Shock Avoidance:
High Bill Alert Programs and Energy Consumption

Grant D. Jacobsen and James I. Stewart*

Last Update: February 2024

Abstract

Utilities have invested billions of dollars in advanced metering infrastructure (AMI) but proceeded slowly to deploy AMI-enabled programs that benefit consumers. High bill alert (HBA) programs, which inform consumers of unusually high usage patterns, offer an avenue for tapping AMI-enabled benefits. We evaluate an HBA program and find that the program reduced mean electricity and natural gas consumption by about 0.5%. The effects were largest at the top of the usage distribution, especially when normalized by pre-program usage, indicating that households experienced fewer expenditure shocks. Welfare estimates depend critically on assumptions related to the forgone value of the conserved energy.

JEL Codes: Q41, D12, L97
Keywords: bill shocks, high bill alerts, energy, inflation, electricity, energy, natural gas, consumer welfare

*Jacobsen (grant.jacobsen@oregonstate.edu) is at Oregon State University and Stewart (jim.stewart@cadmusgroup.com) is at the Cadmus Group. Jacobsen gratefully acknowledges funding received from the Consumer Protection Research Fund. We are thankful for comments received at the Allied Social Science Annual Meeting, the Consumer Protection Symposium,, the Workshop on Environmental Economics and Data Science, the University of Alaska, and the Federal Reserve Bank of Kansas City. This research was reviewed under UO IRB Protocol 01252021.032.

1 Introduction

Generation, transmission, and distribution costs for electric utilities are expected to rise in the coming years. Spurred by federal clean-energy investment tax credits and aggressive new state regulatory mandates to reduce carbon emissions, utilities are retiring coal and oil electricity plants and investing in cleaner sources of power including wind and solar generation (Borenstein, 2012; Borenstein and Kellogg, 2023). Utilities are also modernizing their transmission and distribution systems to support the integration of renewables, to improve the reliability of energy delivery, and to enhance system resiliency against adverse weather events or other natural disasters such as wildfires and earthquakes (USEIA, 2021). Meanwhile, robust worldwide demand for oil and natural gas has coincided with increased international conflict to further elevate energy costs to utilities. Simultaneously, governments have implemented policies to move energy consumption increasingly toward electrification, most notably encouraging the adoption of electric vehicles.

These developments are likely to affect households in the form of higher and more volatile utility energy bills (USEIA, 2022). As a result, unexpectedly high bills or "bill shocks" may become endemic because many households often cannot or do not track their energy consumption, only learning about their expenses after receiving a monthly bill.[1] Bill shocks are more than an annoyance; they can be particularly deleterious for the poor, forcing economically insecure households to reduce expenditures on other goods and services such as food, housing, and medical care to pay utility bills and "to keep the lights on" (Tuttle and Beatty, 2017; Drehobl et al., 2020; Kontokosta et al., 2020).

A possibly fortuitous development with respect to the increasing prevalence of bill shocks is that many utilities have upgraded their metering infrastructure in recent years. These upgrades have the potential to enable households to have better access to information on their usage and expenditures and avoid unexpected surprises that exacerbate difficult trade-offs for household budgets. Most notably, "smart meters," which are digital meters that replace traditional mechanical meters, can measure and communicate usage data at high

[1]A small number of U.S. utilities have experimented with pre-payment plans, which provide the customer with real-time information about their energy consumption. However, pre-payment plans are much more common outside the United States, particularly in the developing world (Jack and Smith, 2020).

frequencies. The appeal to households of this enhanced information flow for managing energy expenditures was reflected in a 2008 report prepared for the US DOE, which reported that "the Smart Grid will empower average energy consumers to a degree unimaginable just a few years ago. Given new awareness, understanding and tools, they'll be able to make choices that save money, enhance personal convenience, improve the environment – or all three (Litos Strategic Communication, 2008)."[2]

The prospect of empowering consumers and enhancing informational flows have justified billions of dollars of public investment in smart meters over the recent decades.[3] These investments have been effective at increasing deployment of smart meters. Estimates indicate that nearly 80% of U.S. households had smart meters by the end of 2022 (Edison Foundation, 2022). Unfortunately, despite claims about the benefits of smart meters for consumers and the widespread deployment of smart meters, the benefits of smart meters appear to be largely untapped by utilities, who are often ill-equipped to leverage the volumes of data provided by the meters or provide access to it to their consumers and, as a result, have often deactivated data-sharing features on smart meters (Mission Data, 2022). Reflecting the untapped potential of smart meters, List et al. (2018) present evidence that the installation of smart meters has not led to statistically significant changes in energy usage for either electricity or natural gas.

Motivated by recent household and regulator concerns about bill shocks and the underutilization of smart meter investments, this paper focuses on a new energy management tool which may tap the potential offered by smart meter investments: high bill alert (HBA) programs. HBA programs alert consumers to unusually large increases in their usage part way through their bill cycle. The alert gives households the opportunity to adjust their energy usage before a bill shock is fully realized, thereby limiting the negative consequences of wasteful or low-value energy expenditures. HBA programs are a new tool but may become as ubiquitous in the utility sector as other commonly adopted behavioral programs,

[2]The opportunity to avoid unexpected bill shocks was further highlighted in Energy.gov article that focused on unusually large electric bills occurring due to undetected equipment failures or malfunctions, concluding that "these types of problems can be avoided through the smart grid (Energy.gov, 2010)."

[3]For example, the 2009 American Recovery and Reinvestment Act allocated \$4.5 billion to grid modernization, which provided support for two major initiatives, the Smart Grid Investment Grad program and the Smart Grid Demonstration Program (SmartGrid.gov, 2023). More recently, the Infrastructure Investment and Jobs Act of 2022 included \$3 billion for Smart Grid Investment Matching Grants (Whitehouse.gov, 2022).

such as home energy reports (Allcott and Kessler, 2019; Allcott, 2015; Allcott and Rogers, 2014; Ayres et al., 2013; Allcott, 2011). Major third-party "implementors" in the utility sector, such the Oracle Corporation, are offering HBA programs to their clients, and some of the nation's largest utilities have initiated HBA programs, including Pacific Gas & Electric, Consumers Energy, Xcel Energy, Salt River Project, and Commonwealth Edison.[4]

Despite ongoing adoptions of HBA programs, we are unaware of peer-reviewed evidence of their effectiveness. This paper fills that gap by analyzing the impacts of a high bill alert (HBA) program run by a large Midwestern electric and gas utility in 2015 and 2016 designed to help households avoid bill shocks. The utility sent alerts part way through monthly billing cycles to dual-fuel households whose electricity or natural gas consumption was on track to exceed normal levels by 30% or more. For evaluation purposes, the program was implemented as a natural field experiment (Harrison and List, 2004), with about 50,000 households randomly assigned to a treatment group and 25,000 households to a control group. The utility delivered electricity alerts to about 31,000 households and gas alerts to about 6,000 households during the 13-month treatment window. In order to assess the persistence of the effects, we also analyze the effects of the program during the eleven months immediately after the program was discontinued.[5]

We report several key findings. First, the HBA program led to mean conservation of about 0.5% of electricity consumption and 0.5% of natural gas consumption during the treatment period. These effects are measured across all treatment group households regardless of whether a household received an alert and capture the average effect of the HBA program on enrolled households. Second, consistent with avoiding bill shocks, the estimates were small and insignificant at the bottom of the usage distribution and large and significant at the top of the usage distribution. This finding holds whether we measure usage in nominal terms or normalize consumption to measure usage relative to each household's

[4]HBA programs can be considered behavioral interventions, which have become common additions to the suite of demand-side management (DSM) programs offered by many utilities, in part because behavioral programs have lower costs and scale more easily than traditional programs that involve upgrading equipment to improve energy efficiency. DSM programs are often motivated or required by regulatory bodies through integrated resources plans or similar processes that require utilities to make choices that are aligned with "the public interest" (Lazar, 2016).

[5]The program was discontinued by the utility to allow for evaluation. It was later deployed universally among the customer base, although that time frame is outside our sample.

pre-program usage. Third, HBA conservation of both energy types persisted after the utility stopped sending HBAs. We find conservation of about 0.5% for electricity and gas in the eleven months after the utility stopped delivery, indicating that households made lasting adjustments (e.g., fixing or replacing energy inefficient equipment). Fourth, motivated by the current focus on distributional effects within the energy sector (e.g., the Biden Administration's Justice40 initiative), we examine whether the effect of the program varied by zip-code income and find little evidence of heterogenous effects by income, indicating the benefits accrued to households in both low- and high-income communities.

We estimate the consumer welfare benefits of the HBA program for the utility's residential customer population in calendar year 2021 were about \$200,000 – if the foregone value from the conserved energy was just marginal to the price of energy – and could have been as high as \$7 million if the foregone value was zero. The two bounds straddle our estimates of program costs, which illuminates an easily overlooked point: not all behavioral conservation is the same from a welfare perspective. The welfare impact depends on the foregone value to consumers of the avoided consumption. For example, if a utility is implementing a program, and the goal is energy conservation, there is greater value if the behavioral mechanism is conservation of low-value, wasteful energy – such as energy that is more likely to be eliminated by the HBA program – than if it is achieved through conservation of energy that may be high-value, such as households setting the thermostat at a less comfortable temperature in response to a home energy report that displays a household's energy consumption relative to their neighbors.

Our findings about the benefits of bill alerts also have relevance for consumers in other markets. Consumers can experience bill shocks whenever it is difficult to track one's consumption and consumption precedes payment, including in health care, internet service, phone service, and credit cards. Recognizing this potential harm, the Federal Communications Commission estimates one in six wireless consumers has experienced a mobile phone bill shock and now maintains a bill shock website (Horrigan and Satterwhite, 2020). Our research suggests that bill alerts may be a useful policy for helping consumers in other markets avoid such shocks. Our research also highlights the importance of "behavioral mechanisms." For example, interventions that aim to help people lose weight may be more valuable if they lead participants to shift to a more nutritious diet, and not just a low-calorie one. Similarly, programs that aim to increase youth sports participation are likely to be more valuable if the behavioral mechanism is reduced social media usage rather than reduce time spent studying or practicing musical instruments. In our case, the value of energy conservation is enhanced if it is achieved through savings that eliminate low-value or wasteful consumption.

In addition to having policy relevance and connecting to an array of consumer markets, this paper also adds to a large and growing literature on information feedback to utility customers about energy consumption and prices.[6] Much of this literature has focused on feedback provided continuously or at regular intervals that encourages customers to conserve energy daily or during specific times of day. Some studies find that providing feedback enhances household responsiveness to pricing or creates incentives for energy conservation (Allcott, 2011; Jessoe and Rapson, 2014; Martin and Rivers, 2018). However, other studies (Burkhardt et al., 2019; Fabra et al., 2021) find no statistically significant effects from information feedback. Our study contributes to this literature by presenting evidence about the effectiveness of personalized feedback delivered at irregular intervals, triggered by specific events (forecasts of high bills), and alerting households to potential economic losses.

Finally, our results contribute to the literature on the persistence of information and behavioral treatments effects. Assessing persistence is important because it indicates whether a behavior or information program is likely to work and be cost-effective in the long run. Citing a large literature on persistence of behavior changes in retail energy and other consumer markets, Brandon et al. (2022) observe that it is rare for the treatment effects of behavior modification and information treatments to persist after treatment ends. Other studies also show post-treatment decay of treatment effects (Allcott and Rogers, 2014; Gilbert and Graff Zivin, 2014). Our results stand out because we document persistence: despite irregular and infrequent delivery of HBAs, electricity and gas conservation fully persisted for at least one

[6]Other examples of interventions in the utility sector connected to consumer behavior include pairing price schedule changes with enhanced information, such as in-home displays (Jessoe and Rapson, 2014; Bollinger and Hartman, 2020; Matisoff et al., 2020; Schneider and Sunstein, 2017); subsidizing smart thermostats (Blonz et al., 2018); subsidizing high efficiency products (Borenstein and Davis, 2016; Jacobsen, 2019; Jacobsen, 2023; Neveu and Sherlock, 2016); and deploying behavioral demand response programs (Brandon et al., 2018).

year after treatment ended. The framing of forecasts of consumption increases as potential losses (Kahneman and Tversky, 1979) may have motivated households receiving HBAs to take long-lasting actions to reduce their consumption and contributed to the persistence of energy savings.

At the outset, we want to emphasize that this paper focuses on evaluating the effect of the HBA *program*, not an alert itself. There are three reasons for this. First, we view the aggregate effect of the program as the more important question, as that is the dimension that has most relevance to program managers and policymakers. Second, due to the randomized nature of enrollment in the program, we can obtain unbiased estimates of the effect of the HBA program. Third, alerts themselves are endogenous and triggered by usage. In practice, the association between an alert and usage will be a function of 1) the factors that triggered elevated usage at the beginning of the billing cycle, such as mechanical failures; and 2) any behavioral response to receiving a bill. Empirically, estimates that use the release of an alert as the treatment variable would capture the effect of both margins and it would not be possible to ascertain the role of each factor.[7]

Despite these challenges, we do, however, use a scaling procedure based on our estimates of program effects to bound estimates of local average treatment effects for households who received alerts. Upper bound estimates are based on the assumption that all alert effects are concentrated in billing cycles when a household received an alert, whereas lower bound estimates are based on the assumption that alerts affect consumption during all billing cycles on or after the household received its first alert. This procedure, which is described in more detail in Section 5.4, indicates that the effect of an electricity alert on electricity consumption is between 0.82% and 4.8% and the effect of a gas alert on gas consumption is between 5.9% and 41.2%.

[7] We present estimates of the relationship between an alert and energy usage in Table A.1 in the Appendix. The results indicate than alert is associated with elevated usage amounts, but only by up to 20% (depending on which fuel is used for the dependent variable and which fuel triggered the alert), which is lower than the projected 30% increase that triggers an alert.

2 High Bill Alerts: Background and Conceptual Model

An HBA is a notification from the utility sent mid-billing cycle that the household's next electricity or natural gas bill is expected to be higher than normal. In our case, the utility sent an HBA by email when a household's monthly energy consumption was projected to exceed its consumption for the same period in the year preceding the start of the HBA program by 30% or more and the bill was projected to be greater than $30. The HBA provided the household with a forecast of the monthly billing amount and an estimate of the difference between the forecasted bill and the bill for the same period in the previous year. The HBA also embedded a link to an online home energy assessment, where the household could receive personalized recommendations for conserving energy.

As consumers are often inattentive to their energy consumption and expenditures (DellaVigna, 2009; Allcott and Greenstone, 2012; Sallee, 2014; Jacobsen, 2015; Gillan, 2018; Gabaix, 2017; and Allcott and Knittel, 2019), it is natural to ask, why they would pay attention to HBAs, especially if the alerts contain similar information to monthly bills? One reason is that, unlike monthly bills, which are delivered on a schedule, HBAs are not delivered regularly. The asynchronous, event-based nature of HBAs raises their salience. A second reason is that the information in HBAs is framed in a way to motivate households to act. The HBA explicitly compares the forecasted consumption to consumption in the same period a year ago, and the difference is framed as a potential loss (e.g., "you're on track to spend $37 more than the same time last year"). Since households tend to be loss-averse and weigh potential losses of a given amount more than equal-sized gains (Kahneman and Tversky, 1979), the framing of HBAs in this way may motivate households to conserve. Generally, HBAs are a tool that enhances the amount of information households receive about their consumption and the information is presented in a way that is both salient and not overwhelming or overly complex.[8] Below, we present a model of how enhanced information is likely to affect energy consumption and consumer welfare.

[8]Jacobsen and Stewart (2022) present evidence that overly complex environments can overwhelm consumers and dampen their behavioral responses.

2.1 Model Setup

Consider a utility-maximizing household with income, *I,* that receives utility, *U(T, z)*, from consumption of energy services, *T,* produced in the home using utility-supplied energy, *e,* with unit price, *p,* and a composite of all other market-based goods and services, z. The household produces energy services with home appliances and equipment according to a constant returns to scale production function (Durbin and McFadden, 1984; Davis, 2008). The price of the market composite good is normalized to one. The household has standard, well-behaved preferences over consumption of market goods and services and home-produced energy services.[9]

At the beginning of each period, the household forms beliefs about the marginal cost of producing energy services $\hat{\mu}$. The marginal cost of energy services depends on the unit price of energy, p, and the efficiency of the production function, θ, with which the home converts energy e into home services such as space heat.[10] For energy services connected to weather (e.g., the cost of cooling the residence to 72 degrees), the marginal cost may also depend on outside weather. Due to efficiency degradation from wear-and-tear and/or mechanical failures, weather fluctuations affecting the amount of energy required to cool or heat a residence to a certain temperature, or changes in the price of electricity, the marginal cost of energy services will vary over time.[11] The household can instantaneously and without cost change its consumption in response to new information about the marginal cost of energy services. However, due to utility monthly billing, the household only learns its energy consumption and the marginal cost of energy services when it receives a bill at the end of the period.

Each period (i.e., a utility billing cycle), the household's constrained utility maximization

[9]Assume the household has locally nonsatiated and strictly convex preferences over T and z. Also, U is well-behaved, continuous, and twice differentiable function of T and z, with $U_k > 0$, $U_{kk} < 0$, $U_{kj \neq k} > 0$, for all T and z in $\{(T, z) : T \geq 0, z \geq 0\}$.

[10]Assume the CRS production function is $T = f(e;\theta)$, where θ is the efficiency with which the home converts energy into energy services and T is such that $f(0;\theta) = 0$, $\partial f/\partial e > 0$, $\partial^2 f/\partial e^2 = 0$, and $\partial f/\partial\theta > 0$ for all $e > 0$.

[11]Examples of efficiency degradation include break down of insulation levels in the home, furnaces developing cracks in heat exchangers or building up dust, and poor maintenance of HVAC equipment.

problem can be written as:

$$\max_{T,z} U(T,z)$$
$$\text{subject to } \hat{\mu}T + z \leq I \tag{1}$$

where the household budget constraint has been written as a function of the household's belief about the marginal cost of home energy services.[12] The household chooses consumption of energy services T at the beginning of the period. After learning about their actual expenditures on energy services at the end of the period, to maintain budget balance, the household spends its remaining income on the composite of market goods and services.[13] The household's utility for the period is a function of the quantities of energy services and the composite market good it consumes.[14]

2.2 Outcomes Absent an HBA Program

Figure 1 illustrates several solutions to the household's utility maximization problem when there has been a change in the marginal cost of producing energy services. We begin with scenario "A," which depicts the households chosen bundle given the composite market good price and the household's belief about the marginal cost of energy services μ_0. The household consumes T_A energy services with the expectation of consuming bundle (T_A, Z_A) and realizing utility U_A. At the end of the period, if the *ex post* marginal cost, μ_p, equals μ_0, the household will consume composite market goods and services Z_A, where the ex-post budget constraint is just satisfied and the marginal rate of substitution between home energy services and market-produced goods is equal to μ_0. Specifically, the household chooses Z_A to satisfy the budget constraint, which also satisfies the ex-ante first order necessary

[12]Pollak and Wachter (1975) show that if the home production technologies are constant returns to scale and there is no joint home production of goods and services, the solution to the household's utility maximization problem can be simplified by rewriting the household budget constraint as a function of consumption of the household energy services and the marginal cost of producing these services using utility supplied energy.

[13]Many aspects of household consumption conform to this model of sequential choice of energy services and market goods and services. For example, households program the home thermostat, select a temperature on the water heater, set the refrigerator temperature, and configure other appliance settings, which then determine the household's future energy consumption and billing charges.

[14]To simplify the analysis, we do not allow household to save and instead require budget balance in each period. In a dynamic multi-period model, savings could be used to insure against future bill shocks and to smooth consumption. Even under such a such a scenario, households would be harmed by informally insuring against unexpected bill shocks because they would need to forego appealing investment opportunities to maintain liquidity.

condition:

$$\frac{\partial U/\partial T}{\partial U/\partial Z} = \mu_0 = \mu_p. \tag{2}$$

When the household's beliefs about μ are correct ($\mu_0 = \mu_p$), the household's consumption of home energy services will be optimal from both *ex ante* and *ex post* perspectives, that is, the bundle the household would select after learning μ is the same one the household chooses before the marginal cost of energy services is revealed.

But what happens if the marginal cost of energy services is different than expected? Suppose the marginal cost turns out to be μ_1, which is greater than μ_0. In this case, depicted in scenario "B" in Figure 1, the *ex post* budget constraint pivots inward relative to the *ex-ante* constraint. Only when the household receives the next energy bill will it learn this, and therefore it does not have a chance to re-optimize its consumption of energy services, which remain at T_A. At the end of the period, the utility-maximizing household learns the *ex post* marginal cost and chooses z_B to satisfy the new (*ex post*) budget constraint and consumes bundle B, (T_B, z_B). At consumption bundle B with ex post relative prices μ_1, the household consumes too much T and e and the marginal rate of substitution between home energy services and market-produced goods is less than the marginal price ratio:

$$\frac{\partial U(T_A, z_B)/\partial T}{\partial U(T_A, z_B)/\partial z} < \mu_1 \tag{3}$$

When the marginal cost of energy services is higher than expected, the household is worse off in two ways. First, the household faces a higher marginal cost of energy services. This means the household pays more for space cooling, hot water, or lighting and is therefore poorer in real terms due to the income effect. Second, the household consumes too much energy services and too little other goods and services from a utility maximization standpoint at the higher marginal cost. Specifically, if the household had held accurate beliefs at the beginning of the billing cycle, they would have chosen the bundle represented by scenario C in Figure 1, which equalizes the ratio of marginal utility to prices across T and z, and corresponds to a level of utility, $U_C(T_C, Z_C)$, that exceeds $U_B(T_B, Z_B)$.

To evaluate the loss in utility from an unexpected change in μ, let $T(\mu, I)$ be the energy services a household with income I would demand *ex post* after learning the marginal cost

μ. Then, by substituting this demand function and the rearranged budget constraint for z into the utility function, the household's reduction in utility from being unable to optimize at the *ex post* relative price can be written as:

$$U(T_C, z_C) - U(T_B, z_B) = U(T(\mu_1, I), I - \mu_1 T(\mu_1, I)) - U(T(\mu_0, I), I - \mu_1 T(\mu_0, I)). \qquad (4)$$

In the right side of this expression, the second term is less than the first – implying a utility loss – because after marginal cost increases to μ_1, the household continues to consume energy services as if marginal cost is μ_0. This expression is the household's welfare loss stemming from its inability to track its energy consumption, which prevents adjustments to its consumption when there is a change in marginal cost of energy services.

2.3 High Bill Alerts as a Partial Solution to Welfare Loss from Incomplete Information

How does the household's consumption and utility change with an HBA program? When the household receives an HBA, it learns new information about the cost of energy services, which, for simplicity, we assume it can use to fully (i.e., accurately) update its beliefs.[15] To see this, assume, without loss of generality, each period has length of one, and an HBA is delivered at time t in (0, 1). At the beginning of the billing cycle, a household believes the marginal cost of energy services is μ_0 and consumes energy services at a rate that would lead to T_A total energy services over the entire billing cycle. In reality, the marginal cost of energy services is high, μ_1, and the household consumes more energy than normal, triggering an HBA. After delivery of the HBA, the household updates its beliefs and reduces its rate of consumption of energy services to a level that would accumulate to T_C over an entire billing period (because this level equalizes the ratio of marginal utilities to prices across T and z). In this case, the household's consumption of energy services is represented by scenario D in Figure 1, and entails a level of consumption of energy services, T_D, equal to $tT_A + (1-t)T_C$, which is less than T_A, and therefore energy consumption decreases relative

[15]We assume the electric utility did not inform the household it was enrolled in an HBA program, consistent with the HBA program we study. This simplifies the analysis because risk averse households unable to track their energy use over the billing cycle may adjust their consumption knowing the utility will monitor for and inform them of large increases.

to a no-alert counterfactual.[16] The household spends its remaining income on market goods and services $z_D = I - \mu_1(tT_A + (1-t)T_C)$.

The change in utility because of the HBA can be represented with the following expression,

$$
\begin{gathered}
U(T_D, z_D) - U(T_B, z_B) = \\
U(tT(\mu_0, I) + (1-t)T(\mu_1, I), I - t\mu_1 T(\mu_0, I) - (1-t)\mu_1 T(\mu_1, I)) - U(T(\mu_0, I), I - \mu_1 T(\mu_0, I)).
\end{gathered} \tag{5}
$$

This expression is positive due to the assumed strict convexity of the household's preferences and the strict quasi-concavity of its utility function.[17] While the household's utility in scenario D is less than it would be under scenario C, it is greater than under scenario B (which is the level that would have accrued absent an HBA) and therefore the HBA reduces the utility loss from monthly billing. Since the HBA increases the household's utility from U_B to U_D, the HBA's impact on consumer welfare can be measured as the dollar expenditures just needed to move the household from U_B to U_D. In Figure 1, this amount is shown by the difference between the two dotted budget constraints with slopes equal to μ_1 and tangent to the U_B and U_D indifference curves. This HBA consumer welfare impact equals $exp(U_D) - exp(U_B)$, where $exp(U)$ is the minimum expenditures needed to achieve utility U at the new, higher energy services marginal cost μ_1 and shown by where the budget constraint achieving utility U intersects the z (numeraire good) axis. In Section 6, we undertake a partial equilibrium analysis to estimate the HBA program consumer welfare benefit from alerting households to large consumption increases.

This simple model shows households incur welfare losses from monthly utility billing when they experience undetected changes in the marginal cost of energy services; and HBAs can improve consumer welfare by alerting households mid-billing cycle to these changes.

[16] Let e_B and e_C be energy consumption when the household consumes, respectively, T_B and T_C energy services. By the properties of the energy services production function, $e_B > e_C$. At bundle D, the household's consumption of energy is $e_D = te_B + (1-t) * e_C < e_B$, which means energy consumption falls after the HBA is delivered.

[17] Because the household's preferences are strictly convex, $U(T,z)$ is strictly quasi-concave. Strict concavity of the utility function implies $U(tT(\mu_0,I) + (1-t)T(\mu_1,I), I - t\mu_1 T(\mu_0,I) - (1-t)\mu_1 T(\mu_1,I)) > \min U(T(\mu_0,I), I - \mu_1 T(\mu_0,I)), U(T(\mu_1,I), I - \mu_1 T(\mu_1,I))$. Since $U(T(\mu_0,I), I - \mu_1 T(\mu_0,I)) < U(T(\mu_1,I), I - \mu_1 T(\mu_1,I))$, it must be that $U(t * T(\mu_0,I) + (1-t) * T(\mu_1,I), I - t * \mu_1 T(\mu_0,I) - (1-t) * \mu_1 T(\mu_1,I)) > U(T(\mu_0,I), I - \mu_1 T(\mu_0,I))$, that is, that the HBA lifts the household's welfare.

The theoretical predictions from the model are that the HBA program should lead to reductions in household demand for energy services and therefore decreased energy consumption and increased consumer utility. We now move on to discuss the HBA experimental design and the empirical analysis, in which we will examine the effect of the program on energy consumption.

3 Experimental Design and Data

In 2015 and 2016, the HBA program under study was deployed by a large vertically integrated, investor-owned electric and gas utility in the Midwest. The utility's motivation for running the program was to benefit households by enabling them to take pre-emptive steps toward lowering their bills and avoiding bill shocks and to reduce the volume of inquiries to its customer service center from households surprised by large bills. To cleanly identify the effects of the HBA program, it was implemented as a large, randomized natural field experiment. In coordination with a third-party program implementer, Opower, the utility randomly assigned about 50,000 households receiving gas and electric service to a treatment group and about 25,000 dual-fuel households to a control group in 2015.

Treatment group households were automatically enrolled in the program, making them eligible to receive the alerts based on their usage levels once the program was initiated. Treatment group households were not informed the utility had enrolled them in an HBA program and only became aware of the alerts when they received the first one. Households who received an alert could opt out of receiving future ones at any time by clicking on a link embedded in the alert email, though in practice very few households (about 1%) opted out. Households in the control group were not eligible to receive HBAs, were not informed of the experiment, and provided the baseline for measuring the energy impacts of the HBA. All treatment group and control group households received gas and electricity service from the utility, had a single meter per fuel type, were on standard (non-time-of-use) rates for gas and electricity, and were not enrolled in the utility's home energy reports program.

HBAs were first issued in June 2015. As noted above, an alert was sent mid-billing cycle via email to a customer when the customer's monthly electricity or gas consumption was on track to be 30% higher than the bill for the same period in the previous year and the bill was forecasted to exceed $30, with both estimates being derived using a proprietary algorithm from the third-party implementer. The alerts were based on gas-specific or electric-specific consumption patterns (e.g., a household's typical electricity usage) and contained gas-specific or electric-specific information (e.g., households were informed that their electricity bill was expected to be unusually large).

Figure 2 shows the numbers of electric and gas bill alerts sent in each month between June 2015 and June 2016 and the cumulative numbers of electric and gas alerts sent since the beginning of the experiment. Over the experiment's 13 months, the utility sent about 100,000 electric or gas alerts or an average of about two per treatment group household. The utility sent more electric than gas alerts, and accordingly, about 31,000 treatment group households received an electric alert, but only 6,000 received a gas alert. According to the program implementer, approximately 59% of the HBA emails sent between June 10, 2015 and October 31, 2015 were opened by the recipient, and 6.4% of HBA emails resulted in the recipient clicking the Home Energy Assessment link.

We collected monthly billing consumption data from June 2014 to May 2017 for all HBA program treatment group and control group customers. This period includes twelve months before treatment, thirteen months when the utility sent alerts, and eleven months after the utility stopped sending alerts. After this window, the utility shifted to universal deployment of the HBA program, which ended our window for evaluation. The availability of billing data for the year before treatment started provides information about the pre-treatment consumption of all households and allows us to estimate the HBA program treatment effects using difference-in-differences panel regression methods. Before analyzing the data, we dropped a small number of monthly observations that have zero electricity consumption or a bill length of forty days or longer (both of which comprise less than 0.1% of the data).

We approached outliers in the consumption data with circumspection. Outliers in the right tail of the distribution have the potential to bias estimates of the HBA program treatment effects, or limit the precision of the estimates, which might warrant dropping them from the analysis. However, to an extent, the HBA program is designed to target outlying large consumption levels, which argues for their inclusion. To be conservative, in our main analysis, we drop the bottom and top 1% of observations based on ADC of electricity and the top 1% of observations for ADC of natural gas (because 0 is a credible value for natural gas), and, as a robustness check, we re-run the analysis with the outliers.[18] As shown in the Appendix, the main results hold when the outliers are included and in some cases get notably stronger (see Tables A.2, A.3, and A.4 and Figure A.1). In total, in the primary data that exclude outliers, there are 4.3 million observations spanning 74,475 households, two-thirds of which are treatment households.

To check the appropriate randomization of the treatment, we compared the consumption of the treatment group and the control group in each pre-treatment billing month. Panel A and Panel B of Figure 3 show monthly means of daily electricity and daily gas consumption and estimates of the differences with 95% confidence intervals.[19] For both energy types, the means are closely aligned, and all differences are close to zero and statistically insignificant. We also performed a chi-square test of the difference between the treatment group and the control group in the distribution of households across zip codes and failed to reject the hypothesis that the distributions were different (p=0.46). Zip code locations are strongly correlated with household housing, demographic, and economic characteristics, so this result provides additional confidence about the validity of the randomization. Overall, we concluded the randomized treatment and control groups were well-balanced on consumption and location characteristics.

4 Econometric Overview

In this section, we present our empirical approach for estimating the impacts of the HBA program on electricity and gas consumption. For all parts of the analysis, we estimate panel regression models of household average daily consumption (ADC) of electricity and natural gas, and the models are estimated with monthly billing data. For some models, we divide the household's ADC by the ADC for the corresponding month in the pre-treatment

[18]In models that examine relative ADC variable, we go through a second stage of drops (i.e. drop the top and bottom 1% again after constructing the relative ADC variables) because the scaling process itself can create outlying values (for example, if the usage during the pre-period happened to be unusually small).

[19]Across the sample, mean ADC for electricity is 21.1 kWh/day (standard deviation = 12.97) and mean ADC for natural gas is 1.94 therms/day (standard deviation = 1.99).

period to create a new variable called "relative ADC."[20] These relative ADC models are estimated with only treatment period and post-treatment period data and provide a way to measure more directly the effect of the HBA program on avoiding "shocks"—i.e., changes in consumption that are different from each household's normal consumption patterns. All models weight observations by the number of days in the billing cycle and cluster standard errors by household.

Our baseline fixed-effects panel regression model, which captures the conditional mean effect of the program on usage, is specified as follows:

$$\text{ADC}_{it} = \beta_1 \text{Treatment}_i \times \text{Post}_t + \alpha_i + \gamma_t + \epsilon_{it} \tag{6}$$

where ADC_{it} is average daily consumption for household i in period t, Treatment_i is an indicator variable for assignment to the treatment group, Post_t is an indicator for the period after the HBA program was initiated, α_i represents a vector of household fixed effects, and γ_t represents a vector of month-of-sample fixed effects coded by the end of the household's billing cycle. The household fixed effects control for time-invariant differences in consumption between households, while the time-period fixed effects control for effects specific to each month and year. The coefficient β_1 indicates the HBA program average treatment effects during the period after the program was initiated. In addition to this model, we also estimate a specification that includes calendar-month-by-household fixed effects (i.e., twelve fixed effects for each household – one for each calendar month) to account for the seasonality in each household's energy consumption. We further estimate a model that buckets the post-treatment windows into two periods: one when the HBA program was active (June 2016-June 2016) and one for after it was discontinued (July 2016-May 2017).

To evaluate the dynamics of the treatment effects in more detail, we estimate an "event study" version of the fixed effects model that interacts the treatment group indicator variable with month-by-year indicator variables.[21] This specification yields an estimate of the

[20]Observations during the period after treatment was initiated are matched to pre-program observations that ended in the same calendar month. Relative ADC is computed by dividing nominal ADC by pre-program ADC.

[21]In these estimates, we only include one fixed effect per household so that we can non-parametrically examine trends leading into treatment. We could not do this with household-by-calendar-month fixed effects

HBA program treatment effect for each month-year of the sample except for the last month of the pre-treatment period, which is omitted, and shows how the program energy impacts evolve for the thirteen months during treatment and eleven months after treatment.

We also deploy quantile regression models that examine the effect of the HBA program at different segments of the usage distribution to estimate whether the program helped households avoid unusually large bill shocks. Specifically, we estimate unconditional quantile regressions (Firpo et al, 2009),[22] which provides estimates of treatment effects at different segments of the consumption distribution.[23] While we estimate quantile models with ADC as the dependent variable, our primary focus is on models that use relative ADC. The reason for this is that, if the effects are large at the top of the distribution in the nominal ADC model, the mechanism could be either a) large users being the most likely to respond to the HBA program or b) representative users avoiding unusually large bills relative to their typical usage levels. In contrast, in the relative ADC models, if larger effects are evident at the top of the distribution, the mechanism must be that households avoided unusually large bills relative to their historical consumption levels. In other words, the relative ADC models are best-suited for testing whether the HBA program allowed households to avoid bill shocks.

Lastly, we estimate models that allow for heterogenous program effects based on the median income in the household's zip code to see if the effect was stronger in communities with high or low incomes. Bill shocks can have the most serious effects on poor households, which may be forced to cut back on expenditures for energy and other essential goods and services to pay utility bills. We estimate heterogeneous income effects by linking each household to its community's median income household based on the household's zip code

because we only have one year of pre-treatment data.

[22]The unconditional quantile regressions were estimated by first transforming the dependent variable using a recentered influence function (RIF) and then running an OLS regression of the transformed data on a treatment indicator, household-by-calendar month fixed effects, and month-by-year fixed effects. See Rios-Avila (2020) for additional details.

[23]An alternative to the unconditional quantile model is the conditional quantile model (Koenker and Bassett, 1978). Both conditional and unconditional quantile regressions control for covariates. The unconditional versus conditional distinction refers to whether the relevant margin in the distribution is determined by the distribution of the dependent variable or the distribution of the dependent variable after netting out the role of covariates. Unconditional quantile models have simpler, and generally more meaningful interpretations than conditional quantile regression models. For brevity, for the remainder of the manuscript, we will just write "quantile" models/estimates – all of which are unconditional.

and five-year median income estimates for zip code tabulation areas (ZCTA) from the 2016 U.S. Census' American Community Survey. There are 100 zip codes represented in our sample with median incomes ranging from about \$30,000 to over \$100,000, with a mean across these medians of about \$67,000. To estimate heterogenous effects, we include an interaction of $\text{Treatment}_i \times \text{Post}_t$ and income in the regression models, and to ease interpretation, we de-mean the income variable using the overall sample mean and scale it by \$10,000s.

5 Results

5.1 Estimating HBA Program Effects on Mean Usage: Overall, by Program Stage, and in an Event Study

We begin by estimating the conditional mean effect of the HBA program on consumption. For each energy type, we deploy three specifications, starting with a basic D-in-D model. The second model incorporates month-by-year and household fixed effects, and the third model, our preferred specification, includes both month-by-year fixed effects and household-by-calendar month fixed effects. The models capture the mean effect of the HBA program across all twenty-four months after the program was initiated, including thirteen months when alerts were actively being released and eleven months after alerts were discontinued. We report the results in Table 1. For electricity, the main treatments effects are stable across specifications and our preferred estimates indicate that the program decreased electricity consumption by .10 kWh or 0.5% of consumption. For gas, treatment effects are also generally stable across specifications, although the inclusion of household fixed effect(s) matters more than in the electricity models, and our preferred estimates indicate a saving of .009 therms, or 0.5% of consumption.

It is notable that, although the utility sent many fewer gas than electric alerts, the HBA program caused similar percentage reductions in gas consumption as electricity. The approximate equality of the percentage treatment effects for electricity and gas despite many fewer households having received gas alerts could be explained by it being comparatively easier or less costly for households to adjust gas than electricity consumption because natural gas is used for a relatively small number of home end uses (space heating, water heating, and cooking) or by households not differentiating between electric and gas alerts and responded to either by reducing consumption of both fuels.

The mean HBA program HBA program's effects in Table 1 are estimated across the treatment and post-treatment periods, but the treatment effects may have changed after the utility stopped sending HBAs. For example, if treated households increased energy conservation behaviors in response to the alerts, they may have experienced backsliding and the treatment effects might have diminished during the post-treatment period (Allcott and Rogers, 2014; Brandon et al., 2022). Table 2 presents models that estimate the effect of the HBA program separately for the period when it was active versus after it had been discontinued. Across both fuel types, the estimated effects are statistically indistinguishable during the period when alerts were active versus after they had been discontinued. These results indicate that the effects of the program persisted in the eleven months after the alerts had been discontinued, potentially due to either persistent behavioral changes or durable investments in energy efficiency.

To examine program dynamics in more detail, we turn to our event study model, which estimates HBA treatment effects for each pre-treatment month (except May 2015, the excluded month), each treatment month from June 2015 to June 2016, and each post-treatment month from July 2016 to May 2017. Since we estimate treatment effects for each month, the confidence intervals are relatively wide, and many intervals include zero. Nonetheless, the patterns in the point estimates are still useful for understanding the program dynamics.

We begin with electricity, which is presented in Figure 4.1. During the pre-treatment period, the estimates were close to zero and statistically insignificant. Once the program began, the effects steadily increased, which is consistent with more households having received at least one alert as the program matured, thereby triggering conservation. After the alerts were terminated, the effects persisted, which aligns with the estimates provided in Table 2, although there is some minor visual evidence that they were tapering off toward the end of the sample.

Turning to gas, which is presented in Figure 4.2, seasonality is evident, as the coefficients display relative increases during winter months. These increases likely stem from sampling noise in the randomization process that becomes most evident during periods of elevated usage due to winter-time space heating. Focusing on the winter months, there is evidence that usage was higher in treatment group households before the program began, but then decreased in subsequent winters both when the HBA program was active (2015) and after it had been discontinued (winter 2016). There is less evidence of an effect during the non-winter months. Collectively, the evidence here is less conclusive, but the pattern in the point estimates is consistent with previous estimates and the electricity event study: the HBA program decreased gas consumption, especially during winter months when usage is greatest, and this decrease persisted after the HBA program was discontinued. The persistence of HBA program energy impacts during the post-treatment period is a notable finding because the effects of behavioral programs in the energy and environmental sectors are often fleeting (Brandon et al., 2022; Allcott and Rogers, 2014; Gilbert and Graff Zivin, 2014; Ferraro and Price, 2013; Jacobsen, 2011).

5.2 Quantile Estimates: Did the HBA Program Help Households Avoid Bill Shocks?

Panel A and Panel B of Table 3 show estimates of the HBA program impacts for different percentiles of the electricity and gas consumption distributions based on quantile regression models. For electricity, the HBA program did not have statistically significant impacts below the 50th percentile of the consumption distribution. However, between the 50th and 90th percentiles, statistically significant effects emerge, and the effects increase toward the tail of the distribution. For example, at the 90th percentile, the HBA program treatment effect is -0.293 kWh per treatment group household per day, which is three times larger than the effect on mean consumption levels. Above the 90th percentile, the estimates are noisy and insignificant. Similarly, for natural gas, the program has statistically significant, negative, and, in general, progressively greater impacts for households above the 25th percentile.[24] Thus, for both fuels, HBAs primarily affect bills involving larger amounts of consumption. As mentioned earlier, these effects could be driven either by large-use households responding most to the HBA program or households across the distribution taking actions to avoid

[24]We are not able estimate effects at the bottom of the natural gas distribution because all values are zeroes.

unusually large bills *relative to their typical usage patterns* – i.e., "bill shocks." To focus on the avoidance of bill shocks, we turn to quantile estimates that use relative ADC as the dependent variable.

Table 4 presents estimates from unconditional quantile regression models that use a household's relative energy consumption as the dependent variable. In these models, large effects are especially evident at the tail of the distribution. For electricity, significant effects begin to emerge around the 25th percentile and trend upward all the way to the 99th percentile. The estimate from the 99th percentile model, which is eight times larger than the estimate at the median, indicates that relative ADC is 2.3 percentage points lower at the 99th percentile of the relative ADC distribution for households enrolled in the HBA program. The natural gas patterns exhibit a similar trend. While the coefficients become insignificant at the tail of the distribution, the point estimates show a steadily increasing trend as the percentile of the distribution increases, capping at the 99th percentile at 1.3 percentage points. These results indicate that the HBA program was effective at reducing the extent to which households experienced bill shocks from elevated usage levels.

It should be noted that the estimates of the stronger effects at the tail of the distribution are conservative because, as discussed earlier, we drop outliers from the main analysis. When we include outliers, the stronger relative effect at the tail of the distribution is especially prominent. As we show in Table A.4, the point estimates for the relative ADC models at the 99th percentile are -11.860 for electricity and -15.33 for natural gas. These are *36* and *44* times stronger than the effects at the median in the same models.

5.3 HBA Program Treatment Effects by Household Income

As described above, the value of the HBA program from a public policy perspective may be enhanced if it were particularly effective in low-income communities. To examine that possibility, we estimate models analogous to those in Table 1, except we include an interaction of the $\text{Treatment}_i \times \text{Post}_t$ with median zip-code level income. We report the results in Table 5. Across specifications, the interaction term is insignificant, which suggests that the effectiveness of the program did not vary by income. We also estimate analogous relative ADC quantile regression models that include a treatment by income interaction term as well. The estimates are reported in the Appendix (Table A.5) and are also insignificant. We conclude that there is little evidence that the HBA program had weaker or stronger effects in lower income areas, suggesting low-income households shared in the benefits of the program. A potential explanation for this null result is that while low-income households may have the most to gain from avoiding expenditure shocks, they may also find it more difficult to undertake conservation, as reflected in the lower rates of efficiency in lower-income households (Drehobl et al., 2020).

5.4 The Effect of an Alert

Before proceeding to estimating welfare benefits, it is worth reiterating, that we are estimating the effect of the HBA program as opposed to the effect of an alert itself. This is partly because the effect of the program may be of primary interest due to program offering being the margin that policymakers and utility managers can influence or control. However, the effect of an alert itself may still be of interest because it enhances understanding of behavioral responses from households in the energy sector.

We cannot directly estimate the effects of alerts by using an alert indicator as the treatment variable because alerts are endogenously triggered, but we can bound the effects through a scaling procedure.[25] In particular, if we assume all alert effects are concentrated in the month in which the alert was issued, we can back out the mean effect of an alert by dividing our preferred estimates of program effects – reported in columns 3 and 6 of Table 1 for gas and electricity, respectively – by the proportion of observations for treated households after the HBA program was initiated that correspond to a billing cycle when an alert was issued. We scale separately for each fuel, reflecting that 9.9% of post-period treatment household billing cycles had an electricity consumption alert and 1.2% of post-period treatment household billing cycles had a gas consumption alert. This scaling procedure produces an upper-bound estimate of the effect of an electricity alert of 1.03 kWh (4.8%) and an upper bound estimate of the effect of a gas alert of .75 therms (41.2%). In contrast, if we assume the effect of an alert occurred during all billing cycles during or after the first alert for each

[25]The assumption in our scaling procedure is that all HBA program effects occur through the issuance of alerts. We believe this is a reasonable assumption in part because the utility did not inform households that they were enrolled in the HBA program, so we can rule out a Hawthorne effect.

energy type was issued to each household, we can back out program effects by scaling by the proportion of post-period treatment household billing cycles occurring on or after each household's first alert (58% of observations for electricity and 8.3% of observations for natural gas). In this case, the scaling procedure produces a lower-bound estimate of the effect of an electricity alert of .17 kWh (0.82%) and a lower bound estimate of the effect of a gas alert of .75 therms (5.9%). The range of these estimates, especially for electricity (0.82%-4.8%), is in alignment with other behavioral responses in the energy sector, such as consumption responses to home energy reports (Allcott and Rogers, 2014) or consumption responses to enrolling in a green electricity program (Jacobsen et al., 2012). The relatively large response for gas (5.9%-41.2%) may be explained by fewer gas alerts being issued in combination with the potential for the effect of electricity alerts to spillover onto gas consumption, or by households being more responsive to gas alerts because they can more easily identify the source of the increase in consumption because gas is used for fewer end-uses.

6 Consumer Welfare Benefits

In the conceptual model of Section 2, we showed an HBA can lift consumer welfare by providing information mid-billing cycle to a household about an increase in its energy consumption. In this section, we use partial equilibrium analysis to estimate the consumer welfare benefit of the HBA program from the delivery of billing alerts. We measure the benefits of the HBA program for consumers as the difference between their avoided expenditures and the foregone value of the conserved energy. Calculating the foregone value is complicated by the likelihood that HBA recipients were initially consuming more energy than would be optimal, at least under full information (Chetty, 2009). In a simplified economic model with complete knowledge of the costs and benefits of energy consumption, the value of the last unit consumed by a household would be equal to the price of the good. If the household conserved energy from this starting point, the foregone value of the first unit of conservation would be equal to the price of energy and each conserved unit thereafter would be slightly decreasing in foregone value. However, if HBA recipients were consuming more energy than would be optimal under full information – potentially because the cost of acquiring information was too high absent the HBA program – then the value of conserved energy could be lower than the price and the net benefit of conservation induced by the HBA program would be greater.

Uncertainty about the foregone value of the conserved energy induced by the HBA program makes it difficult to obtain a point estimate of the benefit of the HBA program to consumers; however, we can bound the impact, which we do by calculating estimates for two scenarios. We estimate a lower bound by assuming the value of the last unit of consumption was exactly equal to the price and then use that, combined with an estimate of market demand elasticities, to calculate the welfare effect. We estimate an upper bound by making the other extreme assumption: the foregone consumption was wasteful and had no value to consumers.

Figure 5 illustrates the intuition behind our approach and displays the inverse market demand curve, $p(e)$, for utility-supplied energy for households and the unit market price for energy, p. We begin with the assumption that the marginal value of the first unit of conserved energy is exactly equal to the price of energy, p_0, which corresponds to a level of consumption equal to e_1. As we know from our empirical estimates, after receiving an HBA, households reduce their energy consumption and we represent that level of consumption with e_0. Graphically, the avoided expenditures are represented by the area *abed* and the foregone value from reduced energy consumption is the area *aced*. The difference is the welfare gain equal to *abc*.[26]

In practice, because the HBA program was designed to detect unusual usage patterns, the alerts may have triggered conservation of wasteful, low-value energy, thereby lifting the gains from conservation. Returning to Figure 5, we depict this conservation as the difference between e_3 and e_2. In this scenario, the welfare gain from reducing consumption by Δe is equal to *fghi*, which is a much larger gain than under the previous scenario due to the lower value of the foregone energy consumption. In an extreme scenario, where the conserved energy had no value to the household (because it was all being used wastefully), the value of the foregone consumption would be the entire expenditure savings, which is graphically equivalent to a rectangle with a high of p_0 and a width equal to the amount of conservation. This area is represented by either *abed* or the equivalently size rectangle that begins with

[26]We assume consumers have quasi-linear utility over energy and other market goods.

the line between f and g and extends down to the x-axis.

We now proceed to calculate bounds for the private consumer welfare benefits. We obtain our lower bound estimate by assuming all HBA program conservation was marginal to the price of energy. In this case, households that conserve in response to HBAs have a foregone value of energy consumption that begins at a level equal to the price of energy and decreases slightly with each unit of conservation. The lower bound can be approximated as:

$$\underline{\Delta W} \approx p_0(e_1 - e_0) - p_0(e_1 - e_0)[1 + (e_1 - e_0)/2\varepsilon e_0] \tag{7}$$

We derive this expression by first approximating the inverse market demand curve p(e) at e_0 using a Taylor's series expansion and then rewriting the approximation as a function of the short-run price elasticity of electricity demand ε and the HBA program energy conservation, $e_1 - e_0$. Second, we obtain an upper bound estimate of the HBA welfare impact by assuming all HBA-induced energy conservation had no foregone value to households. In this case, the gain per unit of conserved energy would equal the price, p_0, and the welfare gain in total would equal $\overline{\Delta W} \approx p_0(e_1 - e_0)$.[27]

To capture the effect of the program at scale, welfare impacts of the HBA program are calculated for 2021 when the utility had made all residential customers eligible to receive HBAs. We assume a short-run price elasticity of demand for electricity and gas equal to -0.1 based on estimates from the literature (Labandeira, Labeaga, and Lopez-Otero, 2017; Burke and Abayasekara, 2018) and an HBA treatment effect equaling 0.5% of consumption per the electricity and gas impact estimates in Table 1. We obtain electricity and gas retail prices, counts of the residential customer populations, and annual energy consumption per customer for 2021 from EIA Form 861 (electricity) and EIA Form 176 (natural gas).

Table 6 reports welfare estimates. Aggregate upper bound estimates – which again are average price times average energy conservation times the number of customers – are about \$6.5 million for electricity, \$2.5 million for natural gas, and \$9 million annually overall. The

[27]The welfare calculations are based on the average energy price, as opposed to the marginal energy price. There are three reasons for this. First, there is evidence that consumers mostly respond to average prices, as opposed to marginal prices (Ito, 2014). Second, over time, reduced expenditures from energy conservation will be translated to reduced bills from consumers based on average prices because the rate-making process requires that regulated utilities recover their costs. Third, we can observe average prices in the available data.

lower bound estimates are much smaller. For electricity, the lower bound estimate at our preferred elasticities equals about \$161,000, and for natural gas, it is about \$62,000. The combined total is about \$223,000, although it falls considerably with small increases in the chosen elasticities.

The utility provided us with its cost of administering the HBA program, which includes fees paid to a third-party program implementer. Based on this information, we estimate the utility's cost of administering the program for one year to be about \$200,000, which, due to regulation, can be expected to be passed down to ratepayers in the form of elevated energy charges.[28] The cost estimate falls about equal to our lower bound estimates of the consumer benefits, and potentially exceeds the benefits if we choose elasticities that exceed .10 in absolute value.[29] However, due to the design of the HBA program, we think it likely induces conservation of low-value energy, and therefore the program's benefits are probably larger than its costs given how close our lower bound estimate falls to our estimated program costs.[30] The range in the estimates underscores the point that the performance of behavioral programs in the utility sector depends critically on the consumer's valuation of the avoided energy consumption.

7 Conclusion

Providing feedback to households on their energy usage holds the promise of creating better outcomes in the energy sector. This paper provides the first empirical evidence that we are aware of regarding high bill alert programs targeted at energy consumers, which are a new tool that utilities can deploy to help households avoid unexpected bill shocks. We find that

[28]This is our best estimate, but there is some ambiguity in how this number should be calculated depending on assumptions about start-up program costs, discounted pricing if the program is packaged as part of a suite of DSM program, and the salary of workers dedicated to operate the HBA program.

[29]Our analysis focuses on consumer welfare. We do not account for any impacts of HBAs on other margins, such as externalities in the form of reduced air pollution emissions from electric power plants or home combustion of natural gas.

[30]A potential omission from our analysis is that we do not consider the possibility that HBAs can cause disutility for certain consumers. For example, Allcott and Kessler (2019) find significant heterogeneity in welfare gains between recipients of home energy reports, as some recipients would have preferred not to receive them. Disutility from receiving HBAs may be less of a concern, however, because, unlike home energy reports, HBAs were targeted at utility customers who experienced large consumption increases and were likely beneficiaries from the alerts.

the HBA program led to mean conservation of about one-half of a percentage point for program participants for both electricity and natural gas consumption, and the conservation persisted for at least one year after the utility ceased sending alerts. Effects were concentrated at the top of the usage distribution – especially when measuring usage relative to each household's pre-program level – indicating that the program was helpful at allowing consumers to avoid unusually large bills.

We evaluate the benefits of the HBA program to consumers and compare them to the costs of administering the program and show that whether the benefits exceed the costs is sensitive to assumptions regarding the foregone value of the conserved energy. This result highlights an important consideration in consumer welfare analysis: the benefits of behavioral programs are a function of the behavioral mechanism by which the primary objective is achieved. If a behavioral program is aimed at encouraging socially-beneficial behavior, there can be added value if the program simultaneously corrects individual-level inefficiencies, such as those caused by incomplete information, rational inattention, or cognitive biases. With respect to energy conservation, which is a goal of many programs in the energy sector, if conservation from a program is achieved by a reduction in consumption of valuable energy services, such as conservation from setting the thermostat at a less comfortable temperature, then the net effects of the program are likely to be of less value than an alternative program – such as the one under study – that yields the same amount of energy conservation but does so by targeting wasted energy due to mechanical breakdowns or inefficiencies in home heating and cooling equipment.

Several questions are prompted by the results in this paper. One is: how will the effect of HBA programs on energy consumption and consumer benefits change as the energy system evolves? While speculative, it may be reasonable to project that the effects of HBA programs will get stronger in the coming years. Changes in the electricity system like time-varying prices, increased at-home charging of electric vehicles, and increasing electrification of energy-using durable goods are likely to add volatility to household bills, thereby creating more opportunities for HBA program alerts to trigger conservation.

More generally, our research raises questions about the ways in which the massive amounts of data being generated and collected on consumer behavior can be used to improve outcomes for these consumers. For example, in the electricity sector, customers historically have received information on their consumption through a monthly bill. With smart meters, there is now the potential to provide customers with nearly continuous information on their usage. On the one hand, providing customers with frequent data on their usage can enhance the amount of information that they have access to when making consumer decisions. On the other hand, overly frequent information provision can overwhelm consumers and reduce the salience of the information households receive. As the amount of data collected by organizations expands, the frequency versus salience tension is likely to be relevant to other areas of the economy as well, such as with data usage alerts for phones, financial alerts for possible fraud based on credit card expenditures, car system warnings to drivers about inattentiveness, or medical alerts about changes in health markers. Each area is likely to have its own unique features that influence how consumers respond to information and how alert-based programs affect their welfare, and we look forward to further research that informs how to optimize the delivery of information in the energy sector and beyond.

8 Appendix

1. Estimates of the Association between Receiving an Alert and Mean ADC (Table A.1)
2. Reproducing Results from Main Analysis with Sample that Includes Outliers (Tables A.2, A.3 and A.4)
3. Quantile Regression Estimates for Relative ADC with Income Interaction (Table A.5)

9 References

Allcott, H. (2011). "Social Norms and Energy Conservation." *Journal of Public Economics* 95(2) 1082-1095.

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

Allcott, H. and Greenstone, M. (2012). "Is There an Energy Efficiency Gap?" *Journal of Economic Perspectives* 26(1) 3-28.

Allcott, H, and Kessler, J.B. (2019). “The Welfare Effects of Nudges: A Case Study of Energy Use Social Comparisons.” *American Economic Journal: Applied Economics* 11(1) 236-76.

Allcott, H, and Knittel, C. (2019). “ Are Consumers Poorly Informed about Fuel Economy? Evidence from Two Experiments.” *American Economic Journal: Economic Policy* 11(1): 1–37.

Allcott, H. and Rogers, T. (2014). “The Short-Run and Long Run-Effects of Behavioral Interventions: Experimental Evidence from Energy Conservation.” *American Economic Review* 10 3003-37.

Ayres, I., Raseman, S., Shih, A. (2013). “Evidence from Two Large Field Experiments that Peer Comparison Feedback Can Reduce Residential Energy Usage.” *The Journal of Law, Economics, and Organization* 29(5) 992–1022.

Blonz, J., Palmer, K., Wichman, C., and Wietelman, D., (2018). “Automation, smart thermostats, and time-varying pricing.” RFF Working Paper 21-20.

Bollinger, B.K. and Hartman, W.R. (2020). “Information versus Automation and Implications for Dynamic Pricing,” *Management Science* 66(1) 290-314.

Borenstein, S. (2012). “The Private and Public Economics of Renewable Electricity Generation.” *Journal of Economic Perspectives* 26(1) 67-92.

Borenstein, S. and Kellogg, R. (2023). “Carbon Pricing, Clean Electricity Standards, and Clean Electricity Subsidies on the Path to Zero Emissions.” *NBER Environmental and Energy Policy and the Economy* 4 125-176.

Borenstein, S. and Davis, L. (2016). “The Distributional Effects of US Clean Energy Tax Credits,” *NBER Tax Policy and The Economy* 30: 191-234.

Brandon, A., List, J.A., Metcalfe, R.D., Price, M.K., and Rundhammer, F. (2018) “Testing for crowd out in social nudges: Evidence from a natural field experiment in the market for electricity.” *Proceedings of the National Academy of Sciences* 116(12) 5293-5298.

Brandon, A., Ferraro, P.J., List, J.A., Metcalfe, R.D., Price, M.K., and Rundhammer, F. (2022). “Do the Effects of Nudges Persist? Theory and Evidence from 38 Natural Field Experiments.” NBER Working Paper 23277.

Burke, P.J., and Abayasekara, A. (2018). “The Price Elasticity of Electricity Demand in the United States: A Three-Dimensional Analysis.” *The Energy Journal* 39(2) 123–146.

Burkhardt, J., Gillingham, K., and Kopalle, P.K. (2019). “Experimental Evidence on the Effect of Information and Pricing on Residential Electricity Consumption.” NBER Working Paper No. 25576.

Chetty, R. (2009). “The Simple Economics of Salience and Taxation.” NBER Working Paper No. 15246.

Davis, L.W. (2008). "Durable Goods and Residential Demand for Energy and Water: Evidence from a Field Trial." *RAND Journal of Economics* 39(2) 530-546.

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

Drehobl, A. and Ross, L. (2016). "Lifting the High Energy Burden in America's Largest Cities: How Energy Efficiency Can Improve Low Income and Underserved Communities." American Council for an Energy-Efficient Economy. Online at: `https://www.aceee.org/sites/default/files/publications/researchreports/u1602.pdf`

Drehobl, A., Ross, L., and Ayala, R. (2020). "How High Are Household Energy Burdens?" The American Council for an Energy-Efficient Economy. Online at: `https://www.aceee.org/research-report/u2006`

Dubin, J.A. and McFadden, D.L. (1984). "An Econometric Analysis of Residential Electric Appliance Holdings and Consumption." *Econometrica* 52(2) 345-362.

Edison Foundation (2022). "Smart Meters at a Glance." The Edison Foundation Institute for Electric Innovation. Online at: `https://www.edisonfoundation.net/-/media/Files/IEI/publications/IEI-Smart-Meters-Infographic_2022.pdf`

Energy.Gov. (2010). "Meters Will Help Customers Avoid High Electric Bills." U.S. Department of Energy - VP 100. Online at: `https://www.energy.gov/articles/vp-100-smart-meters-will-help-customers-avoid-high-electric-bills`

Fabra, N., Rapson, D., Reguant, M., and Wang, J. (2021). "Estimating the Elasticity to Real-Time Pricing: Evidence from the Spanish Electricity Market." *AEA Papers and Proceedings*, 111: 425-29.

Ferraro, P. J. and Price, M. K. 2013. "Using Nonpecuniary Strategies to Influence Behavior: Evidence from a Large-Scale Field Experiment," *Review of Economics and Statistics* 95(1) 64-73.

Firpo, S., Fortin, N.M., and Lemieux, T. (2009). "Unconditional Quantile Regressions." *Econometrica* 77(3) 953-973.

Gabaix, X. (2017). "Behavioral Inattention." NBER Working Paper 24096.

Gilbert, B. and Graff Zivin, J. (2014). "Dynamic Salience with Intermittent Billing: Evidence from Smart Electricity Meters." *Journal of Economic Behavior and Organization* 107 176-190.

Gillan, J.M. (2018). "Dynamic Pricing, Attention, and Automation: Evidence from a Field Experiment in Electricity Consumption" Working Paper.

Harrison, G. and List, J.A. (2004). "Field Experiments." *Journal of Economic Literature* 42 1009-1055.

Horrigan, J. and Satterwhite, E. (2020). "Americans' Perspectives on Early Termination Fees and Bill Shock." Federal Communications Commission. Online at: `https://www.fcc.gov/general/bill-shock`

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

Jack, K. and Smith, G. (2020). "Charging Ahead: Prepaid Metering, Electricity Use, and Utility Revenue." *American Economic Journal: Applied Economics* 12 (2): 134-68.

Jacobsen, G.D. (2011). "The Al Gore Effect: An Inconvenient Truth and Voluntary Carbon Offsets." *Journal of Environmental Economics and Management* 61 67-78.

Jacobsen, G.D. (2015). "Do Energy Prices Influence Investment in Energy Efficiency? Evidence from Energy Star Appliances." *Journal of Environmental Economics and Management* 74 94-106.

Jacobsen, G.D. (2019). "An Examination of How Energy Efficiency Incentives Are Distributed Across Income Groups." *The Energy Journal* 40(6): 171-198.

Jacobsen, G.D. (2023). "Race, Ethnicity, and the Distribution of Energy Efficiency Incentives." *Energy Economics* 130 106933.

Jacobsen, G.D., Kotchen, M.J., and Clendenning, G. (2014). 'The behavioral response to voluntary provision of an environmental public good: Evidence from residential electricity demand." *European Economic Review* 56(5) 946-960.

Jacobsen, G.D. and Stewart, J.I. (2022). "How do consumers respond to price complexity? Experimental evidence from the power sector." *Journal of Environmental Economics and Management* 116 102716.

Jessoe, K. and Rapson, D. (2014). "Knowledge is (Less) Power: Experimental Evidence from Residential Energy Use." *American Economic Review* 104(4) 1417-1438.

Kahneman, D. and Tversky, A. (1979). "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47(2) 263-292.

Koenker, R. and Bassett, Jr. G. (1978). "Regression Quantiles." *Econometrica* 46(1) 33-50.

Kontokosta, C., Reina, V.J., Bonczak, B. (2020). "Energy Cost Burdens for Low-Income and Minority Households," *Journal of American Planning Association* 86(1) 89-105.

Labandeira, X., Labeaga, J.M., and López-Otero, X. "A Meta-Analysis on the Price Elasticity of Energy Demand," *Energy Policy* 102 549-568.

Lazar, J. (2016). "Electricity Regulation in the US: A Guide." The Regulatory Assistance Project. Online at: `https://www.raponline.org/knowledge-center/electricity-regulation-in-the-us-a-guide-2/`

item[]List, J., Metcalfe, R., and Price, M. (2018). Smart Meters: Do Prices Matter to Their Adoption and Do They Save Energy?," working paper.

Litos Strategic Communication. (2008). “What the Smart Grid Means to Americans.” Report prepared for the U.S. Department of Energy under Contract No. DE-AC26-04NT41817. Online at: https://www.energy.gov/oe/articles/what-smart-grid-means-americans

Martin, S. and Rivers, N. (2018). “ Information Provision, Market Incentives, and Household Electricity Consumption: Evidence from a Large-Scale Field Deployment.” *Journal of the Association of Environmental and Resource Economists* 5(1) 207-231.

Matisoff, D.C., Beppler, R., Chan, G., and Carley, S. (2020). “A review of barriers in implementing dynamic electricity pricing to achieve cost-causality,” *Environmental Research Letters* 15 093006.

Mission Data (2022). “Deactivated: How Electric Utilities Turned Off the Data-Sharing Features of 14 Million Smart Meters.” Mission Data Coalition. Online at: http://www.missiondata.io/s/Deactivated_white_paper.pdf

Neveu, A.R. and Sherlock, M.F. (2016). “An Evaluation of Tax Credits for Residential Energy Efficiency,” *Eastern Economic Journal* 42: 63-79.

Pollak, R.A. and Wachter, M.L. (1975). “ The Relevance of the Household Production Function and Its Implications for the Allocation of Time.” *Journal of Political Economy* 83(2) 255-278.

Rios-Avila, F. (2020). “Recentered in influence functions (RIFs) in Stata: RIF regression and RIF decomposition.” *The Stata Journal* 20(1) 51-94.

Rubin, E. and Auffhammer, M. (2023) “Quantifying heterogeneity in the price elasticity of residential natural gas.” *Journal of Association of Environmental and Resource Economists*, in press.

Sallee, J.M. (2014). “Rational Inattention and Energy Efficiency.” *Journal of Law and Economics* 57(3) 781-820.

Schneider, I. and Sunstein, C.R. (2017). “Behavioral considerations for effective time-varying electricity prices,” *Behavioural Public Policy* 1(2) 219-251.

SmartGrid.Gov. (2023). “Recovery Act Smart Grid Program.” U.S. Department of Energy - SmartGrid. Online at: https://www.smartgrid.gov/recovery_act/

Tuttle, C.J. and Beatty, T.K.M. (2017). “The Effects of Energy Price Shocks on Household Food Security in Low-Income Households.” Economic Research Report 233. Economic Research Service.

U.S. Energy Information Administration. (2021). “Major Utilities’ Spending on the Electric Distribution System Continues to Increase.” Online at: https://www.eia.gov/todayinenergy/detail.php?id=48136

U.S. Energy Information Administration. (2022). "EIA Expects Higher and More Volatile Wholesale Electricity Prices in the United States this Winter." Online at: https://www.eia.gov/pressroom/releases/press522.php

WhiteHouse.Gov. (2022). "Biden-Harris Administration Races to Deploy Clean Energy that Creates Jobs and Lowers Costs." The White House. Online at: https://www.whitehouse.gov/briefing-room/statements-releases/2022/01/12/fact-sheet-biden-harris-administration-races-to-deploy-clean-energy-that-creates-jobs-and-lowers-costs/

10 Tables and Figures

10.1 Main Text Tables and Figures

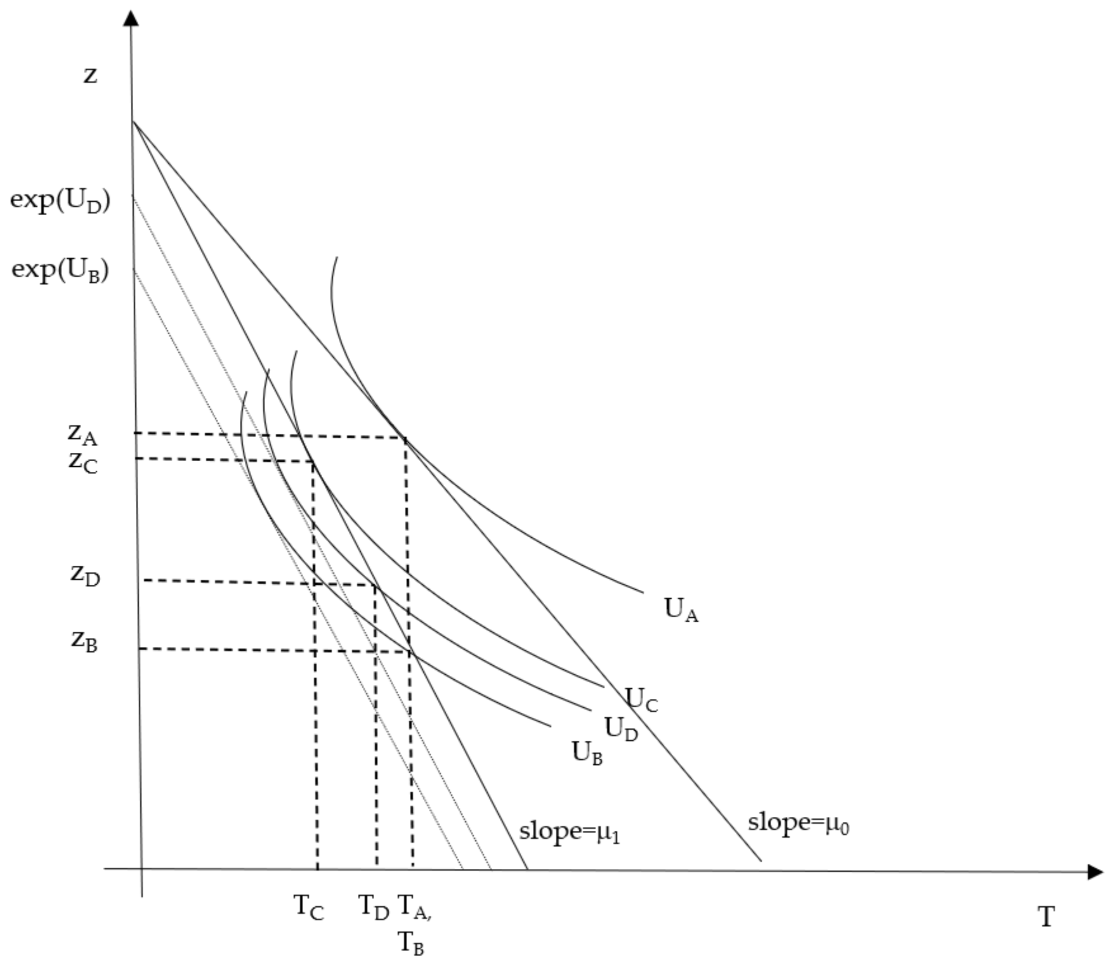

Figure 1: **Household Consumption Bundles and Utility Under Different Scenarios.** T represents consumption of energy services and z represents a composite market good (with price normalized to one). μ_0 represents the *ex ante* expected costs of energy services. Consumption bundle A represents consumption when the *ex ante* expected costs of energy services are realized. Consumption bundle B represents consumption when the *ex post* expected costs of energy services are elevated to μ_1 and an HBA program *is not* active. Consumption bundle C represents the optimal consumption bundle when the *ex post* expected costs of energy services are elevated to μ_1. Consumption bundle D represents consumption when the *ex post* expected costs of energy services are elevated to μ_1 and an HBA program *is* active. The U_D indifference curve exceeds the U_B indifference curve, reflecting the benefits to consumers from the HBA program.

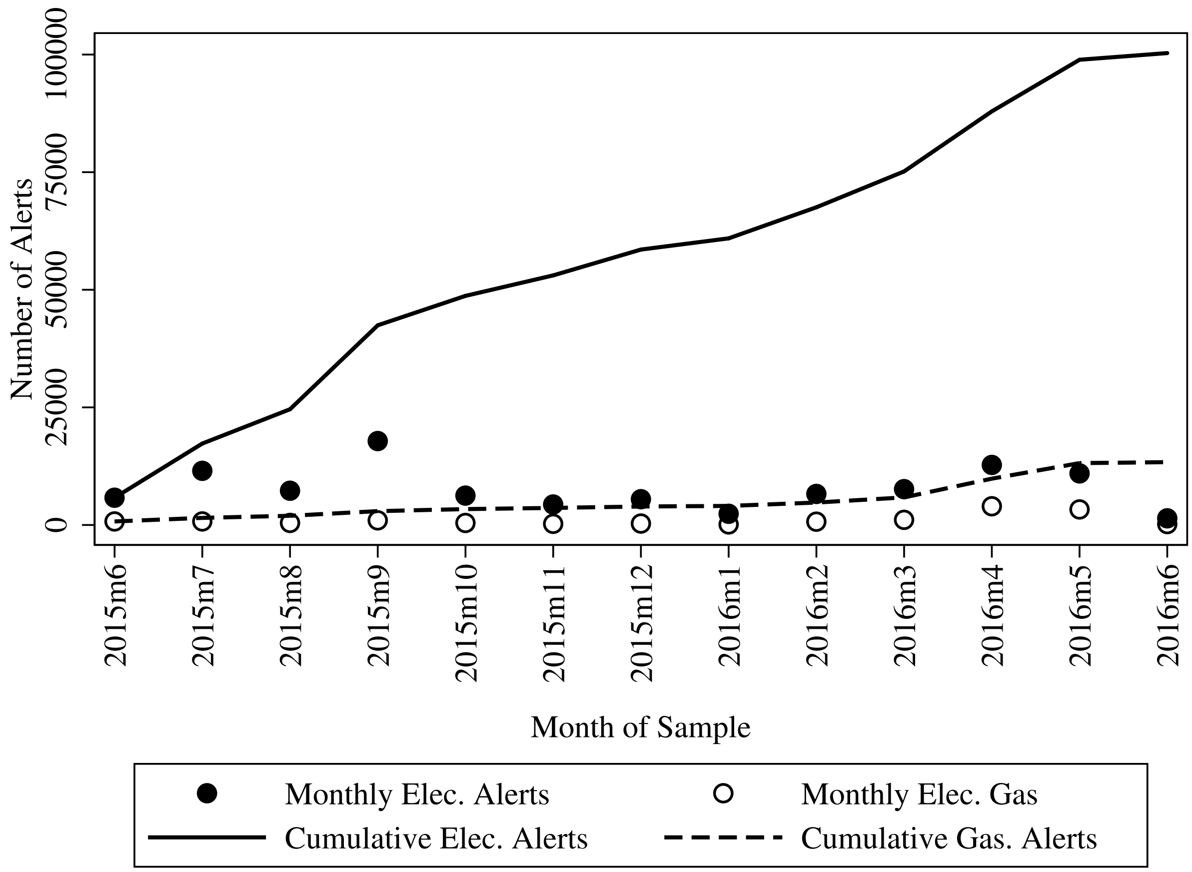

Figure 2: **HBA Alerts Trends.**

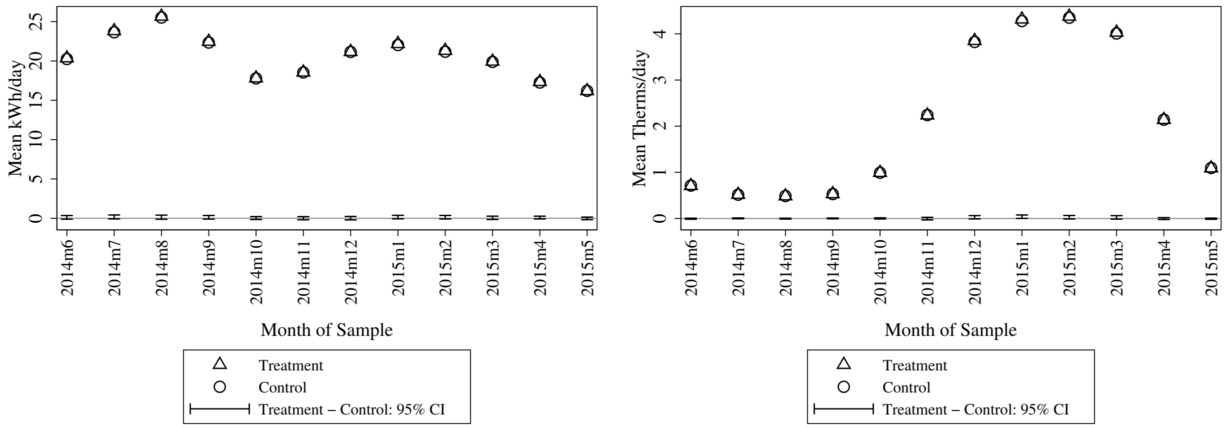

(1) Electricity

(2) Natural Gas

Figure 3: **Mean ADC Prior to the Start of the HBA Program by Experimental Group.** The triangles and circles plot mean daily usage levels for the treatment and control groups, respectively, based on billing cycles that ended during the corresponding month as indicated by the horizontal axis. The brackets indicate the 95% confidence interval for the difference in means usage levels between treatment and control households.

Table 1: Estimates of the Effect of the HBA Program on Mean ADC

	Electricity			Natural Gas		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment × Post-Pd.	-0.105*	-0.110***	-0.103***	-0.014**	-0.009***	-0.009***
	(0.056)	(0.038)	(0.038)	(0.006)	(0.003)	(0.003)
Treatment Indicator	0.079			0.015*		
	(0.091)			(0.009)		
Post-Pd. Indicator	1.218***			-0.325***		
	(0.046)			(0.005)		
Constant	20.334***	21.184***	21.181***	2.180***	1.976***	1.976***
	(0.075)	(0.016)	(0.017)	(0.008)	(0.001)	(0.001)
HH FEs	No	Yes	No	No	Yes	No
HH-by-Cal.-Month FEs	No	No	Yes	No	No	Yes
Month-of-Sample FEs	No	Yes	Yes	No	Yes	Yes
Treat. Eff. as %	-0.49	-0.51	-0.48	-0.79	-0.52	-0.50
R-squared	0.00	0.77	0.91	0.01	0.79	0.96
Observations	2,144,460	2,144,460	2,144,460	2,167,807	2,167,807	2,167,807

Notes: The unit of analysis is a household and an electricity or natural gas bill. The dependent variable is ADC for either electricity (kWh) or natural gas (therms) as indicated in the heading. All models are linear regression models. Temporal fixed effects are coded based on the month at the end of the billing cycle. Standard errors clustered by household. Observations are weighted by the number of days in the billing cycle. One, two, and three stars indicate 10 percent, 5 percent, and 1 percent significance, respectively.

Table 2: Estimates of the Effect of the HBA Program on Mean ADC by Program Stage

	Electricity			Gas		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment × (Jun 2015-June 2016)	-0.093*	-0.095***	-0.090**	-0.016***	-0.009***	-0.010***
	(0.052)	(0.036)	(0.036)	(0.005)	(0.003)	(0.003)
Treatment × (July 2016-May 2017)	-0.118*	-0.130***	-0.121**	-0.012*	-0.009**	-0.008***
	(0.070)	(0.048)	(0.049)	(0.007)	(0.004)	(0.003)
Treatment Indicator	0.079			0.015*		
	(0.091)			(0.009)		
Jun 2015-June 2016 Indicator	0.706***			-0.371***		
	(0.042)			(0.004)		
July 2016-May 2017 Indicator	1.841***			-0.267***		
	(0.058)			(0.006)		
Constant	20.334***	21.184***	21.181***	2.180***	1.976***	1.976***
	(0.075)	(0.017)	(0.017)	(0.008)	(0.001)	(0.001)
HH FEs	No	Yes	No	No	Yes	No
HH-by-cal.-month FEs	No	No	Yes	No	No	Yes
Month-of-Sample FEs	No	Yes	Yes	No	Yes	Yes
1st Year Treat. Eff. as %	-0.44	-0.45	-0.43	-0.90	-0.53	-0.54
2nd Year Treat. Eff. as %	-0.56	-0.62	-0.58	-0.69	-0.53	-0.48
R-squared	0.00	0.77	0.91	0.01	0.79	0.96
Observations	2,144,460	2,144,460	2,144,460	2,167,807	2,167,807	2,167,807

Notes: The unit of analysis is a household and an electricity or natural gas bill. The dependent variable is ADC for either electricity (kWh) or natural gas (therms) as indicated in the heading. All models are linear regression models. Temporal fixed effects are coded based on the month at the end of the billing cycle. Standard errors clustered by household. Observations are weighted by the number of days in the billing cycle. One, two, and three stars indicate 10 percent, 5 percent, and 1 percent significance, respectively.

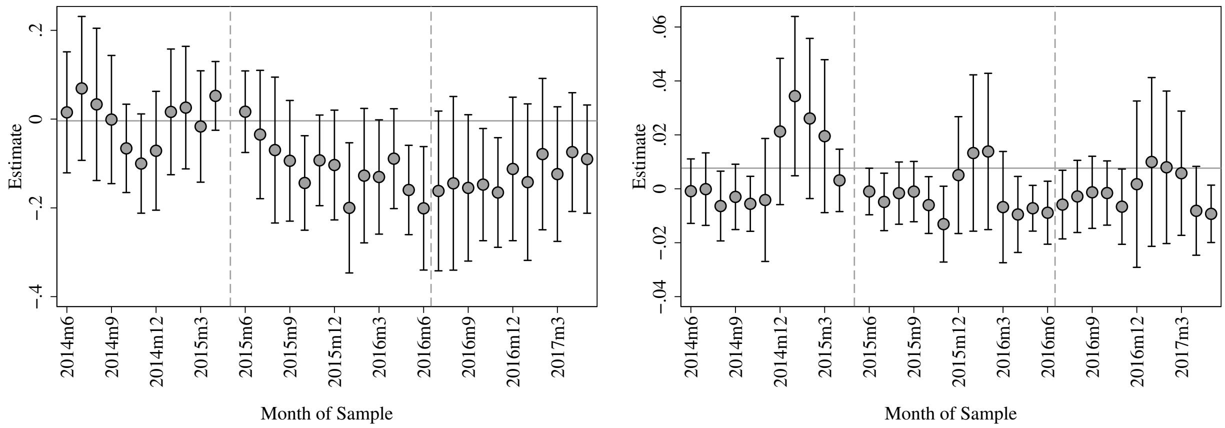

(1) Electricity

(2) Natural Gas

Figure 4: **Estimates of the Effect of the HBA Program on Mean ADC by Month.** Estimates are based on a linear regression of ADC on household fixed effects, month-of-sample fixed effects, and interactions of a treatment indicator with an indicator for each month of the sample, except for the month immediately prior to the launch of the program in June 2015, which is represented by the vertical dashed line to the left. The second vertical dashed line to the right represents the point where alerts stopped being sent after June 2016. The horizontal line represents the mean for the point estimates for all months of the sample prior to the beginning of the program.

Table 3: Quantile Regression Estimates for Nominal ADC

Panel A: Electricity	1%	5%	10%	25%	50%	75%	90%	95%	99%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treatment × Post-Pd.	0.000	0.011	0.047	-0.030	-0.102**	-0.219***	-0.293**	-0.173	0.026
	(0.029)	(0.036)	(0.035)	(0.037)	(0.047)	(0.073)	(0.124)	(0.176)	(0.303)
Constant	3.023***	5.511***	7.555***	11.979***	18.650***	27.890***	39.017***	47.013***	63.264***
	(0.013)	(0.016)	(0.016)	(0.017)	(0.021)	(0.033)	(0.056)	(0.080)	(0.138)
R-squared	0.65	0.77	0.79	0.81	0.81	0.79	0.76	0.73	0.61
Observations	2,007,410	2,007,410	2,007,410	2,007,410	2,007,410	2,007,410	2,007,410	2,007,410	2,007,410
Panel B: Natural Gas	1%	5%	10%	25%	50%	75%	90%	95%	99%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treatment × Post-Pd.	-	-	-0.001	-0.003*	-0.007**	-0.028***	-0.021*	-0.051***	-0.009
	-	-	(0.001)	(0.002)	(0.003)	(0.008)	(0.013)	(0.018)	(0.029)
Constant	-	-	0.128***	0.481***	1.154***	3.086***	5.014***	6.143***	8.121***
	-	-	(0.001)	(0.001)	(0.002)	(0.003)	(0.006)	(0.008)	(0.013)
R-squared	-	-	0.90	0.87	0.91	0.88	0.80	0.74	0.58
Observations	-	-	2,032,096	2,032,096	2,032,096	2,032,096	2,032,096	2,032,096	2,032,096

Notes: The unit of analysis is a household and an electricity or natural gas bill. The dependent variable is ADC for either electricity (kWh) or natural gas (therms) as indicated in the panel headings. All models are unconditional quantile regression models, where the column headings indicate the point in the distribution for which the model is estimated. Natural gas models are not estimated for the first and fifth percentiles due to the concentration of zeroes at the bottom of the natural gas consumption distribution. All models include month-of-sample fixed effects and household-by-calendar month fixed effects. Temporal fixed effects are coded based on the month at the end of the billing cycle. Standard errors are clustered by household. Observations weighted by the number of days in the billing cycle. One, two, and three stars indicate 10 percent, 5 percent, and 1 percent significance, respectively.

Table 4: Quantile Regression Estimates for Relative ADC

Panel A: Electricity	1%	5%	10%	25%	50%	75%	90%	95%	99%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treatment × Post-Pd.	0.140	-0.205	-0.285	-0.355**	-0.338**	-0.466**	-1.224***	-1.738***	-2.275*
	(0.227)	(0.224)	(0.187)	(0.139)	(0.133)	(0.218)	(0.424)	(0.662)	(1.263)
Constant	45.008***	61.846***	71.415***	86.219***	101.454***	121.024***	149.311***	174.271***	240.049***
	(0.187)	(0.184)	(0.153)	(0.114)	(0.109)	(0.180)	(0.350)	(0.547)	(1.049)
R-squared	0.00	0.00	0.01	0.02	0.03	0.03	0.01	0.01	0.00
Observations	1,148,220	1,148,220	1,148,220	1,148,220	1,148,220	1,148,220	1,148,220	1,148,220	1,148,220
Panel B: Gas	1%	5%	10%	25%	50%	75%	90%	95%	99%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treatment × Post-Pd.	-	-	-0.101	-0.202*	-0.326***	-0.357***	-0.657***	-0.738	-1.301
	-	-	(0.181)	(0.117)	(0.103)	(0.109)	(0.246)	(0.472)	(1.130)
Constant	-	-	54.252***	70.571***	89.287***	106.240***	126.061***	145.959***	206.796***
	-	-	(0.147)	(0.095)	(0.085)	(0.090)	(0.204)	(0.392)	(0.939)
R-squared	-	-	0.06	0.16	0.20	0.12	0.06	0.03	0.01
Observations	-	-	1,119,977	1,119,977	1,119,977	1,119,977	1,119,977	1,119,977	1,119,977

Notes: The unit of analysis is a household and an electricity or natural gas bill. The dependent variable is relative ADC (where relative is measured as a comparison to ADC during the same calendar month in the year prior to the start of the HBA program) for either electricity (kWh) or natural gas (therms) as indicated in the panel headings. All models are estimated with the post-treatment sample. Natural gas models are not estimated for the first and fifth percentiles due to the concentration of zeroes at the bottom of the natural gas consumption distribution. All models are unconditional quantile regression models, where the column headings indicate the point in the distribution for which the model is estimated. All models include month-of-sample fixed effects. Temporal fixed effects are coded based on the month at the end of the billing cycle. Standard errors are clustered by household. Observations weighted by the number of days in the billing cycle. One, two, and three stars indicate 10 percent, 5 percent, and 1 percent significance, respectively.

Table 5: Estimates of the Effect of the HBA Program on Mean ADC with Income Interaction

	Electricity			Natural Gas		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment × Post-Pd.	-0.106*	-0.111***	-0.103***	-0.014**	-0.009***	-0.009***
	(0.055)	(0.038)	(0.038)	(0.006)	(0.003)	(0.003)
Treatment Indicator	0.070			0.015		
	(0.090)			(0.009)		
Treatment × Post-Pd. × Income	0.023	0.034***	-0.008	-0.002	-0.000	-0.001
	(0.029)	(0.011)	(0.011)	(0.002)	(0.001)	(0.001)
Post-Pd. Indicator	1.171***			-0.326***		
	(0.045)			(0.005)		
Income	0.866***			0.029***		
	(0.025)			(0.002)		
Constant	20.369***	21.184***	21.181***	2.181***	1.976***	1.976***
	(0.074)	(0.016)	(0.017)	(0.008)	(0.001)	(0.001)
HH FEs	No	Yes	No	No	Yes	No
HH-by-Cal.-Month FEs	No	No	Yes	No	No	Yes
Month-of-Sample FEs	No	Yes	Yes	No	Yes	Yes
Treat. Eff. as %	-0.49	-0.52	-0.48	-0.79	-0.52	-0.50
R-squared	0.02	0.77	0.91	0.01	0.79	0.96
Observations	2,144,460	2,144,460	2,144,460	2,167,807	2,167,807	2,167,807

Notes: The unit of analysis is a household and an electricity or natural gas bill. Income is measured at the zip-code level, de-meaned, and measured in units of $10,000s. The dependent variable is ADC for either electricity (kWh) or natural gas (therms) as indicated in the heading. All models are linear regression models. Temporal fixed effects are coded based on the month at the end of the billing cycle. Standard errors clustered by household. Observations are weighted by the number of days in the billing cycle. One, two, and three stars indicate 10 percent, 5 percent, and 1 percent significance, respectively.

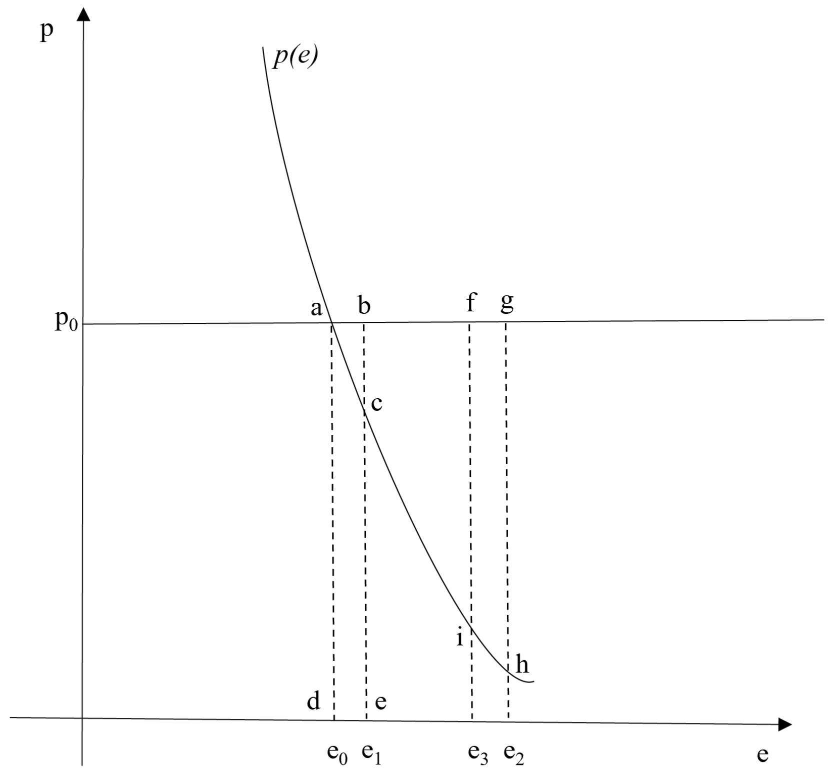

Figure 5: **Market Demand and Energy Price.** Area *abc* represents an estimate of the benefits of the HBA program to consumers when the foregone value of energy consumption is just marginal to the energy price. Area *fghi* represents the benefits of the HBA program to consumers when the foregone value of energy consumption is significantly less than the energy price. Area *abed* represents the benefits of the HBA program to consumers when the foregone value of energy consumption is zero.

Table 6: Upper and Lower Bound Estimates of Consumer Benefits

Assumptions and Parameters:

Annual Electricity Consumption per Household (kwh):	7,820
Average Price of Electricty (cents/kWh):	13.90
Annual Electricity Expenditures per Household (kwh):	1086.98
Annual Electricity Conservation per Customer (kwh)	39.10
Avoided Electricity Expenditures per Household (dollars):	5.43
Residential Electricity Customers:	1,186,195
Annual Natural Gas Consumption per Household (therms):	853
Average Price of Natural Gas (dollars/therm):	1.33
Annual Natural Gas Expenditures per Household (dollars):	1,137.90
Annual Natural Gas Conservation per Customer (therms)	4.26
Avoided Natural Gas Expenditures per Household (dollars):	5.69
Residential Natural Gas Customers:	438,247

Upper Bound Estimates of Consumer Benefits (dollars):

Electricity Benefits per Household	5.43
Natural Gas Benefits per Household	5.69
Electricity Benefits in Aggregate	6,444,851
Natural Gas Benefits in Aggregate	2,493.406
Combined Benefits in Aggregate	**8,940,257**

Lower Bound Estimates of Consumer Benefits (dollars):

Elec. Elas.	Elec. Ben. per HH.	Elec. Ben. Agg.	Gas. Elas.	Gas Ben. per HH	Gas Ben. Agg.	Combined Ben. Agg.
-0.01	1.36	1,611,713	-0.01	1.42	623,352	2,235,065
-0.05	0.27	322,343	-0.05	0.28	124,670	447,013
-0.10	**0.14**	**161,171**	**-0.10**	**0.14**	**62,335**	**223,506**
-0.15	0.09	107,448	-0.15	0.09	41,557	149,004
-0.20	0.07	80,586	-0.20	0.07	31,168	111,753
-0.25	0.05	64,469	-0.25	0.06	24,934	89,403
-0.30	0.05	53,724	-0.30	0.05	20,778	74,502
-0.35	0.04	46,049	-0.35	0.04	17,810	63,859

Notes: Consumption and customer count data taken from EIA-861 for the utility for the year 2021. All dollars are inflation adjusted to the year 2021. Avoided expenditures based on estimates from Table 1. Upper bound estimates assume zero foregone value from conservation and are calculated as the price of electricity, times the change in consumption per household, times the total number of customers. Lower bound estimates assume the foregone value from the first unit of conservation is equal to the price and then declines linearly as a function of the price elasticity. See more discussion in Section 6.

10.2 Appendix Tables and Figures

A.2.1 The Association between Receiving an Alert and Mean ADC

Table A.1: Estimates of the Association between Receiving an Alert and Mean ADC

	ADC				Relative ADC			
Dependent Variable:	Elec.	Gas	Elec.	Gas	Elec.	Gas	Elec.	Gas
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Alert - Elec.	4.482***	-0.007***			32.546***	2.901***		
	(0.032)	(0.002)			(0.169)	(0.131)		
Alert - Gas.			0.846***	0.186***			5.506***	16.457***
			(0.088)	(0.009)			(0.533)	(0.464)
Constant	20.942***	1.973***	21.131***	1.971***	102.826***	91.541***	105.204***	91.651***
	(0.001)	(0.000)	(0.000)	(0.000)	(0.071)	(0.049)	(0.075)	(0.049)
Sample	Pre+Post	Pre+Post	Pre+Post	Pre+Post	Post	Post	Post	Post
HH-by-cal.-month FEs	Yes	Yes	Yes	Yes	No	No	No	No
Month-of-Sample FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Alert Coef. as %	20.87	-0.41	3.94	10.23	-	-	-	-
R-squared	0.91	0.96	0.91	0.96	0.10	0.18	0.03	0.19
Observations	2,144,460	2,167,807	2,144,460	2,167,807	1,143,546	1,085,333	1,143,546	1,085,333

Notes: The unit of analysis is a household and an electricity or natural gas bill. The dependent variable is ADC for either electricity (kWh) or natural gas (therms) as indicated in the heading. The relative ADC models are estimated with the post-treatment sample. All models are linear regression models. The alert variables are indicators for having received an alert of the corresponding type the billing cycle. Temporal fixed effects are coded based on the month at the end of the billing cycle. Standard errors clustered by household. Observations are weighted by the number of days in the billing cycle. One, two, and three stars indicate 10 percent, 5 percent, and 1 percent significance, respectively.

A.3 *Reproducing Results from Main Analysis with Sample that Includes Outliers*

Table A.2: Estimates of the Effect of the HBA Program on Mean ADC - Outliers Included

	Electricity			Natural Gas		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment × Post-Pd.	-0.102	-0.093**	-0.083*	-0.015**	-0.008**	-0.008**
	(0.065)	(0.043)	(0.044)	(0.006)	(0.004)	(0.003)
Treatment Indicator	-0.115			0.009		
	(0.120)			(0.012)		
Post-Pd. Indicator	1.210***			-0.382***		
	(0.054)			(0.005)		
Constant	21.147***	21.854***	21.849***	2.332***	2.086***	2.086***
	(0.101)	(0.019)	(0.019)	(0.010)	(0.002)	(0.001)
HH FEs	No	Yes	No	No	Yes	No
HH-by-Cal.-Month FEs	No	No	Yes	No	No	Yes
Month-of-Sample FEs	No	Yes	Yes	No	Yes	Yes
Treat. Eff. as %	-0.46	-0.42	-0.37	-0.76	-0.42	-0.44
R-squared	0.00	0.78	0.92	0.01	0.76	0.96
Observations	2,187,669	2,187,669	2,187,669	2,189,624	2,189,624	2,189,624

Notes: The unit of analysis is a household and an electricity or natural gas bill. The dependent variable is ADC for either electricity (kWh) or natural gas (therms) as indicated in the heading. All models are linear regression models. Temporal fixed effects are coded based on the month at the end of the billing cycle. Standard errors clustered by household. Observations are weighted by the number of days in the billing cycle. One, two, and three stars indicate 10 percent, 5 percent, and 1 percent significance, respectively.

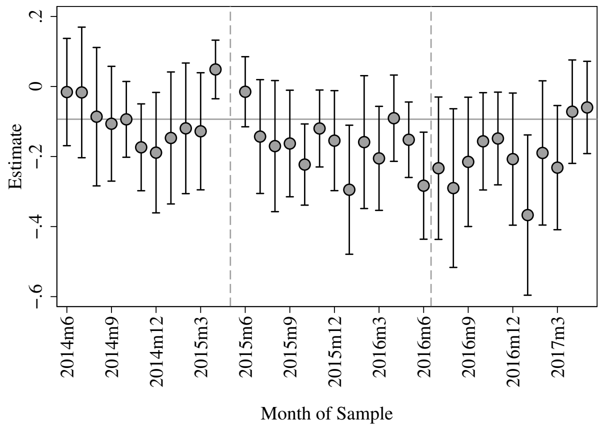

(1) Electricity

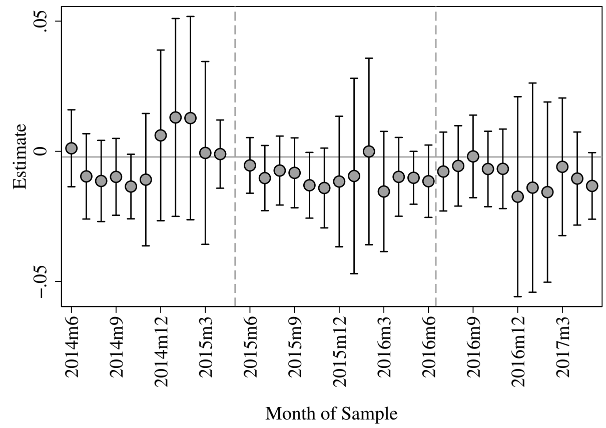

(2) Gas

Figure A.1: **Estimates of the Effect of the HBA Program on Mean ADC by Month - Outliers Included.** Estimates are based on a linear regression of ADC on household fixed effects, month-of-sample fixed effects, and interactions of a treatment indicator with an indicator for each month-of-sample, except for the month immediately prior to the launch of the program in June 2015, which is represented by the vertical dashed line to the left. The second vertical dashed line to the right represents the point where alerts stopped being dispensed after June 2016. The horizontal line represents the mean for the point estimates for all months of the sample prior to the beginning of the program.

Table A.3: Quantile Regression Estimates for Nominal ADC - Outliers Included

Panel A: Electricity	1%	5%	10%	25%	50%	75%	90%	95%	99%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treatment × Post-Pd.	0.003	-0.007	0.041	-0.030	-0.102**	-0.231***	-0.248*	-0.054	-0.011
	(0.043)	(0.039)	(0.036)	(0.037)	(0.047)	(0.074)	(0.131)	(0.205)	(0.693)
Constant	2.207***	5.083***	7.279***	11.856***	18.678***	28.227***	40.209***	49.695***	77.861***
	(0.019)	(0.018)	(0.017)	(0.017)	(0.021)	(0.033)	(0.060)	(0.093)	(0.315)
R-squared	0.70	0.78	0.80	0.81	0.81	0.80	0.78	0.77	0.77
Observations	2,053,498	2,053,498	2,053,498	2,053,498	2,053,498	2,053,498	2,053,498	2,053,498	2,053,498
Panel B: Natural Gas	1%	5%	10%	25%	50%	75%	90%	95%	99%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treatment × Post-Pd.	-	-	-0.000	-0.003*	-0.007**	-0.024***	-0.019	-0.045**	0.039
	-	-	(0.001)	(0.002)	(0.004)	(0.008)	(0.013)	(0.021)	(0.063)
Constant	-	-	0.135***	0.487***	1.182***	3.183***	5.189***	6.515***	9.646***
	-	-	(0.001)	(0.001)	(0.002)	(0.003)	(0.006)	(0.009)	(0.029)
R-squared	-	-	0.90	0.87	0.91	0.88	0.81	0.77	0.74
Observations	-	-	2,055,651	2,055,651	2,055,651	2,055,651	2,055,651	2,055,651	2,055,651

Notes: The unit of analysis is a household and an electricity or natural gas bill. The dependent variable is ADC for either electricity (kWh) or natural gas (therms) as indicated in the panel headings. All models are unconditional quantile regression models, where the column headings indicate the point in the distribution for which the model is estimated. Natural gas models are not estimated for the first and fifth percentiles due to the concentration of zeroes at the bottom of the natural gas consumption distribution. All models include month-of-sample fixed effects and household-by-calendar month fixed effects. Temporal fixed effects are coded based on the month at the end of the billing cycle. Standard errors are clustered by household. Observations weighted by the number of days in the billing cycle. One, two, and three stars indicate 10 percent, 5 percent, and 1 percent significance, respectively.

Table A.4: Quantile Regression Estimates for Relative ADC - Outliers Included

Panel A: Electricity	1%	5%	10%	25%	50%	75%	90%	95%	99%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treatment × Post-Pd.	1.154*	0.064	-0.178	-0.324**	-0.333**	-0.520**	-1.516***	-2.798***	-11.860**
	(0.626)	(0.301)	(0.216)	(0.146)	(0.136)	(0.230)	(0.509)	(0.945)	(5.871)
Constant	32.405***	58.516***	69.818***	85.736***	101.422***	121.628***	152.434***	183.118***	321.230***
	(0.516)	(0.248)	(0.177)	(0.119)	(0.111)	(0.190)	(0.421)	(0.785)	(4.912)
R-squared	0.00	0.00	0.01	0.02	0.03	0.03	0.01	0.01	0.00
Observations	1,193,975	1,193,975	1,193,975	1,193,975	1,193,975	1,193,975	1,193,975	1,193,975	1,193,975
Panel B: Gas	1%	5%	10%	25%	50%	75%	90%	95%	99%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treatment × Post-Pd.	-	-	-0.136	-0.238**	-0.349***	-0.413***	-0.910***	-1.305*	-15.332***
	-	-	(0.177)	(0.116)	(0.104)	(0.113)	(0.312)	(0.680)	(5.563)
Constant	-	-	54.654***	71.020***	89.847***	107.045***	128.921***	153.274***	286.522***
	-	-	(0.144)	(0.095)	(0.085)	(0.093)	(0.259)	(0.566)	(4.651)
R-squared	-	-	0.06	0.17	0.20	0.12	0.06	0.03	0.00
Observations	-	-	1,141,558	1,141,558	1,141,558	1,141,558	1,141,558	1,141,558	1,141,558

Notes: The unit of analysis is a household and an electricity or natural gas bill. The dependent variable is relative ADC (where relative is measured as a comparison to ADC during the same calendar month in the year prior to the start of the HBA program) for either electricity (kWh) or natural gas (therms) as indicated in the panel headings. All models are estimated with the post-treatment sample. Natural gas models are not estimated for the first and fifth percentiles due to the concentration of zeroes at the bottom of the natural gas consumption distribution. All models are unconditional quantile regression models, where the column headings indicate the point in the distribution for which the model is estimated. All models include month-of-sample fixed effects. Temporal fixed effects are coded based on the month at the end of the billing cycle. Standard errors are clustered by household. Observations weighted by the number of days in the billing cycle. One, two, and three stars indicate 10 percent, 5 percent, and 1 percent significance, respectively.

A.4 Quantile Regression Estimates for Relative ADC with Income Interaction

Table A.5: Quantile Regression Estimates for Relative ADC with Income Interaction

Panel A: Electricity	1%	5%	10%	25%	50%	75%	90%	95%	99%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treatment × Post-Pd.	0.132	-0.216	-0.293	-0.359***	-0.336**	-0.456**	-1.200***	-1.702***	-2.228*
	(0.227)	(0.223)	(0.187)	(0.139)	(0.133)	(0.218)	(0.422)	(0.661)	(1.261)
Treatment × Post-Pd. × Income	-0.023	-0.034	-0.054	-0.026	-0.056	-0.134	-0.146	-0.221	0.082
	(0.112)	(0.112)	(0.096)	(0.072)	(0.069)	(0.110)	(0.207)	(0.324)	(0.600)
Income (Zip-code-level, median)	0.688***	0.903***	0.730***	0.328***	-0.133**	-0.802***	-1.895***	-2.785***	-3.835***
	(0.091)	(0.091)	(0.078)	(0.059)	(0.057)	(0.091)	(0.173)	(0.273)	(0.502)
Constant	45.013***	61.852***	71.420***	86.221***	101.454***	121.018***	149.298***	174.252***	240.023***
	(0.187)	(0.183)	(0.153)	(0.114)	(0.109)	(0.179)	(0.348)	(0.546)	(1.047)
R-squared	0.00	0.00	0.01	0.02	0.03	0.03	0.02	0.01	0.00
Observations	1,148,220	1,148,220	1,148,220	1,148,220	1,148,220	1,148,220	1,148,220	1,148,220	1,148,220
Panel B: Gas	1%	5%	10%	25%	50%	75%	90%	95%	99%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treatment × Post-Pd.	-	-	-0.105	-0.204*	-0.328***	-0.358***	-0.656***	-0.737	-1.300
	-	-	(0.180)	(0.117)	(0.103)	(0.109)	(0.246)	(0.472)	(1.130)
Treatment × Post-Pd. × Income	-	-	0.015	-0.001	0.027	-0.006	0.022	-0.131	-0.257
	-	-	(0.092)	(0.059)	(0.052)	(0.055)	(0.125)	(0.239)	(0.584)
Income (Zip-code-level, median)	-	-	0.771***	0.303***	0.326***	0.168***	-0.228**	-0.258	0.047
	-	-	(0.075)	(0.048)	(0.043)	(0.046)	(0.104)	(0.198)	(0.480)
Constant	-	-	54.252***	70.572***	89.288***	106.241***	126.061***	145.958***	206.796***
	-	-	(0.147)	(0.095)	(0.085)	(0.090)	(0.204)	(0.392)	(0.939)
R-squared	-	-	0.06	0.16	0.20	0.12	0.06	0.03	0.01
Observations	-	-	1,119,977	1,119,977	1,119,977	1,119,977	1,119,977	1,119,977	1,119,977

Notes: The unit of analysis is a household and an electricity or natural gas bill. Income is measured at the zip-code level, de-meaned, and measured in units of $10,000s. The dependent variable is relative ADC (where relative is measured as a comparison to ADC during the same calendar month in the year prior to the start of the HBA program) for either electricity (kWh) or natural gas (therms) as indicated in the panel headings. All models are estimated with the post-treatment sample. Natural gas models are not estimated for the first and fifth percentiles due to the concentration of zeroes at the bottom of the natural gas consumption distribution. All models are unconditional quantile regression models, where the column headings indicate the point in the distribution for which the model is estimated. All models include month-of-sample fixed effects. Temporal fixed effects are coded based on the month at the end of the billing cycle. Standard errors are clustered by household. Observations weighted by the number of days in the billing cycle. One, two, and three stars indicate 10 percent, 5 percent, and 1 percent significance, respectively.