

NBER WORKING PAPER SERIES

DYNAMIC SALIENCE WITH INTERMITTENT BILLING:  
EVIDENCE FROM SMART ELECTRICITY METERS

Ben Gilbert  
Joshua S. Graff Zivin

Working Paper 19510  
<http://www.nber.org/papers/w19510>

NATIONAL BUREAU OF ECONOMIC RESEARCH  
1050 Massachusetts Avenue  
Cambridge, MA 02138  
October 2013

We are grateful to SDGE for providing us access to their data. For many helpful comments we thank Ed Barbier, Richard Carson, Ted Groves, Mark Jacobsen, Katrina Jessoe, Jacob LaRiviere, Kevin Novan, Sherrill Shaffer, and numerous seminar participants. The usual disclaimer applies. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w19510.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2013 by Ben Gilbert and Joshua S. Graff Zivin. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Dynamic Saliency with Intermittent Billing: Evidence from Smart Electricity Meters  
Ben Gilbert and Joshua S. Graff Zivin  
NBER Working Paper No. 19510  
October 2013  
JEL No. D03,Q4

**ABSTRACT**

Digital tracking and the proliferation of automated payments have made intermittent billing more commonplace, and the frequency at which consumers receive price, quantity, or total expenditure signals may distort their choices. This category of goods has expanded from household utilities, toll road access and software downloads to standard consumption goods paid by credit card or other "bill-me-later"-type systems. Yet we know surprisingly little about how these payment patterns affect decisions. This paper exploits hourly household electricity consumption data collected by "smart" electricity meters to examine dynamic consumer behavior under intermittent expenditure signals. Households reduce consumption by 0.6% to 1% following receipt of an electricity bill, but the response varies considerably by household type and season. Our results also suggest that spending "reminders" can reduce peak demand, particularly during summer months. We discuss the implications for energy policy when intermittent billing combined with inattention induces consumption cycles.

Ben Gilbert  
University of Wyoming  
1000 E University Ave,  
Laramie, Wyoming 82071  
bgilbe10@uwyo.edu

Joshua S. Graff Zivin  
University of California, San Diego  
9500 Gilman Drive, MC 0519  
La Jolla, CA 92093-0519  
and NBER  
jgraffzivin@ucsd.edu

## 1. Introduction

The neoclassical agent takes account of all present and future expected costs and benefits when making decisions. Yet in many settings consuming a good and spending money on it are separate experiences, and the true costs of one's actions are only revealed ex post and intermittently - often in the form of bills. This expenditure may not be salient to the agent if the market environment makes prices or quantities opaque when consumption is happening, particularly if the agent is billed later for decisions that are made over time. The problem of salience is not simply one of uncertainty; forming rational expectations with uncertainty requires cognitive processing of the distribution of outcomes, while insalient costs facilitate inattention and incomplete cognitive processing (DellaVigna (2009)). An agent may overconsume if he becomes inattentive to his spending when price or quantity information is obscured.

Overconsumption has been demonstrated with static models in the context of toll road use (Finkelstein (2007)), alcohol purchases (Chetty et al. (2007)), bank overdrafts (Stango and Zinman (2011)), and monthly electricity consumption (Sexton (2011)). Yet the problem is a fundamentally dynamic one if actions taken when expenditure is not salient are hysteretic or deplete capital stocks, and if this type of spending is repeated. True expenditure is revealed eventually; how does realizing that one has recently been inattentive and overconsumed - and may soon repeat this behavior - affect present and future behavior? For how long after an expenditure signal (a bill) is received do agents remain attentive? These problems are very different from those faced by an impatient or even present-biased agent who simply places less weight on the future but still acts with full awareness of expected costs.

The number of market environments in which signals about accumulated expenditure are opaque or intermittent is large and growing. Health patients with chronic illnesses make sequential decisions about medical treatments at different points in time than they are billed for those decisions. Accounts allowing downloads of media, games, and software apps to smart phones and tablet PCs are typically linked to credit cards that are billed monthly. Many subscription services and household utilities are moving or have already moved to automatic bank draft or credit card billing systems. Nonmarket decisions are also affected by the relative salience of the true costs of an action. For example, the health costs of over or under consumption of certain types of goods like fatty or sugary foods, harmful drugs, and chronic medicines, are not apparent at the time of consumption.

Problems with salience are the rule rather than the exception in a world in which people pay for a larger share of even basic consumption with credit cards, many additional types

of consumption are billed weeks after consumption occurs, and many types of bill payments are automated. We should not be surprised by a consumer who is shocked at his own consumption or unhappy with the decisions that his past self has imposed on the wellbeing of his current self because a lack of salience can facilitate intertemporal moral hazard<sup>1</sup>. This poses costs and benefits for sellers, however, because opaque prices can increase sales but also demand volatility when customers vacillate between overconsuming while inattentive and underconsuming to compensate for past or future actions.

Dynamic analyses of salience are not as prevalent in the literature as static ones. For example, Finkelstein (2007) shows that toll-road drivers who use EZ-Pass systems drive on toll-roads more frequently and protest fee increases less frequently. Presumably these drivers receive monthly bills, however. Does toll-road use decline after the driver is reminded of her expenditure? This is the question we address in this paper<sup>2</sup>.

Using the motivating example of residential electricity consumption, we answer this question with evidence from a unique panel data set of hourly household-level electricity use for several thousand households in the San Diego suburbs. If electricity expenditures are only salient to consumers when they receive a bill, then consumers should adjust behavior at the point of salience (i.e., the beginning of the billing cycle) but then return to overconsumption as their attention to the price fades. We test this prediction by studying changes in household consumption behavior throughout their monthly billing cycle. Such “recency effects” of intermittently salient information, or effects that decline as time passes, have been found with late payment fees for credit cards in Agarwal et al. (2013) and personalized home energy reports with normative and educational messages in Allcott and Rogers (2012)<sup>3</sup>. In

---

<sup>1</sup>This is in contrast to rational addiction models (e.g. Becker and Murphy (1988)) in which the agents act with full awareness of the intertemporal externality but engage in bad behavior anyway.

<sup>2</sup>Karlan et al. (2010) present a dynamic model with inattention to uncertain future costs (e.g., car repairs or durable goods replacement) rather than periodic inattention to current spending, and Grubb and Osborne (2012) have a dynamic component to learning about preferences in the context of insalient cell phone usage costs. The focus of these models is related to but distinct from what we address in this paper. The rational inattention literature (e.g., Sims (2003)), which has primarily addressed problems in macroeconomics and finance, takes the point of view that agents face constraints in processing multiple flows of uncertain information, but information is nevertheless freely available; agents therefore exhibit inertia and respond sluggishly to aggregate updates but discretely to own shocks. This might be an appropriate explanation for the inertia and variability in electricity consumption but this is beyond the scope of this paper. Within the context of a consumption environment with inertia, we deal with agents periodically exposed to information that is only freely available (or perfectly salient) in discrete intervals, between which agents forget the information.

<sup>3</sup>Allcott and Rogers (2012) also find that reversion to past behavior is reduced when intermittent messages are maintained over a long period of time.

both cases, however, the information treatments are “unusual” in the sense that they are designed to draw the agent’s attention and alter behavior. We show that the same pattern of behavior is present with routine billing.

Our results suggest that households reduce consumption by 0.6% to 1% in the first week of a new electricity billing cycle. These impacts are more pronounced in the summer, particularly if the weather is hot when the bill arrives, and among larger households with more sensitivity to hourly weather fluctuations (which we take as an indication of air-conditioning use). Within the day, reductions persist through the peak hours of afternoon and early evening, providing the first evidence to our knowledge that a price reminder (a bill) can induce peak load reductions in addition to general conservation. We also examine differential responsiveness to billing information across households based on several novel measures of household composition and capital stock.

These results are consistent with several possible theories of salience and inattention. Households could be rationally inattentive and simply forgo the costs of calculating the cost of energy at every instant within the month<sup>4</sup>. On the other hand, limited cognitive processing of opaque price information may systematically bias down the effective price used in decision making, in which case households would consume as if the good is cheap when the price is not salient, and consume the optimal bundle when prices become salient again. This response could be exacerbated if self-aware households expect to become inattentive again in the future and preempt their future selves by underconsuming now<sup>5</sup>.

The rest of the paper proceeds as follows. Section 2 discusses the context of household electricity use and “Smart Grid” infrastructure investment and associated policy implications. Section 3 presents the empirical approach and section 4 discusses the data used in our empirical analysis. Section 5 discusses the results and section 6 concludes.

---

<sup>4</sup>However, it seems under this explanation rationally inattentive consumers should still form rational expectations about prices they haven’t bothered to calculate, in which case consumers should guess their effective prices correctly on average and not systematically overconsume in periods when price information is unavailable.

<sup>5</sup>The repeated nature of the problem could present the opportunity for learning. Although we do not address learning in this paper, it is more likely that an agent who learns about his tendency to be inattentive puts in place better safeguards against eventual bad behavior. This could take the form of reducing the persistence of the expenditures (e.g., changing electrical appliance settings while attentive), or putting himself on an alternative consumption path that takes account of future inattention. Both of these explanations would still lead to consumption cycles that follow price information. Regardless, our sample contains only consumers that have lived in the same residence for at least three years, so any potential learning would have likely already occurred before our study period.

## 2. Background on Smart Electricity Meters

Energy efficiency has been controversially touted as a large untapped source of low-cost pollution abatement, particularly for carbon dioxide. McKinsey Global Institute has argued that cost effective investments in existing technology could reduce energy demand by more than 20% (McKinsey (2007)). Adoption rates remain stagnant for a number of reasons (Jaffe and Stavins (1994)) and more recently attention has turned to behavioral “nudges” for households to alter electricity use habits (Allcott and Mullainathan (2010)). These include goal setting, social norms, and personalized feedback about prices or quantities (e.g., Hsiaw and Harding (2014), Allcott and Rogers (2012), Allcott (2011a)). In a number of recent trials in different settings, savings of between 2% and 5% on average have been achievable with various types of behavioral “nudges”.

Electric utilities are still learning about how to alter customer behavior using “smart grid” infrastructure investments that retrofit the electricity grid with monitoring and communication technology. These tools may improve salience as well as enable more creative energy price structures and signalling schemes. In a few limited cases smart meters have been enabled to communicate real-time information between customers and the utility, such as price signals and other alerts. Eventually these meters may be able to send additional information designed to shape behavior such as appliance-level disaggregated usage, security breaches, or pollution levels. A large piece of the expected social return on these investments is the potential for electricity conservation and reduced emissions.

One frequently discussed cause of inefficient electricity use is the lack of transparent information on prices or quantities available to households. The salience problem in household electricity spending has several important features. First, the price structure itself is complicated and can confuse customers. Most residences face a nonlinear “Increasing Block Tier” (IBT) rate structure in which the marginal price the customer pays for a unit of electricity is raised when the customer consumes above a given threshold within a given month. There may be several usage thresholds, or tiers, which may be unevenly spaced and induce marginal price increases that are nonlinear. A number of studies on this type of rate structure have shown that consumers do not respond to the marginal rate of the tier they are on in any given month, and generally do not fully understand the pricing structure and its implications for minimizing their monthly bill (Kahn and Wolak (2013), Ito (2012), Bushnell and Mansur (2005)). The second source of salience problems with household electricity is that consumption decisions are made in real-time whereas expenditure is only “experienced” monthly when the bill arrives. The opportunity for a person to investigate which pricing

tier they are on occurs only once a month, and ex post after thresholds have been exceeded. The intermittency of quantity and price information makes it nearly impossible to respond to marginal prices at the point of consumption.

Lastly, consumption decisions are disaggregated across appliances without price information for individual uses. Even if the marginal price of a unit of electricity were known with certainty, translating this price into the price of drying laundry or powering a flat screen television for an hour is not a trivial task. This problem occurs partly because of the intermittency in quantity information and partly because of the disaggregated and heterogeneous nature of end uses. We address the dynamic problem in this paper but we acknowledge that solving the disaggregation problem could interact with our results in unexpected ways. For example, if monthly bills were displayed by appliance we might estimate larger effects as households adjust settings on their biggest sources of energy use. After a few monthly cycles of adjustment, however, the reaction to a bill might diminish if households are satisfied with their allocation across appliances. Jessoe and Rapson (2013), for example, show that real-time consumption information provided by in-home displays causes persistent reductions in energy use, which they attribute to households learning about disaggregated uses.

One argument for the smart grid has been that more salient prices and quantities would help consumers make privately optimal electricity consumption choices, which would make market-based policy instruments more efficient. Many authors have studied the effect of improving electricity cost information on subsequent consumption. These studies fall into two main categories. First, the use of “peak prices” in which customers are informed that their prices will spike during certain hours on certain days in order for the utility to shift demand away from peak hours when electricity grids are constrained. Peak-hour or real-time price elasticities have been estimated in a number of field experiments including Jessoe and Rapson (2013), Allcott (2011b), Wolak (2011) and Wolak (2007).

The second category covers the use of behavioral “nudges” or other non-price information treatments designed to induce households to conserve energy, or reduce average consumption for energy efficiency or environmental reasons. There have been a number of studies evaluating the provision of real-time quantity information to households, and most of these focus on reporting the number of kilowatt-hours used or saved rather than making prices or total expenditures visible. Overall savings of 2% to 5% have been reported with these programs (Darby (2006)) but many of the studies involve short-term trials with self-selected participants. Rebound and long-run effects are also of interest. Allcott and Rogers (2012) find that the effects of monthly or weekly information erode over time if the provision of in-

formation ends, providing evidence for a dynamic component of intermittent salience. While Allcott and Rogers (2012) study information nudges specially designed to induce behavioral change and conservation, such as normative messages, we study routine billing and show that simply receiving a bill acts like a nudge. An exception that spans both categories of information treatments is Jessoe and Rapson (2013), who show that households with real-time quantity information respond more to peak price signals than households facing peak prices without the quantity information. This confirms that salience is important for both prices and quantities<sup>6</sup>.

These two categories also differ in the data used to estimate the treatment effects. Hourly data has been used to estimate the demand shifts from peak price instruments, but to our knowledge has not been used to examine the time profile of non-price nudges. Yet the question of what hours of the day that information nudges take hold is important for both private market and public policy efficiency. Utilities want to shift demand away from peak hours in order to avoid the costs of building and maintaining peak power generating plants that are used only a few hours each year to avoid blackouts. In addition, the mix of fuels used to generate electricity varies throughout the day, so the pollution emissions avoided because of responses to nudges will depend on when households actually respond (Graff Zivin et al. (2014)). Our results suggest that nudges can be effective during peak hours, particularly among customers with the greatest peak demand.

The results in this paper raise three important issues for utility smart grid policy and information feedback. First, attention and consumption may have stable steady states that are difficult to alter with only periodic information; our results and those of Allcott and Rogers (2012) and Jessoe and Rapson (2013) suggest that some “nudges” are transitory while others can shift the steady state. Second, information updates will be more or less effective depending on the customer’s consumption habits and appliance stock. Third, improving the salience of prices may help utilities manage peak demand even absent an increase in peak prices.

### 3. Empirical Approach

If the price of electricity is more salient when the bill arrives then we should observe lower consumption following a new bill than during the rest of the month. Let  $y_{it}$  be daily electricity

---

<sup>6</sup>Jin and Leslie (2003) show that quality salience also affects demand in the context of restaurant hygiene ratings.



consumption from a sample of individual households we will use to test for reductions in use during the first week of each billing cycle relative to the rest of the month. We first examine whether simply reminding the household about this cost - the arrival of the bill itself - reduces consumption. We then examine the source of these reductions and how these responses vary across customer segments. In order to estimate the consumption response to the arrival of a bill, we estimate several dynamic panel specifications of the form:

$$\begin{aligned}
y_{it} = & \gamma y_{i,t-1} + \alpha_d D_{it} + CDH_t \cdot \mathbf{Z}_i' \beta_c + HDH_t \cdot \mathbf{Z}_i' \beta_h \\
& + \beta_{c2} CDH_t^2 + \beta_{h2} HDH_t^2 + \beta_{ch} CDH_t \cdot HDH_t + u_{iwm} + \epsilon_{it}
\end{aligned} \tag{1}$$

where  $D_{it}$  is a dummy variable indicating a one-week window of time following a new billing cycle (labeled “post” in the following tables). We don’t observe exactly when the households read their bills, but given that bills are sent out (either by e-mail or U.S. mail) on the morning of the first day of the new billing cycle, households are most likely to have seen a new bill during the first week<sup>7</sup>. Some households may be affected by the price signal outside the one-week window, in which case our estimates may be conservative<sup>8</sup>.

The variables  $CDH_t$  and  $HDH_t$  are Cooling and Heating Degree Hours, respectively. Heating Degree Hours measure the amount of time and the extent to which the outdoor temperature was below 65 degrees Fahrenheit each day, indicating the conditions in which household heating is required to maintain indoor comfort. The differences between 65 degrees and the actual temperature (if temperature is below 65) at each hour of the day are summed over each 24 hour period. Likewise, Cooling Degree Hours are a measurement for exogenous air-conditioning needs with temperatures above 65. The  $u_{iwm}$  term captures household-time fixed-effects. We report models with household and time fixed effects estimated separately as well as household-by-weekday effects, household-by-calendar month effects, and household-month-weekday effects to more flexibly capture household-specific weekly seasonality that may change throughout the year.  $\mathbf{Z}_i$  is a vector of household-specific covariates constructed from the data as described in section 4. We allow a household-specific relationship with outdoor temperature by interacting Cooling and Heating Degree Hours with  $\mathbf{Z}_i$ .

---

<sup>7</sup>Unfortunately our data does not identify whether the bill was sent by email or U.S. post, or whether the customer paid using automatic bill pay or another method.

<sup>8</sup>We check the robustness to different windows of time in section 5.

## 4. Data

Our data was generously provided by San Diego Gas & Electric (SDG&E) and consists of historical consumption and billing information for 30,000 households in Escondido, California, a residential community in northeastern San Diego County. Starting in April, 2009, SDG&E began installing digital smart meters on Escondido residences. These meters gather and store hourly observations on electricity consumption. During our sample period (April 2009 to January 2010), the meters were not enabled to communicate with any devices in the household, so the meters allow the researcher to observe consumption at high frequency while nothing changed for the household or its occupants. We aggregate the hourly data to daily electricity use as the dependent variable in our main analysis, but also present results by hour of day.

We combined this hourly data set with billing information for the same households. The billing information consists of the monthly bills during the sample period as well as three years prior to the installation of the meters, including the dollar amount of the bill, the quantity consumed during the month, and the dates of the billing cycle. We removed any household that did not have the same account-holder for the entire three-year billing history period in order to focus on households that are likely to have stable expectations about their bills. Further, we removed any household that did not have a smart meter for more than seven months so that we could observe responses over multiple billing cycles. The resulting data set consists of 10,826 households.

The period SDG&E uses to calculate the bill for each month ends at midnight on the last day of each billing cycle, and bills are sent by email or U.S. post the following morning. Households in SDG&E have billing cycles that differ by the day of the calendar month at which they begin. Our trimmed sample from Escondido has 12 different billing cycles that start between the 30th of a given calendar month and the 10th of the subsequent calendar month, and they are between 28 and 32 days long, although in the vast majority of cases not more than 31 days.

The households in our sample are in one of three electricity rate classes: standard residential (DR), standard residential low income (DRLI), and general service (A). A detailed description of the rates is provided in Table A.1. The general service A rate applies to small commercial users that are also owner-occupied and are considered “residential” by the utility. In the Escondido community we study this describes small farms, ranches, or orchards

at which the owners also reside<sup>9</sup>. The A rate is a flat rate that varies by summer and winter. The majority of our households, however, are on DR or DRLI rates which have an increasing nonlinear block-tier structure with baseline usage and rate tiers that vary by summer and winter. Rates increase to a new tier each month once the baseline allowed quantity is consumed within that month, and again at 130% and 200% of the baseline. The DRLI rates are slightly higher than DR for baseline consumption but much lower for the higher two tiers (above 130% of baseline). Baseline allowed usage was close to 11 kWh per day during our time frame - well below the average daily usage for both rate groups<sup>10</sup>. While an understanding of the pricing context is useful, a number of studies have found that consumers do not respond to the marginal rate of the tier they are on in any given month (Kahn and Wolak (2013), Ito (2012), Bushnell and Mansur (2005)). In the case that consumers are responding to marginal rates, our empirical results will provide a lower bound estimate of the response to expenditure information since nearly all consumers will be in the lowest price tier when they receive their bill during the first week of their new cycle.

We combine hourly consumption with hourly outdoor temperature data from several weather stations in the Escondido area. We use the weather data in two ways. First, we construct the Heating and Cooling Degree Hour measures<sup>11</sup>. Second, we use the hourly variation in temperature with the hourly variation in household-specific consumption as described below.

Summary statistics for variables used in the analysis are presented in Table A.2. The means of daily and monthly usage and monthly bills (20.7 kWh, 621 kWh, and \$115, respectively) are in the range of what we would expect for single family residences, although the standard deviations and maximum values are large. This reflects the presence of a small number of large residences including small ranches, orchards, or other small agricultural uses in addition to more traditional suburban homes and apartments. The 99th percentiles of these variables (87 kWh, 2371 kWh, and \$592, respectively) are less extreme. The results presented in the next section are very similar in magnitude when the models are estimated after dropping the highest and lowest 1% of users<sup>12</sup>. The summary statistics for Cooling

---

<sup>9</sup>The A rate does not include multiunit housing like apartments, condominiums or townhouses.

<sup>10</sup>A time-of-use (TOU) rate was offered for A rate customers, and some larger customers were permanently switched to TOU after exceeding a usage threshold, but none of these customers are present in the sample used in this paper.

<sup>11</sup>Empirical electricity studies typically use monthly household data rather than daily, and transform raw temperature data into “Heating Degree Days” and “Cooling Degree Days”. Our measure is exactly the same at a higher frequency.

<sup>12</sup>Results from these specifications are available upon request.

and Heating Degree Hours indicate the wide temperature variability in inland San Diego county. A day with Cooling Degree Hours of 107 could indicate a short temperature spike in the afternoon, or 10 hours of the day with temperatures at 75.7 degrees Fahrenheit. Extreme days with Cooling Degree Hours approaching 500 are days in which temperatures exceeded 100 degrees in the afternoon and stayed in the 80’s through the night. Likewise, colder days with Heating Degree Days around 400 can be interpreted as days in which the temperature was in the 40s for most of the 24-hour period, or were in the upper 40’s during the day and reached the 30’s at night.

Because of privacy concerns with smart meter data, we have only a limited set of household-specific covariates and no common identifiers to match the smart meter data to external datasets. The hourly smart meter data itself provides a unique set of household-specific information, however. We take advantage of the hourly variation in energy usage to infer characteristics about each household and use these constructed variables as household covariates in  $\mathbf{Z}_i$ . These include the sensitivity of the household to hourly temperature fluctuations, the volatility of the household’s daily consumption, and the lag persistence in the individual household’s daily consumption series. These variables are meant to capture the durable goods stock such as heating and cooling as well as household energy usage-related preferences. We include all of these variables in  $\mathbf{Z}_i$  primarily to allow household-specific temperature responses, but we also examine how the salience effect varies along several of the dimensions that are measured, such as household size and air-conditioning use.

In order to construct these variables, we use the explanatory power of hourly temperatures for hourly consumption as a proxy for the household building and behavioral characteristics. We construct these by using the hourly data to estimate the following regression independently for each household:

$$\begin{aligned}
 y_{ih} = & \sum_{h=1}^{24} \phi_{ih} \Psi_h + \sum_{m=1}^{10} \beta_{im} Month + \sum_{d=1}^7 \beta_{id} Weekday + \beta_{iC} \mathbb{1}\{Temp_h > 65\} \cdot (Temp_h - 65) \\
 & + \beta_{iH} \mathbb{1}\{Temp_h < 65\} \cdot (65 - Temp_h) + \epsilon_{ih} \quad \forall i
 \end{aligned} \tag{2}$$

where  $\Psi_h$  are hourly dummies and  $\mathbb{1}$  is an indicator for whether the temperature was above or below 65 degrees. The variable “fit” is the household-specific R-squared from this regression. This measures how tightly the hourly consumption in the household is explained by hourly temperature changes and seasonal patterns at the hourly, weekly, and monthly level, which reflects heating, cooling, and insulation properties of the household.

As shown in Table A.2, seasonal and weather variations explain 41% of the variation in hourly consumption for the average household, but this statistic varies considerably across households, from almost no explanatory power (0.015) to high explanatory power (0.99). “Pattern” is a proxy for daytime occupancy and measures the percentage difference between the minimum and maximum  $\phi_{ih}$  for each household - a measure of the average variation in consumption throughout the day for each household. A mean of -0.74 for this variable indicates that for the average household, the lowest consuming hours at night are 74% below the highest consuming peak hours in the afternoon or early evening. We also use  $\beta_{iC}$  and  $\beta_{iH}$  to proxy for the air-conditioning and heating capacity of the household, respectively. We rename these “hotco” and “coldco”, respectively. A mean of 0.027 for “hotco” in Table A.2 says that the average household increases usage by 0.027 kWh in response to a 1 degree temperature increase in a given hour, or that a 10-degree increase over 10 hours would raise consumption by 2.7 kWh for the average household. The “coldco” variable has a symmetric interpretation for cold temperatures below 65 degrees. That the average of “coldco” is negative indicates that electricity is less important for heating than cooling in this region, and that often households experience cold temperatures during nighttime hours in which less electricity is being used on average regardless of weather.

Additionally, we constructed measures of household-specific interday volatility and lag-persistence in energy consumption over the study period. We ran household-specific regressions of daily total consumption on lagged consumption (and Cooling and Heating Degree Hours to capture dependence net of serial correlation in weather) and captured each household’s autoregressive coefficient in a variable called “persist”. Strong persistence might indicate a household that has a high cost of adjustment in energy consumption or one with less flexible, more habit-driven occupants. A mean of 0.27 for the “persist” variable suggests that daily consumption is persistent, but not close to a random walk on average. We also obtained household-specific volatility by calculating the daily raw differences in consumption, and taking the standard deviation of these differences in consumption by household. This standard deviation, labeled “dayonday” in Table A.2, was around 5 kWh for the average household, however there are a small number of highly variable households that skew the distribution of this volatility measure. Finally, we proxy for income by categorizing the households’ 3-year average monthly use by percentile.

## 5. Results

### 5.1. *Saliency and the arrival of the bill*

Results from estimation of equation (1) are reported in Table A.3 and Figure B.1. Table A.3 reports estimates of the bill arrival effect across specifications that offer differing flexibility in household-specific seasonality and temperature responses. Recall from equation (1) that we allow temperature responses and seasonal effects (weekly and monthly) to be household-specific, and that we allow for the possibility of lag dependence in consumption. The first four columns of Table A.3 omit the lagged dependent variable and the last five columns include it. All specifications include a quadratic function of Cooling and Heating Degree Hours as well as the full set of household covariates, but only some of the specifications include interactions. Column (1) reports results from a basic OLS regression in which all the household covariates enter linearly and are not interacted with temperature. Month and weekday dummies that are not household-specific are also included. The coefficient on  $D_{it}$  (labeled as “post” in the tables) is the conditional average effect on consumption of being in the first week following the arrival of a bill. The coefficient reveals a reduction of 0.28 kWh per day or 1.35% of average daily consumption in response to the issuance of a bill.

The remaining columns add temperature flexibility (in the form of household-temperature interactions), seasonal flexibility (in the form of household-specific weekday and month dummies), and lagged dependent variables to capture persistence. These richer specifications lead to a slight reduction in the size of the bill effect, but the size of the reduction is largely consistent across specifications. In column (2) we add household fixed effects to the basic OLS regression in column (1) to control for household-specific characteristics that are not captured by our constructed covariates (which are swept out in all the fixed effects specifications). Column (3) includes household-month and household-weekday effects to allow for heterogeneity in seasonal behavior, and column (4) includes household-weekday-month effects to allow for household-specific weekly patterns to vary across months of the year. Both columns (3) and (4) allow household-specific temperature responses by interacting the full  $\mathbf{Z}_i$  vector with Heating and Cooling Degree Days. These specifications represent more flexible approaches to estimating the regression from column (2). Column (5) adds a lagged dependent variable to the OLS specification from column (1). The coefficient on the lagged variable is highly significant and fairly large (0.83), although ignoring unit-specific heterogeneity (omitting fixed effects) in OLS estimation of dynamic panels is known to be biased upward. Column (6) adds more temperature-household covariate interactions and produces similar results. Columns (7) through (9) add a lagged dependent variable to the more flexible

specifications from columns (2) through (4), respectively.

The estimates from the more flexible specifications range from 0.13 to 0.20 kWh per day, or 0.6% to 1.0% of average daily consumption. The size of the impact is relatively small, but consistent with the recent literature on behavioral nudges in energy consumption. To place this in context, field experiments with the express purpose of reducing consumption have led to changes in demand between 2% and 5%. Throughout the rest of the paper we report results from the specification in column (9) because it offers the most flexibility for capturing household heterogeneity in seasonality and weather responses.

In order to trace out the average response throughout the billing cycle, we estimated the specification in column (9) of Table A.3 using a dummy variable for each day in the billing cycle instead of the “post” variable. Figure B.1 plots a moving average of these coefficients and their confidence intervals. A moving average is helpful to visualize the time path of the effect because the estimates at the billing cycle-day level are quite noisy. The reference observation in this regression is day 0, or the day after the last day of the previous cycle and the day on which the bill for that cycle is sent to the household.

It is important to notice that consumption appears to be 0.1 to 0.3 kWh lower in the first week to 10 days of the billing cycle than during the rest of the month, and consumption rises throughout the month until around the 20th day of the billing cycle. There is a slight decline after the 20th day of the billing cycle. One possible explanation for this is that bills are due 19 days following the beginning of a new cycle. To the extent that some households wait until the due date to pay their bill, and the bill becomes salient again during the payment, we might expect a decline. Another possibility is that the 20th day of the billing cycle corresponds to some households entering a new calendar month<sup>13</sup>. The beginning of a calendar month might affect consumption for several reasons, for example, if calendar months correspond to new paychecks and budgeting cycles because of payroll timing. However, many payroll systems do not issue paychecks on the first of the calendar month and there is considerable

---

<sup>13</sup>In our sample the billing cycle rarely extends beyond 31 days. A 32nd billing cycle day occurred for less than 1% of households and occurred on only three calendar days for those houses (longer billing cycles occasionally happened for more houses in parts of the data that were dropped before estimation). Temperatures happened to be unusually cold on these three days, and consumption for the fraction of houses on the 32nd day of their billing cycle was even lower than what our quadratic polynomial weather model predicts. These three days seem to exhibit weather-driven low consumption outliers. For this reason we dropped any household-day observations on the 32nd day of their individual billing cycle when estimating the cycle-day coefficients plotted in figure B.1. In specifications where they are included, there is a steep drop on day 32 but this is only identified off a small number of observations under unusual conditions. However, these household-day observations were included in all other specifications reported in this paper.

heterogeneity in pay cycles in California. Including a dummy variable for first week of the calendar month does not change the results in Figure B.1 or Table A.3.

### *5.2. Heating and cooling as the adjustment mechanism*

In order to uncover the mechanism by which households respond to reminders, we estimate our model by season, by quintile of our air conditioning demand proxy, and by hour of day (Tables A.4 and A.5, and Figures B.2, B.3 and B.4). The evidence reveals that the bill reminder effect works primarily through reductions in the use of air conditioning at peak times on hot days. This suggests that non-price, information-based “nudges” normally targeted for overall conservation can also address the goals of peak load curtailment for which price-based instruments like TOU or real-time rates are normally used.

In Table A.4 we examine the bill response in different seasons and under different temperature conditions. The largest estimated reductions occur on hot days in the summer, although colder days in the winter also account for a significant portion of the response. Columns (1) and (3) show that the average bill response is about 50% larger in the summer than in the winter. When the “post” variable is interacted with cooling degree and heating degree hours, allowing the response to vary with temperature during the period in which bill response occurs (columns (2) and (4)), the marginal effects of bill arrival are also largest in the summer, when they are driven by warm temperatures. In contrast, the marginal effects in winter are driven by cold temperatures, but the effects are more muted than those found for summer. The marginal effects at the mean cooling and heating degree hours for each season are slightly larger than the coefficients in columns (1) and (3), and the distribution of marginal effects is negatively skewed.

Figure B.2 plots kernel density estimates of the distributions of marginal effects for summer and winter. The summer plot includes the density plots evaluated at relatively hot versus relatively mild days (the 75th and 25th percentile of Cooling Degree Hours), and the winter plot includes the density plots evaluated at relatively cold versus relatively mild days (the 75th and 25th percentile of Heating Degree Hours). The mass of the distribution of these marginal effects lies primarily below zero with a much larger tail of negative responses during the summer. Large changes in Cooling Degree Hours drive the shift during the summer, whereas the winter effects are driven by Heating Degree Hours. These results indicate that to the extent that households respond to the bill arrival, they do so by reducing the use of the heating or cooling systems that they are likely to be running at the time. Because many households in San Diego are heated with natural gas but cooled with electric air conditioning units, the effect on electricity in summer is considerably larger.



To confirm that air conditioning is driving the response, we estimated our model separately for each quintile of our air-conditioning proxy. The results reported in Table A.5 indicate that households that are more sensitive to temperature increases above 65 degrees within a given day (our air-conditioning proxy) are also more responsive to the arrival of a bill. It should be noted that our proxy combines households likely to have larger air-conditioning units and those likely to use air-conditioning more intensively, so we are not able to separate preferences from technology. The effect is larger across quintiles in terms of the estimated coefficient and as a percentage of the average daily consumption in each quintile, with the effect of the arrival of the bill being 0.3% in the lowest quintile and 1.2% in the highest quintile. The absolute effect of a bill arrival is an order of magnitude larger for households in the highest quintile of air conditioning than those in the lowest quintile, and the relative effect is four times as large.

Finally, using the hourly data we show that adjustments primarily come during peak hours in summer months for households with more air conditioning demand (Figures B.3 and B.4). Figure B.3 shows the conditional means of consumption at each hour of the day for summer and winter separately, conditioning on all the covariates in our econometric model. The graph shows that the average summer effect is not only larger, but it occurs during hours in which air-conditioning use and aggregate peak demand are likely to be highest: between noon and 8 p.m. The dip in the size of the effect in the late afternoon may be attributed to air conditioning needs accumulating during the day as houses heat up, or by residents arriving home from school and work, or both. Figure B.4 shows the same information for summer only, plotted by quintile of our air-conditioning measure. The effects are larger during hot hours of the day and for households with a larger air-conditioning proxy, however the highest quintile group reduces its response dramatically between 4 p.m. and 7 p.m. This pattern is consistent with households who want to conserve, that have identified air-conditioning as the most effective mechanism, but who “give-in” to air-conditioning demand during the hottest hours of the day. Alternatively these households may not “pre-cool” their houses in the morning but turn up the air conditioning when they really need it late in the day. Nonetheless, the fact that the overall salience effect is largest during the summer, on the hottest days, during peak hours, among houses with the most weather sensitivity suggests that air conditioning is a major component of adjustment.

### *5.3. Heterogeneity in salience effect across customer segments and time*

The effect of price salience on consumption also varies across customers segmented by home size (or percentile of average use) and rate structure. In Table A.6, we split the sam-

ple by a proxy for income - the quintiles of historical average monthly electricity use. The magnitude of the absolute bill arrival effect increases monotonically across the quintiles, although the responses are not statistically different in percentage terms. Similarly, in Table A.7 we show that responses are smaller in absolute terms for households on the low income rate than the standard residential rate but approximately equal in terms of percentage reductions. However, despite their much larger average consumption, general service rate customers (e.g., small agricultural operations with residences or commercial living facilities) have a comparable absolute response and a much smaller percentage response. This could reflect either greater attention paid to energy costs by profit maximizers.

In Table A.8 we examine whether our “post” variable defined as the first seven days of a new billing cycle properly captures the time in which the bill is seen by the household. We reestimate our preferred base specification from Table A.3 with “post” redefined as the first three days of the billing cycle, the first seven days, 10 days, and 14 days. The effect dissipates as the window of time expands, consistent with the price becoming less salient and the household returning to a consumption path chosen with an opaque price. The last column of this table includes a dummy variable for each of the first three weeks of the billing cycle. The reference observation in this case is a day in the final nine or 10 days of the billing cycle. As shown in Figure B.1, there is a slight consumption decline late in the cycle. The results in Table A.8 confirm this; while conditional average consumption in the first week is 0.14 kWh below the late cycle average, consumption in the third week is highest at 0.14 kWh *above* the final days in the cycle. If additional reminders about energy costs or consumption budgets occur in those final days such as the payment of the bill or the receipt of a paycheck, then it is possible that the average salience effect is closer to the difference between the third week average and the first week, or at least somewhat larger than the 0.1 to 0.2 kWh per day average effect we report in Table A.3.

## 6. Conclusion

In neoclassical demand theory, the costs of decisions are known at the time choices are made. For goods like electricity use, water consumption, and medical treatments for chronic illness, however, expenditure is only “experienced” intermittently while choice occurs at greater frequency. Digital tracking and the proliferation of automated payments have made intermittent cost signals more common for an expanding share of goods, from toll road access and software downloads to standard consumption goods paid by credit card or other “bill-me-later”-type systems. Yet we know surprisingly little about how these payment patterns

affect decisions. This paper exploits hourly household electricity consumption data collected by “smart” electricity meters to examine dynamic consumer behavior under intermittent signals.

Our results suggest that residential energy consumers reduce consumption by 0.6% to 1% in the week after receiving a bill. The response magnitudes are larger for larger households, but the percentage responses are almost the same across household sizes. Most interestingly, reductions continue through the peak hours of the day, providing the first evidence to our knowledge that information signals (as opposed to direct time-of-use pricing) can reduce peak electricity demand. Summertime cooling is a key driver of this peak demand reduction, suggesting that thermostat adjustments are a low cost response to the billing information signal.

For the energy industry specifically, our analysis raises important issues for optimizing the return on smart grid investments. Achieving demand response and conservation through household behavioral change has proven frustrating for policy makers and utilities. Our results provide an explanation: when a bill reminds a household about a component of their budget constraint, initial responses evaporate as salience fades and households revert back to their higher consumption patterns. In the short run, however, the households that are most information-elastic also happen to be the most valuable for electric utilities to target: large households that use a lot of air conditioning on hot days. The smart grid will enable new information technologies and platforms. Once customers have become accustomed to these new formats, will they continue to use the information? Our results suggest that the salience of periodic alerts is real, but small, and that the persistence of electricity demand choices is an important component of price and quantity salience. Another remaining question for future research is how frequently do households need to be “nudged” with feedback, reminders, or normative messages in order to permanently shift the long run consumption path or achieve a given conservation goal. The medical literature offers a cautionary tale with regards to lulling effects that may result if reminders are too frequent (Pop-Eleches et al. (2011)).

Our results raise questions about how utilities might segment their market alerts by individual consumption patterns in addition to demographic and geographic information, and how behavioral “nudges” must be structured to push the household to a new steady state consumption path as opposed to achieving only transitory deviations. Whether these results generalize to markets other than electricity and what they imply for firm pricing models and cash flow management are rich areas for future research.

## References

- Agarwal, Sumit, John Driscoll, Xavier Gabaix, and David Laibson**, “Learning in the Credit Card Market,” *Available at SSRN 1091623*, 2013.
- Allcott, H.**, “Social norms and energy conservation,” *Journal of Public Economics*, 2011.
- **and S. Mullainathan**, “Behavior and energy policy,” *Science*, 2010, *327* (5970), 1204–1205.
- Allcott, Hunt**, “Rethinking real-time electricity pricing,” *Resource and Energy Economics*, 2011, *33* (4), 820–842.
- **and Todd Rogers**, “The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation,” Technical Report, National Bureau of Economic Research 2012.
- Becker, G. S. and K. M. Murphy**, “A theory of rational addiction,” *The Journal of Political Economy*, 1988, pp. 675–700.
- Bushnell, James B and Erin T Mansur**, “Consumption Under Noisy Price Signals: A Study of Electricity Retail Rate Deregulation in San Diego,” *The Journal of Industrial Economics*, 2005, *53* (4), 493–513.
- Chetty, R., A. Looney, and K. Kroft**, *Saliience and taxation: Theory and evidence*, 2007.
- Darby, S.**, “The effectiveness of feedback on energy consumption: a review for DEFRA of the literature on metering, billing and direct displays,” Technical Report, Environmental Change Institute April 2006 2006.
- DellaVigna, Stefano**, “Psychology and economics: Evidence from the field,” *Journal of Economic Literature*, 2009, pp. 315–372.
- Finkelstein, A.**, *E-ZTax: Tax Saliience and Tax Rates*, 2007.
- Graff Zivin, Joshua, Matthew J. Kotchen, and Erin T. Mansur**, “Spatial and Temporal Heterogeneity of Marginal Emissions: Implications for Electric Cars and Other Electricity-Shifting Policies,” *Journal of Economic Behavior & Organization*, 2014, *This issue*.

- Grubb, Michael and Matthew Osborne**, “Cellular service demand: biased beliefs, learning, and bill shock,” 2012.
- Hsiaw, Alice and Matthew Harding**, “Goal Setting and Energy Conservation,” *Journal of Economic Behavior & Organization*, 2014, *This issue*.
- Ito, Koichiro**, “Do consumers respond to marginal or average price? Evidence from non-linear electricity pricing,” Technical Report, National Bureau of Economic Research 2012.
- Jaffe, Adam B. and Robert N. Stavins**, “The energy paradox and the diffusion of conservation technology,” *Resource and Energy Economics*, 5 1994, 16 (2), 91–122.
- Jessoe, Katrina and David Rapson**, “Knowledge is (Less) Power: Experimental Evidence from Residential Energy Use,” April 2013.
- Jin, Ginger Zhe and Phillip Leslie**, “The effect of information on product quality: Evidence from restaurant hygiene grade cards,” *The Quarterly Journal of Economics*, 2003, 118 (2), 409–451.
- Kahn, Matthew E and Frank A Wolak**, “A Field Experiment to Assess the Impact of Information Provision on Household Electricity Consumption,” Technical Report ARB Contract 08-325, California Air Resources Board, <http://arbis.arb.ca.gov/research/rsc/3-8-13/item6dfr08-325.pdf> January 2013.
- Karlan, Dean, Margaret McConnell, Sendhil Mullainathan, and Jonathan Zinman**, “Getting to the top of mind: How reminders increase saving,” Technical Report w16205, National Bureau of Economic Research 2010.
- McKinsey**, “Wasted Energy: How the U.S. Can Reach its Energy Productivity Potential,” Technical Report June 2007 2007.
- Pop-Eleches, Cristian, Harsha Thirumurthy, James Habyarimana, Joshua Graff Zivin, Markus Goldstein, Damien de Walque, Leslie Mackeen, Jessica Haberer, John Sidle, Duncan Ngare et al.**, “Mobile Phone Technologies Improve Adherence to Antiretroviral Treatment in Resource-Limited Settings: A Randomized Controlled Trial of Text Message Reminders,” *Aids*, 2011, 25, 825–834.
- Sexton, S. E.**, “DRAFT: Automatic Bill Payment, Price Salience, and Consumption: Evidence from Residential Electricity Consumption,” 2011.

**Sims, C. A.**, “Implications of rational inattention,” *Journal of Monetary Economics*, 2003, 50 (3), 665–690.

**Stango, Victor and Jonathan Zinman**, “Limited and varying consumer attention: Evidence from shocks to the salience of bank overdraft fees,” Technical Report w17028, National Bureau of Economic Research 2011.

**Wolak, F. A.**, “Residential customer response to real-time pricing: the Anaheim critical-peak pricing experiment,” Technical Report Paper CSEMWP-151 February 14, 2007 2007.

**Wolak, Frank A.**, “Do residential customers respond to hourly prices? Evidence from a dynamic pricing experiment,” *The American Economic Review*, 2011, 101 (3), 83–87.

## Appendix A. Tables

Table A.1: SDG&E Rate Structures 4/15/2009 to 1/31/2009

		Rates (\$/kWh)		
Dates		DR	DRLI	A (flat rate)
4/15/2009	Baseline (11.5 kWh/day)	0.04363	0.04854	0.06544
	101% to 130% of Baseline	0.0638	0.06871	
	131% to 200% of Baseline	0.21818	0.16067	
	Over 200% of Baseline	0.23818	0.16067	
5/1/2009	Baseline (11.8 kWh/day)	0.01692	0.02183	0.07508
	101% to 130% of Baseline	0.03709	0.042	
	131% to 200% of Baseline	0.20379	0.14472	
	Over 200% of Baseline	0.22379	0.14472	
9/1/2009	Baseline (11.8 kWh/day)	0.01692	0.02183	0.07239
	101% to 130% of Baseline	0.03709	0.042	
	131% to 200% of Baseline	0.1988	0.13973	
	Over 200% of Baseline	0.2188	0.13973	
11/1/2009	Baseline (11.5 kWh/day)	0.04455	0.04946	0.06253
	101% to 130% of Baseline	0.06472	0.06963	
	131% to 200% of Baseline	0.21019	0.15268	
	Over 200% of Baseline	0.23019	0.15268	
1/1/2010	Baseline (10.8 kWh/day)	0.06026	0.05867	0.07222
	101% to 130% of Baseline	0.08103	0.07944	
	131% to 200% of Baseline	0.18952	0.1394	
	Over 200% of Baseline	0.20952	0.1394	

*Notes:* During the time span our data covers (April 15, 2009 to January 31, 2010), residential electricity rates changed several times by small amounts. This table describes in detail the rates faced by our sample of households.

Table A.2: Summary statistics

<b>Variable</b>	<b>Mean</b>	<b>Std. Dev.</b>	<b>Min.</b>	<b>Max.</b>
<i>Panel A. Observed Variables</i>				
Daily usage (kWh)	20.7	18.3	0	629
Monthly usage (kWh)	621	509	0	14310
Bill (\$)	115	127	5.07	2955
CDH	107	113	0	503
HDH	118	118	0	444
<i>Panel B. Constructed Variables</i>				
Usage percentile	56.5	27.2	0.396	99.9
Persist	0.27	0.20	-0.61	1.23
Dayonday	5.13	3.49	0.013	51.6
Hotco	0.026	0.032	-0.10	0.28
Coldco	-0.008	0.017	-0.17	0.19
Fit	0.41	0.17	0.015	0.99
Pattern	-0.74	1.14	-106	-0.003

*Notes:* Our sample has 10,826 households observed between 71 and 228 days for a total of 2,914,313 observations for all variables except daily consumption. Occasional technical problems with the smart meter data system caused missing observations on daily consumption for some days during the sample period. The sample size of daily consumption observations is 2,044,222. The variables in panel B were constructed from household-specific monthly, daily, and hourly data treating each household separately. Persist is the household-specific coefficient on lagged consumption. Dayonday is the interday standard deviation of daily consumption. Hotco and Coldco are household specific coefficients on hourly hot and cold temperatures, respectively, and Fit is the  $R^2$  from the regression that produced those coefficients. Pattern is the percentage difference in average hourly use between the highest and lowest hours of the day. These variables are described in more detail in Section 4.



Table A.3: Regression Results for the Average Consumption Response to the Arrival of a Bill

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$y(t-1)$					0.83*** (0.00184)	0.81*** (0.00209)	0.61*** (0.00761)	0.36*** (0.000724)	0.43*** (0.00649)
Post	-0.28*** (0.0215)	-0.19*** (0.0138)	-0.13*** (0.0101)	-0.16*** (0.0107)	-0.14*** (0.0138)	-0.15*** (0.0137)	-0.13*** (0.00996)	-0.14*** (0.0111)	-0.20*** (0.0118)
Temperature*household covariate interactions:									
	NO	NO	YES	YES	NO	YES	NO	YES	YES
Fixed effects:									
Weekday & Month	X	X			X	X	X		
Household		X					X		
Household-Weekday			X					X	
Household-Month			X					X	
Household-Weekday-Month				X					X
$N$	2044222	2044222	2044222	2044222	1606732	1606732	1606732	1606732	1606732
$R^2$	0.529	0.208	0.912	0.310	0.860	0.863	0.520	0.928	0.465

Notes: Robust standard errors are reported in parentheses. The dependent variable is daily total consumption measured in kWh. Post is a dummy variable indicating the first week of a new billing cycle. Household covariates include all those listed in panel B of Table A.2. All regressions include the household covariates and a quadratic in Cooling and Heating Degree Hours.

\*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

Table A.4: Regression Results for the Response to the Arrival of a Bill, by Season and Temperature

	June-Sept.		Oct.-Jan.	
	1	2	3	4
y(t-1)	0.41*** (0.00327)	0.41*** (0.00327)	0.45*** (0.0244)	0.45*** (0.0244)
Post	-0.22*** (0.0165)		-0.14*** (0.0183)	
Post*CDH		-0.0018*** (0.00011)		0.00024 (0.000245)
Post*HDH		0.0013*** (0.00029)		-0.0011*** (0.000112)
Marginal Bill Effect (kWh)		-0.28*** (0.018)		-0.19*** (0.019)
<i>N</i>	932963	932963	673769	673769
<i>R</i> <sup>2</sup>	0.511	0.511	0.278	0.278

*Notes:* Robust standard errors are reported in parentheses. The dependent variable is daily total consumption measured in kWh. Post is a dummy variable indicating the first week of a new billing cycle. All regressions include all household covariates listed in panel B of Table A.2, a quadratic in Cooling and Heating Degree Hours, the interaction of the household covariates with Cooling and Heating Degree Hours, and house-by-month-by-weekday fixed effects.

\*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

Table A.5: Regression Results for the Response to a Bill by Air-conditioning Use

	Households divided by quintile of air-conditioning proxy				
	1	2	3	4	5
Full Sample					
y(t-1)	0.49*** (0.0529)	0.37*** (0.00820)	0.40*** (0.00540)	0.41*** (0.00483)	0.41*** (0.00460)
Post	-0.041** (0.0189)	-0.085*** (0.0157)	-0.18*** (0.0223)	-0.33*** (0.0279)	-0.41*** (0.0396)
Mean consumption (kWh/day)	15.0	13.9	18.1	22.0	34.4
Percent reduction	0.3%	0.6%	1.0%	1.5%	1.2%
<i>N</i>	322217	321874	320924	322890	318827
<i>R</i> <sup>2</sup>	0.260	0.203	0.308	0.436	0.562
Summer Months Only, Allowing Bill Response to Vary with Temperature					
y(t-1)	0.37*** (0.0126)	0.37*** (0.0111)	0.40*** (0.00605)	0.40*** (0.00551)	0.41*** (0.00528)
Post*CDH	-0.00016 (0.000126)	-0.00023* (0.000125)	-0.0011*** (0.000213)	-0.0026*** (0.000281)	-0.0049*** (0.000370)
Post*HDH	0.00048 (0.000493)	0.00036 (0.000423)	-0.00016 (0.000546)	0.00096 (0.000657)	0.0038*** (0.000953)
Marginal Effects (kWh)	-0.015 (0.0215)	-0.03 (0.0202)	-0.19*** (0.0333)	-0.42*** (0.0442)	-0.74*** (0.0599)
Mean consumption (kWh/day)	14.5	14.0	19.0	24.7	40.7
Percent reduction	0.1%	0.2%	1.0%	1.7%	1.8%
<i>N</i>	186468	185957	187078	185991	187469
<i>R</i> <sup>2</sup>	0.148	0.227	0.348	0.475	0.599

*Notes:* Robust standard errors are reported in parentheses. The dependent variable is daily total consumption measured in kWh. Post is a dummy variable indicating the first week of a new billing cycle. All regressions include all household covariates listed in panel B of Table A.2, a quadratic in Cooling and Heating Degree Hours, the interaction of the household covariates with Cooling and Heating Degree Hours, and house-by-month-by-weekday fixed effects.

\*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

Table A.6: Regression Results for the Response to a Bill by Household Size

	Households divided by quintile of average bill				
	1	2	3	4	5
y(t-1)	0.38*** (0.00672)	0.38*** (0.00714)	0.39*** (0.00525)	0.40*** (0.00488)	0.45*** (0.0116)
Post	-0.084*** (0.0106)	-0.16*** (0.0166)	-0.21*** (0.0224)	-0.29*** (0.0283)	-0.36*** (0.0422)
Mean consumption (kWh/day)	7.5	12.6	17.3	23.7	42.3
Percent reduction	1.1%	1.3%	1.2%	1.2%	0.9%
<i>N</i>	321871	321685	320735	321197	321244
<i>R</i> <sup>2</sup>	0.311	0.364	0.403	0.444	0.500

*Notes:* Robust standard errors are reported in parentheses. The dependent variable is daily total consumption measured in kWh. Post is a dummy variable indicating the first week of a new billing cycle. All regressions include all household covariates listed in panel B of Table A.2, a quadratic in Cooling and Heating Degree Hours, the interaction of the household covariates with Cooling and Heating Degree Hours, and house-by-month-by-weekday fixed effects.

\*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

Table A.7: Regression Results for the Response to a Bill by Rate Class

	Households divided by electricity rate		
	Residential	Low Income	General Service
y(t-1)	0.42*** (0.00280)	0.36*** (0.00426)	0.49*** (0.0332)
Post	-0.19*** (0.0146)	-0.16*** (0.0170)	-0.20** (0.0856)
Mean consumption (kWh/day)	21.2	15.8	36.6
Percent reduction	0.9%	1.0%	0.5%
<i>N</i>	1061309	447864	97559
<i>R</i> <sup>2</sup>	0.484	0.425	0.437

*Notes:* Robust standard errors are reported in parentheses. The dependent variable is daily total consumption measured in kWh. Post is a dummy variable indicating the first week of a new billing cycle. All regressions include all household covariates listed in panel B of Table A.2, a quadratic in Cooling and Heating Degree Hours, the interaction of the household covariates with Cooling and Heating Degree Hours, and house-by-month-by-weekday fixed effects.

\*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

Table A.8: Robustness of Regression Results to Different Periods Following Bill Arrival

	1	2	3	4	5
y(t-1)	0.43*** (0.00649)	0.43*** (0.00649)	0.43*** (0.00650)	0.43*** (0.00650)	0.43*** (0.00650)
3 Day	-0.22*** (0.0190)				Week 1 -0.14*** (0.0125)
7 Day		-0.20*** (0.0118)			Week 2 0.095*** (0.0128)
10 Day			-0.18*** (0.0111)		Week 3 0.14*** (0.0122)
14 Day				-0.070*** (0.0103)	
<i>N</i>	1606732	1606732	1606732	1606732	1606732
<i>R</i> <sup>2</sup>	0.465	0.465	0.465	0.465	0.465

*Notes:* Robust standard errors are reported in parentheses. The dependent variable is daily total consumption measured in kWh. 3 Day is a dummy variable indicating the first three days of a new billing cycle. 7 Day, 10 Day, and 14 Day are defined similarly. Column 2 reproduced the results from column 9 of Table A.3, our preferred specification. Column 5 includes a dummy variable for each of the first three weeks of the billing cycle. All regressions include all household covariates listed in panel B of Table A.2, a quadratic in Cooling and Heating Degree Hours, the interaction of the household covariates with Cooling and Heating Degree Hours, and house-by-month-by-weekday fixed effects.

\*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

## Appendix B. Figures




Figure B.1: Conditional consumption throughout the billing cycle.

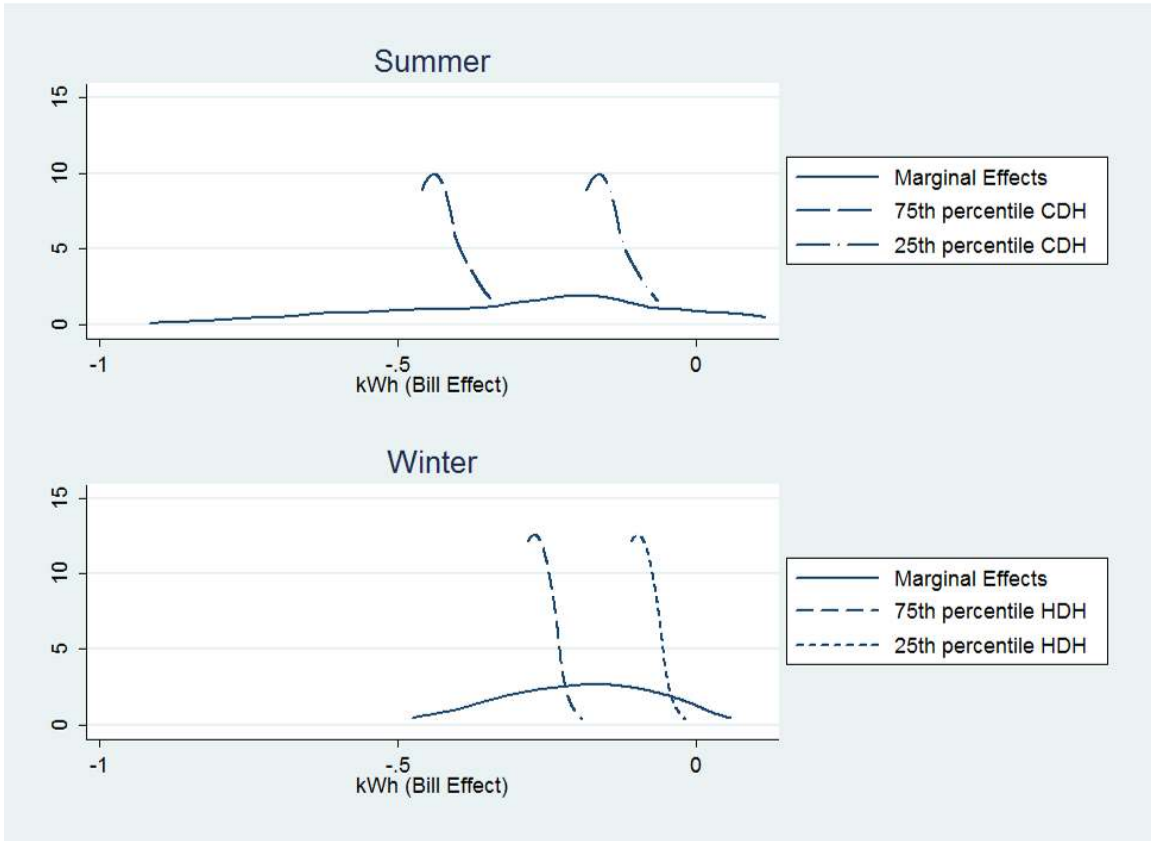


Figure B.2: Marginal effects of bill arrival at different cooling degree hours during the summer and heating degree hours in the winter.



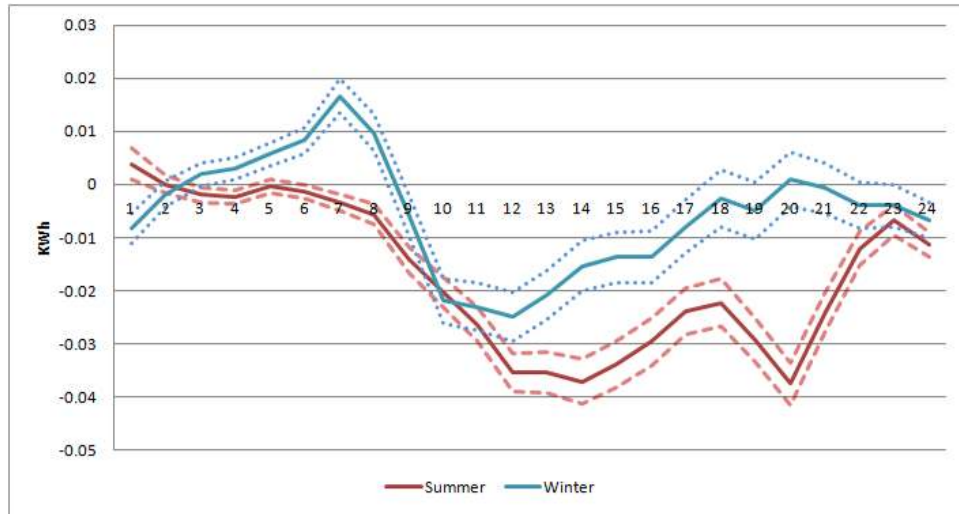


Figure B.3: Conditional response to bill arrival by hour of day in summer and in winter.




Figure B.4: Conditional response to bill arrival by hour of day in summer for each quintile of air-conditioning.