periods before and after the automated enforcement intervention, we find similar levels of water conservation: 163 and 165 gallons, respectively. We also find similar levels of water conservation from HWR during irrigation-prohibited hours (80 gallons per week) as during irrigation-allowed hours (85 gallons per week).

Finally, we show results in table 7 in two panels separately for late May through June and for July through October (2015), across four time blocks within the week: (1) non-irrigation days between 0:00–9:00 and 18:00–24:00, (2) nonirrigation days between 9:00 and 18:00, (3) irrigation days (Tuesday and Saturday) between 0:00–9:00 and 18:00–24:00, and (4) irrigation days between 9:00 and 18:00. This decomposition allows us to compare the impact of HWR on water use for the same set of hours on days when irrigation is permitted and when irrigation is prohibited. The two panels of the table

Table 7. Estimated Effects of Randomized WaterSmart Home Water Reports by Time Block

	Weekly Water Consumption in 2015 (Gallons)			
	Non-Tuesday and Saturday		Tuesday and Saturday	
	0:00–9:00/ 18:00–24:00	9:00–18:00	0:00–9:00/ 18:00–24:00	9:00–18:00
	(1)	(2)	(3)	(4)
A. Late May–June:				
I{HWR}	-33.97** (15.43)	-19.88** (8.60)	-17.33 (12.30)	-7.41* (4.40)
B. July–October:				
I{HWR}	-19.06 (11.76)	-12.31** (5.96)	-36.74** (15.10)	-7.03* (3.63)
Household controls	Yes	Yes	Yes	Yes
Panel A control mean	1,001	499	699	245
Panel B control mean	744	425	855	218
Number of households	16,623	16,623	16,623	16,623
Panel A observations	99,417	99,417	99,349	99,349
Panel B observations	291,855	291,855	291,496	291,496

Note. This table presents estimates of the average intent-to-treat effect of the randomized WaterSmart HWR for weekly water consumption during 2015 for the time blocks within each week indicated by the column titles. Specifically, col. 1 includes the hours from midnight to 9:00 and 18:00 to midnight on days other than Tuesday and Saturday. Column 2 includes hours from 9:00–18:00 on days other than Tuesday and Saturday. Column 3 includes the hours from midnight to 9:00 and 18:00 to midnight on Tuesday and Saturday. Column 4 includes hours from 9:00–18:00 on Tuesday and Saturday. Panel A shows estimates for the period that includes only late May through June, and panel B shows estimates for the period that includes only July through October, after the automated violation notices were sent. Throughout the entire late May through October period, irrigation was allowed only on Tuesdays and Saturdays before 9:00 or after 18:00, i.e., in the hours for col. 3 only. Home water report–treated households each had been sent one HWR as of late May 2015, and monthly reports continued to be sent throughout (and following) October, the end of the legal and technical local summer water season. 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. Standard errors in parentheses are two-way clustered by household and week.

* $p < .1$.

** $p < .05$.

*** $p < .01$.

facilitate making this comparison before versus after the automated irrigation enforcement.²⁶ Qualitatively, the evidence in table 7 shows that HWR treatment effects are concentrated more heavily on nonirrigation days prior to the automated enforcement, and then the

26. Table A4 further disaggregates these estimates across 21 time blocks within the week and shows that, while the conservation from HWR is especially pronounced during irrigation-permitted time periods, the social pressure reduces water use across the hours of the week.

treatment effects moderately shift to occur during irrigation-allowed periods following the automated enforcement. This shift is reflected in the statistical (in)significance of the respective coefficients but is not itself a statistically significant shift in the composition of treatment effect timing. Overall, this new table supports that HWR appear to largely operate through a different behavior channel than that of the automated irrigation enforcement, especially as the HWR effects remain stable during nonirrigation hours (9:00–18:00) of all days of the week across the two regimes.

On the whole, these findings do not show any evidence of a significant change in HWR effects from layering the two interventions, and the estimates reported in tables 6 and 7 serve as a strong further robustness check for additive effects of the two policies. We caveat, however, that these evaluations are comparisons of behavior by the same people at different points in time; while total precipitation was very minimal during the periods we study, there might be seasonality or other time-varying factors determining the response to HWR treatment. Thus, we view our evidence as only being suggestive that the effectiveness of social pressure does not appear to vary with the strength of enforcement for regulations targeting the same or similar behaviors.

4. CONCLUSIONS

In this study, we examine both the impacts of automating the enforcement of day-of-week and time-specific irrigation regulations and of social pressure for water conservation via home water reports. Using hourly data on household water consumption for a city in Southern California, we estimate how these two policy interventions affect residential water usage during an extreme drought.

Using a regression discontinuity design based on the irrigation detection algorithm, we find that automating the enforcement of existing irrigation policies reduces average water consumption by approximately 31% among targeted households. In comparison, we find using a randomized field experiment that the social pressure reduces average water use by about 3%, replicating the findings of Jessoe et al. (2019, 2021). Extrapolating the RD estimates to apply to all treated households, this implies that the total water conservation effects of automated enforcement of irrigation restrictions are about 10-fold those of home water reports on a per-treated-household basis. Exploring how these respective effects vary across the hours of the week, we show that the automated enforcement induces intertemporal substitution from irrigation-prohibited to unregulated periods within the week; in contrast, we find that home water reports cause conservation during all hours of the week, both before and following the application of automated irrigation policy enforcement.

Although irrigation regulations and social pressure have both been studied extensively in the resource conservation literature, our study presents novel findings about the automation of associated enforcement and the efficacy of these two policies when implemented contemporaneously. Given the sparsity of existing literature studying automated enforcement of public policies, our research demonstrates promising potential for technology-enforced regulations to dramatically change consumer behaviors. While

the city we study used automated enforcement only once, we find that the conservation effects persist in magnitude for the duration of the summer water season and underlying irrigation policy period. Moreover, rather than the layering of multiple policies having diminishing marginal returns, these two interventions appear to yield fully additive conservation effects: home water reports successfully encouraged conservation both during time periods when irrigation was allowed and when irrigation was not allowed, to a similar extent both before and after the application of automated irrigation policy enforcement.

Our findings speak to a challenging trade-off faced by utilities and public policy makers. Automated enforcement of regulations produces meaningful resource conservation in this setting, but efficacy is only one consideration. Technology-driven enforcement is imperfect and somewhat invasive, which incurs complaints from offended utility customers.²⁷ There is also a potential concern that strengthened enforcement of regulations might reduce people's intrinsic motivation to conserve resources. The evidence from our study suggests that this potential displacement is minimal in practice, perhaps because irrigation policy and social pressure may work through different behavioral channels. As the technological scope for automated enforcement of policies continues to rapidly grow, the potential lack of popular support will need to be carefully balanced against the large potential social benefits from resource conservation and other behavioral changes. Our study shows that these policy-making trade-offs can reasonably be considered independently of existing or proposed social pressure interventions.

REFERENCES

- Allcott, Hunt, and Judd B. Kessler. 2019. The welfare effects of nudges: A case study of energy use social comparisons. *American Economic Journal: Applied Economics* 11 (1): 236–76.
- Allcott, Hunt, and Todd Rogers. 2014. The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation. *American Economic Review* 104 (10): 3003–37.
- Baker, Jonathan E. 2021. Subsidies for succulents: Evaluating the Las Vegas Cash-for-Grass rebate program. *Journal of the Association of Environmental and Resource Economists* 8 (3): 475–508.
- Becker, Gary S. 1968. Crime and punishment: An economic approach. *Journal of Political Economy* 76 (2): 169–217.
- Bénabou, Roland, and Jean Tirole. 2006. Incentives and prosocial behavior. *American Economic Review* 96 (5): 1652–78.
- Bernedo, Maria, Paul J. Ferraro, and Michael K. Price. 2014. The persistent impacts of norm-based messaging and their implications for water conservation. *Journal of Consumer Policy* 37 (3): 437–52.
- Bhanot, Syon P. 2021. Isolating the effect of injunctive norms on conservation behavior: New evidence from a field experiment in California. *Organizational Behavior and Human Decision Processes* 163:30–42.
- Brelsford, Christa, and Joshua K. Abbott. 2018. How smart are “Water Smart Landscapes”? Working paper, arXiv.org.

27. The water utility ultimately apologized for mailing some automated violation notices to households who were not actually in violation of the irrigation regulations, as discussed in this *Los Angeles Times* article: www.latimes.com/local/california/la-me-water-error-20150720-story.html.

- Brent, Daniel A., Joseph H. Cook, and Skylar Olsen. 2015. 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.
- Brent, Daniel A., Corey Lott, Michael Taylor, Joseph Cook, Kim Rollins, and Shawn Stoddard. 2020. What causes heterogeneous responses to social comparison messages for water conservation? *Environmental and Resource Economics* 77:503–37.
- Browne, Oliver, Ludovica Gasse, and Michael Greenstone. 2020. Do conservation policies work? Evidence from residential water use. Unpublished manuscript.
- California Department of Water Resources. 2013. California water plan update. Vol. 3, chap. 3, Urban water use efficiency. Technical report, water.ca.gov.
- Calonico, Sebastien, Matias D. Cattaneo, and Rocio Titiunik. 2014. Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica* 82 (6): 2295–326.
- Castledine, Anita, Klaus Moeltner, Michael K. Price, and Shawn Stoddard. 2014. Free to choose: Promoting conservation by relaxing outdoor watering restrictions. *Journal of Economic Behavior and Organization* 107 (A): 324–43.
- Christiansen, Jerald Emmet. 1942. *Irrigation by sprinkling*. Vol. 670. University of California, Berkeley, College of Agriculture, Agricultural Experiment Station.
- Evans, Mary F., Scott M. Gilpatric, and Lirong Liu. 2009. Regulation with direct benefits of information disclosure and imperfect monitoring. *Journal of Environmental Economics and Management* 57:284–92.
- Ferraro, Paul J., Juan José Miranda, and Michael K. Price. 2011. The persistence of treatment effects with norm-based policy instruments: Evidence from a randomized environmental policy experiment. *American Economic Review: Papers and Proceedings* 101 (3): 318–22.
- Ferraro, Paul J., and Michael K. Price. 2013. Using nonpecuniary strategies to influence behavior: Evidence from a large-scale field experiment. *Review of Economics and Statistics* 95 (1): 64–73.
- Gallagher, Justin, and Paul J. Fisher. 2020. Criminal deterrence when there are offsetting risks: Traffic cameras, vehicular accidents, and public safety. *American Economic Journal: Economic Policy* 12 (3): 202–37.
- Gillan, James M. 2018. Dynamic pricing, attention, and automation: Evidence from a field experiment in electricity consumption. Working paper 284, Energy Institute at Haas.
- Gilpatric, Scott M., Christian A. Vossler, and Michael McKee. 2011. Regulatory enforcement with competitive endogenous audit mechanisms. *RAND Journal of Economics* 42:292–312.
- Gneezy, Uri, Stephan Meier, and Pedro Rey-Biel. 2011. When and why incentives (don't) work to modify behavior. *Journal of Economic Perspectives* 25 (4): 191–210.
- Hayden, Hanna, and Tsvetan Tsvetanov. 2019. The effectiveness of urban irrigation day restrictions in California. *Water Economics and Policy* 5 (3): 1950001.
- Holladay, Scott, Jacob LaRiviere, David Novgorodsky, and Michael Price. 2019. Prices versus nudges: What matters for search versus purchase of energy investments? *Journal of Public Economics* 172:151–73.
- Hu, Wen, and Anne T. McCart. 2016. Effects of automated speed enforcement in Montgomery County, Maryland, on vehicle speeds, public opinion, and crashes. *Traffic Injury Prevention* 17 (sup1): 53–58.
- Ito, Koichiro, Takanori Ida, and Makoto Tanaka. 2018. Moral suasion and economic incentives: Field experimental evidence from energy demand. *American Economic Journal: Economic Policy* 10 (1): 240–67.
- Jessoe, Katrina, Gabriel E. Lade, Frank Loge, and Edward Spang. 2019. Residential water conservation during drought: Experimental evidence from three behavioral interventions. Working paper, University of California, Davis.
- . 2021. Spillovers from behavioral interventions: Experimental evidence from water and energy use. *Journal of the Association of Environmental and Resource Economists* 8 (2): 315–46.

- List, John A., Robert D. Metcalfe, Michael K. Price, and Florian Rundhammer. 2017. Harnessing policy complementarities to conserve energy: Evidence from a natural field experiment. NBER Working paper 23355, National Bureau of Economic Research, Cambridge, MA.
- McCrary, Justin. 2008. Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics* 142 (2): 698–714.
- Mitchell, David L., and Thomas W. Chesnutt. 2013. 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.
- Olmstead, Sheila M., Michael Hanemann, and Robert N. Stavins. 2007. Water demand under alternative price structures. *Journal of Environmental Economics and Management* 54 (2): 181–98.
- Pellerano, José A., Michael K. Price, Steven L. Puller, and Gonzalo E. Sánchez. 2017. Do extrinsic incentives undermine social norms? Evidence from a field experiment in energy conservation. *Environmental and Resource Economics* 67 (3): 413–28.
- Pratt, Bryan. 2019. A fine is more than a price: Evidence from drought restrictions. Unpublished manuscript.
- Reiss, Peter C., and Matthew W. White. 2008. What changes energy consumption? Prices and public pressures. *RAND Journal of Economics* 39 (3): 636–63.
- Retting, Richard A., Sergey Y. Kyrtychenko, and Anne T. McCartt. 2008. Evaluation of automated speed enforcement on Loop 101 freeway in Scottsdale, Arizona. *Accident Analysis and Prevention* 40:1506–12.