ABSTRACT

Title of dissertation: INFORMATION AND ENVIRONMENTAL POLICY

Casey John Wichman,
Doctor of Philosophy, 2015

Dissertation directed by: Professor Roberton C. Williams III
Department of Agricultural and Resource Economics

Professor Maureen Cropper
Department of Economics

Within this manuscript, I present three distinct essays linked by the commonality of how information is utilized in decision-making and its effect on environmental policy.

In the first essay, I evaluate the price to which consumers respond under complicated billing structures. I exploit a natural experiment to estimate a causal effect of price for residential water customers during the introduction of increasing block rates for a North Carolina utility. Perceived price is identified through a billing anomaly in which changes in marginal and average prices move in opposite directions. Empirical results contribute evidence that residential water customers respond to average price. Average price elasticity estimates vary from -0.43 to -1.14 across the distribution of consumption in triple-difference models, with an estimate of -0.31 in the tightest bandwidth of regression discontinuity specifications.

In the second essay, I examine a causal effect of billing frequency on consumer behavior. I exploit a natural experiment in which residential water customers transitioned exogenously from bi-monthly to monthly billing. I find that customers increase consumption by approximately five percent in response to more frequent information. This result is reconciled in a model of price and quantity uncertainty, where increases in billing frequency reduce the distortion in consumers' perceptions. Using treatment effects as sufficient statistics, I calculate gains in consumer surplus equivalent to 0.5–1 percent of annual water expenditures. Heterogeneous treatment effects suggest increases in outdoor water use.

And, in the final essay, I consider the role of heterogeneous green preferences for private provision of environmental public goods in an asymmetric information context. Under varying degrees of information available to a regulator, I characterize equilibrium properties of several mechanisms. I find incentive compatible Nash equilibria that provide socially optimal public goods provision when the regulator can enforce individual consumption contracts, as well as when reported consumption contracts are supplemented with group penalties. The role of budget balancing is recast as a policy intervention for correcting environmental market failures.

INFORMATION AND ENVIRONMENTAL POLICY

by

Casey John Wichman

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2015

Advisory Committee:
Professor Roberton C. Williams III, Co-chair
Professor Maureen Cropper, Co-chair
Professor Richard E. Just
Professor Lint Barrage
Professor Ginger Zhe Jin

Dedication

To my parents, who supported me unconditionally through times when they, and more importantly I, did not understand what I was doing.

Acknowledgments

This dissertation is a product of the insight, advice, friendship, and tolerance of many.

I thank Rob Williams for his unfailing insight and guidance while allowing me the perfect degree of autonomy throughout my graduate work. Maureen Cropper, for her constant encouragement, acute advice, and absolute support. Richard Just, for his encouragement to progress through my doctoral work efficiently and his always constructive criticism. Lint Barrage, for forcing me to think about the broader implications of my work, for her willingness to entertain all of my wildest research ideas, and for reading the introduction to my job market paper too many times to count.

I am grateful for the advice, criticism, and encouragement of Lars Olson, Anna Alberini, Lesley Turner, Ginger Jin, Bob Chambers, Erik Lichtenberg and for their evaluation of my scholarship since the very beginning of my doctoral work. I am indebted to Laura Taylor and Roger von Haefen for igniting my passion for economic research and for their continued mentorship. And, Chris Timmins, for his willingness to serve as my de facto job market adviser.

My friends and fellow classmates—including Jonathan Eyer, Evan Rogers, Timothy Hamilton, Jane Zhou, Charlene Chi-Johnston, Jikun Wang, Jen He, Yuandong Qi, Min Kim, and countless others—contributed to the quality and rigor of my education, as well as the ability to tolerate the stressfulness of graduate school.

I thank those who enabled and cooperated with my data requests at the Orange Water and Sewer Authority and the City of Durham's Department of Water Management, and for allowing me to use their data in my research. And, Mary Tiger, Shadi Eskaf, Christine Boyle, Glenn Barnes, and Jeff Hughes at the Environmental Finance Center at the University of North Carolina at Chapel Hill for the opportunity to learn about water utility management, and so much more.

My family's love and support was unparalleled in preparing me for the pursuit of a career in inquisitiveness, and it continues to be.

Huckleberry Jefferson and Chloe Moses, the ultimate distractions, served as incredible reminders that leisure is important. Although, they already knew that.

And, I would not be who or where I am today without Ashley Chaifetz. Her tolerance, love, intelligence, patience, thoughtfulness, discipline, stubbornness, and selflessness is an unending inspiration. I am boundlessly fortunate to have her in my life.

Table of Contents

List of Tables

List of Figures

Chapter 1: Introduction

The role of information in economic policy has evolved drastically in the 50 years since Stigler's seminal article on "the economics of information" (Stigler, 1961). And yet, in the context of consumer behavior, many interesting questions remain unanswered on the value of information, how it is used in decision-making, and precisely how it is ignored. Questions of particular interest encompass the degree to which consumers utilize price information in decision-making, whether individuals confronted with changes in the frequency of information received behave as economic models predict, and the value of private information in a regulatory context. While these questions are inherently general, the focus of this research is to address the role of information with applications to environmental and resource economics.

Within this manuscript, I present three distinct essays linked by the commonality of how information is utilized in decision-making and its effect on environmental policy. Specifically, I examine several research questions that explore the use of various forms of information in environmental policy. First, in the context of residential water demand, I analyze the effect of changes in tiered rate structures to ascertain what information customers use to make consumptive decisions. Next, I examine whether consumers billed intermittently for water respond to the frequency at which they receive billing information and the corresponding welfare implications. Lastly, I explore the provision of environmental public goods when consumers display different preferences for the environment.

In the first essay, "Perceived price in residential water demand: Evidence from a natural experiment," I explore the price signal that consumers use to make consumptive decisions. I motivate a simple model in which consumers update this period's consumption based on last period's water bill. This framework aligns itself with a natural experiment to determine empirically whether consumers respond to marginal or average price when facing tiered rate structures. I exploit the introduction of a new rate structure in which changes in marginal and average prices move in opposite directions to estimate a causal effect the rate structure change on demand. The estimated demand response is consistent with the hypothesis that consumers use average price to make consumption decisions.

The second essay, "Information provision and consumer behavior: A natural experiment in billing frequency," explores whether consumers respond to the rate at which they receive a bill for episodic billing for residential water. I posit that the price signal consumers used to make consumptive decisions is distorted if consumers do not pay for an economic good at the time of its consumption. A conditionally random transition from bi-monthly to monthly billing for billing districts within the same utility allows for identification of a causal effect of billing frequency on consumer demand. Empirical results provide strong evidence that consumers increase consumption in response to more frequent billing. Under the notion that more frequent information reduces the distortion in consumer perceptions, I develop a sufficient statistic framework for calculating the welfare effects of changes in billing frequency. Measures of consumer surplus suggest a welfare gain equivalent to 0.5-1% percent of annual expenditures on water.

In the final essay, "Incentives, green preferences, and private provision of environmental public goods," I develop a conceptual framework for understanding the role of pro-environmental preferences for privately provided public goods in an asymmetric information context. In this research, I examine equilibrium properties of several incentive schemes under progressively weaker informational constraints on an environmental regulator. The primary results indicate that socially optimal provision of the public good can be attained when the regulator can construct incentives based on (1) individual consumption and (2) reported consumption paired with observable group output.

Chapter 2: Perceived price in residential water demand: Evidence from a natural experiment

Reprinted from the *Journal of Economic Behavior & Organization*, 107, Casey J. Wichman, Perceived price in residential water demand: Evidence from a natural experiment, 308–323, Copyright (2014), with permission from Elsevier.

2.1 Introduction

Increasing block rate structures are typically adopted by water utility managers to promote conservation among high-users and affordability for low-to-moderate-users. With growing concern for water scarcity, as well as the need for revenue stability at the utility-level, the introduction of block rates for residential water customers is becoming increasingly common. However, upward pressure on the costs of providing water to households necessitates a better understanding of how consumers respond to unclear price signals. This knowledge will help utility managers craft rate structures commensurate with their goals and provide researchers with a framework for studying water demand that better conforms to observed consumer behavior. Specifically, price elasticity of residential water demand is the key parameter of interest because price is used as an instrument of conservation during periods of acute scarcity and utilities are often constrained by zero-profit mandates such that the impact of changes in prices is important for revenue planning (Olmstead et al., 2007).

In this paper, I motivate a conceptual framework for residential water demand that relies on a customer's consumption patterns last month as a heuristic for this month's prices. This identification strategy expands upon previous research by exploiting the assignment of billing cycles for residential customers of Orange Water and Sewer Authority (OWASA) in Chapel Hill, North Carolina during the introduction of increasing block rates. Conceptually, this treatment allows for identification of perceived price for one subset of households relative to nearly identical households on a different billing cycle who face similar weather patterns and utility-specific shocks but may be responding to different price information.

Most of the literature on the price elasticity of water demand suffers from the non-experimental nature of utility pricing, relying on cross-sectional variation across households or municipalities, or variation over time. The former provides a potential avenue for omitted variables to bias results and the latter fails to exploit exogenous, unanticipated changes to the pricing structure for a given household (Klaiber et al., 2014; Olmstead et al., 2007). Further, the co-movement of marginal and average price tend to confound results that examine price perception. In this analysis, three complementary quasi-experimental methods are applied to explore the behavioral response to a change in the rate structure in an attempt to isolate a causal estimate of price elasticity for water customers. Difference-in-difference (DD) methods are applied to capture an average treatment effect for all water customers based on the notion that consumers respond to last month's bill as a proxy for this month's price. Triple-difference (DDD) estimates allow for the identification of heterogeneous price effects by segregating the sample into decile groups that correspond to a household's average historical consumption. Finally, the transition to a block rate structure allows for the exploitation of a discontinuity in price in which changes in marginal and average prices move in opposite directions for a portion of the water use distribution. Fuzzy regression discontinuity (FRD) models that focus on this anomaly reinforce the DDD results and allow for testing whether the identification strategy is plausible in light of potential confounding factors. The complementarity of this suite of quasi-experimental techniques provides strong empirical evidence that residential water consumers respond to average price.

The results of empirical models indicate that the short-run response to the adoption of increasing block rates is a net increase in consumption due to lower prices in the first and second consumption blocks. Further, DDD estimation results lend evidence that customers across most deciles of consumption respond to average price, while customers in both tails of the distribution exhibit no significant response to price. Average price elasticity estimates range from -0.43, for consumption around 3,000 gallons per month, to -1.14, for consumption around 7,000 gallons per month, which are within the range of previous studies (Dalhuisen et al., 2003; Espey et al., 1997). Elasticity estimates calculated with changes in marginal price display either implausibly large magnitudes, wrong signs, or statistical insignificance. Lastly, by exploiting the divergence in average and marginal prices at 7,000 gallons per month after the rate change, FRD results provide further evidence that residential water customers respond to average price.

2.2 Perceived price

Despite an extensive literature estimating the effect of price changes on customer demand under complicated rate structures, the price signal that residential water customers use to make consumption decisions remains unclear to researchers (Nataraj and Hanemann, 2011). Standard economic theory stipulates that consumers should respond to the marginal price for the next unit of water consumed, as well as the marginal price in each block below the final block of consumption (Hewitt and Hanemann, 1995; Olmstead et al., 2007; Strong and Smith, 2010). However, it is plausible that consumers are not aware of the marginal price they face in each of the blocks (Nataraj and Hanemann, 2011) and consumers might not fully understand how their bill is calculated within the tiered rate structures (Nieswiadomy and Molina, 1989). Recent work in electricity demand suggests that consumption patterns better reflect a response commensurate with changes in either expected marginal price (Borenstein, 2009) or average price (Ito, 2014) under tiered rate structures, while results from a quasi-experiment in water demand imply that high-volume residential water customers seem to respond to marginal price (Nataraj and Hanemann, 2011).

2.2.1 Marginal, expected marginal, and average price

In a simple model of demand under block pricing without uncertainty, consider a consumer with quasi-linear utility:

$$u(\mathbf{w}, x) = V(\mathbf{w}) + x \tag{2.1}$$

where $\mathbf{w}$ is a vector of water consumption in each price block for that period and x is a linearly separable numeraire good with its price normalized to unity. Let $I = x + \mathbf{p}'\mathbf{w}$ represent the consumer's budget constraint with $\mathbf{p}$ being the price schedule for water consumption and wealth, I, is determined exogenously. The consumer maximizes her utility subject to the prices she faces in two blocks of water consumption:

$$\max_{\mathbf{w}} u(\mathbf{w}, x) = I + V(\mathbf{w}) - p_1 w_1 - p_2 w_2 \tag{2.2}$$

where w_1 and w_2 are the quantities demanded in each consumption block corresponding to marginal prices p_1 and p_2. The solution to the consumer's problem results in the following piece-wise demand function for water:

$$\mathbf{w}^\star = \begin{cases} w^\star(p_1) & \text{if } w^\star(p_1) \leq k \\ k & \text{if } w^\star(p_1) = k = w^\star(p_2) \\ w^\star(p_2) & \text{if } w^\star(p_2) \geq k \end{cases} \tag{2.3}$$

where k is the "kink point" in the consumer's budget constraint (Ito, 2014).[1] This model of water demand with non-linear budget constraints, formalized by Hewitt and Hanemann (1995), is typically estimated using discrete-continuous choice methods in which the consumer chooses her consumption block and the optimal amount to consume within that block simultaneously. In the discrete-continuous choice framework, marginal prices in each block are necessary pieces of information for the consumer to make her consumption decision. While this model conforms to utility theory, it makes the assumption that a consumer performs a complex utility-maximizing decision based on perfect information of the water provider's rate schedule and her precise level of consumption throughout the billing period (Borenstein, 2009).

Borenstein (2009) relaxes this assumption by introducing a model that allows for consumption decisions to be made in response to local marginal prices in the context of electricity demand. This conceptualization of consumer behavior provides a more intuitive model of demand under block rates while avoiding the restrictive assumptions in Hewitt and Hanemann's framework. In this model, consumers maximize expected utility,

$$\max_{\mathbf{w}} E[u(\mathbf{w}, x)] = I + E[V(\mathbf{w})] - E[p_1 w_1 + p_2 w_2], \tag{2.4}$$

[1] In a model of quasi-linear utility, there are no income effects on water consumption. The income effect, likely to be small for residential water consumption, is typically important in estimating price elasticity for tiered rate structures through intra-marginal rate changes affecting virtual income. See Olmstead et al. (2007) for sufficient treatment of the effect of virtual income on price elasticity in water demand under block rates. Following Borenstein (2009) and Ito (2014), this analysis assumes the effect of income is negligible and proceeds with quasi-linear utility as a plausible model of residential water demand.

where the first-order condition states that a consumer will choose consumption $\mathbf{w}^*$ that sets her marginal utility equal to the expected marginal price, which is a probability-weighted average of the marginal prices in each consumption block. This formulation allows for smooth demand functions, though it still requires complete knowledge and understanding of the utility's rate structure. Borenstein tests an empirical model in which customers set behavioral consumption rules at the start of the period, such as setting the thermostat to a fixed temperature, based on the marginal price they expect to face while allowing for exogenous demand shocks within the consumption period. He finds evidence that electricity customers are more likely to respond to expected marginal price than marginal price in Southern California, but it is possible that they are responding to even less precise information.

Further, several empirical papers have suggested that even expected marginal price places too large of a computational burden on customers who face block rates and that the total bill, or average price, is a more accurate representation of a customer's perceived price (Ito, 2014; Foster and Beattie, 1981; Liebman and Zeckhauser, 2004). In this framework, consumers optimize their consumption in an ad-hoc fashion such that the marginal utility of consumption is equal to the ex-post average price. Additionally, Ito (2014) and Shin (1985) formalize models in which the response to average price is motivated by the costs of obtaining the necessary price information to maximize welfare in the standard framework. Despite these conceptual advances, empirical evidence has not provided a conclusive answer to which price is perceived in residential water demand.

2.2.2 A heuristic approach

In this paper, the restrictive assumptions inherent in Hewitt and Hanemann's (1995) model, and subsequent structural representations of water demand, are relaxed by assuming that the average household does not actively seek out information from the water utility about rate changes nor do they habitually monitor their water use throughout the billing period. Conceptually, this model is plausible because the primary means of communication between the water customer and the utility is the periodic water bill. Thus, this framework implies that the customer does not respond to the price she is charged in the current time period, rather she updates this month's consumption based on her utility bill for water use in the previous month. In effect, this model is nested within Equation 2.4 with the following assumption on the customer's conditional expectation of price:

$$E[p_t|\mathbf{w}_{t-1}] = p_{t-1} + S_t + \mu_t \tag{2.5}$$

where p_t is the price a customer uses to make consumption decisions in period t, $\mathbf{w}_{t-1}$ is the customer's consumption last billing period, S_t allows for known differences in seasonal water usage, and μ_t allows for error in the customer's prediction of her price due to unexpected shocks to water use.

This model of price perception has several attractive features for residential water demand: 1) in the absence of a well-publicized rate change, customers are generally informed about changes to their rate structure through their water bill after the changes have taken place; 2) utility bills rarely include more information than the charges incurred by consumption in each block, so this model assumes imperfect information about the full price schedule and block endpoints; and 3) it is plausible that customers learn about their consumption habits ex-post and alter their behavior for the next billing period. Further, this framework is flexible enough to accommodate consumption behavior that responds to either marginal, expected marginal, or average price. In fact, using last month's as a heuristic for this month's consumption requires minimal information costs. Empirically, many researchers have modeled water demand with lagged price variables in an informal fashion to minimize the effect of endogeneity and avoid contemporaneous correlation with consumption (Arbués et al., 2003; Ito, 2014; Renwick and Archibald, 1998; Wichman et al., 2014).

2.3 Empirical strategy

To avoid common problems of simultaneously determined price variables, I approach the question of perceived price with a quasi-experimental model that produces partial price effects as a function of treatment assignment, rather than explicit inclusion of a price variable in the regression. Nataraj and Hanemann (2011) and Klaiber et al. (2014) are the only researchers to develop a quasi-experimental strategy to estimate a causal effect of price on water demand. Nataraj and Hanemann estimate a regression discontinuity model to exploit the introduction of an additional price block among customers just above and just below the new block cut-off. The treatment effect is interpreted as an elasticity estimate of -0.12 for a price increase of nearly 100% (Nataraj and Hanemann, 2011). But, this interpretation is confounded by the co-movement of average and marginal price. By performing back-of-the-envelope calculations with the price schedule and statistics reported in Nataraj and Hanemann's analysis, the increase in average price is roughly 10% for the difference between treatment and control groups. This price increase translates to an average price elasticity estimate of -1.16 which is a plausible response for high volume water customers with a large proportion of extraneous water use.[2] Klaiber et al. (2014) use seasonal variation in marginal prices to assess heterogeneous responses to price. However, they do not consider alternative measures of price perception. These studies make an important advancement in the literature on perceived price by employing quasi-experimental methods to elicit a causal effect of price while avoiding potential sources of bias.

2.3.1 Treatment assignment

The notion that water customers respond to last month's bill as a proxy for this month's prices allows for the analysis of a unique natural experiment. Since Chapel Hill water customers are segmented into one of three billing cycles, residential customers in different cycles receive their utility bill at different points each month.

[2]Using the difference between Nataraj and Hanemann's estimate of 43 ccf for pre-treatment water-use and 35 ccf for a control estimate within the tightest band of RD models (note: 1 ccf = 100 cubic feet = 748 gallons), average price is calculated at \$1.56 for the treatment group and \$1.42 for the control group. The percent change in consumption after the treatment effect is $-5.1\text{ccf}/43\text{ccf} = -0.119\%$. The average price elasticity was estimated by dividing the percent change of the treatment effect by the percent change in average price between treatment and control groups.

If customers respond only to the information provided in their monthly bill, it is possible that customers in adjacent billing cycles would respond to different price information if a change in the rate structure occurred between the bill dates for each cycle.[3]

To illustrate the assignment of treatment status, consider two identical households who do not actively seek out rate information from the utility—households A and B. Household A is billed at the end of the month and household B is billed in the beginning of the month. If a rate change occurs on the first day of the month, household A will not recognize that rates have changed until they receive their next month's bill. Household B, however, will receive information about the new rate structure from this month's bill and update consumption in the current period based on the new rates. Thus, two otherwise identical households will make decisions on water use during similar time periods based on different price information. The differential consumption behavior between these two households can be interpreted as the effect of the new price structure since both households faced similar weather conditions and the same exogenous demand shocks in the overlapping billing period.

This anecdote describes the assignment of treatment and control status in this analysis. I focus on two periods—the month before and the month after tiered block rates were introduced for Chapel Hill water customers. I restrict my analysis to two

[3]There are roughly 10 days between the bill dates for each of the billing cycles and subsequently anywhere from 1-10 days between the date on which the meter was read and the date on which the bill was generated, depending on the account location within the meter route. On average, sequential cycles have overlapping consumption of 20 days. The average bill length is designed to be between 27 and 30 days and remains consistent throughout the year, however the exact length varies due to the inability to read meters on weekends, among other constraints.

sequential billing periods for two billing cycle groups. Within the time frame of the study, the first cycle group received monthly water bills on August 29th, September 28th, and October 31st in 2007 which results in the number of days between bills at 30 and 33 respectively. The second cycle group received bills on September 9th, October 10th, and November 7th in 2007 resulting in the number of days between bills at 33 and 28 respectively.[4] The possible divergence in billing lengths between the two groups indicate that changes consumption could simply reflect the relatively shorter billing period for the latter group in the second period. However, this effect works against the main findings of the paper, thus making the results presented here conservative estimates of the true effect. Since each bill reflects water use in the preceding month, the first group's October 31st bill will be the first instance in which rates have been calculated under the new rate structure; hence, this group's October consumption will reflect the old rates despite the fact that they are being charged on the new rate structure. Alternatively, customers in the second cycle group will have received their October 10th bill and updated their consumption to reflect the new rates, thus consumption observed on the November 7th bill will be commensurate with the rate change. Households in the first group are designated as control households and households in the second group are designated as treatment households.

[4]While there is variation in the number of days between bills among different billing cycles, it is unlikely that this variation significantly affects the amount of billed consumption within a household's billing period. Bill dates represent the point at which bills are mailed to customers and thus subject to the utility's administrative work schedule. In fact, it is likely that this variation is driven by the inability to mail water bills on the weekend. Discussions with utility officials suggest that billing periods remain relatively uniform over the year in order to maintain consistency in billed amounts from period to period within customers.

2.3.2 Data

For this analysis, a panel of monthly billing data for all residential customers from January 2006 through December 2008 was obtained from OWASA. Each residential billing record contains quantities and charges incurred by water and sewer use. Meters are read monthly and consumption is rounded to the nearest thousand gallons for billing purposes.[5] OWASA assumes the amount of water and sewer usage is identical each month. The billing data are supplemented with temperature and evapotranspiration information from the North Carolina State Climate Office. Lastly, information about rates and outreach to customers regarding rate changes was obtained from OWASA. Summary statistics for consumption, prices, and weather are reported in Table 2.1. Only customers residing in single-family dwelling units who did not change premises within the time frame of the study, as indicated by the billing records, are analyzed. Customers with consumption billed through irrigation meters are not included in the data set.[6] University accounts, which are billed on a separate cycle, are also removed. In the final sample of customers, there are 10,249 household consumption observations in September 2007 and 10,435 in October 2007.

[5] Rounding to the nearest thousand gallons of water consumption results in classical measurement error. Thus, it is likely that there is attenuation present in parameter estimates.

[6] Water billed through an irrigation meter is charged on a separate rate schedule without sewer rates and comprises less than one percent of the billing observations in this sample, thus removing these observations does not exclude households who have in-ground irrigation systems or strong preferences for outdoor water use.

Table 2.1: Summary statistics

	Obs	Mean	Std. Dev.	Min	Max
Volume (1,000 gallons/month)	96,702	5.28	5.73	0	191.00
Total Bill ($/month)	96,702	61.73	60.47	9.84	2,425.38
Average Price ($/1,000 gallons)	96,702	13.77	4.85	3.08	48.45
Marginal Price ($/1,000 gallons)	96,702	8.11	2.28	1.98	17.21
Maximum Temperature (°F)	96,702	82.06	6.04	71.10	94.08
Evapotranspiration (in)	96,702	4.30	0.93	3.00	6.24

Notes: Summary statistics are for September and October household consumption from 2006, 2007, and 2008. Consumption and bill data was obtained from customer billing records. Price variables were calculated from billing records and utility rate sheets. Weather variables were obtained from the NC State Climate Office.

2.3.3 Overview of rate structure change

Increasing block rates were introduced for OWASA customers on October 1st, 2007. Prior to the rate change, residential customers paid a uniform price for both water and sewer usage. After the rate change, the same customers faced an increasing five-block rate structure. In April 2007, OWASA mailed a brochure about the introduction of block rates to customers. In addition, customers were encouraged to attend a public hearing on the new rate schedule with the Board of Directors in May 2007. These rates were introduced to help meet revenue needs as well as to encourage conservation among high-volume users while allowing water bills to remain affordable (Orange Water and Sewer Authority, 2007). Incidentally, below-average rainfall in the summer of 2007 resulted in severe drought conditions throughout the fall of 2007 and into the spring of 2008. To encourage conservation, OWASA implemented voluntary watering restrictions on September 27, 2007 and mandatory watering restrictions on October 18, 2007 which remained in effect until the spring of 2008. Due to the delayed nature of billing cycles, which is central to the identification of this analysis, the treatment group faced a longer exposure to drought restrictions than did the control group. This effect would pose significant concerns for the accuracy of this research if treatment households reduced consumption by more than the control households. However, results show that the treatment effect moves in the opposite direction of the drought restrictions. This finding alleviates concerns about confoundedness and allows for the treatment effects to be interpreted as lower bounds. In addition, FRD models are estimated with data from the treatment group only to examine the robustness of the results to this potential threat to identification.

The rates customers faced before and after the rate change are presented in Table 2.2 and Figure 1. The marginal price for water and sewer rates was $9.17 for all units of consumption prior to October 1st. After the rate change, customers paid $6.14 for the first 3,000 gallons, $8.86 for the next 3,000 gallons, $9.69 for consumption between 7,000 and 11,000 gallons inclusive, $11.62 for consumption between 11,000 and 16,000 gallons inclusive, and $17.21 for consumption beyond 16,000 gallons.[7] Average price is defined as the total volumetric charge for water and sewer use divided by the ex post level of consumption and marginal price is defined as its "local" counterpart in that it is assumed consumers only respond to the highest marginal price they face within the block rate structure. This formulation of average price represents a proxy for the total bill that a customer would observe

[7]Due to rounding of consumption to the nearest thousand gallons, the rate schedule is defined in regards to discrete blocks of one-thousand gallons per month. If a customer displayed 3,499 gallons of monthly water use, for example, she would be billed at the rate for 3,000 gallons.

after their consumption period. Descriptive statistics for both variables are also presented in Table 2.2.[8] A customer who consumed less than 7,000 gallons before and after the rate change saw a decrease in both the marginal and average price she faced in each time period. A customer who used between 7,000 and 14,000 gallons each period, however, faced an increase in marginal price but a decrease in average price. This anomaly occurred because the price paid for the first units of water consumed in the month after the rate change were charged at a lower rate than the uniform price prior to the rate change.

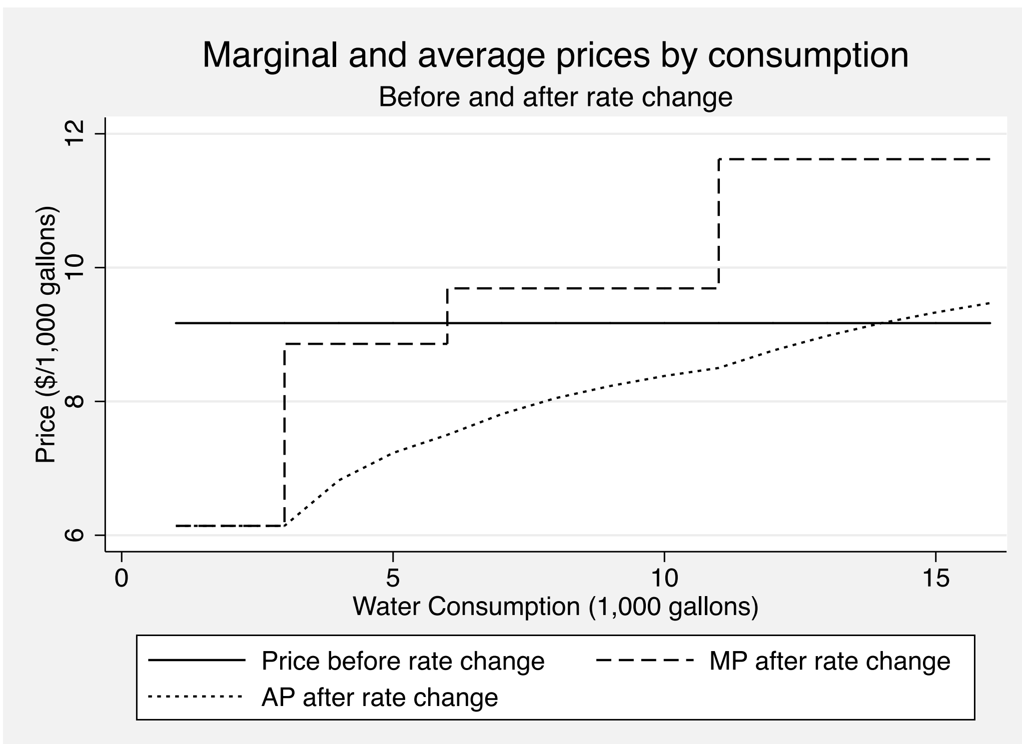

Figure 2.1: Marginal and average prices before and after price change

[8]Volumetric wastewater charges are included in the price variables in the analysis. While the rate structure for water changed from uniform rates to increasing block rates, the charges for wastewater remained uniform throughout the time frame of this study.

Table 2.2: Changes in marginal and average prices before and after rate change

	Before Rate Change		**After Rate Change**			
Volume (gallons)	Marginal Price	Average Price	Marginal Price	Average Price	%Δ MP	%Δ AP
1,000	\$9.17	\$9.17	\$6.14	\$6.14	-33.04%	-33.04%
2,000	\$9.17	\$9.17	\$6.14	\$6.14	-33.04%	-33.04%
3,000	\$9.17	\$9.17	\$6.14	\$6.14	-33.04%	-33.04%
4,000	\$9.17	\$9.17	\$8.86	\$6.82	-3.38%	-25.63%
5,000	\$9.17	\$9.17	\$8.86	\$7.23	-3.38%	-21.18%
6,000	\$9.17	\$9.17	\$8.86	\$7.50	-3.38%	-18.21%
7,000	\$9.17	\$9.17	\$9.69	\$7.81	5.67%	-14.83%
8,000	\$9.17	\$9.17	\$9.69	\$8.05	5.67%	-12.21%
9,000	\$9.17	\$9.17	\$9.69	\$8.23	5.67%	-10.25%
10,000	\$9.17	\$9.17	\$9.69	\$8.38	5.67%	-8.62%
11,000	\$9.17	\$9.17	\$9.69	\$8.50	5.67%	-7.31%
12,000	\$9.17	\$9.17	\$11.62	\$8.76	26.72%	-4.47%
13,000	\$9.17	\$9.17	\$11.62	\$8.98	26.72%	-2.07%
14,000	\$9.17	\$9.17	\$11.62	\$9.17	26.72%	0.00%
15,000	\$9.17	\$9.17	\$11.62	\$9.33	26.72%	1.74%
16,000	\$9.17	\$9.17	\$11.62	\$9.47	26.72%	3.27%

Notes: Marginal price is defined as the dollar amount paid for the next 1,000 gallons of water consumed for both water and sewer use. Average price is defined as the total bill for water and sewer use without inclusion of base service fees divided by the amount of water consumed.

2.3.4 Difference-in-differences

First, I specify a DD model to estimate an overall effect of the rate structure change on consumption. This initial specification, similar to that of Nataraj and Hanemann (2011), takes the form:

$$w_{it} = \beta_1 post_t + \beta_2(treat_i \times post_t) + Z_{it}\theta + \alpha_i + \epsilon_{it} \tag{2.6}$$

where w_{it} is the quantity of water consumed by household i in month t in thousands of gallons, $post_t$ is a dummy variable equal to 1 if the period of consumption is after the rate change and 0 otherwise, $treat_i$ is a dummy variable equal to 1 for treatment households, Z_{it} is a vector of control variables, α_i are the household fixed effects, and ϵ_{it} is the residual error term. If the common trend assumption between treatment and control groups is satisfied, β_2 will represent the average causal effect of the change in rate structure on the consumption of treatment households.

Figure 2.2 illustrates consumption patterns by treatment status for monthly consumption at the 50th percentile. As shown, there is strong co-movement prior to October 2007. At the time of the rate change, represented by the vertical line, the difference in trends between treatment and control households diverges in both magnitude and slope. This graphical analysis evidences common trends prior to the rate change and suggests that the change in the rate structure has a positive effect on consumption for treatment households. While there appears to be a slight deviation in the trends directly before the rate change, that this effect is sustained for the succeeding two billing periods suggests that there is a positive response among treated households.

2.3.5 Difference-in-difference-in-differences

Next, I exploit the assignment of billing cycles by applying a difference-in-difference-in-difference (DDD) strategy similar to Gruber (1994) and Davidoff et al. (2005) to assess the relative response to the introduction of tiered rates among

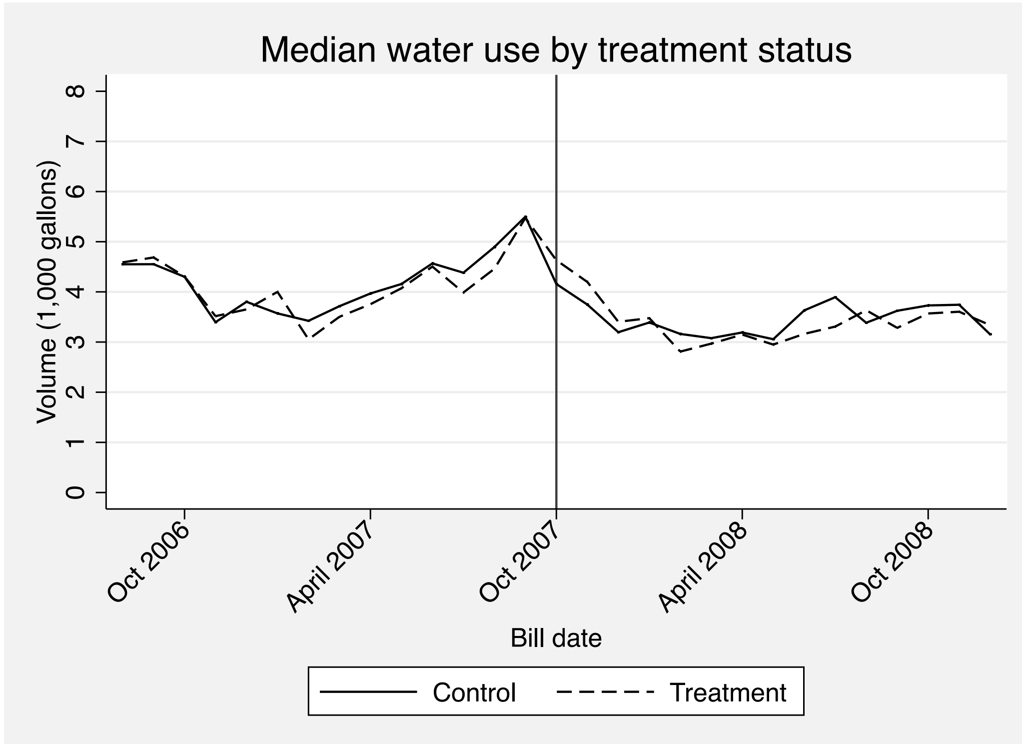

Figure 2.2: Median monthly water consumption for treatment and control households

different levels of consumption. The baseline model is,

$$
\begin{aligned}
w_{it} = \beta_1 post_t + \beta_2(treat_i \times post_t) + \sum_{j=2}^{10} \gamma_j (post_t \times D_i^j) \\
+ \sum_{j=2}^{10} \delta_j (treat_i \times post_t \times D_i^j) + Z_{it}\theta + \alpha_i + \epsilon_{it} \qquad (2.7)
\end{aligned}
$$

where D_i^j is a dummy variable equal to 1 if average fall consumption in the year prior to treatment is in the jth decile and 0 otherwise; the other variables are the same as in Equation 2.6.[9]

[9]The D_i^j term is determined by the decile of mean consumption for September, October, and November consumption in 2006 across all households. This allows the assignment of a decile group that reflects typical, pre-treatment, fall consumption for each household. All models were estimated with decile groups assigned by mean fall consumption in 2006 and 2008 to 1) allow for a more robust assessment of a consumer's typical fall consumption, 2) mitigate the effect of mean reversion

The model illustrated in Equation 2.7 differs from traditional triple-difference models in that it allows for household-specific fixed effects which prevent the inclusion of time-invariant regressors.[10] The intuition of this equation, however, remains the same: β_1 absorbs any utility-specific changes over time, β_2 is the treatment effect for the omitted decile dummy, γ_j controls for decile-specific changes over time, and the set of third-level interaction terms, δ_j, captures the DDD treatment effect of water consumption, conditioned on the decile, relative to control households after the rate structure change. The assumption necessary for clean identification of a treatment effect across deciles is that there be no contemporaneous shock that affects the relative response of treatment households in the same time period as the rate change (Gruber, 1994). This assumption is plausible since treatment status is based on billing cycles that have overlapping consumption periods in a small geographic area serviced by the same utility. Thus, any variation in weather, utility-specific effects such as drought conservation programs, or other exogenous shocks to demand are likely to affect all households uniformly within the decile. As mentioned previously, the main threats to this assumption are 1) the increased exposure of treatment households to drought restrictions and 2) the relative decrease in billing period length among treatment households. These confounding effects are examined through several robustness checks and since positive treatment effects are found, the

for a single month with unusually high consumption, and 3) control for decreasing medium-run trends in water consumption observed at the utility-level. These models produced qualitatively similar results and are available from the author upon request. Pre-treatment decile groups are preferred since all households eventually received the treatment when using 2008 consumption in the assignment of decile groups.

[10]The triple-difference moniker is typically reserved for specifications with an additional group of controls, whereas I simply adopt the terminology to explore heterogeneous treatment effects.

main findings of the DDD models are interpreted as lower bounds of the true treatment effect.

To quantify heterogeneous treatment effects, means for the treatment and control groups are presented in Table 2.3, which displays average consumption in each decile before and after the rate change for treatment and control households. As shown, mean consumption decreased for each decile after the rate change. For the first nine deciles, mean consumption decreased by a smaller amount for treatment households than it did for control households. Only in the tenth decile did consumption decrease by an amount larger than that of control households. These dynamics are captured in Panel C of Table 2.3, which are changes in consumption conditional on consumption decile for treatment households relative to control households in response to the rate change. The DDD effects indicate that for average consumption less than 11,000 gallons, the change in the rate structure increased relative consumption between treatment and control groups. For households in the tenth decile, the rate structure change reduced relative consumption. The standard errors, however, indicate the treatment and control households in the top and bottom 20% of the distribution are statistically similar, though regression will improve upon the precision of these effects.

2.3.6 Regression discontinuity framework

Lastly, I employ a fuzzy regression discontinuity (FRD) framework in the spirit of Nataraj and Hanemann (2011) to examine consumption behavior at a discontinu-

Table 2.3: Difference-in-difference-in-difference average treatment effects by decile

	Before Rate Change			After Rate Change				
Decile	Mean	Std. Err.	Obs	Mean	Std. Err.	Obs	Diff.	Std. Err.
A. *Water consumption for treatment households by decile*								
1	2.578	(0.173)	526	2.297	(0.130)	535	-0.281	(0.214)
2	2.679	(0.120)	405	2.414	(0.101)	408	-0.265	(0.156)
3	3.532	(0.156)	496	3.004	(0.096)	500	-0.528	(0.178)
4	4.198	(0.201)	516	3.662	(0.117)	521	-0.535	(0.224)
5	4.631	(0.143)	583	4.198	(0.112)	586	-0.433	(0.180)
6	5.352	(0.190)	457	4.862	(0.168)	458	-0.490	(0.253)
7	6.947	(0.240)	624	5.978	(0.179)	630	-0.969	(0.296)
8	7.426	(0.227)	476	6.608	(0.161)	479	-0.819	(0.274)
9	9.755	(0.309)	584	7.64	(0.225)	591	-2.116	(0.377)
10	17.535	(0.483)	757	12.37	(0.300)	759	-5.166	(0.553)
B. *Water consumption for control households by decile*								
1	2.677	(0.186)	575	2.133	(0.141)	562	-0.543	(0.231)
2	3.389	(0.192)	453	2.785	(0.156)	442	-0.603	(0.246)
3	4.403	(0.266)	506	3.175	(0.126)	503	-1.228	(0.277)
4	4.953	(0.300)	551	3.666	(0.116)	548	-1.287	(0.294)
5	5.804	(0.241)	556	4.431	(0.156)	559	-1.373	(0.280)
6	6.525	(0.256)	520	5.039	(0.172)	512	-1.486	(0.303)
7	7.652	(0.319)	592	5.676	(0.186)	586	-1.976	(0.357)
8	8.296	(0.298)	426	6.156	(0.179)	423	-2.140	(0.337)
9	10.747	(0.360)	451	8.336	(0.336)	446	-2.411	(0.492)
10	17.005	(0.841)	395	12.03	(0.524)	387	-4.979	(0.966)
C. *DDD Average treatment effect*								
1	0.262	(0.315)						
2	0.339	(0.286)						
3	0.700	(0.322)						
4	0.751	(0.368)						
5	0.940	(0.325)						
6	0.996	(0.394)						
7	1.007	(0.461)						
8	1.321	(0.431)						
9	0.295	(0.608)						
10	-0.187	(1.013)						

Notes: Treatment status is determined by billing cycle. Households are assigned to decile groups based on mean consumption in fall of 2006. DDD mean effects are calculated by subtracting the difference estimates in panel B from those in panel A for each decile group.

ous block endpoint due to the introduction of tiered pricing in Chapel Hill. Nataraj and Hanemann assess the treatment effect for water customers above and below the price kink point for the addition of a third price block at bi-monthly consumption levels of 40 ccf, while I assess one point along the distribution of consumption (7,000 gallons per month) at which the change in marginal and average prices move in opposite directions. The FRD model has the same general form as the DD model developed in Equation 2.6, however, treatment status is predicted by whether average fall consumption in 2006 was observed at or above 7,000 gallons. Customers in this consumption range are expected to face an increase of 5.7% in marginal price and a 14.8% decrease in average price due to the rate change. Thus, the treatment effect, β_2, from local linear regression of Equation 2.6 around 7,000 gallons would be negative if customers respond to marginal price and positive if customers respond to average price, relative to control households below the discontinuity.

Typically, regression discontinuity models require the forcing variable to lie either above or below a particular threshold (Imbens and Lemieux, 2008). In this paper, the forcing variable is whether a water customer's average fall consumption in 2006 (the year before the rate change) exceeds 7,000 gallons. The 7,000 gallon indicator is chosen as the cutoff since that is the point at which marginal and average prices diverge under the new rate schedule. Since the forcing variable is not identified with contemporaneous consumption, it has the benefit of being plausibly exogenous within the study period. But, the imprecise nature of this assignment results in a fuzzy assignment to treatment (i.e., as a household's historical fall consumption approaches the cut-off, the probability of treatment assignment does not jump cleanly from zero to one at the cutoff (Imbens and Lemieux, 2008)).

The appropriateness of an FRD methodology in this context is highlighted graphically in Figure 2.3 which depicts mean water consumption in 50 gallon bins relative to the distance from the 7,000 treatment cutoff under varying bandwidth levels.[11] Additionally, I examine the distribution of the forcing variable according to McCrary (2008) to test for discontinuities at the 7,000 gallon cutoff for the same set of bandwidths. This procedure is performed by fitting kernel densities to either side of the cut-off as well as the entire domain of the forcing variable. For all bandwidths considered, there is an observable jump in consumption at the cutoff which provides evidence that there is a discontinuous change in consumption behavior due to treatment. But, there is no observable change in the distribution of the assignment variable at the point of discontinuity for the density of the forcing variable. While the magnitude of the consumption discontinuity wanes as the bandwidth increases in Panels A, C, E, and G, the graphical analysis presented in Figure 2.3 provides convincing evidence that the regression discontinuity specifications capture a positive and plausible response to the treatment.[12]

[11]Bandwidth refers to the absolute distance (in units of the forcing variable) from the discontinuity within which observations are included in the regression (McCrary, 2008).

[12]For a more detailed discussion on the appropriateness of regression discontinuity (and the difference between RD and FRD), see Imbens and Lemieux (2008); Lee and Lemieux (2010); McCrary (2008), and Angrist and Pischke (2009).

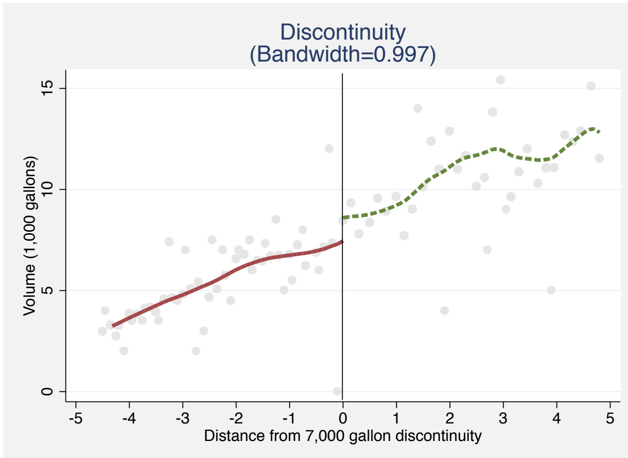

(a)

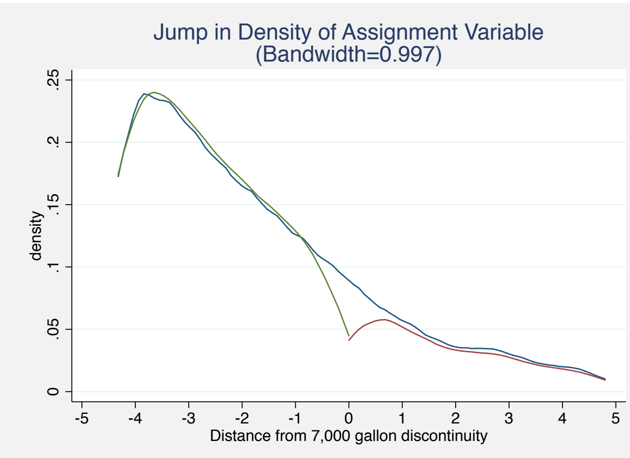

(b)

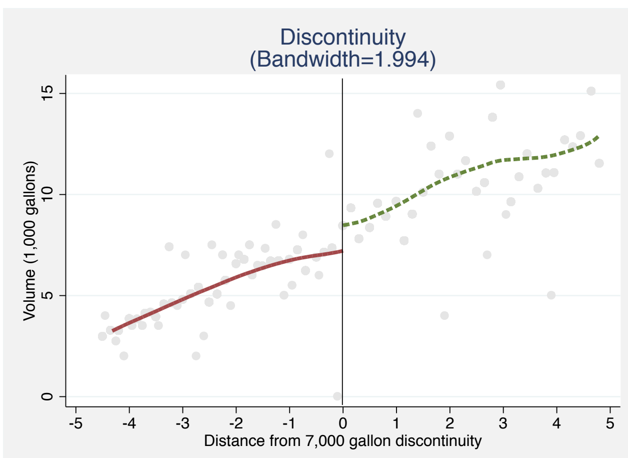

(c)

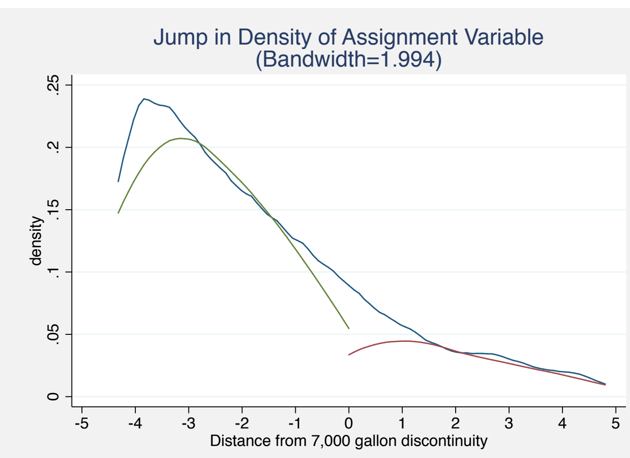

(d)

Figure 2.3: Regression discontinuity consumption at 7,000 gallon cut-off

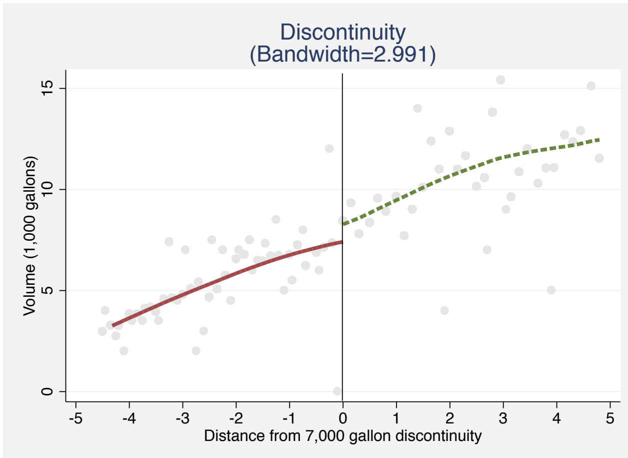

(e)

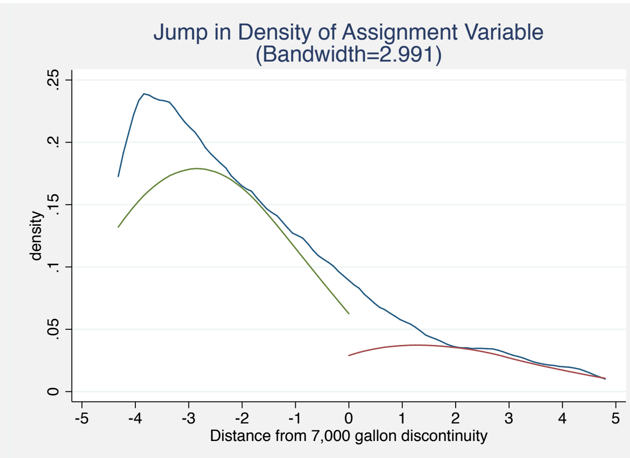

(f)

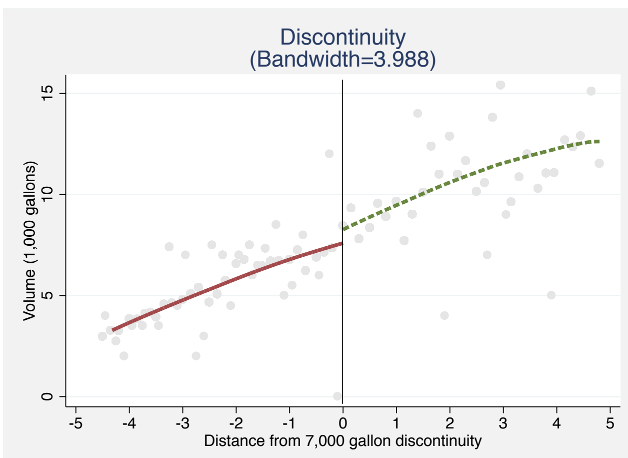

(g)

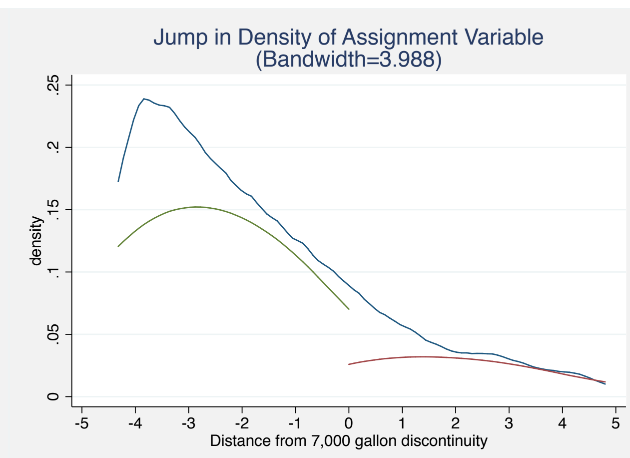

(h)

Figure 2.3: Regression discontinuity consumption at 7,000 gallon cut-off (cont.)

2.4 Results and discussion

In this section, I present empirical results that identify causal effects of price across the distribution of residential water consumption. DD models indicate a positive average treatment effect arising from the change in rate structure. Additionally, the DDD specification provides evidence that most residential water customers respond to average price rather than marginal price as determined by their previous bill. To examine this effect more closely, FRD methods are applied at a point of divergence in marginal and average price and reinforce the notion that customers are indeed responding to average price. Finally, I assess the robustness of these results by varying sets of controls and assigning false treatment status. The results of falsification tests imply that the DDD and FRD models capture a valid response to changes in average price.

2.4.1 Difference-in-difference estimation results

First, I estimate a traditional DD model to capture the overall treatment effect of the rate structure change on consumption. The results of this model are presented in Column 1 of Table 2.4. The treatment effect, estimated at 433 gallons per month, indicates that the rate structure change induced an overall increase in consumption due to the lower rates for the first units of consumption. This effect is significant at the 1% level. While one of the reasons that increasing block rates are typically adopted is to encourage conservation among high users, this result indicates that lower prices among low-users exhibits a perverse effect of raising aggregate consumption. This positive effect also assuages concerns that the treatment effect is estimating the effect of something other than the change in the rate structure since the primary sources of concern—seasonal changes in consumption, conservation initiatives, and divergent billing durations between groups—would affect treatment consumption negatively.

Table 2.4: Difference-in-difference regression results

	(1) Fixed Effects	(2) Random Effects	(3) Random Effects	(4) Random Effects
post	-1.770***	-1.758***	3.410	2.376
	(0.079)	(0.081)	(3.887)	(3.109)
treat		0.177	1.541	1.541
		(0.137)	(1.035)	(1.035)
treat×post	0.433***	0.410***	0.904**	1.747*
	(0.111)	(0.111)	(0.388)	(1.012)
evapotr			3.061	
			(2.301)	
maxtemp				0.632
				(0.475)
Hausman test statistic:	143.845			
(p-value)	(<0.001)			
Observations	21,094	21,094	21,094	21,094
Number of Households	10,654	10,654	10,654	10,654
Within R-Squared	0.070	0.070	0.070	0.070

Notes: Dependent variable is monthly water consumption in thousands of gallons. Fixed effects are at the household level. Robust standard errors are clustered at the household level. *, **, and *** represent significance at the 0.10, 0.05, and 0.01 levels. The Hausman test is adjusted for heteroskedasticity and within-household error correlation.

Since weather varies only at the utility level, temperature and evapotranspiration effects are perfectly collinear with the *post* variable and the household fixed effects in a two-period framework. Thus, I also estimate the DD model in a random effects (RE) framework to assess whether the treatment effect is influenced by weather by exploiting cross-sectional variation.[13] The baseline random effects model in Column 2 of Table 2.4 is statistically similar to Column 1.[14] In Columns 3 and 4, I sequentially introduce evapotranspiration and maximum temperature to avoid further collinearity (since evapotranspiration is a function of temperature levels). Neither of the estimated parameters on weather regressors are statistically different than zero. Since RE models are inconsistent, these estimates should be interpreted with caution. But, the RE models provide justification for moving forward with the fixed effects specification without inclusion of weather controls.

2.4.2 Difference-in-difference-in-difference estimation results

The estimation results of the DDD model described in Section 3 are presented in Column 1 of Table 2.5. The coefficients listed are the third-level interaction effects diagrammed in Equation 2.7. These effects represent the adjusted consumption response for treatment households relative to control households in each decile.[15] For the second through eighth decile interactions, the consumption response is positive and jointly significant at the 1% level, while the ninth decile is positive and significant at the 5% level. Due to aggregate seasonal decreases in water use, this result

[13]To identify a comparable treatment effect in this framework an indicator for treatment households must be included to absorb unobserved variation across treatment and control groups since this variation was previously captured by the household-specific fixed effect.

[14]A Hausman test (adjusted for heteroskedasticity and within-household error correlation) rejects the hypothesis with p-value<0.001 that the strong exogeneity assumption holds, which is necessary for the RE model to be consistent. As such, the RE models only provide suggestive evidence of the implications of weather on treatment effects.

[15]Rather than reporting nominal regression parameters, I report the overall treatment interaction effect for the third-level DDD estimates relative to the omitted category. This allows for a more intuitive interpretation. A full set of nominal regression parameters are available from the author upon request.

implies that the treatment households reduced consumption by an amount less than that of control households in the same decile. The insignificance in the first decile is not surprising since habitually low-use customers cannot cut back consumption in response to a rate increase and they likely have little preference to increase consumption significantly in response to lower rates. Conversely, the coefficient for the tenth decile of consumption is negative, though not significantly different from zero. Intuitively, this decile group faced increases in both marginal and average price so a negative effect is expected.

Mean volume for each decile prior to the rate change is presented in Column 2 of Table 2.5. In order to present results that compare this treatment effect with other studies, own-price elasticity of demand estimates are computed based on the percent change in consumption and the percent change in marginal and average price for treatment relative to control households. Elasticity estimates are presented in Columns 3 and 4 of Table 2.5. Standard errors for the elasticity estimates are simulated with 1,000 draws using the Krinsky and Robb (1986) methodology assuming elasticities are distributed normally and centered at the calculated elasticity estimate with the standard deviation derived from the estimated variance-covariance matrix. For the first two deciles, the marginal and average price elasticity estimates are equivalent because prices are identical for this level of consumption. Marginal price elasticity estimates for the third through seventh decile range from -0.569 to -6.182, while average price elasticity estimates range from -0.734 to -1.103. Within this central region of the consumption distribution, marginal price estimates are implausibly large as all but one lie outside the range of previous elasticity estimates,

Table 2.5: Difference-in-difference-in-difference regression results and elasticity estimates by decile

	(1) Adj. DDD Estimate	(2) Volume (1,000 gal)	(3) MP Elasticity Estimate	(4) AP Elasticity Estimate
treat×post×D1	0.287	2.578	-0.336	-0.336
	(0.208)	(0.173)	[1.015]	[0.890]
treat×post×D2	0.380***	2.679	-0.429	-0.429
	(0.032)	(0.120)	[0.290]	[0.185]
treat×post×D3	0.664***	3.532	-0.569	-0.734
	(0.067)	(0.156)	[0.471]	[0.273]
treat×post×D4	0.807***	4.198	-5.684	-0.750
	(0.088)	(0.201)	[0.592]	[0.305]
treat×post×D5	0.968***	4.631	-6.182	-0.987
	(0.050)	(0.143)	[0.399]	[0.177]
treat×post×D6	1.101***	5.352	-6.088	-0.972
	(0.062)	(0.190)	[0.513]	[0.189]
treat×post×D7	1.134***	6.947	-4.828	-1.103
	(0.091)	(0.240)	[0.523]	[0.194]
treat×post×D8	1.252***	7.426	2.974	-1.139
	(0.075)	(0.227)	[0.488]	[0.158]
treat×post×D9	0.434**	9.755	0.784	-0.514
	(0.189)	(0.309)	[0.326]	[0.212]
treat×post×D10	-0.083	17.535	-0.018	-0.048
	(0.395)	(0.483)	[0.201]	[0.221]
Observations	20,884			
Number of Households	10,542			
Within R-Squared	0.128			

Notes: Dependent variable is monthly water consumption in thousands of gallons. Households are assigned to decile groups based on mean consumption in fall of 2006. Time dummy, second-level decile interactions, and constant term are omitted. Fixed effects are at the household level. Parameter estimates presented in Column 1 are the adjusted third-level DDD interaction terms. Elasticity estimates are calculated for the mean consumption at the decile group and price changes relative to prices prior to the rate change. Robust standard errors in Column 1 are clustered at the household level; standard errors of the mean are presented in parentheses in Column 2. Standard errors presented in brackets in Columns 3 and 4 represent are simulated with 1,000 draws using the Krinsky and Robb (1986) methodology assuming elasticities are distributed normally and centered at the calculated elasticity estimate with standard deviation derived from the estimated variance-covariance matrix. *, **, and *** represent significance at the 0.10, 0.05, and 0.01 levels.

though average price estimates are well within this range (Dalhuisen et al., 2003; Espey et al., 1997).[16]

The coefficients for the eighth and ninth decile are most illuminating with respect to identifying the price to which consumers respond. For this subset of customers, the marginal price increased roughly 5% for treatment households while average price decreased between 7% and 15% relative to control households.The estimated regression parameters for water customers in this group exhibit a response that corresponds to a positive marginal price elasticity, ranging from 0.784 to 2.974, and a negative average price elasticity estimate, ranging from -1.139 to -0.514. Lastly, the result for the tenth decile, though not statistically different from zero, displays plausible elasticity estimates (-0.018 and -0.048) for both marginal and average price, respectively. Collectively the DDD results indicate that consumers generally respond to average price.

2.4.3 Regression discontinuity estimation results

To further analyze the consumption response to changes in price, the FRD analysis focuses on a point in the rate schedule with a discontinuous divergence in marginal and average price for the same level of consumption after the rate change. For consumption levels at 7,000 gallons per month, marginal price increases by 5.67% more than the previous uniform rate, while average price decreases by 14.83% for treatment households relative to the control households. Consider a customer

[16] In a meta-analysis, Espey et al. (1997) find price-elasticity estimates of residential water demand to range from -0.02 to -3.33 with an average of -0.51.

just above the discontinuity—if she responds to marginal price, she would reduce consumption commensurate with the increase in marginal price relative to customers below that level; however, if she responds to average price, then her consumption response will be positive.

In Panel A of Table 2.6, local linear regression estimates of Equation 2.6 around 7,000 gallons are presented. At a bandwidth of 5,000 gallons around the FRD cut-off, the treatment effect is positive and significant at the 1% level. Within 4,000 gallons of the cut-off, the treatment effect tends monotonically toward zero, but remains significantly positive, at 291 gallons per month, within 1,000 gallons of the cut-off. This result indicates that customers just above the discontinuous jump in price increased consumption by 4.54% in response to the change in price relative to control households.[17] Since this effect is positive, it lends credibility to the notion that consumers respond to average, not marginal, price. For the sake of comparison to previous research, this effect is interpreted as a local price elasticity estimate of -0.31, dividing the 4.54% increase in consumption by the 14.83% decrease in average price, for treatment households after the rate change within 1,000 gallons of the cut-off. This elasticity estimate is slightly different than the estimate for similar consumption levels in the DDD model because the control group in the RD model is composed of households on the new price schedule with average fall consumption below the 7,000 gallon cut-off as well as the control households who have not yet observed the new price schedule.

[17]The percent change in consumption is calculated using 6,413 gallons as the pre-treatment mean consumption for treatment households within 1,000 gallons of the cut-off.

Table 2.6: Fuzzy regression discontinuity results

	(1) +/- 5,000 gallons	(2) +/- 4,000 gallons	(3) +/- 3,000 gallons	(4) +/- 2,000 gallons	(5) +/- 1,000 gallons
A. *Regression discontinuity cut-off at 7,000 gallons (Treatment and control)*					
treat×post	0.288***	0.431***	0.351***	0.320***	0.291***
	(0.086)	(0.082)	(0.080)	(0.080)	(0.070)
Observations	16,084	12,909	9,398	6,071	2,944
Within R-Squared	0.079	0.080	0.072	0.072	0.032
B. *Regression discontinuity cut-off at 7,000 gallons (Treatment only)*					
post	-0.270***	-0.105	-0.094	-0.048	0.204***
	(0.083)	(0.078)	(0.075)	(0.073)	(0.062)
Observations	1,447	1,290	1,104	849	471
Within R-Squared	0.017	0.003	0.004	0.002	0.105
C. *False regression discontinuity cut-off at 7,000 gallons (Treatment and control)*					
treat×post	0.093	0.012	-0.093	0.048	0.066
	(0.108)	(0.103)	(0.104)	(0.101)	(0.095)
Observations	15,174	12,330	8,982	5,807	2,822
Within R-Squared	0.006	0.009	0.013	0.003	0.015

Notes: Dependent variable is monthly water use in thousands of gallons. In Panel A, treatment status is assigned to households who satisfy two requirements: 1) were on the new price schedule as determined by their billing cycle and 2) had mean fall consumption in 2006 at or above 7,000 gallons per month. In Panel B, only treatment households are included such that the coefficient on post is simply the FRD treatment effect within treatment households. In Panel C, false treatment status is assigned to households who had mean fall consumption above 7,000 gallons per month for an artificial cut-off two months prior to the rate change. Fixed effects are at the household level. Robust standard errors in parentheses are clustered at the household level. *, **, and *** represent significance at the 0.10, 0.05, and 0.01 levels.

2.4.4 Robustness checks and falsification tests

Since the primary threats to identification in the DDD model are 1) the increased exposure to drought restrictions for the treatment group due to their delayed billing cycle and 2) the duration of billing lengths among treatment and control groups, the regression discontinuity model is estimated on a subsample of treatment households only. This strategy eliminates the potential confounding effect of drought restrictions or bill length since treated households are otherwise identical, though treatment is now predicted strictly by whether fall consumption lies above or below the 7,000 gallon discontinuity in price. Results are shown in Panel B of Table 2.6. Within 5,000 gallons of the cut-off, the treatment effect is negative and significant. As the bandwidth decreases, the treatment effect tends monotonically towards zero, becoming positive and significant within 1,000 gallons of the cut-off. The initial negative effect can be interpreted as habitually low users being poor controls for habitually high users in that they do not display qualitatively similar consumption patterns. As households become more similar in consumption patterns, the treatment effect becomes insignificant. The fact that a positive and significant treatment effect is identified within 1,000 gallons, and statistically similar to the treatment effect in Panel A, is reassuring given the lack of statistical power in the restricted sample. Regardless, this result provides strong evidence that the identification strategy is not confounded by the contemporaneous drought restrictions.

Additionally, to test whether the estimated treatment effects in the regression discontinuity models are valid, I estimate the same FRD model in Section 2.4.3 for two consecutive time periods two months prior to the rate change. Thus, in this falsification test, all households face the exact same uniform rate structure. The results of the false treatment FRD models are presented in Panel C of Table 2.6. Within all bandwidth specifications, this model produces no statistically significant treatment effects and the parameter estimates are small in magnitude relative to the FRD treatment effects found previously. With the lack of a non-zero treatment effect, this falsification test indicates that the true FRD model is estimating a plausible response to a discontinuous change in price.

Finally, the DDD model is estimated in a random effects framework including a set of weather controls. I use a random effects estimator to exploit between household variation which allows for the inclusion of weather parameters without being perfectly collinear with household fixed effects and the time indicators.[18] The first column of Table 2.7 presents the baseline fixed effects estimates discussed in Section 2.4.2 for comparison. Results from a baseline random effects model with no weather controls is presented in Column 2.[19] While the RE analysis is merely suggestive, the inclusion of weather controls in the empirical specification enlarge the coefficients of interest, though all significant treatment effects in Column 1 maintain conventional levels of significance in the RE framework. The weather covariates exhibit insignificant parameters. Overall, these results mimic the intuition of the estimates in Table 2.6 and provide justification that the fixed effects model adequately controls for changes in weather across deciles.

2.4.5 Lingering empirical concerns

This subsection outlines the implications of empirical concerns that could lead to potentially biased estimates and evaluates how these concerns might affect the

[18] Similar to the DD model, the random effects estimates rely on inclusion of baseline treatment effects and additional secondary interactions such that Equation 2.7 must include time-invariant regressors at the household-level. The results of interest are presented in Table 2.7.

[19] The Hausman test statistic (adjusted for heteroskedasticity and within-household error correlation) rejects the consistency of the RE model with p-value<0.001.

Table 2.7: Difference-in-difference-in-difference regression results by decile with random effects and weather variables

	(1) Fixed Effects	(2) Random Effects	(3) Random Effects	(4) Random Effects
treat×post×D1	0.287	0.277	0.764	1.592*
	(0.208)	(0.207)	(0.489)	(0.937)
treat×post×D2	0.380***	0.363***	0.852***	1.680*
	(0.032)	(0.032)	(0.230)	(0.911)
treat×post×D3	0.664***	0.679***	1.164***	1.992**
	(0.067)	(0.067)	(0.262)	(0.885)
treat×post×D4	0.807***	0.784***	1.268***	2.096**
	(0.088)	(0.087)	(0.283)	(0.877)
treat×post×D5	0.968***	0.957***	1.439***	2.267***
	(0.050)	(0.049)	(0.241)	(0.867)
treat×post×D6	1.101***	1.059***	1.542***	2.370***
	(0.062)	(0.062)	(0.254)	(0.888)
treat×post×D7	1.134***	1.082***	1.567***	2.395***
	(0.091)	(0.090)	(0.287)	(0.864)
treat×post×D8	1.252***	1.281***	1.767***	2.595***
	(0.075)	(0.076)	(0.249)	(0.900)
treat×post×D9	0.434**	0.377**	0.861**	1.6890*
	(0.189)	(0.187)	(0.380)	(0.883)
treat×post×D10	-0.083	-0.125	0.364	1.192
	(0.395)	(0.390)	(0.601)	(1.031)
evapotr			3.005	
			(2.717)	
maxtemp				0.620
				(0.414)
Hausman test statistic:	92540.8			
(p-value)	(<0.001)			
Observations	20,884	20,884	20,884	20,884
Number of Households	10,542	10,542	10,542	10,542
Within R-Squared	0.128	0.043	0.043	0.043

Notes: Dependent variable is monthly water consumption in thousands of gallons. Households are assigned to decile groups based on mean consumption in fall of 2006. Second-level interactions, baseline effects, and the constant term are omitted. Fixed effects are at the household level. Parameter estimates presented are the adjusted third-level DDD interaction terms. Robust standard errors are clustered at the household level. *, **, and *** represent significance at the 0.10, 0.05, and 0.01 levels. The Hausman test is adjusted for heteroskedasticity and within-household error correlation.

external validity of this study. First, the most obvious shortcoming of this analysis is the inability to include demographic or household characteristic information to illustrate that demographics are qualitatively similar among treatment and control groups. Since the data used in this analysis were stripped of any unique geographic identifier for consumption records other than the billing cycle, demographic and spatial comparisons between treatment and control groups could not be performed. Thus, I rely on analyzing consumption patterns between billing cycles to argue that the difference between groups is negligible and these data provide a valid quasi-experimental setting. While this argument is admittedly weak, the econometric methods used mitigate the effect of this uncertainty. The DDD methods provide nonparametric control for decile-specific effects and relax the common trend assumption necessary for clean identification of a treatment effect (Gruber, 1994). Moreover, the FRD method further avoids this potential source of bias by comparing households just above and just below a discontinuous jump in prices for both treatment and control households, as well as treatment households only. Since there is reason to believe that demographics and water use are correlated (Nataraj and Hanemann, 2011), as we approach the discontinuity, the effect of heterogeneity in household composition is mitigated. Further, all econometric models control for time-invariant household characteristics. Thus, I contend that these concerns do not contaminate the results, though examining demographic and spatial heterogeneity in this context would be a fruitful area for further research.

Additionally, the identification strategy relies on the staggered nature of utility billing cycles such that drought regulations could potentially confound results due to the treatment group's relatively longer exposure to these policies. Further, the fact that the length of the billing cycle for treatment households becomes relatively shorter than that of control households potentially biases the findings of this paper. These effects, however, work in the opposite direction of the estimated treatment effects biasing the results toward zero. Since the main treatment effects are positive, it is likely that the estimated price elasticities are lower bounds of the true elasticity. The regression discontinuity model estimated only for treatment households allows for this effect to be isolated, and results are robust to this stratification.[20]

Further, the identification strategy relies on plausibly random assignment into treatment and control groups as well as an exogenous forcing variable that delineates the cutoff in FRD models. For the former, I assign treatment status as determined by a customer's billing cycle. Billing cycles are deterministic in that they are constructed for practical convenience such that a utility employee can conveniently drive along a meter route to read meters. While there is spatial correlation between households and demographic composition within billing cycles, there is little observed difference between household water consumption within the decile of consumption for the DDD models. Further, the main regression discontinuity specification includes households from the same billing cycle in both the control and

[20]In addition, I examine DDD mean effects for treatment and control groups in September and October of 2006 prior to the 2007 drought. During this time period, a small uniform rate increase occurred for all customers, which is identified exactly according the treatment-control strategy outlined previously given the staggered nature of utility billing. All of the deciles display similar trends in reductions due to seasonality and the effects and the triple-difference mean effects are primarily small and insignificant except for the upper and lower tails of the distribution. This serves to further alleviate concerns that a confounding effect of drought restrictions or billing period length is biasing results and ultimately strengthens the main findings of this paper. These results are available from the author upon request.

treatment groups, thus mitigating any deterministic bias that might arise from non-random assignment into the treatment group. Additionally, the household itself has no ability to change its billing cycle short of moving across billing cycle boundaries, thus it is not likely that there is any manipulation of the treatment within the time frame of the study. For the latter, since the forcing variable in the RD models exploits the implementation of a five-block tiered price schedule adopted for the utility's revenue goals, conservation among high-users, and affordability among low-income customers, it is possible that the block endpoints were chosen to explicitly align with certain indicators in the distribution of customer consumption. If this is the case, then perhaps the 7,000 gallon discontinuity in price was chosen to penalize household water use above the mean of the total customer base. Since there is no observable discontinuity in the distribution of water consumption (as illustrated in Figure 2.3) it is assumed that this assignment is as good as random.

Finally, though the results of this paper imply that residential water customers respond to average price, it should be noted that this study relies on the introduction of an entirely new price structure and not merely a marginal change in the price level. Thus, it is possible that the variation in consumption reflects uncertainty about the new rate structure. The treatment effect, however, is generally positive. If households were reacting to uncertainty about the new rate structure, the expectation would be a negative or null effect until households learn about how their bill is calculated. Because the positive treatment effect is robust across specifications, it is not likely that consumers are responding to this uncertainty. Lastly, while this analysis was motivated by the structural assumptions of consumer demand under tiered rate schedules, the result that most consumers respond to average price is conditional upon the alternative hypothesis that the price to which consumers respond is a local marginal price, rather than the entire rate schedule as discrete-continuous choice models predict. Thus, the main findings of this research provide evidence that reduced-form models of consumer water demand should incorporate behavior commensurate with changes in average price, though a comparison of elasticity estimates from (quasi-)experimental methods to structural models of water demand is a promising area for future research.

2.5 Concluding remarks

In this paper, a conceptual framework for residential water demand that relies on a customer's consumption patterns last month as a heuristic for this month's prices is developed. Then, by exploiting the introduction of block rates and the assignment of billing cycles for residential customers in Chapel Hill, North Carolina, I identify a causal effect of price across the entire distribution of water customers using DDD techniques. In addition, FRD models are estimated to focus on an anomaly in the rate change in which changes in average and marginal price move in opposite directions for customers with consumption at 7,000 gallons per month.

The results of this analysis imply that residential water customers exhibit consumption patterns commensurate with changes in average, rather than marginal, price. Across the distribution of consumption, price elasticity estimates from DDD models range from -0.43 to -1.14 and households consuming between 4,000 and 7,000 gallons per month are found to be the most responsive to changes price. Additionally, empirical estimates from FRD models identify a positive treatment effect at a point where marginal prices increased and average prices decreased in response to the new rate structure. By exploiting the divergence in average and marginal prices, I obtain further support of the hypothesis that average prices are perceived by residential water customers. Within the tightest bandwidth of regression discontinuity models at 7,000 gallons per month, a local price elasticity is estimated at -0.31.

The results of this paper contribute in several ways to the literature on water demand and conservation, as well as the practitioner's guidebook on managing local water resources. First, this study provides empirical evidence that the price perceived by residential water customers is the average price from a customer's previous bill. This result adds to mounting evidence in the literature that consumers facing complicated pricing structures tend not to respond in a manner that standard utility theory predicts. Second, the introduction of increasing block rates can produce a perverse effect of increasing total demand due to price decreases in the lower blocks. Lastly, this paper provides evidence of heterogenous price elasticity estimates from quasi-experimental methods that support the well-accepted notion that residential water demand is generally price inelastic, but certainly not unresponsive to changes in price.

Chapter 3: Information provision and consumer behavior: A natural experiment in billing frequency

3.1 Introduction

Conventional economic wisdom implies that more information is typically better. For many consumer goods and services, however, the decision to consume an economic good is disconnected from its purchase price. In these contexts, providing consumers with more information may affect their behavior. For consumption of water or electricity, for example, information on consumption costs is limited because billing is infrequent. If this source of limited information distorts the price signal that consumers use to make decisions, then improving the clarity of this signal has implications for consumer welfare.

Whether and how imperfect perception of prices and quantities affects consumer behavior is an empirical question of growing interest. A recent vein of literature suggests that consumers tend to underestimate prices, taxes, and quantities consumed that are transmitted opaquely or allow for customer inattention (Chetty et al., 2009; Grubb and Osborne, 2015). Empirical examples range from behavioral responses to tax-inclusive prices to improving the salience of consumption information through "bill shock" reminders for cell-phone use. A parallel literature on consumer behavior in environmental policy considers the impact of social norms (Ferraro and Price, 2013) and information provision (Jessoe and Rapson, 2014) and shows that informative interventions can reduce consumption and thus serve as an instrument of conservation.

With few exceptions, previous research suggests that various information treatments can be utilized to reduce consumption of economic goods that impose external costs on society. Particularly in highly regulated markets for electricity and water demand where prices are politically difficult to change, finding cost-effective conservation strategies is of topical policy interest. However, no one has examined the effect of changes in the frequency of information on consumer behavior in intermittent choice settings. Further, the majority of the existing research on informative interventions as a tool for conservation stops short of estimating changes in welfare, with Allcott and Taubinsky (2015) being a notable exception for the purchase of energy efficient lightbulbs.

In this paper, I make several contributions to this growing literature with an application to the management of a pertinent environmental resource. I take advantage of a natural experiment in which residential water customers are exposed to exogenous increases in billing frequency within a single water provider's service area in the southeastern United States. I find strong empirical evidence that the provision of more frequent information *increases* water consumption by approximately five percent. This result contrasts the findings of previous work and has significant implications for efficient management of scarce environmental resources.

A powerful implication of this empirical finding is that it raises concerns for the generalizability of informative interventions as a conservation tool in the salience literature, particularly with regard to electricity and water demand management. Similar to repeated interventions, I find a lasting effect within the 4.5-year study period. I show that households transition quickly to a new baseline equilibrium in which long-run treatment effects are consistently larger than the magnitude of the short-run effect.

My empirical results necessitate a closer examination of the mechanism driving consumer behavior in response to more frequent information. To that effect, I motivate a conceptual model of imperfect price perception that reconciles my empirical findings with the current literature on salience and inattention. Based on the notion that consumers are receiving more frequent information about the price of water with the receipt of monthly (versus bi-monthly) bills, the information "treatment" allows consumers to update their perception of price.[1] This framework is general enough to accommodate the findings of previous research since more frequent information nudges consumers closer to the neoclassical ideal of decision-making under perfect information. As a motivating example, a consumer who initially under-perceives the price of electricity can be modeled similarly to a customer who over-perceives the price of water, since more frequent billing will reduce the wedge between her *perceived* price and the actual price.

Within this framework, I develop transparent analytical formulas for calculating changes in welfare associated with more frequent information using treatment

[1] Bi-monthly bills refer to customer bills that are received every two months.

effects as sufficient statistics for consumer demand. Since a consumer who misperceives price (quantity), and thus consumes suboptimally from her perfectly informed self will be better off upon the receipt of new information there are welfare gains from information provision. I show that a reduction in quantity uncertainty driving consumer behavior in response to more frequent information provides a lower (upper) bound for welfare estimates relative to price misperception in the case of initial over- (under-)perception of prices. Consumer surplus measures suggest a welfare gain of approximately 0.5 to 1 percent of annual household expenditures on water that are attributable to the change in billing frequency.

The empirical setting is an exogenous transition of residential water customers from bi-monthly to monthly billing within a single water utility. Beginning in 2011, the City of Durham's Department of Water Management in North Carolina transitioned residential customers in geographically differentiated billing districts to monthly billing over the course of two-and-a-half years. By exploiting the assignment of monthly billing, I estimate an average treatment effect on water consumption due to increased billing frequency at the household level. The primary result is that households billed monthly consume approximately five percent more water than households billed bi-monthly. I show that this effect is robust to unobserved neighborhood effects by examining household consumption before and after the change in frequency within 500 feet of common billing group boundaries. Treatment effects are found to persist over time with a long-run treatment effect implying an approximately ten percent increase in water use. Further, I estimate conditional average treatment effects that indicate an increase in outdoor water use in the summer months. I also find important heterogeneity among baseline water use, lot size, and assessed home value.

From an environmental policy perspective, informative signals are being used increasingly as a regulatory instrument in the context of electricity and water conservation. The findings of this paper suggest that increases in billing frequency can have the perverse effect of increasing consumption. This result is particularly poignant because the efficient price for residential water is its long-run marginal cost of provision (Olmstead and Stavins, 2009; Timmins, 2002). However, since the market price is likely set below its efficient level (Mansur and Olmstead, 2012), the demand response to more frequent information may exacerbate the wedge between privately and socially optimal consumption levels.

In the next section, I motivate a conceptual framework to explore changes in billing frequency. In Section 3, I describe the data used in the analysis and outline the empirical setting. I present a series of quasi-experimental models to estimate a causal effect of information provision on consumer demand as well as heterogeneous treatment effects in Section 4. In Section 5, I discuss the results and implications of the empirical models, while I estimate associated welfare changes in Section 6. The final section concludes.

3.2 Conceptual framework

In this section, I first provide background on informative interventions and their effect on consumer behavior. Next, I develop a model of consumer misperception of price and quantity information, separately, to examine consumer decision making in light of receiving more frequent information. In each mechanism, a utility framework for misperceived prices and quantities is analyzed. Within those frameworks, I construct welfare measurements that rely on treatment effects as sufficient statistics and compare the analytical derivations. Further, I reconcile the difference between a consumer who may respond to more frequent billing through price perception, quantity perception, or both. Lastly, I consider briefly several alternative mechanisms that could drive consumer behavior in this choice setting.

3.2.1 Background

Consider the choice setting in which a consumer is deciding how much water to use in a given billing period.[2] Borenstein (2009), Gilbert and Graff Zivin (2014), Harding and Hsiaw (2014), and Wichman (2014), for example, motivate models of behavior based on prices, quantities consumed, and behavior in previous periods as heuristics for making consumptive decisions in electricity and water demand. Since utility bills are received periodically, the arrival of billing information offers consumers an opportunity to update their consumption in response to external feedback regarding their behavior. A change in the frequency of billing information is particularly relevant in the intermittent choice setting for water use since consumers generally do not know how much water they are using at any point in time, nor how much water an appliance uses and its associated variable costs. Thus, more fre-

[2] While the model presented in this paper is generalizable to many choice settings in which consumption of the economic good and payment for consumption are separated temporally (e.g., cell phone usage, credit card purchases, electricity demand, and so forth), the discussion henceforth will consider water consumption to motivate the empirical setting.

quent billing allows a consumer to better align price signals directly with the usage of appliances or water-intensive behavior.

With a fuzzy link between water consumption and the receipt of a water bill, however, the consumer may not have perfect information about prices and consumption that neoclassical models of consumer demand require. Several papers have documented this behavior theoretically and empirically in different markets. Numerous studies show that: 1) obtaining the relevant information to make perfectly informed decisions is costly (Caplin and Dean, 2014; Sallee, 2014; Shin, 1985); 2) consumers may be inattentive to or unaware of (changes in) prices or taxes (Chetty et al., 2009; Finkelstein, 2009; Houde, 2014; Li et al., 2014; Sexton, 2014); 3) inattention could be a function of attributes that are "shrouded" from consumers (Gabaix and Laibson, 2006); 4) consumers may use heuristics for decision-making when price and quantity information is opaque or uncertain (Borenstein, 2009; Ito, 2014; Wichman, 2014); or 5) consumers may have biased perceptions of prices, expenditures, and consumption (Allcott, 2013; Bollinger et al., 2011; Grubb and Osborne, 2015).[3] Thus, relaxing the notion that consumers respond with perfect information for water use should not be met with much criticism. But, the question remains: how are consumers using price and quantity information to make decisions in intermittent choice settings?

Many researchers examine this question in framed field experiments in the context of water and electricity demand to examine quantity reminders, social norms,

[3]This literature complements research that examines the effect of informative signals of product quality on consumer behavior (c.f., Foster and Just (1989) or Jin and Leslie (2003)).

and other forms of informative interventions (Allcott, 2011; Brent et al., 2014; Ferraro and Price, 2013; Jessoe and Rapson, 2014; Kahn and Wolak, 2013). But, no studies have focused on an information treatment as simple as the provision of more frequent billing information, which is arguably the easiest form of information provision to implement as policy. In the conceptual model presented below, the general static setting is one in which a consumer is planning for water consumption and expenditures conditional on her preferences, technology stock, past usage, and prevailing market prices. Then, the consumer is provided with an information shock—the receipt of a utility bill. The bill allows consumers to learn about their past usage and prices paid for water and update their consumption habits accordingly.[4]

In this framework, the customer's bill serves as a familiar mechanism for receiving information, but provides new price and quantity information to the consumer when it arrives. This treatment mechanism stands in contrast to many recent field experiments in which consumers are given a foreign source of information about their consumption. For example, social comparisons among use and expenditures (Allcott, 2011; Brent et al., 2014; Costa and Kahn, 2013; Ferraro and Price, 2013), educational materials about complicated rate structures (Kahn and Wolak, 2013), real-time feedback on consumption (Gans et al., 2013; Jessoe and Rapson, 2014; Strong and Goemans, 2014), or informative signals on variable costs of durable goods (Allcott and Taubinsky, 2015). The drawback of the natural experiment in this paper is the inability to control the mechanism through which consumers may

[4]An example of a utility bill, which serves as the information "treatment" in this paper, is included as Figure B.1 in Appendix B.

respond, though this is not guaranteed in field or lab experiments (Ludwig et al., 2011), as well as the lack of pure randomization. The primary benefit, however, is that consumers are familiar with their typical utility bill, so it is perhaps more likely that they will simply adjust their behavior along an existing margin. That is, consumers will update whatever decision rule they use to make water consumption decisions, rather than constructing a new rule in response to a foreign information intervention. In this context, consumers are not receiving *more* nor *better* information about their consumption, they are simply receiving the same information more frequently. To the extent that this is true, the natural experiment isolates a consumer response along the frequency of information provided, rather than the quality or quantity of information provided. Thus, any prevailing misperceptions within a consumer's decision-making process are plausibly mitigated with more frequent information of the same type.

In the limit, increasing the frequency of information provided to consumers (i.e., real-time feedback) should tend toward the neoclassical ideal of perfect information.This notion corresponds to a "pure nudge," in the parlance of Allcott and Taubinsky (2015), that corrects the informational failure completely. With less frequent information (e.g., periodic utility bills), however, consumers are more likely to base their consumptive decisions on imperfect information. So, if a consumer receives a more frequent signal about her consumption, she may change her behavior in such a way that aligns more closely, but not perfectly, with standard models of consumer demand. This example provides the groundwork for the conceptual setting in which water customers alter consumption in response to an increase in the frequency of bills to capture information rents. Allowing for *ex post* misoptimizing behavior in this choice setting deviates from the behavioral welfare frameworks of Chetty et al. (2009) and Allcott and Taubinsky (2015) at the cost of specifying a positive model of consumer behavior.

3.2.2 Billing frequency and price misperception

Since most analyses concerned with the effect of information provision on consumer behavior deal solely with quantity reminders, social comparisons, or prices (Ferraro and Price, 2013; Jessoe and Rapson, 2014; Kahn and Wolak, 2013), there is typically an isolated mechanism through which changes in information affect behavior. The receipt of a bill more frequently confounds clean identification of a mechanism through which consumers respond, though as argued in Section 3.2.1, this presumably allows consumers to adjust their behavior along an existing margin. Since a water utility bill contains both price and quantity information, it is not clear whether consumers are responding to more frequent price information or more frequent quantity information (or, perhaps, both).

In this subsection, I develop a model of misperceived price such that a consumer who receives more frequent information about her consumption habits may choose to consume more or less water since she has a more accurate perception of water prices for her consumption in each billing period. Finkelstein (2009), Li et al. (2014), and Sexton (2014) find that consumers tend to misperceive (changes in) prices or taxes that are not salient. Thus, increasing billing frequency is a plausible mechanism to increase the salience of water prices.

Consider a consumer with utility over water consumption (w) and a composite good (x):

$$u = x + aw^{1/\gamma+1} \tag{3.1}$$

where utility is quasilinear in x and preferences over w exhibit constant elasticity of demand. The consumer observes the budget constraint, $M = x + \tilde{p}w$, where her wealth (M) equals expenditures on x with its price normalized to unity and her perceived expenditures on water. The budget constraint is satisfied with equality since any residual consumption is allocated to the composite good. The price a consumer perceives for her consumption of w is defined as $\tilde{p} = \theta p$ where p is the true price and θ is a perception parameter that specifies the degree to which she over or underestimates the true price. A consumer with perfect information is represented by $\theta = 1$. While previous research bounds this parameter from above at unity, I allow for misperceptions to deviate above and below the true price since there is no good theoretical foundation for why an inattentive consumer would always under-perceive prices.

The water provider can manipulate θ by changing information available on an intermittent bill, changing the frequency at which consumers are billed, allowing customers to automatically deduct their bill from their checking account, and so forth. To coincide with the empirical analysis, define θ as a function of billing frequency, BF, and a vector of other fixed characteristics, Z,

$$\theta = f^P(\mathrm{BF}, Z). \tag{3.2}$$

Equation 3.2 captures the essence of the analysis that follows—changes in billing frequency affect demand by altering consumers' perception of price.

As a useful, and plausible, assumption, I restrict information to be weakly welfare improving:

Assumption 1 *More frequent billing information can never make anyone worse off.*

That information is welfare improving simply means that consumers can always choose to be ignorant, and that there are no cognitive costs to ignoring new information. There may, in fact, be cognitive costs to processing this information, but the consumer will only undertake such an action if its benefits exceed costs at the margin. This assumption implies structure for the functional relationship in Equation 3.2. For $\theta > 1$, increases in BF exhibit a negative effect on θ such that increasing billing frequency reduces the distortion between the true price and perceived price. However, I do not restrict θ to be greater than one. For $\theta < 1$, the predicted response to BF is reversed, allowing for increases in information to decrease the wedge between the true price regardless of the direction of price misperception. Assuming f^P is differentiable, we can write these effects on demand via the envelope theorem,

$$\frac{\partial^2 w(\tilde{p})}{\partial\theta\partial\mathrm{BF}} = \frac{\partial w(\tilde{p})}{\partial\mathrm{BF}} \begin{cases} > 0 & \text{if } \theta > 1 \\ = 0 & \text{if } \theta = 1 \\ < 0 & \text{if } \theta < 1 \end{cases} \cdot \tag{3.3}$$

Since θ is unknown to the researcher, it is unclear *a priori* whether a consumer perceives a price that is higher or lower than the true price p. Hence, the demand response to an increase in the frequency of billing information is ambiguous and depends on the initial degree of price misperception. Further, Equation 3.3 implies that the direction of the demand response ($\partial w/\partial\mathrm{BF}$) reveals consumers' initial misperceptions of price.

As an illustrative example, consider a consumer facing a change in the perceived price of water induced by an increase in billing frequency. In Figure 3.1, let DE be the consumer's initial misperceived budget line under the bi-monthly billing regime with a corresponding perceived price, $\tilde{p}_0$. Due to misperceptions of the price of water, she targets the consumption bundle $(x, \bar{w})$ at the beginning of the billing cycle, but consumes the bundle $(\bar{x}, \bar{w})$ since she over-perceives the price of water ($\theta > 1$) and, hence, allocates more consumption to the numeraire good. Upon updating her perception of price through more frequent billing, she learns that her true budget line is DF and chooses $(x^\star, w^\star)$ as the preferred allocation.[5] This movement corresponds to an increase in water consumption in response to a reduction in

[5]Note, however, that misperceptions need not be fully corrected in this framework.

the wedge between her initial misperception and her updated perception of price. Additionally, the preferred bundle indicates an increase in consumer welfare.

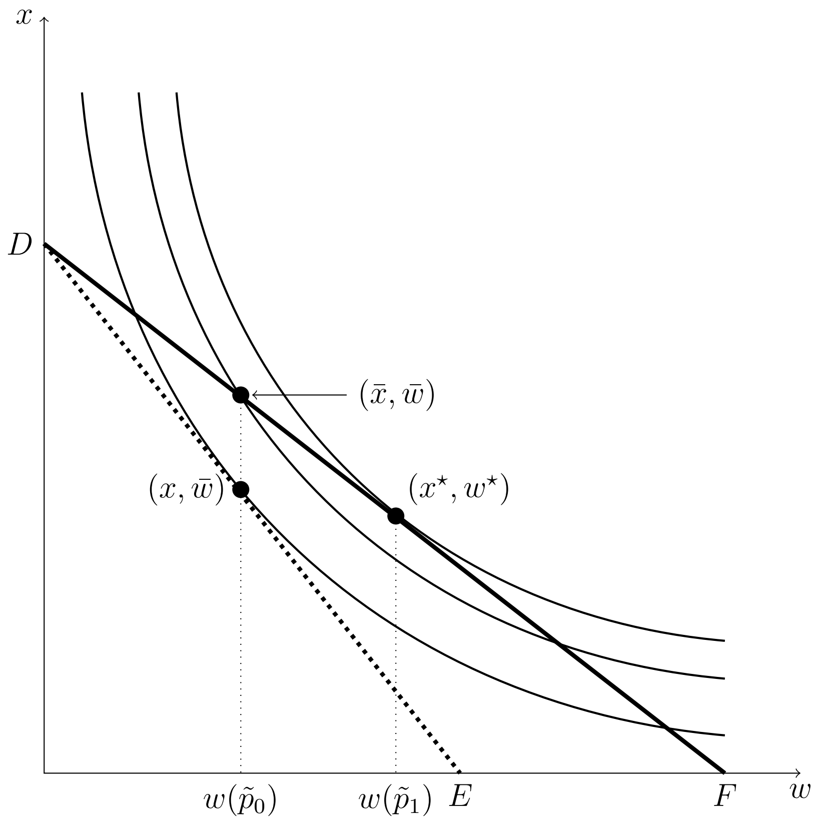

Figure 3.1: The economics of price misperception

3.2.3 Welfare effects from price misperception

Since the misperception of prices from infrequent billing drives a wedge between the actions of a perfectly informed consumer and an inattentive consumer, there are welfare gains from information provision. If θ were observable to the researcher, the change in welfare could be obtained by measuring the area under the demand curve between the implied changes in perceived price. However, since θ is an unknown parameter, I rely on behavioral welfare analyses similar to that of Allcott and Taubinsky (2015), Chetty et al. (2009), and Just (2011) to obtain sufficient statistics from observable or estimable parameters that allow for welfare analysis when consumers may not optimize perfectly.[6]

In particular, consider an increase in the frequency of information, represented by θ_0 to θ_1, that corresponds to a change in perceived price from $\tilde{p}_0$ to $\tilde{p}_1$. Assumption 1 allows us to remain agnostic about the initial misperception of price since a positive demand response to an increase in information implies that consumers display positive θ, and the converse is also true. Under the notion that perceived price tends towards the true price with an increase in billing frequency, I make the following assumption as a useful benchmark:

Assumption 2 *Under the more frequent billing regime, ex post perceived prices are proportional to the true price.*

While strong, this assumption allows the researcher to back out perceived prices from observable changes in demand. Assumption 2 is reflective of the pure nudge assertion in Allcott and Taubinsky (2015) and Chetty et al. (2009), in which the informational treatment is fully corrective; *ex post* proportionality to the true price weakens this assumption. With respect to real-time feedback through smart-metering technology, it is possible that consumers do in fact know the marginal price they are paying at any point within the billing cycle. Strong and Goemans (2014) and Kahn and Wolak (2013), for example, show that consumers tend to optimize

[6]The welfare analysis differs from other work in behavioral welfare analysis in that I specify a positive model of demand that may be driving behavior in this setting whereas Chetty et al. (2009), for example, use price and tax elasticities to identify preferences when consumers make optimization mistakes.

"better" under block rate structures when provided with real-time consumption feedback and educational treatments on how their bill is calculated, respectively. In any case, monthly billing provides an opportunity for price misperception to prevail after the change in billing frequency.

To derive an empirically tractable measure of θ, we can write the constant elasticity of demand with respect to perceived price,

$$\eta^P = \frac{\%\Delta w(\tilde{p})}{\%\Delta \tilde{p}}, \tag{3.4}$$

which holds by definition. The percent change in perceived price can be used to obtain an empirical measure of the change in the perception parameter by observing a change in quantity demanded. In particular, we can write,

$$\%\Delta \tilde{p} = \frac{\tilde{p_0}}{\tilde{p_1}} - 1 = \frac{\theta_0 p}{\theta_1 p} - 1 = \%\Delta\theta, \tag{3.5}$$

since the market price does not change. Combining this expression with Equation 3.4, rearranging, and multiplying through by $\tilde{p_1}$ provides,

$$\Delta \tilde{p} = \Delta\theta \tilde{p_1} = \left(\frac{\%\Delta w(\tilde{p})}{\eta^P} \right) \tilde{p_1}, \tag{3.6}$$

where $\tilde{p_1}$ is the *ex post* perceived price that is proportional to p, that is $\tilde{p_1} = \alpha p$ with α being the degree of *ex post* misperception via Assumption 2. Equation 3.6 states that a change in perceived price is simply a function of the market price, the perceived price elasticity, and the corresponding demand response. This expression is convenient since η^P can be estimated or inferred from other studies and the change in water consumption can be estimated using quasi-experimental techniques. Since Equation 3.6 describes the change in perceptions of price due to a change in an unobserved parameter, I utilize Assumption 2 to provide a reference point (i.e., the observable price) to obtain a price-equivalent, in dollars per unit of water consumption, of the consumer's perceived price.

Using this measure of the consumer's *ex post* perceived price, consumer surplus can be calculated by integrating the demand function between the initial price perceived, $\tilde{p_0}$, and the price perceived after the change in billing frequency, $\tilde{p_1}$,

$$\Delta CS = \int_{\tilde{p_1}}^{\tilde{p_0}} w(\tilde{p})\mathrm{d}\tilde{p} \cong -\frac{1}{2}\Delta\theta\tilde{p_1}\frac{\partial w}{\partial \mathrm{BF}}, \tag{3.7}$$

which is analogous to the Harberger (1964) triangle approximation for deadweight loss since the data in the experiment are not sufficient to estimate a true demand function.[7] Under the assumptions made so far, the treatment effect ($\partial w/\partial\mathrm{BF}$) and the perceived price elasticity of demand (η^P) serve as sufficient statistics for calculating changes in welfare due to a change in billing frequency.[8]

The welfare analysis is illustrated for a stylized example in Figure 3.2. As shown, $\Delta\theta p$ is the change in perceived price that reflects the movement in quantity demanded along a fixed demand curve. Since there is an economic cost borne by

[7]Note that welfare calculations are simplified by assuming quasilinear preferences. As such, equivalent and compensating variation are identical and estimate consumer surplus exactly (Hausman, 1981).

[8]Note, Goulder and Williams (2003) highlight the potential bias from ignoring general equilibrium effects when using this approximation, but under quasilinear utility and the fact that periodic water bills are a generally small portion of a consumer's budget, these effects are likely to be minimal.

consumers initially misperceiving the market price, the shaded area represents the approximate welfare gain from an increase in billing frequency. The change in consumer surplus depicted illustrates the equivalent variation of a change in the perceived price of water since preferences are quasilinear. That is, the amount of income the consumer would need to be given to forgo the change in perceived prices.

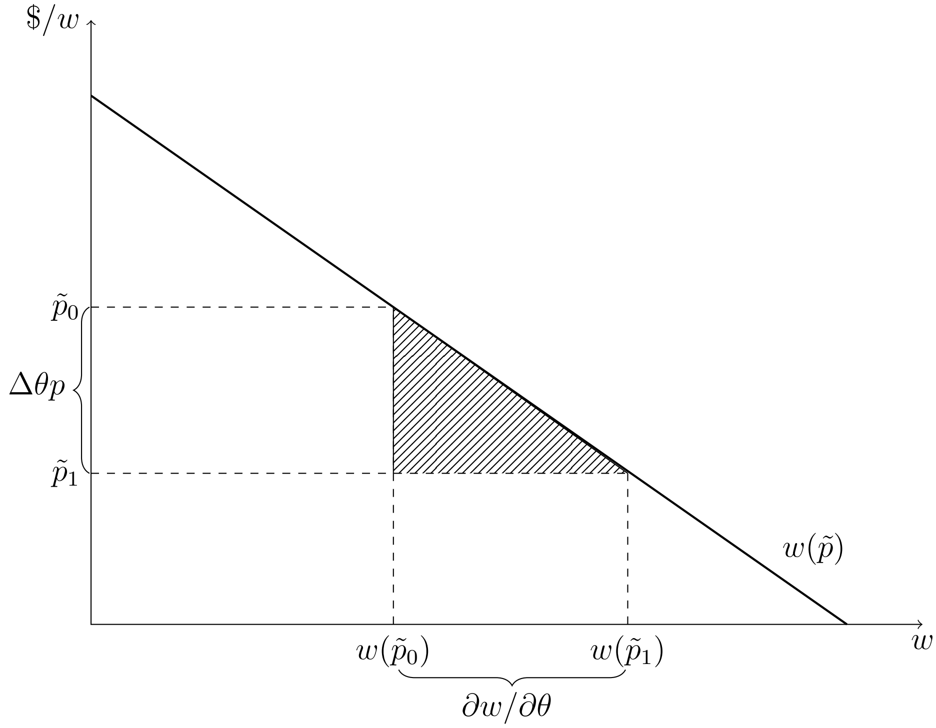

Figure 3.2: A stylized example of the welfare change from an increase in billing frequency for price misperception

3.2.4 Billing frequency and quantity uncertainty

In the previous subsection, I motivated a model in which changes in the frequency of billing induce better price perception. Since a water utility bill contains both price and quantity information, a consumer could be fully aware of the market price, but uncertain about her quantity consumed each period. This precise sort of quantity uncertainty for water consumption is documented in Strong and Goemans (2014). Thus, I develop a similar model of consumption under misperceived quantities, analogous to that of prices, though I abstract from risk preferences and uncertainty in an expected utility framework. Rather, the uncertainty considered within this mechanism is interpreted as quantity salience in the sense that a water customer has imprecise knowledge and imperfect control over her water use within a billing period. In this framework, however, a consumer who receives more frequent information about her consumption habits may alter her water use since she has a better sense of how much water, and for what purpose, she is using in each billing period.

Consider initial consumer utility provided in Equation 3.1, assuming perfect information about prices, augmented by a quantity information parameter,

$$u = x + a(\lambda w)^{1/\gamma+1} \tag{3.8}$$

where λ is a parameter that scales quantity demanded within a billing cycle similar to the perceived price parameter introduced in the previous subsection. In particular, define this term as a function of billing frequency and a vector of other fixed attributes,

$$\lambda = f^{Q}(\mathrm{BF}, Z). \tag{3.9}$$

In this framework, a consumer maximizes Equation 3.8 subject to a budget constraint, $M = x + \lambda pw$. In contrast to imperfect price perception, imperfect quantity perception affects consumer utility by scaling consumption by λ in both the utility function as well as the budget constraint through its effect on anticipated water expenditures. After substitution, the consumer's objective function is written

$$\max_{w}\{M + a(\lambda w)^{1/\gamma+1} - \lambda pw\} \tag{3.10}$$

where first-order conditions imply that *perceived consumer demand*, $\tilde{w}(p)$, can be represented, after some algebra, by

$$\tilde{w}(p) = Ap^{\frac{1-\gamma}{\gamma}}\lambda^{\frac{1-2\gamma}{\gamma}} = \lambda^{\eta^Q-1}w(p) \tag{3.11}$$

where $A \equiv \left(\frac{1+\gamma}{a}\right)^{\frac{1-\gamma}{\gamma}}$, $\eta^Q = (1-\gamma)/\gamma$ is the price elasticity of perceived demand, and $w(p)$ is consumer demand under perfect information. A derivation of demand functions under quantity misperception is presented in Appendix C. Equation 3.11 illustrates that demand is scaled conveniently by the information parameter.

By Assumption 1, we can write the marginal effects of billing frequency on *perceived demand*, $\tilde{w}(p)$, similarly to that of the price parameter,

$$\frac{\partial^2\tilde{w}(p)}{\partial\lambda\partial\text{BF}} = \frac{\partial\tilde{w}(p)}{\partial\text{BF}}\begin{cases} > 0 & \text{if } \lambda > 1 \\ = 0 & \text{if } \lambda = 1 \\ < 0 & \text{if } \lambda < 1 \end{cases} \tag{3.12}$$

which states that increasing quantity information allows a consumer to predict consumption closer to her true consumption. The implication here, then, is that the marginal effect of changes in billing frequency on consumer demand depends critically on the initial perception of quantities. Intuitively, this result implies that increasing the precision of the information with which consumers make decisions decreases the wedge between perceived quantity and the actual quantity consumed.

3.2.5 Welfare effects from quantity uncertainty

Similar to price misperception, quantity misperception allows for a divergence in the behavior of a consumer with perfect information and a consumer who misperceives her consumption. Thus, any policy that decreases the wedge between these two types of consumers will provide welfare gains to the consumer. An analytical calculation for this welfare change is obtained in a similar fashion to that of price misperception, though it is not necessary to estimate the change in λ empirically. Since λ scales consumption multiplicatively, and we observe the demand response directly in the experiment, all of the information necessary to calculate welfare is revealed through observable consumption. Effectively, we observe the movement in quantity demanded, trace out the demand function for a given price elasticity of perceived demand, and recover the prices that reflect different points of consumption along the demand curve.

Within the experiment, we observe a change in information that moves a customer from a bi-monthly billing regime (represented by λ_0) to a monthly billing regime (λ_1). To provide a calculable estimate of welfare changes from these changes in information, we can write the price elasticity of perceived demand as a function of the demand response and market prices,

$$\eta^Q = \frac{\%\Delta\tilde{w}(p,\lambda)}{\%\Delta p} \tag{3.13}$$

and, since we do not observe price changes directly, we can infer them by rearranging Equation 3.13 and multiplying through by p,

$$\Delta p = \left(\frac{\%\Delta\tilde{w}(p,\lambda)}{\eta^Q}\right)p. \tag{3.14}$$

Since the change in quantity demanded is observed in the experiment and η can be estimated or inferred from other studies, Equation 3.14 provides a price change that corresponds to the observed change in billing frequency under quantity misperception. Within this framework, we can use the observable demand response to changes in billing frequency ($\partial\tilde{w}/\partial\mathrm{BF}$) and the price elasticity (η) to calculate consumer surplus in response to the change in billing frequency,

$$\Delta CS^Q = \int_{p_1}^{p_0} \tilde{w}(p)\mathrm{d}p = \lambda^{\eta^Q-1}\int_{p_1}^{p_0} w(p)\mathrm{d}p \cong -\frac{1}{2}\Delta p\frac{\partial\tilde{w}}{\partial\mathrm{BF}}. \tag{3.15}$$

Thus, Equation 3.15 provides an analytical formula from which we can use quasi-experimental estimates to calculate changes in economic welfare.

3.2.6 Reconciling price and quantity misperception

The previous two subsections outlined a simple conceptual approach to calculating welfare changes to consumers under two different frameworks—price misperception and quantity uncertainty. Both frameworks culminate in analytic formulas for measuring welfare from changes in information provision that are estimable or observable. The key ingredients are the demand response to changes in billing frequency, which can be obtained using program evaluation methods, and perceived price elasticities, which can be inferred from other studies if the empirical framework lacks sufficient variation to estimate structural parameters of demand.[9] Conditional on the assumptions made so far, the framework is general enough to accommodate both positive and negative changes in demand in response to increases in information and maps to corresponding changes in (perceived) price.

The primary conceptual contribution thus far is that if consumers over (under) perceive prices, then an increase in information provision will increase (decrease) quantity consumed. Intuitively, the provision of more frequent information allows consumers to mitigate uncertainty in their perception of prices or quantity consumed within a billing period. Thus, for consumers who misperceive prices or quantities, there are welfare gains to increasing the frequency of price and quantity information. The welfare change is predicated on the notion that more frequent information brings consumers closer to the neoclassical ideal of decision-making under perfect information.

[9] Note that *perceived* price elasticities in the price or quantity misperception frameworks are not necessarily equivalent to true price elasticities.

To consider whether the welfare gains of information provision are larger for consumers who respond to information through a quantity mechanism or a price mechanism, it is easy to compare the analytical welfare functions for both types. Recall, ΔCS^P is the welfare change for consumers who misperceive prices,

$$\Delta CS^P \cong -\frac{1}{2}\Delta\theta p \frac{\partial w}{\partial \mathrm{BF}}, \tag{3.16}$$

and ΔCS^Q is the analog for quantity misperception,

$$\Delta CS^Q \cong -\frac{1}{2}\Delta p \frac{\partial \tilde{w}}{\partial \mathrm{BF}}. \tag{3.17}$$

In both formulas, the demand response to the change in information is estimated empirically. This implies that for some fixed demand response to a change in information, $\partial w/\partial\theta$ must equal $\partial\tilde{w}/\partial\lambda$ regardless of whether consumers misperceive prices or quantities. Further, $\Delta\theta p$ under price misperception is equal to $(\%\Delta w(p,\theta)/\eta^P)p$ by Equation 3.6. Whereas, Δp under quantity misperception is equal to $(\%\Delta\tilde{w}(p,\lambda)/\eta^Q)p$ by Equation 3.14. For a common price elasticity, $\eta^Q = \eta^P$, the welfare changes are equivalent since the percent change in demand is the same in both scenarios. Thus, under these conditions, the mechanism through which consumers respond to changes in information provision is immaterial for the calculation of welfare.

However, this result depends critically on Assumption 2, in which consumers respond to prices perfectly in the more frequent information regime. But, if I relax this assumption to allow for consumers to misperceive higher (lower) prices after the increase in information, the welfare effects under price misperception increase (decrease) monotonically with the degree to which prices are misperceived. As an example, consider a consumer who perceives a price 10% greater than the market price in the new billing regime. For a fixed change in consumption, this consumer is effectively located further leftward on the demand curve, where its slope is steeper. So, for the same demand response, the change in perceived prices ($\Delta\theta p$) will be greater than that of a consumer who perceives prices perfectly in the new billing regime. In contrast, the welfare effects of quantity misperception are invariant to the scale of misperception, conditional on perfect price information, for an observed demand response. Since we observe a revealed preference measurement of demand before and after the change in information, price certainty allows us to pin down the demand function explicitly, regardless of the degree of quantity misperception either before or after the provision of new information. As such, assuming consumers respond to information through a quantity mechanism bounds welfare estimates from below (above) if prices are initially over- (under-) perceived, conditional on common price elasticities.

Lastly, it is important to consider that a consumer may respond to increases in information through a joint price and quantity mechanism, or that a sample of households my vary in the mechanism through which they respond to changes in information. It may also be the case that households misperceive prices and quantities in opposite directions. For all of these scenarios, however, the quantity misperception welfare estimates will remain a lower (upper) bound depending on the aggregate demand response. As an example of the latter, let a single consumer under-perceive quantity (i.e., $\lambda < 1$), but over-perceive prices (i.e., $\theta > 1$). The prediction would be that the quantity misperception would decrease consumption upon receipt of more frequent information, while the prediction for price misperception would work in the opposite direction. For a fixed demand response, we observe the net effect of these competing responses. Since we cannot identify which effect is driving the results without additional information, the best approach would be to use the most conservative estimate of welfare, which would assume that the consumer responds to new information only through quantity misperception if the net effect is positive. A similar argument holds for price and quantity misperception that move in the same direction, as well as an aggregate response for heterogeneous populations.

3.2.7 Alternative mechanisms

While the discussion thus far considers only price and/or quantity uncertainty in intermittent choice settings, it is possible that a number of other mechanisms could influence consumer behavior in response to a change in billing frequency.

Since consumers are receiving more frequent bills, it could be the case that budgeting along a fixed time horizon might influence consumption directly. In Appendix A, I present a thought-experiment in a classical demand framework to highlight the differential effects from household budgeting for utility expenditures (e.g., water, electric, and natural gas) across two different time horizons. In this example, a consumer who budgets on a monthly basis (as opposed to an annual basis) may be more sensitive to changes in the frequency of information simply because expenditure shares change upon receipt of a bill in a month-to-month context. This framework requires consumers to be sufficiently short-sighted such that fluctuations in monthly water bills are not smoothed over a longer time horizon.

Additionally, if consumers are risk averse or credit constrained, then variation in billing frequency could alter their behavior in a way that aligns with my empirical findings. Under a risk aversion framework, consumers might use relatively less water in a longer billing period for fear of exceeding their level of budgeted consumption. With more frequent billing, uncertainty in a consumer's expected payment is reduced since she has more precise control over their water usage. Hence, an individual consumer may be less inclined to conserve water. On the other hand, a shorter billing regime might allow for absolute deviations from an expected bill to be proportionally larger than in a longer billing period. In this case, a consumer's aversion to bill "shock" could affect consumption planning and budgeting negatively. Thus, a general model of consumer behavior in this vein could explain the causal increase in consumption in response to more frequent billing information, but does not provide unambiguous theoretical predictions.

Alternatively, since consumers, in general, may face nonlinear pricing structures, risk neutral preferences could produce effects consistent with my empirical results. In the experiment considered, consumers face increasing block rates. With more precise control and knowledge over their consumption in a more frequent billing regime, consumers may be more prone to optimize in accordance with the block rate schedule, which encourages conservation. The effect of block rates on behavior in response to changes in billing frequency is explored empirically, to some degree, in the following section.

Overall, the price and quantity misperception mechanisms provide a plausible story for why consumers may alter behavior in response to more frequent information. That the price, or quantity, of water consumption becomes more salient with more frequent billing aligns well with recent literature (Chetty et al., 2009; Finkelstein, 2009; Gilbert and Graff Zivin, 2014; Sexton, 2014), but it is not the only explanation consistent with the results that I document in the following section.

3.3 Empirical setting

Beginning in December 2011, the City of Durham's Department of Water Management in North Carolina (henceforth, "Durham") transitioned individual billing districts from bi-monthly to monthly billing at different points in time. Primary reasons for the transition include cost-saving from fewer delinquent payments, early leak detection, improving customer service, and reducing administrative costs. In addition, the change in billing frequency was enabled by district-wide installation of automated meter reading devices. The new meters allow for consumption levels to be obtained via radio frequency such that the costs to read meters manually were reduced. Meters were installed for each billing district and, once the installations were completed, the entire district was transitioned to monthly billing. Customers were notified of the transition to monthly billing by mail approximately six weeks before the transition.[10]

[10]A copy of the mailer distributed to customers is included as Figure B.2 in Appendix B.

To make a cost-saving argument to the city council, the water utility used a single billing district as a pilot group to measure changes in administrative costs before and after the transition to monthly billing cycles. After that, billing districts were transitioned to monthly billing according to meter installation and administrative schedules. The order of districts for meter installation (and, subsequently, monthly billing) was chosen to work around billing cycles and other feasibility constraints. According to utility officials, no consideration of billing history, income base of neighborhood, or any other financial indicator was taken into account when choosing which districts to transition.[11]

Given these details, the assignment of monthly billing is plausibly exogenous to the household, conditional on residing within a particular billing district. The household has no ability to manipulate the assignment of billing frequency short of moving across billing cycle boundaries. Within the study period, 12 of 17 billing districts were transitioned according to the timing in Table 3.1. The first district transitioned received their first monthly bill on December 1st, 2011. Figure 3.3 presents a map of the first six billing districts to transition to monthly billing (in Panel A) and all districts transitioned to monthly billing by June 2014 (in Panel B). The entire service area is represented by the union of all billing districts outlined in bold. In Figure 3.4, I present a magnified view of billing district boundaries within neighborhoods. This figure illustrates that the district boundaries are designated in such a manner that neighbors could be consuming water concurrently, but may be billed at different frequencies. Thus, this design allows for the exploitation of

[11] I consider selection issues explicitly in Section 3.5.2.

geography to minimize the concern that differences in neighborhood characteristics might bias results.

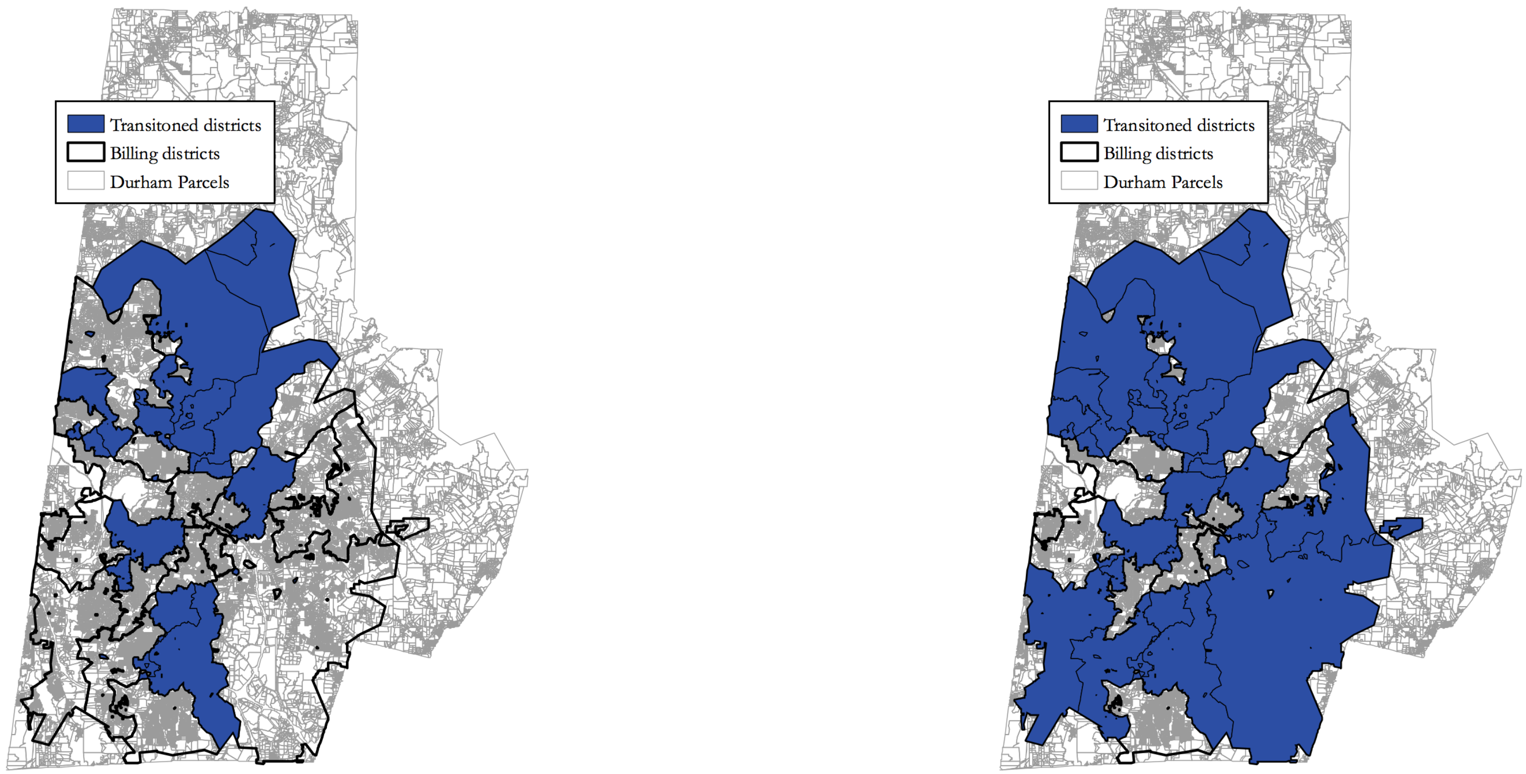

(a) First six billing districts transitioned to monthly billing (before February 2013)

(b) All billing districts transitioned to monthly billing by June 2014

Figure 3.3: Billing districts transitioned to monthly billing over time

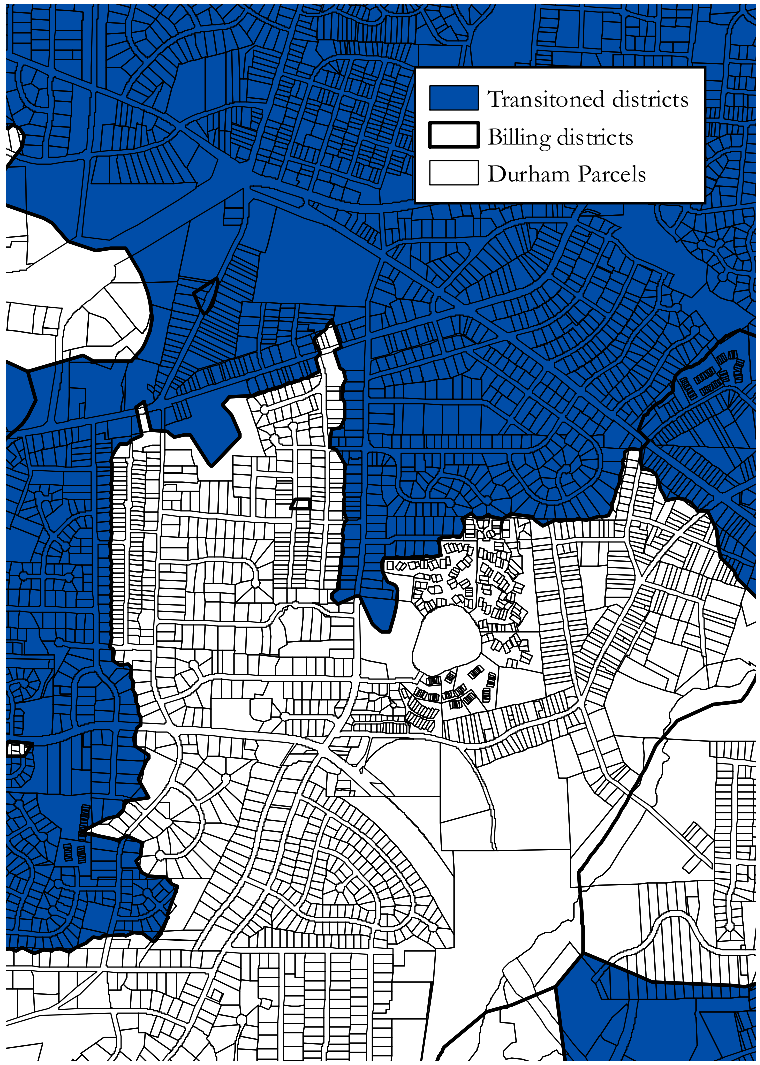

Figure 3.4: Depiction of billing district boundaries within neighborhoods

3.3.1 Data

The primary data used in this analysis are residential billing records for Durham water customers. Included in these data are (bi-)monthly water and sewer use, fixed service fees and volumetric consumption rates, the address of the customer, billing district, and whether a customer has their water bill automatically deducted from a bank account. The billing data were matched by address with geocoded tax assessor data, containing structural characteristics of the home, obtained from Durham County. Each matched residential address was spatially linked to its 2010 Census block as well as billing district polygons provided by Durham. For each household, I determine the nearest billing district, as well as the linear distance from the centroid of the tax parcel its nearest district boundary. Key demographic variables from the 2010 SF1 Census are matched to each household's Census block. Residential premises that changed water billing accounts within the timeframe of the study are removed from the sample—this strategy reduces the impact of renters, who may not pay water bills explicitly. Further, this avoids econometric identification problems when relying on variation within a household over time.[12]

The final sample consists of roughly 59,000 individual household accounts with water bills from February 2009 through June 2014, which implies slightly less than 1.7 million household-by-bi-monthly unique observations. Summary statistics for variables of interest are presented in Table 3.1. The first four columns decompose household characteristics and details on water use by the year in which households transitioned from bi-monthly to monthly billing. Summary statistics in the final column are for the entire sample. Each of the treatment waves are relatively similar across demographic, water use, and housing characteristics, and similar to the sam-

[12]Renters may cause problems for identification if they do not receive a water utility bill. However, this effect would tend to pull any estimated treatment effect towards zero so long as renters did not change their behavior at exact time of the change in billing frequency.

ple mean, with the exception of households that transitioned in 2013 along several dimensions. For this group, home value (a proxy for wealth) is notably larger than that of all other groups. Further, these households tend to have larger homes on larger lots, and are more likely to be located in a Census block with fewer renters and a higher proportion of white residents. Water consumption, however, is statistically similar across all groups. For the typical household in the sample, the mean assessed home value is approximately $186,000 with a standard deviation of $126,000. The average home is on one-third of an acre, 34 years old, roughly 1,800 square feet, with three bedrooms. Within the final sample, households reside in Census blocks in which approximately one-quarter of all homes are renter-occupied. 53% of the sample is white and the average household size is between two and three people. Average bi-monthly water bills for all time periods in the sample are $85 for consumption of 985 cubic feet of water.

Further, I include weather covariates obtained from the North Carolina State Climate Office. The key variables used are mean maximum temperature for a 60-day rolling window that is backwards-looking from the date each individual bill was mailed. The sum of rainfall (in inches) for the same 60-day time window is also calculated.

A final caveat is that under bi-monthly billing, water bills are mailed on a staggered schedule that smooths administrative work and meter reading throughout the year. As an example, a billing district on the odd cycle may receive a bi-monthly bill in March for consumption in January and February. Contrarily, a billing district on the even cycle would receive a bill in April for consumption in February and

Table 3.1: Demographic and water use characteristics among households that transitioned to monthly billing at different points in time

	Summary statistics for households that received first monthly bill in:				
	2011-2012	2013	2014	Never	**Total**
Tax assessor records:					
Assessed value of home	161,248	226,557	179,055	169,726	185,998
	(103,633)	(162,985)	(80,776)	(123,219)	(126,453)
Lot size (acres)	0.32	0.39	0.25	0.31	0.32
	(0.43)	(0.53)	(0.41)	(0.31)	(0.43)
Age of home	33.81	29.72	28.73	44.88	34.47
(years since 2014)	(22.85)	(19.98)	(24.63)	(28.02)	(25.1)
Size of home (square feet)	1639.5	2005.2	1777.6	1708.9	1793.3
	(766.08)	(893.5)	(641.24)	(780.04)	(808.98)
Number of bedrooms	3.02	3.25	3.11	3.04	3.10
	(0.73)	(0.76)	(0.72)	(0.80)	(0.78)
2010 Census (block):					
Percent renters	0.27	0.17	0.23	0.33	0.25
	(0.23)	(0.20)	(0.23)	(0.28)	(0.24)
Percent white	0.46	0.61	0.50	0.47	0.53
	(0.28)	(0.31)	(0.29)	(0.37)	(0.31)
Household size	2.52	2.50	2.51	2.48	2.48
	(0.48)	(0.46)	(0.52)	(0.54)	(0.51)
Billing records:					
Total bi-monthly water	81.96	91.21	89.02	84.13	84.62
bill ($/ccf)	(36.54)	(44.76)	(41.99)	(38.91)	(39.67)
Full sample bi-monthly	977.46	1033.81	986.89	978.50	985.24
water use (cf)	(520.35)	(551.99)	(489.66)	(542.51)	(528.03)
2009-2010 bi-monthly	997.83	1066.01	1004.7	1014.96	1018.17
water use (cf)	(600.74)	(651.73)	(578.80)	(635.39)	(622.24)
No. households:	18,042	15,415	10,589	14,215	58,965
No. billing districts:	5	4	3	5	17
First monthly bill:	12/1/11 7/13/12 10/25/12 11/14/12 12/29/12	1/29/13 2/12/13 3/30/13 11/22/13	1/30/14 4/18/14 5/15/14		

Note: Means and standard deviations (in parentheses) are presented. The first billing district to transition to monthly billing occurred on December 1, 2011, so this district is grouped jointly with districts that transitioned in 2012. The 2010 Census (SF1) data is assigned to the Census block in which the household resides. 2009-2010 bi-monthly water use is used to provide a sense of average consumption among each group prior to the transition to monthly billing (2009-2010 refers to consumption that occurred in the full calendar years of 2009 and 2010).

March. Rather than dealing with these two groups independently, I pool households into two-month cycles corresponding to the date in which bills are received, but allow for each district to retain accurate measures of weather fluctuations within their use period. As such, there are 32 distinct time periods in the study that correspond to two-month windows between February 2009 and June 2014. Additionally, since monthly bills account for consumption during a shorter duration than bi-monthly bills by construction, monthly consumption is aggregated to a bi-monthly level. So, the unit of observation for consumption is a two-month period regardless of whether households are being billed monthly or bi-monthly.

3.3.2 Prices

The transition of households to monthly billing provides a unique natural experiment to identify a causal effect of more frequent information on consumer behavior so long as other factors are not changing at the same time. When households were transitioned from bi-monthly to monthly billing, fixed water and sewer service fees were cut in half as well as block cut-offs in the tiered rate structure. Marginal volumetric rates for consumption remained constant across billing frequencies. Figure 3.5 illustrates the change in the rate structure for monthly and bi-monthly billing. The solid line is the increasing block rate structure used to calculate bi-monthly bills, while the dotted line is used to calculate monthly bills for the 2012-2013 fiscal year. As shown, the marginal prices for consumption do not change between monthly and bi-monthly rate structures, but the quantity blocks for consumption are halved for each price tier.

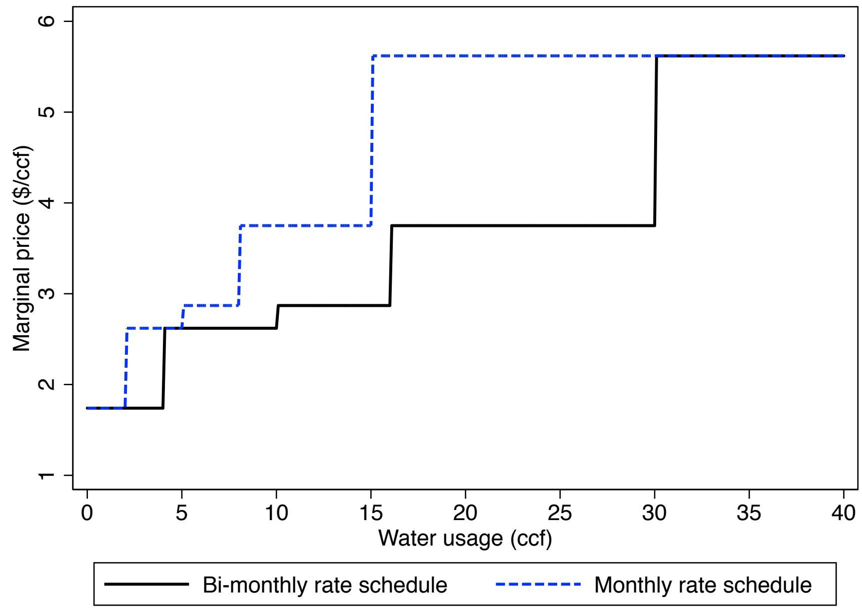

Figure 3.5: Increasing block rate structure before and after transition for monthly and bi-monthly billing

This structure was adopted to ensure that customers transitioned to monthly billing were charged at the same rate as bi-monthly customers. Thus, for the same level of consumption, two monthly bills are equivalent to one bi-monthly bill in dollar amounts. While this is a mechanical interpretation of the notion that prices did not change, the change in the block endpoints could affect consumer behavior. It is not clear, however, in what direction this change might bias results, but the extent to which this bias exists depends on whether consumers know and use the tiered rate information to make decisions. Since the water utility bill includes no information about the block rate structure (see Figure B.1), it is unlikely that consumers are responding to changes in the rate structure itself. As further evidence, I present the empirical density of consumption in Figure 3.6. In this figure, there is no evidence of bunching at the block rate cut-offs for consumption in the calendar year prior to treatment. Further, Wichman (2014) shows that water customers, and Ito (2014) for electricity customers, exhibit behavior that corresponds to changes in average price, or the total bill, when facing increasing block rates, which eases this concern.

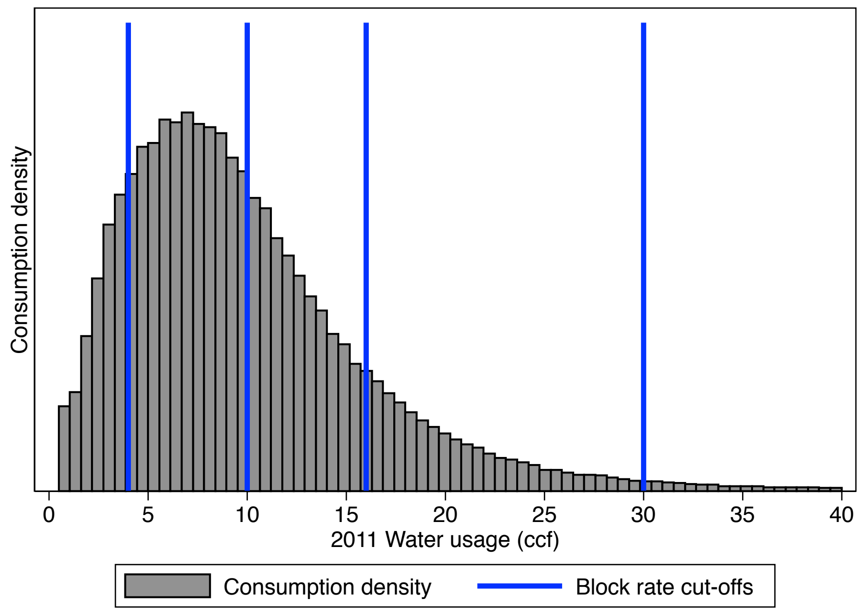

Figure 3.6: Empirical density of bi-monthly water consumption with block rate cut-offs

3.3.3 Technical efficiency of new meters

The primary threat to the unconfoundedness assumption in the empirical strategy is that billing districts were transitioned to monthly billing directly after new meters were installed. It is generally accepted in the water utility industry that as meters age, they fail to register all water that passes through. The typical advertised increase in technical efficiency is around the magnitude of 1-2%. Indeed, metering companies market the efficiency of new meters as a means to capture "non-revenue" water that utilities are treating and distributing, but is not showing up as billed consumption. Given this, it is possible that a switch to new meters may increase water "consumption" mechanically through the adoption of a more technically efficient metering system. Through a suite of empirical tests and robustness checks, I find no evidence that the empirical results are driven by technical efficiency improvements.

3.4 Empirical strategy

The empirical approach I take in this paper identifies demand responses to an increase in billing frequency using quasi-experimental techniques. I regard the transition from bi-monthly to monthly billing as the treatment, whereas households that, at any point in time, are billed on a bi-monthly basis serve as controls.

To estimate the average treatment effect (ATE) of a change in billing frequency empirically, the following log-linear equation is specified,

$$\ln(w_{ijt}) = \alpha + \beta \mathrm{BF}_{jt} + C_t\omega + Z_{ijt}\gamma + \tau_t + \epsilon_{ijt}, \tag{3.18}$$

where w_{ijt} is household i's water consumption in time t in billing district j.[13] BF_{jt} is a dummy variable equal to one if the billing district j is billed monthly at time t and zero otherwise. C_t is a vector of weather variables including mean maximum

[13] In all regression specifications with logged dependent variables of water consumption, I add 1 to bi-monthly water use since there are a small number of true zero use observations in the data set. All models were also estimated with consumption levels as the dependent variable and the results are unaffected. These estimates are available upon request.

temperature (degrees Fahrenheit) and total rainfall (inches) for each billing period. Z_{ijt} is a vector of household and demographic characteristics described in Table 3.1. τ_t is a vector of time controls including some combination of time fixed effects, a linear time trend, and seasonal indicators. ϵ_{ijt} is the residual error term.

Treatment status is assigned to households if they reside in a billing district that is billed monthly at time t. In this specification, β will capture a causal effect of the change in billing frequency on water consumption, conditional on common trends between treatment and control, as well as standard exogeneity assumptions. Since the assignment of treatment is plausibly exogenous, it is worth noting the lack of potential selection bias. Selection into (or out of) treatment is not likely to occur since it would require households to move premises with the intention of sorting along billing district boundaries, which are not publicly observable. Additionally, all households are treated with monthly billing eventually, so even if there was an incentive to move premises to seek (or avoid) monthly billing, a sophisticated consumer would know that the benefit is unlikely to remain for long. Under these assumptions and common trends in pre-treatment water use between treatment and control, β is an unbiased estimate of the ATE and is equivalent to the average treatment effect on the treated (in the absence of selection).

To analyze common trends, I first note that treatment and control households reside in a small geographic area, and sometimes within the same neighborhood (see Figure 3.3, for example). So, there is no reason, a priori, to think that consumption would exhibit different trends, or even different levels, since all households are exposed to common weather, annual rate increases, and other exogenous utility-specific shocks. Further, most control households eventually become treatment households within the study, so households that are treated later in the study serve as controls for households treated earlier in the study. Regardless, I explore mean consumption levels graphically for all households that eventually transitioned to monthly billing relative to households who never transitioned to monthly billing within the study. Figure 3.7 compares mean usage for both treatment and control groups in each time period. The control households (solid trend line) track the treated households (dashed trend line) closely before the first monthly bill represented by the vertical line in December 2011. There are slight deviations from this trend in the peak summer months prior to December 2011, but nothing extreme enough to invalidate common trends.

While a simple linear regression approach presented in Equation 3.18 provides a consistent estimate of the ATE, it is possible that there are unobservable factors within households that influence water consumption. Preferences for the environment or water-intensive durable goods are likely candidates for variables that may affect a household's response to the frequency of water bills, but the available data do not permit for direct controls. As such, I estimate Equation 3.18 with household-specific intercepts that absorb time-invariant unobservable characteristics,

$$\ln\left(w_{ijt}\right) = \alpha_i + \beta \mathrm{BF}_{jt} + C_t\omega + \tau_t + \epsilon_{ijt} \tag{3.19}$$

where α_i is now a household-specific fixed effect. The identifying assumptions for β to remain a causal effect of billing frequency remains the same as before, though

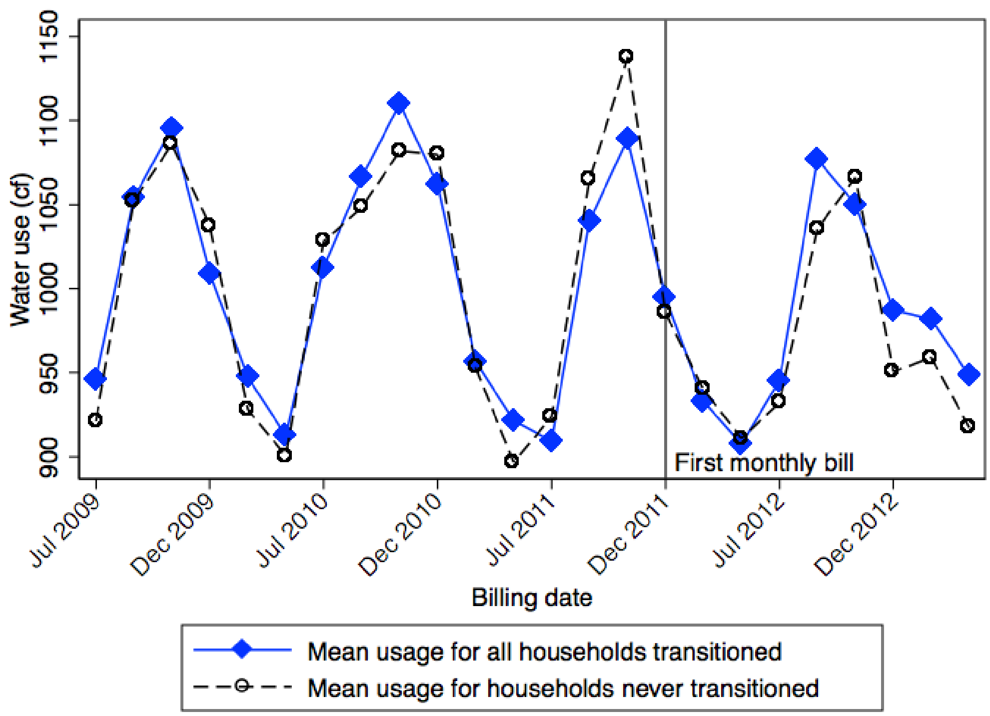

Figure 3.7: Mean bi-monthly consumption over time for households that transitioned to monthly billing and households that never transitioned to monthly billing

we now require strict exogeneity and conditional independence of the household-specific intercepts. The latter assumption indicates that treatment status cannot be correlated with an unobservable, time-varying factor at the household-level. Exogenous assignment of monthly billing across billing districts implies that treatment assignment did not target household characteristics explicitly.

3.4.1 Exploiting billing district boundaries

To assess the robustness of the difference-in-difference estimates, I estimate treatment effects within a narrow window on either side of the billing district boundaries. I consider households within 2000, 1000, and 500 feet of district boundaries.

I first estimate Equation 3.18 in a pooled cross-sectional framework since differencing out unobservables is unnecessary for identification of a local average treatment effect (LATE) in a regression discontinuity framework (Lee and Lemieux, 2010). However, I also estimate Equation 3.19 in this framework since there may be unobservables biasing results in the pooled cross-section. The latter specification, dubbed "difference-in-discontinuity," takes account of the differences over time both within households and across treatment/control groups, while simultaneously limiting the set of included household observations based on distance from the nearest billing district boundaries (Grembi et al., 2012). The intuition behind the identification of a LATE in this scenario is that as one approaches the billing district boundary households become more similar, thus avoiding potential confounding factors in the difference-in-difference framework. Additionally, since there are control households on either side of a given district threshold at some point in time, the LATE identifies the relative difference of a treated household in the same "neighborhood" as a control household over time. The main benefit of this approach is that it is robust to changes in unobservable neighborhood characteristics over time.

To explore the appropriateness of these methods, I plot mean consumption for 2012 in 40 foot bins within 1000 feet of billing district boundaries. In Figure 3.8, treatment households, represented by solid diamonds on the right-hand side of each figure are the households in the first district to transition to monthly billing. Control households, represented by hollow circles on the left-hand side, are all other households. As shown, there is a clear deviation in the quadratic trend at the billing district boundary. While the control households display a more uniform, linear trend, treatment households have a larger variance and their consumption is generally increasing with the distance from the district boundary. The latter effect is likely driven by the initial treatment households residing in billing districts further from the city center and, thus, located on larger lots with higher demands for outdoor water (see Figure 3.3 for treated districts in the northern portion of the billing district, which is more rural than the central region). This differential use among treatment households is not problematic since the difference-in-discontinuity estimator is a consistent estimator of a LATE as the distance from the district threshold decreases in the limit and, as shown, the trends near the boundary threshold are relatively similar (Grembi et al., 2012).

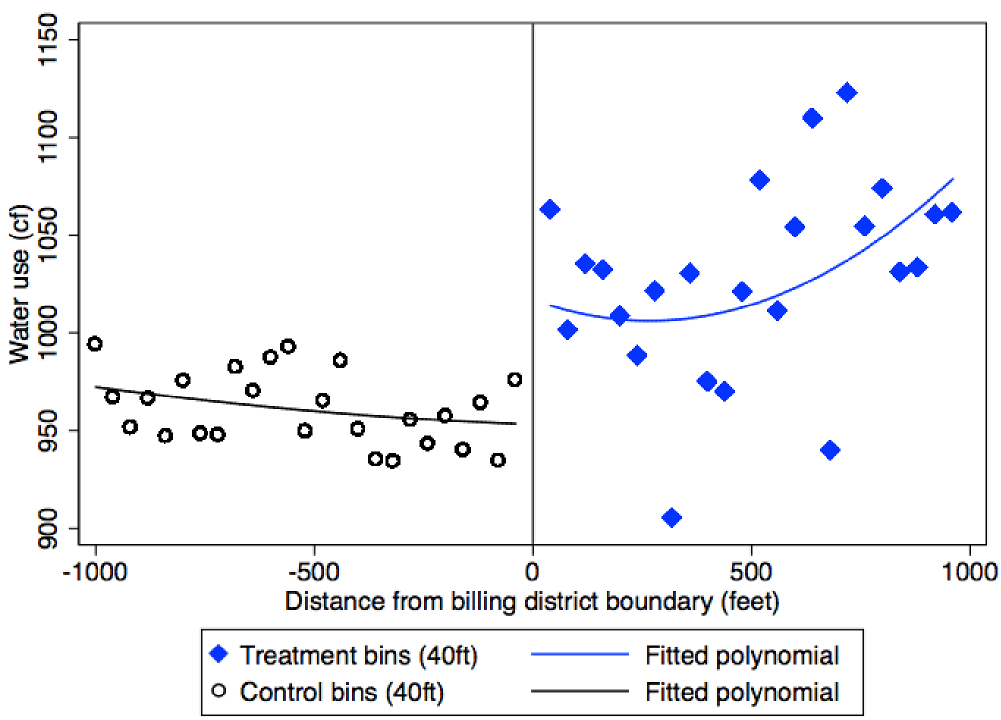

Figure 3.8: Mean consumption in 40-foot bins as a function of distance from district boundaries for consumption during 2012

3.4.2 Dynamic models of adjustment

Since it is possible that any treatment effect observed is simply an initial response to more frequent information, but households revert back to initial levels of consumption, I specify dynamic models that include a lagged dependent variable as a covariate to control for inertia, that could include habit formation, in the consumption process. Similar to Gilbert and Graff Zivin (2014) the general estimating equation takes the form,

$$\ln(w_{ijt}) = \alpha + \beta \mathrm{BF}_{jt} + \zeta \ln(w_{ijt-1}) + C_t\omega + Z_{ijt}\gamma + \tau_t + \epsilon_{ijt} \tag{3.20}$$

where, in an OLS framework, β remains the causal effect of interest in the short-run, which is augmented by an adjustment parameter, ζ. A long-run treatment effect is obtained from the estimates of Equation 3.20, $\beta^\star = \hat{\beta}/(1-\hat{\zeta})$. Additionally, the short-run ATE obtained from estimation of Equation 3.20 provides a lower bound for the treatment effect based on the conditional independence assumption that states that past water usage is not correlated with treatment. The ATE estimated in Equation 3.19 and the ATE estimated in Equation 3.20 bound the treatment effect from above and below, respectively, which provides a useful robustness check of the true effect under different modeling assumptions (Angrist and Pischke, 2009).

Further, since I rely on differencing out time-invariant unobservables for identification of a causal effect with correlated observations within a household, I estimate the panel analog of Equation 3.20. Consistent estimation of this model requires serial independence in the error terms for a given household, which is violated with certainty when using differencing methods to drop individual-specific intercepts (Nickell, 1981). However, I estimate this model for comparison with the OLS specification.

3.4.3 Heterogeneous responses to information provision

I conclude the empirical analysis with an examination of heterogeneity in the estimated treatment effects. In particular, I focus on heterogeneity arising from subpopulations that may respond differentially to more frequent information. Formally, I estimate conditional average treatment effects (CATE) similar to Allcott (2011), Ferraro and Miranda (2013), and Abrevaya et al. (2014) by interacting covariates of interest with the treatment indicator. The CATE specification takes the form,

$$\ln(w_{ijt}) = \alpha_i + \beta \mathrm{BF}_{jt} + \sum_{s \in S} \beta_s (\mathrm{BF}_{jt} \times 1[Z_{ijt}]_s) + C_t \omega + \tau_t + \epsilon_{ijt} \tag{3.21}$$

where β_s captures the conditional average treatment effect for a series of S discretized variables in vector Z_{ijt}, and all else is the same as in Equation 3.19. Specifically, I consider indicators for quintiles of each of the following covariates: pre-treatment mean (summer) consumption in 2009 and 2010, assessed home value as a proxy for wealth (Ferraro and Miranda, 2013), lot size as a proxy for irrigation intensity (Mansur and Olmstead, 2012; Renwick and Green, 2000), square footage of the home, age of the home, as well as the number of bathrooms within a household.[14]

[14]Note, I do not consider quantile treatment effects since I rely on household-specific intercepts to control for time-invariant unobservables. Within a fixed effect framework, quantile regression

Additionally, to expunge any source of endogeneity or mean reversion in the CATE estimation using (2009 and 2010) consumption patterns as conditioning covariates, I restrict the sample to observations beginning in January 2011 for models that explore heterogeneity based on pre-treatment water use.

The conditioning variables chosen in Equation 3.21 either represent key drivers of residential water demand or serve as a proxy for preferences that may influence demand responses. As such, evidence of heterogeneous treatment effects along established margins allow for probing the mechanism through which consumers may be responding.

In addition, I interact the treatment indicator with dummy variables for each season. Since informative interventions are largely used to induce conservation in environmental and resource policy, outdoor water use in the summer months is a key target for reductions in consumption. Finally, I explore the responsiveness of households who are enrolled in automatic bill payment (ABP) throughout the transition to monthly billing. These households provide two useful functions. First, under the notion that ABP customers are less attentive to changes in prices and, correspondingly, billing frequency (Sexton, 2014), a null treatment effect would instill confidence that the empirical results are identifying a behavioral response to the information treatment. Secondly, if these households do not display a response to the treatment, then it provides additional evidence that the ATE is not contaminated by an improvement in efficiency of the metering technology.

requires strong assumptions on rank preservation, which do not hold in this context.

3.5 Empirical results and discussion

In this section, I present results motivated by the empirical models in previous section. First, Table 3.2 presents pooled OLS estimates of Equation 3.18. From Column (1) through (5), I add additional controls successively and analyze the coefficient on BF, the indicator for billing frequency. In each column, I present coefficients and robust standard errors clustered at the household level, as well as the billing district level, to account for serial correlation (Bertrand et al., 2004). The former accounts for the fact that water consumption is correlated within a household, whereas the latter acknowledges that the variation in treatment status arises at the billing district level. For all estimated ATEs, the sign is positive and each coefficient is significant at conventional levels. All other covariates have expected signs and significance. Collectively, the results imply that the transition to monthly billing increased consumer demand between five and nine percent, though the preferred specifications from this set of models are Columns (4) and (5), which control comprehensively for observable factors, exhibit the best fit, and also contribute the most conservative treatment effects.

In Table 3.3, I present results that include fixed effects at the household-level to control for any time-invariant unobservables. The results suggest that any omitted variable bias from the OLS models in the previous table is likely small. The ATE of billing frequency is estimated at 8.4% and 4.6% increases in quantity of water demanded with the inclusion of seasonal and period-of-sample fixed effects, respectively.[15] For comparison, the estimate of 4.6% in Column 2 of Table 3.3 can be compared directly with the estimate of 5.1% in Column 5 of Table 3.2, which could indicate a small positive bias from omitted factors. These estimates, however, are statistically similar. Moving forward, the estimate from Column (2) in Table 3.3 is the preferred estimate. This specification accounts for any time-invariant unobservables while simultaneously controlling for time effects. In terms of economic significance, this estimate is approximately half of the magnitude of water reductions called for during moderate to severe droughts in North Carolina (Wichman et al., 2014).

To examine whether there are unobservable changes in neighborhood characteristics that may bias the preferred ATE, I present regression discontinuity (RD) estimates (for the pooled cross-section) and difference-in-discontinuity (for fixed effects) estimates in Table 3.4. In Panel A of Table 3.4, pooled cross-section models are estimated for the set of households within 2000, 1000, and 500 feet of billing district boundaries. The treatment effect of billing frequency remains stable as the window shrinks, and remains significant at conventional levels.[16] Further, I present results from the preferred fixed effects model, and similarly restrict the sample to the same households within 2000, 1000, and 500 feet of billing district boundaries in Panel B of Table 3.4. The estimates decrease monotonically moving from Col-

[15]For these estimates, a block bootstrap of t-statistics (with 200 draws) was performed both at the household and billing district level to account for serial correlation, as suggested by Bertrand et al. (2004). Inference on the coefficient of interest does not change. Results are available upon request.

[16]The inclusion of a flexible polynomial for distance from the boundary cut-off (up to degree 3) has no effect on the treatment estimate nor the fit of the model. These results are available upon request.

umn (1) to (3), while standard errors increase, but the results remain statistically similar under the more conservative standard errors. The estimate within 500 feet of billing district boundaries indicates a positive effect of billing frequency on water demand that is robust to unobservable time-varying factors at the neighborhood level. Standard errors clustered at the billing district level indicate that this effect is not statistically different from zero, but I note that this effect is statistically similar to that of Columns (1) and (2).

Dynamic models presented in Table 3.5 imply that past water consumption is an important predictor of current water consumption and significantly improves the fit of each model. The treatment effect in the first two columns is smaller than that of the preferred specification (2.1% and 2.6%) but statistically different positive. Additionally, the estimate on the lagged dependent variable suggests that there is a relatively quick adjustment period (between one and two periods) to reach the new equilibrium. A long-run treatment effect is estimated at 10.5% and 6.6% in the OLS and fixed effects models, respectively. Since the fixed effects estimates are inconsistent (Nickell, 1981), the preferred estimate is from the OLS estimator, which suggests a long-run treatment effect that is double the magnitude of the preferred ATE in Table 3.3. As mentioned previously, the short-run effect in Column (1) provides a lower bound on the true treatment effect, whereas the preferred estimate from the fixed effects models (4.6%) bounds the treatment effect from above (Angrist and Pischke, 2009). These results suggest that there is indeed a positive response to increases in information provision in water demand and that this effect appears to strengthen over time.

Results from the dynamic models provide an interesting conclusion relative to previous research. While one-shot interventions of social norms and information provision tend to have an immediate effect that dissipates over time (Allcott and Rogers, 2014; Ferraro et al., 2011; Jessoe and Rapson, 2014), I find empirical evidence that a permanent change in billing frequency induces a short-run increase in consumption and provides a long-term mechanism through which consumers respond to the frequency of billing information. Of particular importance is that the long-term effect is larger than the magnitude of the short-run effect. This result suggests that consumers observe the change in frequency and, after a period or two of adjustment, consistently consume more water than they otherwise would have under a less frequent billing regime. The long-run treatment effect could be indicative of learning or habit formation within the new billing regime. However, an explicit analysis of the durability of these effects over time is left for future work.[17]

Additionally, this dynamic result is consistent with Gilbert and Graff Zivin (2014) who find that electricity consumption decreases in the week after the receipt of a larger than expected utility bill, but reverts back to a baseline level at the end of the month, when the electricity bill is less salient. With respect to my conceptual framework, this result implies that consumers might make similar changes to water use upon the receipt of a bill each month, but revert back to their prevailing misperceptions at the end of the billing cycle. While this result is not encouraging from

[17] An alternative approach to addressing the permanence of the treatment effects would be in an event study framework that captures the time profile of the ATE. However, Sexton (2014) notes that event studies are ill-suited for highly cyclical consumption processes such as water consumption. Additionally, the sequential treatment assignment mechanism in this setting confounds the choice of an appropriate control group from which treatment effects can be identified.

a conservation perspective, it sheds light on how consumers react to a permanent change in the frequency of information received. Additionally, since most randomized field experiments tend to be short-lived, this result adds to the literature on consumer responses to permanent changes in information provision in the case when consumers fully adjust their expectations about future price signals.

3.5.1 Heterogeneous treatment effects

I estimate heterogeneous treatment effects for two reasons—first, to explore any policy relevant heterogeneity in response to information provision and, second, to examine the mechanism through which billing frequency affects behavior.

The first question considered is whether treatment effects vary across seasons. In particular, outdoor water use in the summer is typically the target of conservation campaigns, so this use of water is particularly relevant for considerations of information provision as a conservation mechanism. As shown in Table 3.6, the ATE in summer months (June through August) is an approximately 15 percent increase in use, while other seasons exhibit relatively modest effects. All effects except spring usage are statistically positive. Since summer water usage is comprised of the production of predominantly "nonessential" household goods (e.g., a green lawn, a clean car), this effect is relevant for conservation policy. Regardless of the mechanism through which consumers are responding, this result indicates that more frequent information does not always induce conservation by reminding consumers of their consumption patterns. This result necessitates a closer examination of the ways in which consumer perceptions drive responses to new information before informative interventions can be a robust instrument of conservation.

Next, I examine the CATE of baseline water consumption. In particular, I create indicators for households that reside in each quintile of the consumption distribution for 2009 and 2010 (as well as the distribution of 2009 and 2010 summer consumption) and interact each of these with the treatment indicator. These results are presented graphically in Figure 3.9. As shown in Panel A, there is significant heterogeneity in responsiveness to billing frequency across the consumption distribution. I find a strong decreasing trend in the CATE as pre-treatment consumption increases. Households in the lowest 20th percentile of water consumption exhibit the largest CATE, while households in the highest 20th percentile exhibit a negative response to billing frequency. A similar relationship is observed in Panel B for the distribution of summer consumption as well, though CATE estimates remain positive for all subgroups.

Table 3.2: Baseline difference-in-difference regression results

Dep. Variable: $\ln(w_{ijt})$	(1)	(2)	(3)	(4)	(5)
BF	0.062	0.072	0.087	0.065	0.051
	(0.006)***	(0.006)***	(0.007)***	(0.007)***	(0.005)***
	[0.035]*	[0.036]*	[0.025]***	[0.027]**	[0.020]**
Rain		-0.003	-0.003	-0.003	-0.003
		(0.000)***	(0.000)***	(0.000)***	(0.000)***
		[0.001]*	[0.001]**	[0.001]**	[0.001]**
Maxtemp		0.004	0.004	0.005	0.005
		(0.000)***	(0.000)***	(0.000)***	(0.000)***
		[0.000]***	[0.000]***	[0.001]***	[0.001]***
Home Value			0.003	0.003	0.003
(in $10,000)			(0.001)***	(0.001)***	(0.001)***
			[0.001]**	[0.001]**	[0.001]***
Lot Size			0.022	0.023	0.026
			(0.009)**	(0.008)**	(0.009)***
			[0.014]	[0.013]	[0.014]*
Square Feet			0.014	0.014	0.014
(in 100 ft)			(0.001)***	(0.001)***	(0.001)***
			[0.002]***	[0.002]***	[0.002]***
% Rent			-0.002	-0.003	-0.016
			(0.018)	(0.018)	(0.019)
			[0.037]	[0.038]	[0.035]
% White			0.053	0.051	0.025
			(0.014)***	(0.014)***	(0.017)
			[0.032]	[0.033]	[0.028]
Home Age			-0.002	-0.002	-0.002
			(0.000)***	(0.000)***	(0.000)***
			[0.001]***	[0.001]***	[0.001]***
Household Size			0.232	0.233	0.231
			(0.008)***	(0.008)***	(0.008)***
			[0.017]***	[0.017]***	[0.019]***
Observations	1,694,859	1,694,859	1,684,025	1,684,025	1,684,025
Adj. R-squared	0.002	0.002	0.036	0.037	0.039
Additional controls:					
Time trend	Y	Y	Y	N	N
Season FEts	N	Y	Y	N	N
Time FEs	N	N	N	Y	Y
District FEs	N	N	N	N	Y

Note: Robust standard errors clustered at the household-level in parentheses. Robust standard errors clustered at the billing district in square brackets. Estimation results are from OLS regressions with the dependent variable as log consumption. Constant term omitted. *** p<0.01, ** p<0.05, * p<0.1.

Table 3.3: Fixed effects difference-in-difference regression results

Dependent Variable: $\ln(w_{ijt})$	(1)	(2)
BF	0.084	0.046
	(0.005)***	(0.005)***
	[0.020]***	[0.017]**
Rain	-0.003	-0.004
	(0.000)***	(0.000)***
	[0.001]**	[0.001]***
Maxtemp	0.004	0.005
	(0.000)***	(0.000)***
	[0.001]***	[0.001]***
Households	58,965	58,965
Observations	1,694,859	1,694,859
Within R-squared	0.004	0.006
Additional controls:		
Season fixed effects	Y	N
Time fixed effects	N	Y
Household fixed effects	Y	Y

Note: Robust standard errors clustered at the household-level in parentheses. Robust standard errors clustered at the billing district in square brackets. Results are from linear panel data estimators with log consumption as the dependent variable. Constant term omitted. *** p<0.01, ** p<0.05, * p<0.1.

Table 3.4: Local average treatment effect estimates

Dependent Variable: $\ln(w_{ijt})$	(1) Within 2000ft	(2) Within 1000ft	(3) Within 500ft
Panel A: Pooled cross-section regression-discontinuity results			
BF	0.043	0.044	0.047
	(0.006)***	(0.008)***	(0.011)***
	[0.021]*	[0.022]*	[0.024]*
Observations	1,334,147	812,965	436,377
R-squared	0.039	0.037	0.035
Panel B: Fixed effects difference-in-discontinuity results			
BF	0.045	0.040	0.035
	(0.006)***	(0.007)***	(0.010)***
	[0.018]**	[0.018]**	[0.021]
Number of households	46,645	28,632	15,462
Observations	1,341,595	816,643	437,730
Within R-squared	0.006	0.006	0.006

Note: Robust standard errors clustered at the household-level in parentheses. Robust standard errors clustered at the billing district in square brackets. Results are from local linear (panel) estimators with log consumption as the dependent variable. Each column represents a limited sample of households within 2000, 1000, and 500 feet of a billing district boundary, respectively. In Panel A, all models include full demographic covariates, weather variables, time effects, and billing district fixed effects. In Panel B, all models include weather variables, time effects, and household fixed effects. *** p<0.01, ** p<0.05, * p<0.1.

Table 3.5: Dynamic regression results and partial adjustment estimates

	(1)	(2)
Estimator:	OLS	Fixed Effects
Dependent variable:	$\ln(w_{ijt})$	$\ln(w_{ijt})$
$\ln(w_{ijt-1})$	0.800	0.597
	(0.002)***	(0.003)***
	[0.012]***	[0.018]***
BF	0.021	0.026
	(0.002)***	(0.003)***
	[0.012]*	[0.012]**
Long-run treatment effect:	0.105	0.066
	(0.010)***	(0.006)***
	[0.057]*	[0.030]**
Number of households	—	58,911
Observations	1,548,942	1,558,593
R-squared	0.633	—
Within R-squared	—	0.345
Additional controls:		
Full demographic covariates	Y	N
Billing district fixed effects	Y	N
Household fixed effects	N	Y

Note: Robust standard errors clustered at the household-level in parentheses. Robust standard errors clustered at the billing district in square brackets. All models include weather covariates and time fixed effects. *** p<0.01, ** p<0.05, * p<0.1.

Table 3.6: Heterogeneous treatment effects among seasons and automatic bill payment

Dependent variable: $\ln(w_{ijt})$	(1) Spring	(2) Summer	(3) Fall	(4) Winter	(5) ABP
BF	0.004	0.154	0.068	0.022	-0.015
	(0.006)	(0.008)***	(0.006)***	(0.006)***	(0.015)
	[0.023]	[0.061]**	[0.038]*	[0.022]	[0.021]

Note: Robust standard errors clustered at the household-level in parentheses. Robust standard errors clustered at the billing district in square brackets. Coefficients correspond to interactions with BF from linear panel data estimators with log consumption as the dependent variable and household and time fixed effects. All models include weather covariates. Columns (1) through (4) are estimated within the same model. Column (5) is estimated independently. *** p<0.01, ** p<0.05, * p<0.1.

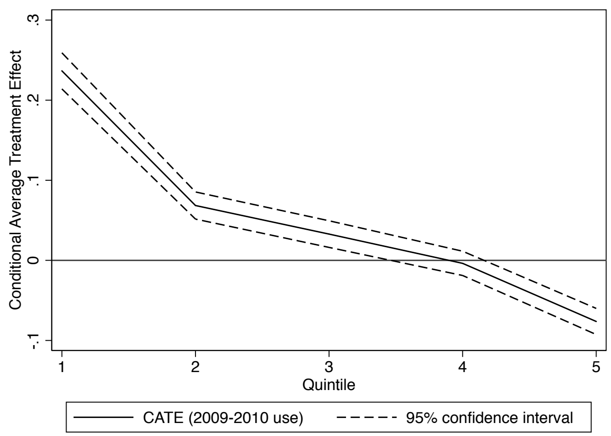

(a) ATE conditional on quintile of 2009-2010 use

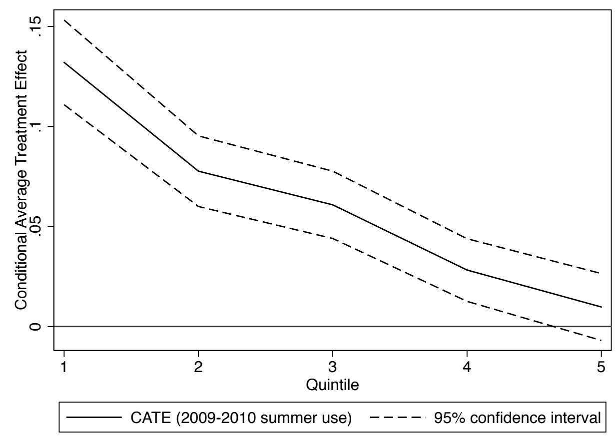

(b) ATE conditional on quintile of 2009-2010 summer use

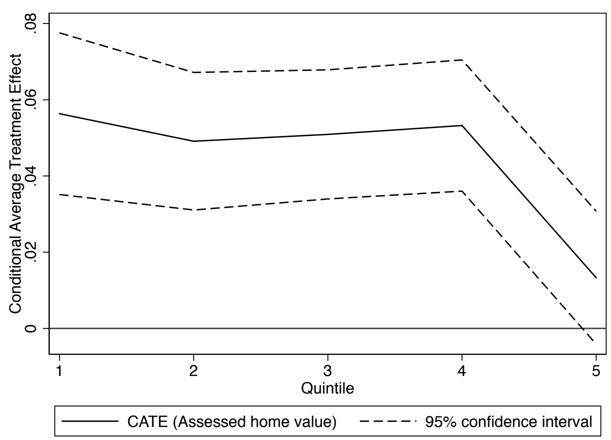

(c) ATE conditional on quintile of assessed value of home

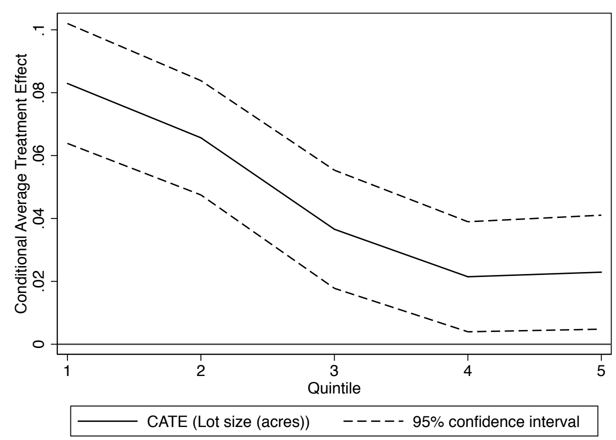

(d) ATE conditional on quintile of lot size (in acres)

Figure 3.9: Conditional average treatment effects for usage, wealth, and lot size

These relationships are consistent with Mansur and Olmstead (2012); Wichman (2014); Wichman et al. (2014), and Klaiber et al. (2014) who show that low-use households are more sensitive to price changes. However, the negative effect observed in Panel A suggests that informative interventions work differently for different subpopulations. The implications that arise from these effects on baseline usage suggest that lower users of water increase consumption by more (in percentage terms) than high users of water. Upon receiving monthly bills, low water users may realize that water is less expensive than they previously thought and increase consumption accordingly. Additionally, the within-sample heterogeneity indicating both a positive and a negative response to the same information treatment emphasizes the importance of reconsidering the mechanism through which consumers assimilate and use information in intermittent choice settings.

Additionally in Figure 3.9, I consider the assessed value of a home as a proxy for a consumer's wealth, as well as lot size as a proxy for preferences for outdoor water use. As shown in Panel C, households in the highest 20th percentile of "wealth" exhibit no significant response to the treatment. High-income households generally have a greater willingness-to-pay for water and, thus, are likely to be less sensitive about changes to their water bill. Since low-income households are the most sensitive to price (Mansur and Olmstead, 2012; Wichman et al., 2014), the heterogeneity among income classes is consistent with the notion that consumers are responding to the change in frequency similarly to a change in price.

In Panel D of Figure 3.9, I present the CATEs for quintiles of a household's lot size as a proxy for outdoor water use preferences (Mansur and Olmstead, 2012; Renwick and Archibald, 1998). As shown, there is a decreasing trend in the estimated CATE as lot size increases, though all estimates remain positive and significant. This relationship can be interpreted similar to the effect of summer water use—more frequent reminders might attenuate the effect for large users of water, while low users of water may increase their usage more in response to the same information.

I also explore the quintile of house size and the number of bathrooms as indicators of preferences for water intensive indoor use and appliances in Figure 3.10. Both of these structural characteristics of a home display similar trends in the first two panels—the higher ends of the distribution exhibit lower responses to the treatment, but similar trends in the lower portion of the distribution. These trends imply that the treatment effect is not driven by preferences for indoor water use. Lastly, I show CATEs for quintiles of the age of the home in the last panel of Figure 3.10. Newer homes exhibit a significantly larger response to the treatment than older homes. While Ferraro and Miranda (2013) find no significant differences among the age of home in response to information and social comparison treatments for water use, they conclude that the margin on which consumers might be responding is a behavioral change rather than a fixed investment. Since older homes are likely more prone to leaky plumbing and older meters, I draw a similar conclusion in that the response is likely a recurring behavioral response to the receipt of a monthly bill.

Collectively, these results suggest that there is substantial heterogeneity in response to treatment and, for large baseline consumers, more frequent provision of consumption information can incur a negative response. For the majority of other conditioning covariates, a positive treatment effect is the predominant result despite within-sample heterogeneity. The results are consistent with several stylized facts about water consumption in response to changes in price and non-pecuniary instruments, suggesting that price misperception may be driving the ATE, particularly for low-use households. This conclusion, however, is speculative. Further research on the mechanism driving these results is needed. Overall, the heterogeneous treatment effects provide evidence of increases in outdoor water use in the summer months in response to increases in billing frequency.

3.5.2 Robustness and sensitivity

I examine the sensitivity and robustness of the primary empirical results in several ways. In Table 3.7, I present regression results that predict the likelihood of a billing district to be transitioned from bi-monthly to monthly billing based on observable household characteristics. In particular, I consider the first district to transition in a probit model, as well as the sequential transition of routes in an ordered probit framework. The former suggests that lot size, household size, and the number of bedrooms increased the probability of the pilot group being chose, while the square footage of the home and number of bathrooms decreased the probability. There is weak evidence that smaller water bills are correlated with initial treatment. Further, the ordered probit reveals that the age of a home is weakly negatively correlated with the order of transition, while all other observable factors are insignificant. These results suggest that treatment selection on observable household characteristics is not a significant concern.

Table 3.7: Predicting the likelihood of billing districts to be transitioned from bi-monthly to monthly billing based on observable household characteristics

	Probit (first district)			Ordered probit (All districts)		
	(1)	(2)	(3)	(4)	(5)	(6)
Home Value ($10,000)	0.008		0.008	-0.010		-0.011
	(0.005)		(0.005)	(0.010)		(0.010)
Lot size (acres)	0.104*		0.110*	0.045		0.076
	(0.054)		(0.061)	(0.127)		(0.137)
Square feet (in 100 ft)	-0.028***		-0.029***	0.013		0.011
	(0.005)		(0.006)	(0.010)		(0.010)
No. bathrooms	-0.261***		-0.251***	-0.011		0.008
	(0.051)		(0.048)	(0.094)		(0.086)
No. bedrooms	0.065***		0.044**	-0.022		-0.023
	(0.018)		(0.018)	(0.033)		(0.036)
% Rent	-0.169		-0.164	-0.329		-0.290
	(0.181)		(0.172)	(0.313)		(0.305)
% White	-0.393		-0.405	0.436		0.489
	(0.266)		(0.270)	(0.456)		(0.467)
Home age	0.005		0.005*	-0.012*		-0.011*
	(0.003)		(0.003)	(0.007)		(0.007)
Household Size	0.306***		0.311***	-0.027		-0.004
	(0.100)		(0.101)	(0.131)		(0.130)
Total Bill ($)		-0.002*	-0.002		0.001	0.001
		(0.001)	(0.001)		(0.002)	(0.002)
Pre-treatment use		0.000	-0.000		-0.000	-0.000
		(0.000)	(0.000)		(0.000)	(0.000)
Pre-treatment summer use		0.000	0.000***		0.000	0.000
		(0.000)	(0.000)		(0.000)	(0.000)
Observations	58,486	55,617	55,322	58,486	55,617	55,322

Note: Robust standard errors in parentheses clustered at the district level. Nominal coefficients are presented here; marginal effects are available upon request. The first 3 columns predict the likelihood of the first district being chosen to transition from monthly billing. The last 3 columns predict the order in which billing districts were transitioned from bi-monthly to monthly billing. *** p<0.01, ** p<0.05, * p<0.1

Further, In the last column of Table 3.6, I show that the effect of an interaction of the treatment indicator and automatic bill payment (ABP) works in the opposite direction of the treatment effect. This provides an empirical test for whether customers who are inattentive to prices and water bills observe the change in billing frequency (Sexton, 2014). The net response of customers enrolled in ABP throughout the transition to monthly billing is negative but not statistically different from zero—indicating a null response among this subgroup of customers.[18] This result instills confidence that the preferred specifications are indeed identifying a response to the change in billing frequency. The lack of a positive effect also lends credence to the notion that the preferred ATE in previous models is not an artifact of mechanical efficiency of the new meters. However, due to selection into ABP, these results should be interpreted with caution.

To examine technical efficiency of the new meters directly, I present the percent of water accounted for, by utility accounting sheets, as a percent of total pumped water in Figure 3.11 in each month over the time period of my study. These data are obtained from Annual Financial Information Reports (AFIR) submitted to the NC State Treasurer. While this measure of system efficiency is an aggregate statistic, a technical efficiency gain from metering technology should display an increasing trend as older meters are replaced. As shown, there is no discernible increase in the system efficiency after the rolling introduction of monthly bills. While this measure includes commercial and industrial water usage, residential consumption is the largest share of water use in Durham and, by June 2014, nearly 35% of all water sold was residential. Of that 35%, the majority was billed through the new meters. So, the flat trend observed after the introduction of the new meters provides some indication that the purported technical efficiency improvement of meters does not

[18]Approximately 5% of the sample was enrolled in ABP.

confound the treatment.

Further, as a final robustness check on the unconfoundedness assumption in light of more efficient metering technology, I exploit the fact that the new meters were installed within a billing district several months prior to the transition to monthly billing. Part of this delay is simply due to feasibility constraints in the installation process, as well as testing of the new meters and the corresponding meter reading software prior to the full transition to monthly billing. As such, for each billing district (or districts, if multiple transitioned in the same time period), I estimate the preferred fixed effects specification on a two-period window directly before the change in billing frequency. Thus, the "treatment" in this test is a dummy variable for the billing district that had new meters installed interacted with a time fixed effect for the latter period. The control group is comprised of all households billed on a bi-monthly basis at that point in time. This falsification test examines whether the new meters had a significant impact on consumption. Formally, the null hypothesis is that there is no consumptive effect of the new meters, while the alternative is that consumption will increase in response to the "treatment." Table 3.8 displays the estimated effect. In every scenario, we fail to reject the null hypothesis. I find that for three of the nine models, the "treatment" effect is negative and significant. All other estimates are not statistically different from zero. This final test provides more confidence in the notion that the new metering technology does not contaminate the treatment effect.

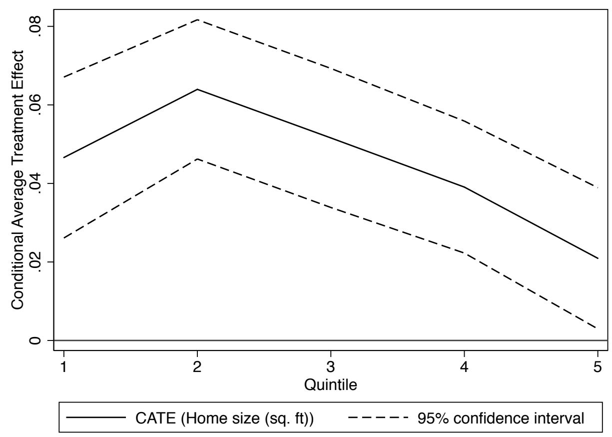

(a) ATE conditional on quintile of house size (in square feet)

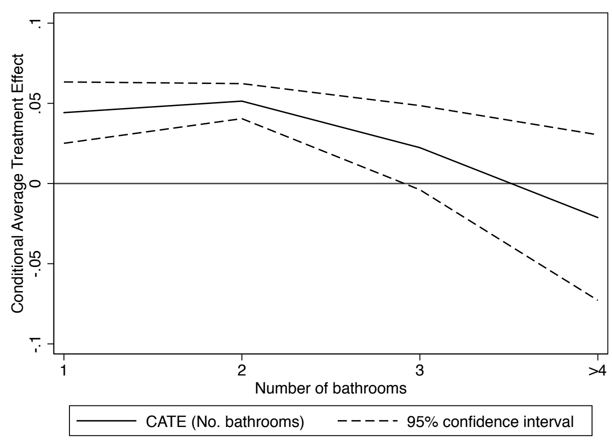

(b) ATE conditional on number of bathrooms

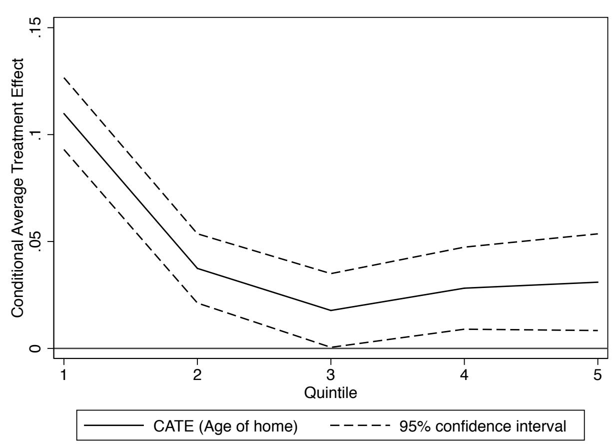

(c) ATE conditional on age of home

Figure 3.10: Conditional average treatment effects for structural characteristics of home

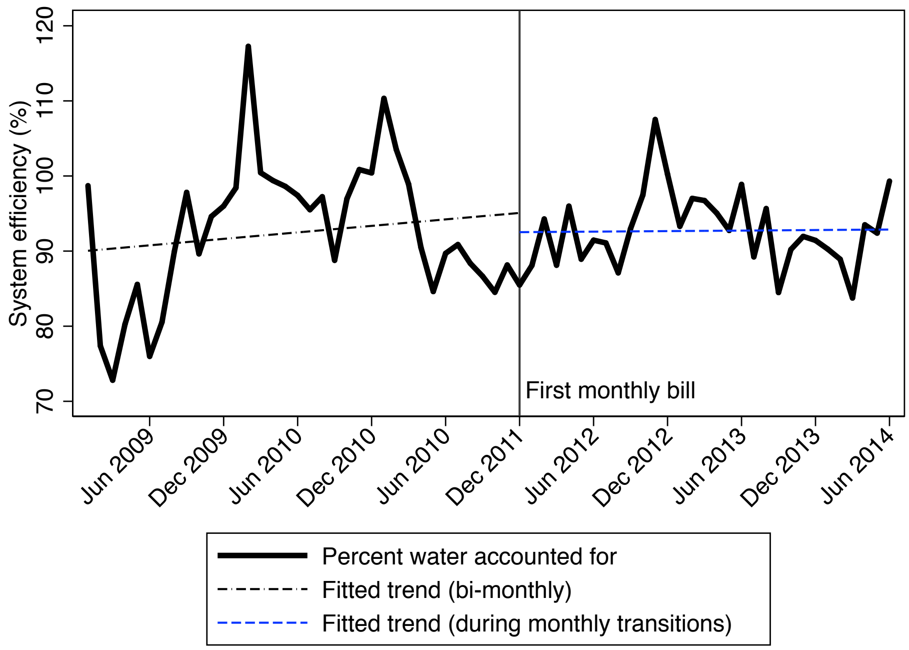

Figure 3.11: Percent water accounted for as a percentage of total pumped water

Table 3.8: Robustness check for two periods directly before transition to monthly billing

Dependent Variable: $\ln(w_{ijt})$	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Billing district(s) used as false treatment:								
	3	9	4 & 10	6 & 15	16	13	17	1	11
"treat"	-0.006	-0.053***	-0.012	0.035	-0.088***	-0.096**	-0.292	-0.061	0.051
	[0.017]	[0.018]	[0.018]	[0.028]	[0.009]	[0.035]	[0.244]	[0.048]	[0.037]
Number of households	53,141	53,676	50,667	43,806	36,603	33,045	28,389	24,247	19,828
Observations	101,981	105,109	98,499	82,860	69,151	61,939	55,685	44,788	38,839
Within R-squared	0.009	0.009	0.022	0.021	0.002	0.019	0.061	0.027	0.019
Additional controls:									
Weather covariates	Y	Y	Y	Y	Y	Y	Y	Y	Y
Time fixed effects	Y	Y	Y	Y	Y	Y	Y	Y	Y
Household fixed effects	Y	Y	Y	Y	Y	Y	Y	Y	Y

Note: Robust standard errors clustered at the billing district in square brackets. Each column represents a difference-in-difference estimate in a two-period model for the coefficient "treat" for the two periods directly prior to the transition to monthly billing. A single time fixed effect in the latter period serves as the mean sample trend from which the coefficient on "treat" is identified. For each subsequent column, any billing district on monthly billing is omitted from the sample to prevent contamination of treatment. For example, moving from Column (2) to (3), billing districts 3 and 9 are not included in the regression for Column (3). *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

3.6 Welfare analysis

The empirical results suggest a significant increase in consumption in response to the change from bi-monthly to monthly billing. As such, there are welfare gains to providing consumers more frequent information if that information reduces the uncertainty in consumers' perception of prices or levels of quantity demanded. This section quantifies the economic benefit that arises from reducing the distortion in perceived prices and quantities, subject to the models and assumptions formulated in Sections 3.2.2 through 3.2.5.

For the welfare analysis, I use results from the preferred difference-in-difference models. The change in consumption considered is a 4.5% increase in water use in response to the change in billing frequency. Welfare estimates are calculated for a common range of short-run elasticities for residential water demand found in the literature that imply that residential water demand is generally inelastic, with modal estimates lying between -0.5 and -0.2 (Arbués et al., 2003; Espey et al., 1997). Additionally, I assume that the price relevant for consumer decision-making under block rate structures is the *ex post* average price as established by Wichman (2014) for water demand and Ito (2014) for electricity demand, which is \$8.60 per hundred cubic feet in the sample for households who never transitioned to monthly billing.

In Table 3.9, I present welfare statistics for a range of elasticities under different modeling assumptions. Particularly, I let the *ex post* perceived price vary from 100% to 150% of the true price. Allowing $\tilde{p_1} = p$ reflects the pure nudge assumption, while $\tilde{p_1} > p$ relaxes this assertion in line with Assumption 2. While the price elasticity may indeed be different for consumers responding to perceived price and the true price, there is not sufficient variation in prices within the natural experiment to estimate these elasticities directly. As such, I present changes in consumer surplus as a function of the perceived price elasticity. Under the assumption that the average consumer responds to prices perfectly after the change in billing frequency (which is equivalent to assuming that consumers respond through a quantity mechanism) with a price elasticity of -0.3, the welfare gain from increased information provision is approximately \$0.28 per month. This amount reflects a \$1.29 decrease in the price paid for water and approximately 0.68% of the average consumer's monthly bill. Taken collectively, a conservative estimate of the change from bi-monthly to monthly billing for Durham, NC water customers resulted in an approximately \$300,000 per annum increase in consumer welfare (that is, \$5.11 per household).[19] From a conservation perspective, the change in billing frequency resulted in an aggregate increase in water consumption of roughly 46,000 cubic feet per year. This increase in consumption is roughly equivalent to the amount of water it would take to fill 520 Olympic-size swimming pools.[20]

3.7 Conclusions

Information provision is a growing topic both in markets where expenditures are made intermittently and where price regulation is politically challenging. In

[19]This back-of-the-envelope statistic is obtained by multiplying the average customer's welfare gain by the approximate number of households in the sample (60,000) by twelve to obtain an annual equivalent.

[20]The volume of an Olympic-size swimming pool is approximately 88,000 cubic feet.

Table 3.9: Changes in consumer surplus from an increase in billing frequency under different modeling assumptions

Price elasticity (η)		-0.2	-0.3	-0.4	-0.5
Panel A: Price mechanism for a -4.5% demand response					
	$\Delta\theta p$ (\$/ccf)	-1.93	-1.29	-0.97	-0.77
$\tilde{p}_1 = 100\% \times p$	Δ Consumer surplus (\$)	0.43	0.28	0.21	0.17
	% Monthly bill	1.01%	0.68%	0.51%	0.41%
	$\Delta\theta p$ (\$/ccf)	-2.13	-1.42	-1.06	-0.85
$\tilde{p}_1 = 110\% \times p$	Δ Consumer surplus (\$)	0.47	0.31	0.23	0.19
	% Monthly bill	1.11%	0.74%	0.56%	0.45%
	$\Delta\theta p$ (\$/ccf)	-2.42	-1.61	-1.21	-0.97
$\tilde{p}_1 = 125\% \times p$	ΔConsumer surplus (\$)	0.53	0.35	0.27	0.21
	% Monthly bill	1.27%	0.84%	0.63%	0.51%
	$\Delta\theta p$ (\$/ccf)	-2.90	-1.93	-1.45	-1.16
$\tilde{p}_1 = 150\% \times p$	Δ Consumer surplus (\$)	0.64	0.43	0.32	0.26
	% Monthly bill	1.52%	1.01%	0.76%	0.61%
Panel B: Quantity mechanism for a -4.5% demand response					
	Δp (\$/ccf)	-1.93	-1.29	-0.97	-0.77
$p(\lambda) = p$	Δ Consumer surplus (\$)	0.43	0.28	0.21	0.17
	% Monthly bill	1.01%	0.68%	0.51%	0.41%

Notes: $\tilde{p}_1$ is the baseline price after the change in billing frequency. The first row in Panel A reflects the pure nudge assumption in that perceived price is equated with the true price with increased billing frequency. Subsequent rows relax this assumption by allowing for misperception to be proportional to the true price. $\Delta\theta p$ is the change in price that reflects the demand response to billing frequency under the price mechanism for different assumptions on price elasticities. Panel B presents the analogous measurements for the quantity mechanism. Δp is the change in the true price corresponding to the demand response under a quantity mechanism. Consumer surplus is approximated according to Equations 3.7 and 3.15, respectively, for the relevant change in demand. The true price used for all calculations is \$8.60/ccf, which is the sample mean average price for households that never transitioned to monthly billing.

this paper, I make several contributions to this literature. I first posit a model of consumer behavior that is consistent with the observed demand response to more frequent billing. I develop a transparent analytical framework to measure economic welfare using limited empirical information. Empirically, I take advantage of a natural experiment in which residential water customers are exposed to more frequent billing information for a water provider in the southeastern US. I find strong evidence that with the provision of more frequent information, consumers increase consumption of water by approximately 5%, which is roughly half the magnitude of water reductions called for during moderate to severe drought in North Carolina. This result is the first documented causal increase in consumption in response to an increase in billing frequency within environmental and resource policy, which is a particularly pertinent result for drought-prone regions. Similar to repeated informative interventions, I find evidence of a persistent long-run treatment effect. Further, heterogeneous treatment effects suggest that the increase in consumption is driven by outdoor water use during the summer.

In practical terms, many water utilities bill customers bi-monthly, or even less frequently, despite growing empirical evidence that more frequent price and quantity signals can encourage conservation. This research provides a counterpoint to the existing empirical literature by identifying a robust positive demand response to an exogenous transition from bi-monthly to monthly billing. This change in behavior is attributed to the fact the consumers billed less frequently observe more opaque price signals, thus driving a wedge between perceived and actual billing information. Increasing the transparency of prices through increased billing frequency results in welfare gains of approximately 0.5 to 1 percent of aggregate expenditures on water use. These economic gains, however, come at the cost of increased consumption of a scarce natural resource. Thus, future research seeking to use increased information provision as a tool of conservation needs to account for potential uncertainty in consumers' perceptions of prices.

This research adds to the broader literature on intermittent billing and inattention for economic goods, and provides a topical counterpoint to mounting empirical evidence that informative interventions can serve as a cost-effective instrument of conservation. While this research sheds light on how consumers may react to a change in the frequency of familiar information, more research needs to be performed to understand the specific mechanism through which consumers assimilate and use information for decision-making in intermittent choice settings.

Chapter 4: Incentives, green preferences, and private provision of environmental public goods

4.1 Introduction

Environmental policy often appeals to an individual's attitude towards the environment by relying on social norms to influence behavior. By encouraging individuals to "do the right thing," voluntary provision of environmental public goods is acquiring a larger role as a policy instrument (Allcott, 2011; Ferraro and Price, 2013; Glaeser, 2014). But, preferences for public goods are difficult to observe and free-riding is a common problem in the provision of public goods. From an economic perspective, these problems are not new. In the context of privately provided public goods, however, several policy-relevant questions deserve our attention. In particular, is there a normative rationale for regulation in the private provision of public goods? Do standard solutions to free-riding problems apply to the class of privately provided goods? Can we say anything new about the role of incentives for voluntarily supplied public goods that correct environmental market inefficiencies?

In this paper, I develop a general model of environmental public goods provision that facilitates optimal provision under heterogeneous preferences for the environment. "Green" preferences are defined informally as the differential benefit that arises both from contributing to and consuming an environmental public good (Kotchen, 2005, 2006).[1] Not surprisingly, these preferences are distributed across the population in an unknown manner. Several studies have attempted to quantify these tastes by revealed preference indicators. Kotchen and Moore (2007) and Jacobsen et al. (2012), for example, examine the conservation behavior of electricity customers who opt-in to green electricity programs. Additionally, Kahn (2007) use the proportion of green party voters within a county to examine fuel-efficient vehicle sales in California. Further, Sexton and Sexton (2014) characterize private signaling benefits that arise from green consumption for the case of hybrid vehicle purchases. Clean electricity, energy efficient and renewable energy technology adoption, and fuel efficient vehicles are salient examples of private goods that provide public environmental benefits.

In contrast to the existing literature in environmental economics, I approach the policy relevance of green preferences as a problem of demand revelation between heterogeneous consumers. Despite the public good being provided at least cost to the consumer, the optimal level of public goods provision is not attained in a Laissez-faire economy. Thus, there are attainable efficiency gains by constructing policy instruments that elicit true preferences for the environment. In this effort, I consider the equilibrium and efficiency properties of several mechanisms that fall within the general class of Clarke-Groves mechanisms under different regulatory scenarios. In

[1] Unlike Kotchen (2006), I consider an impure environmental public good explicitly, rather than modeling the private and public characteristics of an environmental good independently. Doing so disallows the potential for substitution possibilities between public and private characteristics of environmental goods, which is not a central focus of this analysis.

particular, I relax the information available to a regulator progressively. In the mechanisms considered, truthful revelation of preferences is not a dominant strategy equilibrium. But, I find incentive compatible contracts that support Nash equilibria and induce socially optimal public goods provision when the regulator can contract upon either 1) individual provision or 2) individual reported provision paired with observable group output. I show that a contract conditional on individual reported provision alone is not incentive compatible.

Additionally, I reframe the role of budget balancedness in the context of policy incentives for correcting inefficiencies in environmental markets. A topical issue in environmental policy is the role of government intervention to close the "energy efficiency gap." Allcott and Greenstone (2012) outline the potential "win-win" scenario from increases in energy efficiency investments characterized by cost-savings to consumers and lessened environmental damages from electricity generation. Energy efficiency is an ideal example of a privately provided public good upon which to focus. Moreover, recent theoretical work emphasizes the potential cost-effectiveness of quantity over price regulation when heterogeneous preferences preferences for the environment are present (Jacobsen et al., 2014).

In this analysis, I extend the literature on public goods by examining privately provided environmental public goods in a mechanism design context. Several researchers have examined the dynamics of privately provided public goods, however relatively few have found positive equilibrium results (Falkinger, 1996; Falkinger et al., 2000; Kirchsteiger and Puppe, 1997; Varian, 1994). Many of the mechanisms proposed are either difficult to implement in practice (Varian, 1994) or exhibit undesirable equilibrium properties (Kirchsteiger and Puppe, 1997). Falkinger (1996) proposes a simple mechanism that punishes or rewards individuals based on their deviation from the average provision. This mechanism balances the budget, by definition, and almost achieves Pareto efficiency in Nash equilibrium. A test of this mechanism in a laboratory experiment suggests that it performs well in practice (Falkinger et al., 2000). Much of this research, however, relies on strong informational assumptions. I contribute to this literature by examining weaker informational constraints imposed on the regulator. Additionally, I model heterogeneous agents explicitly to provide intuition for incentive compatibility in the case of privately provided environmental public goods.[2]

The model I develop in this paper complements previous research in designing incentive contracts to regulate nonpoint source pollution. In particular, I draw insight from the general formulation of team production by Hölmstrom (1982), and its application to environmental regulation through collective penalties by Meran and Schwalbe (1987) and Segerson (1988). Both rely on group penalties that apply when realized emissions levels exceed some desired level of pollution. The parallels to public goods provision are in constructing contracts conditional on the observability of group provision, and its deviation from individual reports. A primary distinction is that emission levels are determined exogenously for nonpoint source pollution, whereas optimal provision of the public good in this paper is constructed from consumer preferences.

[2]While I motivate this model with environmental preferences, this model applies for the general case of privately provided public goods with heterogeneous preferences.

Further, an additional complementary line of research examines how agents with different preferences match with mission-oriented principals to provide collective goods. Besley and Ghatak (2005) consider how aligning the preferences of consumers and firms, for example, may result in efficient outcomes that deviate from that of a purely competitive economy. In my model, without heterogeneous principals, non-pecuniary incentives drive "motivated" consumer behavior, which in turn allows for a regulator to exploit preferences to induce the optimal levels of public goods provision. In this sense, we can think of heterogeneity in green consumers as differential characteristics of mission-oriented preferences.

The paper proceeds as follows. In the next section, I develop a general model of privately and socially optimal provision for two agents. I consider the efficiency properties and incentive compatibility of various incentive contracts under progressively weaker informational constraints on a regulator in the third section. And, the last section concludes.

4.2 Private provision of environmental public goods

I adopt a general model of private provision of public goods similar to that of Andreoni (1990) and Bergstrom et al. (1986). In general, a consumer gains utility from a private good, x_i, an environmental public good, C, and her own private contributions to the public good, c_i. The public good is funded entirely by private contributions such that $C = c_i + C_{-i}$ where $C_{-i} = \sum_{j \neq i} c_j$. Consumer i's preferences are specified as,

$$V(c_i, C; \theta_i) + x_i \tag{4.1}$$

where θ_i, lying in the space $\Theta \in \mathbb{R}_0^+$, is a green preference parameter known only to the consumer. V is a twice-differentiable, weakly concave preference function common to all consumers. Further, I assume V has increasing differences, which implies that V satisfies the Spence-Mirrlees single-crossing condition (SCC) in $(c_i, C; \theta_i)$.[3] C_{-i} is treated as exogenous by the Nash assumption.

I simplify by restricting agents to two types—a green consumer, denoted by θ_g, and a non-green consumer, denoted by θ_n. I impose the ordering of preference types, $\theta_g > \theta_n$, such that $V(c_g, C; \theta_g) \geq V(c_n, C; \theta_n)$ by the single-crossing condition. Thus, heterogeneity in environmental preferences arises through the individual-specific parameter, θ_i. I model green preferences as tastes for an impure environmental good, rather than a pure public good, to capture the private benefit that green consumers reap from their provision (Kotchen, 2006). Models with heterogeneous preferences for pure environmental goods can be solved with standard mechanisms.

Since V defines a general structure of public goods provision, it is worth noting special cases. In particular, if we place zero weight on the first term of V, we are left only with preferences over C, which is the canonical model of private public goods provision, or "purely altruistic" provision (Bergstrom et al., 1986). Further, by placing some weight on both arguments of V, we have "warm glow" preferences for an impure public good in the sense of Andreoni (1990), in which there is utility

[3]Increasing differences ensures that the marginal benefits of c_i are increasing in θ_i. Mathematically, this is equivalent to $\partial V(c_i, C; \theta')/\partial c_i - \partial V(c_i, C; \theta)/\partial c_i \geq 0$ for all $\theta' > \theta$, since V is differentiable.

derived from the contribution itself that is separate from direct consumption of the public good. Placing no weight on consumption of the public good results in "purely egoistic" preferences. Further, V could take on social-welfare or social-efficiency preferences in the sense that consumers contribute to the public good in the direction of the socially efficient outcome (Charness and Rabin, 2002).

The most natural interpretation of consumer heterogeneity in this framework given the assumptions on V, however, is that of "warm glow" preferences (Andreoni, 1990). That is, consumers share common values for consumption of the public good, but there is variation in their preferences over their own provision.

Consumers face the budget constraint, $w_i = x_i + ac_i$, with constant marginal costs of provision, a, and a non negativity constraint on individual provision $c_i \geq 0$. I assume homogenous costs to focus the analysis on heterogeneity arising from an individual's preferences.[4] The budget constraint is satisfied with equality since any residual wealth is allocated to the numeraire good. Thus, by substitution, we can write net utility for consumer i,

$$u_i = w_i + V(c_i, C; \theta_i) - ac_i. \tag{4.2}$$

And, thus, the consumer's objective function is defined as,

$$\max_{c_i}\{u_i | c_i \geq 0\} = \max_{c_i}\{w_i + V(c_i, C; \theta_i) - ac_i | c_i \geq 0\} \text{ for } i = g, n, \tag{4.3}$$

[4]The main results of this analysis hold for increasing marginal costs, but do not add insight into the problem. Common costs of environmental provision are not realistic, however, abstracting from heterogeneous costs isolates the role of preferences in the provision of public goods; relaxing this assumption is left for future work.

where net utility (u_i) is a function of an individual's wealth and value of the public good less private costs incurred by providing the public good. An individual's valuation of the public good, however, is independent of her wealth as is necessary for application of Clarke-Groves mechanisms in subsequent sections (Green and Laffont, 1977a; Groves and Ledyard, 1976).

First-order conditions for Equation 4.3 implicitly define optimal contributions to the public good in the *private market equilibrium* for interior solutions,

$$\frac{\partial V}{\partial c_i}(c_i, C; \theta_i) - a = 0 \text{ for } i = g, n. \tag{4.4}$$

From the structure of this first-order condition, the single-crossing condition implies the following,

Proposition 1 *In the private market equilibrium, green consumers will provide as least as much of the public good as non-green consumers, that is,* $c_g \geq c_n$.

Proof: Under the assumption that $\theta_g > \theta_n$, $c_g \geq c_n$ follows directly from the single-crossing condition of V in $(c_i, C; \theta_i)$ that implies marginal utility is nondecreasing in type. □

The private market equilibrium defined in Equation 4.4 will be inefficient, however, since it fails to incorporate the external spillovers from any one individual's provision. To see this, we can introduce a social planner whose objective is to maximize social welfare, that is the sum of each individual's net utility. The social planner, who can be thought of as a government, is able to freely distribute income among consumers via lump-sum transfers and faces no budget constraint.[5] The former reflects the assumption of quasilinear preferences, which enables lump-sum transfers between agents, which is important for the mechanisms considered in the sections that follow. The latter implies that the public good is funded entirely by individual contributions. While not realistic in practice, the solution to the social planner's problem will define the optimal level of public goods provision from each type.

The social planner's problem (SPP) takes the following form for a two-person economy,

$$\max_{c_g, c_n} \{u_g + u_n | c_g, c_n \geq 0\} \text{ for } g \neq n. \tag{4.5}$$

Net utility for each type (u_i for $i = g, n$) is given by the functional form in Equation 4.3. The first-order conditions for Equation 4.5 implicitly define optimal contributions to the public good for each type (c_i^* for $i = g, n$) in the *social equlibrium*,

$$[c_g] : \underbrace{\frac{\partial V}{\partial c_g}(c_g, C; \theta_g)}_{g\text{'s private MB}} + \underbrace{\frac{\partial V}{\partial c_g}(c_n, C; \theta_n)}_{g\text{'s external MB}} = a \tag{4.6}$$

$$[c_n] : \underbrace{\frac{\partial V}{\partial c_n}(c_n, C; \theta_n)}_{n\text{'s private MB}} + \underbrace{\frac{\partial V}{\partial c_n}(c_g, C; \theta_g)}_{n\text{'s external MB}} = a, \tag{4.7}$$

ignoring non-negativity constraints and complementary slackness conditions. The first term in each equation is the consumer's private marginal consumption benefit, whereas the second term is the external marginal benefit from an individual's

[5]The budget constraint of the government is revisited in subsequent sections.

provision.

A direct implication of the first-order conditions for the social equilibrium is,

Proposition 2 *In the social equilibrium, (i) green consumers will provide as least as much of the public good as non-green consumers, that is, $c_g^* \geq c_n^*$. And, (ii) the inequality is strict for interior solutions.*

Proof: To prove (i), choose $\bar{c}_g \geq \bar{c}_n$, corresponding to some $\bar{\theta}_g > \bar{\theta}_n$, respectively, as solutions to the SPP

$$\bar{c}_i = \arg\max_{c_i} \left(u_g + u_n | c_i \geq 0\right) \text{ for } i = g, n.$$

Increasing differences implies,

$$\frac{\partial V}{\partial \bar{c}_i}(\bar{c}_i, C; \bar{\theta}_g) - \frac{\partial V}{\partial \bar{c}_i}(\bar{c}_i, C; \bar{\theta}_n) \geq 0 \text{ for } \bar{\theta}_g > \bar{\theta}_n.$$

Thus, we have the following inequalities,

$$\frac{\partial V}{\partial \bar{c}_g}(\bar{c}_g, C; \bar{\theta}_g) - \frac{\partial V}{\partial \bar{c}_g}(\bar{c}_n, C; \bar{\theta}_n) \geq 0$$

and

$$\frac{\partial V}{\partial \bar{c}_n}(\bar{c}_g, C; \bar{\theta}_g) - \frac{\partial V}{\partial \bar{c}_n}(\bar{c}_n, C; \bar{\theta}_n) \geq 0.$$

Adding the previous two inequalities provides,

$$\frac{\partial V}{\partial \bar{c}_g}(\bar{c}_g, C; \bar{\theta}_g) - \frac{\partial V}{\partial \bar{c}_g}(\bar{c}_n, C; \bar{\theta}_n) + \frac{\partial V}{\partial \bar{c}_n}(\bar{c}_g, C; \bar{\theta}_g) - \frac{\partial V}{\partial \bar{c}_n}(\bar{c}_n, C; \bar{\theta}_n) \geq 0.$$

And, by rearranging terms,

$$\frac{\partial V}{\partial \bar{c}_g}(\bar{c}_g, C; \bar{\theta}_g) - \frac{\partial V}{\partial \bar{c}_n}(\bar{c}_n, C; \bar{\theta}_n) \geq \frac{\partial V}{\partial \bar{c}_g}(\bar{c}_n, C; \bar{\theta}_n) - \frac{\partial V}{\partial \bar{c}_n}(\bar{c}_g, C; \bar{\theta}_g) \geq 0. \tag{4.8}$$

Now, let $\bar{c}_g = c_g^*$ and $\bar{c}_n = c_n^*$ such that the previous equation defines relationships within the first-order conditions to the SPP. Equation 4.8 states that both terms that implicitly define the marginal contribution of the green consumer in the social equilibrium are at least as great as that of the non-green consumer. Hence, the green consumer will contribute at least as much to the public good as the non-green consumer, which provides the desired result.

To prove (ii), note that the first two terms in the inequality in Equation 4.8 are strictly positive with $\bar{c}_g, \bar{c}_n > 0$. If $\frac{\partial V}{\partial \bar{c}_g}(\bar{c}_g, C; \bar{\theta}_g) - \frac{\partial V}{\partial \bar{c}_n}(\bar{c}_n, C; \bar{\theta}_n) = 0$, then the marginal benefits of g's own provision would have to equal that of n's. But, this would imply that $\theta_g = \theta_n$, which is a contradiction. Hence, $\bar{c}_g$ must be strictly greater than $\bar{c}_n$. $\square$

Equations 4.6 and 4.7 are insightful for two reasons. First, they internalize the external benefits that arise from any one individual's contribution to the public good. As such, ignoring these benefits results in the well-known fact that the private market equilibrium will under-provide the public good relative to the socially optimal level. Second, any policy that induces privately optimal provision to replicate the socially optimal level of provision will be Pareto improving in the sense of Samuelson (1954). That is, the socially efficient level of private provision of public goods can be expressed by adding the marginal benefits of each consumer and equating them with the marginal cost of provision,

$$\frac{\partial V}{\partial c_g}(c_g, C; \theta_g) + \frac{\partial V}{\partial c_n}(c_n, C; \theta_n) + \frac{\partial V}{\partial c_g}(c_n, C; \theta_n) + \frac{\partial V}{\partial c_n}(c_g, C; \theta_g) = a. \tag{4.9}$$

That the private equilibrium will under-provide the public good is shown simply by noting that the third and fourth terms on the left-hand side of Equation 4.9 are omitted in the private equilibrium and that these terms are nonnegative. Ignoring these terms shifts the aggregate willingness to pay for the public good inward, such that the private equilibrium allocation is less than that of the social equilibrium allocation.

In the private equilibrium, the public good is provided at least cost to each consumer. This is, however, a positive result. The normative solution is to correct this market failure through bargaining or government intervention. As such, free riding is illustrated by the gap between the private equilibrium and the socially optimal provision. Since it is unlikely that consumers possess enough information to bargain in a decentralized fashion, a corrective regulatory instrument is a more realistic device to increase public goods provision. Information limitations are now a problem between the regulator and the consumer, since the consumer holds private information about her preferences. This line of reasoning encapsulates the motivation for designing contracts to elicit true preferences for public goods (Green and Laffont, 1977b; Groves and Ledyard, 1977).

Thus, in the sections that follow, contracts are specified in an attempt to replicate the socially efficient provision of the environmental public good as it accords to the system in Equations 4.6 and 4.7.

4.3 Preference revelation and optimal provision

The primary function of this section is to explore different ways of eliciting information about consumer's preferences through incentives. I approach the information asymmetry between an environmental regulator and a consumer in a mechanism design context. This approach to solving free-riding problems is not new. But, the application of mechanism design in the provision of environmental public goods is novel and the debate on ways to obtain efficient private provision of public goods is not settled (Falkinger et al., 2000).

The revelation principle (Myerson, 1979) allows for a convenient methodological approach within this analysis. Particularly, this principle obviates the need to consider all possible feasible mechanisms, rather we can simply focus our attention on finding direct mechanisms with incentive compatible equilibria. In this section, I rely on Clarke-Groves mechanisms which are truth-revealing in the case of pure public goods and quasi-linear utility (Groves and Ledyard, 1976). Generally, these mechanisms are structured to account for the social externality induced by an individual's reported value for a public project. By internalizing this social benefit, it is a dominant strategy for each consumer to reveal their true value for the public good (Clarke, 1971; Groves and Loeb, 1975). Within the class of privately provided public goods, however, an individual's valuation of the public good is intrinsically linked to the provision from other consumers. Due to the interdependence of agents' valuation and provision of the public good, truth revelation is not a dominant strategy.[6]

This section presents a series of mechanisms that are variants of Clarke-Groves mechanisms adapted to the structure of privately provided public goods. I focus on three similar mechanisms under progressively weaker assumptions on a regulator's informational content. Particularly, I examine cases in which the regulator contracts upon 1) individual provision, 2) reported messages, and 3) reported messages as well as observable aggregate provision. In both the first and third cases, transfers are found such that preferences are revealed truthfully and the socially optimal level of public goods provision is supported by a Nash equilibrium. The second case, while not incentive compatible, motivates the use of transfers conditional on aggregate provision. Finally, the least restrictive contract is scaled up to an I-consumer economy limited to consumers with and without green preferences.

4.3.1 Individually enforceable consumption contracts

Consider a regulator who is assigned the task of maximizing social welfare (as in Equation 4.5) while ignorant of true preferences, θ_i. The role of the regulator

[6] The archetypal Clarke-Groves mechanism defines a transfer for individual i as the sum of everyone else's valuation of the public good evaluated at the socially optimal level with and without i's presence, plus an arbitrary function of everyone else's reported valuation. But, the structure imposed on V in this paper deviates from traditional public goods models since consumer i's valuation is implicitly included in consumers j's valuation of the public good for some $i \neq j$.

is to design an incentive contract, $\mathscr{C}$, for individuals to reveal their true preferences such that the optimal level of public goods provision is obtained. The contract is formally defined as the mapping $\mathscr{C} = [c_i(m_j), T_i(m_j)]$ from Θ into $\mathbb{R}^+ \times \mathbb{R}$ where $c_i(m_j) : \Theta \mapsto \mathbb{R}^+$ are individually enforceable consumption contracts and $T_i(m_j) : \Theta \mapsto \mathbb{R}$ are continuously differentiable transfer functions conditioned on a consumer's strategy space, $m_j \in \{m_g, m_n\}$, which correspond to types. In this mechanism, the regulator can enforce individual contracts but he does not observe the distribution of types. He simply knows that consumers are either green or non-green. In the example of providing energy efficiency as an economic good with private and public environmental benefits, the strong assumption on individually enforceable consumption contracts is justifiable in practice by considering the potential regulatory role of smart-metering technology at the household-level.

The timing of the mechanism in this section works as follows. First, consumers learn their type. Second, the regulator offers a contract conditional on reported types, m_j. Third, consumers report a value of their preference parameters through messages, m_j, that do not necessarily correspond to the value of their true θ_i. Fourth, the contract is executed. This timing represents a standard procedure for implementing direct mechanisms under adverse selection (Laffont and Martimort, 2002).

To facilitate the mechanism, the regulator specifies a payment scheme, $T_i(m_j)$, that individuals receive conditional on their report, m_j. That is, if an individual reports that she is a green consumer by sending m_g, she will incorporate the transfer $T_i(m_g)$ into her net utility function via her budget constraint. Thus, net utility for consumer i sending message j can be written

$$u_i^T = w_i + V(c_i, C; \theta_i) + T_i(m_j) - ac_i \text{ for some } i, j, \tag{4.10}$$

which indicates that a consumer can send a message that does not correspond to her type (i.e., i need not equal j). The consumption bundle that maximizes u_i^T is a set of *individual consumption contracts* defined by $(c_g^T(m_j), c_n^T(m_j))$.

Transfers are specified as

$$T_i(m_j) = \begin{cases} V(c_n(m_{-i}), C; \theta_n) & \text{if } m_j = m_g \\ \\ V(c_g(m_{-i}), C; \theta_g) & \text{if } m_j = m_n \end{cases} \tag{4.11}$$

where T_i is a variant of a Clarke-Groves transfer with individually enforceable consumption contracts. Equation 4.11 indicates the transfer that a consumer receives if she sends the message m_j, which is a function of the report of other consumers, m_{-i}. Individual enforceability ensures that the regulator can observe, and contract upon, individual consumption.

Within this framework, we obtain the following result,

Proposition 3 *Under the transfers in Equation 4.11 and individually enforceable consumption contracts for the two-consumer economy, (i) there is an incentive-compatible Nash equilibrium in which preferences are revealed truthfully and (ii) the socially optimal provision of the public good is obtained.*

Proof: To show (i), assume arbitrary preference parameters $\tilde{\theta}_g > \tilde{\theta}_n$. The regulator solves the SPP for optimal contributions,

$$\tilde{c}_i = \underset{i=g,n}{\arg\max}\{V(\tilde{c}_g, C, \tilde{\theta}_g) + V(\tilde{c}_n, C, \tilde{\theta}_n) - a(\tilde{c}_g + \tilde{c}_n)|\tilde{c}_g, \tilde{c}_n \geq 0\}$$

with $\tilde{c}_g > \tilde{c}_n$, indicating an interior solution.

Under the transfer scheme in Equation 4.11, the payoff for g if she sends m_g is

$$V(\tilde{c}_g, C; \tilde{\theta}_g) + V(\tilde{c}_n, C; \tilde{\theta}_n) - a\tilde{c}_g \tag{4.12}$$

And, the payoff for g if she sends m_n is

$$V(\tilde{c}_n, C; \tilde{\theta}_g) + V(\tilde{c}_g, C; \tilde{\theta}_g) - a\tilde{c}_n \tag{4.13}$$

Assume n sends m_n truthfully. If g prefers to report m_n, this would imply Equation 4.13 is at least as large as Equation 4.12, but this means that $\tilde{c}_n$ cannot be different from $\tilde{c}_g$ in the arbitrary social equilibrium, which is a contradiction. Hence, g can do no better than report m_g since doing so replicates the social optimum. A similar argument holds for n. Thus, reporting truthfully is incentive compatible for both types.

The proof of (ii) is a direct implication of truthful reporting from both types, as the transfers in Equation 4.11 internalize the external marginal benefit of an individual's provision. As such, a consumer's optimal response coincides with solutions to the SPP. □

The intuition for Proposition 3 arises from the fact that the transfers in Equation 4.11 represent the externality arising from each type's provision in the social equilibrium. Since the social optimum maximizes net utility for both types, the transfers are chosen to mimic the social equilibrium without consumer i's presence, thus capturing the external effect of her provision. In this structure, if the other consumer reports truthfully, the individual's maximization problem coincides with the social planner's planner problem. Thus, she can do no better than report truthfully. These best responses conditional on truthful reporting imply socially optimal provision of the public good. This result, however, requires the regulator to fully observe, and contract upon, individual provision, which is a strong assumption in practice.

4.3.2 Reported consumption contracts

To relax individual enforceability, assume that the regulator can only condition transfers on reported levels of provision. Define $\mathscr{C}^R = [c_i^R(m_j), T_i^R(m_j)]$ as a mechanism with the same properties as $\mathscr{C}$ where $c_i^R(m_j) : \Theta \mapsto \mathbb{R}^+$ is a *reported consumption contract*. The practical motivation for this mechanism is that a regulator cannot enforce individual provision under the contract, rather he conditions transfers only on the reported distribution of consumer types. For this mechanism, the timing remains the same as in the previous subsection.

In the two-consumer economy, define individual transfers

$$T_i^R(m_j) = \begin{cases} V(c_n(m_{-i}), C; \theta_n) & \text{if } m_j = m_g \\ \\ V(c_g(m_{-i}), C; \theta_g) & \text{if } m_j = m_n, \end{cases} \tag{4.14}$$

which are incorporated into consumer utility via the budget constraint,

$$u_i^R = w_i + V(c_i, C; \theta_i) + T_i^R(m_j) - ac_i \text{ for some } i, j. \tag{4.15}$$

The consumption bundle that maximizes u_i^R is a set of *reported consumption contracts*, $(c_g^R(m_j), c_n^R(m_j))$.

Under reported consumption and the inability of the regulator to observe individual provision, we have the following result,

Proposition 4 *Under the transfers in Equation 4.14 and reported consumption contracts, reporting truthfully is not incentive compatible.*

Proof: Consider a green consumer's payoff under the transfers in Equation 4.14 if she sends m_g truthfully and acts green,

$$V(c_g, C; \theta_g) + V(c_n, C; \theta_n) - ac_g$$

and her payoff if she send m_n untruthfully, but still acts green,

$$V(c_g, C; \theta_g) + V(c_g, C; \theta_g) - ac_g.$$

By inspection, the latter payoff is strictly larger than the former, hence sending a truthful message is not incentive compatible in this mechanism.□

In this mechanism, the green consumer prefers to report that she is non-green to appropriate a larger payoff. But, it is not optimal for her to act as if she were non-green. Since the regulator cannot enforce individual consumption, g can send message m_n and consume according to her true preferences without penalty. Thus, the mechanism defined by the transfers in Equation 4.14 is not incentive compatible since there are profitable deviations from reporting the truth for the green consumer.

4.3.3 Group and reported consumption contracts

Now consider a regulator who can condition transfers on reported consumption contracts as well as observable aggregate provision of the public good. This mechanism borrows insight from team provision of a public good under moral hazard, which has been motivated primarily by Alchian and Demsetz (1972) and Hölmstrom (1979, 1982). This method has been adopted for environmental applications to provide group incentives for pollution reduction when firm behavior is unobservable (Meran and Schwalbe, 1987; Segerson, 1988; Xepapadeas, 1991). Meran and Schwalbe (1987), for example, construct incentives for meeting an emissions standard when a firm's pollution is not directly observable. Whereas Segerson (1988) extends these incentives to the case in which there is stochasticity in both the (unobservable) abatement actions taken by firms as well as ambient pollution levels. Both rely on group penalties that apply when realized emissions levels exceed some desired level of pollution. But, these models do not explicitly consider markets for public goods.

Define the mechanism, $\mathscr{C}^{GR} = [\{c_i^R(m_j), C(m_j)\}, T_i^{GR}(m_j)]$, with $\{c_i^R(m_j), C(m_j)\} : \Theta \mapsto \mathbb{R}^+$ representing a joint *reported-group consumption contract*. Within this mechanism, a regulator does not observe individual provision, however, he does observe group provision from all consumers, that is the total level of public good provision (e.g., aggregate energy conservation). Given these assumptions on observability, incentives for team provision of a public good apply in a standard moral hazard framework (Alchian and Demsetz, 1972; Hölmstrom, 1979, 1982).

For the mechanism, $\mathscr{C}^{GR}$, the timing is as follows. First, individuals learn about their type. Second, the regulator specifies a menu of transfers conditional on: a) reported messages and b) observable aggregate contributions to the public good. Third, consumers send a message m_j. Fourth, consumers choose privately optimal provision under the contract. Fifth, the contract is executed. The timing of this mechanism follows that of common moral hazard contracts (Laffont and Martimort, 2002).

The primary objective of this mechanism is to elicit true preferences without relying on individual consumption contracts. We augment the transfers in Equation 4.14 by an increasing function τ, with $\tau(0) = 0$, that penalizes deviations from the truth conditional on observable characteristics. Define these new transfer functions,

$$T_i^{GR} = \begin{cases} V(c_n(m_{-i}), C; \theta_n) - \tau(|C(m_j) - C^R|) & \text{if } m_j = m_g \\ \\ V(c_g(m_{-i}), C; \theta_g) - \tau(|C(m_j) - C^R|) & \text{if } m_j = m_n, \end{cases} \tag{4.16}$$

where $\tau(|C(m_j) - C^R|)$ is a function of the absolute difference between observable contracted group provision, $C(m_j)$, and aggregate reported provision, C^R.

In the two-person economy, let $C(m_j)$ represent *group consumption contracts* defined by $C(m_j) = \sum_i c_i(m_j)$ for $i = g, n$ for some contract. While individual consumption contracts are not enforceable, the level of aggregate contracted provision, $C(m_j)$, is observable by the regulator. Finally, define *aggregate reported provision*, C^R, in the two-person economy as the optimal level of provision from both reported consumer types that solves the consumer's problem,

$$C^R = \sum_{i=g,n} c_i(m_j) = \sum_{i=g,n} \arg\max_{c_i} \{V(c_i^R(m_j), C, \theta_i) - ac_i^R(m_j) | c_i \geq 0\} \tag{4.17}$$

where m_j is a consumer's reported type.

The τ function takes advantage of the concept illustrated in the previous subsection—while it is in g's best interest to report untruthfully, she will act according to her true preferences. This notion introduces a wedge between aggregate reported provision, C^R, and aggregate contracted consumption, $C(m_j)$. Since τ is increasing when the two measures of aggregate provision diverge, this "tax" on misreporting provides a condition for which truthful reporting is a best response conditional on truthful reporting of others.

Proposition 5 *Under the transfers in Equation 4.16 and reported-group consumption contracts for the two-consumer economy, (i) there is an incentive-compatible Nash equilibrium in which preferences are revealed truthfully and (ii) the socially optimal provision of the public good is obtained.*

Proof: To prove (i), assume that n sends m_n truthfully. Now, define information rent for g under the transfers in Equation 4.16,

$$\begin{aligned}\Omega_g = &\underbrace{V(c_g, C, \theta_g) + V(c_g, C, \theta_g) - ac_g - \tau(|C(m_j) - C^R|)}_{\text{Payoff if } g \text{ sends } m_n} \\ &- \big[\underbrace{V(c_g, C, \theta_g) + V(c_n, C, \theta_n) - ac_g}_{\text{Payoff if } g \text{ sends } m_g}\big],\end{aligned} \tag{4.18}$$

where the payoff if g sends m_g does not include τ since both consumers have reported truthfully. g's information rent simplifies to,

$$\Omega_g = V(c_g, C, \theta_g) - V(c_n, C, \theta_n) - \tau(|C(m_j) - C^R|). \tag{4.19}$$

Now, choose τ, such that $\tau(|C(m_j) - C^R|) > V(c_g, C, \theta_g) - V(c_n, C, \theta_n)$. This value of τ is sufficient to ensure truthful reporting from g conditional on n reporting truthfully. A similar argument holds for n.

With truthful reporting of both customer types, the rest of the proof for incentive compatibility follows that of Proposition 3.

To prove (ii), note that truthful reporting of both types implies that $\tau = 0$ and thus the individual's problem coincides with the social planner's problem. Hence, socially optimal provision of the public good obtains. □

As an illustration, the payoff for the green consumer reporting truthfully is strictly greater than the payoff from reporting untruthfully if the non-green consumer reports truthfully. Thus, the green consumer, observing this, can do no better than reporting truthfully if the benefit from misreporting is less than the penalty incurred from aggregate provision deviating from contracted provision. Thus, for a sufficiently large τ, the green consumer will report truthfully, which establishes the implementability of an incentive-compatible Nash equilibrium in which both consumers report preferences truthfully.

Socially optimal provision is obtained by noting that with truthful revelation of preferences, aggregate reported provision will be identically equal to aggregate observed provision. Thus, τ will be zero, which results in the consumer's problem coinciding with the social planner's problem.

4.3.4 An I-consumer economy

The previous mechanism $\mathscr{C}^{GR}$ in the two-person economy provides a stylized illustration of the incentives that arise when consumers have heterogeneous preferences for providing a public good. In this section, I generalize the economy to that of I consumers, but restrict consumer types to being either green or non-green. Further, let I_g and I_n, known to the regulator, represent the number of green and non-green consumers, such that the economy is comprised of $I = I_g + I_n$ consumers, with $I_g, I_n > 0$.

First, consider the social planner's problem that characterizes optimal provision of the public good under perfect information,

$$\max_{c_g, c_n} \left\{ I_g \left(w_g + V(c_g, C; \theta_g) - ac_g \right) + I_n \left(w_n + V(c_n, C; \theta_n) - ac_n \right); c_g, c_n \geq 0 \right\} \quad (4.20)$$

resulting in the first-order conditions for an interior optimum,

$$[c_g] : \frac{\partial V}{\partial c_g}(c_g, C; \theta_g) + \left(\frac{I_n}{I_g} \right) \frac{\partial V}{\partial c_g}(c_n, C; \theta_n) = a \quad (4.21)$$

and

$$[c_n] : \frac{\partial V}{\partial c_n}(c_n, C; \theta_n) + \left(\frac{I_g}{I_n} \right) \frac{\partial V}{\partial c_n}(c_g, C; \theta_g) = a \quad (4.22)$$

where the first term in each condition is the private marginal consumption benefit, and the second term is the weighted external marginal benefit of an individual's provision. An immediate implication of the previous first-order conditions is the following result.

Proposition 6 *Socially optimal provision in the I-consumer economy is (i) nondecreasing in the number of other types and (ii) nonincreasing in the number of own types.*

Proof: The proof follows from applying the envelope theorem to the first-order conditions in Equations 4.21 and 4.22. That is, let c_i^* maximize the SPP in the I-consumer economy for consumer type i. Then, $\partial c_i^* / \partial I_i = -I_j I_i^{-2} \partial V(c_j, C, \theta_j) / \partial c_i \leq 0$ for $i \neq j$ since V is nonnegative and $I_i, I_j > 0$, which proves (i). And, $\partial c_i^* / \partial I_j =$

$I_i^{-1}\partial V(c_j, C, \theta_j)/\partial c_i \geq 0$ for $i \neq j$ since V is nonnegative and $I_i, I_j > 0$, which proves (ii). $\square$

This result indicates an important crowding in (out) effect. As the proportion of green consumers, for example, increase in the economy, the burden of providing the public good is decreasing for each individual green consumer. Contrarily, as the proportion of non-green consumers increase, it is optimal for the green consumer to provide more of the public good.

4.3.5 Group and reported contracts in a large economy

Consider the task of a regulator implementing an incentive scheme subject to the information constraints in Section 4.3.3. That is, the regulator observes individual reports and group provision. This subsection considers the ability of a regulator to implement incentive compatible contracts based on reported provision supplemented with a group penalty for deviating from the truth when there are I consumers in the economy. Denote this mechanism $\mathscr{C}^{GR\star} = [\{c_i^{R\star}(m_j), C(m_j)\}, T_i^{R\star}(\hat{\theta}_i)]$. The properties and timing of this mechanism remain the same as that of $\mathscr{C}^{GR}$.

Define the transfer scheme $T_i^{GR\star}(m_j)$ as follows,

$$T_i^{GR\star}(m_j) = \begin{cases} (\hat{I}_g - 1)\big[V(c_g(m_{-i}), C; \theta_g) - ac_g(m_{-i})\big] + & \\ \qquad +\hat{I}_n\big[V(c_n(m_{-i}), C; \theta_n) - ac_n(m_{-i})\big] & \text{if } m_j = m_g \\ (\hat{I}_n - 1)\big[V(c_n(m_{-i}), C; \theta_n) - ac_n(m_{-i})\big] + & \\ \qquad +\hat{I}_g\big[V(c_g(m_{-i}), C; \theta_g) - ac_g(m_{-i})\big] & \text{if } m_j = m_n \end{cases}$$

$$- \tau(|C(m_j) - C^R|) \quad (4.23)$$

where $\hat{I}_i$ for $i = g, n$ is the count of each reported type, which each consumer treats as exogenous since everyone moves simultaneously.

Under these transfers, which internalize both the external marginal benefits of an individual's provision on all other consumers in the economy, there is an analogous result to that of the simpler economy,

Proposition 7 *Under the transfers in Equation 4.23 and reported-group consumption contracts for the I-consumer economy, (i) there is an incentive-compatible Nash equilibrium in which preferences are revealed truthfully and (ii) the socially optimal provision of the public good is obtained.*

Proof: The proof follows that of Proposition 5. To show (i), assume all other consumers report truthfully. Now, fix $\hat{I}_g$ and $\hat{I}_n$ and define information rent for a green consumer under the contract in Equation 4.23,

$$\bar{\Omega}_g = \underbrace{\begin{array}{r} V(c_g, C; \theta_g) - ac_g + (\hat{I}_n - 1)\big[V(c_n, C; \theta_n) - ac_n\big] \\ + \hat{I}_g\big[V(c_g, C; \theta_g) - ac_g\big] - \tau(|C(m_j) - C^R|) \end{array}}_{\text{Payoff if } g \text{ sends } m_n}$$
$$- \underbrace{\begin{array}{r} \Big[V(c_g, C; \theta_g) - ac_g + (\hat{I}_g - 1)\big[V(c_g, C; \theta_g) - ac_g\big] \\ + \hat{I}_n\big[V(c_n, C; \theta_n) - ac_n\big]\Big], \end{array}}_{\text{Payoff if } g \text{ sends } m_g} \tag{4.24}$$

which simplifies to,

$$\bar{\Omega}_g = V(c_g, C; \theta_g) - V(c_n, C; \theta_n) - a(c_g + c_n) - \tau(|C(m_j) - C^R|). \tag{4.25}$$

Now, choose τ such that $\tau(|C(m_j) - C^R|) > V(c_g, C; \theta_g) - V(c_n, C; \theta_n) - a(c_g + c_n)$. This value of τ is sufficient to ensure truthful reporting from g conditional on truthful reporting from all other types. A parallel argument holds for n. Thus, the transfers in Equation 4.23 are incentive compatible and support a Nash equilibrium.

To prove (ii), note that truthful reporting of both types implies that $\tau = 0$ and thus the individual's problem coincides with the social planner's problem. Hence, socially optimal provision of the public good obtains. □

The intuition of this result is analogous to Proposition 5—the transfers in Equation 4.23 internalize the externality of an individual's provision, while a sufficiently high group penalty for misreporting induces truth-revelation.

4.3.6 Budget balancedness

In general, the mechanisms proposed here will not satisfy budget balancing. In other words, since the transfers are structured as subsidies for optimal provision of the public good, the regulator needs to pay consumers the value of the transfers. In this context, it is useful to think of the regulator being a government with a plausibly large budget to finance the transfers. This mechanism, however, will never achieve Pareto efficiency since there will always be a need for the mechanism to be financed from an external source. Thus, the regulator in this context also serves to break the budget balancing constraint as in Holmstrom (1982).

As a practical example, however, the transfers outlined in the previous mechanisms could represent subsidies for household energy efficiency investments. Since the energy efficiency gap is a substantial market inefficiency, there is justification for government intervention to increase the level of public goods provision. In fact, Allcott and Greenstone (2012) tabulate that the U.S. government spends roughly \$3.6 billion annually on demand-side management electricity programs, and \$5.8 billion on energy efficiency tax credits to homeowners, arising from the 2009 economic stimulus package. While this analysis abstracts from the general equilibrium effects of these taxes, these numbers provide a concrete justification for reframing the budget balance constraint as a discussion about correcting environmental market inefficiencies through government intervention.

4.4 Concluding remarks

In this paper, I explore the role of incentives in optimal provision of environmental public goods when agents have heterogeneous preferences. Particularly, I affirm that environmental public goods are underprovided relative to the social optimum in general. I contribute to the literature on public goods provision by identifying equilibrium properties of preference revelation mechanisms under varying degrees of informational restrictions for a regulator. Results are generally positive in that, with a combination of subsidies and credible punishments, there are incentive compatible Nash equilbria that attain socially optimal public goods provision with group incentives.

These results contribute to growing literature in environmental policy that considers the role of voluntary provision of environmental goods. While policy instruments for energy efficiency may include subsidies or tax credits for technology adoption, the optimal policy depends critically on understanding individuals' preferences for providing public goods. Empirical work measuring the distance between current levels of privately provided environmental public goods and the socially optimal level is important for designing effective policies.

Fruitful areas for future work in this line include addressing the role of heterogeneous costs, examining a continuum of preference types, and exploring explicit forms of preferences for public goods. On the conceptual side, a more rigorous analysis of the equilibria proposed in this paper would consider issues of uniqueness and stability.

Appendix A: Budgeting, billing frequency, and consumer demand

As an illustrative motivating example, consider the classic "linear approximate" almost ideal demand system (LA/AIDS) of Deaton and Muellbauer (1980) applied to household utility bundles (e.g., water, electricity, natural gas), written in expenditure share form,

$$S_{it} = \alpha_i + \sum_{j=1}^{n} \gamma_{ij} \ln(p_{jt}) + \beta_i \ln(x_t/P_t^\star), \tag{A.1}$$

where S_{it} is the budget share of good i at budgeting time t, x_t is total expenditures on household utilities, and $P_t^\star$ is a Stone (1954) price index given by the sum of share-weighted prices,[1]

$$\ln(P_t^\star) = \sum_{i=1}^{n} S_{it} \ln(p_{it}). \tag{A.2}$$

I index the system of share equations by a "budgeting time" (t) to explore the time horizon on which households make decisions. In standard models of demand the window of budgeting time is chosen arbitrarily based on the time unit of the available data. However, many consumer demand frameworks in water and electricity

[1]The benefits of the AIDS model presented here are well-known—it provides a first-order approximation to any demand system, it satisfies the axioms of choice exactly, and imposes neither separability nor homotheticity on preferences. As an example in the environmental literature with implications for welfare, see West and Williams (2004).

rely on monthly variation (or infra-monthly variation), to explore within-household decisions in response to price and quantity information. In models of perfect information, the budgeting time is immaterial since consumers are assumed to be sufficiently far-sighted to budget for household utility expenditures perfectly. By relaxing this assumption, the role of billing frequency becomes a question of intra-annual decision making. Further, when household utility expenditures on water, electricity, and natural gas are not made at the same time, the time horizon on which consumers make decisions may directly influence demand.

As a thought experiment, assume that a representative household receives monthly bills for consumption of electricity (Q_E) and natural gas (Q_G), while bills for water use (Q_W) may arrive monthly or bi-monthly. The prices for each unit of consumption are given by P_{it} for $i = W, E, G$ in any budgeting time, t, and remain fixed. For a consumer who makes decisions on a monthly budgeting horizon, which may be a reasonable assumption since aggregate utility expenditures comprise a relatively small portion of total expenditures, define the budget shares S_{it} for $i = W, E, G$.[2] By exponentiating (and suppressing t subscripts), we can write Equation A.2 as

$$P^{\star} = P_W^{S_W} P_E^{S_E} P_G^{1-S_W-S_E} \tag{A.3}$$

so that it is apparent that the price index used to normalize expenditures in Equation A.1 depends not only on the price level, but also the budget share of each expenditure in the consumption bundle. It follows that a consumer who is making consumption

[2] As additional motivation, empirical researchers rely almost exclusively on exploiting month-to-month variation in water and electricity use to identify price and income elasticities, as well as the effect of non-pecuniary interventions to induce conservation.

decisions over household utilities may respond differently to a bundle of *monthly* bills for all three goods than a consumer who receives monthly bills for electricity and gas, but *bi-monthly* bills for water. In the event of receiving a bi-monthly water bill alongside monthly electric and gas bills, the expenditure share of water will be higher than that of its monthly counterpart simply because a bi-monthly bill is larger by construction. Thus, two identical households consuming water at the same price at different billing frequencies may exhibit different price responsiveness in a classical demand system along a fixed budgeting horizon.

To illustrate this thought experiment further, I estimate the demand system above with observed water consumption and prices, and simulated electric and natural gas demand. I collapse water billing data in the calendar year of 2012 for a subset of households and construct monthly and bi-monthly billing equivalents. I make the assumption that water and gas are billed monthly throughout the year, which is true for Durham, NC. If households are billed monthly for all three goods, the mean expenditure shares for water, electric, and gas are 0.200, 0.544, and 0.256, respectively.[3] If water is billed bi-monthly, however, the expenditure shares are 0.327, 0.457, and 0.215, while price levels remain the same in each case. Estimating demand systems for both scenarios provides an own-price elasticity estimate of -1.367 in the bi-monthly scenario relative to -1.224 for the monthly billing (the former is more than 11% larger). Both parameters are precisely estimated such that the elasticities are statistically different. Elasticity estimates and standard errors are presented in Table A.1 and the elasticity formula is given in Equation A.6. The

[3] See Table A.1 for summary statistics and select estimation results.

key takeaway from this exercise is simply to show that under constant prices, a consumer planning expenditures on a monthly budgeting horizon may respond inconsistently with neoclassical demand models of perfect information when moved from a bi-monthly to a monthly billing scenario. While these parameters are simply illustrative, and should not be used for policy conclusions, they provide motivation for examining the effect of changes in billing frequency on consumer misperception of prices and quantities in decision making.

To estimate the LA/AIDS model presented in the previous subsection, I collapse Durham's billing data to an annual level for all billing districts for the calendar year of 2012. I then determine monthly and bi-monthly averages of quantities consumed and expenditures for each household. Quantities demanded and average prices are derived directly from the billing data set. I use average monthly electricity and average annual natural gas consumption in North Carolina obtained from EIA and AGA tables for 2012. Natural gas consumption is converted to monthly use. I append this information with average market prices for residential electricity and gas use in North Carolina for 2012.

To induce variation in both electricity and natural gas consumption and prices, as well as preserve a negative correlation between consumption and prices, I sample N random draws from the multivariate normal distribution,

$$\begin{pmatrix} Q_E \\ P_E \end{pmatrix} \sim \mathscr{N}\left[\begin{pmatrix} 1{,}077 \text{ kWh/month} \\ 0.1019 \text{ \$/kWh} \end{pmatrix}, \begin{pmatrix} 10{,}000 & -0.2 \\ -0.2 & 0.0001 \end{pmatrix}\right] \tag{A.4}$$

Table A.1: LA/AIDS simulation results for monthly and bi-monthly utility budgeting scenarios

	Bi-monthly water (monthly electricity and gas)	Monthly Utilities
Mean expenditure share:		
Water	0.327	0.200
	(0.104)	(0.078)
Electricity	0.457	0.544
	(0.075)	(0.060)
Gas	0.215	0.256
	(0.036)	(0.031)
Mean price:		
Water ($/cf)	10.25	10.25
	(6.09)	(6.09)
Electricity (cents/kWh)	0.1018	0.1018
	(0.001)	(0.001)
Gas ($/mmbtu)	12.18	12.18
	(0.71)	(0.71)
Price elasticity of water demand	-1.367	-1.224
	[0.021]	[0.013]
Observations:	12,879	12,879

Note: Standard deviations in parentheses. Estimated standard errors in square brackets. Price elasticity estimates shown are functions of parameters from estimating a simultaneous LA/AIDS demand system for water and electricity demand (excluding natural gas). Electricity and natural gas consumption and prices are simulated as described in this appendix. Water consumption is annualized data for 2012 for households who never transitioned to monthly billing. Elasticity estimates are meant to be purely illustrative and not to be interpreted for policy conclusions.

to represent simulated electricity demand, and

$$\begin{pmatrix} Q_G \\ P_G \end{pmatrix} \sim \mathscr{N}\left[\begin{pmatrix} 4.2083 \text{ 1,000cf/month} \\ 12.19 \text{ \$/1,000cf} \end{pmatrix}, \begin{pmatrix} 0.1 & -0.2 \\ -0.2 & 0.5 \end{pmatrix}\right] \tag{A.5}$$

to represent simulated natural gas demand. The number of draws, N, corresponds to the number of households in the limited "control" sample (13,645).

Equation A.1 is estimated for two share equations (water and electricity, with natural gas omitted) in a seemingly unrelated regression framework, imposing homogeneity and symmetry on the demand system. The parameter of interest is the uncompensated own-price elasticity of demand for water, which is given by its AIDS counterpart,

$$\eta_{ij} = -\delta_{ij} + \gamma_{ij}/S_i + \beta_i \alpha_i / S_i - \frac{\beta_i}{S_i} \sum_k \gamma_{kj} \ln(P_k) \tag{A.6}$$

where δ_{ij} is the Kronecker delta that equals one if $i = j$ and zero otherwise, the γ and β parameters are estimated in the system, S_i is the sample mean expenditure share for good i, and P_k are sample mean prices.[4]

[4]Green and Alston (1990) outline the potential bias of using AIDS elasticity formulas with LA/AIDS parameters, however the proper LA/AIDS formulas for elasticities include a full system of own- and cross-price elasticities that are not likely to be meaningful for the simulated electricity and gas markets. As such, I use the AIDS formula as a useful, and strictly illustrative, approximation.

Appendix B: Additional figures

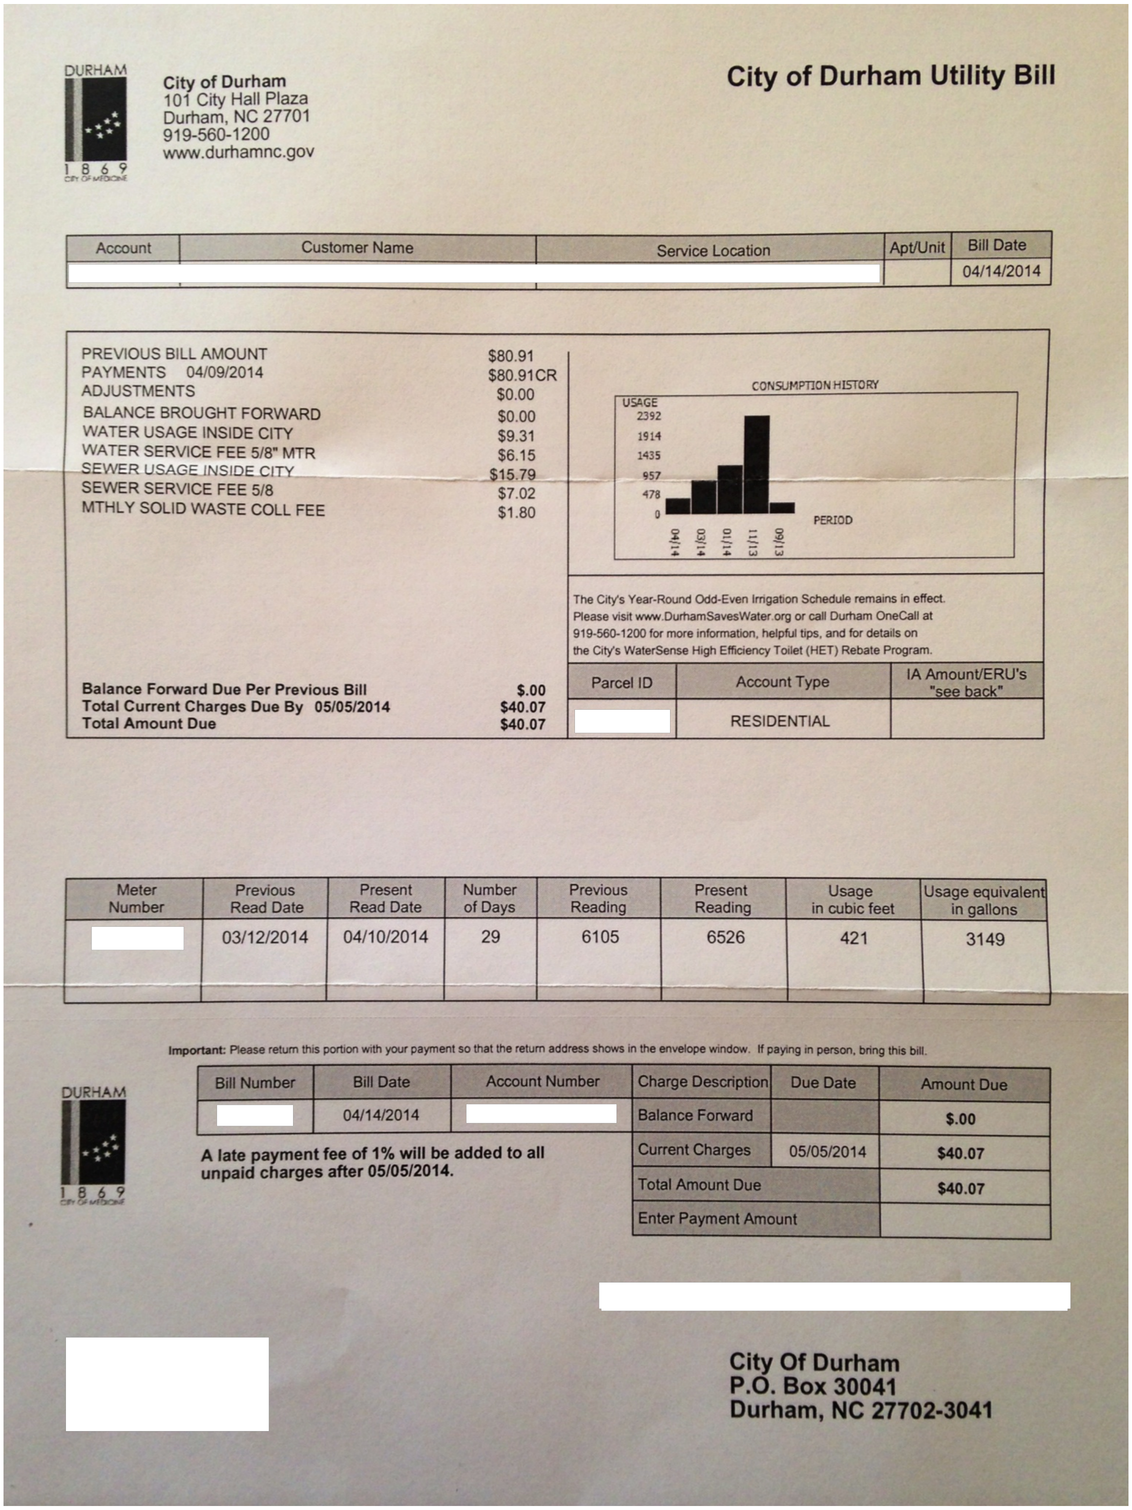

DURHAM 1869 CITY OF MEDICINE

City of Durham
101 City Hall Plaza
Durham, NC 27701
919-560-1200
www.durhamnc.gov

City of Durham Utility Bill

Account	Customer Name	Service Location	Apt/Unit	Bill Date
[redacted]				04/14/2014

PREVIOUS BILL AMOUNT	$80.91
PAYMENTS 04/09/2014	$80.91CR
ADJUSTMENTS	$0.00
BALANCE BROUGHT FORWARD	$0.00
WATER USAGE INSIDE CITY	$9.31
WATER SERVICE FEE 5/8" MTR	$6.15
SEWER USAGE INSIDE CITY	$15.79
SEWER SERVICE FEE 5/8	$7.02
MTHLY SOLID WASTE COLL FEE	$1.80
Balance Forward Due Per Previous Bill	**$.00**
Total Current Charges Due By 05/05/2014	**$40.07**
Total Amount Due	**$40.07**

The City's Year-Round Odd-Even Irrigation Schedule remains in effect. Please visit www.DurhamSavesWater.org or call Durham OneCall at 919-560-1200 for more information, helpful tips, and for details on the City's WaterSense High Efficiency Toilet (HET) Rebate Program.

Parcel ID	Account Type	IA Amount/ERU's "see back"
[redacted]	RESIDENTIAL	

Meter Number	Previous Read Date	Present Read Date	Number of Days	Previous Reading	Present Reading	Usage in cubic feet	Usage equivalent in gallons
[redacted]	03/12/2014	04/10/2014	29	6105	6526	421	3149

Important: Please return this portion with your payment so that the return address shows in the envelope window. If paying in person, bring this bill.

DURHAM 1869 CITY OF MEDICINE

Bill Number	Bill Date	Account Number
[redacted]	04/14/2014	[redacted]

A late payment fee of 1% will be added to all unpaid charges after 05/05/2014.

Charge Description	Due Date	Amount Due
Balance Forward		**$.00**
Current Charges	05/05/2014	**$40.07**
Total Amount Due		**$40.07**
Enter Payment Amount		

[redacted]

[redacted]

City Of Durham
P.O. Box 30041
Durham, NC 27702-3041

Figure B.1: Example of first monthly water bill for the City of Durham Water Utility.

DURHAM

1869

CITY OF MEDICINE

CITY OF DURHAM
DEPARTMENT OF WATER MANAGEMENT
101 CITY HALL PLAZA • DURHAM, NC 27701
919-560-4381 • FAX 919-560-4479

September 4, 2012

Customer Name
Customer Address
Customer City, State Zip

Service Address:

Dear Valued Customer:

The City of Durham is transitioning to billing for water/sewer services on a monthly basis. For the past several years you have been receiving bills every other month. Starting in October, you will begin receiving a bill monthly.

This will benefit you by reducing the amount you need to pay at one time, and by shortening the period when leaks or other problems may be discovered.

Another change is that the City will no longer be sending out "friendly reminder" letters if your payment is not received prior to the due date. In that case, you will see a past due balance in bold letters at the top of the bill. If your payment for any prior month is not received by the due date for that bill, your water service may be disconnected even though your current bill is not yet due.

The City will still send disconnection letters and provide telephone reminders prior to disconnection for nonpayment. To make sure you receive these notices, please notify the City at once if you have any change in your mailing address or phone number.

If you have any questions or concerns, please call [redacted] or e-mail [redacted]. We appreciate this opportunity to improve our service to you.

Sincerely,

Department of Water Management

City of Durham

Good Things Are Happening In Durham

Figure B.2: Example of monthly billing notification received at least six weeks before transition to monthly billing.

Appendix C: Derivation of demand functions under quantity misperception

The consumer's problem is

$$\max_{w}\{x + a(\lambda w)^{1/\gamma+1}\} \text{ subject to } M = x + \lambda p w, \tag{C.1}$$

which provides the following necessary condition for an interior solution,

$$\left(\frac{a}{\gamma+1}\right)(\lambda w)^{-\gamma/\gamma+1} = \lambda p. \tag{C.2}$$

Equation C.2 can be rearranged to represent perceived demand as

$$\tilde{w}(p,\lambda) = Ap^{\frac{1-\gamma}{\gamma}}\lambda^{\frac{1-2\gamma}{\gamma}}\text{, where } A \equiv \left(\frac{1+\gamma}{a}\right)^{\frac{1-\gamma}{\gamma}}, \tag{C.3}$$

which shows that demand is scaled multiplicatively by a function of λ since $w(p) = Ap^{\frac{1-\gamma}{\gamma}}$. Thus, the perceived demand function can be written $\lambda^{\frac{1-2\gamma}{\gamma}}w(p)$.

Further, we can derive the price elasticity of perceived demand,

$$\eta = \frac{\partial \tilde{w}(p,\lambda)}{\partial p}\frac{p}{\tilde{w}(p,\lambda)} = \left(\frac{1-\gamma}{\gamma}\right)Ap^{\frac{1-\gamma}{\gamma}-1}\lambda^{\frac{1-2\gamma}{\gamma}}\frac{p}{\tilde{w}(p,\lambda)} = \frac{1-\gamma}{\gamma}, \tag{C.4}$$

which shows that elasticity is constant across prices and levels of λ.

Combining Equations C.4 and C.3 allows for perceived demand to be written succinctly,

$$\tilde{w}(p, \lambda) = \lambda^{\frac{1-2\gamma}{\gamma}} w(p) = \lambda^{\eta-1} w(p). \tag{C.5}$$

Bibliography

Abrevaya, Jason, Yu-Chin Hsu, and Robert P. Lieli, "Estimating conditional average treatment effects," *Journal of Business and Economic Statistics*, 2014, *Forthcoming.*

Alchian, Armen A. and Harold Demsetz, "Production, Information Costs, and Economic Organization," *The American Economic Review*, 1972, *62* (5), 777–795.

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

_ , "The Welfare Effects of Misperceived Product Costs: Data and Calibrations from the Automobile Market," *American Economic Journal: Economic Policy*, August 2013, *5* (3), 30–66.

_ **and Dmitry Taubinsky**, "Evaluating Behaviorally-Motivated Policy: Experimental Evidence from the Lightbulb Market," *American Economic Review*, 2015, *Forthcoming.*

_ **and Michael Greenstone**, "Is there an energy efficiency gap?," *Journal of Economic Perspectives*, 2012, *26* (1), 3–28.

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

Andreoni, James, "Impure altruism and donations to public goods: A theory of warm-glow giving," *The Economic Journal*, 1990, *100* (401), 464—477.

Angrist, Joshua D. and Jörn-Steffen Pischke, *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton University Press, 2009.

Arbués, Fernando, Mara Ángeles Garca-Valiñas, and Roberto Martnez-Espiñeira, "Estimation of residential water demand: A state-of-the-art review," *The Journal of Socio-Economics*, March 2003, *32* (1), 81–102.

Bergstrom, Theodore, Lawrence Blume, and Hal Varian, "On the private provision of public goods," *Journal of Public Economics*, 1986, *29*, 25–49.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How Much Should We Trust Differences-in-Differences Estimates?," *The Quarterly Journal of Economics*, February 2004, *119* (1), 249–275.

Besley, Timothy and Maitreesh Ghatak, "Competition and incentives with motivated agents," *American Economic Review*, 2005, *95* (3), 616–636.

Bollinger, Bryan, Phillip Leslie, and Alan Sorensen, "Calorie Posting in Chain Restaurants," *American Economic Journal: Economic Policy*, February 2011, *3* (1), 91–128.

Borenstein, Severin, "To what electricity price do consumers respond? Residential demand elasticity under increasing-block pricing," *University of California, Berkeley, Working Paper*, 2009.

Brent, Daniel A., Joseph Cook, and Skylar Olsen, "Norms and water conservation: How do effects vary across and within utilities?," *Working Paper*, 2014.

Caplin, Andrew and Mark Dean, "Revealed Preference, Rational Inattention, and Costly Information Acquisition," Working Paper 19876, National Bureau of Economic Research January 2014.

Charness, Gary and Matthew Rabin, "Understanding social preferences with simple tests," *The Quarterly Journal of Economics*, 2002, *117* (3), 817–869.

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

Clarke, Edward H., "Multipart pricing of public goods," *Public Choice*, 1971, *11* (1), 17–33.

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

Dalhuisen, Jasper M., Raymond JGM Florax, Henri LF de Groot, and Peter Nijkamp, "Price and income elasticities of residential water demand: A meta-analysis," *Land Economics*, 2003, *79* (2), 292–308.

Davidoff, Amy, Linda Blumberg, and Len Nichols, "State health insurance market reforms and access to insurance for high-risk employees," *Journal of Health Economics*, 2005, *24* (4), 725–750.

Deaton, Angus and John Muellbauer, "An almost ideal demand system," *American Economic Review*, 1980, pp. 312–326.

Espey, M., J. Espey, and W. D. Shaw, "Price elasticity of residential demand for water: A meta-analysis," *Water Resources Research*, June 1997, *33* (6), 1369–1374.

Falkinger, Josef, "Efficient private provision of public goods by rewarding deviations from average," *Journal of Public Economics*, 1996, *62*, 413–422.

_ **, Ernst Fehr, Simon Gächter, and Rudolf Winter-Ebmer**, "A Simple Mechanism for the Efficient Provision of Public Goods: Experimental Evidence," *The American Economic Review*, 2000, *90* (1), 247–264.

Ferraro, Paul J. and Juan José Miranda, "Heterogeneous treatment effects and mechanisms in information-based environmental policies: Evidence from a large-scale field experiment," *Resource and Energy Economics*, 2013, *35* (3), 356–379.

_ **and Michael K. Price**, "Using nonpecuniary strategies to influence behavior: Evidence from a large-scale field experiment," *The Review of Economics and Statistics*, 2013, *95* (1), 247–264.

_ **, Juan Jose Miranda, and Michael K. Price**, "The Persistence of Treatment Effects with Norm-Based Policy Instruments: Evidence from a Randomized Environmental Policy Experiment," *American Economic Review, Papers and Proceedings*, 2011, *101* (3), 318–322.

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

Foster, William and Richard E. Just, "Measuring welfare effects of product contamination with consumer uncertainty," *Journal of Environmental Economics and Management*, 1989, *17*, 266–283.

Gabaix, Xavier and David Laibson, "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets," *The Quarterly Journal of Economics*, May 2006, *121* (2), 505–540.

Gans, Will, Anna Alberini, and Alberto Longo, "Smart meter devices and the effect of feedback on residential electricity consumption: Evidence from a natural experiment in Northern Ireland," *Energy Economics*, 2013, *36*, 729–743.

Gilbert, Ben and Joshua S. Graff Zivin, "Dynamic salience with intermittent billing: Evidence from smart electricity meters," *Journal of Economic Behavior & Organization*, 2014, *107*, 176–190.

Glaeser, Edward L., "The supply of environmentalism: Psychological interventions and economics," *Review of Environmental Economics and Policy*, 2014, *8* (2), 208–229.

Goulder, Lawrence H. and Roberton C. Williams, "The substantial bias from ignoring general equilibrium effects in estimating excess burden, and a practical solution," *Journal of Political Economy*, 2003, *111* (4), 898—927.

Green, Jerry and Jean-Jacques Laffont, "Characterization of satisfactory mechanisms for the revelation of preferences for public goods," *Econometrica*, 1977, *45* (2), 427–438.

_ **and** _ , "On the revelation of preferences for public goods," *Journal of Public Economics*, 1977, *8* (1), 79–93.

Green, Richard and Julian M. Alston, "Elasticities in AIDS models," *American Journal of Agricultural Economics*, 1990, *72* (2), 442–445.

Grembi, Veronica, Tommaso Nannicini, and Ugo Troiano, "Policy responses to fiscal restraints: A difference-in-discontinuities design," Technical Report 3999, CESifo Working Paper: Public Finance 2012.

Groves, Theodore and John Ledyard, "Some limitations of demand revealing processes," *Public Choice*, 1976.

_ **and** _ , "Optimal Allocation of Public Goods: A Solution to the "Free Rider" Problem," *Econometrica*, 1977, *45* (4), 783–809.

_ **and Martin Loeb**, "Incentives and public inputs," *Journal of Public Economics*, 1975, *4* (3), 211–226.

Grubb, Michael D. and Matthew Osborne, "Cellular service demand: Biased beliefs, learning, and bill shock," *American Economic Review*, 2015, *105* (1), 234–271.

Gruber, Jonathan, "The incidence of mandated maternity benefits," *The American Economic Review*, 1994, pp. 622–641.

Harberger, Arnold C., "The Measurement of Waste," *American Economic Review, Papers and Proceedings*, 1964, *54* (3), 58–76.

Harding, Matthew and Alice Hsiaw, "Goal setting and energy conservation," *Journal of Economic Behavior & Organization*, 2014, *107*, 209–227.

Hausman, Jerry A., "Exact Consumer's Surplus and Deadweight Loss," *American Economic Review*, 1981, *71* (4), 662–676.

Hewitt, Julie A. and W. Michael Hanemann, "A Discrete/Continuous Choice Approach to Residential Water Demand under Block Rate Pricing," *Land Economics*, 1995, *71* (2), 173–192.

Hölmstrom, Bengt, "Moral Hazard and Observability," *The Bell Journal of Economics*, 1979, *10* (1), 74–91.

— , "Moral Hazard in Teams," *The Bell Journal of Economics*, 1982, *13* (2), 324–340.

Houde, Sébastien, "How Consumers Respond to Environmental Certification and the Value of Energy Information," Working Paper 20019, National Bureau of Economic Research March 2014.

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

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

Jacobsen, Grant D., Matthew J. Kotchen, and Michael P. Vandenbergh, "The behavioral response to voluntary provision of an environmental public good: Evidence from residential electricity demand," *European Economic Review*, 2012, *56* (5), 946–960.

Jacobsen, Mark, Jacob LaRiviere, and Michael Price, "Public Goods Provision in the Presence of Heterogeneous Green Preferences," *NBER Working Paper w20266*, 2014.

Jessoe, Katrina and David Rapson, "Knowledge is (less) power: Experimental evidence from residential energy use," *American Economic Review*, 2014, *104* (4), 1417–1438.

Jin, Ginger Zhe and Phillip Leslie, "The Effect of Information on Product Quality: Evidence from Restaurant Hygiene Grade Cards," *The Quarterly Journal of Economics*, 2003, *118* (2), 409–451.

Jr., Henry S. Foster and Bruce R. Beattie, "On the Specification of Price in Studies of Consumer Demand under Block Price Scheduling," *Land Economics*, 1981, *57* (4), 624–629.

Just, Richard E., "Behavior, Robustness, and Sufficient Statistics in Welfare Measurement," *Annual Review of Resource Economics*, 2011, *3* (1), 37–70.

Kahn, Matthew E., "Do greens drive Hummers or hybrids? Environmental ideology as a determinant of consumer choice," *Journal of Environmental Economics and Management*, 2007, *54* (2), 129–145.

— **and Frank A. Wolak**, "Using Information to Improve the Effectiveness of Nonlinear Pricing: Evidence from a Field Experiment," *Working Paper*, 2013.

Kirchsteiger, Georg and Clemens Puppe, "On the possibility of efficient private provision of public goods through government subsidies," *Journal of Public Economics*, 1997, *66* (489–504).

Klaiber, H. Allen, V. Kerry Smith, Michael Kaminsky, and Aaron Strong, "Measuring Price Elasticities for Residential Water Demand with Limited Information," *Land Economics*, 2014, *90* (1), 100–113.

Kotchen, Matthew J., "Impure public goods and the comparative statics of environmentally friendly consumption," *Journal of Environmental Economics and Management*, 2005, *49* (2), 281–300.

_ , "Green markets and private provision of public goods," *Journal of Political Economy*, 2006, *114* (4), 816—834.

_ **and Michael R. Moore**, "Private provision of environmental public goods: Household participation in green-electricity programs," *Journal of Environmental Economics and Management*, 2007, *53* (1), 1–16.

Krinsky, Itzhak and A. Leslie Robb, "On approximating the statistical properties of elasticities," *The Review of Economics and Statistics*, 1986, *68* (4), 715–719.

Laffont, Jean-Jacques and David Martimort, *The theory of incentives: The principal-agent model*, Princeton University Press, 2002.

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

Li, Shanjun, Joshua Linn, and Erich Muehlegger, "Gasoline taxes and consumer behavior," *American Economic Journal: Economic Policy*, 2014, *6* (4), 392–342.

Liebman, Jeffrey B. and Richard Zeckhauser, "Schmeduling," *Unpublished manuscript*, 2004.

Ludwig, Jens, Jeffrey R Kling, and Sendhil Mullainathan, "Mechanism Experiments and Policy Evaluations," *Journal of Economic Perspectives*, 2011, *25* (3), 17–38.

Mansur, Erin T. and Sheila M. Olmstead, "The value of scarce water: Measuring the inefficiency of municipal regulations," *Journal of Urban Economics*, May 2012, *71* (3), 332–346.

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

Meran, Georg and Ulrich Schwalbe, "Pollution Control and Collective Penalties," *Journal of Institutional and Theoretical Economics*, 1987, *143* (4), 616–629.

Myerson, Roger B., "Incentive compatibility and the bargaining problem," *Econometrica*, 1979, *47* (1), 61–74.

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

Nickell, Stephen, "Biases in dynamic models with fixed effects," *Econometrica*, 1981, *49* (6), 1417–1426.

Nieswiadomy, Michael L. and David J. Molina, "Comparing Residential Water Demand Estimates under Decreasing and Increasing Block Rates Using Household Data," *Land Economics*, 1989, *65* (3), 280–289.

Olmstead, Sheila M. and Robert N. Stavins, "Comparing price and nonprice approaches to urban water conservation," *Water Resources Research*, 2009, *45* (4).

_ , W. Michael Hanemann, and Robert N. Stavins, "Water demand under alternative price structures," *Journal of Environmental Economics and Management*, 2007, *54* (2), 181–198.

Orange Water and Sewer Authority, "Our new water and sewer rates," 2007.

Renwick, Mary E. and Richard D. Green, "Do residential water demand side management policies measure up? An analysis of eight California water agencies," *Journal of Environmental Economics and Management*, 2000, *40*, 37–55.

_ and Sandra O. Archibald, "Demand side management policies for residential water use: Who bears the conservation burden?," *Land Economics*, August 1998, *74* (3), 343–359.

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

Samuelson, Paul A., "The Pure Theory of Public Expenditure," *The Review of Economics and Statistics*, 1954, *36* (4), 387–389.

Segerson, Kathleen, "Uncertainty and incentives for nonpoint pollution control," *Journal of Environmental Economics and Management*, 1988, *15* (1), 87–98.

Sexton, Steven E., "Automatic bill payment and salience effects: Evidence from electricity consumption," *The Review of Economics and Statistics*, 2014, *In press.*

_ and Alison L. Sexton, "Conspicuous conservation: The Prius halo and willingness to pay for environmental bona fides," *Journal of Environmental Economics and Management*, 2014, *67* (3), 303–317.

Shin, Jeong-Shik, "Perception of Price When Price Information Is Costly: Evidence from Residential Electricity Demand," *The Review of Economics and Statistics*, November 1985, *67* (4), 591–598.

Stigler, George J., "The economics of information," *The Journal of Political Economy*, 1961, *69* (3), 213–225.

Stone, Richard, "Linear Expenditure Systems and Demand Analysis: An Application to the Pattern of British Demand," *The Economic Journal*, 1954, *64* (255), 511–527.

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

_ **and V. Kerry Smith**, "Reconsidering the economics of demand analysis with kinked budget constraints," *Land Economics*, 2010, *86* (1), 173–190.

Timmins, Christopher, "Measuring the dynamic efficiency costs of regulators' preferences: Municipal water utilities in the arid west," *Econometrica*, 2002, *70* (2), 603–629.

Varian, Hal, "Sequential contributions to public goods," *Journal of Public Economics*, 1994, *53* (2), 165–186.

West, Sarah E. and Roberton C. Williams, "Estimates from a consumer demand system: implications for the incidence of environmental taxes," *Journal of Environmental Economics and Management*, May 2004, *47* (3), 535–558.

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

_ **, Laura O. Taylor, and Roger H. von Haefen**, "Conservation policies: Who responds to prices and who responds to prescription?," Working Paper 20466, National Bureau of Economic Research 2014.

Xepapadeas, A. P., "Environmental Policy under Imperfect Information: Incentives and Moral Hazard," *Journal of Environmental Economics and Management*, 1991, *20*, 113–126.