Enduring effects of changes in billing frequency: Evidence from urban water use

Casey J. Wichman*

January 3, 2017

DRAFT PREPARED FOR 2017 ASSA CONFERENCE
PRELIMINARY AND INCOMPLETE.
PLEASE DO NOT CITE.

Abstract

Little is known about the way consumers update their behavior in intermittent choices settings. I exploit a natural experiment in which residential water customers were switched from bi-monthly to monthly billing to assess the long-term responses to increases in billing frequency. I find that customers increase water consumption, but the estimated average treatment effect weakens over time and its profile differs based on historical consumption patterns and preferences for outdoor water use. Inattentive consumers respond to the treatment by a smaller magnitude and for a shorter duration. Notably little heterogeneity is found among wealth groups in the dynamic response to monthly, which suggests that more frequent billing may not aid in smoothing expenditures.

JEL Codes: D12, D61, H42, L95, Q21, Q25
Keywords: information provision, billing frequency, natural experiment, water demand, dynamic treatment effects

*Resources for the Future. *Author email:* wichman@rff.org. *Author address:* 1616 P Street NW, Washington, DC 20036. Rob Williams, Grant Jacobsen, and participants at the 2015 APPAM conference, the 2016 AERE conference, and 2016 Camp Resources provided helpful comments and suggestions.

1 Introduction

Little is known about the way consumers update their behavior in intermittent choice settings. In this paper, I exploit a natural experiment in which residential water customers were transitioned from bi-monthly to monthly billing to assess dynamic responses to increases in billing frequency.

Previous research has found significant changes in electricity and water use induced by nonpecuniary incentives (Allcott, 2011; Ferraro and Price, 2013; Brent et al., 2015), although the mechanism by which consumers alter their behavior is unclear. Some behavioral interventions influence consumption through moral channels (e.g., Ferraro and Price (2013)), while other informative interventions allow consumption or prices to become more salient (e.g., Jessoe and Rapson (2014); Kahn and Wolak (2013); Wichman (2015)).

The durability of the effects induced by behavioral and informative interventions remain even more ambiguous. Allcott and Rogers (2014) find that social comparisons reduce electricity consumption in the short run, but consumers become increasingly numb to these treatments over time. Bernedo et al. (2014) find, however, that the effect of a one-shot behavioral treatment can be observed up to six years after treatment for water conservation efforts, with evidence suggesting that both habitual changes and structural changes to the home drive this effect. Contrastingly, Ito et al. (2015) observe dissipating short-run reductions in electricity use for randomized moral suasion treatments, but longer-lasting effects from economic incentives, positing a mechanism of habit formation.

In this paper, I build on previous research by examining consumption in response to changes in the frequency of bills received by residential water customers for a large water utility in the Southeastern US. The empirical setting is a conditionally exogenous transition of residential water customers from bi-monthly to monthly billing for a single water utility. Beginning in 2011, the City of Durham's Department of Water Management in North Carolina transitioned residential customers in geographically differentiated billing districts to monthly billing over the course of three and a half years. By exploiting the assignment of monthly billing, I estimate a causal increase in consumption in response to more frequent information. The average treatment effect, however, masks important heterogeneity over time that better highlights the dynamic response of economic agents in response to repeated interventions. Treatment effects are found to persist but weaken over time to effectively no response after 12 months.

Because water customers increase consumption at a decreasing rate over time, my empirical results suggest a mechanism other than price by which consumers respond to changes in billing frequency. Rather, I interpret these effects as the product of habitual heuristics in which consumers are repeatedly reminded of how much water they use and, importantly, how inexpensive water is relative to other budgeted monthly expenditures.

I extend the literature further by examining heterogeneity in the dynamic response to billing frequency among socioeconomic and structural characteristics of the household. My results suggest that historically low users of water and households on larger lots display the largest dynamic response to increases in billing frequency. Notably, little heterogeneity is found among wealth groups in the dynamic response to more frequent billing. This finding suggests that more frequent billing does not aid in smoothing expenditures, which was a purported motivation for the adoption of monthly billing. This contradicts recent studies (e.g., Jack and Smith (2015)) that suggest increasing the frequency of billing via prepaid metering for electricity may help to avoid backing up against credit constraints with lumpy billing practices. I find no evidence consistent with the hypothesis that more frequent bills, and thus smaller billed amounts, allow constrained consumers to smooth income over the year.

Overall, this research makes contributions along several dimensions. This paper provides the first causal analysis of long-run behavior in response to changes in billing frequency for economic goods purchased intermittently. Second, I find empirical evidence that economic agents increase consumption at a decreasing rate in response to more frequent billing information—both of which are novel findings in program evaluations of price and non-price interventions for water and electricity consumption. Third, the effects for inattentive consumers (as proxied by enrollment in automatic bill payment) exhibit a dampened profile of treatment effects, where effects are smaller in magnitude and persist for a shorter duration. And fourth, this analysis suggests that more frequent billing has no meaningful effect on helping credit constrained consumers smooth expenditures throughout the year.

In the next section, I describe the data used in the analysis and outline the empirical setting. I present a series of quasi-experimental models to estimate a causal effect of billing frequency on consumer demand as well as dynamic responses to treatment in Section 3. In Section 4, I discuss the results and implications of the empirical models. The final section concludes.

2 Empirical setting

Beginning in December 2011, the City of Durham's Department of Water Management in North Carolina (henceforth, "Durham") transitioned individual billing districts from bimonthly to monthly billing at different points in time. Primary reasons for the transition include cost-saving from fewer delinquent payments, early leak detection, improving customer service, and reducing administrative costs. In addition, the change in billing frequency was enabled by district-wide installation of automated meters. The new meters allow for consumption levels to be obtained via radio frequency such that the costs to read meters manually were reduced. Meters were installed for each billing district and, once the installations were completed, the entire district was transitioned to monthly billing. Customers were notified of the transition to monthly billing by mail approximately six weeks before the transition.[1]

To make a cost-saving argument to the city council, the water utility used a single billing district as a pilot group to measure changes in administrative costs before and after the transition to monthly billing cycles. After that, billing districts were transitioned to monthly billing according to meter installation and administrative schedules. The order of districts for meter installation (and, subsequently, monthly billing) was chosen to work around billing cycles and other feasibility constraints. According to utility officials, no consideration of billing history, income base of neighborhood, or any other financial indicator was taken into account when choosing which districts to transition.[2]

Given these details, the assignment of monthly billing is plausibly exogenous to the household, conditional on residing within a particular billing district. The household has no ability to manipulate the assignment of billing frequency short of moving across billing cycle boundaries. Within the study period, billing districts were transitioned according to the timing in Table 1. The first district transitioned received its first monthly bill on December 1, 2011. Figure 1 presents a series of maps of the districts that switched to monthly billing. The entire service area is represented by the union of all billing districts outlined in bold. In Figure 2, I present a magnified view of billing district boundaries within neighborhoods. This figure illustrates that the district boundaries are designated in such a manner that neighbors could be consuming water concurrently, but may be billed at different frequencies. Thus, this design allows for the exploitation of geographic discontinuities to minimize the concern that selection into treatment is nonrandom and that unobserved changes in neighborhood characteristics might bias results.

[1] A copy of the mailer distributed to customers is included as Figure A.2 in the appendix.

[2] In Wichman (2015), I present regression results that predict the likelihood of a billing district to be transitioned from bi-monthly to monthly billing based on observable household characteristics. In particular, I consider the first district to transition in a probit model, as well as the sequential transition of routes in an ordered probit framework. The former suggests that lot size, household size, and the number of bedrooms increased the probability of the pilot group being chose, while the square footage of the home and number of bathrooms decreased the probability. There is weak evidence that smaller water bills are correlated with initial treatment. Further, the ordered probit reveals that the age of a home is weakly negatively correlated with the order of transition, while all other observable factors are insignificant. These results help to establish conditional exogeneity of treatment.

2.1 Data

The primary data used in this analysis are residential billing records for Durham water customers. Included in these data are (bi-)monthly water and sewer use, fixed service fees and volumetric consumption rates, the address of the customer, billing district, and whether a customer has her water bill automatically deducted from a bank account. The billing data were matched by address with geocoded tax assessor data, containing structural characteristics of the home, obtained from Durham County. Each matched residential address was spatially linked to its 2010 Census block as well as billing district polygons provided by Durham. For each household, I determine the nearest billing district, as well as the linear distance from the centroid of the tax parcel to its nearest district boundary. Key demographic variables from the 2010 SF1 Census are matched to each household's Census block. Residential premises that changed water billing accounts within the timeframe of the study are removed from the sample—this strategy reduces the impact of renters, who may not pay water bills explicitly. Further, this avoids econometric identification problems when relying on variation within a household over time.[3]

The final sample consists of roughly 59,000 individual household accounts with water bills from February 2009 through June 2015, which implies slightly more than 2 million unique consumption observations. Summary statistics for variables of interest are presented in Table 1. The first four columns decompose household characteristics and details on water use by the year in which households transitioned from bi-monthly to monthly billing. Summary statistics in the final column are for the entire sample. All 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 that transitioned in 2013 along several dimensions. For this group, home value (a proxy for wealth) is notably larger than that of all other groups. Further, these households tend to have larger homes on larger lots, and are more likely to be located in a Census block with fewer renters and a higher proportion of white residents. Water consumption, however, is statistically similar across all groups. For the typical household in the sample, the mean assessed home value is approximately \$186,000 with a standard deviation of \$126,000. The average home is on one-third of an acre, 34 years old, roughly 1,800 square feet, with three bedrooms. Within the final sample, households reside in Census blocks in which approximately one-quarter of all homes are renter-occupied. Fifty-three percent of the sample is white and the average household size is between two and three people. Average bi-monthly water bills for all time periods in the sample are \$85 for

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

consumption of 985 cubic feet of water.

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

A final caveat is that under bi-monthly billing, water bills are mailed on a staggered schedule that smooths administrative work and meter reading throughout the year. As an example, a billing district on the odd cycle may receive a bi-monthly bill in March for consumption in January and February. Contrarily, a billing district on the even cycle would receive a bill in April for consumption in February and March. Rather than dealing with these two groups independently, I pool households prior to treatment 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. After treatment, I use monthly consumption in the month in which it occurred. For both billing regimes, I divide the total consumption within the period by the number of days in the billing cycle so that the primary outcome of interest is daily average consumption regardless of whether households are being billed monthly or bi-monthly.

2.2 Prices

At the same time households were switched to monthly billing, fixed water and sewer service fees were halved and volumetric block cut-offs in the tiered water rate structure were halved as well. Marginal volumetric rates for consumption remained constant across billing frequencies. Sewer usage has a constant marginal price. Figure A.3 illustrates the change in the rate structure for monthly and bimonthly billing. The solid line is the increasing block rate structure used to calculate bimonthly 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 bimonthly rate structures, but the quantity blocks for consumption are halved for each price tier.

This structure was adopted to ensure that customers transitioned to monthly billing were charged at the same rate as bimonthly customers. Thus, for the same level of consumption, two monthly bills are equivalent to one bimonthly bill in dollar amounts. This is a mechanical interpretation, however, and the change in the block endpoints could affect consumer behavior. To see this, consider an extreme version of the average water customer. She is extreme because she consumes no water in the first month and 10 ccf in the second month. Under bi-monthly billing, her total bill is \$82.12. Under monthly billing her first monthly bill is \$12.05

(i.e., fixed service fees only) and her second bill is \$74.80, totaling \$86.85 for the two-month period, an *increase* in expenditures on water. The difference in the billed amounts for the same quantity consumed is a result of the nonlinear rate schedule: highly variable month-to-month consumption results in larger inframarginal price changes. The percent change in the average price of water for this extreme customer is +5.6%. If we assume a common elasticity of -0.3, we would expect this change in average price to *reduce* consumption by 1.6% (or, 2.4% if we look only at changes in average volumetric prices). These percentages tend towards zero as the monthly consumption levels converge. This exercise shows that the nonlinear rate schedule can have an effect on price signals, although they work in work in the opposite direction as the treatments effects identified below. Further, the likelihood of this type of behavior being representative for Durham water customers is virtually zero.

Of course, the extent to which this bias exists also depends on whether consumers know and use the tiered rate information to make decisions. Because the water utility bill includes no information about the block rate structure (see Figure A.1), it is unlikely that consumers are responding to changes in the rate structure itself.[4] Further, Wichman (2014) and Ito (2014) show that water and electricity customers, respectively, exhibit behavior that corresponds to changes in average price, or the total bill, when facing increasing block rates.

2.3 Meters

Durham installed automated meter reading (AMR) technology prior to switching billing districts to monthly billing. Based on conversations with utility representatives, the order of districts was chosen for new meter installation based on geographic convenience (see, e.g., Panel B of Figure 1), difficulty of reading meters, district size, and working around the odd-even billing schedule. Once a district reached 90% saturation with new meters, the district was ready to be switched to monthly billing.

The new meters allow for consumption levels to be obtained via radio frequency such that the costs to read meters manually were reduced. Meters were installed for each billing district and, once the installations were completed, the entire district was transitioned to monthly billing. In the water industry, meters may fail to register all water that passes through them over time. Meter replacement offers an avenue through which consumption could increase mechanically; that is, the treatment effect may also include the increase in accuracy of the meter. This "mechanical efficiency" improvement depends on the degree of

[4]As further evidence, I present the empirical density of consumption in Figure A.4 in the appendix. In this figure, there is no evidence of bunching at the block rate cut-offs for consumption in the calendar year prior to treatment.

inaccuracy of the meter being replaced, water pressure, appropriate meter sizing, and a host of other factors that are, at present, unobservable. Although manufacturers often tout the benefits of improved meter reading accuracy, there is mixed evidence of this effect in the industry and engineering literature (see, e.g., Boyle et al., 2013; Lovely, 2010; Barfuss et al., 2011; Criminisi et al., 2009; Arregui et al., 2006).

Many of the frequently cited benefits of smarter water meters for utilities are captured by Ritchie (2011), "...utilities across the US are proving that AMI can drive down costs in unaccounted water, plus other important areas, such as energy, labor, conservation, capital investment, forecasting, billing, and customer service." Unaccounted-for water is a common metric of lost revenue. A case study in Leesburg, VA, found that system-wide meter replacement helped reduce unaccounted-for water from upwards of 23 percent down to less than 5 percent (Shoemaker, 2009). They also found that the older meters were under-registering water consumption for households. But age of the meter is not the only factor that matters for efficiency. In McKinney, TX, a utility-wide installation of new AMR meters revealed that many of the newly installed residential meters were undersized for their purpose and only registering a fraction of actual water use (Dobbie et al., 2003). Further, Britton et al. (2013) suggest that "there is still limited understanding of meter accuracy when considering the starting or minimum registrations levels (Qs), therefore water with a flow rate that is below the Qs flow rate of the meter cannot be measured." All of this is to suggest that although meter accuracy potentially changed for Durham households within the study period, there is not a clear direction of this bias, nor is the potential bias a static issue.

Further, the City of Durham's primary goals for installing AMR meters included reduced payment delinquency, earlier leak detection, and reduced administrative costs from manual meter reading. No mention of improved meter accuracy was cited in publicly available documents, which arguably would be a boon in city council deliberations for a revenue-conscious municipal water utility.[5]

3 Empirical strategy

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

[5]See, e.g., `https://durhamnc.gov/ArchiveCenter/ViewFile/Item/1203` and `https://durhamnc.gov/2983/Water-Meter-Replacement-Automated-Meter-`.

3.1 Event study

To explore any differential trends in the treatment group that may invalidate the identification strategy, I use an event study that plots coefficients from the following regression,

$$w_{ijt} = \sum_{S=-12}^{S=30} \gamma_S 1[\Delta_{ijt} = S]_{ijt} + W_t'\gamma_W + X_i'\gamma_X + \lambda_t + \alpha_j + \varepsilon_{ijt}. \tag{1}$$

where w_{ijt} is average daily water consumption in cubic feet (cf) for household i in district j at time t. Δ_{ijt} denotes the distance, in time, from when a billing district was transitioned to monthly billing, with $\Delta_{ijt} = 0$ denoting the period in which monthly billing was enacted. C_t is a vector of weather controls, X_i is a vector of household characteristics, and λ_t and α_j are period-of-sample and route fixed effects. The set of γ_S coefficients and 95 percent confidence intervals are plotted for 12 months before before and 30 months after treatment in Figure 3. Standard errors are clustered at the billing district level.

As shown in Figure 3, the time fixed effects appear to control reasonably well for seasonality, although there is evidence of slight seasonal patterns prior to treatment. At the time of transition to monthly billing, there is a dramatic increase in consumption that remains throughout the duration of the study period. The results of this exercise provide suggestive evidence that there is a significant, long-term trend in consumption attributable to the change in billing frequency. These estimates, however, do not control for the nonrandom selection of billing districts or for unobserved household heterogeneity.

3.2 Empirical models

To improve upon the estimates in Equation 1, I specify a difference-in-difference framework in which households billed on a monthly basis are treated, whereas households who are billed bi-monthly are comparisons. Due to the staggered introduction of monthly billing for different billing districts, the coefficient on the treatment indicator BF_{jt} is identified by changes in consumption for treated households relative to similar comparison households. The model takes the form,

$$w_{ijt} = \beta BF_{jt} + W_t'\gamma_W + X_i'\gamma_X + \lambda_t + \varepsilon_{ijt}. \tag{2}$$

where BF_{jt} equals one if district j is billed monthly in period t, and zero otherwise. All other variables are defined as in Equation 1, except that X_i includes household fixed effects in some specifications.

With conditional exogeneity of treatment assignment and common trends between treatment and comparison groups, γ_B will capture the causal effect of changes in billing frequency. Common trends in the outcome variable for treatment and comparison groups are presented in Figure 4. As shown, the consumption patterns for two sets of treatment and comparison groups are nearly identical leading up to the date at which the first district to transition received its first monthly bill. Thus, the estimate of β from Equation 2 provides an average treatment effect (ATE) of the change in billing frequency on water consumption. See Wichman (2015) for further discussion.

Conditional exogeneity of treatment, however, is more difficult to establish. One could make this assumption by noting the fact that billing districts were chosen to transition to monthly billing somewhat arbitrarily. Even arbitrary choice allows for nonrandom differences across billing districts among covariates that may affect the consumption response to changes in billing frequency. So, I adopt a border discontinuity model that exploits the fact that neighbors could be consuming water at different billing frequencies (see Figure 1). This approach is effectively a regression discontinuity design across the border of billing districts that jointly exploits the flexibility of the panel data in controlling for unobserved household heterogeneity.

I estimate Equation 2 for households within 2000, 1000, and 500 feet of the border discontinuity. The estimate of β is thus a local average treatment effect (LATE) in the neighborhood of the discontinuity. This spatial panel RD design avoids concern of nonrandom selection into treatment. Further, this approach controls for unobserved changes in neighborhood characteristics over time. In Figure 5, I present mean average daily consumption in 40-foot bins as a function of distance to the border boundary for three time periods. The first panel shows no change in consumption across the boundary. The second and third, representing the following three years after the change in billing frequency, show a discontinuous increase in consumption that can plausibly be attributed to the change in billing frequency.

Of course, the primary threat to identification of a LATE is that nothing else is changing at the border discontinuity. In Figure 6, I present observable household characteristics that influence water consumption. For each panel, mean statistics for each household are plotted as a function of its distance to the border discontinuity for the first 9 districts to transition. For the majority of observed covariates, the distribution moves smoothly through the border discontinuity. For the percentage of white residents within a Census block, however, there is a shift in the trend near the discontinuity. It is, however, difficult to come up with a hypothesis that would explain how race (conditional on other covariates in X_i) affects water use. As such, these distribution tests provide evidence that a regression discontinuity design is appropriate in this scenario.

3.3 Dynamic treatment effects

In addition to estimating parameters that summarize the overall treatment effect, I adopt an approach to estimate the dynamics of the treatment effect itself. This specification mimics the intuition of an event study, but provides a stronger causal foundation because it is based upon the spatial panel RD design. Specifically, I re-estimate Equation 2 to obtain a set of "dynamic" LATEs by decomposing the local average treatment effect as follows,

$$\sum_{S=0}^{30} \beta_S \cdot 1(BF_{jt-S}).$$

where BF_{jt-D} is a set of lagged treatment indicators for D periods after the initial transition to monthly billing. The the series of β_S estimates provides a time profile of the ATE that I refer to as the dynamic local average treatment effect (D-LATE). I estimate 30 periods of β_S coefficients after initial treatment, which provides a profile of monthly treatment effects for two-and-a-half years after initial treatment.

3.4 Heterogeneity

In addition to looking at the aggregate dynamic ATEs, I estimate Equation 2 on several subsamples of the data to explore heterogeneity in dynamic treatment effects. Specifically, I focus on quartiles of three observable characteristics of consumers. First, I segregate households into quartiles of historical consumption patterns. To construct these quartiles, I use mean water consumption before treatment occurred in 2009 and 2010. Next, I explore assessed value of the home as a proxy for household wealth. Lastly, I explore lot size, in acres, as it is often used as a proxy for preferences for outdoor water use (Wichman et al., 2016).

To explore whether the change in billing frequency is salient to the customer (Sexton, 2015), I estimate models for subsamples of customers who are enrolled in automatic bill payment throughout the transition to monthly billing.

4 Empirical results and discussion

In this section I briefly discuss the empirical findings and implications of the models outlined above.

First, in Table 2, I present average treatment effects corresponding to Equation 2 with various household and Census block-level controls. In columns (2) through (4), I shrink the window around the border discontinuity progressively. The primary coefficient of interest remains robust to this stratification and suggests an increase in water consumption of approximately 2.1 cubic feet per day, which is similar to the estimates in Figure 3.

Next, I present the same set of models, but I add household fixed effects. The preferred estimates are in the last column of Table 3. As shown, the preferred LATE is approximatelyca cubic feet per day, which translates to a 5.3 percent increase in consumption (relative to a pre-treatment mean of 12 cf per day) in response to the switch to monthly billing. The large change in the estimates between OLS and panel models suggests that controlling for unobservable household heterogeneity plays a more important role than nonrandom selection into treatment (as indicated by moving from the Column (1) to (4) in each of the results tables).

Although these ATE estimates are virtually identical to that of Wichman (2015), I note that the present analysis uses a different outcome variable, an additional year of data than what was used previously, five additional billing districts have been transitioned to monthly billing within this sample, and there are changes in the length of time post-transition for each billing district. As such, there are numerous ways in which the treatment effect may have been altered. Regardless, similar results are produced and they encourage the exploration of longer-run impacts.

What is notable about these estimates is that a simple change in billing frequency, a non-price attribute, can induce changes in behavior that are commensurate with the percent change in consumption due to targeted behavioral interventions in water demand (Ferraro and Price, 2013; Brent et al., 2015). Of course, the estimates presented here work in the opposite expected direction of conservation initiatives and the natural experiment here was not intended to affect consumption. What is lacking from the estimates in Table 3, however, is how these treatment effects play out over time.

To examine the dynamics of the treatment effects, I plot dynamic LATE coefficients in Figure 7. As shown, the estimated coefficient is plotted as a function of its distance in time from initial treatment. The figure can be interpreted as follows. The first point is the treatment effect in period 0—that is, the immediate effect of a household receiving its first monthly bill. The second point is the treatment effect in period 1, one month after initial treatment; and so on. Conceptually, the long-run average treatment effect from the dynamic models is a weighted average of the dynamic treatment effects presented in Figure 7.

The dynamic LATEs in Figure 7 suggest that there is an immediate increase in response to the first monthly bill, and this effect declines over the course of 12–16 months, at which point the treatment effect becomes insignificant. This result is consistent with the notion that consumers may see a small billed amount and react to that initial information treatment. But, as they become accustomed to the lower billed amounts, they adjust to the new billing regime over time.

To shed light on the mechanism driving this dynamic response to monthly billing, I repeat the exercise above for subsets of the full sample. Specifically, I plot dynamic treatment effect coefficients for different quartiles of historical consumption. In Figure 8, the smallest consumptive group is shown to exhibit the most immediate uptick in consumption, and maintains the largest treatment effects (in percentage terms) over the 30 months following initial treatment.

I perform the same exercise for household wealth, as proxied by the assessed value of the home gathered from county tax assessor data. The dynamic treatment effects are presented in Figure 9. Unlike the consumptive heterogeneity, there is very little difference among households residing in homes of different value. The dynamic treatment effects track the profile of the pooled D-LATE estimates. Recent studies (e.g., Jack and Smith (2015)) suggest that increasing the frequency of billing, via prepaid metering, for electricity may help to avoid backing up against credit constraints with lumpy billing practices. I find no evidence consistent with the hypothesis that more frequent bills, and thus smaller billed amounts, allow constrained consumers to smooth income over the year.

I also explore the size of a consumer's lot to determine whether increases in outdoor water use are driving the dynamic increase in consumption. In Figure 10, I plot dynamic treatment coefficients for different quartiles of lot sizes. Previous literature in water demand and conservation has used lot size as a proxy for irrigation intensity, with larger lots indicating a greater willingness to pay for outdoor water use. I find that the dynamic treatment effect is most pronounced initially for households on lots in the largest three quartiles of the distribution. The first quartile displays a dynamic LATE that is lower than that of the other quartiles over the first 12 months after treatment. This result is consistent with insight from Wichman et al. (2016) who show important heterogeneity for larger lots that can be targeted by outdoor watering restrictions.

Finally, I examine the effect of automatic bill payment on the primary empirical results for two reasons: (1) to explore differential response among consumers who are more likely to be inattentive, and (2) to assess the meter accuracy argument from a different angle. In

Table 4, I show that the static effect of changes in billing frequency is statistically zero. This provides an empirical test for whether customers who are inattentive to prices and water bills observe the change in billing frequency (Sexton, 2015). As this result pertains to rationally inattentive consumers (Sallee, 2014), those who enrolled in automatic bill payment made an active choice to be inattentive in this setting, whereas the other customers did not. To the extent that ABP proxies for rational inattentiveness, the divergence in behavior of these types of consumers suggests that the average non-ABP consumer is aware of the change in billing frequency and adjusts her behavior accordingly. This result instills confidence that the preferred specifications are indeed identifying a response to the change in billing frequency. The lack of a positive effect also lends credence to the notion that the preferred ATE in previous models is not an artifact of mechanical efficiency of the new meters.

Further, I explore the dynamic response to monthly billing for automatic bill payment customers in Figure 11. As shown, there is an immediate uptick in consumption following the switch to monthly billing, although it is about 1 cf/day less than that of "attentive" consumers. The dynamic ABP effect follows a similar profile to the overall effect, but it dissipates quicker and becomes largely insignificant after 4 months, with some noise in the trend. Overall, this results suggests a dampened reaction to the switch to monthly billing that is more heavily concentrated in the first few months of treatment, whereas attentive consumers adjust to the change over the course of a year. Due to selection into ABP, however, generalizing these results should be done with caution.

5 Discussion

Overall, there are several important takeaways from this research. First, nonprice treatments, such as changes in the frequency of billing information, can induce behavior similar in magnitude to social comparison treatments. The effect identified here, however, is opposite in sign. Further, the dynamics in the treatment effect are similar to that of nonprice policies—that is, information treatments alter behavior in the short run, but they are not permanent (Allcott and Rogers, 2014; Bernedo et al., 2014; Ito et al., 2015). This result suggests that increased information frequency (or, a treatment that alters the subjective perception of the total billed amount) induces behavior that is not consistent with a simple change in price. Because of this distinction, it is possible that price policies have more attractive benefits in the long run than do information treatments for water conservation. Within rigid institutions that govern dynamic water prices and the need for short-run reductions in demand during drought, however, nonprice policies can be effective demand-management tools to reduce consumption. The sign of the LATE in this paper highlights the danger of simply using more information as a conservation tool. To improve the predictive power of our economic models in intermittent choice settings, we need to account for the fact that inattentive consumers are just that: inattentive.

In this paper, I exploit a natural experiment in which residential water customers were switched from bi-monthly to monthly billing to assess the dynamic response to increases in billing frequency. I find that customers increase water consumption, but the estimated average treatment effect masks important dynamic responses observed for 12–18 months after the initial treatment. I estimate dynamic treatment effects that weaken over time and differ based on historical consumption patterns and preferences for outdoor water use. Notably, little heterogeneity is found among wealth groups in the dynamic response to more frequent billing, which suggests that more frequent billing does not aid in smoothing expenditures.

This research contributes to the general literature on consumer behavior under inattention, with implications for the role of informative interventions as policy instruments. More specifically, I highlight the importance of understanding the mechanism through which consumers assimilate and use information to make decisions in intermittent choice settings. Lastly, this research provides a cautionary tale for electric and water utility managers seeking to use more frequent information as a policy tool.

References

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

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

Arregui, FJ, E Cabrera, R Cobacho, and J Garcia-Serra, "Reducing apparent losses caused by meters inaccuracies," *Water Practice and Technology*, 2006, *1* (4), wpt2006093.

Barfuss, Steve, Michael C Johnson, and Martilyn A Neilson, "Accuracy of In-Service Water Meters at Low and High Flow Rates.," *Water Research Foundation (EPA)*, 2011, pp. 1–303.

Bernedo, María, Paul J. Ferraro, and Michael Price, "The Persistent Impacts of Norm-Based Messaging and Their Implications for Water Conservation," *Journal of Consumer Policy*, 2014, *37* (3), 437–452.

Boyle, Thomas, Damien Giurco, Pierre Mukheibir, Ariane Liu, Candice Moy, Stuart White, and Rodney Stewart, "Intelligent metering for urban water: a review," *Water*, 2013, *5* (3), 1052–1081.

Brent, Daniel A., Joseph Cook, and Skylar Olsen, "Social comparisons, household water use and participation in utility conservation programs: Evidence from three randomized trials," *Journal of the Association of Environmental and Resource Economists*, 2015.

Britton, Tracy C, Rodney A Stewart, and Kelvin R O'Halloran, "Smart metering: enabler for rapid and effective post meter leakage identification and water loss management," *Journal of Cleaner Production*, 2013, *54*, 166–176.

Criminisi, A, CM Fontanazza, G Freni, and G La Loggia, "Evaluation of the apparent losses caused by water meter under-registration in intermittent water supply," *Water Science and Technology*, 2009, *60* (9), 2373–2382.

Dobbie, Corey, Scott Durham et al., "Field Report–AMR: It's More Than Just Data Collection (PDF)," *Journal-American Water Works Association*, 2003, *95* (11), 50–53.

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

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

_ **, Takanori Ida, and Makoto Tanaka**, "The Persistence of Moral Suasion and Economic Incentives: Field Experimental Evidence from Energy Demand," Working Paper 20910, National Bureau of Economic Research January 2015.

Jack, B. Kelsey and Grant Smith, "Pay as you go: Prepaid metering and electricity expenditures in South Africa," *American Economic Review, Papers and Proceedings*, 2015, *105* (5), 237–41.

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

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

Lovely, Lori, "Measuring and managing: Conserving water through more efficient meter reading," *Water Efficiency*, March/April 2010, pp. 15–21.

Ritchie, Ed, "Call it a stampede: On the trail of the AMR/AMI revolution and maximum value data acquisition," *Water Efficiency*, November/December 2011, pp. 24–29.

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

Sexton, Steven E., "Automatic bill payment and salience effects: Evidence from electricity consumption," *The Review of Economics and Statistics*, 2015, *97* (2), 229–241.

Shoemaker, Randy, "Metering and more: How Leesburg, VA, slashed water losses," *Water Efficiency*, July/August 2009, pp. 41–47.

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

_ , "Information provision and consumer behavior: A natural experiment in billing frequency," 2015. Resources for the Future Discussion Paper 15-35.

_ **, Laura O. Taylor, and Roger H. von Haefen**, "Conservation policies: Who responds to prices and who responds to prescription?," *Journal of Environmental Economics and Management*, September 2016, *79*, 114–134.

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[a]	2014[b]	**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)
Number of bathrooms	1.77	2.05	1.94	1.74	1.87
	(0.61)	(0.65)	(0.57)	(0.70)	(0.65)
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 bimonthly 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 bimonthly	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 bimonthly	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
Number of billing districts:	5	4	3	5	17
	12/1/11	1/29/13	1/30/14	6/23/14	
	7/13/12	2/12/13	4/18/14	7/23/14	
Date of first monthly bill:	10/25/12	3/30/13	5/15/14	8/5/14	
	11/14/12	11/22/13		9/4/14	
	12/29/12			9/5/14	

Note: Means and standard deviations (in parentheses) are presented. The first billing district to transition to monthly billing occurred on December 1, 2011, so this district is grouped jointly with districts that transitioned in 2012. The 2010 Census (SF1) data is assigned to the Census block in which the household resides. 2009-2010 bimonthly 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).
2014[a] and 2014[b] refer to the first and second wave of billing transitions in 2014.

Table 2: Baseline OLS model results

	(1)	(2)	(3)	(4)
		<2000 ft	<1000 ft	<500 ft
	Full	from	from	from
	Sample	discontinuity	discontinuity	discontinuity
BF	2.341***	2.203***	2.128***	2.096***
	(0.371)	(0.364)	(0.329)	(0.277)
totassessvalue	0.013	0.010	0.003	0.009
	(0.009)	(0.010)	(0.007)	(0.011)
mapacres	0.097	0.163*	0.221	0.085
	(0.067)	(0.078)	(0.169)	(0.186)
sqft	0.090***	0.088***	0.101***	0.102***
	(0.016)	(0.017)	(0.012)	(0.015)
pctrent	0.198	0.202	0.057	-0.185
	(0.203)	(0.245)	(0.335)	(0.421)
pctwhite	-0.213	-0.153	-0.147	-0.391
	(0.182)	(0.178)	(0.203)	(0.361)
houseage	-0.009***	-0.011***	-0.010***	-0.008**
	(0.003)	(0.003)	(0.003)	(0.003)
hhsize	1.581***	1.521***	1.487***	1.368***
	(0.142)	(0.148)	(0.154)	(0.185)
maxtemp	0.012	0.015	0.019*	0.020*
	(0.009)	(0.009)	(0.010)	(0.010)
rain	-0.005	-0.002	-0.007	-0.013
	(0.020)	(0.019)	(0.019)	(0.019)
Observations	2,072,422	1,633,025	998,916	537,788
R-squared	0.094	0.093	0.089	0.085

Notes: All models include household and month-of-sample fixed effects. Weather covariates included. Robust standard errors clustered at the billing district level in parentheses. Dependent variable is daily water consumption in cubic feet (cf). Constant term omitted. *** p<0.01, ** p<0.05, * p<0.1.

Table 3: Baseline panel model results

	(1) Full Sample	(2) <2000 ft from discontinuity	(3) <1000 ft from discontinuity	(4) <500 ft from discontinuity
BF	0.685**	0.581**	0.555**	0.541**
	(0.239)	(0.237)	(0.223)	(0.212)
Observations	2,086,488	1,642,061	1,003,800	539,601
R-squared	0.030	0.030	0.029	0.028
Number of households	58,780	46,496	28,535	15,412

Notes: All models include household and month-of-sample fixed effects. Weather covariates included. Robust standard errors clustered at the billing district level in parentheses. Dependent variable is daily water consumption in cubic feet (cf). Constant term omitted. *** p<0.01, ** p<0.05, * p<0.1.

Table 4: Panel results for subset of automatic bill payment customers

	(1) <500 ft from discontinuity
BF	0.083
	(0.215)
Observations	26,604
R-squared	0.020
Number of households	677

Notes: Model includes household and month-of-sample fixed effects. Weather covariates included. Robust standard errors clustered at the billing district level in parentheses. Dependent variable is daily water consumption in cubic feet (cf). Constant term omitted. *** p<0.01, ** p<0.05, * p<0.1.

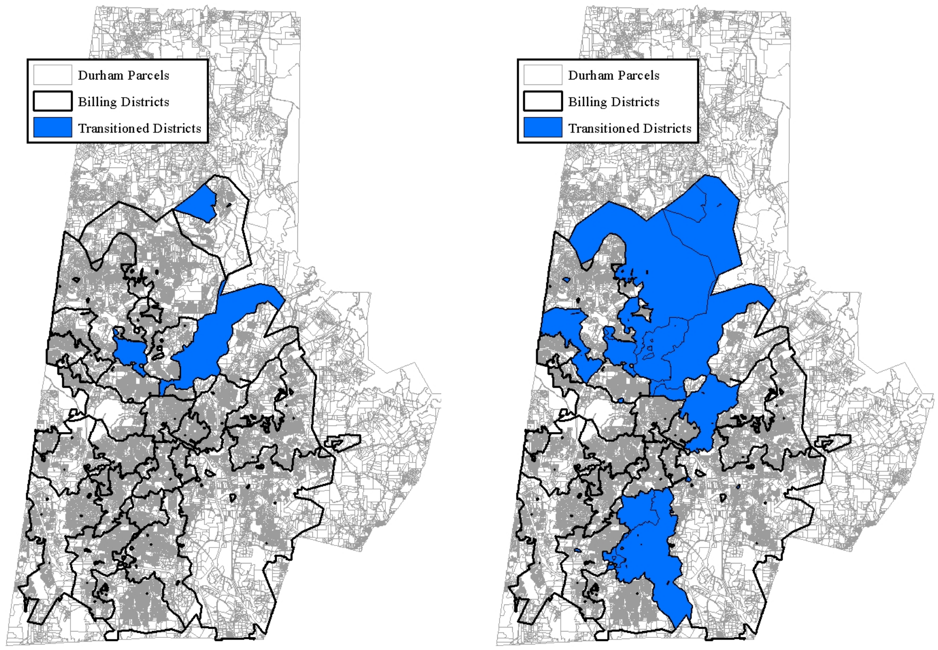

(a) Monthly billing by end of 2011 **(b)** Monthly billing by end of 2012

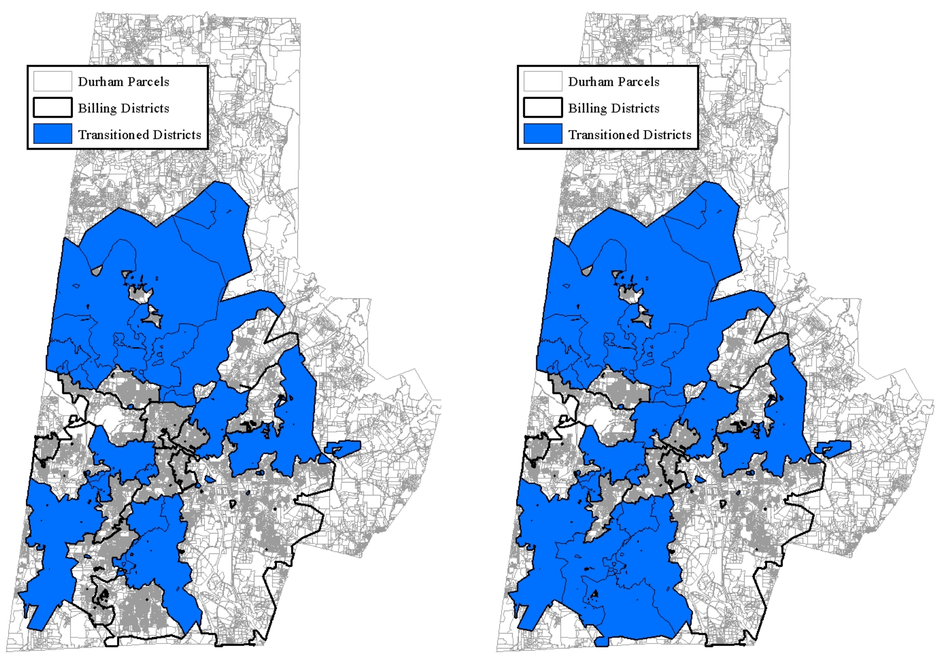

(c) Monthly billing by end of 2013 **(d)** Monthly billing by end of sample

Figure 1: Billing districts transitioned to monthly billing over time.

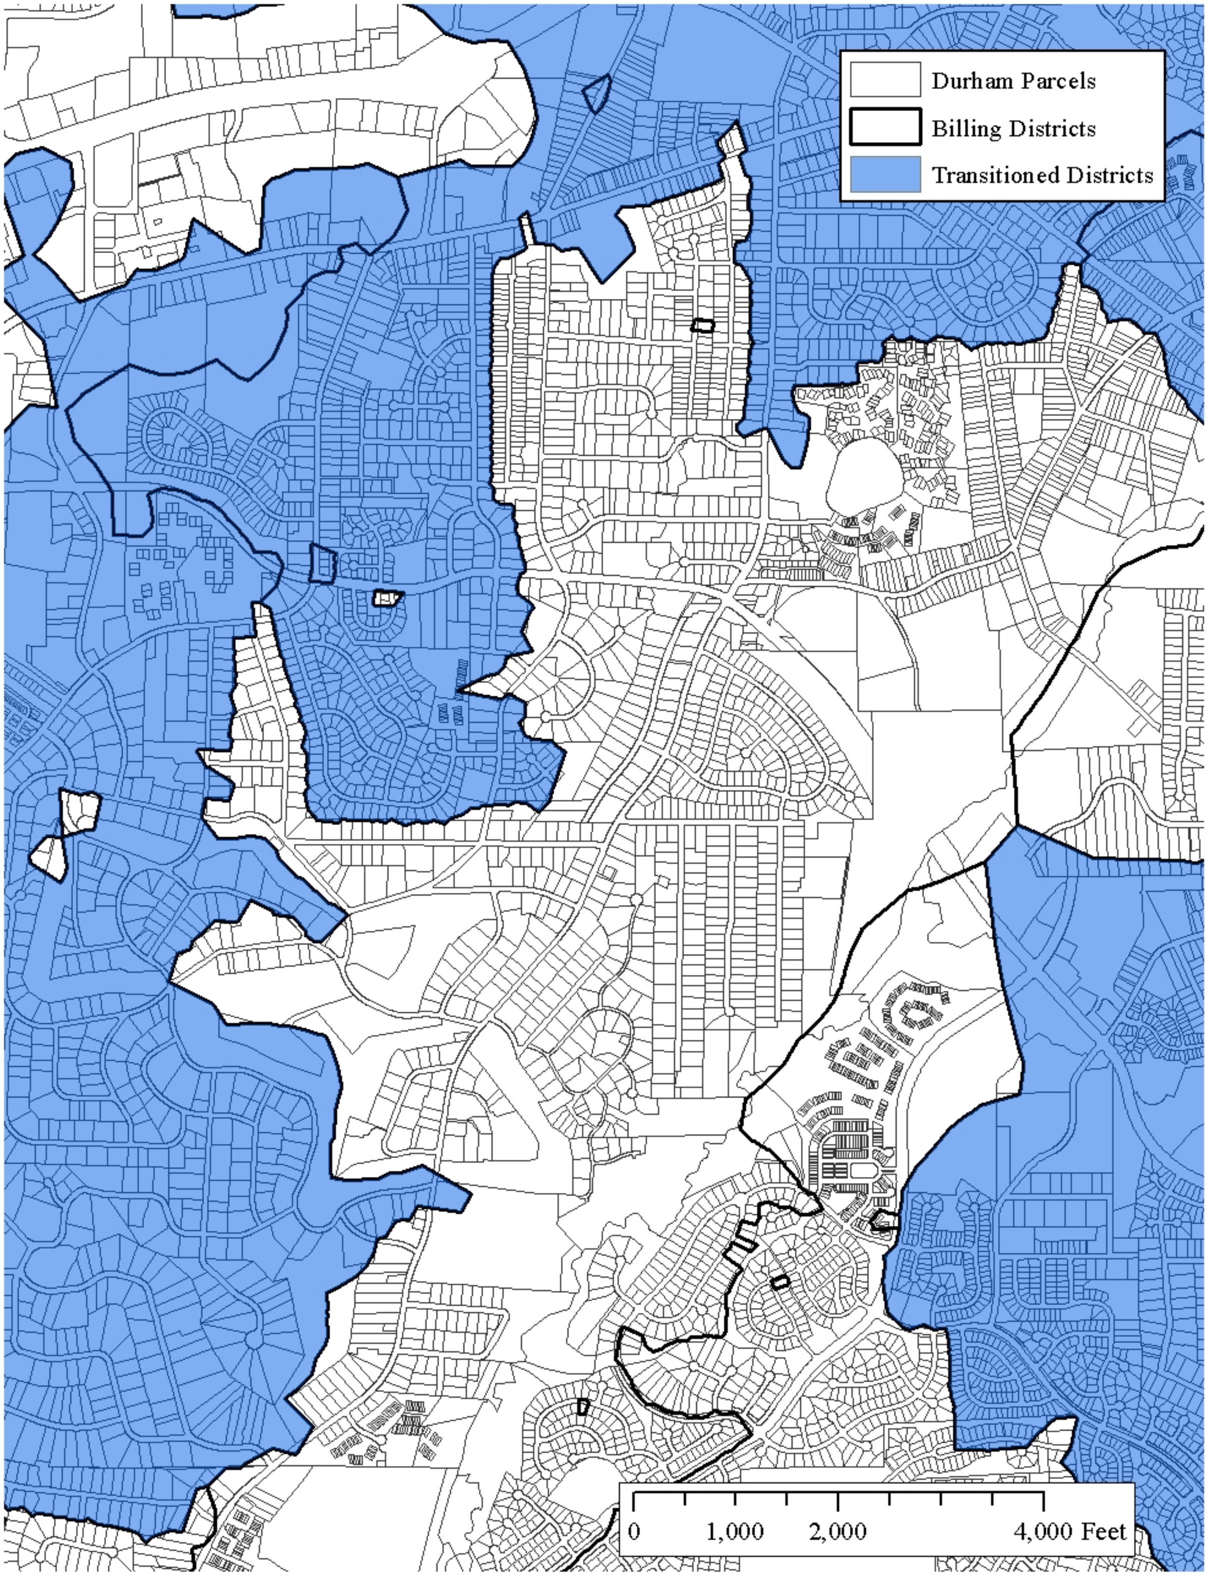

Figure 2: Snapshot of billing group boundaries within neighborhoods.

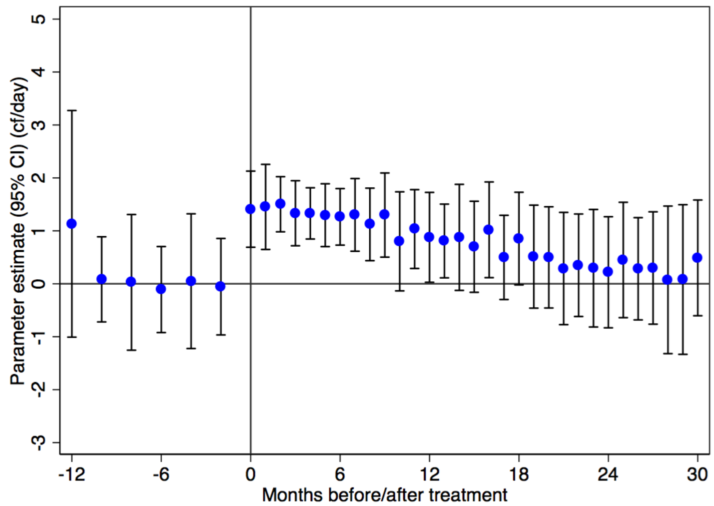

Figure 3: Event study for changes in billing frequency. Dots represent coefficient estimates as a function of the distance in time from treatment. Error bars represent 95% confidence intervals.

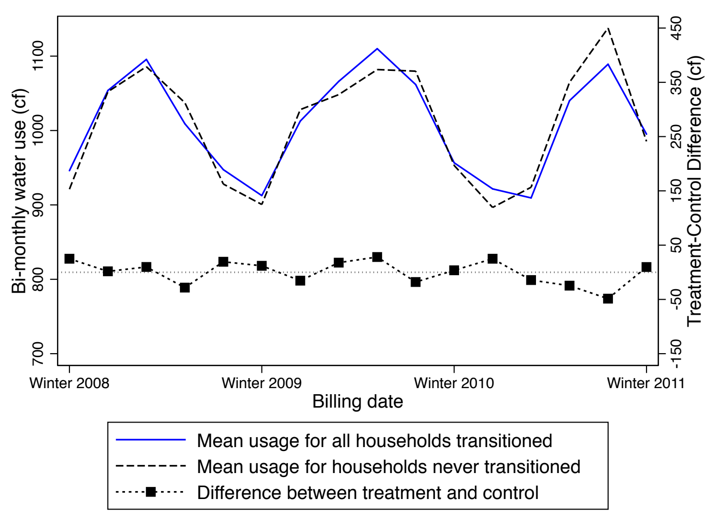

Figure 4: Mean bimonthly consumption over time for households that transitioned to monthly billing and households that never transitioned to monthly billing.

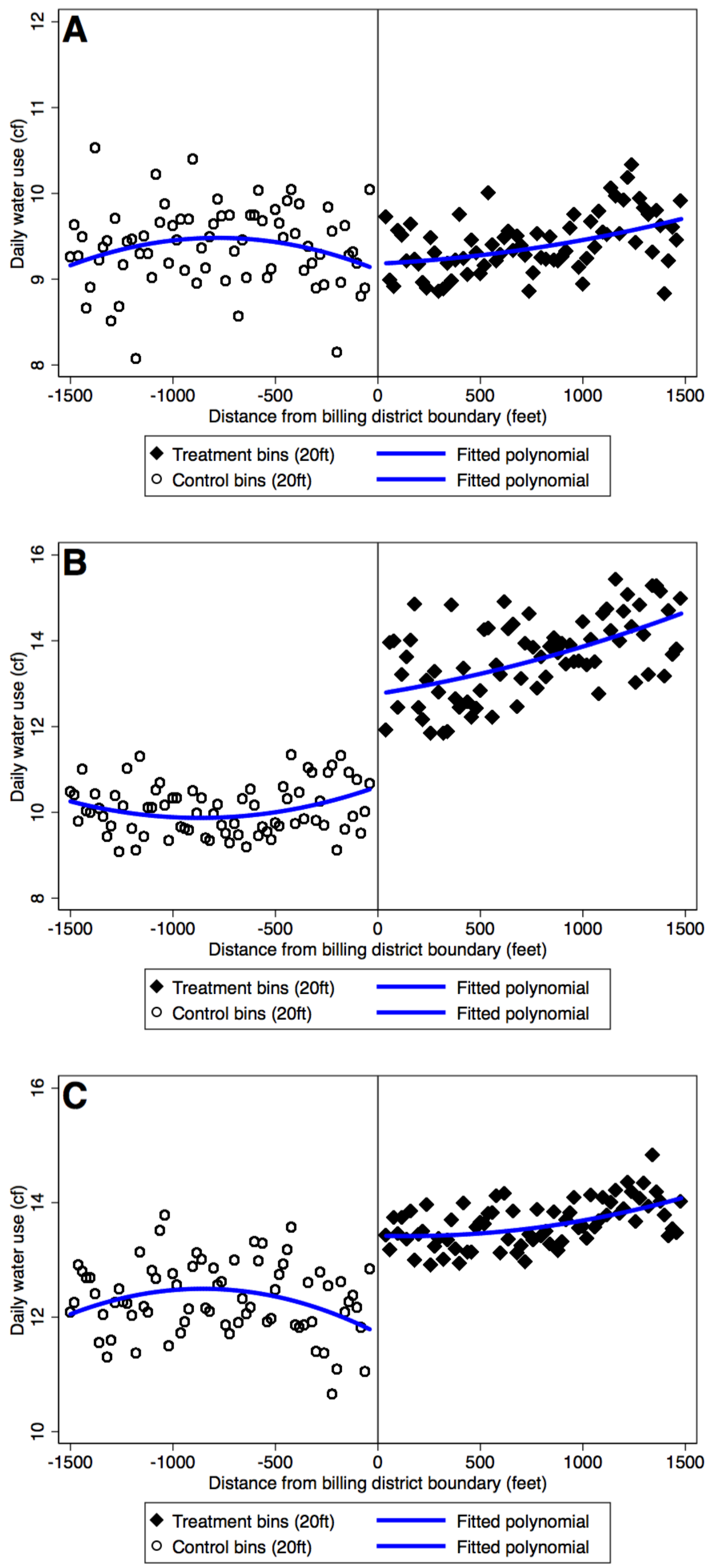

Figure 5: Mean consumption as a function of distance from the billing district boundary for the year before transition to monthly billing, the year after the transiton, and the 2nd and 3rd years after the transition

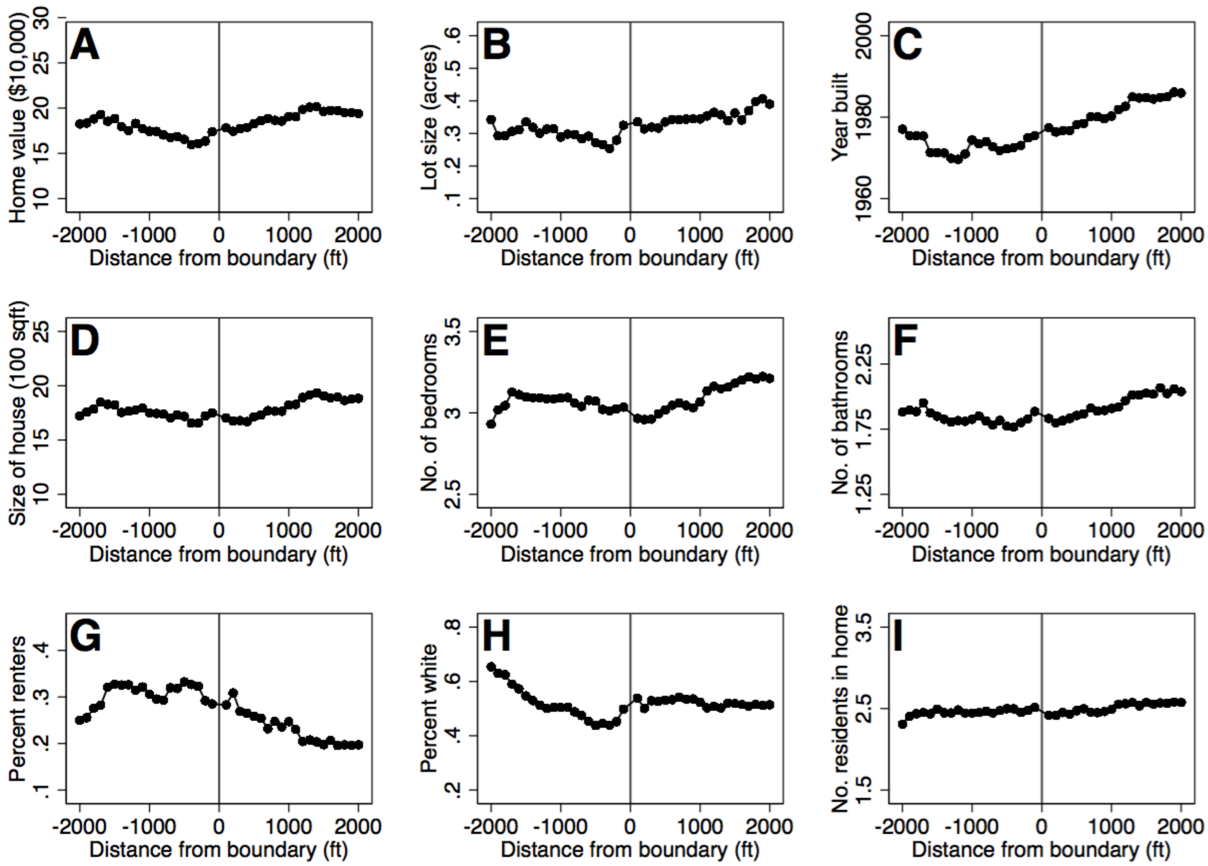

Figure 6: Mean structural and socioeconomic characteristics as a function of distance from the billing district boundary

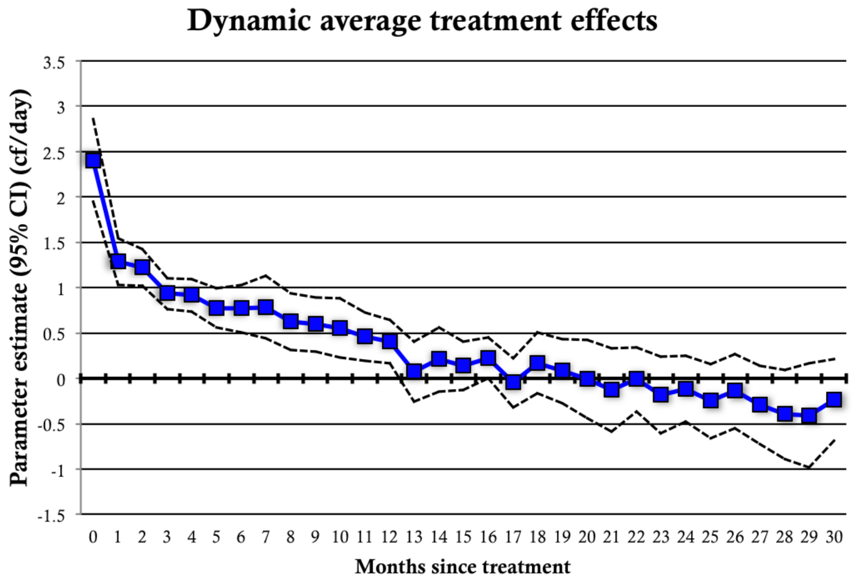

Figure 7: Dynamic average treatment effects for changes in billing frequency. The trend shows how the average treatment effect evolves over time.

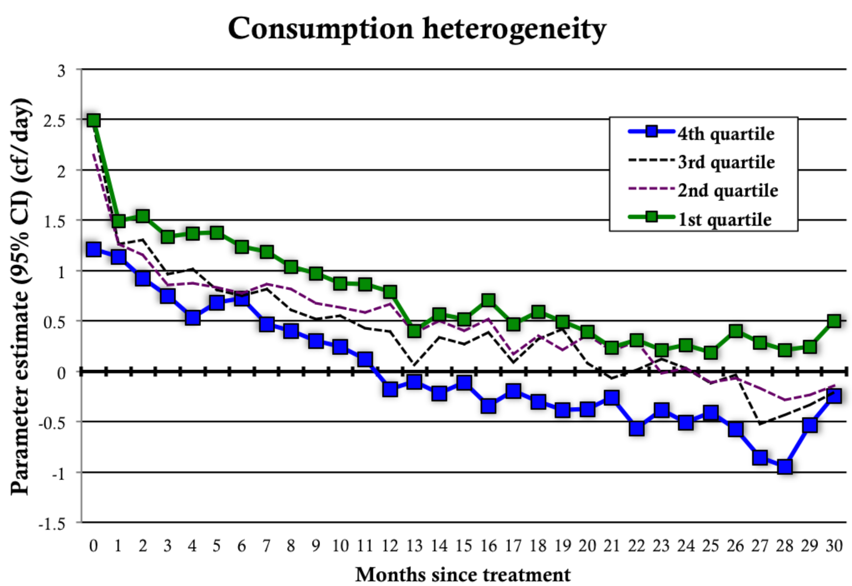

Figure 8: Dynamic average treatment effects for changes in billing frequency conditional on consumption heterogeneity. The trend shows how the average treatment effect evolves over time for each quartile.

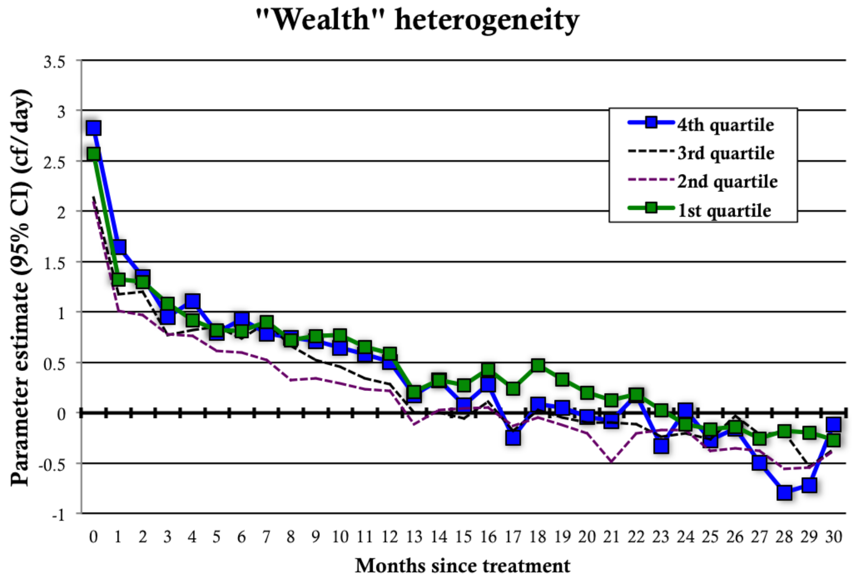

Figure 9: Dynamic average treatment effects for changes in billing frequency conditional on "wealth" heterogeneity. The trend shows how the average treatment effect evolves over time for each quartile.

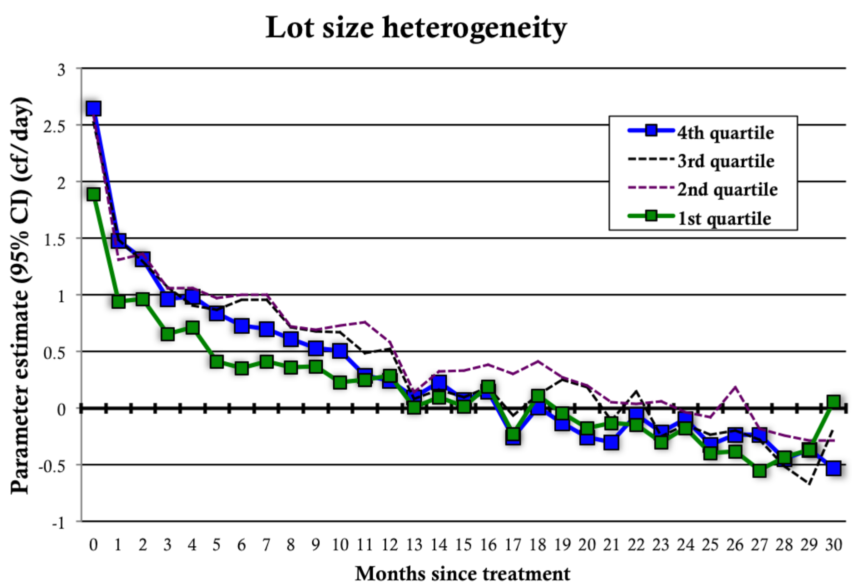

Figure 10: Dynamic average treatment effects for changes in billing frequency conditional on acreage heterogeneity. The trend shows how the average treatment effect evolves over time for each quartile.

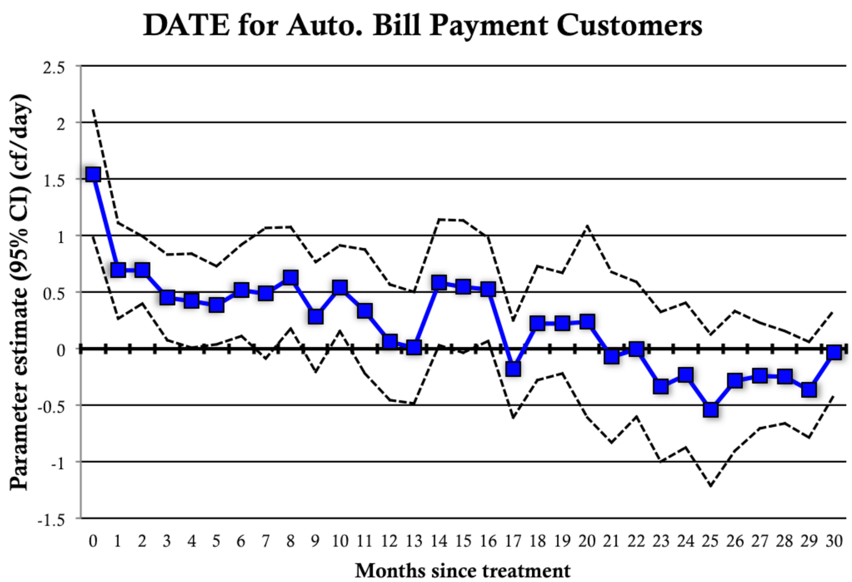

Figure 11: Dynamic average treatment effects for changes in billing frequency for automatic bill payment customers. The trend shows how the average treatment effect evolves over time.

A Appendix A

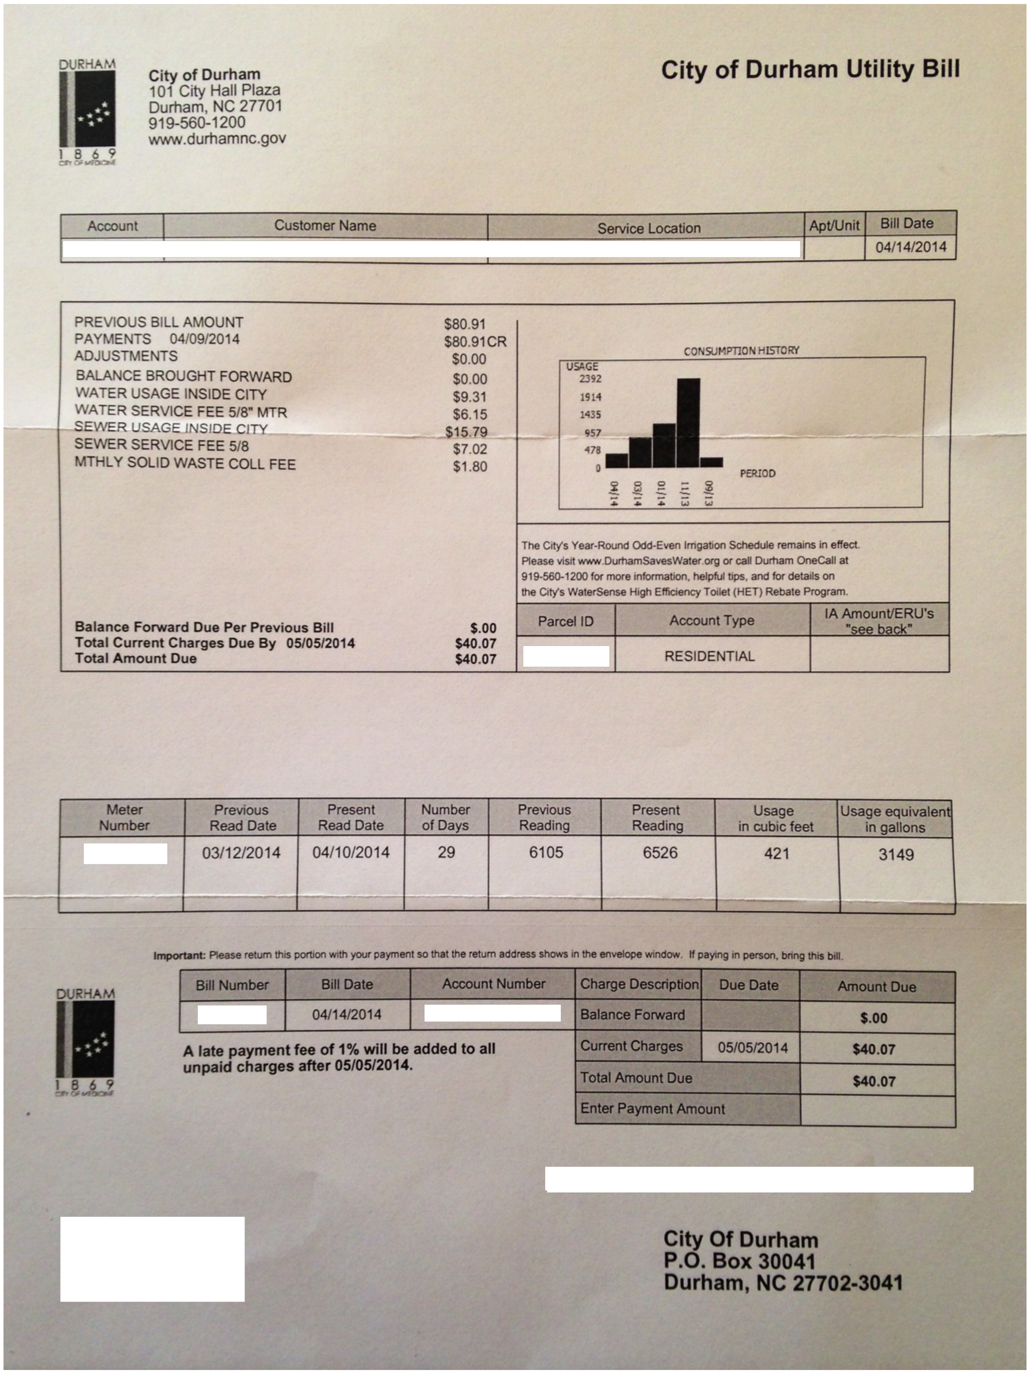

City of Durham
101 City Hall Plaza
Durham, NC 27701
919-560-1200
www.durhamnc.gov

City of Durham Utility Bill

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

PREVIOUS BILL AMOUNT	$80.91
PAYMENTS 04/09/2014	$80.91CR
ADJUSTMENTS	$0.00
BALANCE BROUGHT FORWARD	$0.00
WATER USAGE INSIDE CITY	$9.31
WATER SERVICE FEE 5/8" MTR	$6.15
SEWER USAGE INSIDE CITY	$15.79
SEWER SERVICE FEE 5/8	$7.02
MTHLY SOLID WASTE COLL FEE	$1.80

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

Balance Forward Due Per Previous Bill	**$.00**
Total Current Charges Due By 05/05/2014	**$40.07**
Total Amount Due	**$40.07**

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

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

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

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

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

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

[redacted]

[redacted]

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

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

DURHAM

1869
CITY OF MEDICINE

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

September 4, 2012

Customer Name
Customer Address
Customer City, State Zip

Service Address:

Dear Valued Customer:

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

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

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

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

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

Sincerely,

Department of Water Management

City of Durham

Good Things Are Happening In Durham

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

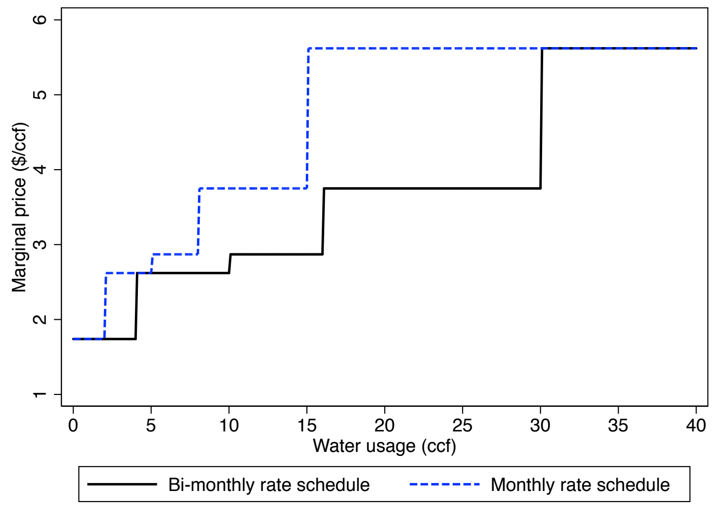

Figure A.3: Increasing block rate structure before and after transition for monthly and bi-monthly billing. As shown, quantity blocks are halved for monthly billing relative to bi-monthly billing while marginal prices remain constant within the rate structure.

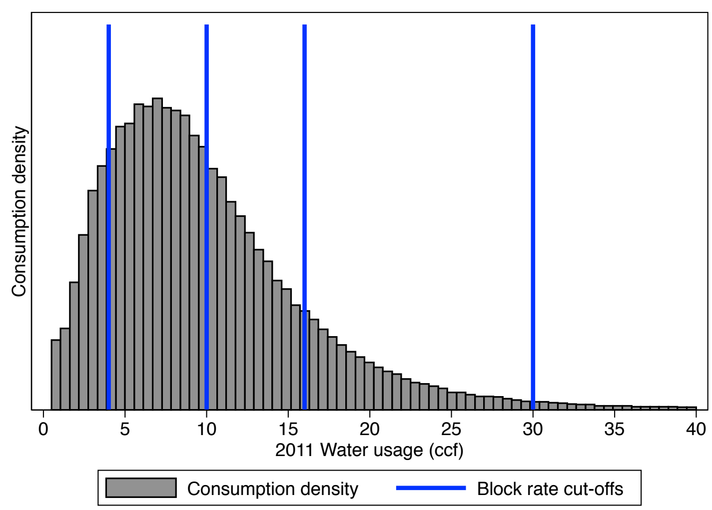

Figure A.4: Empirical density of bi-monthly water consumption with block rate cut-offs.