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

Casey J. Wichman†
University of Maryland, College Park

January 9, 2015

Abstract

In this study, I examine the causal effect of information provision on consumer behavior. When consumers make decisions with less than real-time information, they may misperceive price and quantity consumed. I exploit an exogenous transition from bi-monthly to monthly billing for residential water customers to explore the demand response to more frequent information. In contrast to previous studies, I find that customers *increase* consumption by approximately five percent as billing information increases. This result is reconciled in a model of price or quantity uncertainty, where the increase in billing frequency allows consumers to improve their perception of price and quantity consumed. Within this framework, I calculate welfare changes using treatment effects as sufficient statistics. My results suggest that increases in information provide consumers with welfare gains equivalent to 0.5 to 1 percent of annual expenditures on water. I find important heterogeneity in the response to billing frequency, suggesting an increase in nonessential uses of water.

JEL Codes: D12, D61, H42, L95, Q21, Q25
Keywords: information provision, billing frequency, price perception, quantity uncertainty, natural experiment, water demand, water conservation, welfare

*I am grateful for discussions with Rob Williams, Lint Barrage, Maureen Cropper, Richard Just, Chris Timmins, Kerry Smith, Jacob LaRiviere, Steve Sexton, Lesley Turner, Laura Taylor, Roger von Haefen, Anna Alberini, and Ginger Jin that greatly improved this manuscript. Additionally, I thank seminar and conference participants at the 2014 Southern Economics Association, the 2014 AAEA Annual Meetings, the 11th Annual International Water Resource Economics Consortium meetings, the CU Boulder Environmental and Resource Economics Workshop, the Heartland Environmental and Resource Economics Workshop, Duke University, and North Carolina State University for their insightful comments and suggestions.

†*Author email:* wichman@umd.edu. *Author address:* 2200 Symons Hall, Department of Agricultural and Resource Economics, University of Maryland, College Park.

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 or quantity 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 open empirical question of growing interest. A recent vein of literature suggests that consumers tend to underestimate prices and quantities consumed that are transmitted opaquely or allow for customer inattention (Hossain and Morgan, 2006; Chetty et al., 2009; Finkelstein, 2009; Allcott, 2011a; Li et al., 2012; Grubb and Osborne, 2015; Sexton, 2014; Grubb, 2014; Jessoe and Rapson, 2014)[1]. 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 (Schultz et al., 2007; Allcott, 2011b; Ferraro et al., 2011; Ferraro and Price, 2013; Ferraro and Miranda, 2013; Costa and Kahn, 2013; Allcott and Rogers, 2014; Schultz et al., 2014; Brent et al., 2014) and information provision (Gans et al., 2013; Jessoe and Rapson, 2014; Gilbert and Graff Zivin, 2014) and shows that informative interventions can reduce consumption and thus serve as an instrument of conservation. Generally, the provision of information in electricity and water demand is found to be effective in reducing consumption immediately, though evidence suggests that the effect of one-time interventions wanes over time.

With a few notable exceptions, previous empirical findings imply that various forms of information provision can be utilized to reduce consumption of economic goods that impose external costs on society. Additionally, 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. The majority of the extant literature on informative interventions as a tool for conservation, however, stops short of estimating changes in welfare.

This paper makes several conceptual and empirical 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 drought-prone 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 the majority of extant studies and has significant implications for efficient management of scarce environmental resources (c.f., Jessoe and Rapson (2014); Gilbert and Graff Zivin (2014); Sexton

[1] See DellaVigna (2009) for a broad survey of research in this field.

(2014)). 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.

These surprising empirical results necessitate a closer look at the mechanism driving consumer behavior in response to more frequent information. To that effect, I motivate two independent, but not mutually exclusive, conceptual models of imperfect price perception and imperfect quantity perception that reconcile my empirical findings with the current literature on salience and inattention.[2] Based on the notion that consumers are receiving more frequent information about both prices and quantities with the receipt of monthly (versus bi-monthly) bills, I posit a simple theory of price and quantity uncertainty that depends on a consumer's initial misperceptions.[3] The framework is general enough to accommodate the findings of previous research since more frequent information brings consumers closer to the neoclassical ideal of perfect information about prices and levels of consumption. 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, based on the notion that more information will simply 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 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 of my study is a transition of residential water customers from bi-monthly to monthly billing within a single water utility. Beginning in 2011, the City of Durham Water Utility in North Carolina transitioned residential customers in geographically differentiated billing districts to monthly billing over the course of four-and-a-half years. During the transition, customers billed monthly were charged exactly half of bi-monthly service fees while marginal prices for each unit of consumption remained the same. This transition provides an ideal natural experiment for examining consumer behavior in response to more frequent information since billing frequency is the only factor changing among neighbors in different billing districts. Treatment and control households faced identical marginal prices, weather patterns, and other contemporaneous utility-specific shocks.

Billing data for approximately 59,000 residential water customers between February 2009 and June 2014 are matched with tax assessor and Census data. By exploiting the assignment of monthly

[2]Note, however, that the attribution of a specific mechanism to consumer behavior in this empirical setting is nearly impossible. Thus, I motivate two plausible models of demand, but I cannot refute the likelihood of other mechanisms driving these results.

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

billing, I estimate an average treatment effect on water consumption due to increased billing frequency. The primary result is that households billed monthly consume roughly 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 nonessential water use. I also find important heterogeneity among other dimensions, such as baseline water use, lot size, and assessed value of the home.

From an environmental policy perspective, informative signals are being used increasingly as a regulatory instrument, particularly in the context of electricity and water conservation. The findings of this paper suggest that changes in billing frequency can have the perverse effect of increasing consumption. This result is particularly poignant because standard economic theory stipulates that the efficient price for residential water is its long-run marginal cost of provision (Timmins, 2002; Olmstead and Stavins, 2009; Olmstead, 2010; Mansur and Olmstead, 2012). Since the market price is likely set below its efficient level, however, the demand response to increased information provision may exacerbate the wedge between privately and socially optimal consumption levels. Finally, for water utility managers, a 1999 survey suggests that approximately 37% of U.S. water utilities bill customers bi-monthly or even less frequently (Jordan and Albani, 1999). Since results from previous research suggest that information provision can potentially induce conservation, utilities may adopt more frequent billing policies in response to increasingly scarce water resources. But, my empirical results imply that increases in information may not always aid in conservation.

The paper proceeds as follows. In the following 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. Then, 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.

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. I conclude with a reconciliation of the difference between a consumer who may respond to more frequent billing through price perception, quantity perception, or both.[4]

2.1 Background

Consider the choice setting in which a consumer is deciding how much water to use in a given billing period.[5] Borenstein (2009), Gilbert and Graff Zivin (2014), Harding and Hsiaw (2014), and Wichman (2014), for example, motivate decisions based on prices, quantities, and behavior in previous periods as heuristics for making consumptive decisions in electricity and water demand. These decisions are largely governed by preferences, habitual behavior, and consumption rules. For example, at the beginning of a billing period, a household may be particularly concerned with conserving water and thus commit to taking shorter showers, turning off the faucet while brushing teeth, and only running the dishwasher when it is completely full to maximize efficiency. Since utility bills are received periodically, the arrival of price and quantity information offers consumers an opportunity to update their consumption in response to external feedback regarding the effectiveness of their habits. 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 any given appliance uses and, hence, the variable costs of using that appliance. Thus, more frequent billing allows a consumer to better align price and quantity 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 household may not have perfect information about prices and quantities 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 (Shin, 1985; Sallee, 2013; Caplin and Dean, 2014); 2) consumers may be inattentive to or unaware of (changes in) price, tax, or quantity information (Sexton, 2014; Chetty et al., 2009; Finkelstein, 2009; Li et al., 2012; Stango and Zinman, 2014; Grubb, 2014; Hossain and Morgan, 2006; Houde, 2014); 3) inattention to prices or quantity consumed 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; Grubb and Osborne,

[4] 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 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.

[5] 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.

2015; Bollinger et al., 2011).[6] 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?

To examine this question, the ideal experiment would be to choose households at random and provide some with quantity information, others with price information, and yet others with both, at different time intervals. Then, by tracking their consumption in response to the change in information, one could learn how consumer demand changes in response to the information relative to some control group. Indeed this is what many researchers do as purely experimental designs in the context of water and electricity demand to examine quantity reminders, social norms, and so forth (Allcott, 2011b; Allcott and Taubinsky, 2013; Ferraro et al., 2011; Ferraro and Price, 2013; Kahn and Wolak, 2013; Jessoe and Rapson, 2014; Brent et al., 2014). But, no studies have focused on an information treatment as simple as the provision of more frequent billing information. 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.[7] 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 change their consumption habits accordingly.

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 is in contrast to many recent field experiments in which consumers are given a foreign source of information about their consumption (e.g., social comparisons among use and expenditures (Ferraro and Price, 2013; Allcott, 2011b; Costa and Kahn, 2013; Brent et al., 2014), carbon content of electricity consumption (LaRiviere et al., 2014), educational materials about complicated rate structures (Kahn and Wolak, 2013), real-time feedback on consumption (Jessoe and Rapson, 2014; Gans et al., 2013; Strong and Goemans, 2014), or informative signals on variable costs of durable goods (Allcott and Taubinsky, 2013)).[8] The drawback of the natural experiment in this paper is the inability to control the mechanism through which consumers may 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, whatever decision rule they use to make water consumption decisions, they will adjust this rule accordingly rather than constructing a new rule in response to a foreign information intervention). In this scenario, consumers technically are not receiving *more* information about their

[6]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))

[7]While this discussion focuses exclusively on a static (and deterministic) problem, the joint motivation of learning and habit formation implies a dynamic component to consumer demand. This is explored empirically in Sectons 4 and 5, but is left as a theoretical exercise for future research.

[8]An important distinction of this paper is that it does not consider information provision that appeals to consumers' social preferences in the spirit of Levitt and List (2007), rather all informative signals can be used strictly to maximize private benefits despite externalities that may be induced by any one individual's consumption.

usage, they are 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 converge to the neoclassical ideal of perfect price and quantity information, however the cognitive costs of processing that information may be large. With less frequent information provision (e.g., monthly or bi-monthly utility bills), consumers are more likely to base their consumptive decisions on imperfect information since the perception of price and quantity information is crude. So, if a consumer receives a more frequent, informative 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. The intuition driving the welfare framework in this paper hinge on the notion that when consumers are given more frequent information about prices and quantities, the uncertainty in their perception of prices paid and/or quantity consumed is reduced.

2.2 A price mechanism for billing frequency

Since most analyses concerned with the effect of information provision on consumer behavior deal solely with quantity reminders, social comparisons, or prices (Jessoe and Rapson, 2014; Ferraro and Price, 2013; 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 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. This framework is supported empirically by Sexton (2014); Finkelstein (2009), and Li et al. (2012), for example, who find that consumers tend to misperceive prices (or taxes) that are not salient.

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

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

where utility is quasilinear in x and preferences over w exhibit constant elasticity of demand.[9]

[9] Quasi-linearity in water demand is a relatively innocuous assumption since water expenditures comprise a small

The consumer faces a budget constraint, $M = x + \tilde{p}w$, where a consumer's wealth (M) equals her 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 predetermined 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 (e.g., Sexton (2014) and Finkelstein (2009)), I allow for misperceptions to be both above and below the true price.

Intuitively, θ can be interpreted as an exogenously set information parameter capturing the price information used by the consumer, but unobserved by the researcher. In particular, the water utility can manipulate θ by changing information available on an intermittent bill, changing the frequency with 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 other fixed characteristics, Z,

$$\theta = f^P(\mathrm{BF}, Z). \tag{2}$$

Equation 2 captures the essence of the analysis that follows—changes in billing frequency affect price perception through changes in the frequency of receiving a utility bill.

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

Assumption 1 *More frequent information about prices or quantities can never make anyone worse off.*

That information is always 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. While innocuous, this assumption implies structure for the functional relationship in Equation 2. For $\theta > 1$, BF is assumed to have a negative effect such that increasing billing frequency decreases the distortion between the true price and perceived price. However, I do not restrict θ to be greater than one, which would imply a consistently higher perceived price than the true price. For $\theta < 1$, the predicted demand 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 more formally as

$$\frac{\partial f^P(\cdot)}{\partial \mathrm{BF}} \begin{cases} < 0 & \text{if } \theta > 1 \\ = 0 & \text{if } \theta = 1 \\ > 0 & \text{if } \theta < 1 \end{cases} \tag{3}$$

portion of annual income.

which states that increasing the frequency of information through billing frequency allows a consumer to perceive a price closer to the true price.

The consumer's problem defines the optimal level of water consumption,

$$w(\tilde{p}) = \arg\max_{w}\{M + aw^{1/\gamma+1} - \theta pw\}. \tag{4}$$

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.

As an illustrative example of the implications of misperceived prices, consider a consumer facing a change in perceived prices induced by an increase in billing frequency. In Figure 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 perceived the price of water 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. This movement corresponds to an increase in water consumption in response to a reduction in the wedge between her initial misperception and her updated perception. Additionally, the preferred bundle indicates an increase in consumer welfare.

2.2.1 Welfare effects from price misperception

Since the misperception of true prices from infrequent billing may drive a wedge between the actions of a perfectly informed consumer and an inattentive consumer, there are welfare gains from increasing information provision. If $\hat{\theta}$ were observable to the researcher, the change in welfare could be obtained through a simple exercise by integrating the demand function to measure the area under the demand curve between the implied price changes. However, since $\hat{\theta}$ is an unknown parameter, I rely on methods similar to that of Finkelstein (2009), Chetty (2009), Chetty et al. (2009), and Just (2011) to obtain sufficient statistics from observable or estimable parameters that allow for welfare analysis without estimation of a complete demand system. Particularly, I rely on Chetty et al. (2009) for inspiration in deriving customer's misperceptions of price given observable demand responses.

In particular, consider an increase in information, θ_0 to θ_1 via more frequent billing, that corresponds to a change in perceived price from $\tilde{p}_0$ to $\tilde{p}_1$, holding all else equal. 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, via Assumption 1. 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, θ_1, consumers respond to the true price* *such that*

$$w(p, \theta_1) = w(p) = \arg\max_w \{M + aw^{1/\gamma+1} - pw\}.$$

Intuitively, this assumption implies that consumers do not misperceive prices under more frequent billing. While seemingly strong, this assumption is plausible since consumers are provided with more information about their consumption as it occurs. With respect to real-time pricing and 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), for example, show that water customers with in-home real-time meter displays increase consumption within the tiered rate structure if their consumption is below the next marginal price block. In any case, monthly billing still provides an opportunity for price misperception, and I relax this assumption in the sections that follow by assuming that ex post misperceptions are proportional to the true price.

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

$$\eta^P = \frac{\%\Delta w(p, \theta)}{\%\Delta \tilde{p}}, \tag{5}$$

which holds by definition. The percent change in perceived price can be used to obtain an empirical measure of the change in the information parameter from observing a change in perceived price from $\tilde{p_1}$ to $\tilde{p_0}$, which corresponds to the observed change in quantity demanded.[10] In particular, we can write,

$$\%\Delta \tilde{p} = \frac{\tilde{p_0}}{\tilde{p_1}} - 1 = \frac{\theta_0 p}{\theta_1 p} - 1 = \%\Delta\theta, \tag{6}$$

since the market price does not change.

Combining this expression with Equation 5, rearranging, and multiplying through by p (corresponding to the price *after* the change in information) provides,

$$\Delta\tilde{p} = \Delta\theta p = \left(\frac{\%\Delta w(p, \theta)}{\eta^P}\right) p, \tag{7}$$

which states that for a change in θp, the equivalent 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 7 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 initial perceived price, consumer surplus can be calculated by integrating the demand function between the initial price perceived, $\tilde{p_0}$, and the price perceived

[10]Note that this price change is moving in the opposite direction from the experiment (i.e., from monthly billing to bi-monthly billing). The direction of this change is immaterial for the welfare analysis that follows, though this is a more convenient way to represent the movement in prices.

after the change in billing frequency, $\tilde{p_1}$,

$$\Delta CS^P = \int_{\tilde{p_1}}^{\tilde{p_0}} w(\tilde{p})\mathrm{d}\tilde{p}, \tag{8}$$

which, since the data in the experiment are not sufficient to estimate a true demand function, can be approximated by

$$\Delta CS^P \cong -\frac{1}{2}\Delta\theta p \frac{\partial w}{\partial \theta} \tag{9}$$

which is analogous to the Harberger (1964) triangle approximation for deadweight loss[11]. Under the assumptions made so far, the treatment effect ($\partial w/\partial\theta$) and perceived price elasticity of demand (η^P) serve as sufficient statistics for calculating changes in welfare due to a change in billing frequency.

The welfare analysis is illustrated for a stylized example in Figure 2. As shown, $\Delta\theta p$ is the change in perceived price that reflects the movement in quantity demanded along a fixed demand curve. The initial perceived price level is higher than the price level after the change in billing frequency (reflecting Assumption 2). Further, the change in demand, $\partial w/\partial\theta$, is assumed to increase with a decrease in perceived price. Since there is an economic cost borne by 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.

2.3 A quantity mechanism for billing frequency

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 1, assuming perfect information about prices, augmented by a quantity information parameter,

$$u = x + a(\lambda w)^{1/\gamma+1} \tag{10}$$

[11]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)

where λ is a parameter that scales *perceived quantity* 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,

$$\lambda = f^Q(\text{BF}, Z). \tag{11}$$

By Assumption 1, we can write the marginal effects of billing frequency on λ similarly to that of the price parameter,

$$\frac{\partial f^Q(\cdot)}{\partial \text{BF}} \begin{cases} < 0 & \text{if } \lambda > 1 \\ = 0 & \text{if } \lambda = 1 \\ > 0 & \text{if } \lambda < 1 \end{cases} \tag{12}$$

which states that increasing quantity information allows a consumer to predict consumption closer to her true consumption.

In this framework, a consumer maximizes Equation 10 subject to a budget constraint, $M = x + \lambda pw$. In contrast to imperfect price perception, imperfect quantity perception affects consumer utility both 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

$$\max_{w} \{M + a(\lambda w)^{1/\gamma+1} - \lambda pw\} \tag{13}$$

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

$$\tilde{w}(p, \lambda) = Ap^{\frac{1-\gamma}{\gamma}} \lambda^{\frac{1-2\gamma}{\gamma}} = \lambda^{\eta^Q - 1} w(p) \tag{14}$$

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 (see Appendix B for a derivation). Equation 14 illustrates that demand is scaled conveniently by the information parameter. Envelope conditions provide the following comparative statics result for demand responses with respect to changes in information,

$$\frac{\partial^2 u}{\partial \tilde{w} \partial \lambda} = \frac{\partial \tilde{w}}{\partial \lambda} = (\eta^Q - 1)\lambda^{\eta^Q - 2} w(p) < 0, \tag{15}$$

which states that increases in λ decrease consumer demand. But, recall that λ itself is a function of parameters (as in Equation 12) such that for initially over-perceived quantities ($\hat{\lambda} > 0$), λ is decreasing in information provision, and for initially under-perceived quantities ($\hat{\lambda} < 0$), λ is increasing in information provision. 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.

2.3.1 Welfare effects from quantity misperception

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,

$$\eta^Q = \frac{\%\Delta\tilde{w}(p,\lambda)}{\%\Delta p} \tag{16}$$

and, since we do not observe price changes directly, we can infer them by rearranging Equation 16 and multiplying through by p,

$$\Delta p = \left(\frac{\%\Delta\tilde{w}(p,\lambda)}{\eta^Q}\right) p. \tag{17}$$

Since the change in quantity demanded is observed in the experiment and η can be estimated or inferred from other studies, Equation 17 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 information ($\partial\tilde{w}/\partial\lambda$) and the price elasticity (η) to calculate consumer surplus in response to the change in billing frequency,

$$\Delta CS^Q = \int_{p_1}^{p_0} \tilde{w}(p)\mathrm{d}p = \lambda^{\eta^Q-1}\int_{p_1}^{p_0} w(p)\mathrm{d}p, \tag{18}$$

which can be approximated by

$$\Delta CS^Q \cong -\frac{1}{2}\Delta p\frac{\partial\tilde{w}}{\partial\lambda} \tag{19}$$

Thus, Equation 19 provides an analytical formula from which we can use quasi-experimental estimates to calculate changes in economic welfare.

2.4 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.[12] 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.

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,

$$\Delta CS^P \cong -\frac{1}{2}\Delta\theta p\frac{\partial w}{\partial\theta}, \tag{20}$$

and ΔCS^Q is the analog for quantity misperception,

$$\Delta CS^Q \cong -\frac{1}{2}\Delta p\frac{\partial\tilde{w}}{\partial\lambda}. \tag{21}$$

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 7. Whereas, Δp under quantity misperception is equal to $(\%\Delta\tilde{w}(p,\lambda)/\eta^Q)p$ by Equation 17. 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

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

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. Effectively, 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 Data and empirical setting

The primary data used in this analysis are from the City of Durham Water Utility in North Carolina (henceforth, "Durham"). Included in these data are monthly water and sewer use, fixed service fees and volumetric consumption rates, the date each bill was sent, the address of the customer, billing cycle, 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 obtained from Durham County. Each matched residential address was spatially linked to its 2010 Census block as well as billing district polygons provided by the Durham. For each household, I determine the nearest billing district, as well as the linear distance from the centroid of the tax parcel to the nearest billing 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 allows to reduce the impact of renters, who may or may not have water bills included in their rent. Further, this avoids econometric identification problems when relying on variation within a household over time.[13]

The final sample consists of roughly 59,000 individual household accounts with water use 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 1. The first four columns break up household characteristics and details on water use by the year in which households transitioned from bi-monthly to monthly billing, which is discussed in detail in the following subsection. 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 sample mean) with the exception of households 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 more white residents. Water consumption, however, is qualitatively similar across all groups. For the typical household in the sample, the mean assessed value of the home is around \$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 are likely to reside in Census blocks that have a proportion of renter occupied homes that comprises one-quarter of all residents. Around 53% of the sample is white, and the average household size is between two and three people. Average bi-monthly bills for all time periods in the sample are \$85 for consumption of 985 cubic feet of water use.

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.

3.1 Transition to monthly billing

Beginning in December 2011, Durham began to sequentially transition billing districts from bi-monthly to monthly billing. Primary reasons cited for the transition included cost-saving from fewer delinquent payments, early leak detection, improving customer service, and reducing administrative costs. In addition, the switch in billing frequency coincided with 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 are reduced. Meters were installed for each billing district and, once the installations were completed, the entire district was transitioned to monthly billing. To make a cost-saving argument to the City Council, the water utility

[13]Renters may cause problems for identification if water consumption is included in their rent and they do not receive a utility bill. 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.

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 routes for new meter installation (and, subsequently, monthly billing) was chosen to work around billing cycles and other feasibility constraints. Initially, districts were chosen in a geographically convenient order but no consideration of billing history, income base of neighborhood, or any other financial indicator was taken into account.[14]

In Table C.1 in Appendix C, 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 transitioned first, 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. But, these characteristics may simply be an artifact of the geographic location of the district, rather than explicit selection. 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.

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 (which would remove them from the data set explored in this paper). Within the study period, 12 of 17 billing districts were transitioned according to the timing in Table 2. The first district transitioned received their first monthly bill on December 1st, 2011. Figure 3 presents a map of the billing districts transitioned to monthly billing prior to July 2013. For Durham, the water utility service area is represented by the union of all billing districts outlined in bold. The shaded billing districts are the eight that transitioned to monthly billing prior to June 30th, 2013, though four additional districts transitioned after June 2013. Of the 17 districts, five never transitioned to monthly billing within the sample.[15] All residential districts were scheduled to be billed on monthly basis by September 2014. Figure 4 presents a "close-up" view of billing districts within neighborhoods. This figure illustrates that the billing districts 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 geography to minimize the concern that fundamental differences in neighborhood characteristics might bias results.

[14] These institutional details were obtained through discussions with utility officials in charge of the meter installation and billing transition.

[15] Households in billing district 2 received their first monthly bill at the end of June 2014, which is the last month of the billing data—as such, this district is considered as a non-transitioned district since there exist no post-transition observations.

An important caveat is that under bi-monthly billing, water bills are mailed on a staggered schedule. This staggered nature 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 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 bills are aggregated to a bi-monthly level for both before and after comparisons within a household as well as cross-sectional comparisons of households billed at different frequencies. So, the unit of observation for consumption is a two-month period regardless of whether households are being billed monthly or bi-monthly.

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 identical across billing frequencies. Figure 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. This structure was imposed 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 value. 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, it is unlikely that consumers are responding to changes in the price structure itself. 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.

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. This effect tends to be more pronounced for larger meters used in commercial and industrial applications, but is also true for residential meters. Indeed, metering companies market the efficiency of new meters as a means to capture "non-revenue" water that utilities are treating and distributing, but 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 new metering technology increases billed consumption mechanically.

First, I present the percent of water accounted for, by utility accounting sheets, as a percent of total pumped water in Figure 6 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 being billed through new metering technology. So, the flat trend observed after the introduction of the new meters provides some indication that the purported technical efficiency of meters does not confound the treatment. There is, however, a noticeable decline in the variance of system efficiency after December 2011, which could indicate that monthly billing paired with metering technology improved Durham's capacity to manage demand and detect leaks more quickly, but this should not affect a consumer's response to a change in billing frequency.

As further empirical tests that are outlined in detail in the sections that follow, I examine a subset of households that have their water bills drafted automatically from their checking accounts and are assumed to be less responsive to changes in price (Sexton, 2014) and find no significant change in consumption in response to monthly billing. Further, since meters were installed over the course of several months leading up to the change in monthly billing, I analyze the "treatment" of meter installation in the two periods directly before the switch to monthly billing. And, I find nothing to support the notion that technical efficiency of metering technology is confounding my empirical results.

4 Empirical strategy

The empirical approach I take in this paper is to identify demand responses to increased information provision using quasi-experimental techniques. In particular, 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. In particular, I use a generalized difference-in-difference method to allow for the transition of billing districts to monthly billing to vary over time. By exploiting the rich variation in consumption both cross-sectionally as well as temporally, I estimate a series of models controlling for various demographics, household characteristics, and spatial and temporal fixed effects to estimate an average treatment effect. Overall, I find that customers increased consumption in response to more frequent information between two and seven percent, with the preferred estimate being a 4.6% increase in quantity demanded. Further, to assess the validity of the difference-in-difference estimates to potential differences among treatment and control neighborhoods over time, I restrict the difference-in-difference estimator to a local linear regression within a narrow window around billing district boundaries. This difference-in-discontinuity estimator provides a local treatment effect in light of potential bias that might arise from changes in time-varying unobservable neighborhood characteristics that might be correlated with treatment status. Results are robust to this stratification within 500 feet of district boundaries.

Further, I consider that the change in billing frequency may induce an initial, but impermanent, response to more frequent information until households revert to their initial habitual behavior. To address this, I estimate partial adjustment models and find that the long-run treatment effect is larger in magnitude than the short-run effect and that the main effect of information on consumption remains positive and significant.

Additionally, I explore heterogeneous treatment effects among historical consumption levels, summer use, lot size, and wealth to provide suggestive evidence from where the increase in consumption is arising. The treatment effect is substantially larger in the summer months, which suggest an increase in outdoor water use. Results also imply a decreasing trend in the treatment effect as pre-treatment water use increases. Additionally, households with less expensive homes and smaller lots tend to increase consumption by a larger percentage.

As an ad-hoc test that consumers are indeed responding to the change in billing frequency, I explore the effect of automatic bill payment (via electronic draft payments) on the treatment effect. Households who are enrolled in automatic bill payment do not respond to the increase in billing frequency, as inattention suggests.

4.1 Difference-in-differences

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

$$\ln(w_{ijt}) = \alpha + \beta \mathrm{BF}_{jt} + C_t\omega + Z_{ijt}\gamma + \tau_t + \epsilon_{ijt} \tag{22}$$

where w_{ijt} is household i's water consumption in time t in billing district j.[16] 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 temperature (degrees Fahrenheit) and the total rainfall (inches) for each billing period. Z_{ijt} is a vector of household and demographic characteristics described in Table 1. τ_t is a vector of various time controls including time fixed effects, a linear time trend, and seasonal indicators. Finally, ϵ_{ijt} is the residual error term that likely contains within-household, and within-route, serial correlation as well as an unknown form of heteroskedasticity (Bertrand et al., 2004).[17]

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 will provide an unbiased estimator of the average treatment effect.

To analyze common trends, I first note that treatment and control households reside in a small geographic area, and even sometimes within the same neighborhood (see Figure 4 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. This graphical analysis presented in Figure 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) almost perfectly 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. Further, while only suggestive, there is a slight increase in consumption for treatment households after the change in billing frequency. But, only one of these routes (containing roughly 3000 households) transitioned in December 2011, so it is not surprising that there is no noticeable evidence of a treatment effect until after July 2012, when other routes started

[16]In all regression specifications with logged dependent variables of water consumption, I add 1 to bi-monthly water use since there is a small number of "true" zero use observations in the data set.

[17]Estimated standard errors account for serial correlation within households, as well as within routes, which is the source of variation in treatment source.

transiting to monthly billing.

While a simple linear regression approach presented in Equation 22 provides an consistent estimate of an average treatment effect, it is possible that there are unobservable differences 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 increases in information provision, but the available data do not permit for direct controls. As such, I estimate Equation 22 with household-specific intercepts that absorb time-invariant unobservable characteristics that may affect water use,

$$\ln(w_{ijt}) = \alpha_i + \beta \mathrm{BF}_{jt} + C_t \omega + \tau_t + \epsilon_{ijt} \tag{23}$$

where the α_i is now a household-specific intercept and all else is as in Equation 22. The identifying assumptions for β to remain a causal effect of increases in billing frequency remains the same as before, though we now require strict exogeneity in our regressors and the following structure on the orthogonality of treatment status and unobservable household characteristics, conditional on a vector of all exogenous covariates $X_{ijt} = [C_t, \tau_t, Z_{ijt}]$,

$$E[w_{0ijt}|\alpha_i, X_{ijt}, \mathrm{BF}_{jt}] = E[w_{0ijt}|\alpha_i, X_{ijt}] \tag{24}$$

where w_{0ijt} is counterfactual consumption in potential outcomes notation. Intuitively, this assumption implies that treatment status cannot be correlated with an unobservable, time-varying factor at the household-level. Conditional randomization of monthly billing across billing districts implies that treatment assignment did not target explicit household characteristics.

4.2 Exploiting billing district boundaries

To assess the robustness of the difference-in-difference methodology, I exploit the geographic billing district boundaries by estimating treatment effects within a narrow window on either side of the boundaries. I consider households within 2000, 1000, and 500 feet from the district boundaries. I first estimate Equation 22 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 23 under the notion that there may be unobservables biasing results in the pooled cross-section. The latter specification results in identifying a treatment effect using differences between within-household variation for treated households and control households across the treatment threshold (i.e., the district boundary). Thus, this estimation strategy is dubbed "difference-in-discontinuity" because it 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 local average treatment effect in this scenario is that as one approaches the billing district boundaries 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.[18]

To explore these effects graphically, I plot mean consumption for 2012 in 20 and 40 foot bins within 1000 feet of billing district boundaries. In Figure 8, treatment households, represented by solid diamonds on the right-hand side of each figure are the households who transitioned to monthly billing in or before 2012. Control households, represented by hollow circles on the left-hand side of each figure are all other households. Panels A and B group households into 40-foot bins and fit a linear and quadratic trend, respectively. Panels A and B are analogous for 20-foot bins of consumption. In all panels, there is a clear deviation at the billing district boundary. While the control households display a more uniform, linear trend, treatment households appear to be more noisy and 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 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 should not be problematic as the difference-in-discontinuity estimator is a consistent estimator of a local average treatment effect as the distance from the district threshold decreases in the limit, and the trends nearby the boundary threshold are relatively similar (Grembi et al., 2012). Each panel displays a treatment effect at the boundary ranging, roughly, from 30-80 cubic feet of water consumption, which is consistent with the estimated treatment effects.

4.3 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 in the consumption process. The general estimating equation takes the form,

$$\ln(w_{ijt}) = \alpha + \beta \mathrm{BF}_{jt} + \zeta \ln(w_{ijt-1}) + C_t\omega + Z_{ijt}\gamma + \tau_t + \epsilon_{ijt} \tag{25}$$

where, in an OLS framework, β remains the causal effect of interest in the short-run which is augmented by an adjustment parameter, ζ. We can then determine a long-run treatment effect from the estimates of Equation 25, $\beta^\star = \hat{\beta}/(1-\hat{\zeta})$. Additionally, the short-run average treatment effect obtained from estimation of Equation 25 provides a lower bound for the treatment effect

[18]Note, the term "neighborhood" is used loosely to represent the window around billing district boundaries.

based on the following conditional independence assumption,

$$E[w_{0ijt}|w_{ijt-1}, X_{ijt}, \mathrm{BF}_{jt}] = E[w_{0ijt}|w_{ijt-1}, X_{ijt}], \tag{26}$$

in the absence of omitted variables. Thus, the ATE estimated in Equation 23 and the ATE estimated in Equation 25 bound the treatment effect from above and below, respectively, which provides a useful robustness check of the ATE 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 non-independent observations within a household, I estimate

$$\ln(w_{ijt}) = \alpha_i + \beta \mathrm{BF}_{jt} + \zeta \ln(w_{ijt-1}) + C_t\omega + \tau_t + \epsilon_{ijt}, \tag{27}$$

which is the panel analog of Equation 25. 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.[19]

4.4 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 information. Formally, I estimate conditional average treatment effects (CATE) similar to Allcott (2011b), Ferraro and Miranda (2013), and Abrevaya et al. (2014) by interacting covariates of interest with the treatment indicator. The CATE specification takes the form,

$$\ln(w_{ijt}) = \alpha_i + \beta \mathrm{BF}_{jt} + \sum_{s \in S} \beta_s (\mathrm{BF}_{jt} \times 1[Z_{ijt}]_s) + C_t\omega + \tau_t + \epsilon_{ijt} \tag{28}$$

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 23. 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; Renwick and Archibald, 1998; Baerenklau et al., 2014), square footage of the home, and the age of home; as well

[19] In addition, I consider partial adjustment models in first-differences with an Anderson-Hsiao estimator (Anderson and Hsiao, 1981), as well as in a GMM context with an Arellano-Bond estimator (Arellano and Bond, 1991). However, both of these models rely on untenable assumptions for the data generating process. In particular, water consumption today (t) is serially correlated with consumption last period ($t-1$), two periods previous ($t-2$), as well as all periods before that for a given household. Thus, the identifying assumptions for these models are impractical. Results are available from the author upon request.

as the number of bathrooms within a household.[20]

The conditioning variables chosen in Equation 28 either represent key drivers of residential water demand or serve as a proxy for preferences that may influence demand responses. In light of the former, evidence of heterogeneous treatment effects along established margins allow for probing the mechanism through which consumers may be responding. For example, Wichman et al. (2014) find that low-income households are more responsive to price changes in water demand, while non-price quantity mechanisms are more effective at targeting changes in water use for this subpopulation. As such, a similar response would provide suggestive evidence that households are responding along a price or quantity margin. For the latter, in addition to exploring the response mechanism (i.e., price or quantity perception), I consider the CATE for the number of bathrooms as a proxy for preferences towards water-intensive appliances, for example, to examine the physical source of the treatment effect.

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, nonessential water use in the summer months is the 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. I limit this indicator to those households who enrolled in ABP prior to the transition to monthly billing and remained on ABP until the end of the study. These households provide two useful functions. First, under the notion that ABP customers are less attentive to changes in prices and, thus, billing frequency (Sexton, 2014), a null treatment effect would instill confidence that the empirical results are identifying a behavioral response to the frequency of information received. 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. Since, in order to counteract that positive effect, ABP customers would need to choose to decrease consumption in response to more frequent information.

5 Empirical results and discussion

5.1 Main results

In this section, I present results motivated by the empirical models in previous section. First, Table 3 presents pooled OLS estimates of Equation 22. From Column (1) through (5), I add additional

[20]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 requires strong assumptions on rank preservation, particularly that the rank of the potential outcome for a given household is the same under treatment and non-treatment. In the context of water demand, I show empirically that households respond to information provision differentially and, thus, may not reside in the same quantile of consumption before and after treatment. Additionally, to expunge any source of endogeneity in the CATE estimation using (2009 and 2010) consumption patterns as conditioning covariates, I restrict the sample to observations beginning in January 2011.

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 route level, to account for any serial correlation in error terms (Bertrand et al., 2004). The former accounts for the fact that water consumption is serially correlated within a household, whereas the latter acknowledges that the variation in treatment status arises at the billing district level. For all estimated average treatment effects, 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 largest adjusted R-squared, and also exhibit the most conservative treatment effects.

In Table 4, 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 average treatment effect of billing frequency is estimated precisely at 8.4% and 4.6% increases in quantity of water demanded without, and with, the inclusion of period-of-sample fixed effects, respectively.[21] For comparison, the estimate of 4.6% in Column 2 of Table 4 can be compared directly with the estimate of 5.1% in Column 5 of Table 3, 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 4 will be the preferred causal effect of the average treatment effect of changes in billing frequency on water consumption. This specification accounts for any time-invariant unobservables while simultaneously controlling for time effects.

To examine whether there are unobservable changes in neighborhood characteristics in the pooled difference-in-difference models, I present regression discontinuity (RD) estimates (for the pooled cross-section) and difference-in-discontinuity (for fixed effects) estimates in Tables 5 and 6. These models control for time-varying unobservables at the neighborhood level. In Table 5, pooled cross-section models are estimated for the set of households within 2000, 1000, 500 feet of billing district boundaries. The treatment effect of billing frequency remains stable as the window shrinks, and remains significant at conventional levels.[22] 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 Table 6. The estimates decrease monotonically as we move from Column (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

[21] 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.

[22] 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.

level indicate that this effect is not statistically different from zero, but note that this effect is statistically similar to that of Columns (1) and (2).

Dynamic models presented in Table 7 imply that past water consumption is an important predictor of current water consumption and significantly improves the fit of each model. Interestingly, the treatment effect in the first two columns is smaller than that of the preferred specification (2.1% and 2.6%) but statistically different from zero. 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 would be from the OLS estimator, which suggests a long-run treatment effect that is double the magnitude of the preferred ATE in Table 4. 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 but, more importantly, that this effect appears to strengthen over time. Additionally, the final two columns in Table 7 present Anderson-Hsiao and Arellano-Bond estimates, which overstate a plausible effect of the treatment on consumption (27.6% and 24.7%, respectively). Since these estimators rely on strong assumptions about the serial correlation of the data generating process, I refrain from using these estimates for any policy conclusions.

Results from the dynamic OLS model, however, 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 (Ferraro et al., 2011; Allcott and Rogers, 2014; 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 an increase in the frequency of billing information. Of particular importance is that the long-term effect is roughly double the magnitude of the short-run effect. This result suggests that consumers observe the change in the frequency of information 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. In particular, this result suggests that consumers' misperceptions of prices or quantity may be mitigated permanently through the shift to more frequent billing. However, an explicit analysis of the durability of these effects over time is left for future work.

Additionally, this result is consistent with Gilbert and Graff Zivin (2014) who find that electricity consumption decreases in the week after the receipt of a 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 context, this result implies that consumers might make similar changes to water use upon the receipt of a bill each month, but perhaps revert back to their misperceptions at the end of the billing cycle. While this result is not encouraging from 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 and quantity signals.

5.2 Heterogeneous treatment effects

I interpret results from the heterogeneous treatment effect models 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. To probe the mechanism through which consumers respond, I make use of several stylized facts from the empirical literature on the demand-side management of residential water—**(SF1)** Price sensitivity decreases with wealth/income (Mansur and Olmstead, 2012; Renwick and Archibald, 1998; Olmstead and Stavins, 2009); **(SF2)** High volume consumers are least sensitive (if at all) to changes in price (Mansur and Olmstead, 2012; Wichman, 2014; Wichman et al., 2014; Klaiber et al., 2014); **(SF3)** Responsiveness to non-price interventions (e.g., social comparisons or conservation initiatives) is increasing in baseline water use (Ferraro and Price, 2013; Ferraro and Miranda, 2013; Wichman et al., 2014). These stylized facts allow for suggestive attribution of the mechanism (i.e., price, quantity, or perhaps something else) through which consumers respond to changes in billing frequency.

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 8, the ATE in summer (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, etc.), this effect is important. A consumer who over-perceives prices or quantities is likely to consume more extraneous water upon mitigating uncertainty in her initial perception. That is, with more frequent information, more nonessential water is used. Regardless of the mechanism through which consumers are responding, this result indicates that information provision does not always induce conservation by reminding consumers of their consumption patterns. This result necessitates a closer examination of the ways in which consumers' perceptions drive responses to new information before information provision 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 9. As shown in Panel A, there is significant heterogeneity in responsiveness to billing frequency across the consumption distribution; there is a strong decreasing trend 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.

These relationships are consistent with **SF2**, suggesting that perhaps a price mechanism is driving consumer behavior. However, the negative effect observed in Panel A suggests that information may work as a quantity mechanism (**SF3**) for larger users of water, but not for high *summer* consumption. This result suggests that perhaps large users are not sensitive to the change in perceived price, but they are more salient to quantity consumed. 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 (or, that they were using less water than they thought), and increase consumption accordingly. Additionally, the within-sample heterogeneity of a positive and a negative response to the same information treatment emphasizes the importance of reconsidering the role of information as a conservation mechanism.

Further, 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. CATEs for quintiles of each of these variables are presented graphically in Figure 10. As shown in Panel A, there is considerable homogeneity in the CATEs, estimated at roughly 5 percent for the lowest 80 percent of households, while the highest 20 percent of households exhibit no significant response. The latter effect is consistent with **SF1**, which suggests wealthy households may respond through a price mechanism, but there is no reason to believe that an analogous fact is not true for quantity as well. High income households 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), heterogeneity in the bottom quintile is expected. The lack of heterogeneity for this subgroup casts some doubt on the hypothesis that price responsiveness is driving this relationship. Further, Wichman et al. (2014) show that non-price conservation policies are relatively homogenous across income classes, which aligns well with a quantity mechanism for the bottom 80 percent of the wealth distribution.

Additionally in Panel B of Figure 10, I present the estimated CATEs for quintiles of a household's lot size. Since lot size is often used as a proxy for outdoor water use preferences (Renwick and Archibald, 1998; Mansur and Olmstead, 2012), **SF2** would suggest a small or null response for larger lots if a price mechanism is driving the results. As shown, there is a decreasing trend in the estimated CATE as lot size increases, though all estimates remain positive and significant. However, the average household in each subgroup is still increasing consumption in response to more frequent billing. This relationship can be interpreted similar to the effect of summer water use—more frequent reminders about quantity consumed might attenuate the effect for large users of water, while the low users of water may increase their usage more in response to the same information.

Moreover, I explore the quintile of house size and the number of bathrooms as indicators of preferences for water intensive indoor use and appliances in Figure 11. Both of these structural characteristics of a home display similar trends in Panels A and B—the higher ends of the distribution exhibit lower responses to the treatment, but similar trends in the lower portion of the distribution. While these trends do not provide evidence of a response mechanism (though both measures are correlated with income, which corroborates a price mechanism through **SF1**), they provide a useful indication that there is no evidence of the treatment effect being driven by preferences for indoor water use. This result suggests that nonessential water use is likely driving the full sample ATE.

Lastly, I show CATEs for quintiles of the age of the home in Figure 12. In this figure, the lowest quintile corresponds to the newest homes and the highest quintile corresponds to the oldest homes. Interestingly, 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, 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 to be in need of renovations, 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 result, however, is merely speculative. Further research on the mechanism driving these results is needed. Overall, the heterogeneous treatment effect estimates provide evidence of increases in nonessential water use in response to increases in billing frequency.

5.3 Robustness tests

Additionally, I examine the robustness of the empirical results in several ways. In the last column of Table 8, I show that the effect of an interaction of the treatment effect and automatic bill payment 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. 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. Thus, the treatment effect can be attributed to changes in billing frequency with confidence. In effect, this result shows that if consumers who are enrolled in automatic bill payment do not observe bills in as much detail, they do not respond to changes in the frequency of the bill, which corresponds with recent research in electricity demand (Sexton, 2014), except that it moves consumption in the opposite direction.

As an additional test of robustness, I examine the ATE explicitly for the subset of customers enrolled in ABP under different modeling assumptions to re-examine the effect of more technically efficient metering. In Table C.2, I estimate four models on the subsample of automatic billing customers. The first two columns show pooled OLS estimates of the ATE without and with billing district fixed effects, respectively. In both cases the effects are not significant with standard errors clustered at the billing district level, though the second column is significant at the 10 percent level when standard errors are clustered at the household level. Further, Column (3) presents estimates of the preferred model with household fixed effects accounting for serial correlation at the household and billing district level. The estimated coefficient is 2.8%, but the p-values are 0.193 and 0.513, respectively, indicating a large variance around the ATE for this group. The lack of a strong, statistically significant effect lends credence to the notion that the preferred ATE in previous models is not an artifact of mechanical efficiency of the new meters. Further, these results contribute empirical evidence that those households enrolled in ABP are not salient to changes in the frequency of billing, which adds to the literature of salience and inattention.

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 simply a dummy for the billing district that had new meters installed interacted with a time fixed effect for the latter period. Thus, a false difference-in-difference treatment provides an empirical test of 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 hypothesis is that consumption will increase in response to the "treatment." Table C.3 depicts the estimated treatment effect. In every scenario, we fail to reject the null hypothesis. Interestingly, and somewhat perversely, I find that for three of the nine models, the "treatment" effect is negative and significant. All other estimates are not statistically different than zero. This final test provides more confidence in the notion that the new metering technology does not contaminate the treatment effect.

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 so long as that information reduces the uncertainty in consumer perception of prices or levels of quantity demand. This section quantifies the economic welfare that arises from reducing the distortion in perceived prices and quantities.

For the welfare analysis, I use results from the preferred difference-in-difference models. I calculate welfare changes for both price and quantity mechanisms. 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.[23] 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.

I present a graphical representation of welfare changes due to changes in billing frequency under different modeling assumptions in Figure 13. In this figure, consumer surplus is plotted as a function of the price elasticity of demand as in Equations 9 and 19. For each trend line, different baseline assumptions are made. Perfect price information is assumed for the lowest line (for both the price and quantity mechanism), reflecting Assumption 2 (that is, the market price is perceived perfectly after the change in billing frequency). For the other trend lines, I relax this assumption by allowing the price perceived after the change in billing frequency to remain 10%, 25%, and 50% percent higher than the actual price. In all models, consumer surplus declines as price sensitivity increases. Intuitively, if a consumer exhibits a small price elasticity (e.g., -0.1), their demand curve is much steeper than that of a consumer with a relative larger price elasticity (e.g., -0.8). Thus, to maintain a constant consumptive response, a larger change in price is necessary for the less elastic consumer. Hence, the corresponding welfare gain will be larger for small price elasticities of demand.

Summary statistics for these calculations for select price elasticities are also presented in Table 9. To focus the analysis, I interpret the range [-0.5, -0.2] as the most relevant short-run elasticities for policy conclusions. This range is consistent with estimates in Wichman (2014) and Wichman et al. (2014) who estimate short-run elasticities for geographically and demographically similar households using quasi-experimental methods and instrumental variables, respectively, though these estimates are consistent with other studies under different modeling assumptions (Espey et al., 1997; Arbués et al., 2003).

In Table 9, I present welfare statistics for a range of elasticities under different modeling assumptions. While the price elasticity may indeed be different for consumers responding to perceived price (within a price mechanism) and the true price (under a quantity mechanism), the numerical results are identical for the first row of Panel A and the first row of Panel B. In Section 2.4, I argued that the quantity mechanism provides a lower bound for welfare calculations conditional on initial over-perception of prices. Under the assumption that the average consumer responds within a quantity mechanism with a price elasticity of -0.3, the welfare gain from increased information

[23] For meta-analyses of price elasticities in water demand, see Espey et al. (1997) and Arbués et al. (2003).

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).[24] From a conservation perspective, the change in billing frequency resulted in an increase in water consumption of roughly 46-thousand 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.[25]

7 Conclusions

Information provision is a growing topic both in markets where expenditures are made intermittently and where price regulation is politically challenging. In this paper, I make several contributions to this literature. I first posit models of consumer misperception that provide a micro-foundation that are 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 information for a water provider in the southeastern US through a change in billing frequency. I find strong empirical 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 one of the first documented causal increases in consumption in response to an increase in information provision within environmental and resource policy, which is a particularly pertinent result for drought-prone regions. Similar to repeated information treatments, I find that the effect is not impermanent within the 4.5-year study period—dynamic models suggest a quick transition to a new baseline equilibrium that maintains a consistently larger long-run treatment effect. Further, heterogeneous treatment effects suggest that the increase in consumption is driven by extraneous, outdoor water use.

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 billing frequency through an exogenous transition from bi-monthly to monthly billing for a North Carolina utility. This change in behavior is attributed to the fact the consumers billed less frequently observe more opaque price and quantity signals, thus driving a wedge between perceived and actual billing information. Increasing the transparency of prices and quantities through increased billing frequency results in welfare gains of approximately

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

[25] I approximate the volume of an Olympic-size swimming pool at 88,000 cubic feet of water.

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 both prices and quantities.

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 billing settings.

References

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

Hunt Allcott. Consumers' perceptions and misperceptions of energy costs. *American Economic Review, Papers and Proceedings*, 101(3):98–104, 2011a.

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

Hunt Allcott. The welfare effects of misperceived product costs: Data and calibrations from the automobile market. *American Economic Journal: Economic Policy*, 5(3):30–66, August 2013.

Hunt Allcott and Todd Rogers. The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation. *American Economic Review*, Forthcoming, 2014.

Hunt Allcott and Dmitry Taubinsky. The lightbulb paradox: Evidence from two randomized experiments. Working Paper 19713, National Bureau of Economic Research, December 2013.

T. W. Anderson and Cheng Hsiao. Estimation of dynamic models with error components. *Journal of the American Statistical Association*, 76(375):598–606, 1981.

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

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

Manuel Arellano and Stephen Bond. Some tests of specification for panel data: Monte carlo evidence and an application to employment equations. *The Review of Economic Studies*, 58(2):277–297, 1991.

Kenneth A. Baerenklau, Kurt A. Schwabe, and Ariel Dinar. The residential water demand effect of increasing block rate water budgets. *Land Economics*, 90(4):683–699, 2014.

Marianne Bertrand, Esther Duflo, and Sendhil Mullainathan. How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics*, 119(1):249–275, February 2004.

Bryan Bollinger, Phillip Leslie, and Alan Sorensen. Calorie posting in chain restaurants. *American Economic Journal: Economic Policy*, 3(1):91–128, February 2011.

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

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

Andrew Caplin and Mark Dean. Revealed preference, rational inattention, and costly information acquisition. Working Paper 19876, National Bureau of Economic Research, January 2014.

Raj Chetty. Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annual Review of Economics*, 1:451–487, 2009.

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

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

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

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

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

Paul J. Ferraro 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*, 35(3):356–379, 2013.

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

Paul J. Ferraro, 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*, 101(3):318–322, 2011.

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

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

Xavier Gabaix and David Laibson. Shrouded attributes, consumer myopia, and information suppression in competitive markets. *The Quarterly Journal of Economics*, 121(2):505–540, May 2006.

Will Gans, 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*, 36:729–743, 2013.

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

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

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

Michael D. Grubb. Consumer inattention and bill-shock regulation. *The Review of Economic Studies*, 2014.

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

Arnold C. Harberger. The measurement of waste. *American Economic Review, Papers and Proceedings*, 54(3):58–76, 1964.

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

Jerry A. Hausman. Exact consumer's surplus and deadweight loss. *American Economic Review*, 71(4):662–676, 1981.

Tanjim Hossain and John Morgan. ...plus shipping and handling: Revenue (non) equivalence in field experiments on eBay. *The B.E. Journal of Economic Analysis & Policy*, 5(2), 2006.

Sébastien Houde. How consumers respond to environmental certification and the value of energy information. Working Paper 20019, National Bureau of Economic Research, March 2014.

Koichiro Ito. Do consumers respond to marginal or average price? evidence from nonlinear electricity pricing. *American Economic Review*, 104(2):537–563, 2014.

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

Ginger Zhe Jin and Phillip Leslie. The effect of information on product quality: Evidence from restaurant hygiene grade cards. *The Quarterly Journal of Economics*, 118(2):409–451, 2003.

Jefferey L. Jordan and Rick Albani. Using conservation rate structures. *Journal – American Water Works Association*, 91(8):66–73, 1999.

Richard E. Just. Behavior, robustness, and sufficient statistics in welfare measurement. *Annual Review of Resource Economics*, 3(1):37–70, 2011.

Matthew E. Kahn and Frank A. Wolak. Using information to improve the effectiveness of nonlinear pricing: Evidence from a field experiment. *Working Paper*, 2013.

H. Allen Klaiber, V. Kerry Smith, Michael Kaminsky, and Aaron Strong. Measuring price elasticities for residential water demand with limited information. *Land Economics*, 90(1):100–113, 2014.

Jacob LaRiviere, Michael K. Price, Scott Holladay, and David Novgorodsky. Prices vs. nudges: A large field experiment on energy efficiency fixed cost investments. *Working Paper*, June 2014.

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

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

Shanjun Li, Joshua Linn, and Erich Muehlegger. Gasoline taxes and consumer behavior. Working Paper 17891, National Bureau of Economic Research, 2012.

Jens Ludwig, Jeffrey R Kling, and Sendhil Mullainathan. Mechanism experiments and policy evaluations. *Journal of Economic Perspectives*, 25(3):17–38, 2011.

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

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

Sheila M. Olmstead. The economics of managing scarce water resources. *Review of Environmental Economics and Policy*, 4(2):179–198, 2010.

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

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

Mary E. Renwick 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*, 40:37–55, 2000.

James M. Sallee. Rational inattention and energy efficiency. Working Paper 19545, National Bureau of Economic Research, October 2013.

P. W. Schultz, A. Messina, G. Tronu, E. F. Limas, R. Gupta, and M. Estrada. Personalized normative feedback and the moderating role of personal norms: A field experiment to reduce residential water consumption. *Environment and Behavior*, 2014.

P. Wesley Schultz, Jessica M. Nolan, Robert B. Cialdini, Noah J. Goldstein, and Vladas Griskevicius. The constructive, destructive, and reconstructive power of social norms. *Psychological Science*, 18(5):429–434, 2007.

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

Jeong-Shik Shin. Perception of price when price information is costly: Evidence from residential electricity demand. *The Review of Economics and Statistics*, 67(4):591–598, November 1985.

Victor Stango and Jonathan Zinman. Limited and varying consumer attention evidence from shocks to the salience of bank overdraft fees. *Review of Financial Studies*, February 2014.

Richard Stone. Linear expenditure systems and demand analysis: An application to the pattern of british demand. *The Economic Journal*, 64(255):511–527, 1954.

Aaron Strong and Chris Goemans. Quantity uncertainty and demand: The case of water smart reader ownership. *The B.E. Journal of Economic Analysis & Policy*, 14(3), January 2014.

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

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

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

Casey J. Wichman, 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.

Tables

Table 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 (years since 2014)	33.81	29.72	28.73	44.88	34.47
	(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 bill	81.96	91.21	89.02	84.13	84.62
($/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)
Number of households:	18,042	15,415	10,589	14,215	58,965

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).

Table 2: Timing of monthly billing transitions for each billing district

Billing district	Date of first monthly bill	Number of households
Transitioned:		
3	12/1/11	3154
9	7/13/12	3501
4	10/25/12	3349
10	11/14/12	3727
6	12/29/12	4311
15	1/29/13	3134
5	2/12/13	3688
16	3/30/13	3711
13	11/22/13	4882
17	1/30/14	4199
1	4/18/14	1508
11	5/15/14	5676
Never transitioned:		
2	6/23/14*	2217
7	N/A	3204
8	N/A	2088
14	N/A	3317
18	N/A	3299

Note: Billing district 2 is designated as "never transitioned" since no post-transition data is available for this district (that is, the study period ends in June 2014).

Table 3: Baseline difference-in-difference regression results

Dependent 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 fixed effects	N	Y	Y	N	N
Time fixed effects	N	N	N	Y	Y
Billing district fixed effects	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. The difference in observations between Columns (1)-(2) and (3)-(5) results from missing values for demographic or household characteristics. *** p<0.01, ** p<0.05, * p<0.1.

Table 4: 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
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 5: Pooled cross-section regression-discontinuity regression results

Dependent Variable: $\ln(w_{ijt})$	(1) Within 2000ft	(2) Within 1000ft	(3) Within 500ft
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
Additional controls:			
Full demographic covariates	Y	Y	Y
Weather covariates	Y	Y	Y
Time fixed effects	Y	Y	Y
Billing district fixed effects	Y	Y	Y
Household fixed effects	N	N	N

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 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. Constant term omitted. *** p<0.01, ** p<0.05, * p<0.1.

Table 6: Fixed effects difference-in-discontinuity regression results

Dependent Variable: $\ln(w_{ijt})$	(1) Within 2000ft	(2) Within 1000ft	(3) Within 500ft
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
Additional controls:			
Time fixed effects	Y	Y	Y
Household fixed effects	Y	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 local linear panel data 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. Constant term omitted. *** p<0.01, ** p<0.05, * p<0.1.

Table 7: 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
Weather covariates	Y	Y
Time fixed effects	Y	Y
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. Column (1) presents OLS regression results with demographic covariates and a lagged dependent variable. Column (2) presents results from a linear dynamic panel. *** p<0.01, ** p<0.05, * p<0.1.

Table 8: Heterogeneous treatment effects among seasons and automatic bill payment

Dependent Variable: $\ln(w_{ijt})$	(1)	(2)
BF	0.022	0.049
	(0.006)***	(0.005)***
	[0.022]	[0.018]**
BF×spring	-0.019	
	(0.005)	
	[0.034]	
BF×summer	0.131	
	(0.008)***	
	[0.063]*	
BF×fall	0.046	
	(0.005)***	
	[0.053]	
BF×abp		-0.064
		(0.015)***
		[0.020]***
Number of households	58,965	58,965
Observations	1,694,859	1,694,859
Within R-squared	0.007	0.006
Additional controls:		
Weather covariates	Y	Y
Time fixed effects	Y	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 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 Assumption 2, in which 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 9 and 19, 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.

Figures

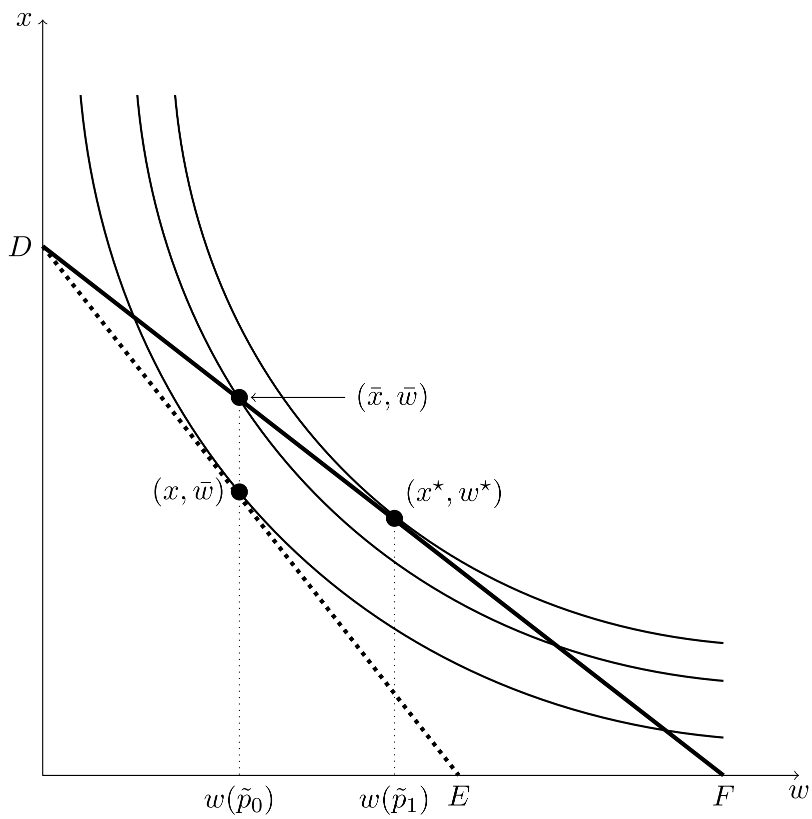

Figure 1: An example of a customer who (mis)perceives budget constraint DE and targets consumption bundle $(x, \bar{w})$ before the change in billing frequency. The actual bundle consumed, however, is $(\bar{x}, \bar{w})$. The change in billing frequency reflects a decrease in the perceived price of water from $\tilde{p}_0$ to $\tilde{p}_1$. With more frequent information, the customer realizes that she was over-perceiving water prices ($\tilde{p}_0$) and thus under-consuming water relative to her private optimum. This change in perception allows her to reallocate consumption along her true budget constraint, DF. So, she increases water consumption to $w(\tilde{p}_1)$, which allows her to obtain the preferred allocation $(x^{\star}, w^{\star})$.

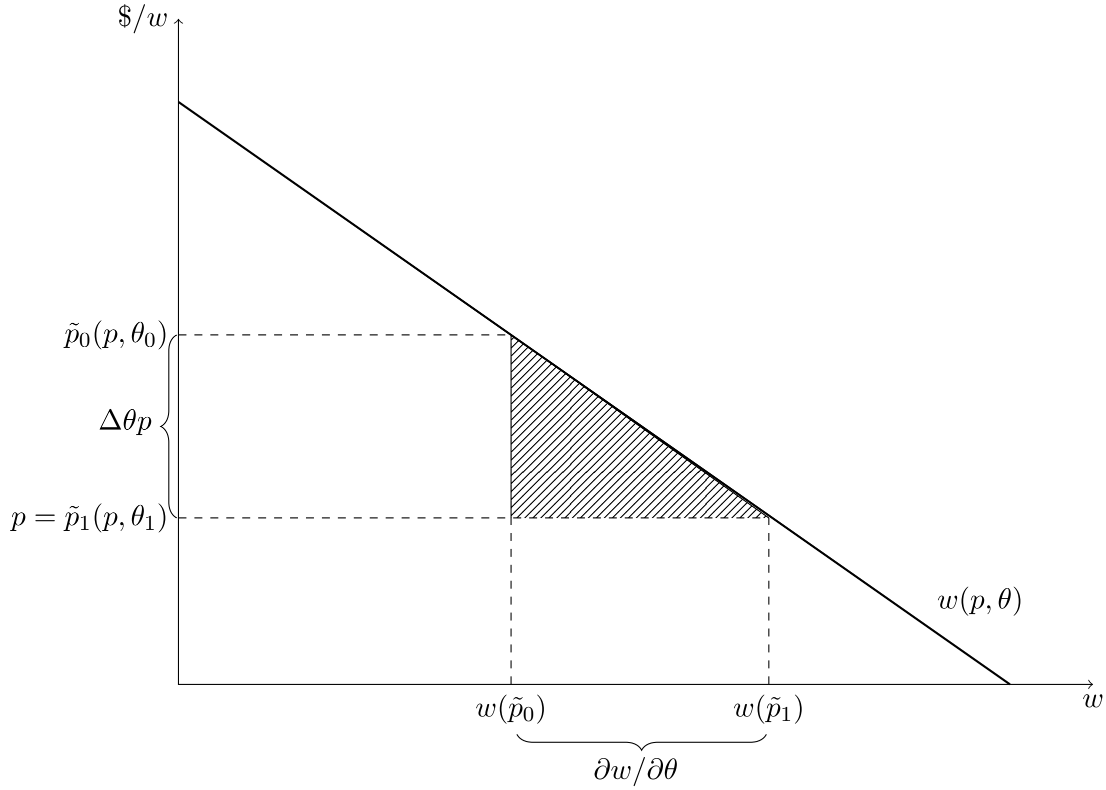

Figure 2: A stylized example of the welfare change from an increase in billing frequency (from θ_0 to θ_1). The shaded area represents the change in consumer surplus from the change in perceived prices ($\Delta\theta p$).

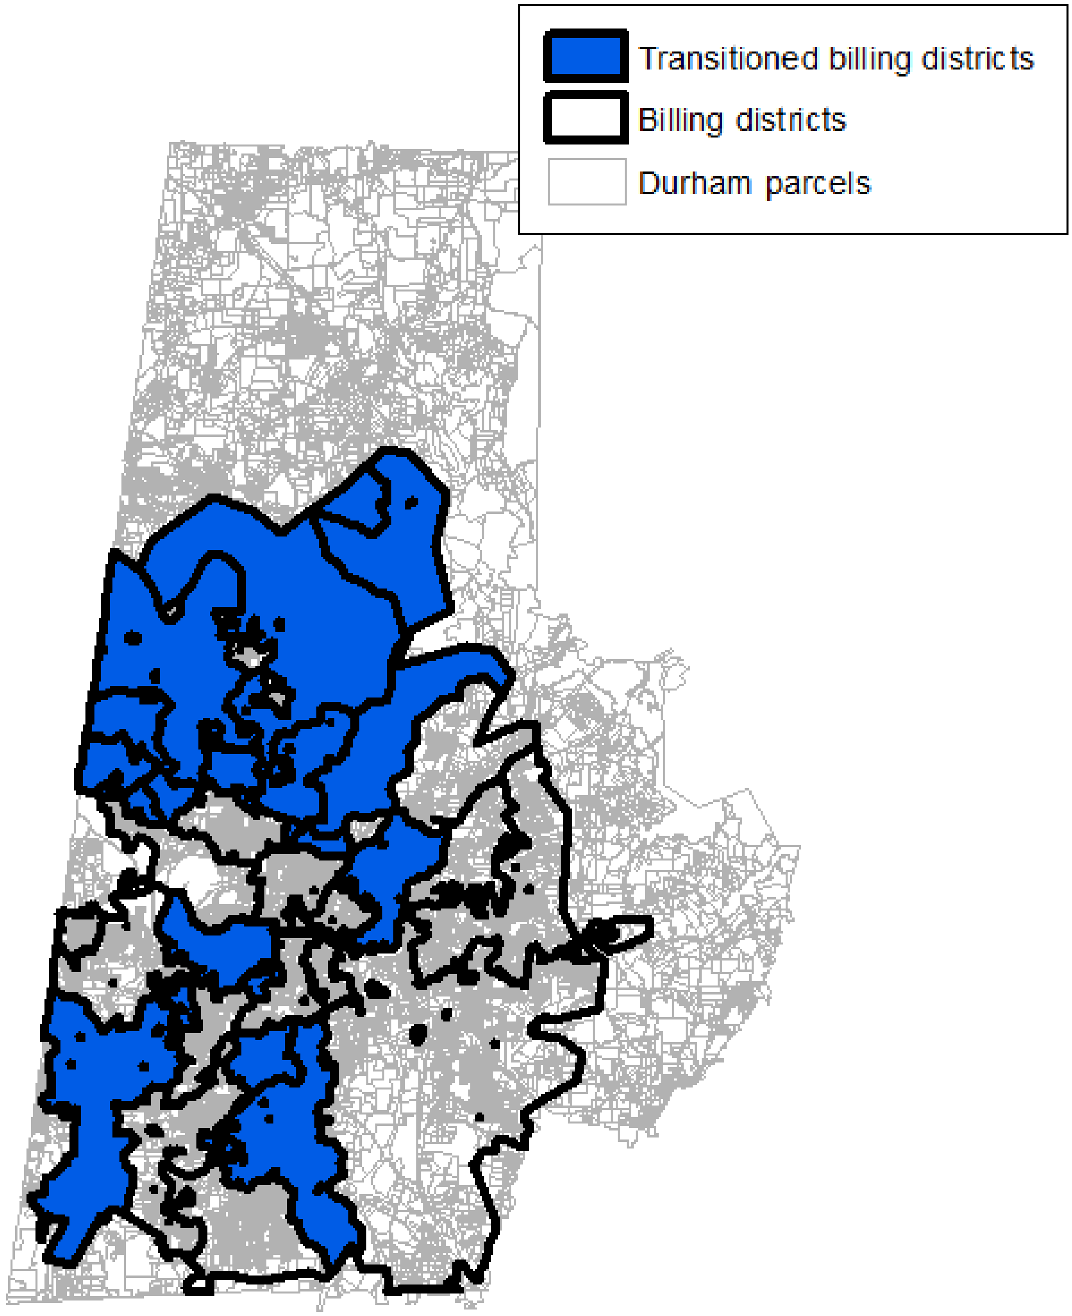

Figure 3: All billing districts transitioned to monthly billing by June 30th, 2013. Four additional districts transitioned to monthly billing within the timeframe of the study.

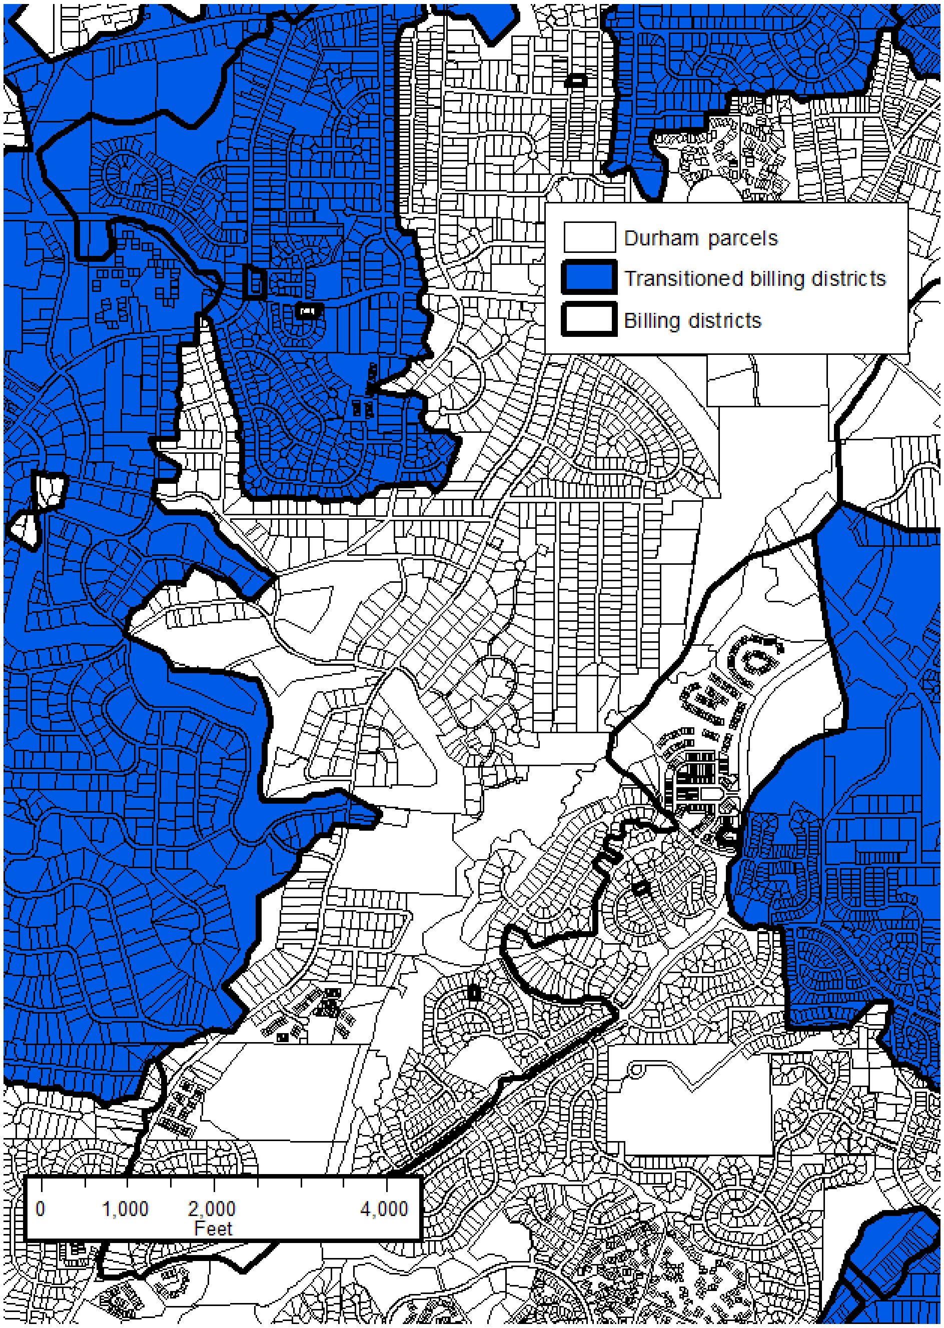

Figure 4: Depiction of billing group boundaries within neighborhoods.

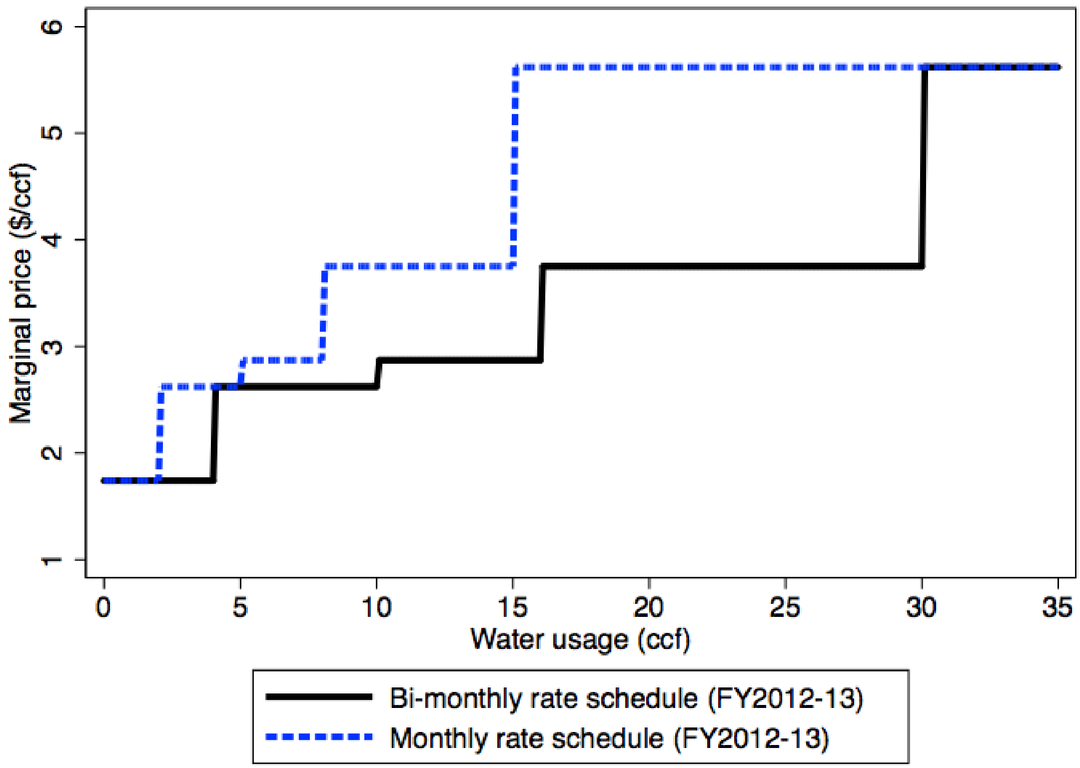

Figure 5: Increasing block rate structure before and after transition for monthly and bi-monthly billing. As shown, quantity blocks are halved for monthly billing relative to bi-monthly billing while marginal prices remain constant within the rate structure.

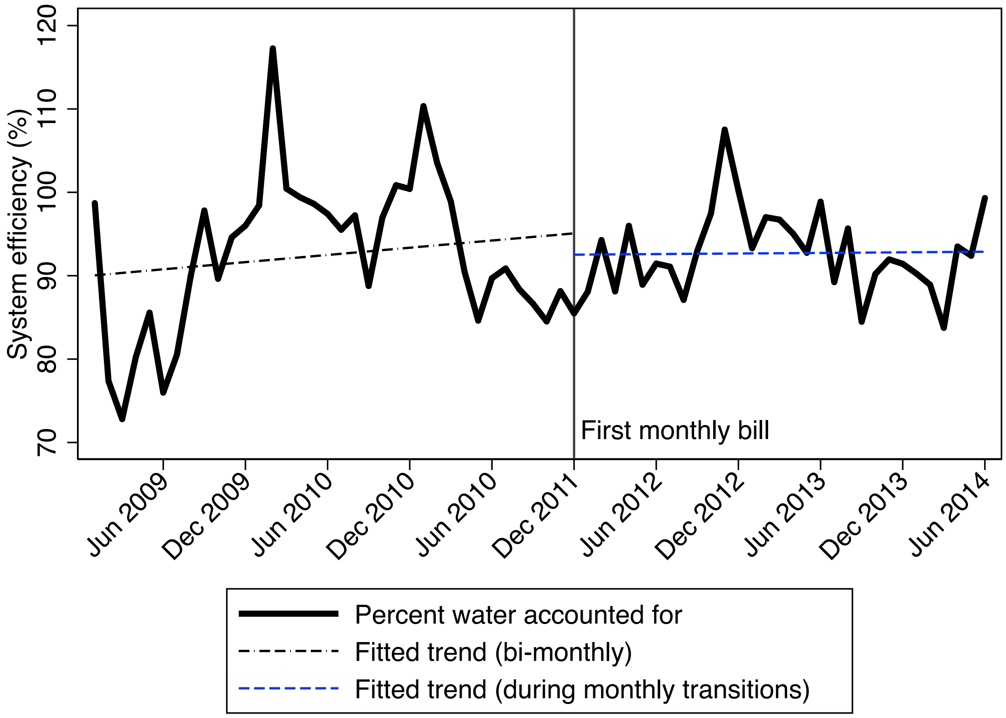

Figure 6: Percent water accounted for as a percentage of total pumped water. As shown, there is no discernible trend, increasing or decreasing, in percent water accounted for after the transition to monthly billing within the study period. Source: Annual Finanicial Information Reports (AFIR) submitted to the NC State Treasurer.

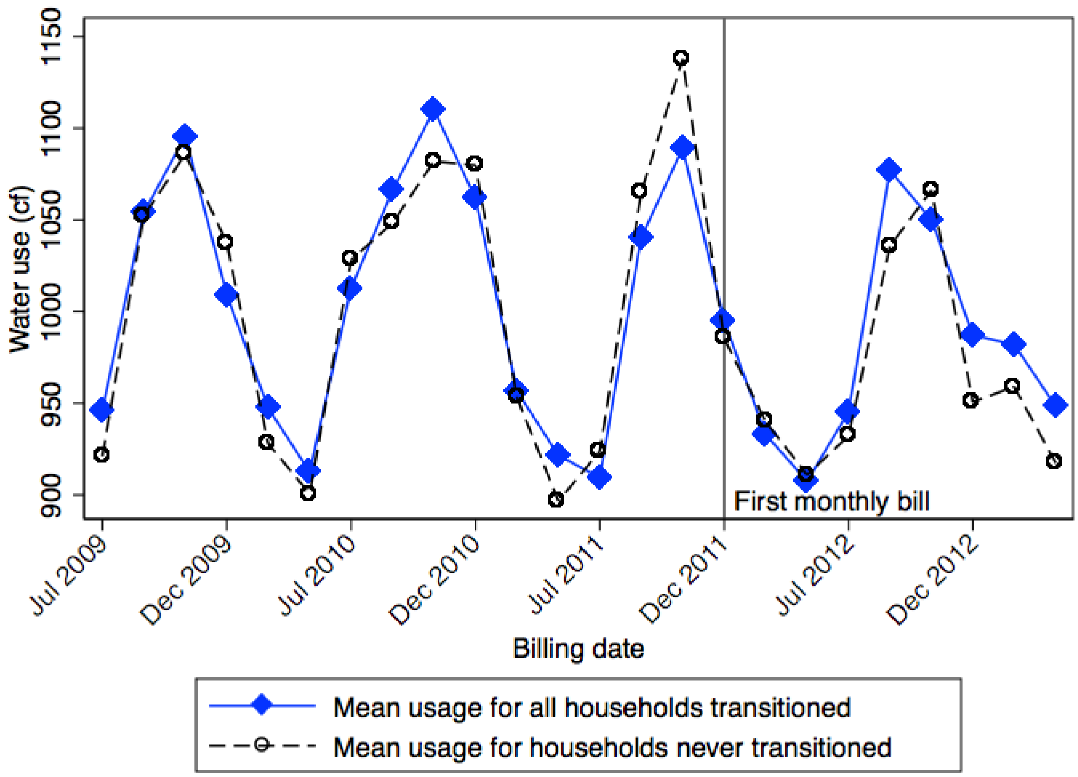

Figure 7: Mean bi-monthly consumption over time for households that transitioned to monthly billing and households that never transitioned to monthly billing.

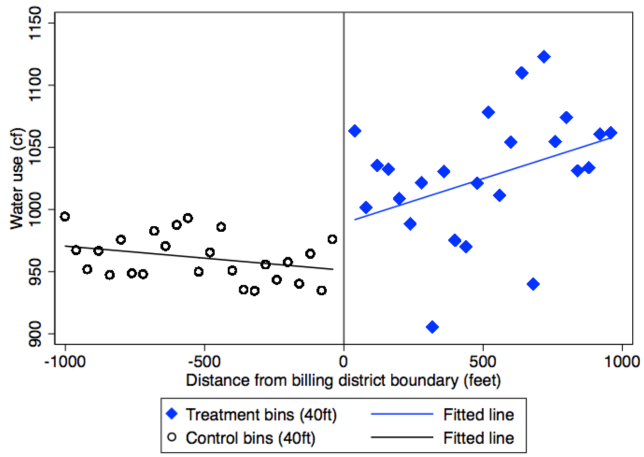

(a) 40-foot bins with linear fit

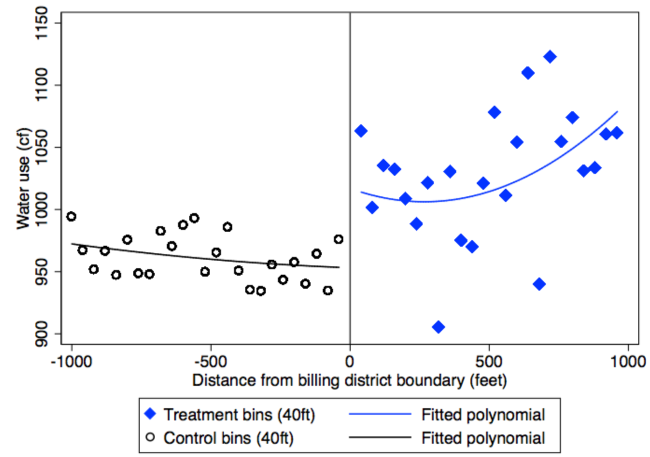

(b) 40-foot bins with quadratic fit

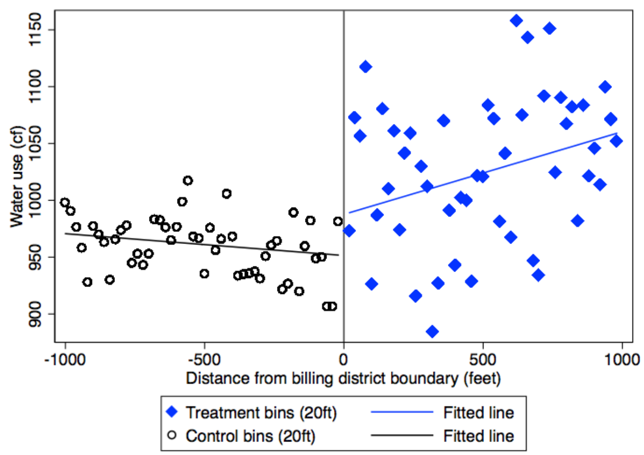

(c) 20-foot bins with linear fit

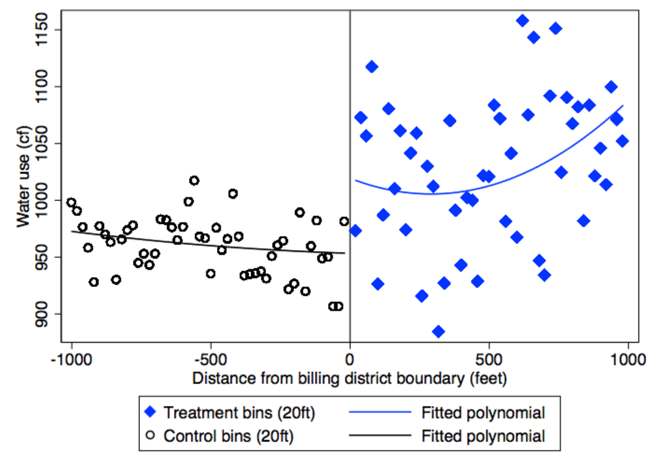

(d) 20-foot bins with quadratic fit

Figure 8: Mean consumption for 20- (40-)foot bins as a function of distance from district boundaries for consumption during the calendar year of 2012. Treated households are those that reside in districts transitioned to monthly billing within 2012, while control households are in all other districts.

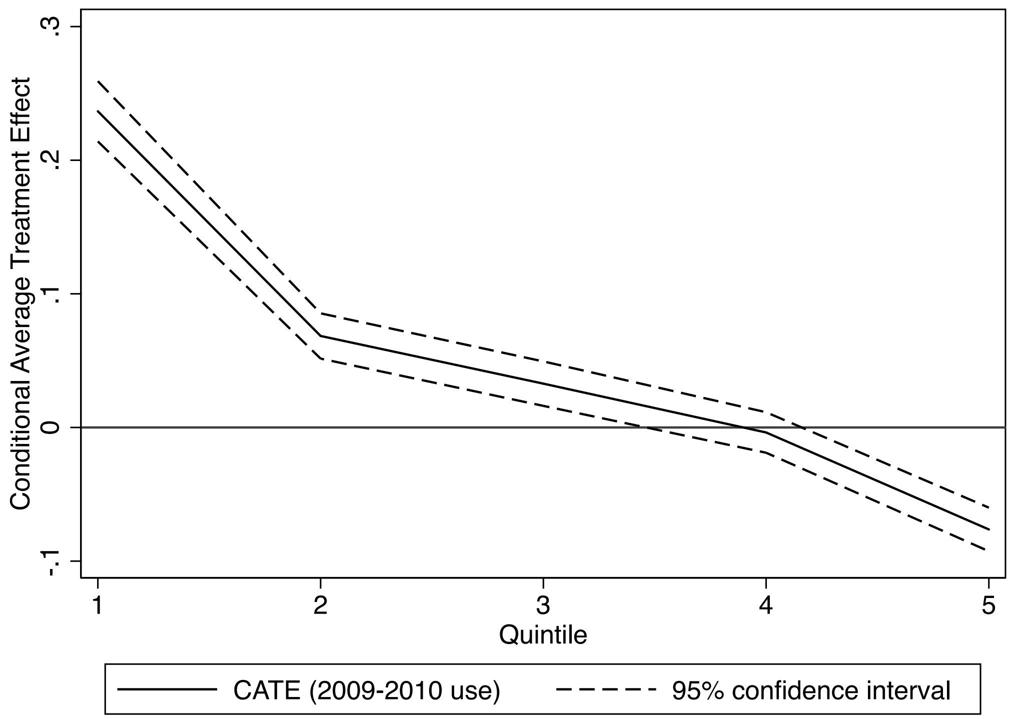

(a) ATE conditional on quintile of 2009-2010 consumption

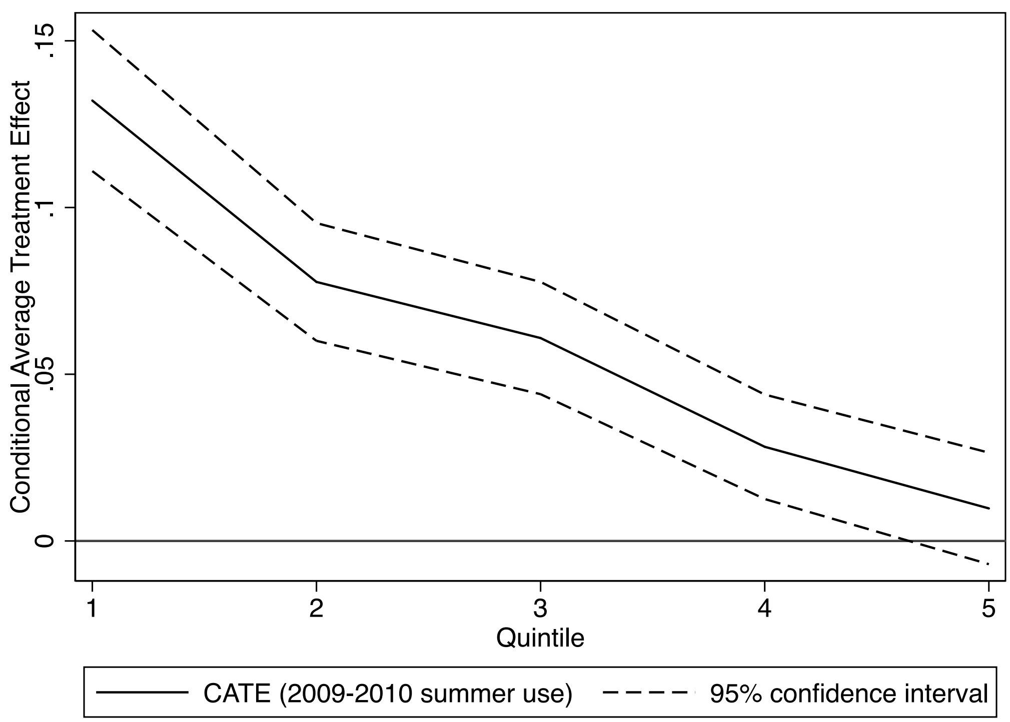

(b) ATE conditional on quintile of 2009-2010 summer consumption

Figure 9: Conditional average treatment effects for baseline consumption patterns.

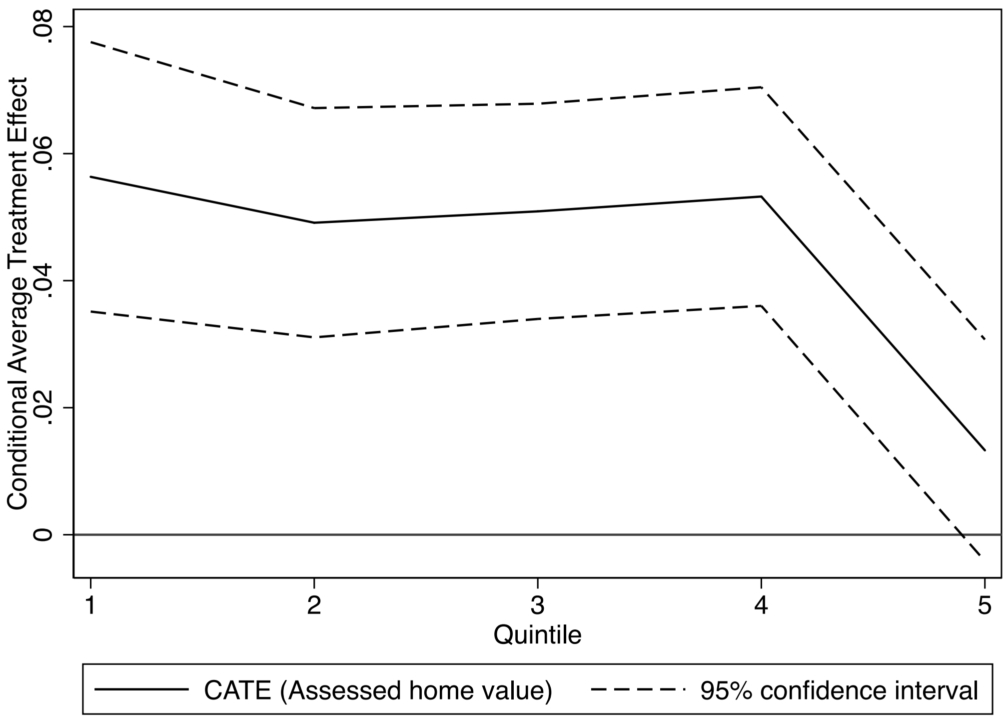

(a) ATE conditional on quintile of assessed value of home

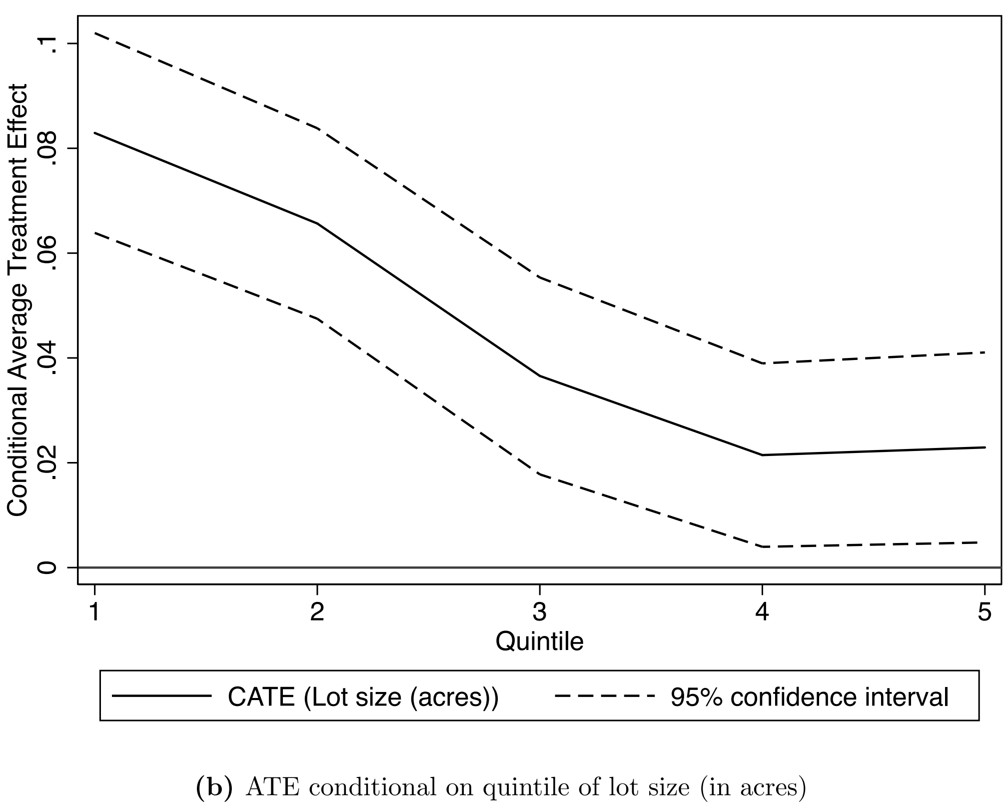

(b) ATE conditional on quintile of lot size (in acres)

Figure 10: Conditional average treatment effects for home value and lot size.

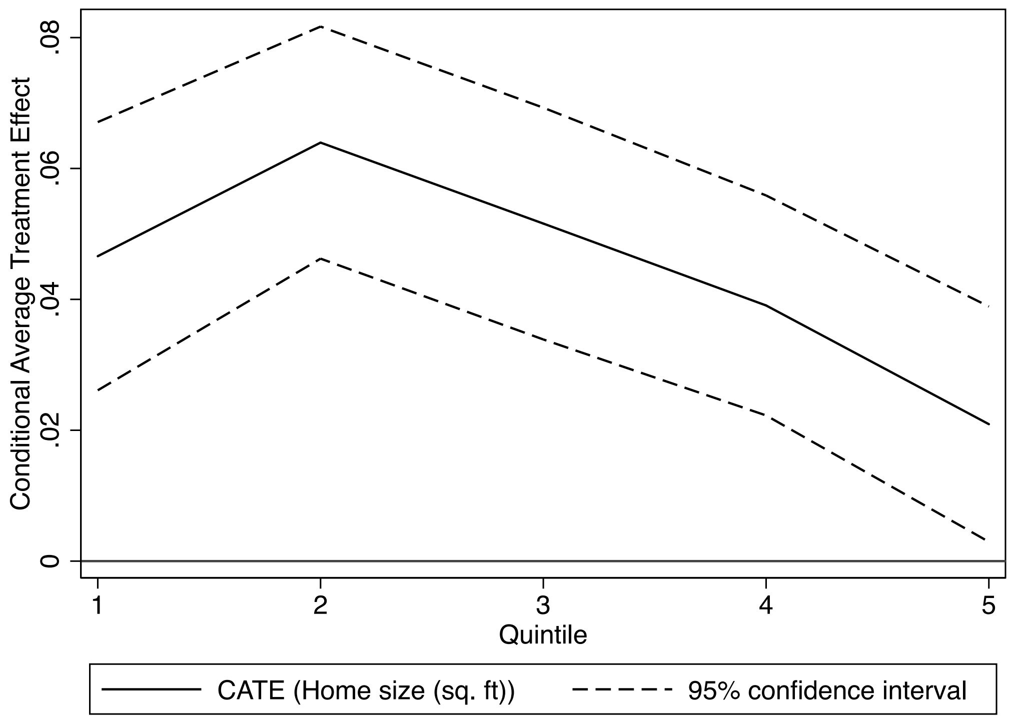

(a) ATE conditional on quintile of house size (in square feet)

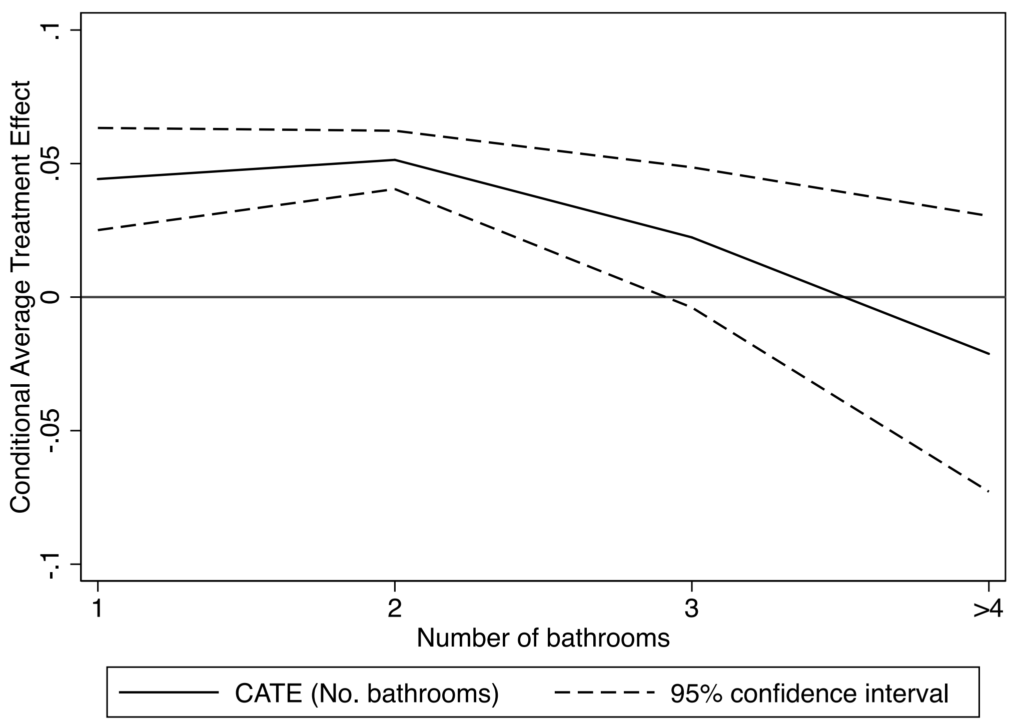

(b) ATE conditional on number of bathrooms

Figure 11: Conditional average treatment effects for structural characteristics of home.

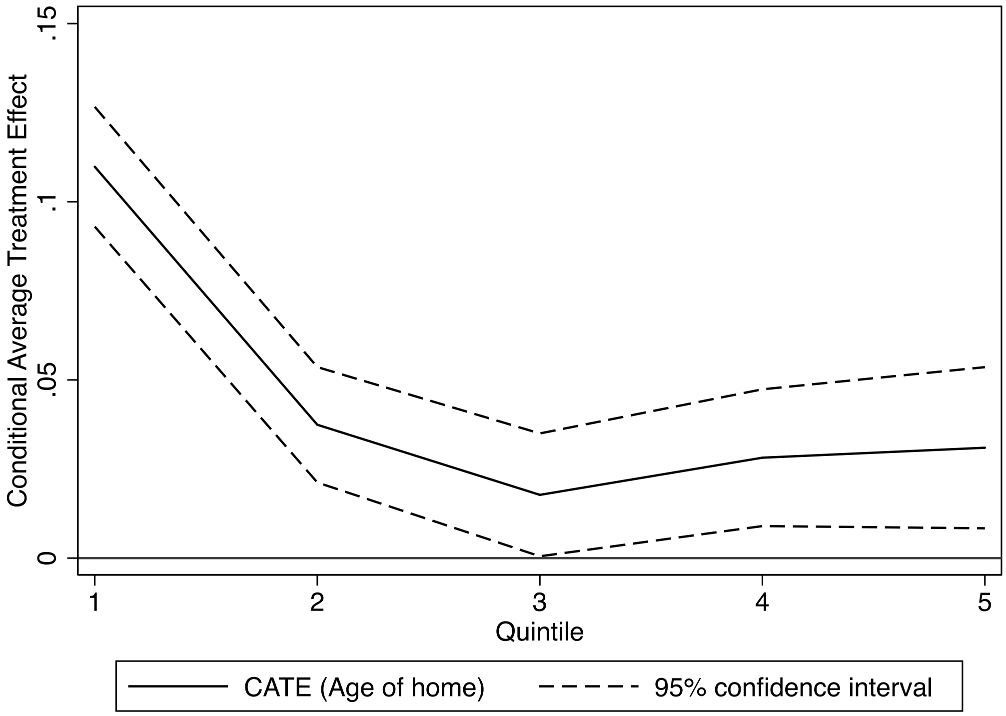

Figure 12: ATE conditional on quintile of age of home.

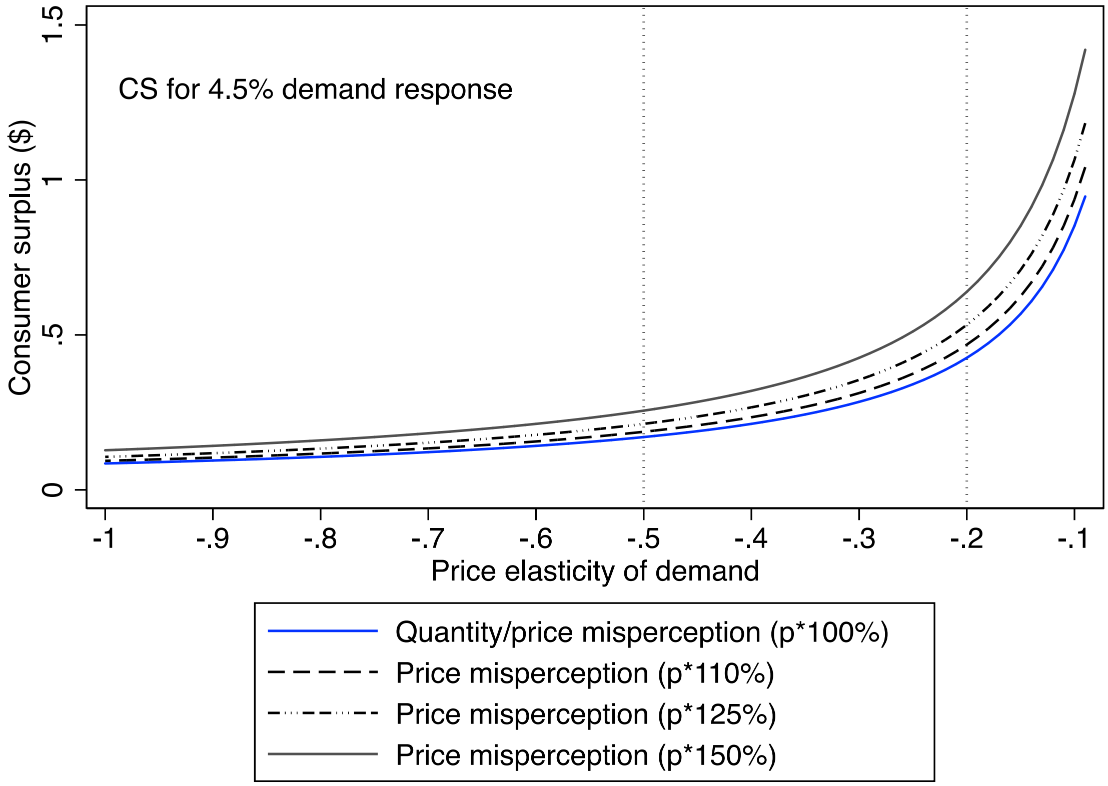

Figure 13: Excess burden from changes in billing frequency as a function of price elasticity for quantity and price misperception under different price scenarios.

Appendices

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[26],

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

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 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 $i = W, E, G$.[27] 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 decisions over household utilities may respond differently to a bundle of *monthly* bills for all three goods than a consumer who receives

[26] 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).

[27] 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.

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.[28] 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 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.

A.1 LA/AIDS simulation procedure

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,

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

to represent simulated electricity demand, and

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

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

Table A.1: LA/AIDS 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 Appendix A. 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 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.[29]

[29] 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.

B Derivation of demand functions under quantity misperception

The consumer's problem is

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

which provides the following necessary condition for an interior solution,

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

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

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

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,

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

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

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

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

C Additional results

Table C.1: Predicting the likelihood of billing districts to be transitioned from b-monthly to monthly billing based on observable household characteristics

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

Table C.2: Automatic bill payment subsample results

Dependent Variable: $\ln(w_{ijt})$	OLS (1)	OLS (2)	FE (3)
BF	0.002	0.036	0.028
	(0.037)	(0.021)*	(0.021)
	[0.074]	[0.044]	[0.041]
Number of households	—	—	2,167
Observations	66,272	66,272	66,384
R-squared	0.075	0.074	—
Within R-squared	—	—	0.012
Additional controls:			
Full demographic covariates	Y	Y	N
Weather covariates	Y	Y	Y
Time fixed effects	Y	Y	Y
Billing district fixed effects	N	Y	N
Household fixed effects	N	N	Y

Note: Robust standard errors clustered at the household-level in parentheses. Robust standard errors clustered at the household-level in square brackets. *** p<0.01, ** p<0.05, * p<0.1.

Table C.3: 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

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 regression 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.