UC Santa Barbara
UC Santa Barbara Electronic Theses and Dissertations

Title
The Effect of Financial and Social Incentives on Water Conservation

Permalink
https://escholarship.org/uc/item/5d29r3wm

Author
Lott, Corey Lott

Publication Date
2017

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

University of California
Santa Barbara

The Effect of Financial and Social Incentives on Water Conservation

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

Doctor of Philosophy
in
Economics

by

Corey J. Lott

Committee in charge:

Professor Olivier Deschênes, Co-Chair
Professor Paulina Oliva, Co-Chair
Professor Kyle Meng

June 2017

The Dissertation of Corey J. Lott is approved.

Professor Kyle Meng

Professor Paulina Oliva, Committee Co-Chair

Professor Olivier Deschênes, Committee Co-Chair

June 2017

The Effect of Financial and Social Incentives on Water Conservation

Copyright © 2017

by

Corey J. Lott

I dedicate this research to my father, William K. Lott.

Acknowledgements

First and foremost, I would like to thank my family for their constant support during my educational journey. I am grateful to my mother, Julee Lott, and my sister, Morgan Lott and Fred Waititi for being patient and understanding about the long hours and late nights spent working on this research.

I am grateful to the members of my committee for their direction and advice. Paulina Oliva taught me to be more careful and rigorous with the way I approach my research and writing. She provided detailed guidance and feedback on my empirical strategy and modeling approach. Olivier Deschênes offered direction on where to take projects and big picture advice. Kyle Meng helped me craft more compelling economic arguments and provided needed encouragement during the job market process.

I would also like to thank Shawn Stoddard at Truckee Meadows Water Authority. Without him this research would not have been possible. Shawn ensured that I had unfettered access to data resources and project opportunities through TMWA. He also instilled in me the importance of careful coding and project documentation in Stata. Finally, Shawn helped me obtain financial support through TMWA during my graduate research.

I am grateful to Kimberly Rollins for encouraging me to pursue my Ph.D. at UCSB. Kim has also provided valuable career advice at every step of this journey and has encouraged me when I really needed it.

Lastly, I would like to thank my classmates from UNR and UCSB for help and encouragement throughout my doctoral studies. Christine Braun provided essential guidance with Matlab simulations. Matt Wibbenmeyer and Sarah Bana supported me throughout the job market. I would also like to thank the members of the Environmental Economics lunch seminars for their feedback on my presentations of this work.

Curriculum Vitæ
Corey J. Lott

Education

2017	Ph.D., Economics, University of California, Santa Barbara
2013	M.S., Economics, University of California, Santa Barbara.
2011	M.S., Resource Economics and Mathematics, University of Nevada, Reno
2011	B.A., Economics and Mathematics, University of Nevada, Reno

Research Fields

Environmental Economics, Applied Microeconomics, and Behavioral Economics

Committee Members

Olivier Deschênes (Co-Chair), Paulina Oliva (Co-Chair), Kyle Meng

Honors and Fellowships

2016	Chancellors Sustainability Graduate Research Internship, UCSB
2015	Outstanding Teaching Assistant, Economics Dept., UCSB
2014	Janet A. Alpert Fellowship, UCSB
2012-2013	Economics Block Grant Fellowship, UCSB
2010-2011	Jerry and Betty Wilson Mathematics Memorial Scholarship, UNR

Research Experience

2015	Field Experiment Coordinator, Truckee Meadows Water Authority
2011-2012	Research Analyst, Center of Resource Economics, Department of Economics, UNR
2008-2010	Research Assistant for Thomas Harris, University Center for Economic Development, UNR

Teaching Experience

2014-2016	Economics 10A (Intermediate Microeconomics), Head Teaching Assistant, UCSB

2013-2014	Economics 10A and Cost-Benefit Analysis (Masters level),Teaching Assistant, UCSB
2011	Statistics 152 (Introductory Statistics), Instructor, UNR
2010-2011	Mathematics 352 (Intermediate Probability/Statistics), Teaching Assistant, UNR

Conferences and Presentations

2017	University of Idaho, Department of Agricultural Economics and Rural Sociology, Laurits Christensen Associates
2016	Camp Resources, AERE Summer Conference, UCSB Environmental Lunch and Development Group, ASSA Annual Meeting
2015	UCSB Environmental Lunch and Development Group, Occasional Workshop in Environmental and Resource Economics
2014	AAEA Annual Meeting, UCSB Environmental Lunch, Nevada Water Resources Association Annual Meeting, UC Center for Energy and Environmental Economics Summer School, International Water Resource Economics Consortium Annual Meeting
2013	AAEA Annual Meeting
2012	W2133 Meeting: Benefits and Costs of Natural Resources Policies Affecting Public and Private Lands
2010	AAEA Annual Meeting

Abstract

The Effect of Financial and Social Incentives on Water Conservation

by

Corey J. Lott

Water utility managers employ a variety price and non-price policies to induce conservation as a way to defer expensive capacity expansion and to overcome scarcity during periods of drought. Economists tend to focus on financial motivations for consumers, but there is a growing literature in economics and other related fields that explores the potential for other sources of motivation to be harnessed to achieve policy outcomes. Knowing the relative effectiveness and the trade-offs associated with such approaches can help planners to make more efficient policy.

This dissertation investigates the effect of financial and social incentives on water conservation by studying multiple policy approaches used by a water utility in the Reno metropolitan area to promote conservation in the context of a historic drought in the southwestern U.S. In the first chapter, I investigate how providing better price information on monthly utility bills can promote consumer response to financial incentives for water conservation. I provide empirical evidence of consumer inattentiveness to prices, and discuss how this behavior can undermine price schedules that are designed in part to promote conservation. In the second chapter, I focus once again on consumer inattentiveness to prices and consumption behavior, but I analyze a treatment that lowers price salience rather than increases it. I examine enrollment in automatic bill payment and paperless billing programs, which can encourage consumers to become inattentive to recurring charges and thus less aware when changing consumption habits lead to increased water bills. These programs have the unintended consequence of encouraging consumers

to be less attentive to prices, which further undercuts the response to conservation incentives built into price schedules. In the final chapter, I turn from financial incentives for water conservation to how social incentives can be used to promote conservation using a field experiment to test the impact of providing normative comparisons on conservation during a drought.

Consumers often face complex pricing schemes when making purchasing decisions. Inefficiencies arise when such schemes are not fully understood. Electric and water utilities use increasing block tariffs (IBTs) to promote conservation. However, recent evidence suggests that consumers may not respond to the marginal price under these price systems. The first chapter of this dissertation investigates price salience as a possible mechanism by leveraging quasi-exogenous variation in the level of price-related information provided to households on their monthly water bills. I exploit a merger between the two water utilities in Reno that occurred in 2015. Before the merger, one group of households received information on monthly bills about the IBT price schedule, while the other did not; after the merger, both groups received price information. Using a difference-in-differences approach, I find that providing consumers with bills containing better price information leads to a more than 3% decrease in average water consumption, suggesting that salience is important. I develop a model of consumer decision-making under price uncertainty that predicts consumers will have a differential response to information based on where their consumption is relative to the block tariff threshold. When estimating the effect of information by historical consumption level, I find empirical evidence consistent with these predictions.

The second chapter investigates how automatic and paperless billing enrollment affect residential water demand using a difference-in-differences approach. I find that enrolling in automatic bill payment (ABP) leads to a 2-3% increase in average water consumption while paperless (PL) enrollment leads to a 1% increase in water consumption. I also shed light on the relationship between income effects and changes in consumption after enrollment by estimating heterogeneous treatment effects by quintiles of the distribution of property values. I find that average treatment effects do not significantly vary across the income distribution, which indicates that the increase in consumption from ABP enrollment is due to consumer inattentiveness rather than changes in income. I also consider how enrollment in ABP and PL can affect subsequent participation in voluntary conservation programs. ABP customers consume an additional 6% more than non-ABP customers during when voluntary drought restriction are implemented. Overall, there is evidence that enrollment in these programs promotes consumer inattentiveness to prices as well as utility communications, which undermines financial and social incentives designed to promote conservation.

The final chapter investigates how normative appeals for water conservation drive behavioral change using a large-scale, randomized field experiment. Using a new social comparison that reduces the correlation between pre-treatment consumption and the difference from the peer group, we isolate the impact of the normative component of the message. The strength of the message, defined as a household's performance relative to a peer group, is a primary driver of social comparisons' efficacy, consistent with social comparisons imposing a moral cost on excess consumption. Relative to a nudge that highlights financial savings, social comparisons generate less persistent water savings and are more dependent on sending multiple mailers. While moral motivations are likely driving behavioral change in response to normative appeals, there are opportunities to improve welfare by designing messages that prompt consumers to address mis-optimization of water consumption. This chapter is joint work with Daniel Brent, Joe Cook, Kimberly Rollins, Shawn Stoddard, and Michael Taylor.

Contents

List of Figures

List of Tables

Chapter 1

The Effect of Price Information on Consumer Behavior Under Nonlinear Tariffs: Evidence from a Water Utility Merger

1.1 Introduction

Consumers often face complex pricing schemes when making purchasing decisions, which can lead to inefficiencies if consumers mis-optimize when such schemes are not fully understood. One reason why consumers might have difficulty optimizing in this context is that prices are likely to be less salient under complex price structures, since consumers must navigate an entire price structure, requiring them to be knowledgeable about more than one marginal price. In addition, prices may not be transparent in many settings, which can exacerbate this problem (DellaVigna, 2009; Finkelstein, 2009; Sexton, 2015). In settings where market failures are involved, consumer inattentiveness may have social consequences.

Electric and water utilities throughout the United States and in many parts of the world use increasing block tariffs (IBTs) in part to promote conservation. While economic theory predicts that consumers should respond to marginal price under these price systems, recent empirical evidence suggests that consumers do not respond to marginal price when facing IBTs for electricity (Ito, 2014) or water consumption (Ito, 2013; Wichman, 2014). One implication of this behavior is that complex prices can be welfare reducing (Ito, 2014). Moreover, there may be implications for climate change and management of scarce water resources when policymakers have a limited ability to use price schedules to induce conservation.[1]

This chapter explores whether better information on a complex price schedule can induce residential water conservation. Using a combination of theory and empirics, I demonstrate that the answer depends jointly on the shape of the price schedule and the distribution of residential water consumption. Price information treatments that fail to consider these aspects may actually lead to *increased* residential water use.

I develop a model of consumer decision in the presence of price uncertainty under IBTs, which predicts that consumers will have a differential response to price information based on the location of their pre-treatment consumption relative to the block tariff thresholds. Providing price information on monthly bills decreases information costs and lowers price uncertainty for consumers, who will re-optimize consumption. The direction of the treatment effect depends on the interaction of the ex-ante distribution of consumption and the price schedule. Price information leads to an overall reduction in water consumption if there are more households above than below the thresholds. If not, price information may actually lead to *increased* water consumption. This model contrasts previous models of consumer inattentiveness, which assume consumers systematically misperceive price signals (DellaVigna, 2009; Chetty, Loney and Kroft, 2009; Wichman, 2016). These models conclude that a reduction (increase) in consumption after a price

[1] Electric and water utilities in the U.S. are regulated natural monopolies, and as such are mandated to set prices at cost recovery levels. Many utilities have turned to IBTs to meet simultaneous objectives of cost recovery and setting prices that are equitable and which encourage conservation. However, it is difficult to meet the cost recovery mandate while setting marginal prices at sufficiently high levels that reflect social costs or scarcity values.

information treatment indicates that consumers under-perceived (over-perceived) prices before the treatment.

I empirically test this model by leveraging quasi-exogenous variation in the level of price-related information provided to households on their monthly water bills. Specifically, I exploit a merger between two water utilities that use IBTs in the same metropolitan area. Before the merger, one group of households received a complete breakdown of how water use charges are determined from the price schedule, while the other did not receive any price information on monthly bills. After the merger both groups receive the same bills with better price information. The merger transitions one group of households to price information on monthly bills while holding rates fixed at pre-merger levels for both groups. Using a difference-in-differences approach, I find that providing consumers with price information on monthly water bills leads to an overall 3% or greater decrease in average water consumption, suggesting that price salience is important. Moreover, I validate the theoretical predictions from this model by finding differential treatment effects for consumers with historical water consumption levels below and above the price thresholds. Households above the threshold decrease water use, while those below increase water use. This suggests that while price information can be used to promote conservation, policy design should be careful in considering how the price schedule interacts with the ex-ante distribution of water consumption. To the best of my knowledge, this is the first quasi-experimental evidence that providing consumers with simple price information on monthly bills significantly impacts demand under IBTs.

In addition to price salience, there are other explanations for why consumers do not respond to marginal price under IBTs (Borenstein, 2009). Quantity salience is arguably a major impediment for consumers to make optimal decisions in this context, since consumers must have perfect knowledge of consumption choices throughout the billing period in order to successfully achieve consumption targets. Several studies find evidence that higher frequency feedback about consumption leads to decreases (Jessoe and Rapson, 2014; Strong and Goemans, 2014) or even increases in demand (Strong and Goemans, 2015; Wichman, 2016). Another issue concerns the cognitive burden that complicated price structures place on consumers. I do not focus on these important issues, but instead I provide evidence that price salience is also important in this context. Moreover the natural experiment analyzed in this chapter provides variation over price information but not quantity information, making it uniquely suited for analysis of the relationship between consumer behavior under IBTs and price salience.[2]

There is less evidence that price information can significantly affect consumption decisions under IBTs. Kahn and Wolak (2013) find evidence that educating electricity consumers about IBTs leads to an overall reduction in electricity demand through a field experiment with two California utilities, however the treatment effect confounds three sources of information that could impact consumer behavior under IBTs: in addition to providing price information, the treatment also lowers the cognitive burden for the consumer and provides information about the cost of operating specific appliances in the home.[3] By contrast, the simple, low-cost intervention of providing price information on monthly bills that I study, does not require voluntary participation in an education program, or adoption of a costly in-home display as in the case of quantity salience-based interventions.

Another related paper finds evidence that consumers increase water consumption in response to more frequent billing (Wichman, 2016). The author argues that a change from bimonthly to monthly billing provides information to households that allows them to

[2] It is unlikely that quantity salience and cognitive difficulties would differentially impact the households in this study area, since they receive similar consumption history information once per month and there are no bill calculator tools available through either utility in the study area.

[3] The treatment group in this field experiment undergoes a 30 minute education program that provides households with information about how their electricity bill is calculated, the marginal price paid based on recent electricity consumption, and how this information can be combined to estimate the cost of various end-uses of electricity in the home.

update their perception of the IBT price schedule. My natural experiment, by contrast, provides actual variation in price information provided on monthly bills, making it more suited to provide evidence of how consumers respond to price information when facing IBTs for water consumption.

The remainder of this chapter is organized as follows. First, I develop a model of consumer decision making under price uncertainty to motivate how we would expect consumers to respond to price information under IBTs. Following that I provide background about the empirical setting. I then develop an empirical test based on these model implications to investigate whether consumer response to price information is consistent with predictions from the model developed in this chapter. I also investigate whether consumer response is consistent with an alternative model of consumer behavior. After that I outline the main empirical strategy, and discuss these results. Finally, I discuss how the insights from the model developed in this chapter can be used to predict overall consumer response to price information and provide concluding remarks.

1.2 Decision Making Under Price Uncertainty

In this section, I develop a model of consumer decision-making under price uncertainty to model behavior in the case where consumers face IBTs but do receive price information on monthly bills. DellaVigna (2009) and Chetty, Loney and Kroft (2009) develop models in which consumers systematically misperceive prices in situations where prices are simple but not transparent. However, in settings where prices are complex, consumers may respond to expected marginal price if they do not have complete information about the price schedule. This is not an information asymmetry problem in that utilities intentionally keep price information private. When it comes to electricity and water consumption, consumers have the ability to become fully informed about prices, because price information is publicly available but not necessarily easy to find. Therefore, the time cost of acquiring price information will make it optimal for some consumers to remain uninformed.

Utilities can lower information costs for consumers by providing easy to understand information about the price structure on monthly bills. Doing so will cause some consumers to re-optimize, making consumption choices that are more consistent with their optimum under full information. I contrast how consumption choices compare under full price information and price uncertainty to illustrate how price information impacts consumption choices and how these impacts may vary with pre-treatment consumption levels.

Assume that consumers have preferences for water consumption, q, and a numeraire good, z. Consumers maximize utility subject to a budget constraint, which is a function of water charges, $B(q)$, and income, I. Assume consumers fully exhaust income, then the budget constraint can be rewritten as $z = I - B(q)$, where the price of numeraire consumption is normalized to unity.

For simplicity, consider a two tier IBT, illustrated in Figure 1.1, with marginal prices p_1 and p_2 and tier threshold k. Water charges are a piecewise function of quantity consumed, which depends on all three of these features of the price schedule.

$$B(q; k, p_1, p_2) = \begin{cases} p_1\, q, & \text{if } q \leq k \\ p_1\, k + p_2\,(q - k), & \text{if } q > k \end{cases}$$

Figure 1.1: Two tier increasing block tariff

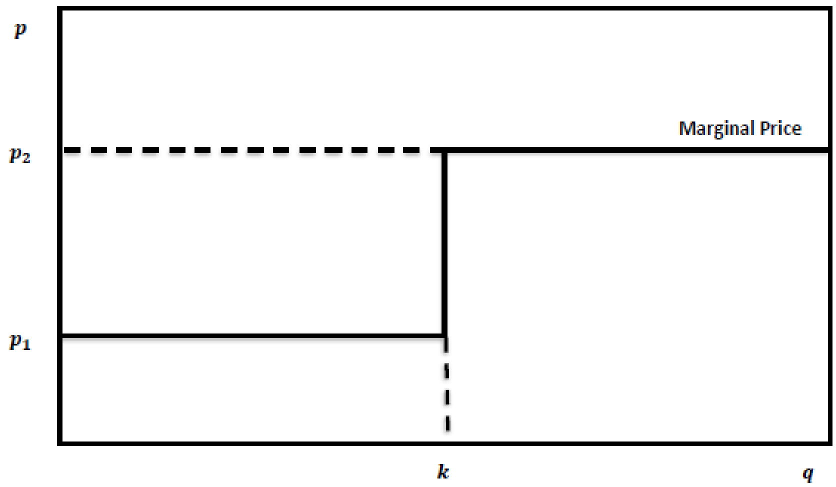

1.2.1 Full Price Information

Under full information the consumer's utility maximization problem is as follows:

$$\max_q U(q, I - B(q; k, p_1, p_2)) \tag{1.1}$$

Assume that the solution occurs at an interior optimum. The piecewise water charges introduce a kink into the budget constraint so that the optimum either occurs at a tangency point along one of the segments of the piecewise budget constraint or at the kink point. Denote the informed consumer's optimal consumption choice by q_I^*. The first order conditions are as follows:

$$q_I^* = \begin{cases} q^*(p_1), \text{ if } \frac{U_1(q^*, I - p_1 q^*)}{U_2(q^*, I - p_1 q^*)} = p_1 \\ k, \qquad \text{if } \frac{U_1(q^*, I - p_1 q^*)}{U_2(q^*, I - p_1 q^*)} < k < \frac{U_1(q^*, I - p_1 k - p_2 (q^* - k))}{U_2(q^*, I - p_1 k - p_2 (q^* - k))} \\ q^*(p_2), \text{ if } \frac{U_1(q^*, I - p_1 k - p_2 (q^* - k))}{U_2(q^*, I - p_1 k - p_2 (q^* - k))} = p_2 \end{cases} \tag{1.2}$$

1.2.2 Price Uncertainty around Tier Threshold

Consumers might face uncertainty about several aspects of the price schedule. For simplicity, assume that uninformed consumers know the price levels, p_1 and p_2, but are uncertain about k, the point at which water consumption becomes more costly.[4] Consumers likely have knowledge of the fact that they face an IBT rate schedule. Moreover, consumers have a rough idea of the price level from the average price, which can be obtained from the the total water use charges and the quantity consumed on the bill. Therefore it is the tier thresholds, about which consumers have the least information. Next, assume the consumer knows their current consumption throughout the billing cycle so that the only source of uncertainty in the consumer's decision problem is k.[5]

To model this decision, suppose that consumer i has beliefs about the location of the tier threshold, where $0 < \tilde{k}_i < \infty$ is a random variable representing these beliefs. Assume that $\mathbb{E}[\tilde{k}_i] = k$ and $Var(\tilde{k}_i) = \sigma_i^2$, since some consumers may be more informed about pricing than others and thus would have less uncertainty about the price schedule. I assume that price information decreases uncertainty for consumers, however I focus my analysis on the extreme case where price information on monthly bills resolves all uncertainty (i.e. $Var(\tilde{k}_i) = 0$). Let $\tilde{k}_i$ have a distribution, $F(\tilde{k}_i)$, that is defined on support $(0, \infty)$. The uninformed consumer maximizes expected utility, choosing optimal consumption q_U^*:

$$\max_q \mathbb{E}[U(q, I - B(q; \tilde{k}, p_1, p_2))] = \max_q \left[\int_0^\infty U(q, I - B(q; \tilde{k}, p_1, p_2)) dF(\tilde{k}) \right] \tag{1.3}$$

[4] Alternatively, we can model uncertainty over p_1, p_2, or $p_2 - p_1$. These models lead to different predictions about consumer heterogeneity. To model uncertainty about p_1, p_2, and k requires a joint distribution defined over the support $0 < p_1 < p_2 < \infty$, $0 < k$.

[5] Strong and Goemans (2014) develop a model of consumer decision making under quantity uncertainty and consider the implications of IBTs in this context. They find that consumers below the tier threshold will under-consume relative to full information and consumers above the tier threshold will either over-consume relative to full information or the effect will be ambiguous depending on risk preferences. These results will only exacerbate my findings.

where

$$B(q;\tilde{k},p_1,p_2) = \begin{cases} p_1 q, & \text{if } q \leq \tilde{k} \\ p_1\tilde{k} + p_2(q-\tilde{k}), & \text{if } q > \tilde{k} \end{cases}$$

Explicitly incorporating the piecewise water charges leads to the following:

$$\begin{aligned} &\max_q \mathbb{E}[U(q, I - B(q;\tilde{k},p_1,p_2))] = &(1.4)\\ &\max_q \left[P(\tilde{k} \geq q)\mathbb{E}[U(q, I-p_1 q)|\, \tilde{k} \geq q] + P(\tilde{k} < q)\mathbb{E}[U(q, I - p_1\tilde{k} - p_2(q-\tilde{k}))|\, \tilde{k} < q]\right] = \\ &\max_q \left[U(q, I-p_1 q) + F(q)\left[\int_0^q U(q, I - p_1\tilde{k} - p_2(q-\tilde{k}))dF(\tilde{k}) - U(q, I - p_1 q)\right]\right] \end{aligned}$$

The optimum is no longer a piecewise function, since the budget constraint is now a smooth and concave function.[6] To gain intuition into why this is the case, Figure 1.2 presents results from simulations for a Quasilinear utility function. Figure 1.2 (a) shows the actual piecewise budget constraint, $B(q;k,p_1,p_2)$, and expected budget constraints, $\mathbb{E}[B(q;\tilde{k}_i,p_1,p_2)]$, for various levels of uncertainty. In the case of quasilinear preferences, expected utility maximization is equivalent to utility maximization with respect to the expected budget constraint.[7] For higher levels of uncertainty, the budget constraint becomes less curved and as a result there is a larger distortion between the expected budget constraint and the kinked budget constraint near the tier threshold.

Figure 1.2 (b) compares consumption choices under price uncertainty to full information consumption choices across the distribution of preferences. This figure illustrates the three main takeaways from this model: 1) uninformed consumers under-consume just below the tier threshold relative to fully informed consumers, 2) uninformed consumers over-consume just above the tier threshold relative to fully informed consumers, and 3) the distortion between informed and uninformed consumption choices dissipates as distance from the tier threshold increases. Moreover, Figure 1.2 (b) shows that the

[6]See the Appendix A for First Order Conditions.

[7]Although I present results for quasilinear utility, these results generalize for additively separable utility functions.

Figure 1.2: Comparison of Expected Water Charges and Consumption Choices

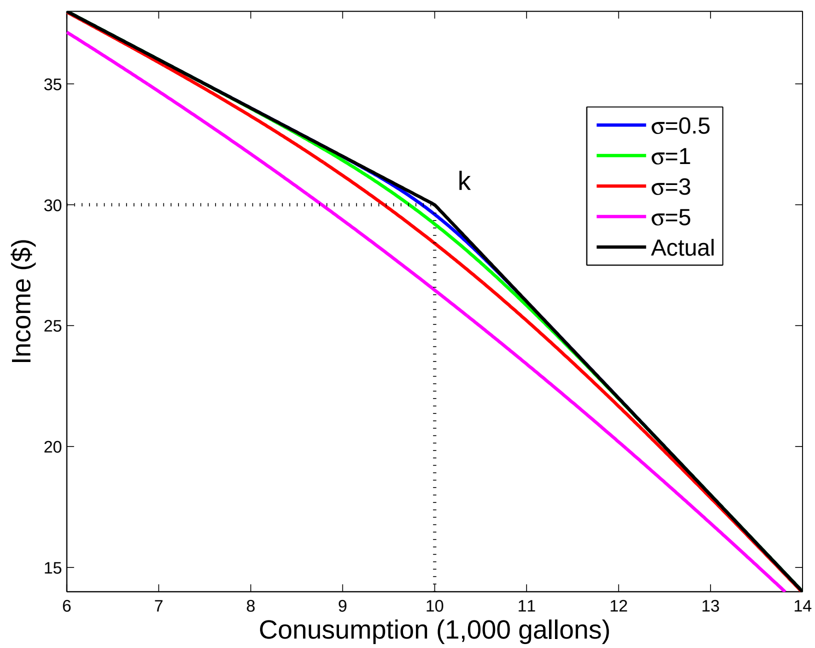

(a) Budget Constraint and Expected Budget Constraints

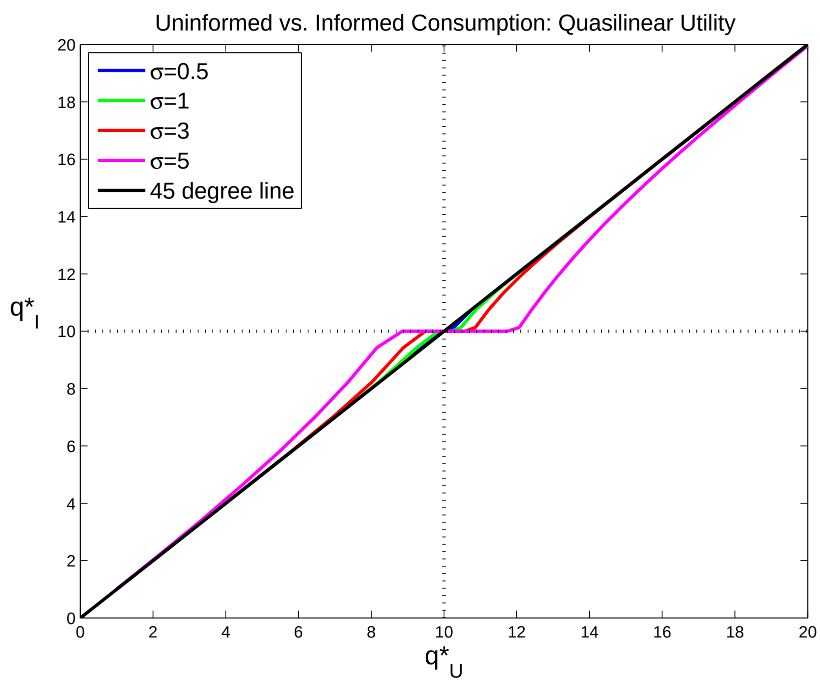

(b) q_I^* vs. q_U^*

Note: These simulations assume $k = 10{,}000$ gallons, $p_1 = \$2/1{,}000$ gallons, $p_2 = \$4/1{,}000$ gallons, quasi-linear preferences: $U_i(q, I - B(q; \tilde{k}_i, p_1, p_2)) = \alpha_i\, log(q) + I - B(q; \tilde{k}_i, p_1, p_2)$, and truncated Normal beliefs: $F(\tilde{k}_i) = \frac{\Phi\left(\frac{\tilde{k}-k}{\sigma_i}\right) - \Phi\left(\frac{-k}{\sigma_i}\right)}{1-\Phi\left(\frac{-k}{\sigma_i}\right)}$. Figure 1.2 (a) simulates the actual budget constraint and expected budget constraints corresponding to different levels of uncertainty. Figure 1.2 (b) simulates optimal choices for informed consumers and uninformed consumers over a range of α_i's. The Appendix presents additional simulation results.

difference between informed and uninformed consumption grows linearly as distance from the tier threshold decreases.

Now consider the effect of price information on consumption, which would lead an uninformed consumer to re-optimize consumption. This model suggests that consumers who were consuming just below the tier threshold will increase consumption in response, consumers who were consuming just above the tier threshold will decrease consumption, and consumers far from the tier threshold will have little or no change in consumption. Moreover, we might expect the effect of price information to increase linearly as distance from the tier threshold decreases. I use these predictions to develop an empirical strategy with testable hypotheses in section 1.4. I then develop a similar empirical strategy to test whether consumer response is consistent with an alternative model of consumer behavior, where consumers respond to average price rather than expected marginal price in the absence of price information (Ito, 2014; Kahn and Wolak, 2013). Following that, I consider how the direction of the overall treatment effect will depend on how the IBTs interact with the ex-ante distribution of consumption.

1.3 Background

This research takes advantage of a merger between the two major water utilities in the Reno metropolitan area in Northern Nevada. Prior to the merger these utilities provided different levels of price information to households on their monthly utility bills. The merger incorporates households from the "low information" utility into the service territory of the "high information" utility. As part of this process, low information households transitioned to the same monthly bills that high information households received throughout the study period, which included a full breakdown of water use charges and associated components of the price structure.[8] Importantly, the impetus for this merger is unrelated to differences in households served by either utility.[9] In 2014, the utilities announced that they would consolidate in order to take advantage of "integration opportunities that enhance economies of scale and/or other efficiencies" (Staff, 2009). In January 2015, the merger was completed and the high information utility began providing water service to both groups of households.

There are several aspects of this natural experiment that make it ideal for analyzing whether providing price information about IBTs affects consumer behavior. The utility merger lends itself to a difference-in-differences approach, where the treatment group are low information households that transition from not having the price schedule reported on monthly utility bills before the merger to having their water use charges broken down by tier consumption along with the associated marginal prices. The control group are high information households who get the full breakdown of water use charges throughout the study period. Moreover, as a stipulation of the merger, all rates are held fixed at pre-merger levels for a period of at least two years after the consolidation. This allows for an examination of whether low information households re-optimize water consumption with respect to the same rate schedule after receiving a price information treatment, relative to high information households, who experience no change in price information or rate schedule.

[8] In January 2014, the merger was announced. In December 2014, all of the low information households received a special mailer explaining the change in their service provider. Finally, in January 2015, low information households started receiving the same water bills as the high information households.

[9] Before 2015, municipal water service in the Reno metropolitan area was provided by Truckee Meadows Water Authority (TMWA), the high information utility, and Washoe County Department of Water Resources (DWR), the low information utility. TMWA is a not-for-profit, community-owned water utility that is overseen by elected officials and citizen appointees. DWR is a division of the Washoe County local government. TMWA provided water service for the majority of residents in the metropolitan area prior to 2015 while DWR was responsible for areas outside of TMWA's service area, including several outlying regions that experienced high levels of growth.

1.3.1 Billing Information

The primary difference in monthly bills before the merger relates to how water use charges are presented. Figures 1.3 and 1.4 show components of the water bills for the low information and high information utilities. The low information utility provides total water use and total water use charges to households without any mention of the IBTs used to determine these charges. The high information utility provides a complete breakdown of water use charges including water consumption, marginal prices, tier thresholds limits, and total water use charges for each tier. Both utilities provide information about their rates on the Internet, however this information is only provided on monthly bills for high information households, substantially lowering the cost of acquiring this information.

There are other dimensions on which consumers can receive information that could affect water use decisions. Empirical evidence finds that providing households with high frequency feedback about consumption significantly impacts demand (Jessoe and Rapson, 2014; Wichman, 2016; Strong and Goemans, 2015). However, the treatment and control group in this study receive similar quality information along these other dimensions. Both groups of households receive comparable information about their consumption history on monthly bills; other bill components such as total monthly charges, fixed fees, and surcharges are summarized in a similar fashion. Moreover, meters are read and water consumption is billed on a monthly basis for both utilities, who report total water consumption at the monthly level, making household groups exposed to the same frequency of consumption information. Finally, both utilities provide price schedule information through a link from their websites.

Figure 1.3: Low Information Utility Monthly Bill

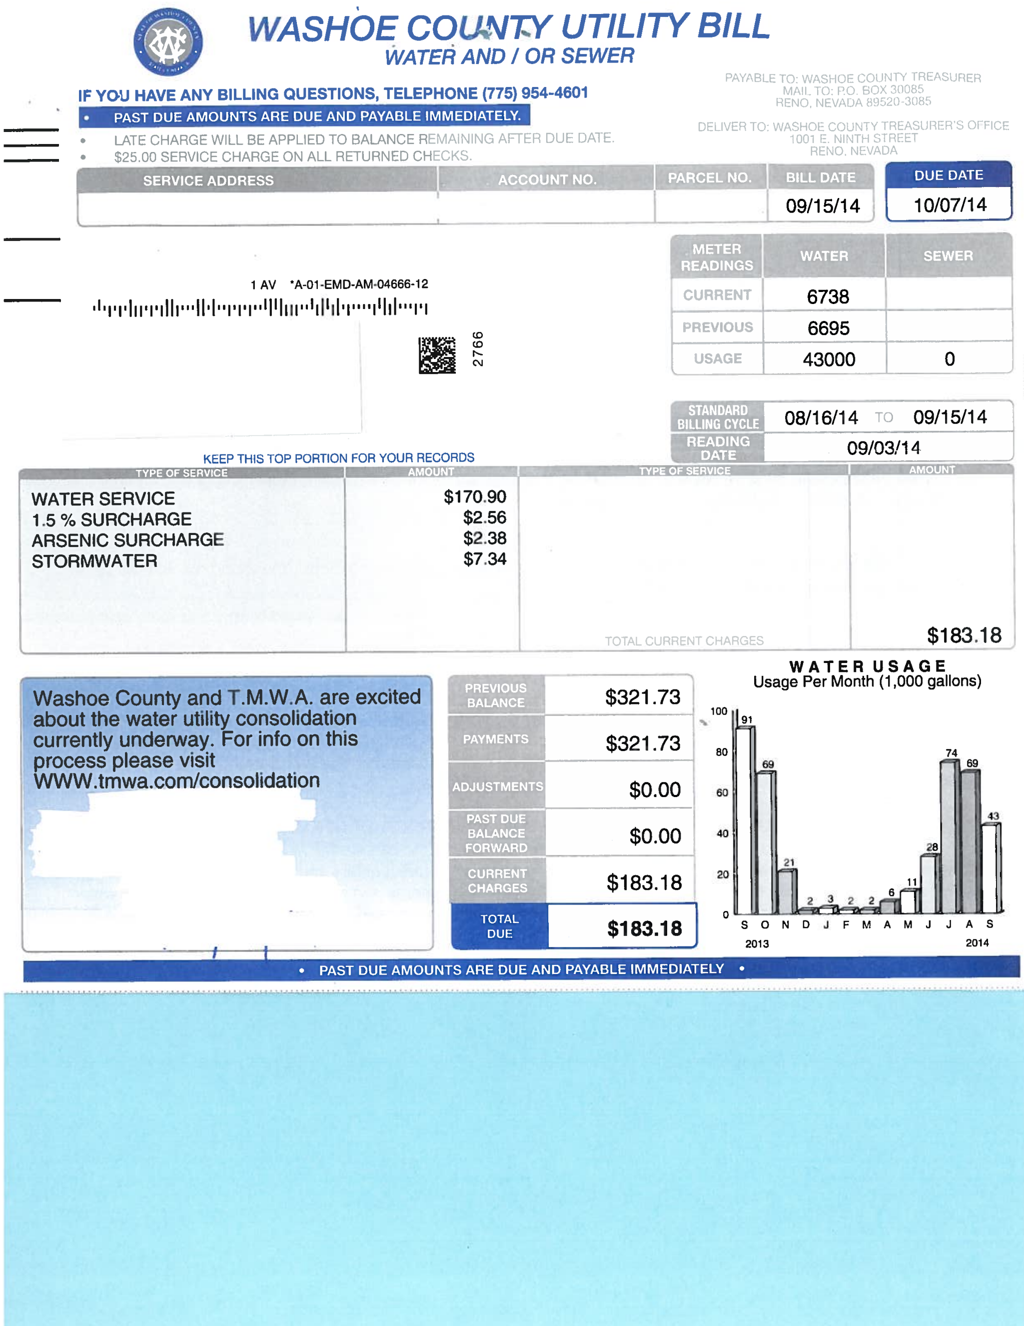
WASHOE COUNTY UTILITY BILL
WATER AND / OR SEWER

IF YOU HAVE ANY BILLING QUESTIONS, TELEPHONE (775) 954-4601

- PAST DUE AMOUNTS ARE DUE AND PAYABLE IMMEDIATELY.
- LATE CHARGE WILL BE APPLIED TO BALANCE REMAINING AFTER DUE DATE.
- $25.00 SERVICE CHARGE ON ALL RETURNED CHECKS.

PAYABLE TO: WASHOE COUNTY TREASURER
MAIL TO: P.O. BOX 30085
RENO, NEVADA 89520-3085

DELIVER TO: WASHOE COUNTY TREASURER'S OFFICE
1001 E. NINTH STREET
RENO, NEVADA

SERVICE ADDRESS	ACCOUNT NO.	PARCEL NO.	BILL DATE	DUE DATE
			09/15/14	10/07/14

1 AV *A-01-EMD-AM-04666-12

2766

METER READINGS	WATER	SEWER
CURRENT	6738	
PREVIOUS	6695	
USAGE	43000	0

STANDARD BILLING CYCLE	08/16/14 TO 09/15/14	
READING DATE	09/03/14	

KEEP THIS TOP PORTION FOR YOUR RECORDS

TYPE OF SERVICE	AMOUNT	TYPE OF SERVICE	AMOUNT
WATER SERVICE	$170.90		
1.5 % SURCHARGE	$2.56		
ARSENIC SURCHARGE	$2.38		
STORMWATER	$7.34		
		TOTAL CURRENT CHARGES	$183.18

Washoe County and T.M.W.A. are excited about the water utility consolidation currently underway. For info on this process please visit WWW.tmwa.com/consolidation

PREVIOUS BALANCE	$321.73
PAYMENTS	$321.73
ADJUSTMENTS	$0.00
PAST DUE BALANCE FORWARD	$0.00
CURRENT CHARGES	$183.18
TOTAL DUE	$183.18

• PAST DUE AMOUNTS ARE DUE AND PAYABLE IMMEDIATELY •

Note: The low information utility provided customers with the total consumption over the course of the month (usage), the total water charges (water service) which include the fixed monthly fee and consumption charges, other fixed monthly charges, and a one year water use history graphic.

Figure 1.4: High Information Utility Monthly Bill

TRUCKEE MEADOWS WATER AUTHORITY
www.tmwa.com

Service Address:

Page	1
Account No.	
Billing Date	08/07/15
Next Meter Read Date	09/03/15

Previous Statement Balance	Payments	Balance Forward (=)	Additional Activity (+ or -)	New Water Charges (+)	NEW Charges Past Due On	New Balance (=)
$154.57	$154.57 CR	$0.00	$0.00	$178.54	08/28/15	$178.54

If "Balance Forward" is not a credit, this amount may be past due. Please pay immediately to avoid collection activity.

Previous Statement Balance		$154.57	
Payments	Payment - THANK YOU	$154.57 CR	
	Total Payments	**$154.57 CR**	
Balance Forward			
	Total Balance Forward		**$0.00**
New Water Charges	Water Charges (see detail on back)	$175.90	
	Regional Water Mgmt Fee - 1.5%	$2.64	
	Total New Water Charges		**$178.54**
	New Balance		**$178.54**

Keep up the good work in saving 10% on your water use compared t o 2013! We are on track to achieve our goal of saving 5,000 acre feet of water - that's 1,629,266,900 gallons. For 10 easy ways to save 10%, visit tmwa.com/save.

Customer Service
775-834-8080
Monday - Friday
7:30 a.m. - 5:30 p.m.

Payment Options
Please visit www.tmwa.com
or call 775-834-8080.

Main Office and Payment Location
1355 Capital Blvd., Reno, NV 89502
Monday - Friday
8:00 a.m. - 5:00 p.m.

TI015080 7-12549-000000622

Water Charges
Avg Use per Month (1,000 gallons)

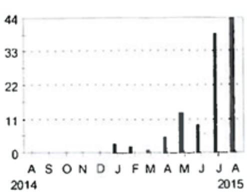

USAGE HISTORY	No. DAYS	GALLONS	AVG USE PER DAY	AVG COST PER DAY
THIS MONTH	30	44,000	1,467	5.86
LAST MONTH	32	39,000	1,219	4.76
LAST YEAR	0			0.00

Meter	Rate Code	No. of Days	From Date	To Date	Prior Reading	Current Reading	Reading Type	Factor	Billed Usage
03970534	RMWD1	30	07/06	08/05	6869	6913	Actual	1	44

Residential Customer Charge	3/4" Service	$17.43
Tier 1 Water Usage Charge	7 THOUSAND GALLONS X 2.62000	$18.34
Tier 2 Water Usage Charge	14 THOUSAND GALLONS X 3.27000	$45.78
Tier 3 Water Usage Charge	20 THOUSAND GALLONS X 3.93000	$78.60
Tier 4 Water Usage Charge	3 THOUSAND GALLONS X 5.25000	$15.75
TOTAL CURRENT CHARGES - Water Charges		**$175.90**

Note: The high information utility provided customers with the total consumption over the course of the month (billed usage), the total water charges (total current charges) broken down into the fixed monthly fee (residential customer charge) and consumption charges broken down by tier including tier consumption and the marginal price, other fixed monthly charges, and a one year water use history graphic.

Table 1.1: Property Characteristics and Demographics

	Low Info	High Info	Difference
Consumption/Day (1,000 gallons)	0.50	0.40	0.11**
Consumption/Day Summer (1,000 gallons)	0.79	0.60	0.19**
Consumption/Day Winter (1,000 gallons)	0.28	0.24	0.04*
Marginal Price ($/1,000 gallons)	2.87	2.35	0.52**
Year Built	1997.2	1983.3	13.8***
Appraised Value ($1,000)	296.54	270.12	26.43
Yard Size (acres)	0.46	0.57	-0.11
Bedrooms	3.38	3.32	0.06
% Male	49.91	49.82	0.09
Median Age	40.80	38.25	2.55
% White	86.04	78.71	7.33***
% Black	1.28	2.23	-0.95***
% Hispanic	10.64	19.35	-8.71***
Avg. HH Size	2.68	2.61	0.07
% Owner Occupied	81.33	64.97	16.36***
Income/Capita	40,400	32,698	7,702*
Median Home Value	392,334	309,819	82,515*
% College Degree	44.64	37.42	7.22

Note: The ACS data are taken from the 2006-2010 ACS 5 year estimates. ACS data has associated sampling error (reported as margins of error), which has not been accounted for in these balance tables. The reported significant differences therefore are likely overstated. Robust standard errors are clustered at the meter reading route level for water use and property characteristics, and at the census tract level for demographic information. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

1.3.2 Data

The primary data used for this analysis are single family residential billing records from TMWA and DWR that span the period from January 2010 to August 2016, which include monthly water consumption, the billing dates corresponding to each record, rate information, meter reading routes, and the spatial location of each water service. The unit of analysis is the household, which is defined to be an unique water service-customer combination. Only customers that remained at the same water service from at least two years before treatment until the end of the study period were included in the final dataset. These sample requirements help limit sorting between utilities and remove some renters from the sample, who may not pay water bills directly and tend to move more frequently. The data were spatially matched with geocoded data containing structural characteristics for each home from the Washoe county assessor and demographic information the American Community Survey (ACS) corresponding to 2010 Census tracts. I used NOAA weather station daily data to create monthly weather averages of temperature and precipitation that correspond to the billing dates for each water service-bill combination.

The final sample consists of an unbalanced panel of 60,833 households, 46,472 from the high information utility and 14,361 from the low information utility, for a total of 4,460,397 observations. Table 1.1 presents summary statistics for the final sample. Although there are significant differences in demographics, property characteristics, and water consumption between low information and high information households, my empirical strategy does not rely on balance between these two groups.

1.4 Heterogeneous Response to Information

In this section, I build an empirical strategy that can be used to test whether consumer response to price information on monthly utility bills is consistent with the three main takeaways from the model. This approach relies on the assumption that household preferences remain constant over the study period. Under this assumption, providing consumers with price information can identify consumer re-optimization by comparing household consumption before and after the price information treatment. My empirical approach uses the location of pre-treatment consumption relative to the tier thresholds to identify heterogeneity in consumption response. Since low and high information households face different IBT rate schedules (i.e. the tier thresholds are located at different

Figure 1.5: Ex-ante distribution of consumption vs. IBT: Low Information Utility

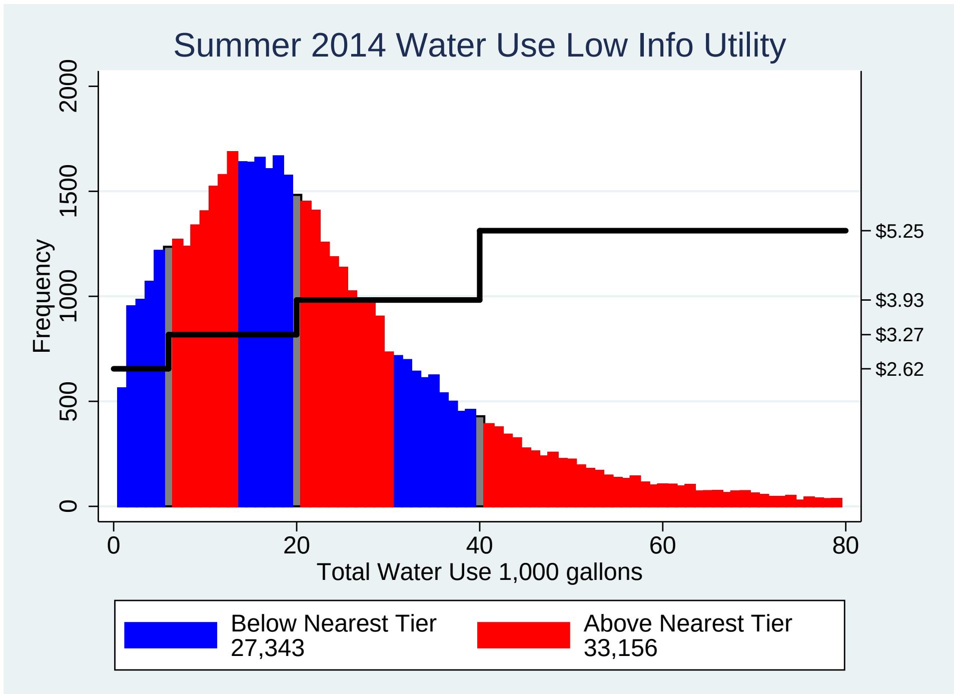

Note: This figure compares the histogram of summer (May-September) 2014 consumption (the last year before the merger) for the Low Information utility to the price schedule. This indicates the relative number of households below the nearest tier to those above the nearest tier.

consumption levels and marginal prices are also different.), I focus this analysis on the low information households only.[10]

The first step of this approach is to separate all pre-treatment observations based on location with respect to the nearest tier threshold. Figure 1.5 shows the interaction of summer 2014 water consumption and the IBT schedule for low information households. The blue consumption levels are observations that are below the nearest tier threshold in 2014, and the red consumption levels are observations that are above the nearest tier threshold in 2014. Households below the nearest tier threshold will experience a differ-

[10]Low and high information households that make comparable pre-treatment consumption decisions will be located on different segments of their respective rate schedules. It is unlikely that households with similar consumption or similar bill amounts have the same preferences.

ential treatment effect compared to households above the nearest tier threshold. Next, I explicitly control for distance from tier threshold, since households with consumption that is far away from any tier threshold should have a smaller response than households with consumption that is near a tier threshold.[11] This approach is similar to the difference-in-discontinuities design developed by Grembi, Nannicini and Troiano (2012).[12] Finally, for simplicity, I use a linear distance specification to match the function of response decay portrayed in Figure 1.2 (b). The actual function of decay, however, may be non-linear and the functional form could vary below and above the tier thresholds. The difference-in-discontinuities model is as follows:

$$ln(Y_{it}) = \beta_1\, b_{is}\, \mathbb{1}_{\text{Post}} + \beta_2\, a_{it} \mathbb{1}_{\text{Post}} + \beta_3\, d_{is}\, b_{is}\, \mathbb{1}_{\text{Post}} + \beta_4\, d_{is}\, a_{is}\, \mathbb{1}_{\text{Post}} + \alpha_i + \lambda_t + W_{it}\gamma + \varepsilon_{it} \quad (1.5)$$

where Y_{it} is monthly water consumption divided by the days on the water bill for household i in billing period t, b_{is} is an indicator for household i equal to one if pre-treatment consumption in month s is below the nearest tier threshold and zero otherwise, a_{is} is an indicator for household i equal to one if pre-treatment consumption in month s is above the nearest tier threshold and zero otherwise, d_{is} is distance from the nearest tier threshold of pre-treatment consumption for household i, in month s, $\mathbb{1}_{\text{Post}}$ is an indicator for billing periods after the merger, α_i are household fixed effects, λ_t are a full set of month-by-year fixed effects, and W_{it} are a vector of weather controls including average temperature and precipitation in inches, and ε_{it} is the idiosyncratic error term. In order to account for possible correlation in the errors within neighborhoods, I cluster the

[11] I also estimate a discrete analogue of this model that uses consumers with pre-treatment consumption that is beyond some bandwidth distance from any tier threshold as a controls for consumers in one of two treatment groups: 1) pre-treatment consumption is just below a tier threshold or 2) pre-treatment consumption is just above a tier threshold. See Appendix A for these results and more details.

[12] The equivalent identification assumption to those described in Grembi, Nannicini and Troiano (2012) is that household preferences are constant across the merger.

standard errors at the meter reading route level.[13]

If information leads consumer just below the tier threshold to increase consumption and households just above the tier threshold to decrease consumption, then $\beta_1 > 0$ and $\beta_2 < 0$. Secondary predictions from the model are that $\beta_3 < 0$ and $\beta_4 > 0$, since the effect of information on consumption should decrease as distance from the tier threshold increases.

1.4.1 Results

I use three different approaches to determine treatment status, which is based on some measure of consumption before the merger. I allow treatment status to vary for each household by calendar month, since consumption exhibits highly seasonal patterns. This implies that the same household could be below the nearest tier threshold during a given month, and above that tier threshold the next month. First I use mean consumption within household-billing month across all pre-treatment years (2010-2014) where:

$$b_{is} = \begin{cases} 1, & \text{if } \frac{1}{5}\sum_{y=2010}^{2014} Y_{isy} < k_{\text{nearest}} \\ 0, & \text{otherwise} \end{cases}$$

$$a_{is} = \begin{cases} 1, & \text{if } \frac{1}{5}\sum_{y=2010}^{2014} Y_{isy} > k_{\text{nearest}} \\ 0, & \text{otherwise} \end{cases}$$

for k_{nearest} representing the location in thousands of gallons of the nearest tier threshold. My second specification uses this approach, but limits the pre-merger period to 2013 and 2014 only. Finally, I use consumption in 2014, the year before the merger, to determine treatment status. The advantage of using the year before treatment is that it more closely represents household preferences for consumption at the time of the merger. However,

[13] See Section 1.5 for more details on this clustering approach.

Table 1.2: Low Information Heterogeneous Response to Information Difference-in-differences Results

	(1) 2010-2016	(2) 2013-2016	(3) 2014-2016
β_1	-0.0143	0.0059	0.0133***
	(0.0084)	(0.0051)	(0.0043)
β_2	-0.0365***	-0.0188***	-0.0088*
	(0.0082)	(0.0051)	(0.0044)
β_3	-0.0001	-0.0011	-0.0016***
	(0.0010)	(0.0007)	(0.0005)
β_4	0.0021***	0.0010***	0.0008**
	(0.0003)	(0.0003)	(0.0003)
Mean Use Below	20.88	20.85	20.25
Mean Use Above	28.62	29.50	29.17
Mean Dist Below	4.15	4.26	4.46
Mean Dist Above	7.19	7.54	7.67
Within R-squared	0.0048	0.0034	0.0035
Households	14,361	14,352	14,344
Observations	468,897	284,166	212,532

Note: The dependent variable is log of monthly water consumption divided by the number of days on the bill. The coefficients can be interpreted as the percentage change in consumption. The sample is limited to summer water use (May-September) for Low Information households only. Controls for all regressions include household FEs, month-by-year FEs, Average Temperature, and Precipitation in Inches. Robust standard errors clustered at the meter reading route level are reported in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

using only one year of pre-treatment data makes this approach more vulnerable to mean reversion. Therefore I use three approaches for comparison. Moreover, I limit the study period to summer months when the information effect is most dramatic and I limit the study period to 2013-2016 for model 2 and 2014-2016 for model 3.

Table 1.2 presents results from equation 1.5. Models 2 and 3 have positive signs for the below tier group and the coefficient is significant in model 3. All three models yield negative and significant coefficients for the above tier group. These results indicate

Figure 1.6: Predicted Response to Information by Distance from Tier Threshold

Note: This figure plots the predicted response to price information for consumers below and above the nearest tier threshold from Equation 1.5.

that after controlling differentially for distance from tier threshold, there is a significant discontinuity in the consumption response to price information on monthly bills. Moreover, this discontinuity has the predicted directional effect of increasing consumption for households just below the tier threshold and decreasing consumption for households just above the tier threshold. Furthermore, the distance coefficients have the expected signs, which indicate that the below tier treatment effect becomes less positive as distance below the tier threshold increases and the above tier treatment effect becomes less negative as distance above the tier threshold increases. My preferred specification in column 3 indicates that being below a tier threshold in 2014 leads to a 1.3% increase in consumption after the merger, and being above a tier threshold in 2014 leads to a 0.88% decrease in consumption after the merger.

Figure 1.6 shows the predicted coefficients from this model. They are consistent with the predicted behavior in Figure 1.2 (b), where the positive difference in the case of below tier consumers and the negative difference in the case of above tier consumers converge toward the 45 degree line. These results provide evidence that consumer response to price information on monthly bills is consistent with price uncertainty around the tier thresholds.

1.4.2 Predicting the Direction of the Treatment Effect

The interaction between the ex-ante distribution of consumption and the rate schedule can be used to provide insight into the direction of the overall treatment effect. There are three factors that combine to determine whether price information is expected to lead to an overall increase or decrease in consumption. First, the treatment effect will depend on whether there are relatively more households below or above the tier thresholds. Second, the relative proximity of households below versus above the tier thresholds will influence the share of each group that is expected to respond strongly to information. Third, the function of decay in the treatment effect might be different for households above the tier threshold than households below the tier threshold as Figure 1.6 suggests.

Figure 1.5 shows the distribution of summer consumption in 2014 for the low information households along with the IBT rate schedule. Due to the location of the tier thresholds combined with the fact that water consumption generally has a right-tailed distribution there are more observations above the nearest tier threshold than below the nearest tier threshold. There are approximately 27,000 observations below the nearest tier threshold and 33,000 observations above the nearest tier threshold. However, much of the right tail is located at a farther distance from the tier threshold, which would lead some of these observations to be less affected by price information. The average distance from the nearest tier threshold is four thousand gallons for the below group and seven thousand gallons for the above group. Overall, examining the interaction between the ex-ante distribution of consumption and the rate schedule is an important first step in predicting how price information is likely to affect average consumption.

1.4.3 Alternative Model of Consumer Behavior

Figure 1.7: Low Information Utility Average Total Price and Marginal Price Schedules

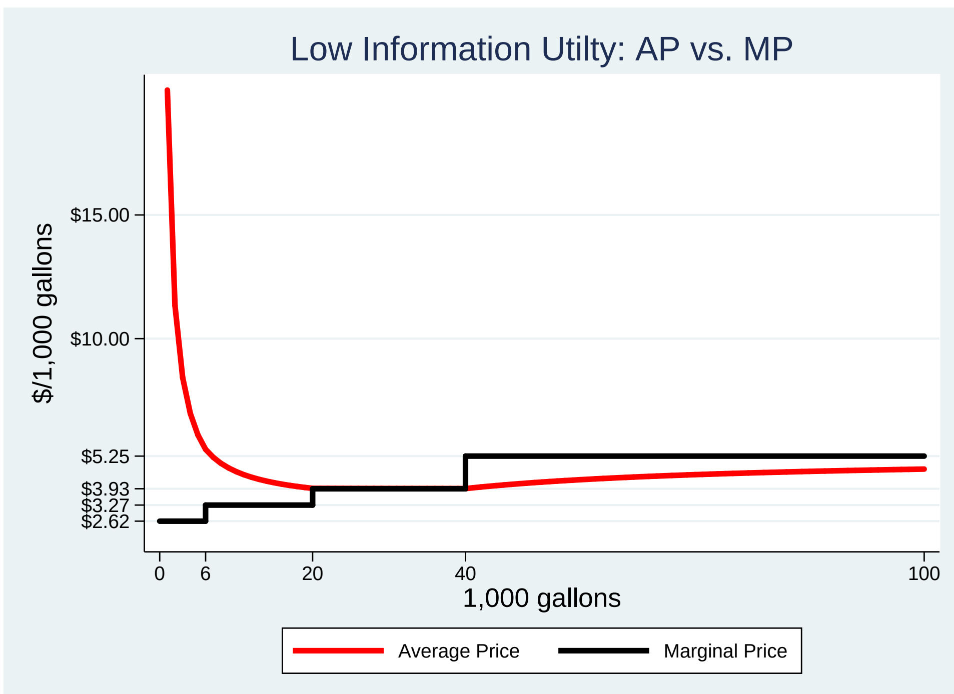

As an additional robustness check to provide further support that consumer response is consistent with the model of consumer behavior developed in this chapter, I show that consumer response is not consistent with predictions from an alternative model of consumer behavior under IBTs. Kahn and Wolak (2013) and Ito (2014) argue that consumers respond to average price rather than marginal price in the absence of adequate price information about IBTs. This model assumes that consumers have no awareness

Table 1.3: Alternative Model of Low Information Heterogeneous Response to Information Difference-in-differences Results

	(1)	(2)	(3)
	2010-2016	2013-2016	2014-2016
β_1	-0.0002	0.0272***	0.0233***
	(0.0068)	(0.0059)	(0.0062)
β_2	0.0466***	0.0315***	0.0214***
	(0.0062)	(0.0047)	(0.0051)
R-squared	0.8230	0.8564	0.8634
Households	14,361	14,361	14,361
Observations	468,944	284,427	213,340

Note: The dependent variable is log of monthly water consumption divided by the number of days on the bill. The coefficients can be interpreted as the percentage change in consumption. The sample is limited to summer water use (May-September) for Low Information households only. Controls for all regressions include household FEs, month-by-year FEs, Average Temperature, and Precipitation in Inches. Robust standard errors clustered at the meter reading route level are reported in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

of the fact that they face an IBT price schedule, but rather a continuously increasing price of water consumption. To illustrate this issue, Figure 1.7 shows the IBTs for the low information households and the average price schedule, which is the total of all water charges (this was called "water service" on the low information utility bills) divided by consumption.[14] If consumers do indeed respond to average price rather than expected marginal price before receiving price information on monthly bills, then providing consumers with this information should lead to a re-optimization that is consistent with a transition from average price response to marginal price response.

Figure 1.7 indicates that 1) for $Y < 20$, $AP > MP$, 2) for $20 \leq Y < 40$, $AP = MP$, and 3) for $Y \geq 40$, $AP < MP$. This motivates a similar approach to the one used in section 1.4. Consumers with $AP = MP$ before the merger will serve as the control group,

[14]Total water charges include water consumption charges and a fixed monthly fee. I do not incorporate other other surcharges from the bill, although doing so does affect the results. This definition of average price is consistent with how Kahn and Wolak (2013) define average price. Ito (2014), however, does not include fixed monthly fees in his concept of average price. I can also reject that consumer response is consistent with re-optimization from this second definition of average price.

because they should not change consumption in response to price information under this model of consumer behavior. By contrast, the two treatment groups with $AP > MP$ and $AP < MP$ before the merger should have a differential response to price information. The model is as follows:

$$lnY_{it} = \beta_1 \mathbb{1}_{AP>MP,s} \mathbb{1}_{\text{Post}} + \beta_2 \mathbb{1}_{AP<MP,s} \mathbb{1}_{\text{Post}} + \alpha_i + \lambda_t + W_{it}\gamma + \varepsilon_{it} \tag{1.6}$$

where $\mathbb{1}_{AP>MP,s}$ is an indicator for households with $AP > MP$ in month s before the merger, $\mathbb{1}_{AP<MP,s}$ is an indicator for households with $AP < MP$ in month s before the merger, and $\mathbb{1}_{\text{Post}}$ is an indicator for billing periods after the merger. If $\beta_1 > 0$ and $\beta_2 < 0$, then consumer behavior is consistent with a transition from average price to marginal price response.

For this test of alternative models of consumer behavior, I used the same three methods of defining pre-treatment consumption as with previous results. These results, shown in Table 1.3, indicate that consumer response is not consistent with this alternative hypothesis of consumer behavior under IBTs. Although the coefficients for the below treatment group are positive and significant in two out of three specifications, the coefficients for the above treatment group are consistently positive and significant across all three specifications. These results provide further support for the model of consumer behavior under IBTs described in this chapter.

1.5 Main Empirical Strategy

Next I estimate the overall average treatment effect (ATE) of providing price information on monthly utility bills using a difference-in-differences strategy, where households served by the low information utility before the merger are in the treatment group, households served by the high information utility are in the control group, and treatment occurs when the merger takes place. The standard difference-in-differences approach regresses the log of average daily water consumption for household i in billing period t on a set of household fixed effects, month-by-year fixed effects, and an indicator for low information households after the merger as follows:

$$ln(Y_{it}) = \beta\, \mathbb{1}_{\text{Low Info, Post}} + \alpha_i + \lambda_t + W_{it}\gamma + \varepsilon_{it} \tag{1.7}$$

where Y_{it} is monthly water consumption divided by the days on the water bill for household i in billing period t, $\mathbb{1}_{\text{Low Info, Post}}$ is an indicator variable equal to one for low information households in billing periods after the merger and zero otherwise, α_i are household fixed effects, λ_t are a full set of month-by-year fixed effects, and W_{it} are a vector of weather controls including average temperature and precipitation in inches, and ε_{it} is the idiosyncratic error term. I allow for robust standard errors, which I cluster at the meter reading route level.[15]

In addition to the standard exogeneity assumption, this approach assumes that low information and high information households have parallel trends in average consumption. The merger causes a parallel shift in consumption trends, which allows for identification of the ATE, β, as the difference between consumption trends after the merger relative to the difference before the merger. There are several concerns for this identification strategy, which include differential trends between low information and high information households and other confounding policies. I discuss these issues in detail and enhance my empirical strategy to address these concerns.

[15]Previous studies have used meter reading route fixed effects to control for unobserved neighborhood-specific characteristics that may influence water consumption (Ferraro, Miranda and Price, 2011). Standard errors are also likely to be correlated within meter reading routes. Some possible sources of correlation could include spatial variation in water use due to regional weather patterns as well as economic shocks to consumers residing in the same housing developments. There are 95 meter reading routes, which ensures that there are a sufficiently high number of clusters.

Figure 1.8: Differences in Pretreatment Seasonality and Trends

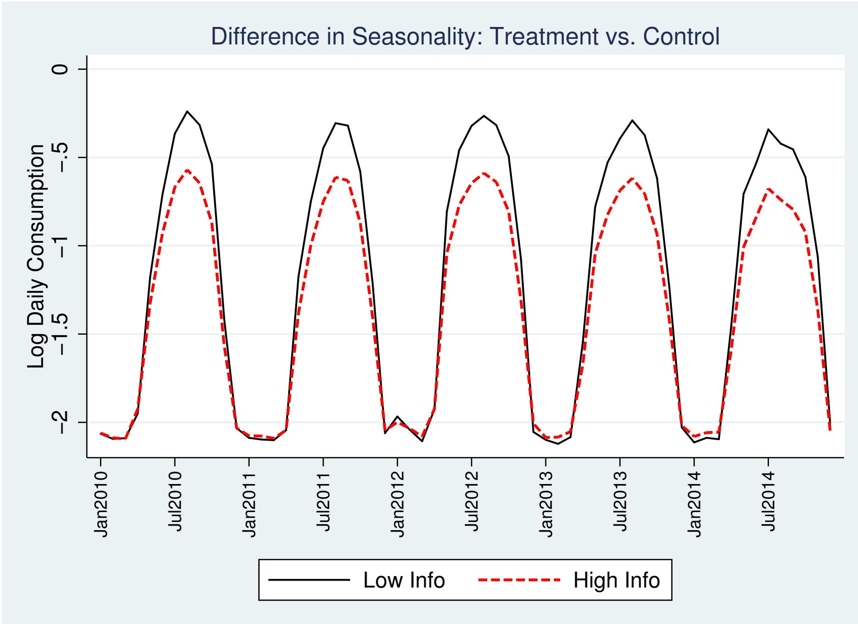

(a) Average Daily Water Use: Control vs. Treatment

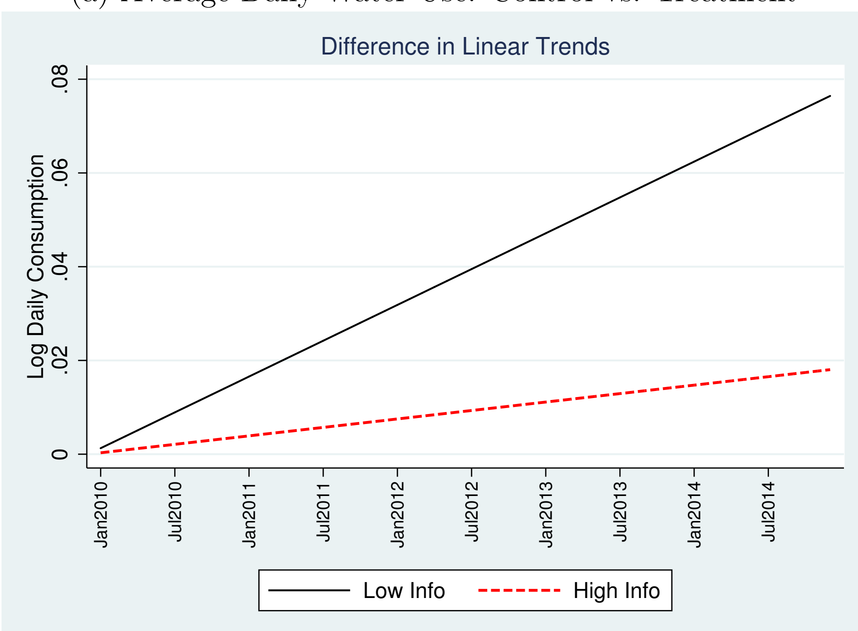

(b) Differences in Estimated Linear Trends: Control vs. Treatment

Note: These figures demonstrate differences in pre-treatment trend and seasonality between low and high information households. Figure (a) plots average daily consumption (monthly consumption divided by the number of days on the water bill) for both groups. Figure (b) plots estimated linear time trends for low and high information households from a model that regresses the log of consumption divided by the number of days on the water bill on separate linear time trends, separate month fixed effects, weather controls, and household fixed effects. The sample is limited to pre-merger time periods. Robust standard errors are clustered at the meter reading route level.

There are indeed differences in pre-treatment trends as well as seasonality between the low information and high information households. Figure 1.8 (a) shows average monthly consumption divided by the days on the water bill for low information and high information households, grouped by billing months over the pre-merger period. Although both utilities have highly seasonal water consumption patterns, there are differences in seasonal water consumption between the two utilities. Average consumption is much higher for low information households during the summer and exhibits more dramatic peaks in the height of the irrigation season. This motivates the addition of separate month fixed effects for the low information households. Next, I allow for linear differences in trend between treatment and control. Figure A.2 in Appendix A estimates a more flexible difference in trends over the pre-treatment period and indicates that a linear difference in trends is probably appropriate. Figure 1.8 (b) shows the estimated linear trends for each utility over the pre-merger periods. Consumption for the High Information group has a small, but insignificant increasing trend over this period. By contrast, low information consumption has a significant increasing linear trend. Given these significant differences, a model that does not allow for separate linear trend and seasonality would understate the impact of the merger on water consumption.

The identification assumption in a difference-in-differences model that allows for a separate linear trend for the low information households is that low information and high information households have parallel growth in average consumption rather than parallel paths of average consumption. Furthermore, the identification assumptions could be further relaxed to allow for differences in seasonality between low and high information households that is fixed over time. The relaxed difference-in-differences model is as follows:

$$ln(Y_{it}) = \beta\, \mathbb{1}_{\text{Low Info, Post}} + \alpha_i + \lambda_t + \mu\, \mathbb{1}_{\text{Low Info}} * t + \sum_{s=1}^{12} \nu_s \mathbb{1}_{\text{Low Info,Month}\, s} + W_{it}\gamma + \varepsilon_{it} \quad (1.8)$$

where $\mathbb{1}_{\text{Low Info}} * t$ is an indicator for low information households interacted with a linear time trend and $\mathbb{1}_{\text{Low Info,Month}\, s}$ is an indicator equal to one for low information households if their billing month occurs during the sth calendar month. Conditional on these relaxed assumptions β is the causal effect of providing price information to low information households on monthly utility bills.

1.6 Main Empirical Results

Column 1 of Table 1.4 presents results from the full study period from 2010-2016. The overall treatment effect is negative and significant at the 1% level. These results suggest that providing price information on monthly bills to low information households leads to a 3.2% decrease in average consumption. This amounts to a 500 gallon per household per month reduction on average. I also separately estimate the results using only the summer months from May to September in Table 1.5, which suggests a larger treatment effect of a 4% decrease in average consumption, amounting to more than 1,000 gallons per household per month. The fact that the treatment effect is larger during the summer is not surprising, since consumers have a much stronger incentive to monitor water consumption during periods where high levels of irrigation significantly drive up the total cost of water consumption.

Another issue for identification pertains to drought restrictions that occurred during the summer of 2014, before the merger, and during the summer of 2015, after the merger.[16] During this period there was an historic drought throughout the Southwestern

[16] Drought restrictions were implemented late July-September 2014 and April-September 2015.

Table 1.4: Main Difference-in-differences Results

	(1) All Periods	(2) No Drought Periods	(3) No 2014	(4) No 2015	(5) No 2014-2015
β (All Periods)	-0.0319***	-0.0409**	-0.0250*	-0.0531**	-0.0521**
	(0.0121)	(0.0162)	(0.0130)	(0.0223)	(0.0257)
Low Info Trend	0.0009***	0.0010***	0.0007**	0.0010***	0.0009**
	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0004)
Low Info Avg. Use	15.72	15.18	15.71	15.72	15.71
Within R^2	0.0089	0.0090	0.0082	0.0089	0.0081
Households	60,833	60,833	60,833	60,833	60,833
Observations	4,460,397	3,897,390	3,742,734	3,743,096	3,025,433

Note: The dependent variable is log of monthly water consumption divided by the number of days on the bill. The coefficients can be interpreted as the percentage change in consumption. Results based on the full sample. Controls for all regressions include household FEs, month-by-year FEs, Average Temperature, Precipitation in Inches, and a separate trend and month FEs for Low Information households. Robust standard errors clustered at the meter reading route level are reported in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

U.S., which prompted a drastic drought response throughout the region.[17] In response to the drought, the high information utility requested voluntary watering reductions from all customers during the summer months amounting to a 10% reduction in a household's monthly water consumption compared to the same month in 2013.[18] The voluntary reductions were implemented midway through summer 2014, before the merger, and throughout summer 2015, after the merger. The low information utility took no significant drought action during 2014, however the drought was a widely discussed issue on the local and national news media. It is reasonable to assume that most households in the Reno metropolitan area were aware of the ongoing drought. The concern with these voluntary watering reductions, henceforth "drought restrictions," is that they may disproportionately affect water consumption for high information households before the merger. Moreover, the low information households may respond differently to drought

[17]For example, in neighboring California, urban utilities were required to reduce demands by 25%.

[18]The high information utility had regular drought information inserts included with the monthly bills and a city-wide public relations campaign that used multiple media formats to reach residents including news media, billboards, and social media. The low information utility, by contrast, had very little outreach.

Table 1.5: Main Difference-in-differences Results: Summer Months

	(1) All Periods	(2) No Drought Periods	(3) No 2014	(4) No 2015	(5) No 2014-2015
β (Summer)	-0.0400***	-0.0644**	-0.0400**	-0.0570**	-0.0627**
	(0.0140)	(0.0274)	(0.0176)	(0.0221)	(0.0284)
Low Info Trend	0.0025***	0.0030***	0.0025**	0.0011***	0.0030***
	(0.0008)	(0.0011)	(0.0010)	(0.0004)	(0.0011)
Low Info Avg. Use	25.27	25.18	25.36	25.27	25.36
Within R^2	0.0074	0.0064	0.0069	0.0073	0.0068
Households	60,833	60,833	60,833	60,833	60,833
Observations	1,918,766	1,489,547	1,617,200	1,617,423	1,315,857

Note: The dependent variable is log of monthly water consumption divided by the number of days on the bill. The coefficients can be interpreted as the percentage change in consumption. Results based on a sample that is restricted to summer months (May-September). Controls for all regressions include household FEs, month-by-year FEs, Average Temperature, Precipitation in Inches, and a separate trend and month FEs for Low Information households. Robust standard errors clustered at the meter reading route level are reported in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

restrictions as new customers of the high information utility in 2015, being unaccustomed to those types of communications. The concern is that the drought restrictions may confound the treatment effect in either direction, depending on the direction of any differential response to this policy.

I investigate possible confounding due to the drought in several ways. First, I provide results that do not include any drought restriction periods that occurred during 2014 and 2015. Next, I sequentially drop all periods in 2014, 2015, and both years. These specifications allow for better balance in seasonality between pre-treatment and post-treatment consumption, since the drought restrictions only take place during the summer months and consumption is highly seasonal. Columns 2-5 of Tables 1.4 and 1.5 present regression results that exclude drought restriction months in the summers of 2014 and 2015.

Column 2 excludes only the months in 2014 and 2015 during which drought restrictions were in effect. This leads to a comparable treatment effect of a 4.1% decrease in average consumption. The difference is even more striking during the summer, when the treatment effect climbs to a reduction of 6.4%.

I also run three additional specifications in Columns 3-5 omitting all months in 2014, all months in 2015, and both years altogether respectively. The specifications produce significant results. Omitting 2014 attenuates the treatment effect, but it is still negative and significant. This model suggests that information leads to a 2.5% reduction in the difference in average consumption between low and high information households after the merger. Omitting 2015 leads to a 5.3% reduction in the difference in average consumption and omitting 2014 and 2015 leads to a 5.2% reduction in the difference in average consumption. The important takeaway is that all of these specifications yield a treatment effect that is statistically significant and similar in magnitude to the main results, which suggests that the estimated treatment effect is not being driven by these other policies. I also consider the robustness of the results to other conservation programs offered by the high information in Appendix A.

1.6.1 Event Study

The main regression model estimated in the previous section assumes that low information and high information households have parallel growth paths, or a linear difference trends. In practice there could be non-linear differences in trends, which would invalidate the identification assumptions. To rule out higher order differences in pre-treatment trends, I perform the following event study analysis that accounts for a linear difference in trends as well as separate seasonality between low information and high information households:

$$ln(Y_{it}) = \alpha_i + \sum_{k=2010}^{2016} \beta_k \mathbb{1}_{\text{Low Info,Year}\,k} + \lambda_t + \mu \mathbb{1}_{\text{Low Info}} * t + \sum_{s=1}^{12} \nu_s \mathbb{1}_{\text{Low Info,Month}\,s} + W_{it}\gamma + \varepsilon_{it} \tag{1.9}$$

Figure 1.9: Event Study Plot

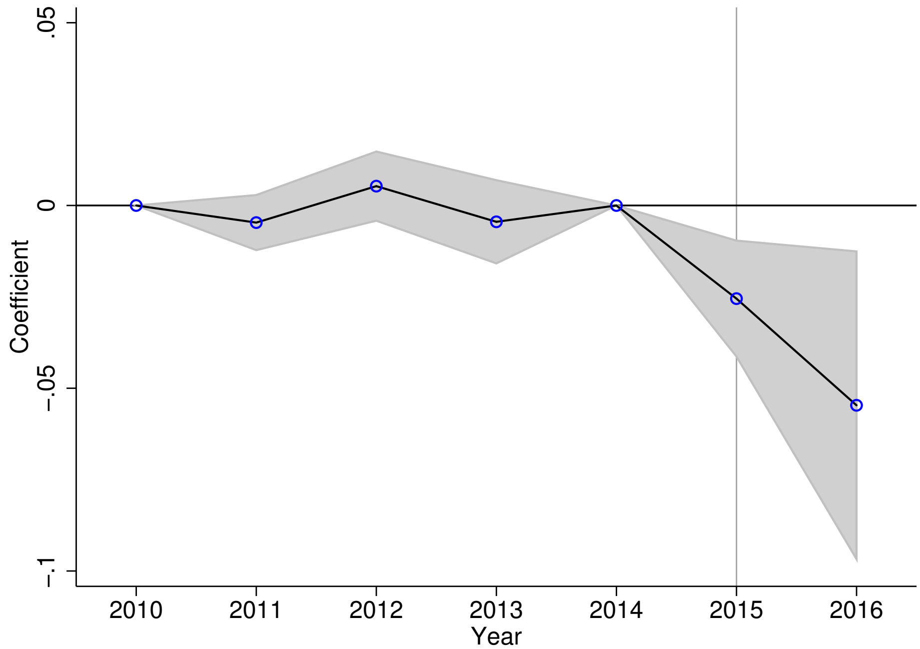

Note: This event study plot uses the last year before the merger (2014) as the reference year. It demonstrates that model 1.8 adequately controls for differences in pre-treatment trends and seasonality between low information and high information households and that there is a significant reduction in the difference in consumption after the merger. Robust standard errors are clustered at the meter reading route level. Figure A.3 in Appendix A shows that without a linear difference in trends, there are significant differences in pre-treatment trends relative to 2014.

where $\mathbb{1}_{\text{Low Info,Year}\,k}$ is a set of year fixed effects for the low information utility. I use 2014, the year before the merger, as the reference period and plot the β_k's which represent the difference between treatment and control relative to the reference period. A convincing event study would indicate that there is no difference between the pre-treatment coefficients (2010-2014) and a significant difference in the post-treatment coefficients, compared to the estimated difference in 2014.

Figure 1.9 estimates an event study that uses 2014 as the reference year.[19] After including a separate linear trend for the low information utility and separate month fixed effects, there are no remaining differences in pre-treatment coefficients, relative to 2014. By contrast, there are significant negative differences in 2015 and 2016 relative to

[19] I also estimate placebo tests in Table A.2 in Appendix A

2014.[20] These results suggest that allowing for a separate linear trend and seasonality adequately controls for differences in pre-treatment trends between low information and high information households.

1.6.2 Treatment Effect Heterogeneity

Finally, to shed light on who responds to price information, I look at heterogeneity in the estimated treatment effect. This analysis investigates whether certain subpopulations respond differently to price information on monthly bills. I estimate a conditional average treatment effect (CATE), which interacts the treatment indicator with an indicator for a given subpopulation (Abrevaya, Hsu and Lieli, 2015). For pre-treatment variables that are continuous, households are grouped into subpopulations using quintiles of the distribution. The CATE model is as follows:

$$ln(Y_{it}) = \sum_{q=1}^{5} [\beta_q \, x_{iq} \mathbb{1}_{\text{Low Info, Post}} + \lambda_{qt}] + \alpha_i + \mu \, \mathbb{1}_{\text{Low Info}} * t + \sum_{s=1}^{12} \nu_s \mathbb{1}_{\text{Low Info,Month } s} + W_{it}\gamma + \varepsilon_{it} \tag{1.10}$$

where x_{iq} is an indicator equal to one if household i is in quintile q and zero otherwise, $\mathbb{1}_{\text{Low Info, Post}}$ is an indicator for low information households after the merger, and λ_{qt} are a full set of quintile-by-month-by-year fixed effects. β_q is then the effect of providing price information on monthly utility bills to households in quintile q.

I focus on quintiles of pre-treatment summer water consumption, assessed property value as a proxy for wealth (Ferraro and Miranda, 2013), and yard size as a proxy for landscaped area (Renwick and Green, 2000). These variables are a way to look at variation in response to information over variables that are proxies of preferences for water consumption.

[20] The 2010 coefficient drops out of any specifications that include a linear trend and month fixed effects.

First, I use quintiles of average pre-treatment summer water consumption from 2010-2014 to estimate how the treatment effect varies across the distribution of water users. The results in Column 1 of Table 1.6 indicate that water users in the highest quintile of summer water consumption are driving the treatment effect. For these households, the treatment effect is a 1.7% decreases in consumption relative to households who used a similar quantity of water before the merger. These results suggest that households with more available margins on which to decrease consumption and who stand to gain the most in terms of saving money from water consumption due to the high marginal price they face have strongest response to price information on monthly bills, as we would expect.

Next, I estimate CATEs using quintiles of yard size in Column 2 of Table 1.6. The results are significant for only the lowest and highest quintiles of yard size. These results indicate that low information households with the smallest and largest yards reduced consumption by 3% relative to high information households with comparable sized yards. In light of the fact that yard size and property value are positively correlated, these results make some sense. Households with small yards tend to be associated with lower levels of income. We would expect these households to respond more to price information than high income households. By contrast, households with larger yards have more margins available to cut back on consumption.[21] Therefore, these results suggest that households with more margins of adjustment in consumption respond to the signal that they can save money on their bills by decreasing consumption.

Finally, the results for quintiles of appraised value in Column 3 of Table 1.6 are significant for the first, second, and fourth quintiles. Low information households in the lowest quintile of property values decrease consumption by 7.5% after receiving price informa-

[21] Making small adjustments in irrigation watering, for example, can lead to substantial decreases in consumption. This effect scales up with the size of the lawn.

Table 1.6: CATE Difference-in-differences Results

	(1) Pre-treatment Use	(2) Yard Size	(3) Appraised Value
β_1	-0.0160	-0.0280*	-0.0747***
	(0.0114)	(0.0144)	(0.0186)
β_2	-0.0245	-0.0070	-0.0402**
	(0.0181)	(0.0103)	(0.0192)
β_3	-0.0034	-0.0104	-0.0219
	(0.0143)	(0.0130)	(0.0169)
β_4	-0.0106	-0.0077	-0.0407**
	(0.0110)	(0.0165)	(0.0156)
β_5	-0.0174*	-0.0312***	-0.0201
	(0.0101)	(0.0105)	(0.0138)
Mean Use Quintile 1	5.52	12.23	14.23
Mean Use Quintile 2	12.07	17.61	19.93
Mean Use Quintile 3	17.71	19.41	19.59
Mean Use Quintile 4	24.55	22.92	23.15
Mean Use Quintile 5	44.32	32.14	34.09
Within R-squared	0.0071	0.0070	0.0073
Households	60,806	60,771	60,771
Observations	1,918,412	1,917,341	1,917,341
All Nonzero: F	2.4075193	2.2160527	5.4343096
All Nonzero: p	.04215667	.05892997	.00019313
All Equal: F	2.2269449	1.331074	5.3786146
All Equal: p	.0718246	.26407663	.00060085

Note: The dependent variable is log of monthly water consumption divided by the number of days on the bill. The coefficients can be interpreted as the percentage change in consumption. The sample is limited to summer water use (May-September). Controls for all regressions include household FEs, month-by-year FEs, Average Temperature, Precipitation in Inches, and a separate trend and month FEs for Low Information households. Robust standard errors clustered at the meter reading route level are reported in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

tion on monthly bills, relative to high information households with comparable property values. By contrast, low information households in the second and fourth quintiles reduce consumption by 4%. These results are reasonable, since we would expect lower income levels to be more motivated to watch their spending on utility bills. Receiving better price information helps cue these households to the fact that they can save money by conserving water.

Overall, the CATE analysis indicates that the treatment effects are being driven by high water users, households with large yards, and low income households. These results suggest that price information leads to a decrease in consumption for households who either have more margins of conservation available or low income households with a stronger need to keep utility bills at a minimum.

1.7 Conclusion

In this chapter, I present quasi-experimental evidence that providing price information on monthly bills for residential water consumers who face IBTs for water consumption leads to a more than 3% decrease in average water consumption, amounting to 500 gallons or more of water saved per household per month on average. These results aggregate to roughly 86 million gallons of water saved over the course of a year for the low information households. These results are robust to non-linear differences in pre-treatment trends and other concurrent conservation policies. I also show that low income households and households with more margins of adjustment in consumption are more likely to respond to price information treatments.

This research demonstrates that price salience is a possible mechanism behind empirical evidence showing that consumers do not respond to marginal price under IBTs. I also argue that economic theory predicts consumers will have a heterogeneous response to price information depending on the location of their pre-treatment consumption relative to the tier thresholds. I provide empirical evidence that consumer response to price information on monthly utility bills is consistent with these predictions. My model results contrast previous theories of consumer inattentiveness, which predict consumers systematically underestimate or overestimate price when prices are nontransparent.

This research dispenses cautionary advice for policymakers who want to use information to improve conservation under IBTs. Although I find evidence that price information can be used as a conservation instrument under IBTs, policymakers should be careful to ensure that this will be the case. Depending on the interaction between the ex-ante distribution of consumption and the rate schedule, price information could increase or decrease average consumption or have no effect.

Price information is a low cost policy tool compared to many of the other programs that have been pursued to promote energy and water conservation, such as high efficiency appliance rebates, "cash for grass" programs, and in-home displays. Price information treatments can be welfare-improving when the overall treatment effect leads to conservation. This contrasts other approaches that employ social incentives to promote conservation, which can impose moral costs on consumers (Allcott and Kessler, 2015). I take up this issue in more detail in the third chapter of this dissertation. Finally, consumers' heterogeneous response to price information can be used to target households who consume just above a tier threshold to achieve maximum effectiveness of prices as a conservation instrument.[22]. This will help guarantee that price information has a net effect of a decrease in consumption.

[22]This is similar to how targeting is done with social norms messages (Ferraro and Price, 2013)

Chapter 2

Automatic Billing, Paperless Billing, and Consumer Inattentiveness: Evidence from Water Demand

2.1 Introduction

Consumer inattentiveness leads to mis-optimization in a variety of settings.[1] In particular, consumers have a tendency to over-consume when facing prices and fees that are well known but less salient such as sales taxes (Chetty, Loney and Kroft, 2009), shipping and handling fees (Brown, Hossain and Morgan, 2010; Hossain and Morgan, 2006), and road tolls (Finkelstein, 2009). Recent empirical evidence suggests that enrollment in automatic billing services can lead to significant increases in consumption of toll roads (Finkelstein, 2009) and electricity (Sexton, 2015). These authors argue that automatic billing promotes consumer inattentiveness to prices. The rationale for these findings is that when consumers enroll in automatic billing, they no longer need to read their bills in

[1] See DellaVigna (2009) for a survey of the literature on decision errors. There are several studies that focus on consumer inattention to prices and fees. Brown, Hossain and Morgan (2010) and Hossain and Morgan (2006) find that consumers are less responsive to changes in shipping and handling fees than to changes in auction prices on eBay. Busse, Silva-Risso and Zettelmeyer (2006) find that consumers are more responsive to rebates for car purchases than dealer discount promotions. Chetty, Loney and Kroft (2009) find that displaying sales tax inclusive prices leads to significant reductions in demand. DellaVigna and Pollet (2009) find that stock prices are less responsive to Friday earnings announcements relative to other weekdays.

order to pay them, which can make consumers unaware of increased spending associated with increased consumption.

These results have wide-reaching implications as automatic bill payment (ABP) becomes more popular and widely used by consumers. As of 2010, it was estimated that two-thirds of U.S. consumers had at least one recurring bill set-up to pay automatically and amongst Internet-connected households 41% of all recurring bills were paid automatically (Fiserv, 2010). Moreover, consumption behavioral will only become more decoupled from spending as technology is developed that enable consumers to shop using mobile apps which allow purchases to be made with the push of a button, hiding the payment transaction entirely. Another trend in billing, called paperless billing (PL), requires that consumers forgo paper bills in favor of electronic bills that are emailed or accessible through an on-line account. Currently, 54% of U.S. consumers are estimated to receive at least one electronic bill, while 25% of U.S. consumers have gone completely paperless with their billing (Fiserv, 2016). With PL billing, there is not a clear driver of changes in consumption as there is with ABP, since consumers who use PL but not ABP are still required to view their monthly bill in order to make a payment. However, it is possible that PL promotes consumer inattentiveness as well. Perhaps consumers do not take as much care to review their monthly charges when they are in an electronic format. As such, there is no theoretical prediction for a consumption response to enrollment in PL.

ABP and PL both have direct environmental benefits in the form of lower carbon emissions and demand for lumber from decreased shipping of payments and bills as well as decreased printing of bills. Moreover, ABP affords a variety of direct benefits to consumers and retailers (Mastercard, 2006; Visa, 2006). Benefits to consumers include convenience as well as savings on postage costs and and foregone late payment fees. Automatic billing reduces costs for service providers due to increased certainty of on-time payments and lower transaction costs. Paperless billing affords some convenience to consumers by allowing them to have an electronic record of their billing history, but the major savings from paperless billing accrue to service providers in the form of reduced printing costs and postage fees. It is for this reason that many service providers require consumers to enroll in PL in order to enroll in ABP. Therefore previous research on ABP effects estimates the combined effect of both programs (Sexton, 2015). Moreover, to get an accurate estimate of the net environmental benefits of ABP and PL, consumer behavior must be taken into account. If consumers increase electricity and water consumption after enrolling in the programs, it undermines the environmental benefits of these policies.

This chapter investigates how automatic and paperless billing affect residential water demand. Water consumption is another setting where automatic billing has the potential to offset conservation efforts similar to electricity consumption. Using billing data from a water utility in the Reno metropolitan area and a fixed-effects approach, I estimate the effects of enrollment in automatic and paperless billing on water demand. This approach takes advantage of a unique feature of the billing system that allows consumers to independently enroll in PL and ABP to produce separate estimates of these two effects. I find that ABP enrollment leads to a more than 2% increase in average water consumption and PL enrollment leads to a more modest 1% increase in water consumption. Moreover, both programs have effects that do not manifest until after three years of enrollment. I also estimate the ABP and PL effect by quintiles of the distribution of property values to shed light on the role that income plays in explaining the change in consumption after enrollment. Sexton (2015) discusses concerns that the ABP effect might actually be an income effect rather than a consumer inattentiveness effect, although the author does not have access to detailed account records to further investigate this issue.

To the best of my knowledge, this is the first study to estimate the impact of ABP enrollment on water consumption. Moreover, this is the first study to investigate the relationship between income, ABP enrollment, and the subsequent impacts on consumption by estimating heterogeneous treatment effects by quantiles of appraised property values. I find that the ABP effect is distributed across all quantiles of property value indicating that the effect is not driven by consumers at higher income levels. By contrast, I find that the lowest quintile of property values drives the PL effect, suggesting that paper bills help low income families to keep better track of consumption.

I also consider how enrollment in ABP and PL can affect subsequent conservation response to voluntary watering restrictions during a drought that took place from 2014-2015. I find that ABP customers' consumption is 6% higher on average than average consumption for non-ABP accounts during the drought restriction periods. This suggests that enrollment in these programs has an additional unintended consequence in that consumers become inattentive to conservation appeals from the utility. A major implication of this result is that ABP has the potential to undercut environmental policies that use non-pecuniary incentives to promote conservation for both energy and water. For example, the well-known Opower studies that use social comparisons to promote conservation obtain a 2% decrease in electricity consumption (Allcott, 2011) and similar effects have been found from social comparisons for water consumption (Brent et al., 2016; Brent, Cook and Olsen, 2015; Ferraro, Miranda and Price, 2011). The ABP effect would more than offset these conservation policies.

This research contributes to a literature focused on consumer inattentiveness to water and energy costs. There is empirical evidence that consumers do not respond to marginal prices for electricity (Ito, 2014) or water (Wichman, 2014) consumption when utilities use complex nonlinear tariffs. Other research shows the importance of providing adequate price information to improve salience in this context (Lott, 2017; Kahn and Wolak, 2013; Gaudin, 2006; Pellerano et al., 2015; McRae and Meeks, 2016). Enrolling in ABP can further exacerbate consumer mis-optimization in this context if prices become less salient as consumers view utility bills less frequently. Boampong (2016) looks at ABP enrollment for electric utilities and finds that ABP customers have less elastic demands in response to price changes. The evidence presented in this chapter, taken together with other research in this area, suggests that ABP has the potential to undercut both pecuniary and non-pecuniary conservation policies.

2.2 Background and Data

Truckee Meadows Water Authority (TMWA) is the primary water utility in the Reno metropolitan area in Northern Nevada. TMWA has offered ABP since its inception in 2003 and PL since the fall of 2009. Currently, 15% of TMWA's residential single family customers are enrolled only in ABP, 6% are enrolled only in PL, and 6% are enrolled in both programs. ABP has gradually gained popularity over time; only 2% of the customers in this study signed up during the first year that ABP was offered and only 7% had signed up by 2010, when PL became available. Less than 2% of TMWA residential customers signed up for PL by the end of 2010, and of these 34% were already signed up for ABP. The unique feature of this billing system is that customers may enroll in either program independently. There is no requirement for consumers to switch to PL when initiating ABP. This requirement is put into place by many service providers to decrease costs from both ABP and PL. However, these statistics demonstrate that, when given the choice, many consumers prefer to continue receiving paper bills. This suggests that many consumers still have a preference for non-electronic personal record-keeping. To encourage customers to enroll in ABP, TMWA offers a 21 cent rebate per bill to "pass savings onto (their) customers."[2] By contrast, there is no rebate for signing up for PL. If self-selection into ABP is a concern for analysis, the enrollment incentive might help

[2]This rebate has been in place since the program began.

mitigate this self-selection.

The primary data used for this analysis are TMWA's single family residential billing records from 2003 to 2017, which include monthly water consumption, the billing dates corresponding to each record, rate information, meter reading routes, and the spatial location of each water service. The unit of analysis is the customer account. Only accounts that have at least one year of billing history are included in the final dataset, which helps reduce the number of renters in the sample who tend to move more frequently and may not pay water bills directly.[3] This final sample, which I refer to as the "full sample" throughout the remainder of this chapter, includes 149,513 accounts: 22,409 ABP only, 6,512 PL only, 7,491 ABP and PL, and 113,101 controls. Altogether, the final sample is an unbalanced panel of account-by-month records for a total of 9,630,007 observations.

A major issue for analysis arises due to a large number of accounts that enroll in automatic billing at approximately the same time as the customer account is initiated with TMWA. For these nearly 60% of ABP and PL accounts there is not a sufficient pre-treatment history, which is important for identification in a difference-in-differences framework. To address this issue, I estimate specifications using both the full sample as well as a more limited sample that excludes the accounts that have less than one year of pre-treatment history. I refer to this sample as the "preferred sample" throughout the remainder of this paper. This sample includes 8,576 ABP only accounts, 3,624 PL only accounts, and 2,491 ABP and PL accounts.

The data are spatially matched with geocoded data containing structural characteristics for each home from the Washoe county assessor and demographic information the

[3]In addition, accounts that have significant gaps in their billing history are excluded from the final dataset. Also excluded are outlier observations and residential bills that have a multi-family designation in the assessor records. Outlier observations are defined as having water consumption greater than 8 times the interquartile range of consumption for all single-family records within each billing month.

Table 2.1: Descriptive Statistics

	Control	ABP Only	PL Only	Both	N	F	p
Log(Consumption/Day)	-1.3115	-1.3528	-1.3746	-1.3852	9,621,702	38.42	0.0000
Consumption/Day (1,000 gal)	0.4090	0.4133	0.3809	0.3930	9,621,702	20.73	0.0000
Water Charges	38.0862	38.6137	36.0148	37.0865	9,621,702	19.06	0.0000
Marginal Price	2.3135	2.3003	2.3117	2.3127	9,621,702	2.76	0.0463
Average Price	4.2803	4.5091	4.3604	4.5011	9,621,702	44.93	0.0000

(a) Billing Characteristics

	Control	ABP Only	PL Only	Both	N	F	p
Account Duration	64.5968	73.5741	68.6062	67.8381	179,876	32.55	0.0000
Year Built	1984.3	1987.2	1987.3	1987.6	177,676	10.75	0.0000
Appraised Value ($1,000)	195.4463	226.5386	194.8001	218.9732	179,816	14.73	0.0000
Yard Size (Acres)	0.2223	0.2313	0.1888	0.2030	179,833	7.72	0.0001
Bedrooms	3.1936	3.2764	3.2059	3.2648	179,833	11.15	0.0000

(b) Property Characteristics

	Control	ABP Only	PL Only	Both	N	F	p
Perc. Male	49.9225	49.7520	49.8850	49.7768	179,560	4.26	0.0072
Median Age	37.8656	39.0270	37.6293	38.7596	179,560	8.06	0.0001
Perc. White	79.0661	80.9839	79.8392	81.1794	179,560	13.49	0.0000
Perc.Black	2.1800	1.9296	2.1088	1.9131	179,560	12.00	0.0000
Perc. Hispanic	18.8066	15.8727	17.2630	15.2675	179,560	15.05	0.0000
Avg. HH Size	2.6241	2.5881	2.6254	2.5740	179,560	8.58	0.0000
Perc. Owner Occupied	65.7875	68.4060	66.9126	67.8492	179,560	6.44	0.0005
Income/Capita	32,286.24	35,554.76	32,633.74	35,327.61	179,560	13.14	0.0000
Median Home Value	310,196.4	339,404.4	315,186.9	338,216.3	179,560	14.33	0.0000
Perc. College Degree	37.2680	41.1809	38.7078	42.0268	179,560	20.74	0.0000
Male Employ. Rate	92.4341	93.2444	92.8033	93.4325	179,559	9.11	0.0000

(c) Census Demographics

Note: The ACS data are taken from the 2006-2010 ACS 5 year estimates. ACS data has associated sampling error (reported as margins of error), which has not been accounted for in these balance tables. The reported significant differences therefore are likely overstated. Robust standard errors are clustered at the meter reading route level for billing and property characteristics, and at the census tract level for demographic characteristics. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

American Community Survey (ACS) corresponding to 2010 Census tracts. I used NOAA weather station daily data to create monthly weather averages of temperature and precipitation that correspond to the exact billing dates for each customer bill. Table 2.1 presents sample means for the full sample divided into the 1) control group, 2) ABP only accounts, 3) PL only accounts, and 4) accounts the enroll in both ABP and PL at some point. In terms of water use, ABP accounts use more water and PL account use less water than control accounts. ABP accounts are associated with higher property values and newer, larger properties, while paperless accounts are associated with lower property values and smaller properties—more in line with the control group. In terms of demographics, ABP accounts are associated with older, white, and more educated consumers, with smaller household sizes, higher rates of employment, and higher per capita incomes and median home values. PL accounts are similar to control accounts across these demographics. These sample statistics suggest that there is selection into ABP based on income and related characteristics, while there is less selection based on income for PL accounts. Although there are significant differences in demographics, property characteristics, and water consumption between these four groups based on an F-test, account fixed effects will control for observable and unobservable differences in time-invariant household characteristics.

2.3 Empirical Strategy

I identify the average effect of enrolling in ABP or PL programs on water consumption using a fixed-effects approach, which controls for selection on time-invariant unobservables. I first perform an event study analysis to rule out concerns about differences in pre-treatment trends between accounts who enroll in these programs and accounts who never enroll in either program. I regress the log of average daily water consumption for account i in billing period t on a set of account fixed effects, month-by-year fixed-effects, and indicators for ABP and PL account before and after enrollment in these programs.

The model is as follows:

$$ln(Y_{it}) = \sum_{g=1}^{m} \left(\beta_g^A\, x_{itg}^A + \beta_g^P\, x_{itg}^P\right) + \sum_{h=1}^{n} \left(\beta_h^A\, x_{ith}^A + \beta_h^P\, x_{ith}^P\right) +$$

$$\alpha_i + \lambda_t\ + W_{it}\gamma + f(\text{Account Duration}) + \varepsilon_{it} \tag{2.1}$$

where Y_{it} is monthly water consumption divided by the days on the water bill for account i in billing period t, α_i is an account-specific fixed effect which controls for observed and unobserved time-invariant heterogeneity, λ_t is a full set of month-by-year fixed effects which control for common shocks to consumption, W_{it} are weather controls that are matched to the exact dates of each bill, f(Account Duration) is a cubic function of the account duration at time t, and ε_{it} is an idiosyncratic error term.[4] I estimate clustered standard errors by meter reading route to account for possible correlation in the errors within neighborhoods.[5]

Interest centers on the β^j's for $j = A, P$, which estimate the average difference between ABP (PL) accounts and the control group before and after enrollment relative to the difference during reference period, which is the first bill after enrollment. x_{itg}^A and x_{itg}^P group ABP and PL observations into years before enrollment, and x_{ith}^A and x_{ith}^P group ABP and PL observations into years after enrollment.[6] I focus on a frame of five

[4]I include a cubic function of account duration to allow for consumption to vary non-linearly over the life cycle of an account. We would expect consumption to increase quickly during the early period of a new account, when for example landscape is being established. Consumption growth should eventually taper off and even decrease over time as landscape becomes established and households invest in efficient appliances.

[5]Previous studies have used meter reading route fixed effects to control for unobserved neighborhood-specific characteristics that may influence water consumption (Ferraro, Miranda and Price, 2011). Standard errors are also likely to be correlated within meter reading routes. Some possible sources of correlation could include spatial variation in water use due to regional weather patterns as well as economic shocks to consumers residing in the same housing developments. There are 95 meter reading routes, which ensures that there are a sufficiently high number of clusters.

[6]Let t_0 be the date of ABP enrollment. Then x_{itg}^A is an indicator variable equal to one for ABP account i if billing period $t \in [t_0 - 12 * g, t_0 - 12 * (g-1))$ and zero otherwise. x_{ith}^A an indicator variable equal to one for ABP account i if billing period $t \in [t_0 + 12 * g, t_0 + 12 * (g-1))$ and zero otherwise. x_{itg}^P

years before and after enrollment to allow the effect of enrollment in these programs to accumulate over a longer time horizon. A convincing event study would indicate that there is no significant difference between the pre-treatment coefficients relative to the reference period and a significant difference in the post-treatment coefficients. I plot the β^A's and β^P's from this single specification in separate event studies. For both of these event studies, I use the preferred sample that requires treatment account have at least one year of billing history prior to enrollment.

Next, I use the difference-in-differences analogue of the previous model to investigate various issues that impact the estimated effect of enrollment in ABP and PL. The model is as follows:

$$ln(Y_{it}) = \beta_i^A\, x_{it}^A + \beta_i^P\, x_{it}^P + \alpha_i + \lambda_t\, + W_{it}\gamma + f(\text{Account Duration}) + \varepsilon_{it} \tag{2.2}$$

where β_i^A is the account-specific treatment effect for enrollment in ABP and β_i^P is the effect for PL enrollment. This framework allows the behavioral response to ABP and PL enrollment to vary by household, where some households may be more prone to inattentiveness than others. This model allows for estimation of unique treatment effects based on account characteristics, which is useful for utility managers who have access to detailed information about their customers.

In order for $\mathbb{E}[\beta_i^j]$ to be a valid estimator of the average treatment effect, β^j, we need to make several assumptions. First, the account-specific treatment effect is uncorrelated with deviations from the average propensity to receive treatment, $\mathbb{E}[\beta_i^j|w_{it} - \bar{w}_i]$ (Wooldridge, 2010). This requires that enrollment decisions are not systematically related to expected treatment effects, which is reasonable in this setting since consumers do not enroll in ABP or PL with the intention of increasing consumption. It is more likely that consumers enroll, because they have a preference for convenience. Moreover,

and x_{ith}^P are similarly defined for PL enrollment.

the 21 cent per month incentive should help mitigate any self-selection bias. Identification of the average treatment effect also requires strict exogeneity of treatment and an overlap assumption (Wooldridge, 2010; Rubin, 1990; Rosenbaum and Rubin, 1983). Strict exogeneity could be violated in this setting if there are correlations between treatment status in any time period and deviations in unobservable time-varying household characteristics. Sexton (2015) argues that strict exogeneity is likely satisfied due to the low number of accounts that dis-enroll from ABP.[7] To sum, since households enroll in ABP and PL for reasons unrelated to treatment and since treatment status is persistent, β^j is the causal effect of enrollment in ABP or PL on water consumption.

Next, I estimate heterogeneous treatment effects to examine how the ABP and PL effects vary by subgroups of interest. I estimate conditional average treatment effects (CATE) following Abrevaya, Hsu and Lieli (2015) by interacting x^A_{it} and x^P_{it} with subgroup indicators. The CATE model is as follows:

$$ln(Y_{it}) = \sum_{q=1}^{Q} \left(\beta^A_q\, x^A_{itq} + \beta^P_q\, x^P_{itq} + \lambda_{tq}\right) + \alpha_i + W_{it}\gamma + f(\text{Account Duration}) + \varepsilon_{it} \quad (2.3)$$

where x^A_{itq} is an indicator for account i in quantile q enrolled in automatic billing during billing period t, x^P_{itq} is an indicator for account i in quantile q enrolled in paperless billing during billing period t, and λ_{tq} are a full set of month-by-year-by quantile fixed effects. Therefore, β^A_q is the effect of being in quantile q and enrolling in ABP and β^P_q is the effect of being in quantile q and enrolling in PL.

Of particular interest, is whether or not the ABP effect is being driven by income rather than consumer inattentiveness. In other words, if enrollment coincides with changes in income, and changes in income also lead to changes in preferences for water consumption, then the observed effect will be due to changes in preferences rather

[7] In this setting, on 2% of accounts dis-enroll in ABP or PL.

than inattention to monthly utility bills. Using property values as a proxy for wealth (Ferraro and Miranda, 2013), I investigate how the effect of enrolling in ABP or PL on water consumption varies by quantiles of appraised value. Test whether the estimated coefficients are significantly different across quintiles of the property value distribution (Crump et al., 2008). A failure to reject the null hypothesis will provide evidence against income and in favor of consumer inattentiveness driving the treatment effects.

I also consider how these effects relate to indoor and outdoor water consumption to shed light on the mechanisms of changes in consumption. I estimate CATE by quantiles of yard size as a proxy for landscaped area (Renwick and Green, 2000), which will indicate whether the overall effect is being driven by changes in outdoor irrigation. If consumers are using more water to achieve a more attractive landscape, and do not react to subsequent increases in their water bills, we would expect this effect to be magnified for larger yards. In another approach, I look at how the baseline effects from model 2 vary seasonally by restricting the data to the summer irrigation months (May-September). This will further indicate whether the ABP and PL effects operate through changes in outdoor consumption habits.

Lastly, I analyze how ABP and PL enrollment affect performance in a voluntary utility-wide drought restriction program by including an indicator for time periods during which this program was in effect as well as interaction terms with the ABP and PL terms. This specification will provide evidence of whether enrollment in ABP and PL leads to lower participation in community-wide conservation efforts. This is important from a policy perspective, since it can potentially impact utility managers' ability to communicate with consumers and ultimately could diminish the effectiveness of non-pecuniary conservation policies such as normative messaging campaigns (Allcott, 2011; Ferraro, Miranda and Price, 2011).

2.4 Results

Estimation results provide evidence that ABP and PL enrollment both lead to significant increases in water consumption, although the effects are more pronounced for ABP as we might expect. Since the dependent variable is the log of consumption, we can interpret the coefficients as the percentage change in consumption. Results from the event study analysis are shown in Figure 2.1. For both ABP (Panel a) and PL (Panel b), the event studies indicate that there are no significant differences in pre-treatment trends and there are significant increases in consumption after enrollment in either program. Since this analysis relies heavily on pre-treatment comparisons, I limit my sample to accounts that have at least two years of pre-treatment history, although the results are similar for one year of pre-treatment history (see Figure B.1 in Appendix B). Furthermore these results suggest that significant increases in consumption do not occur until several years after enrollment. ABP leads to a significant 2% increase in consumption after four years, and PL leads to a significant 1% increase in consumption after two years which increases to 2% after four years. These results are further corroborated by duration results presented in Appendix B. We would expect that if consumer inattentiveness is driving the increase in consumption that it would take time for the change to develop. Therefore these duration results are consistent with this mechanism of change as opposed to a change in preferences for water consumption that coincides with the enrollment decision.

Next, I examine the interplay between ABP and PL enrollment to gain insight into what is driving the increase in consumption. In Tables 2.2 and 2.3, I estimate a model that only includes an ABP effect (Column 1), only a PL effect (Column 2), both effects (Column 3), and an interaction term in addition to both effects (Column 4).[8] The first two models produce significant effects for ABP and PL respectively, however including

[8]See Appendix B for results using the full sample.

Figure 2.1: Event Study

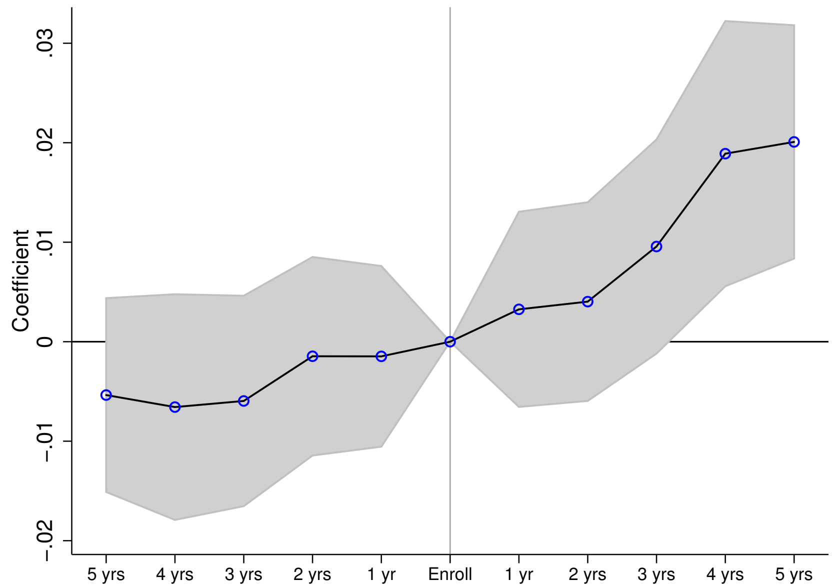

(a) Event Study ABP Enrollment

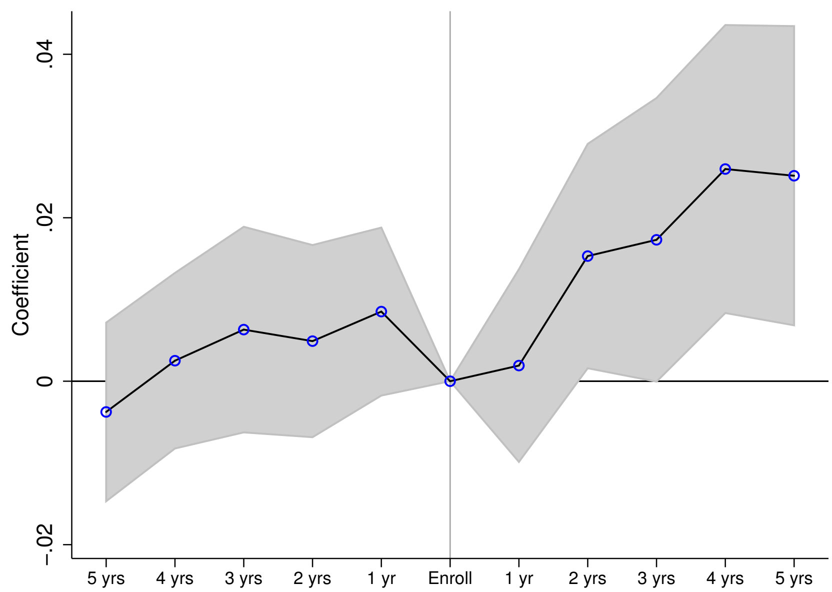

(b) Event Study PL Enrollment

Note: This event study plot uses the year of enrollment as the reference year. It demonstrates that the approach used in this chapter adequately controls for differences in pre-treatment trends and that there is a significant increase in consumption after enrollment. The sample is limited to accounts with at least 2 years of pre-treatment data. Robust standard errors are clustered at the meter reading route level.

Table 2.2: Baseline Results: 2003-2017

	(1)	(2)	(3)	(4)
ABP	0.0238***		0.0232***	0.0228***
	(0.0035)		(0.0036)	(0.0036)
PL		0.0106**	0.0050	0.0039
		(0.0040)	(0.0044)	(0.0045)
Enrolled ABP and PL				0.0033
				(0.0083)
Weather Controls	Yes	Yes	Yes	Yes
f(Act Duration)	Yes	Yes	Yes	Yes
HH FE's	Yes	Yes	Yes	Yes
Period FE's	Yes	Yes	Yes	Yes
Within R-squared	0.0112	0.0110	0.0112	0.0112
Households	158,168	171,282	154,328	154,328
Observations	8,628,508	9,396,989	8,513,972	8,513,972

Note: The dependent variable is log of monthly water consumption divided by the number of days on the bill. The coefficients can be interpreted as the percentage change in consumption. The results are based on the preferred sample, which requires at least 1 year of pre-treatment data. Controls for all regressions include household FEs, month-by-year FEs, Average Temperature, Precipitation in Inches, and a cubic function of account enrollment duration. Robust standard errors clustered at the meter reading route level are reported in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

both terms in the same model leads the PL effect to become insignificant. Moreover, the interaction term in the fourth model is not significant. These results suggest that ABP leads to a 2% increase in consumption, while PL leads to a 1% increase in consumption, which is smaller but comparable to the 4% increase in electricity consumption found by Sexton (2015) for joint ABP and PL enrollment. Overall, these results suggest that ABP drives the increase in consumption, since only the ABP effect is robust to all specifications. Moreover, the slightly smaller ABP effects in Columns 3 and 4 with similar magnitude standard errors suggest that a model that does not account for differences in PL enrollment might overstate the effect of ABP enrollment when enrollment decisions are independent.

Next, I examine some of the self-selection issues that are a concern for causal estimates of the ABP and PL effects. I interact the ABP and PL effects with an indicator variable

Table 2.3: Baseline Results: Summer Months 2003-2017

	(1)	(2)	(3)	(4)
ABP	0.0322***		0.0316***	0.0309***
	(0.0045)		(0.0045)	(0.0047)
PL		0.0132**	0.0032	0.0011
		(0.0057)	(0.0058)	(0.0062)
Enrolled ABP and PL				0.0064
				(0.0109)
Weather Controls	Yes	Yes	Yes	Yes
f(Act Duration)	Yes	Yes	Yes	Yes
HH FE's	Yes	Yes	Yes	Yes
Period FE's	Yes	Yes	Yes	Yes
Within R-squared	0.0150	0.0147	0.0150	0.0150
Households	154,168	167,066	150,449	150,449
Observations	3,579,451	3,900,024	3,533,037	3,533,037

Note: The dependent variable is log of monthly water consumption divided by the number of days on the bill. The coefficients can be interpreted as the percentage change in consumption. The results are based on the preferred sample, which requires at least 1 year of pre-treatment data. The sample is also limited to summer months (May-September). Controls for all regressions include household FEs, month-by-year FEs, Average Temperature, Precipitation in Inches, and a cubic function of account enrollment duration. Robust standard errors clustered at the meter reading route level are reported in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

for accounts that only ever enroll in ABP, an indicator for accounts that only ever enroll in PL, and an indicator for accounts that enroll in both programs at some point during their history. Table 2.4 presents these results in Columns 1 and 2. Once we separate the effects based on these subgroups, we can see that the ABP persists at 2-3% while the PL effect becomes small and insignificant. In Columns 3 and 4, I examine the effect of dis-enrollment in ABP and PL. These results suggest that dis-enrollment leads households to cut consumption, although by a smaller amount than the increase during enrollment. There are no significant effects for PL. Figure 2.2 estimates event studies for dis-enrollment these programs. These figures indicate that household consumption drops back down after dis-enrollment, which suggests that customers might realize that they are paying higher bills than they thought when they resume manual payment. However, these results are consistent with both the ABP effect operating through both

Table 2.4: Expanded Results: 2003-2017

	4 Groups		Disenroll	
	(1) Preferred Sample	(2) Full Sample	(3) Preferred Sample	(4) Full Sample
ABP, no PL	0.0227***	0.0306***		
	(0.0038)	(0.0035)		
PL, no ABP	0.0048	0.0051		
	(0.0052)	(0.0049)		
ABP, Enrolls Both	0.0256***	0.0280***		
	(0.0074)	(0.0049)		
PL, Enrolls Both	0.0037	0.0093*		
	(0.0071)	(0.0053)		
ABP			0.0221***	0.0282***
			(0.0036)	(0.0030)
PL			0.0055	0.0075*
			(0.0044)	(0.0038)
After ABP Dis-enrollment			-0.0184*	-0.0177**
			(0.0096)	(0.0079)
After PL Dis-enrollment			0.0020	0.0041
			(0.0176)	(0.0130)
Weather Controls	Yes	Yes	Yes	Yes
f(Act Duration)	Yes	Yes	Yes	Yes
HH FE's	Yes	Yes	Yes	Yes
Period FE's	Yes	Yes	Yes	Yes
Within R-squared	0.0112	0.0109	0.0112	0.0109
Households	154,328	178,886	154,328	178,886
Observations	8,513,972	9,620,712	8,513,972	9,620,712

Note: The dependent variable is log of monthly water consumption divided by the number of days on the bill. The coefficients can be interpreted as the percentage change in consumption. The preferred sample requires at least 1 year of pre-treatment data. Controls for all regressions include household FEs, month-by-year FEs, Average Temperature, Precipitation in Inches, and a cubic function of account enrollment duration. Robust standard errors clustered at the meter reading route level are reported in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

consumer inattentiveness and self-selection based on changes in income households into the program. This prompts a more in-depth examination of the interaction between treatment and income.

Table 2.5 presents CATE estimates based on quintiles of appraised property value and yard size. First, I examine the effect of ABP and PL enrollment on consumption by quintiles of appraised property value. In terms of income driving any self-selection into

Figure 2.2: Dis-enrollment Event Study

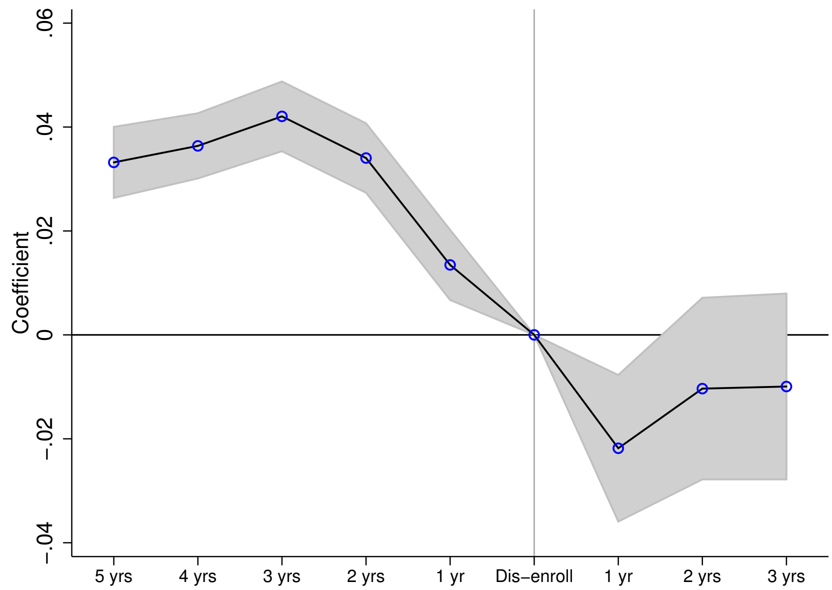

(a) Event Study ABP Dis-enrollment

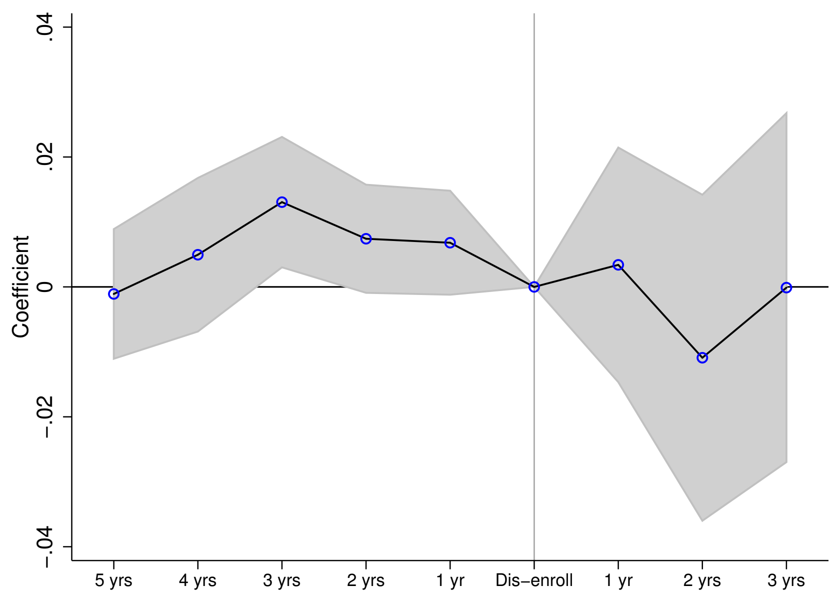

(b) Event Study PL Dis-enrollment

Note: This event study plot uses the year of dis-enrollment as the reference year. It demonstrates that household consumption decreases significantly after dis-enrollment. Estimates limited to three years post-dis-enrollment due to lack of observations outside of this time window. Robust standard errors are clustered at the meter reading route level.

ABP or PL enrollment, the share of ABP accounts within each quintile is comparable and ranges from 10% in the lowest quintile to 15% in the highest quintile.[9] For PL, enrollment ranges from 5% to 7%.[10] The results for ABP enrollment, shown in the first two columns of Table 2.5, suggest the effects are comparable across the distribution of property values. All ABP coefficients except for the fourth quintile of appraised value are statistically significant and all coefficients are positive. The effects by quintile range from 3% in the lowest quintile just over 1.5% in the highest quintile. Moreover, the estimated effects are not significantly different between quintiles. The effects are mostly insignificant for PL and often negative, except for the lowest quintile of property values. These results suggest that inattention drives the ABP effect rather than changes in income—the ABP effect is the largest for the lowest quintile of property values if anything. Moreover, the PL effect is being driven by low income households, which suggests that low income household might be particularly prone to inattention. Perhaps conventional paper bills are a more important resource for low income households to keep track of consumption and utility spending.

Next, I investigate how these effects very across quintiles of yard size in columns 3 and 4 of Table 2.5. There is a fairly even share of ABP and PL accounts within each quintile of yard size.[11] However, there are significant differences in the effect of ABP enrollment on consumption by property value quintiles. The results indicate that accounts with larger yards are driving the ABP effect. In fact the highest two quintiles are significantly different from the lowest two quintiles based on the actual quantity increases in water consumption (see Appendix B for results that use Consumption/Day as the dependent variable). The effects are generally insignificant for PL, except for the top quintile of yard

[9]The exact proportion from lowest to highest quintiles are: 10%, 12%, 13%, 15%, 15%.

[10]The exact proportion from lowest to highest quintiles are: 5%, 6%, 7%, 7%, 6%.

[11]The exact proportion of ABP accounts from lowest to highest quintiles of yard size are: 13%, 13%, 13%, 14%, 13%, and the proportion of PL accounts is 7%, 7%, 6%, 7%, 6%.

Table 2.5: Heterogeneity Results: 2003-2017

	Appraised Value		Yard Size	
	(1)	(2)	(3)	(4)
	Preferred Sample	Full Sample	Preferred Sample	Full Sample
Quintile 1*ABP	0.0296***	0.0324***	0.0167**	0.0252***
	(0.0097)	(0.0072)	(0.0064)	(0.0054)
Quintile 2*ABP	0.0215***	0.0283***	0.0123	0.0238***
	(0.0072)	(0.0058)	(0.0082)	(0.0069)
Quintile 3*ABP	0.0175***	0.0201***	0.0251***	0.0289***
	(0.0056)	(0.0052)	(0.0071)	(0.0060)
Quintile 4*ABP	0.0094	0.0184***	0.0220***	0.0223***
	(0.0059)	(0.0052)	(0.0068)	(0.0058)
Quintile 5*ABP	0.0160***	0.0248***	0.0329***	0.0394***
	(0.0057)	(0.0058)	(0.0067)	(0.0057)
Quintile 1*PL	0.0175	0.0198*	0.0097	0.0108
	(0.0135)	(0.0117)	(0.0103)	(0.0080)
Quintile 2*PL	-0.0007	0.0020	0.0118	0.0133
	(0.0118)	(0.0093)	(0.0103)	(0.0089)
Quintile 3*PL	-0.0045	-0.0011	0.0175	0.0146
	(0.0062)	(0.0078)	(0.0117)	(0.0095)
Quintile 4*PL	-0.0006	0.0001	0.0050	0.0058
	(0.0079)	(0.0068)	(0.0107)	(0.0093)
Quintile 5*PL	0.0094	0.0037	-0.0190*	-0.0125
	(0.0075)	(0.0068)	(0.0103)	(0.0079)
Weather Controls	Yes	Yes	Yes	Yes
f(Act Duration)	Yes	Yes	Yes	Yes
HH FE's	Yes	Yes	Yes	Yes
Period FE's	Yes	Yes	Yes	Yes
Within R-squared	0.0114	0.0111	0.0115	0.0112
Households	154,287	178,832	154,295	178,848
Observations	8,512,734	9,619,047	8,513,104	9,619,776

Note: The dependent variable is log of monthly water consumption divided by the number of days on the bill. The coefficients can be interpreted as the percentage change in consumption. The preferred sample requires at least 1 year of pre-treatment data. Controls for all regressions include household FEs, month-by-year FEs, Average Temperature, Precipitation in Inches, and a cubic function of account enrollment duration. Robust standard errors clustered at the meter reading route level are reported in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

size which has 2% decrease in consumption. Overall, these results suggest that changes in water use behavior might be related to outdoor water consumption, since irrigation is positively correlated with yard size. These results are corroborated by results shown in panel (b) of Table 2.2 that restrict the data to summer irrigation months from May to September. The pattern of results is similar to the overall results in panel (a), however the magnitude of the coefficients appears to be larger, which provides further evidence of the interplay between outdoor consumption habits and the increase in consumption for ABP and PL accounts. The implication of these findings is that ABP enrollment is likely to lead to more dramatic increases in consumption in areas where irrigation represents a large share of total water consumption.

Finally, I examine how enrollment in ABP and PL affects subsequent response to voluntary drought restrictions that were implemented by TMWA during the summers of 2014 and 2015. From July-September 2014 and May-September 2015, TMWA asked all customers to reduce consumption by 10% relative to their consumption in 2013 (the last year before the drought). I include an indicator for these time periods to allow consumption to vary during the conservation campaign. I then interact this indicator with my ABP and PL terms to estimate the interaction effect. These results are shown in Table 2.6. Overall, there was dramatic response to the voluntary restrictions, with an estimated 3-4% reduction in average consumption.[12] The ABP and PL effects are significant and comparable to previous results. The main effects of interest are the interaction terms between ABP and PL and the drought campaign. The ABP term implies that ABP accounts consumed an additional 6% more water relative to control accounts during the drought in addition to the more than 2% difference in consumption during non-drought periods. The results are not significant for PL. Overall, these results suggest that

[12] Utility managers reported a more than a 15% reduction in system-wide water consumption during 2015. Some of this reduction is absorbed in the time fixed-effects in this model.

Table 2.6: Drought Restriction Policy Results: 2003-2017

	(1) Preferred Sample	(2) Full Sample	(3) Preferred Sample	(4) Full Sample
Drought Restrictions	-0.0335	-0.0396*	-0.0335	-0.0401*
	(0.0232)	(0.0233)	(0.0232)	(0.0233)
ABP	0.0170***	0.0223***		
	(0.0035)	(0.0027)		
PL	0.0053	0.0099***		
	(0.0043)	(0.0037)		
ABP, Drought	0.0596***	0.0823***		
	(0.0111)	(0.0077)		
PL, Drought	-0.0089	-0.0338***		
	(0.0106)	(0.0077)		
ABP only			0.0166***	0.0225***
			(0.0037)	(0.0032)
PL only			0.0067	0.0097**
			(0.0051)	(0.0047)
ABP, Both			0.0204***	0.0223***
			(0.0074)	(0.0048)
PL, Both			0.0015	0.0095*
			(0.0074)	(0.0054)
ABP only, Drought			0.0596***	0.0870***
			(0.0124)	(0.0084)
PL only, Drought			-0.0104	-0.0256***
			(0.0130)	(0.0089)
ABP, Both, Drought			0.0559***	0.0601***
			(0.0211)	(0.0121)
PL, Both, Drought			-0.0041	-0.0203
			(0.0229)	(0.0131)
Weather Controls	Yes	Yes	Yes	Yes
f(Act Duration)	Yes	Yes	Yes	Yes
HH FE's	Yes	Yes	Yes	Yes
Period FE's	Yes	Yes	Yes	Yes
Within R-squared	0.0113	0.0112	0.0113	0.0112
Households	154,328	178,886	154,328	178,886
Observations	8,513,972	9,620,712	8,513,972	9,620,712

Note: The dependent variable is log of monthly water consumption divided by the number of days on the bill. The coefficients can be interpreted as the percentage change in consumption. The preferred sample requires at least 1 year of pre-treatment data. Controls for all regressions include household FEs, month-by-year FEs, Average Temperature, Precipitation in Inches, and a cubic function of account enrollment duration. Robust standard errors clustered at the meter reading route level are reported in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

inattentiveness translates into lower participation in utility-wide conservation campaigns (or that households who enroll in ABP are less apt to respond to these programs). The major implication of these results is that ABP enrollment may not only undermine utility efforts to promote conservation in general, but it can seriously hamper utility measures taken to cut demands when facing potential water shortages.

2.5 Conclusion

This chapter investigates the impact of ABP and PL enrollment on water consumption, finding robust and significant results indicating that ABP enrollments leads to a 2-3% increase in average water consumption. There is also evidence that PL enrollment can increase consumption by as much as 1%, but this result is less consistent across different specifications. Furthermore, I demonstrate that this increase in consumption after ABP enrollment is not being driven by an income effect by estimating heterogeneous treatment effects by quintiles of appraised property value, finding that average treatment effects are not different across quintiles. These results provide further evidence that ABP enrollment contributes to consumer inattentiveness to utility costs. By contrast, my results indicate that that low income households drive the PL effect. Therefore maintaining conventional paper bills might low income households keep track of consumption. The implication of this finding is that utility policies that require customers to "go paperless" in order to gain the benefits of enrolling in ABP might disproportionately negatively impact low income households.

I also find that ABP and PL customers are less responsive to voluntary drought restrictions. My results suggest that ABP accounts use an additional 6% more water during drought restriction periods compared to non-ABP customers. Overall, my results indicate that consumer inattentiveness that is promoted by ABP and PL enrollment has the potential to undermine both pecuniary and non-pecuniary conservation policies.

Despite the potential negative environmental impacts associated with ABP enrollment, there ways in which utilities can mitigate these impacts. First, utilities could text reminders each month to ABP customers with some basic information about their current bill amount as compared to historical averages. This could help cue consumers to pay attention to recent increases in their consumption. Another approach, could be to send customers on ABP quarterly or biannual letters notifying them of changes in consumption over that period. Utilities could also use alternative communication approaches to ensure that ABP customers are aware of conservation programs and drought restrictions, such as traditional mail communications.

Chapter 3

Are Normative Appeals Moral Taxes? Evidence from a Field Experiment on Water Conservation

3.1 Introduction

A number of high-quality, randomized experiments have established that information treatments that compare an individual household's electricity or water consumption with that of a peer group cause significant reductions in consumption.[1] In general, most of these social comparison studies have also found that the average treatment effects are driven largely by effects among customers with higher pre-treatment consumption (Allcott, 2011; Dolan and Metcalfe, 2015; Ito, Ida and Tanaka, 2015; Ferraro and Price, 2013; Ferraro and Miranda, 2013; Brent, Cook and Olsen, 2015). Why and how do these comparisons work? One explanation is that these comparisons invoke social norms: a consumer sees that his water consumption is higher than similar neighbors and feels guilty about his "overconsumption." In response to such a message, the consumer may impose a "moral tax" on consumption. Conversely, consuming less than one's peers can lead

[1] See among others Allcott (2011); Allcott and Rogers (2014); Ayres, Raseman and Shih (2013); Costa and Kahn (2013) for energy and Ferraro, Miranda and Price (2011); Ferraro and Price (2013); Brent, Cook and Olsen (2015) for water.

consumers to impart a "moral subsidy" (Levitt and List, 2007; Ferraro and Price, 2013; Allcott and Kessler, 2015). Framing the norm as a tax is appealing to economists, because it can be incorporated into welfare analyses of various price and non-price approaches for achieving conservation in the face of scarcity (as in the case we discuss) or reductions in associated consumption externalities. It also has the logical extension that the behavioral response should vary according to the level of such a moral tax, which we refer to as the strength of the normative message. When the moral tax is higher (one is consuming *much* more than one's peers), the effect of receiving information highlighting this should be larger than for those receiving a smaller, or negative, moral tax.

A second behavioral explanation draws from the more conventional household production approach (Becker, 1965), and recognizes that social comparisons also provide households with financially-useful information. Households informed that their water use is 50% higher than similar neighbors might also receive a signal that they can achieve the same utility from using water or electricity at lower costs by installing efficient appliances or making behavioral changes (Ferraro and Price, 2013). This updating of beliefs and re-optimization more generally fits the emerging rubric of correcting "internalities", or failures of consumers to internalize all of the private costs of their actions (Allcott, Mullainathan and Taubinsky, 2014; Allcott and Sunstein, 2015; Allcott and Taubinsky, 2015). Prominent examples of internalities in the water and energy sectors include: imperfect information about the costs of water/energy consumption (Allcott and Taubinsky, 2015); dynamic inconsistencies in decision-making (Allcott, Mullainathan and Taubinsky, 2014); lack of salience of infrequent or automatic billing (Sexton, 2015; Wichman, 2016); and confusion about nonlinear price structures (Ito, 2014; Wichman, 2014). Which of these two behavioral explanations best fits the pattern of results found in social comparisons?

The answer to this question is important from a welfare perspective. Information treatments to re-optimize and correct internalities are always Pareto-improving in markets without externalities, since some households can achieve the same utility at lower cost, and all other households are no worse off.[2] In markets with negative externalities, information treatments operating through re-optimization that cause an aggregate decrease in consumption, such as most water and energy social comparisons, will have even larger positive welfare gains due to reducing external costs. In contrast, if a social comparison works purely by imposing a moral tax, there may be a significant fraction of households for whom the level of their moral tax exceeds the correct Pigouvian tax. Social comparisons might thus induce over-abatement and reduce welfare overall. Allcott and Kessler (2015) demonstrate theoretically that the social welfare effects of a nudge can vary substantially depending on the prevailing behavioral mechanism.

Testing these behavioral mechanisms has been empirically challenging because of a design feature of existing social comparisons: customers with above-average pre-treatment consumption are much more likely to receive a message that they are consuming more than their comparison group. In fact, water use and the type of message received are perfectly correlated in studies where the comparison group is simply a utility-wide average consumption. These customers receive a higher moral tax, but they also receive a stronger signal to re-optimize consumption (i.e. their water bill is much higher than their neighbors' bills). Why? In our setting in Reno, Nevada, as in many cities in arid climates, summer (May-October) household water use is driven predominantly by outdoor water use on lawns and gardens. Consumption in the summer is generally four times higher than in the winter. Thus, the most impactful water conservation decisions are related to landscape changes (i.e. xeriscaping) and improvements in irrigation efficiency. Optimizing the amount and timing of irrigation water can achieve the same green landscape with less water. If a financial "re-optimization" channel predominates, a signal

[2]This argument neglects the typically small cost of generating the information, which could in theory be passed along in higher water or electricity rates.

that household water use is higher than similar neighbors' is likely a signal that outdoor water use is driving this difference. A re-optimization of outdoor water use, for example by installing an irrigation timer, may lower one's bill substantially.[3] In other words, with a traditional social comparison, high water users are faced with a higher moral tax but are also very likely the ones who might save the most money at the lowest utility-cost when re-optimizing.

Our primary contribution is the introduction of a novel social comparison treatment that allows us to better isolate a moral tax behavioral channel by decoupling pre-treatment water usage from the type of message sent to households, and thus the level of moral tax imposed. We do this by framing the comparison not in levels of consumption (i.e. thousands of gallons of water used), as in traditional social comparisons, but in percentage reductions from a prior baseline compared to the percentage reduction of similar neighbors. More concretely, the water utility in Reno (Truckee Meadows Water Authority, or TMWA) asked all customers to reduce their water consumption in summer 2015 by 10% compared to summer 2013 in response to a drought (discussed more below). Our novel information treatment informs customers by what percentage they have reduced water consumption compared to 2013 as compared to the reduction achieved by similar neighbors. We argue that this treatment still imposes a moral tax by signaling whether one is "doing their part" to manage the drought, but provides much less financially-useful information.[4]

[3]These are devices that, if used properly, can be set to apply the correct amount of water at the correct time and thus keep a lawn healthy without over-watering and unnecessarily increasing one's water bill.

[4]To see this, imagine you are told that you are using 6,000 gallons per month and that your neighbors with a similar lot size used 4,000 gallons. This might indicate that you could reduce water consumption by 2,000 gallons, and thus your bill, while being able to maintain a yard that looks similar to your neighbors. If, on the other hand, you are told that similar neighbors reduced their consumption by 8% compared to 2013, while you reduced by only 4%, you would not learn much to update your beliefs about whether you are optimizing how you use water: you don't know the baseline consumption level of the comparison household.

We analyze this new social comparison in a utility-scale field experiment among TMWA customers that sent one of five different mailers to single-family customers in TWMA's service area. Approximately 4,300 households were included in each of the five treatment groups, with 21,552 in the control group. Two of these include normative appeals: one is a traditional social comparison in terms of total gallons used by the household relative to a peer group, while the other uses the social comparison in percentage terms as described above. The third treatment provides rate information and frames conservation in terms of households' expected monetary savings. All of the mailers in the study describe TMWA's goal for each household to use 10% less water for each month of the summer of 2015 relative to the summer of 2013 to cope with a temporary drought. In this chapter, we focus on three of the five mailers.

Each of the three treatments generate statistically significant average treatment effects (ATEs) of roughly 1.5%. Although the individual ATEs are not statistically different from each other, we find several patterns in our results which are consistent with social comparisons operating *at least partly* by raising the moral cost of consumption. First, we find that the strength of the normative appeal, defined as the difference between a customer's consumption (or percentage reduction) and that of similar neighbors, is a strong driver of differences in responses to social comparisons. Importantly, this result holds when the strength of the normative appeal is decoupled from pre-treatment consumption. The conservation rate social comparison shows the same pattern of response for low and high water users based on pre-treatment consumption, although the effect is magnified for high users.[5] Second, linking our experimental results with survey data

[5]Randomization was balanced by design on pre-treatment water consumption, but we could not know a household's conservation rate in advance. The conservation rate comparison treatment was not therefore guaranteed to be balanced on percentage reduction and conditional average treatment effects should be interpreted cautiously. As we demonstrate below, however, the treatment achieved remarkable balance in conservation rates among baseline water use. In other words, some low baseline water users reduced their consumption by a higher amount than their peers, and some high water users reduced by a smaller fraction than their peers.

collected after the conclusion of the field experiment, we find that pro-social households were more responsive to the social comparison treatments, similar to Costa and Kahn (2013) and Bolsen, Ferraro and Miranda (2014).

Third, by randomizing whether households received one or two treatment letters and examining treatment effects after mailers stop arriving, we find that framing conservation in terms of monetary savings leads to more persistent treatment effects than our normative social comparisons. The two normative treatments largely replicate the "action and backsliding" pattern found by Allcott and Rogers (2014) (though in less temporal granularity), where the initial effect of a mailer wanes over time before increasing upon the receipt of a new mailer. The financially-oriented treatment causes a persistent effect and a second mailer has no additional effect. This is consistent with the findings by Ito, Ida and Tanaka (Forthcoming) that pecuniary incentives lead to more persistent effects than moral suasion. It is also consistent with consumers being more likely to re-optimize to address internalities when cued with monetary savings information, especially if this re-optimization involves capital investments (e.g. low-flow toilets or xeriscaping) or more permanent behavioral changes (e.g. irrigation timers). Consumers cued with normative appeals may draw on actions that are intent-oriented (Attari, 2014), but which are more likely to be transient (e.g. shorter showers). Our survey results also support this: consumers who reported previous conservation actions are less responsiveness to the monetary treatment, suggesting that those households had already re-optimized and exhausted any low-cost conservation opportunities.

Previous studies have attempted to investigate whether different types of nudges trigger different mechanisms, but none has conclusively distinguished between moral and financial motivations to normative appeals. Ferraro and Price (2013) compare social comparisons and generic pro-social appeals, both based on a moral motivation, to show that social comparisons generate greater conservation responses; however, they do not consider non-moral based mechanisms, such as addressing consumption internalities. Ferraro and Miranda (2013) use the traditional quantity-based social comparison and attempt to distinguish between financial versus moral motivations by comparing treatments responses by households just above and below thresholds in increasing block rate pricing structures, but fail to find that financial gain influences conservation response.[6] Pellerano et al. (2015) find weak evidence that adding information about financial savings to social comparisons crowds out moral motivations, as postulated by Gneezy, Meier and Rey-Biel (2011). Allcott and Kessler (2015) find that 35% actually have negative willingness to pay for continuing to receive social comparison, indicating that moral costs are a likely mechanism. However, their study takes place after the households had already received one year's worth of social comparisons, so the novel information component that helps correct internalities may have already been exhausted.

3.2 Background & Experimental Design

In 2015, in response to low snowpack and expected drought conditions during the summer irrigation season, TMWA launched a major media campaign through print, radio, TV, social media, and billboard messages requesting each TMWA customer to use 10% less water from May through September 2015, relative to their water use during the same months in 2013. The comparison year was 2013 because TMWA had asked for a 10% reduction during the latter part of the prior summer (July through September of 2014). Such a conservation request was uncommon in the region prior to 2014, however; the last time TMWA had requested customers to reduce water consumption to address drought was in 1992. To complement the media campaign, we worked with TMWA to design a randomized control trial to test the effectiveness of five different personalized

[6] Differences in savings between tiers tend to be small and not all consumers are aware of marginal changes in the price (Ito, 2014; Wichman, 2014).

letters mailed to residential single family households.

The TMWA request for a 10% system-wide reduction was met and surpassed. The conservation request ended officially at the end of September. In November-December the utility surveyed a sample of its customers (with email addresses) with the primary goal of assessing the effectiveness of their media campaign. The survey included questions about what customers did to conserve water, whether their actions were considered to be impositions, what prevented them from doing more to conserve water, and what more they would have been willing to do if called upon to do so. The survey responses included 2,544 households that were in our experimental groups, approximately 10% of the sample size in our experiment. In their survey sample, TMWA included all households from our experimental sample that had provided TMWA with an e-mail address. We were able to use these data to provide corroborating evidence regarding customers' motivations and mechanisms of response to the treatment letters they received. We discuss relevant TMWA survey outcomes in this chapter in the context of interpreting our main experimental results.[7]

3.2.1 Description of Treatments

This article focuses primarily on three of the five treatments (denoted T1 through T5): a rate information treatment focused on financial savings (T3) and two social comparison treatments (T4 and T5). However, since the treatments of interest contain components of the first two treatments we present a brief description of all five treatments (Table 3.1; Appendix C C includes example components of the five mailers). Every letter began: "Because of the extended drought in Northern Nevada, we are asking all of our customers to reduce water use by at least 10% this summer compared to summer 2013 - the last summer before TMWA started asking for summer water use reductions." All letters also

[7] A summary of selected survey results is available through the TMWA staff report (Christman, 2016).

included the statement: "Since TMWA customers use on average about four times more water in summer than in the winter, we expect that for most customers the easiest way to achieve this reduction is to adjust outdoor watering."

Treatment 1 provided households with six tips that the TMWA media campaign publicized for how to reduce outdoor water consumption, similar to Ferraro, Miranda and Price (2011). This letter was not customized to report on individual household water use. The six tips were also printed on the reverse side of the other four mailers (T2-T5).

Treatment 2 augmented the generic tips with personalized information about the customer's water use, with a title introducing the letters that read: "Below is your customized water use report." The T2 letter included a figure that displayed the customer's water use in thousands of gallons (kgal) for May through September of 2013 and also their water use in 2015 for each month from May up to the last month billed before the letter was sent out (Figure C.1 in Appendix C C shows the mailer). This figure and accompanying descriptive text was included with Treatments 2 through 5. Therefore all our treatments of interest include water conservation tips and personalized historical water use information.

Treatment 3 (T3) contained the same components as T2, with the additional message "Saving water saves you money", a figure displaying (a) the rate structure with tiers and price for each tier, (b) the customer's water use in kgal within TMWA's increasing-block rate structure for the last month billed in 2015, and (c) the upcoming month's target of 10% less water than the same month in 2013 within the rate structure. The letter also provided the monetary savings that the customer could expect from meeting this goal (see Figure C.2 in Appendix C C). We refer to this treatment as the rate treatment and financial treatment throughout the remainder of this chapter.

Treatment 4 (T4) provided the same information as T2 with the additional message

Table 3.1: Information Included in the Five Treatments

Information Components Included in Mailers*	Treatment (Mailer)				
	T1	T2	T3	T4	T5
The message "Helping our region deal with drought: What you can do" with a sheet showing 6 low cost tips to reduce outdoor water use	X	X	X	X	X
The message "What is your 10% Goal?" with a water use bar graph showing in 1000's of gallons the home's: • 2013 water use for May through September, • Target water use for each month in 2015 as 10% less, • 2015 actual monthly water use, up to last month billed.		X	X	X	X
The message "Saving water saves you money" with rate structure ** graphic showing home's water use in 1000's of gallons by tier/price: • for last month billed in 2015, • for same month in 2013, • for target goal of 10% less water used relative to 2013			X		
The message "How does your water use compare?" with a comparison of customer's 2015 last month billed water use in 1000's of gallons with similar neighborhood homes.				X	
The message "Are you doing your part?" with a comparison of customer's 2015 last month billed water use in terms of % change from 2013 with similar neighborhood homes.					X

*See Figures C.1-C.4 in Appendix C C for an example of each information component of each treatment. T2 through T5 Mailers included the title: Below is your customized water use report.
**TMWA had recently merged with two other small regional utilities, Washoe County and South Truckee Meadows Groundwater Irrigation District. In 2015 all customers were subject to the rate structures that they had before the merger. Customers receiving this treatment were shown the tariff structure relevant to them.

"How does your water use compare?" and a figure comparing the customer's total water use in thousands of gallons for the last billed month to the median water use of a peer group consisting of single-family residences in their neighborhood with similar yard size and number of bedrooms. This treatment essentially reproduces the standard social comparison used in the OPower studies on energy and the Cobb-County (Ferraro, Miranda and Price, 2011) and Watersmart (Brent, Cook and Olsen, 2015) experiments in water (see Figure C.3 in Appendix C C).

Treatment 5 (T5) included the same information as T2 with the additional message and figure providing a similar comparison between households as in T4, but instead expressed in terms of relative percent performance towards achieving the 10% goal compared to 2013 water use (Figure C.4 in Appendix C C). This treatment decouples the strength of the normative appeal from the level of water use, since low water users may not have conserved water and high users may have conserved a large amount in percentage terms, enabling us to isolate the impact of the strength of the message after conditioning on water consumption. The comparison group was identified in the same way as Treatment 4.

We also include injunctive norms for T4 and T5 in the form of a message. The message "Keep up the good work" was included for residences that had met their 10% goal in the last month billed. Households that did not meet their 10% goal in the previous month received the message "As a reminder TMWA is asking all customers to do their best to save at least 10% this summer. Please do your part to help with drought." We did not use emoticons or "smiley faces" as in Schultz et al. (2007).

3.2.2 Randomization

Our sample frame included 42,703 eligible[8] single family homes. We then randomly assigned each of these households to either the control group or one of the five treatment groups, with randomization blocks defined by billing cycles, rate schedule and frequency of recorded meter data (i.e. monthly, daily, or hourly, though all customers only receive monthly usage totals). Appendix C provides more detail on our randomization procedure and the process of generating the mailers. In total, 21,151 treatment households were assigned to receive mailers (Table 1). In addition, we randomized whether households received one or two mailers. A total of 7,086 households were assigned to receive a single

[8]Specifically, we included homes that (i) had metered water service; (ii) used enough water during at least one month of the 2013 irrigation season to exceed the tier 1 limit (6,000 gals), indicating some outdoor water use; (iii) had lived at their current residence since April 2013, and therefore had summer 2013 bills for comparison; (iv) had a billing address that corresponded with the residential service address to eliminate rental occupants and other users who may not pay for water or have limited control over water use at the residence; (v) had a 2-inch service main or smaller, excluding unusually large water users; (vi) live within one of the targeted bill cycle regions (some regions were excluded because they had a low number of single-family households, see Appendix C); and (vii) had nonzero water use during each month of the 2013 irrigation season (May-September) and pre-treatment months during the 2015 irrigation season (May-July) to exclude homes that were unoccupied for an extended period of time.

Table 3.2: Total Treated Households by Month/Treatment

Treatment Type	July Only	August Only	Both Months	Total
T1: Tips Sheet Only	1,420	1,410	1,402	4,232
T2: Tips + Water Use History	1,414	1,411	1,412	4,237
T3: Tips + History + Rate Information	1,413	1,411	1,410	4,234
T4: Tips + History + Social Norms (Gallons)	1,419	1,416	1,396	4,231
T5: Tips + History + Social Norms (Percent)	1,420	1,403	1,394	4,217
Total	7,086	7,051	7,014	21,151

mailer in July (using June consumption as the last month billed), 7,051 received a single mailer in August, and 7,014 received mailers in both July and August.

Table 3.3 shows that, in aggregate, the randomization achieved very strong balance on observables. Additionally, Tables C.1-C.3 in Appendix C show that the experimental sample is balanced on pre-treatment consumption for each treatment, across treatments, and within deciles of pre-treatment consumption. Figure C.5 graphically displays the densities of pre-treatment consumption for the pooled treatment, each of the five individual treatments, and the control group. In addition to achieving balance on average pre-treatment consumption, Figure C.5 shows the treatments are balanced across the full distribution of pre-treatment consumption. The graphical evidence is formalized by nonparametric Kolmogorov-Smirnov tests (Table C.4 in Appendix C) that fail to reject the null of equality of distributions for pre-treatment consumption across the control and the pooled treatment as well as each treatment individually. Our sample is well balanced by design, which allows us to make valid inferences for the conditional average treatment effects within subgroups, particularly subgroups that are functions of the normative appeal, which depends on pre-treatment consumption.

Table 3.3: All Treatments Balance on Observables

	Control Mean	Treatment Mean	Difference	(p-value)
2013 Water	23.56	23.55	0.02	0.89
2015 Water	16.90	16.99	-0.09	0.37
Summer Water	21.63	21.70	-0.07	0.52
Winter Water	8.15	8.15	-0.01	0.88
Year Built	1,987.61	1,987.67	-0.05	0.77
Appraised Value	214.85	214.87	-0.02	0.99
Bedrooms	3.37	3.37	0.01	0.43
Lot Acre	0.27	0.27	-0.00	0.91
Yard Acre	0.22	0.22	-0.00	0.95
Build Sq. Ft.	1,985.10	1,991.07	-5.98	0.41
Bathrooms	2.19	2.20	-0.00	0.56

Note: 21,552 Control Observations and 21,151 treated observations, p-value is based on two-sided t-test

3.3 Methodology

The primary variable of interest is monthly household consumption, obtained from TMWA billing records, expressed in average gallons per day (GPD). We calculate GPD by dividing total billing cycle usage by the number of days in that billing period to avoid problems with billing periods of different lengths. The specification in our regression analysis uses "normalized GPD" as the main dependent variable; every customer's GPD is divided by the average control group consumption across the experimental period (July-September 2015) following Allcott (2011). This allows the regression coefficients to be interpreted as the average percent change in consumption, while preserving the treatment effect of very high water users, which the logarithmic transformation of consumption would dampen. Our specification is:

$$y_{it} = \alpha + \gamma_l T_{i,l} + \mathbf{fix_{it}} + \epsilon_{it} \tag{3.1}$$

where y_{it} is normalized GPD, $T_{i,l}$ is a dummy variable for the pooled treatment and each of the five treatment letters ($l = Pooled, 1, 2, ..., 5$), and x_{it} is a vector of control variables. We restrict our sample to the post-intervention period, which comprises the billing months of August, September, and October 2015. While treatment is exogenous by virtue of the randomization, including control variables increases the precision of the estimates. All regressions therefore include average consumption during irrigation seasons prior to the intervention, billing cycle and month fixed effects, and average daily temperature and days of precipitation during the billing cycle. We matched daily weather data from the NOAA weather station at Reno-Tahoe Airport to the exact dates of each customer's water bill to calculate the weather variables. Robust standard errors are clustered at the household level.

3.3.1 Identifying the Effect of Difference from the Peer Group

If normative appeals work by increasing the moral cost of consumption, then we should see stronger treatment effects for households who received a "stronger" social norms and thus a higher moral cost. We define the magnitude of the moral cost as the difference between a household's performance (consumption level or percentage conservation) and the performance of the relevant comparison group. This is one of the cases nested within the consumer utility model of social comparisons presented in Allcott and Kessler (2015), where consumers face a constant marginal moral cost for "inappropriate" consumption. Households consuming further above the norm receive a stronger normative appeal.

As discussed in the introduction, a key feature affecting the interpretation of behavioral responses to social comparisons is that consumers with high pre-treatment consumption are more likely to receive a stronger normative (and negative) message, but they are also likely to have better opportunities to reduce water use at low welfare cost. Figure 3.1 illustrates the existing correlation between baseline water use and the strength of the social norm in traditional social comparisons (our Treatment 4), and demonstrates that our experimental design successfully decouples the two in our new treatment (T5). In each of the two panels (panel (a) for the comparison in gallons and panel (b) for the comparison in percentage reduction, or conservation rate), we partition households into quartiles of baseline water use (displayed on the x-axis). Within each of those quartiles, we further partition households into quartiles of the difference between a household's level of consumption (or conservation rate) and that of its peer group. Since comparisons are based on median consumption (or percent conservation) within the peer group, the first two quartiles (Q1:Much better and Q2:Better) are households who are doing better than their peer group, and the upper two quartiles (Q3:Worse and Q4:Much Worse) represent households who are doing worse.

Panel (a) of Figure 3.1 shows that for the traditional social comparison in gallons most low water users (Q1 on the x-axis) consume less water than their peer group: roughly 90% of the consumers with the lowest water consumption were informed that they used less than their neighbors (Q1 + Q2 of Difference from Peer group (kgal)).[9] Likewise, most high users (Q4 on the x-axis) received a message telling them they used more water than their peers.

This is not the case for our social comparison in percentage terms (Treatment 5): a substantial fraction of low users conserved less than their peers and many high users conserved more than their peers (Figure 3.1, Panel (b)). Even among households in the

[9] The only reason why there are some low users who are above their peer group in the traditional social comparison (T4) is that the norm is based on a peer group - defined by households in the same meter route who have similar number of bedrooms and yard size (above/below the median). By comparison, Ferraro and Price (2013) compare household consumption to the full sample median, producing a treatment where the strength of the descriptive norm is perfectly correlated with pre-treatment consumption. Therefore a household with a high-water-use peer group can be above the median or 75th percentile of the sample-wide distribution of pre-treatment consumption, but still consume less than the peer group.

Figure 3.1: Strength of Normative Message by Quartiles of Pretreatment Consumption

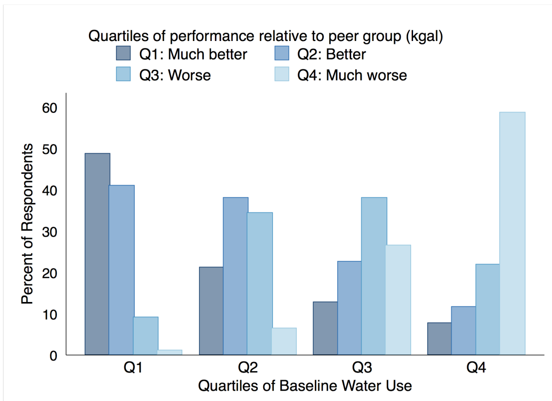

(a) Norm in kgal (T4)

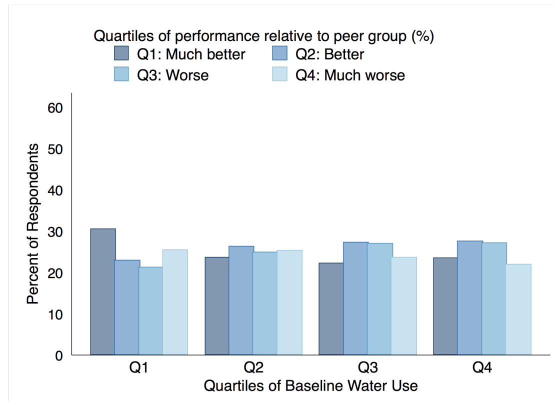

(b) Norm in conservation rate (T5)

Note: The graph displays the percentage of respondents receiving messages divided up by quartiles of the performance relative to the peer group within each quartile of pre-treatment consumption. The x-axis displays the quartiles of pre-treatment consumption and the y-axis displays the percentage of respondents receiving a given message. The performances relative to the norm are designated by the different colored bars. The performance relative to the peer group is defined based on quartiles of the difference between a household's consumption (panel (a) - T4) or conservation rate (panel (b) - T5) and the peer group's consumption.

bottom quartile of pre-treatment water consumption, some reduced their water consumption by less than the median conservation rate in their peer group and thus received a strong normative appeal. Likewise, some households with high pre-treatment consumption reduced consumption by a larger percentage than their peer group. The distribution of norms within each quartile of pre-treatment consumption is remarkably balanced for the conservation rate comparison treatment. Figure 3.1 uses all data from the experimental sample, but the same general pattern of consumption holds if we restrict the sample to the treatment group, the control group, or any combination of treatment and control for individual months of the sample.

We incorporate the information content of the mailers by estimating conditional average treatment effects (CATEs), where we condition on the difference between pre-treatment consumption and the peer group median. The content of the mailers depends on recent water use, which can introduce endogeneity into the estimated treatment effects for treatment groups that receive multiple mailers. The first mailing received only depends on pre-treatment water use in the months immediately preceding the intervention, but additional mailers include information from the first month of treatment. For households randomly-selected to receive two letters, we therefore use only the first month of post-treatment data, before the second letter was received. Since treatment is randomized across the distribution of pre-treatment water use, CATEs provide valid inference - the results can be interpreted as causal treatment effects in the same style as studies that condition on pre-intervention consumption (Allcott, 2011; Ferraro and Miranda, 2013; Brent, Cook and Olsen, 2015).

The CATE model is defined as

$$y_{it} = \alpha + \sum_{c=1}^{k} \gamma_{l,c} T_{i,l} \times C_{i,c} + \sum_{c=1}^{k} \theta_c C_{i,c} + \mathbf{fix_{it}} + \epsilon_{it} \tag{3.2}$$

In this model we are concerned with $\gamma_{l,c}$, which is the CATE for letter l in subgroup c. $T_{i,l}$ is an indicator for whether a household was treated with letter l and $C_{i,c}$ is an indicator for whether a household falls into subgroup c of the conditioning variable $C_{i,c}$. The presence of $C_{i,c}$ accounts for the sample-wide differences in consumption for subgroup c. When we condition based on the performance relative to the peer group, $C_{i,c}$ equals a set of dummies indicating the quartiles of the difference between a household's level of consumption (or percentage conservation) and that of its peers in the month directly preceding treatment. Since control households received no treatment, we calculate the normative appeal that they would have received had they been treated. This accounts for unobserved factors within each subgroup that are common to the treatment and control groups. In the case of the difference from the comparison group in kgal (T4), $C_{i,c}$

controls for the fact that high users generally receive comparisons above their peer group in the gallons comparison (T4) and issues such as mean reversion in the conservation rate comparison (T5). Since treated households only receive one realization of the difference from their peer group, based either on their July or August water consumption, we randomly assign each control household the difference from their peer group based on either July or August in the same proportion as the treated households.[10] We also run specifications where the difference variables for the control households change over time based on the realization in the previous month that produce very similar results.

3.4 Results

3.4.1 Base Results

We begin by reporting the average treatment effects pooling the three treatments of interest, and then briefly discuss each treatment individually. Column 1 of Table 3.4 shows that the average treatment effect (ATE) pooling all three treatments is slightly greater than a 1.5% reduction in consumption. For reference, the generic tips treatment (T1) had no statistically-significant impact on conservation, and the ATE of the tips plus historical information (T2) treatment was slightly less than 1% and statistically significant. Overall, our pooled ATE is a smaller effect than commonly reported for social comparisons: Opower's interventions typically reduced energy consumption by about 2%, and both Ferraro and Miranda (2013) and Brent, Cook and Olsen (2015) find average reductions in consumption of approximately 5%. However, these results should

[10] One third of the treatment group received a single mailer in July, one third received a single mailer in August, and one third received mailers in July and August. Therefore two-thirds of the controls have a comparison based on July consumption and one-third based on August consumption. Moreover, we drop all observations after the first post-intervention month for the one-third of control households that are randomly assigned to stand in for the two-mailer treatments.

be considered in the context of an extensive utility-wide water conservation campaign during the second year of a severe drought. Additionally, given that the aforementioned studies on water examine some of the first interventions using social comparisons for water conservation, the lower treatment effects are consistent with the findings of Allcott (2015) that initial sites often have higher average treatment effects than subsequent sites.

Column (2) breaks down the treatments individually. Each treatment generated statistically significant reductions in consumption and the point estimates are all very close to each other. Columns (3)-(5) reproduce the ATE for each letter in separate regressions using the individual treatment group and the control. This is simply to demonstrate that both the point estimates and the standard errors are almost identical whether we use the entire sample with three treatment dummies or restrict the sample to one treatment and the control. Restricting the sample simplifies the presentation of the results. All subsequent regressions also include controls for temperature, precipitation, bill cycle fixed effects, month fixed effects, and pre-treatment consumption.

Next, we show that results for all of our treatments are consistent with previous research that finds the level of pre-treatment consumption is positively correlated with the conservation response to social comparisons. The results are based on the same model in equation 3.2 but the conditioning variables are quartiles of pre-treatment consumption. Figure 3.2 estimates CATEs by quartiles of pre-treatment consumption for each individual treatment. Each panel reflects the results of one regression of the CATEs; the shaded bars are the point estimates and the error bands are the 95% confidence intervals. Each of the three treatments primarily generates statistically significant savings among high users. Next, we analyze whether the treatment effect for the social comparison treatments (T4 and T5) is driven by pre-treatment consumption or the strength of the normative message.

Table 3.4: Base Regression

	(1) Full Sample (T3-T5)	(2) Individual Treatments	(3) Rate (T3)	(4) Social Comp. kgal (T4)	(5) Social Comp. % (T5)
All Treatments	-1.550***				
	(0.304)				
Rate (T3)		-1.604***	-1.596***		
		(0.466)	(0.466)		
Social Comp. kgal (T4)		-1.455***		-1.451***	
		(0.446)		(0.446)	
Social Comp. % (T5)		-1.591***			-1.591***
		(0.473)			(0.473)
Weather Controls	Yes	Yes	Yes	Yes	Yes
Bill Cycle FEs	Yes	Yes	Yes	Yes	Yes
Month FEs	Yes	Yes	Yes	Yes	Yes
Baseline Water	Yes	Yes	Yes	Yes	Yes
Households	33,937	33,937	25,532	25,529	25,510
Observations	96,759	96,759	74,530	74,494	74,449

Note: The dependent variable is normalized average daily water consumption; the coefficients can be interpreted as a percentage change in consumption. Robust standard errors clustered at the household level are reported in parenthesis. *** p<0.01, ** p<0.05, * p<0.1

3.4.2 Difference from the Peer-Group Descriptive Norm

Our first CATE specification defines strong normative appeals as mailers that informed households that they were *above* their peer group: either using more water than their peer group (T4) or a lower percentage reduction than their peer group (T5). Our second specification estimates CATEs based on quartiles of the difference from the peer group, providing more variation in the performance relative to the peer group.

Column (1) of Table 3.5 shows the CATEs based on being above or below the peer group for the social comparison in gallons. All of the savings essentially come from households who are informed that they are consuming more than their peers. The same pattern holds when we estimate CATEs within quartiles of the difference from the comparison group (Column (2) of Table 3.5). There is a monotonic relationship between the

Figure 3.2: Conditional Average Treatment Effects by Quartiles of Pre-treatment Consumption

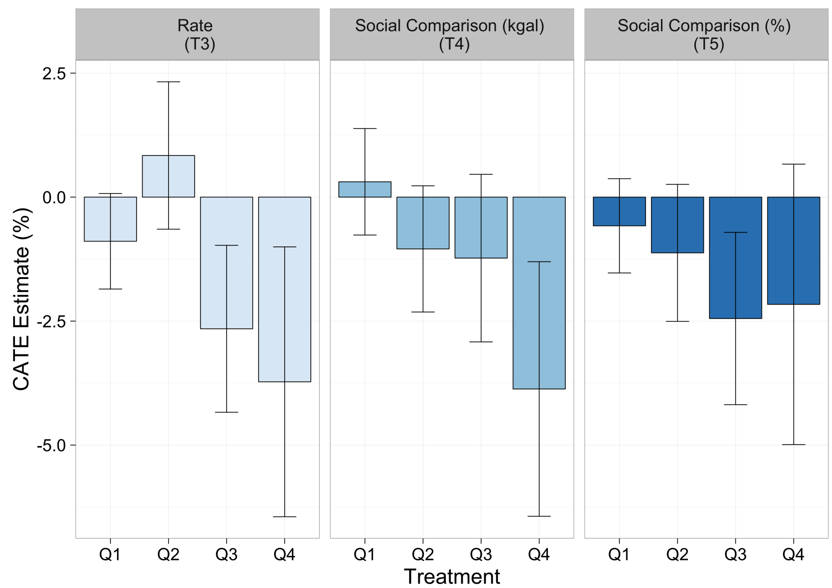

Note: Each bar graph represents the output of one regression where the dependent variable is normalized average daily water consumption. The bars are the point estimates of the CATEs for each quartile of pre-treatment consumption, and the error bars are 95% confidence intervals constructed from robust standard errors clustered at the household level. All regressions include controls for temperature, precipitation, bill cycle fixed effects, month fixed effects, and pre-treatment consumption.

difference from the peer group and the estimated treatment effects.[11] As described above, it is impossible to determine whether these results are due to the moral cost imposed by the message or due to the fact that high users have more scope to correct internalities.

Columns (3)-(4) of Table 3.5 estimate the CATEs based on the difference from the comparison group framed in terms of percentage conservation, which separates the performance relative to the peer group from pre-treatment consumption. The pattern of results is similar: households treated with a message telling them they conserved less than their peers save water, whereas those who conserved more than their neighbors do

[11]Negative differences, which reflect households who are doing better than their peers, are associated with insignificant treatment effects that are either positive or small and negative. Positive differences, which reflect households who are doing worse than their peers, are associated with significant and negative treatment effects, which increase with the difference from peer group.

not. In fact, the results in column (4) using quartiles of the difference from the peer group are even starker. Among households treated with the social comparison in percentage terms, those who were conserving much less than their peers before receiving the letter (Treat*Q4 Norm) saved over 5.7%. This is compared to a 1.3% increase (though not statistically significant) among households who reduced consumption by a much larger percentage than their peers (Treat*Q1 Norm). These results are consistent with social comparisons operating through an increase in the moral cost of consumption: larger increases in the moral cost lead to larger reductions in consumption. However, these results do not rule out the second behavioral channel: stronger messages indicate more potential to correct internalities.

To further disentangle the CATE's from the level of consumption, we divide the sample based on below-median baseline water users (column 5) and above-median users (column 6) and estimate the CATE's within each subsample. We do not perform the analysis for the social comparison in gallons due to the strong correlation between the difference from the peer group and pre-treatment consumption. The pattern of results is the same for both low users who likely have fewer low-cost opportunities to correct internalities and high users who likely have more opportunities. This is strong evidence that the message operates by increasing the moral cost of consumption.

Moreover, higher levels of pre-treatment consumption act to magnify the results. The savings for households in the upper two quartiles of the difference from peer group (worse performance than peers, Treat*Q3 Norm and Treat*Q4 Norm) are roughly twice as large for high users compared to low users. This suggests that there is a synergistic effect between the moral cost of consumption and the motivation to correct internalities in consumption. Alternatively, similar conservation actions, such as changing the irrigation controller, are scaled up for high users so that the same action leads to more conservation. There is no statistically-significant increase in consumption among those who outperform

Table 3.5: Difference from Peer Group

	Norm in Gallons (T4)		Norm in % (T5)			
	(1)	(2)	(3)	(4)	(5)	(6)
	All	All	All	All	Low Users	High Users
Treat*Below Peer	0.205		0.184			
	(0.644)		(0.766)			
Treat*Above Peer	-2.208**		-3.768***			
	(0.914)		(0.907)			
Treat*Q1 Norm		0.347		1.267	1.423	1.310
		(0.901)		(1.166)	(0.980)	(2.119)
Treat*Q2 Norm		0.0579		-0.117	0.477	-0.624
		(0.874)		(0.918)	(0.884)	(1.521)
Treat*Q3 Norm		-1.692*		-2.900***	-1.637**	-3.877***
		(1.014)		(0.898)	(0.828)	(1.465)
Treat*Q4 Norm		-2.627*		-5.670***	-4.358***	-7.092***
		(1.528)		(1.368)	(1.313)	(2.349)
Observations	68,933	68,873	68,929	68,834	33,129	35,477

Note: The dependent variable is normalized average daily water consumption, and the sample is restricted to the month after a household's first mailer. CATEs are estimated based on the difference (above/below) of household consumption relative to the peer group. Column headers All represents the whole sample, and Low/High Users restrict the sample to households above or below median pre-treatment consumption. All regressions include controls for temperature, precipitation, bill cycle fixed effects, month fixed effects, and pre-treatment consumption. Robust standard errors clustered at the household level are reported in parenthesis. *** p<0.01, ** p<0.05, * p<0.1

their peers. Overall, these results suggest that social comparisons appear to operate at least partly through increasing the moral cost of consumption for households performing worse than their peers.

3.4.3 Survey Evidence for Behavioral Channels

To provide additional evidence of the behavioral mechanisms of response to our treatments, we link our data to responses from a utility-sponsored survey conducted after the intervention (Christman, 2016). Approximately 1,500 households in our experimental sample answered the survey, and although the survey respondents are not representative of the greater service area (overall they used less water prior to treatment and lived in smaller, less expensive homes), the survey sample is balanced across treatments and the control group for key variables (Tables C.6-C.8 show balance statistics for the survey sample). We pool the two social comparison treatments to focus on how differences in attitudes influence the treatment effects for the rate information treatment, which frames conservation as a financial savings, compared to the social comparisons, which focus more heavily on moral suasion.

Table 3.6 shows the results of regressions that interact the treatment effects with indicators for households that answered yes to questions related to their motivations for using water or capacity to conserve. In Table 3.6 "Survey Variable" corresponds to a dummy indicating that the respondent answered yes to the question indicated in the column header. Summary statistics for the survey variables are available in Appendix C. The first column simply replicates the treatment effects for the survey sample; the point estimates are similar, but slightly higher and not statistically significant.[12] The next three columns (2-4) represent pro-social motivations. These households indicated that they would be willing to let their lawn go brown if the drought continued (column (2)); thought that water scarcity was a serious concern (column (3)); and indicated that they saved water to help the environment and/or their community (column (4)). The interaction terms for all three pro-social indicators are negative and often significant, indicating that pro-social households were more responsive to *both* types of treatments. Recall also that the financially-oriented treatment also contained a reminder about TMWA's 10% goal and an exhortation to do their part.

The next two columns (5-6) represent households who make their water decisions based on grounds other than water scarcity, environmental, or community concerns. These households indicated that they should be able to use water as they choose (columns (5)) or saved water in order to save money (column (6)). We argue that households that

[12]Pooling the three treatments together does produce a statistically significant treatment effect in the survey sample.

Table 3.6: Survey Evidence of Behavioral Channels

		Pro-social Motivations			Alternative Motivations		Capacity
	(1) Baseline	(2) Lawn Brown	(3) Water Scarcity	(4) Environment & Community	(5) Use How I Want	(6) Save Money	(7) Already Conserved
Rate (T3)	-1.965	-0.631	0.536	3.772	-1.216	-2.739	-5.134*
	(1.920)	(2.091)	(2.460)	(5.230)	(2.271)	(3.243)	(2.633)
Social Comp. (T4-T5)	-2.573	-1.556	0.00760	5.433	-2.812	-5.448*	-2.421
	(1.571)	(1.636)	(2.095)	(4.472)	(1.804)	(2.807)	(1.992)
Survey Variable		4.139	1.472	0.0162	0.847	-2.152	-0.0708
		(3.856)	(1.784)	(2.134)	(2.007)	(1.826)	(1.768)
T3*Survey Variable		-13.52***	-6.496*	-6.696	-3.135	1.183	7.285*
		(5.065)	(3.925)	(5.628)	(4.204)	(4.046)	(3.840)
T4-T5*Survey Variable		-12.71**	-6.709**	-9.323*	1.489	4.285	-0.433
		(5.427)	(3.087)	(4.780)	(3.511)	(3.403)	(3.126)
Households	1,536	1,536	1,536	1,536	1,536	1,536	1,536
Observations	4,371	4,371	4,371	4,371	4,371	4,371	4,371

Note: The dependent variable is normalized average daily water consumption, and the sample is restricted to survey respondents. The first column replicates the ATE for the rate information treatment and pooled social comparisons. Columns (2) - (7) interact the treatment variable with dummy variables if the respondent answered that survey question positively: "Survey Variable" refers to the dummy variable indicated by the column headers. All regressions include controls for temperature, precipitation, bill cycle fixed effects, month fixed effects, and pre-treatment consumption. Robust standard errors clustered at the household level are reported in parenthesis. *** p<0.01, ** p<0.05, * p<0.1

want the freedom to use water as they choose and save water for financial reasons will be less responsive to the social comparison treatments. This is indeed the case: while the results are not significant (nor significantly different) relative to the rate information treatment, the financially motivated households are less responsive to treatment. In the last column, we show results for households that have already conserved and likely have lower capacity to correct internalities. The rate information treatment is significantly less effective for households with previous conservation efforts, which suggests that these households might have previously over-perceived the marginal cost of water. By contrast, for households who have not already conserved, rate and monetary savings information leads to a 5% reduction in consumption that is significant at the 10% level. There is essentially no impact of prior conservation on the savings from social comparisons.

Overall, these survey results indicate that pro-social households are more responsive to both treatments, which corroborates our main finding that social comparisons operate in part by imposing a moral cost on consumption. There is also evidence that the rate information treatment is relatively more effective at prompting households to correct internalities than the social comparisons, a point that we investigate further in the next section.

3.4.4 Persistence & Additionality

To further investigate the mechanisms of response to our treatments, we analyze the variation in the persistence of the treatment effects. We posit that increasing the moral cost of consumption has a temporary impact on consumers. This is consistent with previous studies that the magnitude of the treatment effect decays as time from the last mailer increases, although treatment effect are often detectable long after treatment is discontinued (Ferraro and Miranda, 2013; Bernedo, Ferraro and Price, 2014; Allcott and Rogers, 2014; Brandon et al., 2017). Moreover, consumers who experience an increased moral cost may seek to adjust consumption by making temporary behavioral changes, whereas consumers seeking to correct internalities are more likely to undertake investments that will have more lasting conservation effects.[13] Under this reasoning, treatments that operate through increasing moral cost should have less persistent treatment effects than treatments that prompt consumers to correct internalities. Treatments that generate similar patterns of savings over time are likely to operate through similar behavioral mechanisms. To this end, we compare the persistence of the treatment effects for the treatment focused on rate information and monetary savings (T3) and the social comparisons (T4-T5). The rate information treatment directly prompts consumers to address

[13]This is not to say that moral cost cannot induce investments in efficiency. Brandon et al. (2017) provide evidence that OPower treatments induce savings after the initial tenants who received the reports leave, which is consistent with investments in energy efficiency.

internalities in consumption in order to save money, whereas our previous results suggest that social norms operate in part through a moral motivation. We also examine whether there is a differential effect of sending an additional letter on these treatments. We expect to see less of an effect of additionality on nudges that operate primarily through prompting consumers to correct internalities in consumption. If a consumer receives useful information that helps them re-optimize, seeing similar information a second time should not generate additional savings. In order to understand the implications of the timing and number of treatments we run the following regression,

$$\begin{aligned} y_{it} = \alpha + \gamma_{1,l} 1\{First\ Mailer, First\ Month_{it,l}\} + \\ \gamma_{2,l} 1\{First\ Mailer, Second\ Month_{it,l}\} + \\ \gamma_{3,l} 1\{First\ Mailer, Third\ Month_{it,l}\} + \\ \gamma_{4,l} 1\{Second\ Mailer_{it,l}\} + \mathbf{fix_{it}} + \epsilon_{it} \end{aligned} \tag{3.3}$$

where the treatments are now broken down by the months since the first letter was received in addition to a dummy for whether there was a second mailer ($\gamma_{4,l}$). We analyze three months from the first letter since this corresponds to the end of the summer demand season (October). Therefore, γ_2 and γ_3 test whether the treatment was persistent and γ_4 represents the additionality of the second mailer.[14] Table 3.7 presents the results for the rate and social comparison treatments.

The rate information treatment is very persistent with similar magnitude treatment effects in each of the three months following the first letter, and a second letter generates no additional savings. By contrast, the social comparison treatments (T4 and T5) are not persistent; however these treatments achieve significant additionality from the second mailer. The savings from the first treatment decreases by roughly 50% or more per month. Sending another mailer increases conservation by almost 2%. The results for the

[14] The second mailer affects consumption in two months: September and October.

Table 3.7: Persistence & Additionality of Second Mailers

	All		Below Median		Above Median	
	(1) Rate	(2) Social Comp.	(3) Rate	(4) Social Comp.	(5) Rate	(6) Social Comp.
1st Letter, 1st Month	-1.395**	-1.474***	-0.00465	-0.536	-2.680***	-2.371***
	(0.542)	(0.403)	(0.565)	(0.375)	(0.903)	(0.692)
1st Letter, 2nd Month	-1.899***	-1.108**	-0.241	-0.839**	-3.326***	-1.374*
	(0.592)	(0.449)	(0.565)	(0.417)	(0.998)	(0.762)
1st Letter, 3rd Month	-1.499**	-0.304	-0.305	-0.773	-2.850**	0.149
	(0.739)	(0.554)	(0.677)	(0.547)	(1.250)	(0.917)
2nd Letter	0.158	-1.922***	-0.453	-0.321	0.743	-3.456***
	(0.759)	(0.558)	(0.694)	(0.522)	(1.296)	(0.951)
Households	25,532	29,722	12,305	14,311	13,227	15,411
Observations	74,530	85,586	35,923	41,206	38,607	44,380

social comparisons are consistent with the Allcott and Rogers (2014) results where the pattern of action and backsliding stems from a model of consumption cues, and indicate similar behavioral motivations for response to social comparisons of energy consumption. The same pattern holds when we estimate regressions for each of the social comparisons individually.

We also divide the sample based on whether the household was above or below the median of pre-treatment consumption to compare the persistence of treatment for low and high users. We see high users within both treatments have a pattern of savings that is consistent with the overall treatment effect pattern. High users drive the treatment effect, which is not surprising, since most of the water savings in all of our five treatments come from high users. The fact that social comparisons are not persistent and require additional mailers even among high users with greater potential to correct internalities, provides further evidence that households are responding to social comparisons due to increased moral costs.

The differences in persistence and additionality across treatments can be interpreted in the framework of intent-oriented actions and impact-oriented actions developed by Attari (2014). Intent-oriented actions include behavioral changes such as turning off the

Table 3.8: Mechanisms of Water Conservation: Investment and Changes in Behavior

	(1) Invested in Efficiency	(2) Behavioral Changes
Rate (T3)	0.0528	0.0256
	(0.0354)	(0.0240)
Social Comparisons (T4-T5)	-0.0759***	-0.00339
	(0.0228)	(0.0188)
HOA w/ Landscape Restrictions	Yes	Yes
Already Invested in Traditional Landscape	Yes	Yes
Already Invested in Water Efficiency	Yes	Yes
Observations	1,536	1,536

water when brushing teeth, while impact-oriented actions include investments such as replacing a toilet. Households perform intent-oriented actions with the intention to help the environment, while impact-oriented actions are focused on achieving more substantial conservation. The results suggest that nudges that operate by prompting consumers to correct internalities (T3) lead to impact-oriented actions, while nudges that increase the moral cost of consumption (T4 and T5) lead to intent-oriented actions.

Indeed TMWA survey results corroborate these findings. A series of questions asked households what types of actions they took to reduce water consumption. We focused on responses to questions that correspond to our treatment period. The types of actions can be divided into two general categories: investments in efficiency such as repairing leaks or replacing lawn with xeriscaping and behavioral changes such as taking shorter showers. We generate indicator variables for whether a household engaged in any investment or any behavioral change and use these as the dependent variables in a linear probability model estimated with ordinary least squares. The results, presented in Table 3.8, show that households treated with social comparisons are significantly less likely to invest in efficiency compared to households that received the rate information treatment. There is no statistical difference in the relative propensity to engage in behavioral changes. This provides further evidence of the different mechanisms that are operating for the rate and social comparison treatments. The results are consistent with households correcting internalities in response to rate information and making short-term cuts in consumption in response to social comparisons.

3.4.5 Robustness

We run several robustness tests to assure that our results are not actually being driven by spurious correlation or the discrete effects of being above or below the norm. The treatment period started in August 2015 utilizing July 2015 data on the mailers. Our first robustness check runs falsification tests by using the content generated by June 2015 data. We generate the difference from the peer group using June 2015 for both treated and control households to see if the types of households that were above or below the peer group were actually driving the results as opposed to the actual content of the mailers. This is a valid falsification test of our results due to the fact that variation in seasonal consumption leads to variation in household performance relative to a peer group. Some households can be below their peer group earlier in the summer, but have much higher peak water use than the peer group later in the summer. Table 3.9 shows results using the simulated mailer content. The results are substantially different than Table 3.5. Households above and below the false peer group have statistically significant CATEs in the conservation rate information treatment (column (3)), and neither of the CATEs are statistically significantly different from each other. The CATE for households below the peer group in the traditional social comparison is not statistically significant, but this is likely due to the fact that they are low users and their real letter also indicated that they were below their peer group. The specification with quartiles of the norm in percentage terms (column (4)) shows that three out of the four CATEs generate savings over 2% and are significant at the 5% level. In general, households above the false peer group do save more water, but this is likely due to serial correlation in water consumption; households above the peer group in June are more likely to be above the peer group in July and August.

The regressions presented in Table 3.5 have different numbers of observations for each household. Households treated with one mailer in July have three observations after treatment (August, September, and October), households treated with one mailer in August have two observations (September and October), and households treated with two mailers only have the observation immediately following their initial letter. In order to test if the difference in the number of observations impacts the results we drop all households treated with two mailers and only use the first two observations after the first mailer. The results are very similar to our preferred specification, which incorporates all valid observations, and are available in Table C.9 in Appendix C. We also run the same set of regressions limiting the sample to the first month after a household receives a mailer. This sample includes observations immediately after a household's first mailer for households treated with two mailers. The results, presented in Table C.10 in Appendix C, also display the same general pattern where the magnitude of the treatment effect depends on the difference from the peer group.

Since the mailers also present information on historical consumption and reference the utility-wide 10% goal we also run regressions for quartiles of the difference between an household's conservation rate and the 10% goal. The distance from the 10% goal does affect the CATEs, but this is at least in part due to the correlation with pre-treatment water use and the conservation rate. The results, presented in Table C.11 in Appendix C, do not show the monotonic relationship that we observe for CATEs based on the difference from the peer group, and they vary across treatments. Lastly, we test for the discrete effects for moving above a peer group or failing to meet the 10% goal in a regression discontinuity design. We find no effect of moving above the peer group in

Table 3.9: Falsification Test for Difference from Peer Group

	Norm in Gallons (T4)		Norm in % (T5)	
	(1)	(2)	(3)	(4)
Treat*Below Peer	-0.660		-1.184	
	(0.636)		(0.786)	
Treat*Above Peer	-1.794*		-2.181**	
	(0.938)		(0.934)	
Treat*Q1 Norm		-0.148		-2.274**
		(0.910)		(1.093)
Treat*Q2 Norm		-1.166		-0.238
		(0.881)		(1.090)
Treat*Q3 Norm		-2.197**		-1.629*
		(1.108)		(0.990)
Treat*Q4 Norm		-1.290		-2.708*
		(1.512)		(1.615)
Observations	68,933	68,930	68,929	68,896

Note: The dependent variable is normalized average daily water consumption, and the sample is restricted to the month after a household's first mailer. CATEs are estimated based on the quartiles of the difference of household consumption relative to the peer group. Robust standard errors clustered at the household level are reported in parenthesis. *** p<0.01, ** p<0.05, * p<0.1

either gallons or percentage terms, nor do we find any effect of moving slightly above the 10% goal. This is consistent with the findings of Allcott (2011) for social comparisons in energy and with unpublished results from the project reported in Brent, Cook and Olsen (2015). Appendix C describes the regression discontinuity design in more detail and presents both graphical evidence and the regression discontinuity estimates based on Calonico, Catteneo and Titiunik (2015).

3.5 Conclusions

We designed and implemented an experiment using multiple treatments in a policy environment where a municipal water utility was running a well-publicized campaign asking every customer to reduce their monthly consumption by 10%. We find that two versions of a social comparison and a third treatment providing information on monetary savings lead to a 1.5% reduction in water consumption relative to the control group. This reduced consumption was *in addition* to the overall system-wide reduction achieved of over 15%. While there may be some concern that the 10% goal reduces the external validity of our experimental results, calls for uniform reductions are a common utility practice as evidenced by Governor Brown's 2015 call for reducing water consumption in California by 25%. We also note that there was a utility-wide appeal for conservation in the well-known water conservation experiments in Cobb County examined by Ferraro and Price (2013). The treatments generated effects that were similar in magnitude and statistically indistinguishable from each other.

We introduced a new normative appeal method that decouples pre-treatment consumption from the difference from the peer group by framing savings in terms of percentage achievement to the target conservation goal of 10%. The strength of the normative message is a strong driver of variation in the treatment effect over low and high consumption levels for our new conservation rate social comparison treatment, although treatment effects are magnified for larger water users. These results are consistent with an interpretation that social comparisons impose a moral cost on consumption, and the size of the moral tax depends on the magnitude of the difference between household performance and that of a peer group. Both social comparison treatments display the pattern of "action and backsliding" shown by Allcott and Rogers (2014) indicating that they operate through a similar behavioral mechanism; however, we find a different pattern over time for the treatment that highlights financial savings. In contrast to the social comparisons, a single mailer for the financial savings treatment is persistent over the study period, with no additional conservation generated by an additional mailer. This is consistent with households responding to the social comparisons with intent-oriented actions that fade over time, whereas the financial treatment may lead to re-optimizing water use to correct internalities, as in a household production framework, where responses are more persistent (Attari, 2014). Survey evidence shows that households in the financial savings information treatment are more likely to invest in water efficiency relative to households treated with social comparisons.

The behavioral mechanisms underlying the response to social comparisons are important when evaluating economic welfare effects of programs designed to promote voluntary reduction in water consumption. Use of customized information in mailings that prompt consumers to re-optimize consumption to address internalities will generally produce higher welfare gains than those that impose a moral cost on consumption (Allcott and Kessler, 2015). Individuals operating with moral motivations may act to reduce existing internalities, thereby limiting welfare-reducing effects of normative appeals. However, there is likely variation in the available internalities across regions due to differences in whether irrigation is necessary to maintain healthy landscape as well as variation in the extent to which utilities have promoted investment in efficient fixtures and appliances. In regions similar to the area serviced by TMWA, which experiences seasonal outdoor water use that is 4 to 5 times higher than indoor use alone, consumers may have ample opportunities to correct internalities by investing in irrigation efficiency and water efficient landscape. Similarly, households in areas which have experienced fewer instances of calls for water conservation are likely to have room to exploit internality correction, since consumers have not been previously prompted to invest in efficient appliances. Thus, prior experience with conservation is likely to impact the welfare effects of normative appeals for water and energy conservation. Understanding the behavioral mechanisms also helps policymakers select interventions that meet specific objectives. For urban water managers there are important distinctions between interventions that temporarily reduce water consumption during a drought and permanently reduce consumption after water supplies have recovered. Therefore, understanding different nudges prompt different types of actions expands the toolbox for policymakers to address the different forms of water scarcity that varies across regions.

A practical contribution of our work is a more sophisticated understanding of the motivations that underlie responses to normative appeal campaigns to improve welfare outcomes of future campaigns. Utilities use normative messaging campaigns to reduce energy or water consumption ostensibly to improve welfare in situations where unpriced consumption externalities cause the social marginal cost to be greater than the retail price. Normative messaging campaigns that also improve welfare for a household whose current consumption does not in fact maximize their own welfare, by cuing the correction of internalities, generates both private and external benefits. Conservation campaigns that focus on addressing internalities, as opposed to imposing a moral cost, have the potential to increase the welfare benefits of these tools more than reliance on moral appeals alone, along with generating more persistent treatment effects.

Appendix A

The Effect of Price Information on Consumer Behavior Under Nonlinear Tariffs: Evidence from a Water Utility Merger

A.1 First Order Conditions

The first order conditions from the expected utility maximization are as follows:

$$
\begin{aligned}
&U_1(q_U^*, I - p_1 q_U^*) - p_1\, U_2(q_U^*, I - p_1 q_U^*) \\
&\quad + f(q_U^*) \left[\int_0^{q_U^*} U(q_U^*, I - p_1 \tilde{k} - p_2(q_U^* - \tilde{k})) dF(\tilde{k}) - U(q_U^*, I - p_1 q_U^*) \right] \\
&\quad + F(q_U^*) \left[\int_0^{q_U^*} \Big[U_1(q_U^*, I - p_1 \tilde{k} - p_2(q_U^* - \tilde{k})) \right. \\
&\quad - p_2\, U_2(q_U^*, I - p_1 \tilde{k} - p_2(q_U^* - \tilde{k})) \Big]\, dF(\tilde{k}) + U(q_U^*, I - p_1 q_U^*) \\
&\quad \left. - U_1(q_U^*, I - p_1 q_U^*) + p_1\, U_2(q_U^*, I - p_1 q_U^*) \right] = 0
\end{aligned} \tag{A.1}
$$

where $U_1(\cdot) = \frac{\partial U}{\partial q}$ and $U_2(\cdot) = \frac{\partial U}{\partial z}$.

For Quasilinear utility, which represents risk neutral preferences over income, the problem becomes:

$$\max_q \Bigg[\alpha_i \, log(q) + I - p_1 \, q + \left[\frac{\Phi\left(\frac{q-k}{\sigma_i}\right) - \Phi\left(\frac{-k}{\sigma_i}\right)}{1 - \Phi\left(\frac{-k}{\sigma_i}\right)} \right] (p_2 - p_1) \left[k + \sigma_i \left[\frac{\phi(\frac{-k}{\sigma_i}) - \phi(\frac{q-k}{\sigma_i})}{\Phi(\frac{q-k}{\sigma_i}) - \Phi(\frac{-k}{\sigma_i})} \right] - q \right] \Bigg] \tag{A.2}$$

The first order conditions are as follows:

$$\frac{\alpha_i}{q^*_{U_i}} + (p_2 - p_1) \left[\left[\frac{\frac{1}{\sigma_i} \phi(\frac{q^*_{U_i}-k}{\sigma_i})}{1 - \Phi(\frac{-k}{\sigma_i})} \right] \left[k + \sigma_i \left[\frac{\phi(\frac{-k}{\sigma_i}) - \phi(\frac{q^*_{U_i}-k}{\sigma_i})}{\Phi(\frac{q^*_{U_i}-k}{\sigma_i}) - \Phi(\frac{-k}{\sigma_i})} \right] - q^*_{U_i} \right] + \left[\frac{\Phi(\frac{q^*_{U_i}-k}{\sigma_i}) - \Phi(\frac{-k}{\sigma_i})}{1 - \Phi\left(\frac{-k}{\sigma_i}\right)} \right] \left[\left[\frac{(\Phi(\frac{-k}{\sigma_i}) - \Phi(\frac{q^*_{U_i}-k}{\sigma_i}))\, \phi'(\frac{q^*_{U_i}-k}{\sigma_i}) - (\phi(\frac{-k}{\sigma_i}) - \phi(\frac{q^*_{U_i}-k}{\sigma_i}))\, \phi(\frac{q^*_{U_i}-k}{\sigma_i})}{(\Phi(\frac{q^*_{U_i}-k}{\sigma_i}) - \Phi(\frac{-k}{\sigma_i}))^2} \right] - 1 \right] \right] - p_1 = 0 \tag{A.3}$$

Figures A.1 (a) and A.1 (b) show the equivalence of maximizing expected utility and maximizing utility with respect the expected budget constraint in the case of Quasilinear utility. For a consumer with preferences that lead to a tier 1 consumption choice under full information, in Figure A.1 (a), the uninformed consumption choice is roughly two thousand gallons less than the informed consumption choice in this example. For a consumer with preferences that lead to a tier 2 consumption choice under full information, in Figure A.1 (b), the uninformed consumption choice is roughly two thousand gallons more than the informed consumption choice in this example.

For constant-relative risk aversion (CRRA) utility, which represents risk averse preferences over income, the problem becomes:

Figure A.1: Expected Utility Maximization for Quasilinear Utility

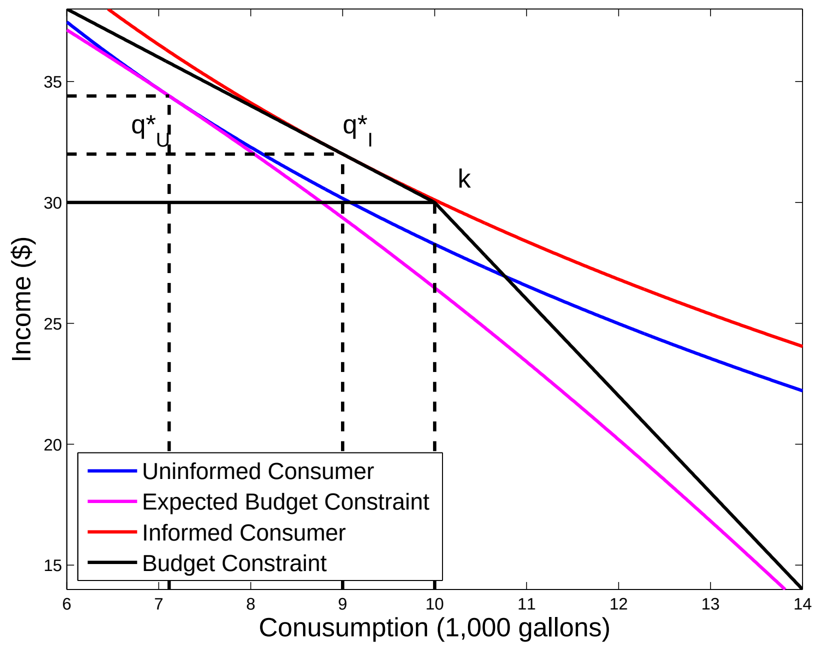

(a) q_I^* vs. q_U^* below tier threshold

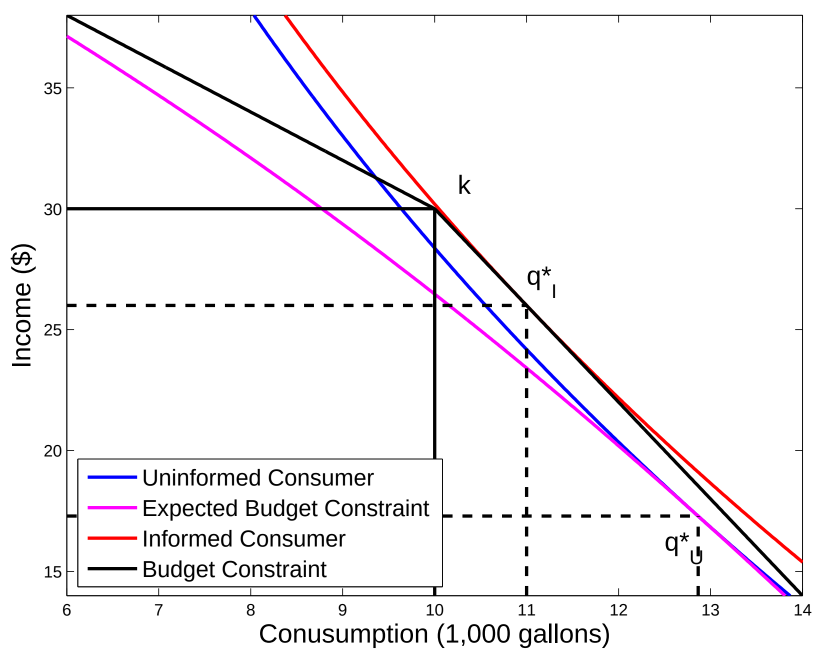

(b) q_I^* vs. q_U^* above tier threshold

Note: These simulations assume $k = 10,000$ gallons, $p_1 = \$2/1,000$ gallons, $p_2 = \$4/1,000$ gallons, quasi-linear preferences: $U_i(q, I - B(q; \tilde{k}_i, p_1, p_2)) = \alpha_i \, log(q) + I - B(q; \tilde{k}_i, p_1, p_2)$, and truncated Normal beliefs: $F(\tilde{k}_i) = \frac{\Phi\left(\frac{\tilde{k}-k}{\sigma_i}\right) - \Phi\left(\frac{-k}{\sigma_i}\right)}{1 - \Phi\left(\frac{-k}{\sigma_i}\right)}$. Figure A.1 (a) compares informed an uniformed choices for a consumer who prefers to consume below the tier threshold under full information. Figure A.1 (b) shows a similar comparison for consumption choices above the tier threshold.

$$U(q, I - B(q;\tilde{k}, p_1, p_2)) = \alpha_i \ \ log(q) + \begin{cases} \frac{1}{1-\theta}(I - B(q;\tilde{k}, p_1, p_2))^{(1-\theta)}, & \text{if } \theta > 0, \theta \neq 1 \\ ln(I - B(q;\tilde{k}, p_1, p_2)), & \text{if } \theta = 1 \end{cases}$$

$$\begin{aligned} &\max_q \Bigg[\alpha_i \ \ log(q) + \frac{1}{1-\theta}(I - p_1\, q)^{(1-\theta)} + \\ &\left[\frac{\Phi\left(\frac{q-k}{\sigma}\right) - \Phi\left(\frac{-k}{\sigma}\right)}{1 - \Phi\left(\frac{-k}{\sigma}\right)}\right] \left[\int_0^q \frac{1}{1-\theta}(I - p_1\tilde{k} - p_2(q - \tilde{k}))^{(1-\theta)} d\Phi\left(\frac{\tilde{k} - k}{\sigma}\right) - \right. \\ &\left.\frac{1}{1-\theta}(I - p_1\, q)^{(1-\theta)}\right]\Bigg] \end{aligned} \tag{A.4}$$

First order conditions:

$$\begin{aligned} &\frac{\alpha_i}{q} - p_1(I - p_1\, q)^{-\theta} + \\ &\left[\frac{\frac{1}{\sigma}\,\phi(\frac{q-k}{\sigma})}{1 - \Phi(\frac{-k}{\sigma})}\right] \left[\int_0^q \frac{1}{1-\theta}(I - p_1\tilde{k} - p_2(q - \tilde{k}))^{(1-\theta)} d\Phi\left(\frac{\tilde{k} - k}{\sigma}\right) - \right. \\ &\left.\frac{1}{1-\theta}(I - p_1\, q)^{(1-\theta)}\right] + \left[\frac{\Phi\left(\frac{q-k}{\sigma}\right) - \Phi\left(\frac{-k}{\sigma}\right)}{1 - \Phi\left(\frac{-k}{\sigma}\right)}\right] \\ &\left[\int_0^q -p_2\,(I - p_1\tilde{k} - p_2(q - \tilde{k}))^{-\theta} d\Phi\left(\frac{\tilde{k} - k}{\sigma}\right)\right. \\ &\left. + \frac{1}{1-\theta}(I - p_1\, q)^{(1-\theta)} - p_1(I - p_1\, q)^{-\theta}\right] = 0 \end{aligned} \tag{A.5}$$

A.2 Heterogeneous Response to Information

This section presents a simple difference-in-differences approach to estimating differential response to information based location of pre-treatment consumption relative to the tier thresholds. This model is based on the third takeaway from the theoretical model, which predicts that consumers with pre-treatment consumption that is far away from any tier threshold will not respond to price information on monthly bills. Households with pre-treatment consumption outside of some bandwidth distance from the tier thresholds serve as a control group for the households that are within this bandwidth distance. These households control for any changes in demand that result from other conservation policies and exogenous shocks to consumption. For a given bandwidth, treatment households within the bandwidth are separated into whether they consume below and above the tier threshold. Let $\bar{d}$ be this bandwidth distance. Household i will be assigned to the below tier treatment group during billing month s if

$$b_{is} = \begin{cases} 1, & \text{if } k - \bar{d} \leq Y_{is} < k \\ 0, & \text{otherwise} \end{cases}$$

and the above treatment group if

$$a_{is} = \begin{cases} 1, & \text{if } k \leq Y_{is} < k + \bar{d} \\ 0, & \text{otherwise} \end{cases}$$

and the control group otherwise. The regression model is as follows:

$$ln(Y_{it}) = \alpha_i + \beta_1 \, b_{is} \, \mathbb{1}_{\text{Post}} + \beta_2 \, a_{is} \, \mathbb{1}_{\text{Post}} + \lambda_t + W_{it}\gamma + \varepsilon_{it} \tag{A.6}$$

where b_{is} is an indicator for household i equal to one if pre-treatment consumption in month s is below the nearest tier threshold and zero otherwise, a_{is} is an indicator for household i equal to one if pre-treatment consumption in month s is above the nearest tier threshold and zero otherwise, and $\mathbb{1}_{\text{Post}}$ is an indicator for billing periods after the merger. If information leads households just below the tier threshold to increase consumption and households just above the tier threshold to decrease consumption then $\beta_1 > 0$ and $\beta_2 < 0$. The drawback of this empirical strategy is that it is dependent on choosing the correct $\bar{d}$ that separates consumers who should be affected by price information from those who should not be affected by price information.

Table A.1 presents results from equation A.6. For this specification $\bar{d} = 4$ (thousand gallons), however the results are similar for bandwidths of three and five thousand gallons. There are consistently positive signs for the below tier treatment group across all three models and the coefficients are significant in models 2 and 3. These results suggest that being just below a tier threshold leads to 1% increase in consumption after the merger relative to the control group. The signs for the above tier treatment group are consistently negative across all models and significant in models 1 and 2. These results suggest that being just above the tier threshold leads to slightly under a 1% decrease in consumption relative to the control group.

Table A.1: Alternative Specification for Low Information Heterogeneous Response to Information

	(1) 2010-2016	(2) 2013-2016	(3) 2014-2016
β_1	0.0051	0.0112**	0.0117***
	(0.0042)	(0.0049)	(0.0038)
β_2	-0.0102**	-0.0069*	-0.0043
	(0.0042)	(0.0035)	(0.0027)
Within R-squared	0.0044	0.0031	0.0034
Households	14,361	14,361	14,361
Observations	468,944	284,427	213,340

Note: The dependent variable is log of monthly water consumption divided by the number of days on the bill. The coefficients can be interpreted as the percentage change in consumption. The sample is limited to summer water use (May-September) for Low Information households only. Controls for all regressions include household FEs, month-by-year FEs, Average Temperature, and Precipitation in Inches. Robust standard errors clustered at the meter reading route level are reported in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.3 Additional Robustness Checks

To further rule out concerns about remaining differences in pre-treatment trends after including a separate linear trend and seasonality for the low information utility, I perform two placebo tests. For the first test, I restrict the study period from 2010 to 2013, before the merger is announced. I assume that the placebo policy occurs in January 2013. A negative and significant estimated treatment effect would indicate that low information average consumption is already decreasing relative to high information average consumption before the merger, causing the estimated treatment effect to overstate the true effect of providing price information on monthly utility bills on water consumption. In a second placebo test, I restrict the study period from 2010 to 2014, and assume the placebo policy occurs in January 2014. A negative and significant estimated treatment effect indicates that news of the merger motivates low information households to decrease consumption, possibly due to anticipation of higher prices or stricter water use policies after the utility merger.

Column 1 of Table A.2 presents the results from the first placebo test. I leave 2014 out of this first placebo test to prevent drought restrictions from confounding the analysis. The placebo test indicates that there are no significant differences in pre-treatment trends before 2014. Column 2 of Table A.2 uses data from 2010-2014 and assumes that a placebo policy occurred in January 2014. Once again, there is no significant placebo effect. These results further confirm that the model adequately controls for differences in pre-treatment trends between low information and high information households.

In addition to the drought restrictions, another possible confounding policy includes several conservation programs offered by the high information utility.[1] At the time of

[1] Programs include providing conservation information, free water audits to detect leaks, and water waste penalties. Water waste penalties are rarely assessed. The utility mainly focuses on educating consumers about water waste and relies on positive encouragement rather than financial penalties.

Table A.2: Additional Robustness Checks for Main Difference-in-differences Results

	(1) 2014 Placebo	(2) 2013 Placebo	(3) Conservation	(4) Field Experiment
β	0.0019	-0.0094	-0.0403***	-0.0371***
	(0.0061)	(0.0091)	(0.0122)	(0.0122)
β*Other Treatment			0.0381***	0.0195***
			(0.0058)	(0.0067)
Low Info Linear Trend	0.0009**	0.0012***	0.0010***	0.0009***
	(0.0004)	(0.0004)	(0.0003)	(0.0003)
Low Info Avg. Use	15.71	15.65		
No Treat Avg. Use			15.32	15.21
Treat Avg. Use			18.15	16.52
Within R-squared	0.0110	0.0106	0.0089	0.0088
Households	60,833	60,810	60,833	60,833
Observations	3,266,388	2,548,715	4,460,397	4,460,397

Note: The dependent variable is log of monthly water consumption divided by the number of days on the bill. The coefficients can be interpreted as the percentage change in consumption. Results based on subsamples as indicated by column titles. Controls for all regressions include household FEs, month-by-year FEs, Average Temperature, Precipitation in Inches, and a separate trend and month FEs for Low Information households. Robust standard errors clustered at the meter reading route level are reported in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

consolidation, all low information households had access to these conservation programs, which were not previously available to them. A difference-in-differences strategy, therefore, estimates the joint effect of providing price information on monthly bill and conservation programs participation on low information households. To address this issue, I estimate a separate treatment effect for low information households who participate in conservation programs at any point after the merger using conservation program records from the high information utility. This specification estimates the differential effect of participation in conservation programs on low information households after the merger. In doing so, I will determine whether the main treatment effect of providing price information on monthly utility bills is actually being driven by the availability of other programs offered by the high information utility. If the treatment effect for households who do not participate in conservation programs is comparable in magnitude to the overall treatment effect estimated in Table 1.4, this will indicate that the results are not driven by voluntary participation in conservation programs.

The results, in Column 3 of Table A.2, are similar to the results for the full sample, in Table 1.4. The estimated treatment effect is negative and significant. Moreover, the coefficient is slightly larger in magnitude if anything. Households who participated in conservation programs used significantly more water after the merger, which only serves to attenuate the overall affect if anything. These results suggest that the availability of conservation programs after the merger does not drive the treatment effect.

Another potential confounder is the treatment of a subset of households during the course of a field experiment that was conducted by the high information utility during summer 2015. The high information utility sent letters to 23,000 low and high information households to encourage them to comply with the drought restrictions. The purpose of these letters was to test the relative effectiveness of various sources of technical information and other normative messages on household conservation.[2] To ensure that the field experiment treatment effect is not driving the change in consumption after the merger, I estimate one set of specifications that exclude treatment households from the sample.

Results in Column 4 of Table A.2 suggest that the field experiment does not dramatically effect the results. This is likely due to the fact that both low information and high information households were treated during the course of this field experiment. The estimated coefficients are almost identical in magnitude to the estimates in Table 1.4. Moreover, the treatment households used significantly more water after the merger relative to non-treatment households. These results suggest that the field experiment, if anything, attenuates the measure effect of providing price information on monthly utility bills.

[2] See Brent et al. (2016) for more information about the details of this field experiment.

A.4 Additional Figures

Figure A.2: Differences in Estimated Linear Trends: Control vs. Treatment

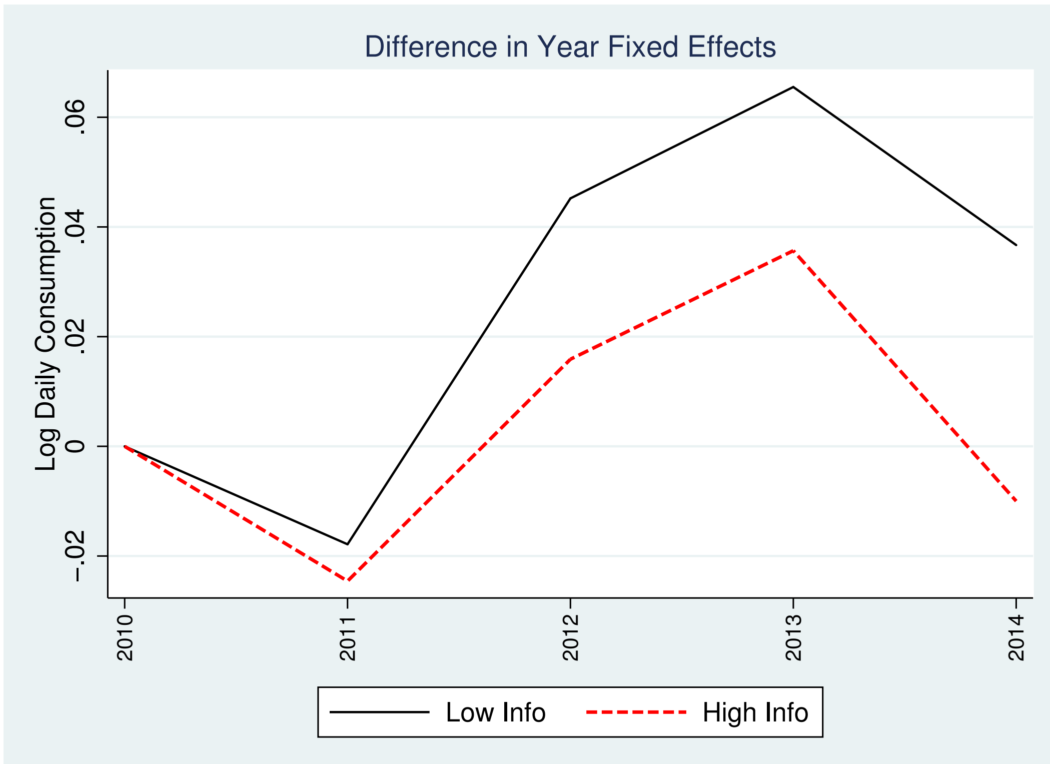

Note: This figure plots estimated year fixed effects for low information and high information households from a model that regresses the log of consumption divided by the number of days on the water bill on separate year and month fixed effects for low information and high information households, weather controls, and household fixed effects. The sample is limited to time periods that occur before the merger. Robust standard errors clustered at the meter reading route level

Figure A.3: Event Study Plot

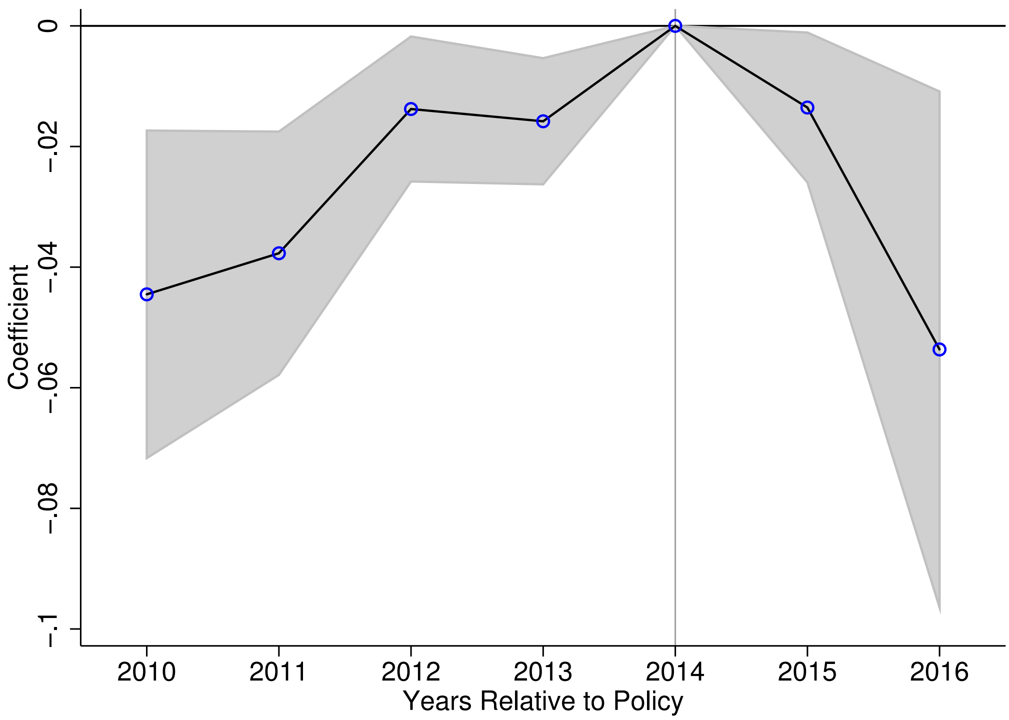

Note: This event study plot uses the last year before the merger (2014) as the reference year. It demonstrates that model 1.7 does not adequately control for differences in pre-treatment trends between low and high information households. Robust standard errors are clustered at the meter reading route level.

Appendix B

Automatic Billing, Paperless Billing, and Consumer Inattentiveness: Evidence from Water Demand

B.1 Duration Results

To analyze the effect of enrollment duration in ABP and PL on consumption I estimate 1) a model that includes a polynomial function of ABP and PL enrollment and 2) a nonlinear function of the treatment by enrollment length. The first model is as follows:

$$ln(Y_{it}) = g(\text{ABP Duration}) + h(\text{PL Duration}) + \alpha_i + \lambda_t + W_{it}\gamma + f(\text{Account Duration}) + \varepsilon_{it} \tag{B.1}$$

where g(ABP Duration) is a cubic function of the duration on ABP and h(PL Duration) is a cubic function of the duration on PL.

To estimate the second specification, I divide the treatments into 1) enrolled for less than 1 year, 2) enrolled for 1-3 years, and 3) enrolled for more than 3 years. The second specification is as follows:

$$\begin{aligned} ln(Y_{it}) =& \beta_1^A\, x_{it}^A\, \mathbb{1}_{1year} + \beta_2^A\, x_{it}^A\, \mathbb{1}_{1-3years} + \beta_3^A\, x_{it}^A\, \mathbb{1}_{3years} \\ &+ \beta_1^P\, x_{it}^P\, \mathbb{1}_{1year} + \beta_2^P\, x_{it}^P\, \mathbb{1}_{1-3years} + \beta_3^P\, x_{it}^P\, \mathbb{1}_{3years} \\ &+ \alpha_i + \lambda_t + W_{it}\gamma + f(\text{Account Duration}) + \varepsilon_{it} \end{aligned} \tag{B.2}$$

where $\mathbb{1}_{1year}$ is an indicator equal to one if account i is enrolled less than 1 year, $\mathbb{1}_{1-3years}$ is an indicator equal to one if account i is enrolled 1-3 years, and $\mathbb{1}_{3years}$ is an indicator equal to one if account i is enrolled at least 3 years.

Table B.1 presents results in columns 1-4 that estimate the effect of enrollment duration. For ABP and PL enrollment duration is positively related to consumption. Moreover, results in columns 3 and 4 suggest this is a linear effect. These results indicate that the average ABP account, which is enrolled for 40 billing periods, increases consumption by 2% relative to non-ABP accounts. The average PL account, which has an enrollment of on 26 billing periods increases consumption by only 0.5% relative to non-PL accounts.

Table B.1: Enrollment Duration Results: 2003-2017

	(1) Preferred Sample	(2) Full Sample	(3) Preferred Sample	(4) Full Sample	(5) Preferred Sample	(6) Full Sample
Duration ABP	0.0005***	0.0004***	0.0007***	0.0010***		
	(0.0001)	(0.0001)	(0.0003)	(0.0002)		
Duration PL	0.0002*	0.0003***	0.0000	0.0003		
	(0.0001)	(0.0001)	(0.0006)	(0.0004)		
Duration ABP Sq.			0.0000	-0.0000**		
			(0.0000)	(0.0000)		
Duration ABP Cubed			-0.0000	0.0000		
			(0.0000)	(0.0000)		
Duration PL Sq.			0.0000	0.0000		
			(0.0000)	(0.0000)		
Duration PL Cubed			-0.0000	-0.0000		
			(0.0000)	(0.0000)		
ABP $<$ 1 year					-0.0138**	0.0013
					(0.0068)	(0.0060)
ABP 1-3 years					0.0096	0.0291***
					(0.0072)	(0.0059)
ABP $>$ 3 years					0.0334***	0.0352***
					(0.0048)	(0.0037)
PL $<$ 1 year					-0.0216**	-0.0181**
					(0.0094)	(0.0076)
PL 1-3 years					-0.0082	-0.0009
					(0.0074)	(0.0055)
PL $>$ 3 years					0.0166***	0.0166***
					(0.0057)	(0.0050)
Weather Controls	Yes	Yes	Yes	Yes	Yes	Yes
f(Act Duration)	Yes	Yes	Yes	Yes	Yes	Yes
ABP Mean Duration	40.34	42.60	40.34	42.60		
PL Mean Duration	25.57	23.60	25.57	23.60		
HH FE's	Yes	Yes	Yes	Yes	Yes	Yes
Period FE's	Yes	Yes	Yes	Yes	Yes	Yes
Within R-squared	0.0112	0.0109	0.0112	0.0110	0.0112	0.0110
Households	154,328	178,886	154,328	178,886	154,328	178,886
Observations	8,513,972	9,620,712	8,513,972	9,620,712	8,513,972	9,620,712

Note: The dependent variable is log of monthly water consumption divided by the number of days on the bill. The coefficients can be interpreted as the percentage change in consumption. The preferred sample requires at least 1 year of pre-treatment data. Controls for all regressions include household FEs, month-by-year FEs, Average Temperature, Precipitation in Inches, and a cubic function of account enrollment duration. Robust standard errors clustered at the meter reading route level are reported in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Columns 5 and 6 of Table B.1 group the effects based on duration of 1 year or less, 1-3 years, or greater than 3 years. For accounts enrolled in ABP less than 1 year, there is a significant decrease in consumption of more than 1%. For accounts enrolled 1-3 years the effect is positive but not significant. However, for accounts enrolled longer than three years the effect is over 3%. My results are similar to findings reported by Sexton (2015), who finds that there is a negative and insignificant effect on electricity consumption during the first 4 months of enrollment in ABP, and no positive and significant effect until at least 8 months of enrollment. The pattern of results is different for PL, where the effect is a significant 2% decrease for accounts enrolled in PL less than 1 year, the effect is not significant for accounts on PL 1-3 years, and the effect is an almost 2% increase in consumption for accounts enrolled at least 3 years. Overall these results show that the ABP and PL effects are being driven by households who are enrolled for longer periods of time—especially more than three years where consumption increases by at least 5% years for both programs combined. These duration results are consistent with consumer inattentiveness driving the increase in consumption after enrollment. It seems reasonable to expect that consumers would gradually become more inattentive to the cost of water consumption over time.

B.2 Additional Tables and Figures

Table B.2: Baseline Results: 2003-2017

	(1)	(2)	(3)	(4)
ABP	0.0310***		0.0303***	0.0296***
	(0.0030)		(0.0029)	(0.0031)
PL		0.0126***	0.0068*	0.0044
		(0.0038)	(0.0038)	(0.0041)
Enrolled ABP and PL				0.0058
				(0.0061)
Weather Controls	Yes	Yes	Yes	Yes
f(Act Duration)	Yes	Yes	Yes	Yes
HH FE's	Yes	Yes	Yes	Yes
Period FE's	Yes	Yes	Yes	Yes
Within R-squared	0.0109	0.0109	0.0109	0.0109
Households	178,886	178,886	178,886	178,886
Observations	9,620,712	9,620,712	9,620,712	9,620,712

(a) Summer Months

	(1)	(2)	(3)	(4)
ABP	0.0353***		0.0348***	0.0331***
	(0.0035)		(0.0034)	(0.0036)
PL		0.0115**	0.0049	-0.0014
		(0.0052)	(0.0051)	(0.0053)
Enrolled ABP and PL				0.0151*
				(0.0081)
Weather Controls	Yes	Yes	Yes	Yes
f(Act Duration)	Yes	Yes	Yes	Yes
HH FE's	Yes	Yes	Yes	Yes
Period FE's	Yes	Yes	Yes	Yes
Within R-squared	0.0147	0.0146	0.0147	0.0147
Households	174,459	174,459	174,459	174,459
Observations	3,990,335	3,990,335	3,990,335	3,990,335

(b) Summer Months

Note: The dependent variable is log of monthly water consumption divided by the number of days on the bill. The coefficients can be interpreted as the percentage change in consumption. Controls for all regressions include household FEs, month-by-year FEs, Average Temperature, Precipitation in Inches, and a cubic function of account enrollment duration. Robust standard errors clustered at the meter reading route level are reported in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.3: Heterogeneity Results: 2003-2017

	Appraised Value		Yard Size	
	(1) Preferred Sample	(2) Full Sample	(3) Preferred Sample	(4) Full Sample
Quintile 1*ABP	0.0112***	0.0104***	0.0052***	0.0060***
	(0.0032)	(0.0020)	(0.0018)	(0.0014)
Quintile 2*ABP	0.0077***	0.0086***	0.0050*	0.0063***
	(0.0026)	(0.0023)	(0.0028)	(0.0022)
Quintile 3*ABP	0.0098***	0.0093***	0.0115***	0.0109***
	(0.0027)	(0.0025)	(0.0021)	(0.0018)
Quintile 4*ABP	0.0070***	0.0072***	0.0097***	0.0081***
	(0.0025)	(0.0022)	(0.0031)	(0.0026)
Quintile 5*ABP	0.0063**	0.0101***	0.0161***	0.0187***
	(0.0030)	(0.0033)	(0.0040)	(0.0036)
Quintile 1*PL	0.0078**	0.0078**	0.0041	0.0039
	(0.0037)	(0.0034)	(0.0034)	(0.0027)
Quintile 2*PL	0.0013	0.0017	0.0033	0.0042
	(0.0033)	(0.0026)	(0.0032)	(0.0029)
Quintile 3*PL	-0.0048	-0.0025	0.0046	0.0040
	(0.0032)	(0.0032)	(0.0041)	(0.0032)
Quintile 4*PL	-0.0024	-0.0006	0.0021	0.0037
	(0.0032)	(0.0028)	(0.0034)	(0.0030)
Quintile 5*PL	0.0081**	0.0051	-0.0047	-0.0036
	(0.0035)	(0.0038)	(0.0057)	(0.0042)
Weather Controls	Yes	Yes	Yes	Yes
f(Act Duration)	Yes	Yes	Yes	Yes
HH FE's	Yes	Yes	Yes	Yes
Period FE's	Yes	Yes	Yes	Yes
Within R-squared	0.0067	0.0065	0.0071	0.0068
Households	154,287	178,832	154,295	178,848
Observations	8,512,734	9,619,047	8,513,104	9,619,776

Note: The dependent variable is log of monthly water consumption divided by the number of days on the bill. The coefficients can be interpreted as the percentage change in consumption. The preferred sample requires at least 1 year of pre-treatment data. Controls for all regressions include household FEs, month-by-year FEs, Average Temperature, Precipitation in Inches, and a cubic function of account enrollment duration. Robust standard errors clustered at the meter reading route level are reported in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.4: Conservation Programs Results: 2003-2017

	(1) Preferred Sample	(2) Full Sample	(3) Preferred Sample	(4) Full Sample
ABP	0.0162***	0.0253***		
	(0.0041)	(0.0033)		
PL	0.0011	0.0031		
	(0.0053)	(0.0044)		
ABP, Conservation	0.0238***	0.0180***		
	(0.0072)	(0.0057)		
PL, Conservation	0.0140	0.0147*		
	(0.0101)	(0.0081)		
ABP, no PL			0.0158***	0.0267***
			(0.0045)	(0.0039)
PL, no ABP			-0.0031	-0.0025
			(0.0060)	(0.0055)
ABP, Enrolls Both			0.0180**	0.0249***
			(0.0089)	(0.0058)
PL, Enrolls Both			0.0057	0.0095
			(0.0092)	(0.0064)
ABP only, Conservation			0.0256***	0.0195***
			(0.0085)	(0.0068)
PL only, Conservation			0.0308**	0.0334***
			(0.0136)	(0.0123)
Both, Conservation			0.0318*	0.0195
			(0.0183)	(0.0120)
Both, Conservation			-0.0121	-0.0031
			(0.0186)	(0.0120)
Weather Controls	Yes	Yes	Yes	Yes
f(Act Duration)	Yes	Yes	No	No
HH FE's	Yes	Yes	Yes	Yes
Period FE's	Yes	Yes	Yes	Yes
Within R-squared	0.0112	0.0109	0.0111	0.0108
Households	154,328	178,886	154,328	178,886
Observations	8,513,972	9,620,712	8,513,972	9,620,712

Note: The dependent variable is log of monthly water consumption divided by the number of days on the bill. The coefficients can be interpreted as the percentage change in consumption. The preferred sample requires at least 1 year of pre-treatment data. Controls for all regressions include household FEs, month-by-year FEs, Average Temperature, Precipitation in Inches, and a cubic function of account enrollment duration. Robust standard errors clustered at the meter reading route level are reported in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure B.1: Event Study

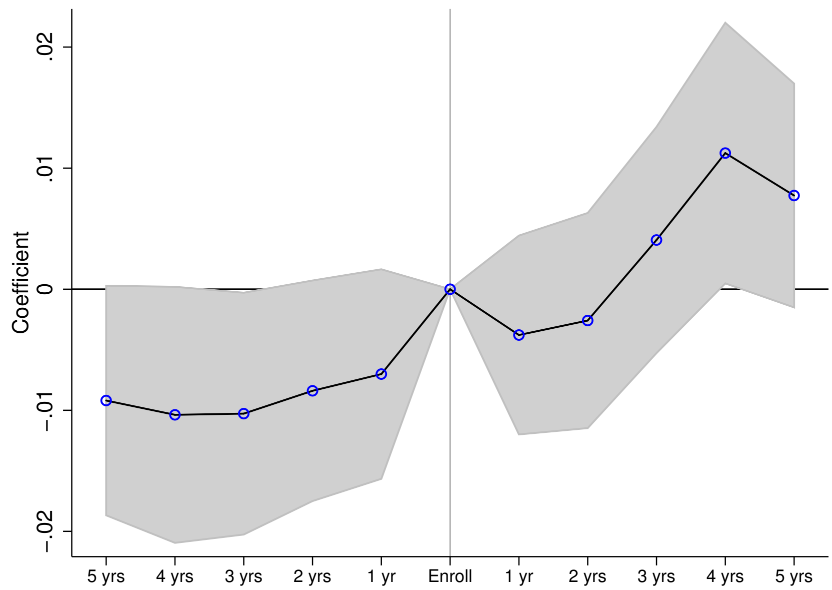

(a) Event Study ABP Enrollment

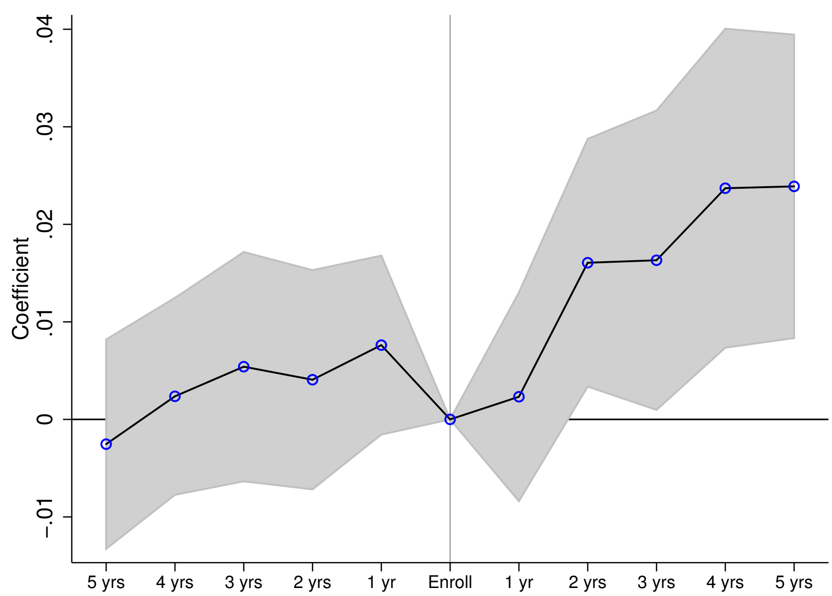

(b) Event Study PL Enrollment

Note: This event study plot uses the year of enrollment as the reference year. It demonstrates that the approach used in this chapter adequately controls for differences in pre-treatment trends and that there is a significant increase in consumption after enrollment. The sample is limited to accounts with at least 1 year of pre-treatment data. Robust standard errors are clustered at the meter reading route level.

Appendix C

Are Normative Appeals Moral Taxes? Evidence from a Field Experiment on Water Conservation

C.1 Randomization and Implementation

The randomization used a procedure of quasi-pairwise matching within blocking groups. This method first defines a set of blocks within which the randomization occurs. The blocking procedure ensures that assignment to a treatment is balanced within certain groups of interest. Blocks were defined by billing cycles, rate schedule and frequency of recorded meter data (i.e. monthly, daily, or hourly, though all customers only receive monthly usage totals). Within each block we ordered all observations on average water consumption in summer 2013 in sets of five households. We randomly assigned each household to one of five experimental samples that correspond to the five treatments (regardless of the ultimate assignment to treatment group vs. control group). This ensures a similar distribution of 2013 water consumption within each of the five experimental samples.

Next, within each of the five experimental samples we repeated the procedure to assign households to one of three possible timing treatments (single letter in July, single letter in August, or two letters repeated in July and August), or the control group. The same blocking structure was used within each experimental sample and then households were re-ordered based on summer 2013 water consumption and in sets of 12: two households are randomly assigned to each of the three timing treatments and six households are assigned to the control.[1]

The process for generating and mailing letters was as follows:

- 1-2 days after the most recent month's consumption data is loaded into the billing system we pull this information into Stata and using a set of pre-programmed routines use it to generate the graphics and data for the mail merge.

- A mail merge is performed in Microsoft Word using the generated data and graphics.

- PDF's of the letters are emailed to Digiprint

- Digiprint prints and ships the the letters within 1-3 days of receiving the electronic files.

The average time from the data upload to letters shipment was 2 days with a maximum of 8 days during this study. We had attrition during the study of about 1.5 percent of the treatment customers; 142 customers dropped out of the study in July and 211 customers dropped in August. This attrition was likely due to customers closing accounts or billing data (meter reading) errors. Furthermore, the mailers did not generate a very large increase in call center volume; out of the 23,213 customers we attempted to reach with this pilot we estimate that only 43 contacted the call center. Most of the customers who called the call center just wanted to ask clarification questions about the information in their letter; only 26 wanted further assistance beyond what the call center

[1] Due to the unequal size of the blocking groups, some timing treatments were oversampled, thereby creating some balance issues. We corrected these by identifying and dropping the oversampled observations after the conclusion of the field experiment (3,677 households: 2,025 control, 1,652 treatment). All balance tables and regression results reflect the corrected sample.

representatives could provide; and only five customers ended being truly upset by the pilot program.

C.2 Treatment Figures

Figure C.1: Treatment 2 - Historical water use information

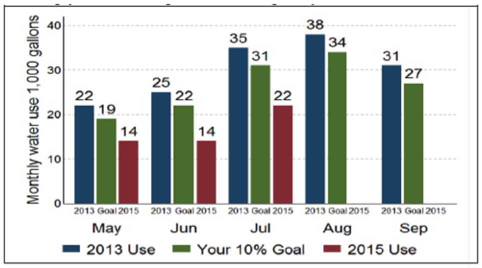

What is *your* 10% goal? The graph to the left shows your household's total water use in the past few months (red) compared to your water use in summer 2013 (blue). Water use on your July 2015 bill was **37% lower** than on your July 2013 bill. To do your part to help the community conserve water this summer, your total water use for August and September should be **at least 10%** lower than in 2013 (green). Your August goal amounts to saving **133 gallons per day** compared to how you used water in August 2013.

Figure C.2: Treatment 3 - Rate structure information

Saving water saves you money. The graph on the right shows how you **could save $13.00** if you reduce your water use on your August bill by 10% compared to 2013. The price you pay for each gallon of water increases as total water use increases. For most homes, indoor use falls within the first tier at the lowest price. Summer outdoor water use brings homes into the higher prices associated with tiers 2 and higher. The graph shows where your water use from your July 2015 bill (red) falls in the price schedule. It also shows your August 2013 water use (blue) and your 10% reduction goal for August 2015 (green).

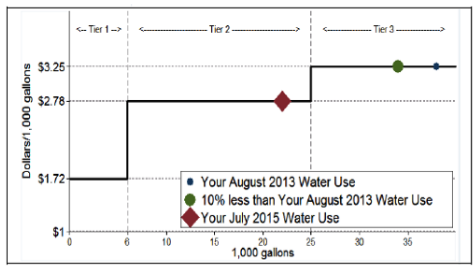

Figure C.3: Treatment 4 - Social comparison, reported in thousands of gallons

How does your water use compare?
The graph on the right shows your water use from your July bill compared to similar properties in your area. You used **1,000 gallons less** than your neighbors with similar properties.

You saved 35% on your July bill compared to 2013.
Keep up the good work!

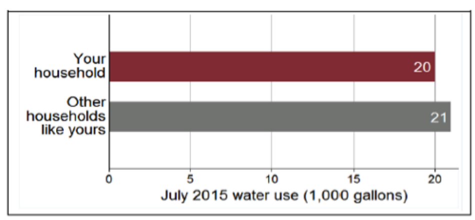

Figure C.4: Treatment 5 - Social comparison, reported as progress towards TMWA's 10% conservation goal

Are you doing your part? The graph on the right shows your change in water use form July 2013 to July 2015 compared to similar properties in your neighborhood. Your neighbors used 19% less water last month compared to 2013.

You saved 23% on your July water bill compared to 2013.
Keep up the good work!

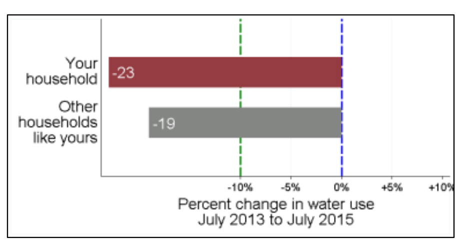

C.3 Additional Balance Tests

Table C.1: Balance on Pretreatment Consumption for each Letter

	Control Mean	Treatment Mean	Difference	(p-value)
Letter 1	21.69	21.81	-0.11	0.66
Letter 2	21.56	21.69	-0.12	0.63
Letter 3	21.63	21.66	-0.03	0.90
Letter 4	21.62	21.66	-0.04	0.86
Letter 5	21.65	21.70	-0.05	0.84

Note: p-values are based on two-sided t-tests

Table C.2: Balance on Pretreatment Consumption across Letters

	Mean 1	Mean 2	Difference	(p-value)
Letter 1 vs Letter 2	21.75	21.62	0.13	0.49
Letter 1 vs Letter 3	21.75	21.65	0.10	0.58
Letter 1 vs Letter 4	21.75	21.64	0.11	0.55
Letter 1 vs Letter 5	21.75	21.67	0.08	0.68
Letter 2 vs Letter 3	21.62	21.65	-0.02	0.89
Letter 2 vs Letter 4	21.62	21.64	-0.02	0.93
Letter 2 vs Letter 5	21.62	21.67	-0.05	0.78
Letter 3 vs Letter 4	21.65	21.64	0.01	0.96
Letter 3 vs Letter 5	21.65	21.67	-0.03	0.89
Letter 1 vs Letter 5	21.64	21.67	-0.03	0.85

Note: p-values are based on two-sided t-tests

Table C.3: Balance within Deciles of Pretreatment Consumption

	Control Mean	Treatment Mean	Difference	(p-value)
Decile 1	7.96	7.91	0.05	0.20
Decile 2	11.06	11.06	-0.00	0.97
Decile 3	13.50	13.45	0.04	0.02
Decile 4	15.70	15.68	0.02	0.32
Decile 5	17.95	17.94	0.01	0.59
Decile 6	20.36	20.38	-0.02	0.41
Decile 7	23.25	23.23	0.02	0.38
Decile 8	26.81	26.83	-0.02	0.62
Decile 9	32.39	32.32	0.07	0.27
Decile 10	48.24	48.13	0.11	0.77

Note: p-values are based on two-sided t-tests

Figure C.5: Distributions of Pretreatment Consumption Across Treatment Status

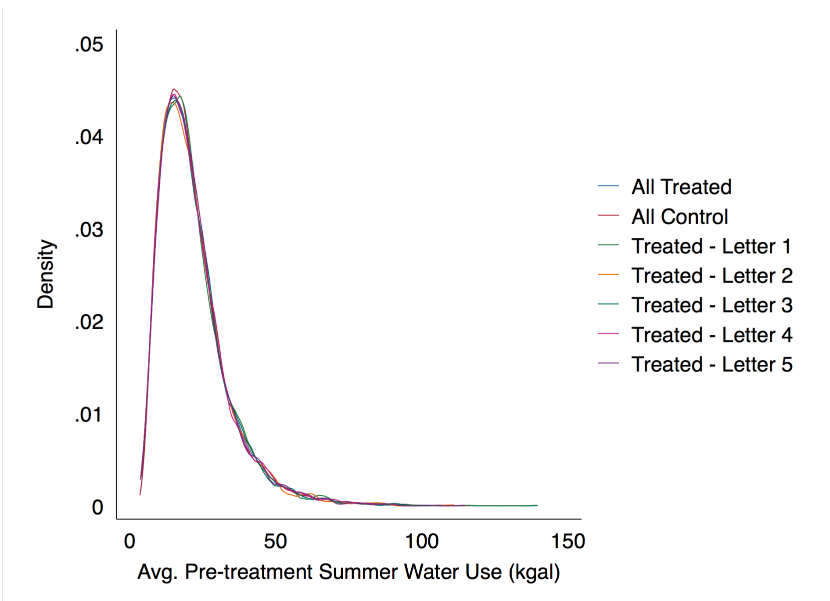

Note: The lines are kernel density estimates of pre-treatment consumption for all treated households, all control households, and households within each of the five treatment groups.

Table C.4: Kolmogorov-Smirnov Tests

	D-statistic	(p-value)
All Treatment vs Control	0.01	0.88
Letter 1 vs Control	0.01	0.89
Letter 2 vs Control	0.01	0.87
Letter 3 vs Control	0.01	0.86
Letter 4 vs Control	0.01	0.99
Letter 5 vs Control	0.01	0.76

Note: p-values are based on the combined D-statistic

C.4 Survey Information

Table C.5: Summary Statistics for Survey Variables

	N	Mean	Standard Deviation
Lawn Brown	1949	0.08	0.28
Not Enough Water	1949	0.38	0.49
Environment/Community	1949	0.85	0.35
Use How Want	1949	0.26	0.44
Save Money	1949	0.66	0.47
Already Conserved	1949	0.38	0.49
Invest Water Efficiency	1949	0.22	0.41
Behavioral Change	1949	0.88	0.33
Invest Traditional Landscape	1949	0.31	0.46
Already Invested Efficiency	1949	0.50	0.50

Note: All variables are binary indicators.

Table C.6: Balance on Pretreatment Consumption across Letters within Survey Sample

	Control Mean	N	Treatment Mean	N	Difference	p-value
Rate (T3) vs.						
Social Comparison kgal (T4)	21.26	361	22.36	361	-1.09	0.18
Rate (T3) vs.						
Social Comparison % (T5)	21.26	361	21.23	361	0.03	0.97
Social Comparison kgal (T4) vs.						
Social Comparison % (T5)	22.36	426	21.23	426	1.13	0.18
Rate (T3) vs.						
Social Comparisons (T4 and T5)	21.26	361	21.82	361	-0.56	0.46

Note: Survey Sample, p-value is based on two-sided t-test

Table C.7: Balance on Pretreatment Consumption for Survey Respondents vs. Non-Respondents

	Non-Survey Mean	Survey Mean	Difference	p-value
2013 Water	23.58	22.86	0.72	0.05
2015 Water	17.00	15.48	1.53	0.00
Summer Water	21.71	20.66	1.05	0.00
Winter Water	8.15	7.47	0.68	0.00
Year Built	1,987.68	1,986.20	1.47	0.00
Appraised Value	215.30	200.68	14.62	0.00
Bedrooms	3.37	3.32	0.05	0.02
Lot Acre	0.27	0.24	0.03	0.00
Yard Acre	0.23	0.20	0.03	0.00
Build Sq. Ft.	1,989.19	1,958.02	31.17	0.13
Bathrooms	2.19	2.17	0.02	0.18

Note: Non-survey sample of 41,110, Survey sample of 1,664; p-value is based on two-sided t-test.

Table C.8: Balance on Pretreatment Consumption for Treatment and Control within Survey Sample

	Control Mean	Treatment Mean	Difference	p-value
2013 Water	24.04	23.57	0.46	0.45
2015 Water	16.42	16.14	0.28	0.54
Summer Water	21.76	21.43	0.33	0.54
Winter Water	8.82	7.66	1.16	0.25
Year Built	1,989.42	1,988.34	1.08	0.19
Appraised Value	225.01	222.67	2.34	0.75
Bedrooms	3.38	3.31	0.07	0.05
Lot Acre	0.28	0.27	0.01	0.44
Yard Acre	0.23	0.23	0.01	0.45
Build Sq. Ft.	2,039.40	2,020.15	19.25	0.56
Bathrooms	2.20	2.20	0.00	0.99

Note: Survey Sample of 954 Controls and 994 Treatments; p-value is based on two-sided t-test.

C.5 Robustness for Difference from Peer Group

Table C.9: Difference from Peer Group - First Two Months for Single Mailers

	Norm in Gallons (T4)		Norm in % (T5)			
	(1)	(2)	(3)	(4)	(5)	(6)
	All	All	All	All	Low Users	High Users
Treat*Below Peer	-2.140***		-0.294			
	(0.817)		(0.971)			
Treat*Above Peer	-1.505		-3.151***			
	(1.177)		(1.179)			
Treat*Q1 Norm		-2.770**		0.907	0.694	1.562
		(1.109)		(1.490)	(1.050)	(2.793)
Treat*Q2 Norm		-1.801		-0.703	0.520	-2.030
		(1.125)		(1.134)	(1.200)	(1.828)
Treat*Q3 Norm		-2.174		-2.682**	-1.546	-3.758*
		(1.362)		(1.194)	(1.138)	(1.946)
Treat*Q4 Norm		-0.630		-4.715***	-3.263	-6.651**
		(1.898)		(1.804)	(2.165)	(2.768)
Observations	66,159	66,099	66,157	66,063	31,811	34,036

Note: The dependent variable is normalized average daily water consumption, and the sample is restricted to the first two months after a mailer for households that only receive one mailer. CATEs are estimated based on the difference (above/below) of household consumption relative to the peer group. Column headers All represents the whole sample, and Low/High Users restrict the sample to households above or below median pre-treatment consumption. All regressions include controls for temperature, precipitation, bill cycle fixed effects, month fixed effects, and pre-treatment consumption. Robust standard errors clustered at the household level are reported in parenthesis. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table C.10: Difference from Peer Group - First Month After Mailer

	Norm in Gallons (T4)		Norm in % (T5)			
	(1) All	(2) All	(3) All	(4) All	(5) Low Users	(6) High Users
Treat*Below Peer	-0.424		-1.248*			
	(0.561)		(0.710)			
Treat*Above Peer	-2.299***		-2.408***			
	(0.841)		(0.761)			
Treat*Q1 Norm		0.0149		-1.123	-0.177	-2.106
		(0.855)		(1.068)	(0.704)	(2.045)
Treat*Q2 Norm		-0.665		-0.978	-0.661	-1.079
		(0.712)		(0.881)	(0.720)	(1.488)
Treat*Q3 Norm		-1.431*		-1.672**	-0.902	-2.195*
		(0.861)		(0.800)	(0.888)	(1.241)
Treat*Q4 Norm		-3.253**		-3.420***	-1.554	-5.406**
		(1.446)		(1.246)	(1.243)	(2.186)
Observations	46,603	46,568	46,581	46,525	22,416	23,956

Note: The dependent variable is normalized average daily water consumption, and the sample is restricted to the first month after a mailer. CATEs are estimated based on the difference (above/below) of household consumption relative to the peer group. Column headers All represents the whole sample, and Low/High Users restrict the sample to households above or below median pre-treatment consumption. All regressions include controls for temperature, precipitation, bill cycle fixed effects, month fixed effects, and pre-treatment consumption. Robust standard errors clustered at the household level are reported in parenthesis. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

C.6 Effects of the community-wide 10% Goal

As a robustness check we analyze how the treatment effect varies based on whether households met their personal 10% savings goal. All three treatments (T3-T5) notify the household whether they met their personal 10% savings goal for the previous month. We estimate CATEs to based on how far away the household was from their 10% savings goal. The specification is based on equation 3.2 where the conditioning variables are indicators for the quartiles of difference between a household and its 10% goal in the month before treatment. One key distinction between quartiles based on the difference from the peer group and the difference from the personal goal is the interpretation within each quartile. The peer group is based on the median consumption or conservation rate so the quartiles

Table C.11: Conditional Average Treatment Effects: Quartiles of Difference from the 10% Goal

	(1) Rate (T3)	(2) Social Comp. kgal (T4)	(3) Social Comp. % (T5)
Treat*Q1 Goal	-0.737	-1.984	0.896
	(1.260)	(1.236)	(1.366)
Treat*Q2 Goal	-1.689*	-2.447***	-3.456***
	(0.888)	(0.909)	(0.975)
Treat*Q3 Goal	-4.465***	-1.762*	-3.087***
	(0.854)	(0.978)	(0.842)
Treat*Q4 Goal	-2.754*	-4.752***	-6.499***
	(1.425)	(1.045)	(1.238)
Observations	68,860	68,840	68,834

Note: The dependent variable is normalized average daily water consumption, and the sample is restricted post treatment data dropping all observations after the second mailer. The CATEs are based on quartiles of the difference from the 10% goal. All regressions include controls for temperature, precipitation, bill cycle fixed effects, month fixed effects, and pre-treatment consumption. Robust standard errors clustered at the household level are reported in parenthesis. *** p<0.01, ** p<0.05, * p<0.1

are centered at zero. However, the median households conserved by *more* than 10% so the third quartile in the difference from the goal contains some households than saved more than 10% and some households that saved less than 10%. We examine the discrete effect of moving just below the 10% goal with a regression discontinuity below. For this analysis the conservation rate for households within: *Q1 Goal* $\gg$ 10%, *Q2 Goal* $>$ 10%, *Q3 Goal* $\lesseqgtr$ 10%, *Q4 Goal* $<$ 10%. Similar to the analysis on the difference from the peer group, we assign control households values for the conditioning variables even though they did not receive a letter, and we drop all observations following a household's second mailer. Whether a household met the goal is balanced across treatment and control since it is based on pre-treatment consumption.[2] Table C.11 shows that there is heterogeneity based on the 10% goal, but the patterns are different than the CATEs based on the strength of the social comparison message. For example, in the rate information

[2]The p-values based on the two-sided t-tests for differences in the July goal and August goals across treatment status are 0.54 and 0.35 respectively.

treatment in column (1) three out of the four quartiles are negative and statistically significant with the largest effects in the third quartile. The pattern most closely mirrors the CATEs for the difference from the peer group for the social comparison in percentage terms (3). This is because the difference from the peer group is highly correlated with the difference from the 10% when the comparison metric is the percentage savings.

C.7 Discrete Effects Moving Above the Peer Group and 10% Goal

It is important to distinguish the continuous difference from peer group consumption from the discrete injunctive norm defining appropriate behavior (Schultz et al., 2007). We consider two separate effects. First, we test whether performing slightly worse than one's peers has an effect on consumption, and second we test for the discrete effect of just barely missing the 10% conservation goal. In our setting the descriptive injunctive norm is based on whether a household met the 10% goal. Therefore if a household less than 10% it received the message, "Please do your part to help with the drought.", while a household that saved more than 10% was told, "Keep up the good work!". It is also possible that the household considered their performance relative to their peer group as an additional categorical norm. The results in Table 3.5 contain *both* the effect of the discrete injunctive norm and the continuous descriptive norm.

To isolate the effect of the injunctive norm, we employ a regression discontinuity (RD) design (Imbens and Lemieux, 2008; Lee and Lemieux, 2010), analyzing behavior on either side of the injunctive category similar to Allcott (2011).[3] In the RD analysis we restrict

[3] Allcott (2011) showed little impact of moving into one of the three distinct categories in the Home Energy Report ("Great", "Good", or "Below Average") in a regression discontinuity design. In that study a household is assigned the category "Great" if they consume below the 20th percentile of peer consumption, "Good" if they consume below the average of peer consumption, or "Below Average" if

Figure C.6: Effect of Moving Above the Peer Group

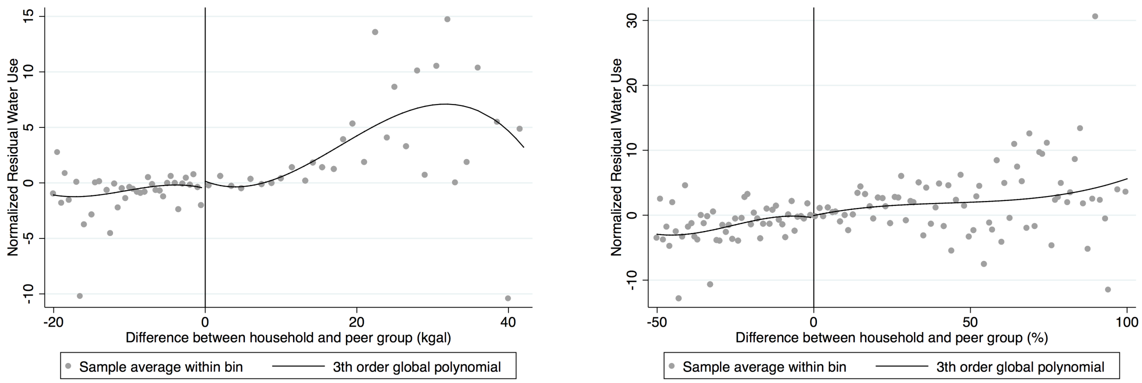

(a) Norm in kgal (T4)

(b) Norm in percentage reduction (T5)

Note: The dependent variable is residual normalized consumption and the units are percentage terms. The discontinuity is based on the moving above the peer group's conservation rate.

the sample to households treated with a social comparison (T4 and T5) and examine the effect of being just above the peer group. The running variable is the difference in performance between a household and the peer group in gallons for Treatment 4 and in percentage reduction for Treatment 5. The dependent variable in both is residual normalized consumption based on a regression of normalized consumption on weather, month fixed effects, and household fixed effects, following the approach of Allcott (2011). The RD estimation assumes that factors varying with the difference from the peer group, such as pre-treatment water and the strength of the descriptive norm, are the same for households just above and below their peer group. Since some households were above their peer group but saved more than 10%. Similarly some households were below their peer and saved less than 10%. We repeat the analysis dropping these households and the results are very similar.

We begin with graphical evidence of differences in consumption near the peer group

they consume above the average of peer consumption.

Table C.12: Regression Discontinuity Estimates of the Moving Above the Peer Group

	(1) Kgal (T4)	(2) % (T5)
Conventional	1.326	-0.159
	(0.897)	(1.037)
Bias-corrected	1.329	-0.212
	(0.897)	(1.037)
Robust	1.329	-0.212
	(1.063)	(1.234)
Observations	3,842	2,869

Note: The rows are three separate RD estimators and the appropriate standard errors according to Calonico, Catteneo and Titiunik (2014). The dependent variable is residual normalized consumption. The discontinuity in column (1) is based on the moving above the peer group's consumption (T4), and in column (2) is based on the moving above the peer group's year-on-year change in consumption (4).

(Figure C.6), as is standard in RD approaches.[4] For both social comparisons the graphical evidence in Figure C.6 suggests that moving above the peer group does not affect consumption.

The graphical evidence is corroborated in the RD estimates for moving above the peer group. We use three different RD estimators developed by Calonico, Catteneo and Titiunik (2014): the conventional, bias-corrected, and bias corrected with robust standard errors. In all specifications the impact for moving above the peer group is small and not statistically significant (Table C.12). The RD estimates show that the effects in Table 3.5 are driven by the *distance* from the peer group as opposed to simply being above or below the peer group. There is either no effect associated with adding a negative injunctive norm.

We repeat this exercise to see if moving slightly below above the 10% goal influences consumption. The analysis is the same as reported above except the running variable is

[4]The graphs are generated with a data-driven approach using spacing estimators to generate the bin sizes in the plots (Calonico, Catteneo and Titiunik, 2015). The points on the graph are the average normalized residual consumption within each bin, and the lines are the fitted values of separate third-order polynomial regressions on either side of the distance threshold (zero).

Figure C.7: Effect of Moving Above the Peer Group

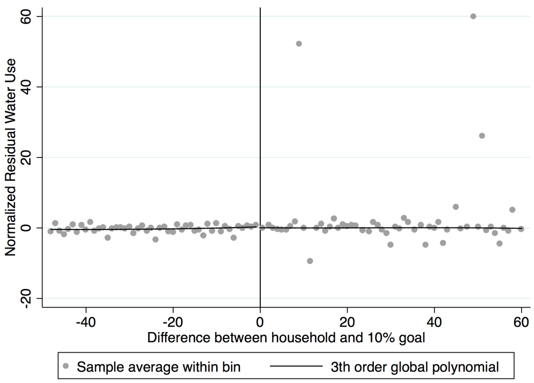

(a) Rate (T3)

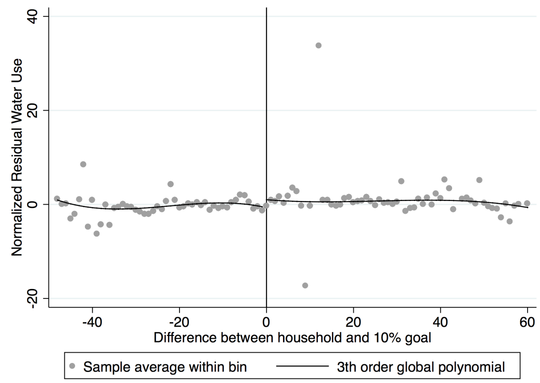

(b) Social Comparison kgal (T4)

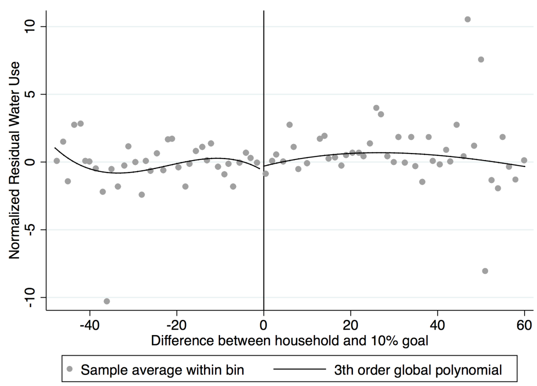

(c) Social Comparison (T5)

The dependent variable is residual normalized consumption and the units are percentage terms. The discontinuity is based on the moving above the consumption threshold that constitutes the household's 10% goal. Data from all treatments are included.

the year-on-year percentage change in water consumption and the threshold is the -10%. To be consistent with the analysis in the main text we subtract 10% off the running variable and such that the threshold is at zero and year-on-year changes of less than 10% are positive and more than 10% are negative. Figure C.7 graphs residual consumption on the y-axis with the year-on-year percentage change in water consumption (minus 10%) on the x-axis. There is no visual evidence of a change in consumption right at the threshold.

This is corroborated with the RD estimates (Table C.13) for each of the treatments.

Table C.13: Regression Discontinuity Estimates of the Failing to Meet the 10% Goal

	(1) History (T2)	(2) Rate (T3)	(3) Social kgal (T5)	(4) Social % (T5)
Conventional	0.238	-0.308	1.022	-0.047
	(1.104)	(1.193)	(1.253)	(1.146)
Bias-corrected	0.338	-0.387	1.262	0.058
	(1.104)	(1.193)	(1.253)	(1.146)
Robust	0.338	-0.387	1.262	0.058
	(1.312)	(1.415)	(1.443)	(1.338)
Observations	3,472	3,568	3,409	3,269

Note: The rows are three separate RD estimators and the appropriate standard errors according to Calonico, Catteneo and Titiunik (2014). The dependent variable is residual normalized consumption. The discontinuity in column (1) is based on moving above the consumption threshold that constitutes the household's 10% goal. The columns show the pooled treatment and each of the individual treatments.

Bibliography

Abrevaya, Jason, Yu-Chin Hsu, and Robert Lieli. 2015. "Estimating Conditional Average Treatment Effects." *Journal of Business & Economic Statistics*, 33(4): 485–505.

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

Allcott, Hunt. 2015. "Site Selection Bias in Program Evaluation." *Quarterly Journal of Economics*, 130(3): 1117–1165.

Allcott, Hunt, and Cass R Sunstein. 2015. "Regulating internalities." *Journal of Policy Analysis and Management*, 34(3): 698–705.

Allcott, Hunt, and Dmitry Taubinsky. 2015. "Evaluating Behaviorally-Motivated Policy: Experimental Evidence from the Lightbulb Market." *American Economic Review*, 1–80.

Allcott, Hunt, and Judd B. Kessler. 2015. "The Welfare Effects of Nudges: A Case Study of Energy Use Social Comparisons." *NBER Working Paper Series*, 21671.

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): 3003–3037.

Allcott, Hunt, Sendhil Mullainathan, and Dmitry Taubinsky. 2014. "Energy policy with externalities and internalities." *Journal of Public Economics*, 112: 72–88.

Attari, Shahzeen Z. 2014. "Perceptions of water use." *Proceedings of the National Academy of Sciences*, 111(14): 5129–5134.

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

Becker, Gary S. 1965. "A Theory of the Allocation of Time." *Economic Journal*, 493–517.

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: 437–452.

Boampong, Richard. 2016. "The Effects of Automatic Bill Payment on Price Sensitivity and Electricity Consumption." *Manuscript.*

Bolsen, Toby, Paul J. Ferraro, and Juan Jose Miranda. 2014. "Are Voters More Likely to Contribute to Other Public Goods? Evidence from a Large-Scale Randomized Policy Experiment." *American Journal of Political Science*, 58(1): 17–30.

Borenstein, Severin. 2009. "To What Electricity Price Do Consumers Respond? Residential Demand Elasticity Under Increasing-Block Pricing." *Manuscript.*

Brandon, Alec, Paul Ferraro, John A. List, Robert D. Metcalfe, Michael Price, and Florian Rundhammer. 2017. "Do the Effects of Social Nudges Persist? Theory and Evidence from 38 Natural Field Experiments." *NBER Working Paper Series*, 23277.

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, Corey Lott, Michael Taylor, Joseph Cook, Kimberly Rollins, and Shawn Stoddard. 2016. "Are Normative Appeals Moral Taxes? Evidence from a Field Experiment on Water Conservation." *Manuscript.*

Brown, Jennifer, Tanjim Hossain, and John Morgan. 2010. "Shrouded Attributes and Information Suppression: Evidence from the Field." *Quarterly Journal of Economics*, 125(2): 859–876.

Busse, Meghan, Jorge Silva-Risso, and Florian Zettelmeyer. 2006. "$1,000 Cash Back: The Pass-Through of Auto Manufacturer Promotions." *American Economic Review*, 96(4): 1253–1270.

Calonico, Sebastian, Matias D. Catteneo, and Rocio Titiunik. 2014. "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs." *Econometrica*, 82(6): 2295–2326.

Calonico, Sebastian, Matias D. Catteneo, and Rocio Titiunik. 2015. "Optimal Data-Driven Regression Discontinuity Plots." *Journal of the American Statistical Association*, 110(512): 1753–1769.

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

Christman, Laine. 2016. "Presentation on the 2015 Summer Drought Campaign Survey results."

Costa, Dora L., and Matthew E. Kahn. 2013. "Energy Conservation Nudges and Environmentalist Ideology: Evidence From a Randomized Residential Electricity Field Experiment." *Journal of the European Economic Association*, 11(3): 680–702.

Crump, Richard K., V. Joseph Hotz, Guido W. Imbens, and Oscar A. Mitnik. 2008. "Nonparametric Tests for Treatment Effect Heterogeneity." *The Review of Economics and Statistics*, 90(3): 389–405.

DellaVigna, Stefano. 2009. "Psychology and Economics: Evidence from the Field." *Journal of Economic Literature*, 47(2): 315–372.

DellaVigna, Stefano, and Joshua M. Pollet. 2009. "Investor Inattention and Friday Earnings Announcements." *The Journal of Finance*, 64(2): 709–749.

Dolan, Paul, and Robert Metcalfe. 2015. "Neighbors, Knowledge, and Nuggets: Two Natural Field Experiments on the Role of Incentives on Energy Conservation."

Ferraro, Paul J., and Juan José Miranda. 2013. "Heterogeneous Treatment Effects and Mechanisms in Information-Based Environmental Policies: Evidence from a Large-Scale Field Experiment." *Resource and Energy Economics*, 35(3): 356–379.

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.

Ferraro, Paul J., Juan Jose 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*, 101(3): 318–322.

Finkelstein, Amy. 2009. "E-ztax: Tax Salience and Tax Rates." *The Quarterly Journal of Economics*, 124(3): 969–1010.

Fiserv. 2010. "2010 Billing Household Survey: Consumer Survey of Offline and Online Billing and Payment Practices." 255 Fiserv Drive, Brookfield, Wisconsin.

Fiserv. 2016. "Eighth Annual Billing Household Survey: Insights on Consumer Billing and Payment Speed, Security and Satisfaction." 255 Fiserv Drive, Brookfield, Wisconsin.

Gaudin, S. 2006. "Effect of price information on residential water demand." *Applied Economics*, 38(4): 383–393.

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.

Grembi, Veronica, Tommaso Nannicini, and Ugo Troiano. 2012. “Effect of price information on residential water demand.” *IZA Discussion Paper*, 6952.

Hossain, Tanjim, and John Morgan. 2006. “...Plus Shipping and Handling: Revenue (Non) Equivalence in Field Experiments on eBay.” *Advances in Economic Analysis and Policy*, 6(2): 1–27.

Imbens, Guido W., and Thomas Lemieux. 2008. “Regression discontinuity designs: A guide to practice.” *Journal of Econometrics*, 142(2): 615–635.

Ito, Koichiro. 2013. “How Do Consumers Respond to Nonlinear Pricing? Evidence from Household Water Demand.” *Manuscript.*

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

Ito, Koichiro, Takanori Ida, and Makoto Tanaka. 2015. “The Persistence of Moral Suasion and Economic Incentives: Field Experimental Evidence from Energy Demand.” *E2e Project Working Paper Series*, 017.

Ito, Koichiro, Takanori Ida, and Makoto Tanaka. Forthcoming. “Moral Suasion and Economic Incentives: Field Experimental Evidence from Energy Demand.” *American Economic Journal: Economic Policy.*

Jessoe, Katrina, and David Rapson. 2014. “Knowledge is (Less) Power: Experimental Evidence from Residential Energy Use.” *American Economic Review*, 104(4): 1417–1438.

Kahn, Matthew E., and Frank A Wolak. 2013. “Using Information to Improve the Effectiveness of Nonlinear Pricing: Evidence from a Field Experiment.”

Lee, David S., and Thomas Lemieux. 2010. “Regression Discontinuity Designs in Economics.” *Journal of Economic Literature*, 48(June): 281–355.

Levitt, Steven D., and John A. List. 2007. “What do laboratory experiments measuring social preferences reveal about the real world?” *The Journal of Economic Perspectives*, 21(2): 153–174.

Lott, Corey. 2017. “The Effect of Price Information on Consumer Behavior Under Nonlinear Tariffs: Evidence from a Water Utility Merger.” *Manuscript.*

Mastercard. 2006. “Revealing Attitudes on Recurring Payments.” Purchase, New York.

McRae, Shaun, and Robyn Meeks. 2016. "Price perception and electricity demand with nonlinear tariffs." *Manuscript.*

Pellerano, José A, Michael K Price, Steven L Puller, and Gonzalo E Sánchez. 2015. "Do Extrinsic Incentives Undermine Social Norms? Evidence from a Field Experiment in Energy Conservation." *Environmental and Resource Economics*, 1–16.

Renwick, Mary E., and Richard D. Green. 2000. "Do Residential Water Demand Side Management Policies Measure Up? An Analysis of Eight California Water Agencies." *Journal of Environmental Economics and Management*, 40: 37–55.

Rosenbaum, Paul R., and Donald B. Rubin. 1983. "The Central Role of the Propensity Score in Observational Studies for Causal Effects." *Biometrika*, 70(1): 41–55.

Rubin, Donald B. 1990. "Formal Models of Statistical Inference for Causal Effects." *Journal of Statistical Planning and Inference*, 25: 279–292.

Schultz, P Wesley, Jessica M Nolan, Robert B Cialdini, Noah J Goldstein, and Vladas Griskevicius. 2007. "The constructive, destructive, and reconstructive power of social norms." *Psychological Science*, 18(5): 429.

Sexton, Steven E. 2015. "Automatic Bill Payment and Salience Effects: Evidence from Electricity Consumption." *Review of Economics and Statistics*, 97(2): 229–241.

Staff, Truckee Meadows Water Authority. 2009. "2010-2030 Water Resource Plan." Truckee Meadows Water Authority.

Strong, Aaron, and Chris Goemans. 2014. "Quantity Uncertainty and Demand: The Case of Water Smart Reader Ownership." *B.E. Journal of Economic Analysis and Policy*, 14(3): 669–694.

Strong, Aaron, and Chris Goemans. 2015. "The Impact of Real-Time Quantity Information on Residential Water Demand." *Water Resources and Economics*, 10: 1–13.

Visa. 2006. "Visa Payment Panel Study: 2006 Payment Trends Summary."

Wichman, Casey J. 2014. "Perceived price in residential water demand: Evidence from a natural experiment." *Journal of Economic Behavior & Organization*, 1–16.

Wichman, Casey J. 2016. "Information Provision and Consumer Behavior: A Natural Experiment in Billing Frequency." *Resources for the Future Discussion Paper*, DP 15-35-REV.

Wooldridge, Jeffrey M. 2010. *Econometric Analysis of Cross Section and Panel Data.* . Second ed., Cambridge, Massachusetts:The MIT Press.