NBER WORKING PAPER SERIES

KNOWLEDGE IS (LESS) POWER:
EXPERIMENTAL EVIDENCE FROM RESIDENTIAL ENERGY USE

Katrina Jessoe
David Rapson

Working Paper 18344
http://www.nber.org/papers/w18344

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2012

We would like to thank Hunt Allcott, Jim Bushnell, Colin Cameron, Scott Carrell, Michael Carter, Meredith Fowlie, Koichiro Ito, Stephen Holland, Hilary Hoynes, Michael Price, Nancy Rose and Burkhard Schipper. This paper benefited from helpful comments by seminar participants at the CU Boulder, NBER, UC Davis, UC Energy Institute, and UCE3. Thanks to United Illuminating personnel for their important role in implementing this project. Tom Blake and Suzanne Plant provided excellent research assistance. Research support for this project was provided by UCE3 and United Illuminating. All errors are our own.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

Knowledge is (Less) Power: Experimental Evidence from Residential Energy Use
Katrina Jessoe and David Rapson
NBER Working Paper No. 18344
August 2012
JEL No. L94

ABSTRACT

This paper presents experimental evidence that information feedback dramatically increases the price elasticity of demand in a setting where signals about quantity consumed are traditionally coarse and infrequent. In a randomized controlled trial, residential electricity customers are exposed to price increases, with some households also receiving displays that transmit high-frequency information about usage and prices. This substantially lowers information acquisition costs and allows us to identify the marginal information effect. Households only experiencing price increases reduce demand by 0 to 7 percent whereas those also exposed to information feedback exhibit a usage reduction of 8 to 22 percent, depending on the amount of advance notice. The differential response across treatments is significant and robust to the awareness of price changes. Conservation extends beyond the treatment window, providing evidence of habit formation, spillovers, and greenhouse gas abatement. Results suggest that information about the quantity consumed facilitates learning, which likely drives the treatment differential.

Katrina Jessoe
Department of Agricultural and Resource Economics
University of California, Davis
One Shields Avenue
Davis, CA 95616
kkjessoe@ucdavis.edu

David Rapson
Department of Economics
University of California, Davis
One Shields Avenue
Davis, CA 95616
dsrapson@ucdavis.edu

1 Introduction

Traditional economic models assume that agents know the full choice set available to them, can costlessly compute strategy-specific payoffs, and have preferences that allow them to maximize their net economic benefit. Starting with Simon (1955) studies began to question these assumptions, recognizing that information may be costly to acquire and that there are limitations to human cognition and attention. Recent empirical studies have supported this view, showing that these features affect outcomes in market settings.[1] Much of the existing literature has focused on understanding what and how information is presented, emphasizing the role of price. However, in many settings there may also be a noisy signal about quantity, and less is known about the implications of this form of uncertainty on the efficiency of consumer decisions.[2] In this study we create an experimental setting in which granular and accurate information on quantity is exogenously provided, allowing us to evaluate the effect of information on price elasticity. Providing simple real-time information to residential electricity users on their consumption increases the price elasticity of demand three-fold.

The residential electricity sector is an ideal market to study this question for three reasons. First, this is an enormous industry known for its regulatory difficulties (e.g. market power) and external costs (e.g. pollution).[3] Second, the quantity of electricity consumption is shrouded from residential customers, in part because of coarse and infrequent billing.[4] Since electricity is inherently an input to household production rather than a final consumable, it lives in the background of our lives and makes the relevant unit of measure (kilowatt-hour) meaningless to most. Third, electricity customers traditionally exhibit low response to prices (Reiss and White 2005, Allcott 2011a, Ito 2011) for reasons that are difficult to disentangle. It may be that, all else equal, low-elasticity households are unresponsive to prices because they are fully informed about how

[1] Consumers treat taxes featured and not featured in the posted price differently, becoming more price elastic when taxes are salient (Chetty et al. 2009). They are inattentive to opaque “add-on” costs such as shipping and handling fees (Hossain and Morgan 2006) and stock market earnings announcements, with inattention to financial news depending on the timing of reports (DellaVigna and Pollet 2008) and the number of competing stimuli (Hirshleifer et al. 2009). See DellaVigna (2008) for a comprehensive review of field evidence from this literature.

[2] Empirical work related to the role of information about quantity shows that consumers rely on heuristics (exhibiting left-digit bias) when considering mileage of used cars in their purchase decision (Lacetera et al 2012) and do not fully incorporate available information on cell phone usage when making choices about plans and future usage (Grubb and Osborne 2012).

[3] There are $370 billion per year in retail sales. This amounts to roughly two-thirds of US public expenditure on primary and secondary education in 2009. (U.S. Department of Energy 2010; National Center for Education Statistics 2011).

[4] Prices may also be confusing. Complex rate structures make it difficult for customers to know the marginal price. Recent work suggests that when faced with a non-linear price schedule, consumers often rely on a heuristic, responding more to average prices (Ito 2011). More frequent information on electricity prices has been shown to increase the short-run price elasticity (Allcott 2011a).

much it costs to use household appliances, equate marginal utility to marginal cost, and are rationally price inelastic. One alternative explanation, which we explore, is that the low elasticity is partly explained by the lack of full information about either price or quantity. If this is the case, policy interventions that provide information about price and quantity may lead consumers to make more efficient choices.

The ideal experiment to identify the incremental effect of information feedback on consumer price elasticity has not previously been implemented. This article describes results from just such an experiment. We design and implement a randomized controlled trial (RCT) in which households are exposed to exogenous price changes and (randomized) real-time feedback on both the price and quantity of electricity. To isolate the effect of price changes on demand, all treatment households were exposed to two- or four-hour long pricing events during which the price of electricity increased by 200 to 600 percent.[5] In addition, some of the households were given in-home displays (IHDs) that provide real-time information on electricity prices and aggregate quantity consumed. This in-home technology enables customers to inform themselves about electricity usage and prices at almost zero marginal cost. The IHD installations occurred months before the first pricing event, and interested customers could use the devices to learn about the how electricity actions translate into usage.

Home Area Network (HAN) devices[6], such as the displays evaluated in this study, are being tested by utilities in a number of pilot settings, often in conjunction with dynamic pricing. These pilots have led to several industry research studies, some of which share features of ours. However, utility-run experiments are often constrained by the desire to avoid treating households differently, which can be perceived as unfair. From a research design perspective, this leads to selection bias since households are able to affect the treatment that they receive.[7] This study is designed to obtain internally consistent estimates of the role of information, which in our setting includes feedback on both quantity and price.[8]

Results from the experiment support the hypothesis that information about quantity plays an important role in consumer decisions. Conditional on changes in price, households exposed to information feedback consis-

[5]This price change is a variant of what the industry refers to as "dynamic pricing". The increment is consistent with the magnitude of peak fluctuations in the wholesale electricity market.

[6]The Home Area Network (HAN), the household-facing layer of technology, transmits real-time information on electricity usage and prices from the meter to both the utility and household, offering consumers the potential to access more information and exert more control over their energy use.

[7]Examples can be found in Faruqui and Sergici (2010), Faruqui et al. (2012), and Herter (2012).

[8]To our knowledge, the only other peer-reviewed study that studies both information and price in electricity is Allcott (2011a). Information in that setting is in the form of a price orb that changes color as prices increase, but gives no feedback about quantity. Another contrast with this study is that the price orb program designed by Commonwealth Edison was not free from selection concerns.

tently exhibit price elasticities that are roughly three times as large as those without feedback. Households in the price-only group reduce their usage by between 0 and 7 percent, on average, during pricing events. In contrast, those exposed to the same price changes but who also have IHDs, exhibit much larger usage reductions of 8 to 22 percent. In an industry in which a 1 to 3 percent behavioral response is considered meaningful, the incremental contribution of real-time information is remarkable.[9] We use responses collected from household surveys to show that the information effect is robust to several possible explanations about the underlying mechanism.

A common hypothesis is that households are inattentive to electricity prices because they are not aware of them, perhaps because price changes are not salient. We attempt to control for this in the experimental design, but nonetheless test if the information treatment differential can be explained by an increase in the awareness of price for households with IHDs. While awareness of price events causes households to be more responsive overall, it does not explain why informed households are more price elastic than their uninformed counterparts. Instead, our empirical evidence suggests that experience with IHDs facilitates consumer learning, helping households to make better decisions when confronted with high prices. Consumers who look at the information device more frequently are significantly more responsive to price. This is consistent with a framework in which experience with the IHDs allows consumers to better understand how their day-to-day actions translate into electricity usage. It also fits with earlier empirical work showing that frequent market experience can reduce market anomalies (List 2003, Feng and Seasholes 2005).

The dynamic effects of pricing events, both within an event day and on non-event days, cause our estimates to be a lower bound of the primary treatment effects as well as of the information treatment differential. Within a day, the conservation effect observed during pricing events spills over into adjacent hours for households with information feedback. In the longer run, an evaluation of trends in usage over the days of the summer reveals that households in both the price and price+IHD group are forming habits, significantly conserving energy even when events are not occurring. The persistence of behavioral change adds another dimension to our understanding of how consumers respond, providing further evidence of learning. Additionally, the combined dynamic effects imply a social benefit in the form of greenhouse gas abatement.

The main result of this paper implies that incomplete information may be causing large efficiency losses,

[9] Some evidence of the 1-3 percent reduction in electricity usage is shown in Jacobsen et al 2010, Allcott 2011b, Harding and Rapson 2012. The Opower social norms treatment in particular is widely cited in both academia and policy circles.

suggesting a potential role for government intervention to correct the market failure.[10] Further exacerbating inefficiency in this industry is the well-documented mismatch between wholesale and retail prices.[11] If consumers were to face a critical peak pricing (CPP) program similar to the one we implement, information feedback via an IHD would augment the efficiency gain. To assess the net benefit of a large-scale IHD deployment, we provide an estimate of the payback period from private and public investment in the feedback devices. In this context, an IHD pays for itself privately in 6-20 years on average. The social benefits of deploying IHDs to one-third of U.S. households yield an upper bound payback period of 7-9 years.

The paper proceeds as follows. Section 2 presents a simple conceptual framework to motivate the empirical approach. Section 3 describes the experimental setting and research design, and Section 4 lays out the data. Our empirical methods and results are presented and discussed in Section 5. Section 6 describes dynamic effects and Section 7 follows with a discussion of welfare and policy implications. Section 8 concludes.

2 Conceptual Framework

Some theoretical features of our setting motivate the empirical analysis. Consider a household that derives utility from the use of energy services such as air conditioning or illumination. The quantity of energy inputs (kWh) required to produce these services is not well understood. While consumers receive monthly bills summarizing aggregate monthly electricity usage, these bills do not detail energy usage by service, nor do they display the energy-intensity of these services. As a result, it is difficult if not impossible for even the most attentive consumer to perform a marginal cost-benefit analysis comparing energy services. Further, the disconnect between perceived usage and services causes consumers to hold incorrect beliefs about how changes in the production of energy services impact energy usage.

Assume that consumers have full knowledge about the marginal benefit of household energy services but are uncertain about the marginal costs. This initial uncertainty may be about the price of electricity, the quantity of electricity required to produce a service, or both. Ex ante (before the introduction of feedback), consumers equate the expected marginal benefit of each activity to its expected marginal cost. The expected marginal cost of an electricity service will depend on a consumer's prior beliefs about how each potential service maps into electricity usage. Of course, the prior may be incorrect; some activities may be more

[10] Greenwald and Stiglitz (1986) demonstrate that economies characterized by incomplete information are not, in general, constrained Pareto efficient.

[11] Examples include Borenstein 2002, Borenstein 2005, and Borenstein and Holland 2005.

energy-intense than others, relative to expectations.

Once a consumer is exposed to feedback, information about the mapping between usage and services is revealed. She will begin to learn how choices about electricity actions translate into usage and expenditure. Ex post, this information may reveal "mistakes" in optimization (even conditional on price), inducing adjustments. Over time, prior beliefs will be updated to incorporate information contained in the feedback, and in the limit the posterior distribution will converge to the true mapping.

As an example of the general case, consider a consumer whose electricity use is comprised of a fixed and variable component. The fixed portion of usage comprises baseline household services such as the continuous operation of the refrigerator and deep freezer, and some baseline level of temperature control and illumination. The discretionary component represents the margins of adjustment that are perceived to be available should a customer want to increase or decrease her electricity consumption.[12] Let X denote overall electricity use,

$$X = \bar{h} + h \tag{1}$$

where $\bar{h}$ and h describe the fixed and variable component respectively. Define the share of total usage that is discretionary as $\alpha \equiv \frac{h}{X}$, and assume that uncertainty only exists with respect to the value of α. For example, households may perceive standby power (also know as "phantom load") as fixed, whereas it is in fact variable.[13] This will cause them to under-estimate their true α.

We now posit certain features that likely describe the relationship between the price elasticity of demand for electricity, ε, and α. Let the demand elasticity be defined as a function $f(\alpha)$,

$$f(\alpha) = \varepsilon \equiv \frac{\partial logX}{\partial logp} \tag{2}$$

where p is the price of electricity. First, $f()$ will be monotonically increasing in the perceived level of α. As $\alpha \to 0$, demand becomes perfect inelastic since the perceived share of discretionary usage approaches zero. On the other extreme, as $\alpha \to 1$, all usage is perceived as discretionary, leaving all margins of adjustment available to the consumer when responding to price changes. As to the curvature of $f()$, it is intuitive that for "high enough" values of α, $f()$ will be locally concave. A higher fraction of discretionary use leads to

[12]In reality, the distinction between discretionary and non-discretionary usage is blurred. At a high enough price, all consumption is discretionary.

[13]Standby power is the electricity usage attributable to appliances that are plugged in and not turned on, but are still drawing load. "Phantom load" is estimated to account for 10 percent of U.S. residential electricity use. See http://standby.lbl.gov/ for more details.

more margins on which to respond to price changes, but the incremental contribution of these margins is decreasing.

To determine how behavior will change as information about the "true" value of α is revealed, one must consider prior beliefs about α. Denote by α^0 the true share of discretionary usage, and by $\tilde{\alpha}$ the consumer's uninformed belief. Depending on the initial relationship between the uninformed belief and the true share of discretionary usage, the effect of information on the price elasticity of demand is ambiguous. If $\tilde{\alpha} > \alpha^0$, then information that moves consumer beliefs closer to α^0 will decrease elasticity; if $\tilde{\alpha} < \alpha^0$, then information will cause the elasticity to increase. However, the magnitude of these effects is asymmetric. If the distribution of prior beliefs is symmetric around the truth (that is $E(\tilde{\alpha}) = \alpha^0$), then the average demand elasticity will increase in the presence of new information. Even this simplified and very stylized framework reveals the complexity of factors that may predict the role of information in price response. Thus, we believe an empirical inquiry is appropriate to answer the question: does information increase price elasticity?

3 Research Design

In partnership with a utility, we designed and implemented a randomized controlled trial that introduced short-term price increases and real-time information to a sample of residential electricity customers in Connecticut. The treatment events occurred in July and August of 2011. In this section, we describe the empirical setting, our sample, and the research design.

The United Illuminating Company (UI) is an investor-owned utility supplying 320,000 customers in the New Haven and Bridgeport areas of Connecticut. Currently, households in UI's territory are on one of two pricing plans: a flat rate of \$0.21/kWh or a time-of-use (TOU) rate which charges a "peak" price of \$0.27/kWh on weekdays from noon until 8 pm, and an off-peak price of \$0.18/kWh at all other hours. UI has been utilizing Smart Grid technology for decades. In the late 1970s, UI replaced traditional analog meters with automated meter readers (AMR) for all of their commercial and industrial customers, enabling it to become one of the first utilities to offer TOU pricing.[14] In 2010 the Department of Public Utility and Control (DPUC) provided funding (i) for the installation of the most recent generation of smart meters called Advanced Metering Infrastructure (AMI) and (ii) to test the potential of smart technologies, including

[14] AMRs are an early generation of the smart meter that allows for one-way communication of information, from the home to the utility.

in-home displays and other HAN devices to shed load and conserve energy. It is within the context of this initiative that we designed the experiment.

To be eligible for participation in the pilot a customer needed to reside in a townhouse or single family home, sign and return an end-user agreement indemnifying UI against litigation risk, and have a broadband internet connection.[15] To recruit such households into the HAN pilot, UI contacted 60,000 customers that had enrolled in paperless billing, indicating the likely presence of Internet in their home. As a participation incentive, the utility offered households $40: $20 upon completion of a pre-survey prior to assignment to treatment and $20 upon completion of a survey once the pilot ended. Of the households that agreed to participate in the pilot, 437 form the sample for this study.[16]

3.1 Treatments

We randomly assigned households to one of three groups: control, price, and price plus information ("price + IHD"). The cell size of each of these groups was determined by power calculations and cost of treatment. A total of 207 households were assigned to the control group, the least costly group. These households (and all others in the pilot) received a mailing notifying them that they were in the pilot and informing them of their group assignment. They were also mailed an energy conservation pamphlet documenting "101 Ways to Conserve Electricity".

Price Treatment: We assigned 130 households to a pricing treatment. These households experienced 2 types of pricing events that varied in magnitude of the price increase and timing of event notification. The first event type, "DA", provided day-ahead notification that the per-kWh price of electricity would be increased by $0.50. These events mimicked a pricing policy that a utility might use to incentivize electricity conservation when high temperatures are expected in the following afternoon. The second type of event, "TM", sent notification of a $1.25 increase in the per-kWh price of electricity thirty minutes before an event. A utility might implement this type of policy to reduce immediate risk in grid stability due to an unexpected decrease in generation (say, due to the failure of a baseload generating plant).[17] Ex-ante we cannot predict to

[15]Some aspects of device-to-utility communication were configured to occur wirelessly via the Internet. To maintain the integrity of our randomization, it was thus important that *all* of the participants have wireless. Any household that upon recruitment had a wireline broadband connection instead of a wireless connection was given a wireless router.

[16]This study represents a portion of a larger project that tests other aspects of the HAN. In total, 1,152 households participated in the project.

[17]On July 22, 2011, a prolonged heat wave on the east coast was occurring and the peak wholesale price on the

which type of event households will be more responsive. While TM events send a much stronger price signal, households may not be able to respond to the price change within the short window of advance notification.

Between July 2011 and August 2011, three DA and three TM pricing events occurred. While all events occurred during peak hours, there was variation in the length and timing of events. Table (1) provides a description of each event including the the start time, event duration, and the mean and high temperature of the day. Households received notification of these pricing events by email, text message or phone call, depending on their preference.[18]

Due to regulatory constraints, these price changes could not be reflected in the customer's monthly electricity bill. In addition, the utility wanted to ensure that no household was made worse off due to the price changes. To navigate these constraints while still achieving the intended marginal incentives, we provided customers assigned to the price treatment with an off-bill account initially credited with $100. The difference between the regulated price and the event price was multiplied by the quantity of electricity consumed during each event time period, and that amount was added to or subtracted from the household's off-bill account balance. By design, the marginal incentive facing consumers was equivalent to a change in their per kWh price. At the conclusion of the experiment, any balance remaining in the account was the customer's to keep.[19]

Price + Information Treatment: The 100 households assigned to this treatment group experienced the pricing program described above but also received real-time information about electricity prices, usage and expenditure. This information was provided via an in-home display (IHD), a portable device (which can be mounted on a wall or placed on a counter) that displays in real-time electricity consumption, prices and bills. Figures 1 and 2 provide photos of the IHDs used in this study. These households also received computer software enabling them to log in to a web portal to monitor and view electricity usage, prices and expenditure.

The main difference between this treatment and the price treatment is that customers are able to view in real time the price of electricity, electricity usage and their estimated monthly bill-to-date. The device removes

New England Independent System Operator's (NEISO) spot market climbed to nearly $0.60/kWh. If households on a flat rate faced a DA event on that day, the retail price would have been $0.71/kWh. A significant disruption of supply would have compounded strain on the grid, and prices may have easily approached levels on the order of magnitude of our experimental price changes.

[18] Upon enrollment in the pilot, households were asked to choose how they would like to be informed of pricing events should they be assigned to this treatment.

[19] The unique implementation of price changes may raise concerns about construct validity. However, the central result of the paper relates to the differential response to prices between households with and without information feedback. Any concerns about construct validity would apply equally to these two groups.

the information acquisition costs involved in translating electricity consuming actions into electricity usage and expenditure. The ability to view real-time usage provides customers with the opportunity to learn which appliances are heavy electricity users and which are not, thereby enabling them to more fully optimize in response to price changes.

Households received the IHDs and professional ("white glove") installation of them free of charge. The latter ensured that displays were set up correctly and activated, and maximized the likelihood that participants understood how to use them. A trained technician affiliated with the vendor scheduled an on-site appointment with households in this treatment group and installed the IHD. Aside from the installation and set-up of the IHD, the vendor provided no services.

4 Data

In this section we discuss descriptive statistics, the effectiveness of the randomization, and address potential compliance-related concerns.

To collect data on energy usage, meters capable of gathering high-frequency data on household electricity usage were installed in all participating households.[20] During the experiment, electricity usage of all participants (including those that dropped out of the study, i.e. "non-compliers") was collected at 15-minute intervals. The collection of these data for non-compliers allows us to estimate the intention-to-treat effect and the treatment effect on treated households. In addition to the usage data, the utility also provided us with information on a customer's rate class (flat rate or TOU rate), and whether he had central air conditioning.

All households in the experiment completed two surveys: one prior to assignment to treatment (the "pre-survey") and another upon completion of the pilot (the "post-survey"). These surveys collected data on household demographic characteristics, housing unit characteristics, appliance ownership, and conservation-related actions. In the pre-survey, we also asked respondents about their propensity to be at home during the day. In the post-survey, we surveyed households on their awareness and understanding of pricing events, and the frequency with which they checked their IHDs. These questions allow us to investigate potential mechanisms and explanations for our primary estimates.

[20]Households that already had the older and less advanced Automated Meter Reading (AMR) devices did not receive the new AMI installations. Instead, their meters were adjusted to collect the same high frequency data as the AMI meters.

4.1 Balance and Compliance

Table 2 presents descriptive statistics by treatment; Panel A includes all households who initially agreed to participate in the study and Panel B is restricted to households who completed the pilot ("compliers"). Thirty-eight households (approximately 9 percent) were non-compliers. These households either moved, requested to be removed from the study or failed to arrange an installation appointment for the IHD. Due to the difficulty of scheduling installation appointments for households assigned to the price+IHD group, attrition is not balanced across treatment groups. Of the 100 households assigned to the price+IHD treatment, 28 did not complete the study. In contrast, compliance was high in the non-technology groups (control and price-only). Only 4 of the 207 households assigned to the control group and 6 of the 130 households assigned to the price treatment moved or requested to be removed from the study. The concern raised by this asymmetry is that the success or failure to schedule an installation appointment is systematically related to the desire or ability to respond to treatment. Below we discuss how our empirical specifications account for this concern. First, though, we describe the sample.

As shown in Panel A of Table (2), peak hourly electricity usage averages at 1.42, 1.53 and 1.38 kWh in the control, price and price+IHD treatments, respectively. A comparison of mean usage between each treatment and control indicates that both peak and off-peak usage are statistically indistinguishable between groups if the sample includes households initially assigned to each treatment (see columns labeled "Difference"). When the sample is limited to compliers, mean off-peak usage significantly differs between households in the price and control groups.

Table (2) also reports home ownership, household income, square footage, and the age of a home. The sample varies across survey questions; notably many households failed to answer questions about the square footage and the age of home. In the initial sample, the average home is 1600 square feet in size and 54 years old. Approximately 77 percent of households own their home and median annual income is between $70,000 and $75,000, indicating that the sample in this study is wealthier than both the national average and the overall population served by the utility. In both the initial and final sample, demographic and household characteristics are balanced between each treatment and control group, with one exception. For both the initial and final sample, households in the price treatment tend to live in larger homes (t=2.10).

To formally test whether the randomization achieved a balance in observables across households, we estimate a linear probability model in which we regress each treatment indicator on the only pre-existing observable variables available for every household in our sample: mean peak electricity usage and rate class (flat rate versus time-of-use rate). While the surveys allow us to obtain much more detailed information on household characteristics, we do not use the survey data in our randomization checks, since survey compliance was not 100 percent and their inclusion would confound the interpretation.[21] We discuss survey compliance below, since our analysis of potential mechanisms explaining the treatment effect relies on these data.

The columns labeled "Initial Group" in Table (3) show results, where the sample is comprised of the control group and the price group in column 1, and the control group and price+IHD group in column 2. Neither electricity usage nor rate class is statistically significant in explaining assignment to the initial price and price+IHD treatments. These results indicate that households in each group are balanced on observables, providing evidence that households were randomly assigned to treatments. Still, to control for potential unobservable differences across households in the empirical specifications, we include an array of non-parametric controls including household fixed effects and aggregate time effects.

Attrition occurred in each group, but disproportionately so for the price+IHD group, simply because households in this treatment needed to successfully schedule and complete an installation appointment. Noncompliance may be systematic, and one can imagine that the same factors determining attrition may also impact response to treatment. For example, consider households with no one home during working hours. These households will have more difficulty scheduling an installation appointment, and may also be less likely to respond to price increases that occur during working hours.[22] If this is the case, then our estimated effect of the treatment may partly reflect that the households most capable of responding to price events are more likely to be in the price+IHD treatment.

We test for asymmetric attrition by estimating a linear probability model that regresses an indicator variable set equal to 1 for compliers and 0 otherwise on electricity usage and rate class (for each group separately). The coefficient estimates will be significant if attrition is systematically correlated with the observable regressors. Results are presented in columns 3 and 4 of Table (3). In column 3 the sample is comprised of households initially assigned to the price treatment and in column 4 the sample is restricted to households initially assigned to price+IHD. The significant coefficient on both peak usage and rate class for Price+IHD households suggests that selected attrition may be occurring. Again, household fixed effects will strip out

[21]We initially estimated a linear probability model that included survey demographics as regressors. Though it did not raise serious concerns about the randomization, we choose not to present it due to the impossibility of separating non-random assignment from survey attrition when attempting to interpret the estimates.

[22]In our study, installation appointments occurred Monday-Friday. No technologies were installed on the weekends.

any time-invariant unobservables (such as rate class). Still, in our analysis we use intention-to-treat (ITT) and treatment-on-the-treated (ToT) estimators to account for asymmetries in non-compliance.

While this research was designed to estimate the information effect, survey responses provide a starting point for understanding some of the mechanisms that may be driving this effect. If survey compliance were 100 percent, we could implement the same ITT and ToT approaches used in our main specifications, but it is not. Nearly all households complete the pre-survey, but there is substantial attrition across all treatment groups in completing the post-survey. Table (4) presents results from regressing pre- and post-survey compliance indicators on observables. With one exception, observables are not significant in explaining survey non-compliance. Post-survey compliers in the price treatment group have lower mean peak kWh usage than survey non-compliers. In our discussion of mechanisms, we later return to the issue of survey non-compliance and how it impacts the interpretation of our results.

5 Results: Information and Price Elasticity

We begin by plotting raw mean hourly electricity usage by treatment group on each of the six event days, as seen in Figures (3)-(8). The commonalities between these events are evident: households exposed to information feedback exhibit visibly lower usage during treatment events than price-only and control households. In these figures, mean 15-minute interval consumption is averaged across all households in each treatment group. The shaded area marks the period during which a pricing event occurred. In Figure (3), for example, a \$0.50 pricing event was held between 12 pm and 4 pm on July 21.[23] In the hours preceding an event, hourly electricity usage in the three groups is approximately the same, though this appears to change over time.[24] Once the event begins, we observe a divergence in usage between the price+IHD treatment and the other two groups. Households in the price+IHD treatment use less electricity (compared to the control group) in each hour of the pricing event. Interestingly, we do not observe significant differences in hourly usage between price and control households, despite households in the price-only group also receiving a price

[23] Recall that the only price change to which any household is exposed during this period is the experimental price change. All households are charged a time-invariant non-experimental price between the hours of 12pm and 8pm, including households on TOU (for whom peak hours are 12pm to 8pm as well).

[24] There is visual evidence of potential habit formation affecting non-event-hour usage, as seen in Figures (5), (7) and (8). Similarly, one may be curious about load shifting. Aside from the pre-treatment hour for the price-only group in Figure (4) and the post-treatment hour for the price+IHD group in Figure (6), the plots suggest that behavioral change is consistent with electricity conservation. In the section on Dynamic Effects, we empirically test for habit formation as well as load shifting.

increase.

Clear patterns in behavior emerge from the plots, but as shown in the summary statistics there is a great deal of cross-sectional heterogeneity. We control for these in empirical specifications that condition out time-invariant household characteristics and aggregate time shocks. Our preferred specification is a difference-in-differences model in which the dependent variable is the natural log of energy usage by household i in 15 minute interval t,

$$k_{it} = \beta_0 + \sum_{g \in \{P, P+I\}} \beta_g D_{it}^g + \gamma_i + \delta_{hd} + \mu_{it} \tag{3}$$

The explanatory variables of interest are the treatment group indicators, D_{it}^g, which are equal to 1 if household i is in group g, and if a pricing event occurs for i in interval t. Controls include household fixed effects (γ_i) and hour-by-calendar date dummies (δ_{hd}). As such, the coefficients of interest are being identified off of variation within households over time. Standard errors are clustered at the household to account for correlation across all observations within a household.

Tables (5) and (6) present the intention-to-treat (ITT) estimates and treatment effect on treated (ToT) estimates, respectively. The strength in these approaches is that they overcome concerns about selection, allowing us to retrieve internally-valid estimates. In each table we first show results from our preferred specification in which we control for calendar-hour and household fixed effects. Column 1 displays the average treatment effect of all pricing events. Column 2 highlights the effect of DA notification pricing events that increase the price per kWh by \$0.50 and column 3 presents the overall effect of TM notification pricing events that increase the price per kWh by \$1.25. To highlight the role of household and time controls, we also show results omitting household fixed effects (col. 6) and both household and time controls (cols. 4-5). In all specifications the sample is comprised of all peak hours in July and August, with the exception of column 4 where the sample is restricted to pricing event intervals.

5.1 Intention-to-Treat

In the ITT estimator, D_{it}^P and D_{it}^{P+I} in equation (3) denote initial assignment to the price and price+IHD treatments, irrespective of whether the household dropped out or was removed from the study. The coefficient in this specification are consistent estimates of the average percentage change in electricity usage from assignment (though not necessarily compliance) to the price and price+IHD treatments during pricing events. Results from the ITT regression are also informative due to the well-documented difficulties encountered by electric utilities in engaging customers in their various initiatives. They offer utilities and their regulators a measure of the overall benefits available were they to implement pricing events (with or without information displays).

As one might expect, an increase in electricity prices generally leads to a decrease in usage. However, neither the plots of daily consumption nor specifications without household fixed effects immediately reveals this pattern, underscoring the important role played by household fixed effects in the empirical specification.[25] The treatment differential seen in the daily plots is confirmed here: households with IHDs are significantly more responsive to price changes than those without. When DA and TM events are pooled together, households with an IHD reduce usage by almost 14 percent and are three times as responsive to pricing events as those without IHDs. In response to DA events, households in the price group reduce usage by 7 percent as compared to a 17 percent reduction for households in the price+IHD treatment cell. Coefficient point estimates and the 10 percentage point difference between them are both significant with over 90 percent confidence. These results provide strong evidence that the cumulative effect of real-time information feedback can meaningfully increase the price elasticity of demand.

Results also suggest that individuals are more responsive, both in economic and statistical significance, to pricing events that occur with advance notice. This is true despite the fact that the price increase in TM events is more than twice that for DA events ($1.25 as compared to $0.50). Even with strong financial incentives, with only 30 minutes of warning individuals may not have the ability to respond to a price change. While usage of households in both treatment groups is not statistically different from zero, the 9 percentage point differential response is estimated with over 85 percent confidence.

5.2 Treatment on Treated

The ToT specification identifies the effect of the price and price+IHD treatments on compliers, while accounting for the effects of heterogeneous attrition across groups. These estimates are most relevant for testing if customers that face a sharp reduction in information acquisition costs will be more price elastic. To implement the ToT approach, we use initial treatment assignment as instruments for final treatment status and estimate the model using two-stage least squares. These instruments are both valid and strong. Initial group assignment was randomized, and compliance in each group was high: 98 percent, 95 percent,

[25] Household fixed effects soak up approximately half of the variation in the data, as can be seen by the R-square statistics.

and 72 percent in control, price-only, and price+IHD, respectively.

Estimates of ToT specifications with various controls are presented in Table (6). The qualitative pattern mirrors the ITT results, with a strong treatment differential between price-only and price+IHD households. The magnitudes are slightly larger as a result of estimating the percentage change in usage from treatment rather than initial assignment to treatment. Similar to earlier results, after controlling for hour-by-calendar day and household fixed effects, households with IHDs are three times as responsive to pricing events as those without, where this treatment differential is present with 95 percent confidence. Again for DA notification events, households in both treatment groups reduce usage in response to price increases. Households are less responsive to TM events than events with longer notice.

Results from both ITT and ToT approaches suggest that while households reduce electricity usage in response to price increases, the absence of information feedback inhibits customers from fully responding to them. In terms of elasticities, prices increased by approximately 200 percent during CPP events, and by 600 percent during DR events. In both cases the implied arc elasticities are low (less than -0.12). However, it is unclear as to what information is actually conveyed by these statistics; the range of the price increases may contain regions of highly elastic demand as well as regions of inelastic demand, allowing for the possibility that lower (absolute value) price changes may induce the same absolute behavioral response. In any case, these estimates are a lower bound for reasons that we will discuss in the section on dynamic effects.

5.3 Secondary Results and Robustness

The primary result in this paper is that information feedback measurably increases price responsiveness. Given the randomization, our two main empirical approaches yield internally-valid point estimates. We now turn to explore the mechanisms influencing these results. What is it about peoples' interaction with the IHDs that leads to the large treatment differential? Several behavioral and practical hypotheses are reasonable, and we test some of them in this section. To do this, we augment the high-frequency meter data with responses collected during our pre- and post-surveys.

Due to our reliance on survey data, the sample for this analysis is restricted to the intersection of initial treatment assignment and either pre- or post-survey compliance. As a result, the estimates may be vulnerable to systematic differences in survey non-compliance across groups, and should be interpreted with this caveat in mind. However, given the high survey response rate, we believe these results provide a starting point to understanding consumer behavior. We hope that insights from this section form the basis for future research, including experiments designed to distinguish between subtle behavioral mechanisms (which ours was not).

One hypothesis is that IHDs may increase awareness of electricity prices and price changes, causing consumers to overcome inattention. Alternatively, the treatment effect may be explained by differences in household composition across treatments cells, notably the frequency with which households are at home or have central air conditioning. An additional hypothesis is that experience with the IHDs facilitates some sort of "learning" about energy choices and the mapping from usage to expenditure. This last mechanism is consistent with a framework in which information about quantity contributes to the treatment differential. In this section we present evidence against the price awareness hypothesis, show that central AC and the frequency of being at home are correlated with consumer response, and suggest that learning through experience plays a role in the treatment differential.

5.3.1 Awareness

A necessary and sufficient condition for consumer optimization is the equality of marginal utility and marginal cost. When there is doubt about the price of a good, the optimization is imperfect since consumers are forced to approximate or to rely on heuristics. Economists are interested in price salience to the extent that it causes consumers to be aware of prices. That is, salience is viewed as a sufficient condition for awareness. In this study, we skip directly to the central issue of whether households are aware of prices, and the impact of this awareness on behavior.

The experimental design sought to make electricity price changes salient to all consumers who were exposed to them. The utility sent all customers in the price and price+IHD treatments either a text message, email or phone call either the day before or 30 minutes preceding each pricing event. This message alerted households that a pricing event was going to occur and informed customers as to the price. However, customers may not have been aware of pricing events if for some reason these notifications were not received. To get a sense of the effectiveness of our notification efforts and to quantify the importance of awareness, we asked in our follow-up survey if customers were "aware of pricing events while they were occurring?".

Conditional on answering this question, column 1 ("Number of HHs") of Table (7) shows that 89 and 95 percent of respondents in the price-only and price+IHD groups, respectively, were aware of pricing events while they were occurring. The awareness of these events was high, suggesting that the treatment events were front-of-mind. If the differential response exists because IHDs increase awareness of price events, then we would expect to see two things. The awareness of events would be higher in the price+IHD group, and conditional on awareness, households in the two groups would respond equivalently. To test this hypothesis, we estimate

$$k_{it} = \beta_0 + \sum_{g \in \{P+I\}} \beta_g D_{it}^g * A_i + \gamma_i + \delta_d + \sigma_{hw} + \mu_{it} \tag{4}$$

in which we interact the treatment indicator D_{it}^g with A_i, a variable set equal to 1 if a household responded that it was aware of pricing events. The overall results ("All events") are presented in column 2 of Table (7).

It is clear that awareness of the event is important, as households aware of treatment events when they occur exhibit a larger response. For households with an IHD, usage reduces by 19 percent for "aware" households, as compared to 7 percent for "unaware" households. We may reject the null that these coefficient estimates are equal with 85 percent confidence. Households in the price-only group appear to be more responsive if they are aware of pricing events, though we cannot reject the null that the point estimates or the treatment differential are significantly different from zero.

However, awareness alone cannot explain the role of the IHD in increasing responsiveness. Conditional on awareness, if we compare the coefficient estimates across treatments groups we find that households with IHDs are almost three times as responsive to pricing events. Equality of these coefficients is rejected with 95 percent confidence. The differential response suggests that IHDs trigger or facilitate a mechanism beyond price awareness, and it is other factors that allow customers to be more responsive to price changes.

5.3.2 Experience

One possible factor is that IHDs facilitate learning. Physical interaction with the IHD informs households about the mapping between electricity consuming actions and usage. That is, if a consumer views the display before and after turning on her air conditioner, she will notice that it consumes a large amount of electricity. This knowledge, if accumulated for multiple appliances, may better equip households to respond to price changes. As evidenced by the economic and statistically significant treatment effects for households in the price-only group, even without IHDs, households are aware of some of the margins of adjustment available to them. However, they may not understand the full choice set of energy saving behaviors, and the quantity of electricity usage associated with these.

To test if frequent experience with the IHD increases price responsiveness, we use survey responses. In the post-survey, we ask households "How many times per week did you look at the IHD in the first month that it was activated?"[26] Table (8) shows that most households frequently experimented with their IHD in the first month that they received it. Conditional on answering the question, only 6 percent of households did not look at it and approximately 65 percent of households looked at it more than 5 times. This suggests that households with an IHD are actively engaging with it, even before the start of pricing events.

In columns 2-4, we present results from the estimation of equation 4, where A_i is now a vector of indicator variables describing how frequently a household looks at the IHD. Households who engage most frequently with the IHD are significantly more responsive to price changes than others.[27,28] This suggests that more frequent experience with the IHDs facilitates learning, perhaps about the quantity of electricity consumption. This is consistent with households gaining a fuller understanding about what is discretionary and what is not, or about the mapping of behavior into electricity usage.

5.3.3 Home During Daytime

To remain in the study, households assigned to the information treatment needed to successfully schedule and install an IHD. A concern is that households who are not at home are less likely to have an IHD installed, and may be less responsive to events. While our ITT and ToT estimators allow us to achieve internal consistency, we are still interested to see how the response to pricing events varies by how frequently an individual is at home. To do this, we rely on pre-survey data that asked households if someone was generally at home during the day, after school or after work, and create indicator variables that we interact with assignment to treatment.[29] Given that almost all households completed the pre-survey, including those who failed to have

[26]Note that IHDs were installed two to four months before the first pricing event occurred, leaving ample time for households to interact with them before the treatment events occurred.

[27]It is possible that the direction of causation is such that those inherently more responsive to price also look at the IHDs more frequently. However, if this inherently more responsive cohort exists, then we would expect them to occur with similar frequency in the price-only group. Given that we observe a differential response across those with and without IHDs, there is little evidence to support a hypothesis about reverse causality.

[28]Households reporting to never look at the displays have large measured responses as well. However, these are few in number and there are several reasonable explanations. It is possible that the respondent is not the person responsible for the energy decisions that lead to the large measured treatment effects. Also, perhaps instead of using the IHD, these households log into the web portal, or already were familiar with all the margins of adjustment available to them (whereby the IHD provides no additional value). Or, when notified of pricing events, these households may simply leave the premises for the duration of the price increase.

[29]It should be noted that this survey response does not reflect whether households were actually at the home during pricing events, but rather provides a coarse measure on the frequency with which households tend to be home during the day.

IHDs installed, the first specification in which the at-home dummy variables are interacted with treatment assignment are very similar to an ITT estimator. In a separate specification, we use a ToT estimator.

As shown in the first two columns of Table (9), the feedback group has fewer households with someone home during the day than price-only, both before and after accounting for attrition. Slightly over 50 percent of the households assigned to the price-only group have someone home during the day, as compared to 40 percent for the price+IHD treatment. Interestingly, when we restrict the sample to compliers this pattern remains, implying that our concerns about installation-related attrition may not be as closely correlated to participants' schedules. Conditional on completing the pilot, 44 percent of households in the feedback treatment and 52 percent of households in the price treatment tend to have someone home during the day. If households with someone home during the day are more responsive to pricing events, then the differential response attributable to information feedback will be, if anything, attenuated.

Results suggest that having someone present in the household is not driving the differential effect between those with and without feedback. Conditional on someone tending to be home during the day, households with an IHD are more responsive to price changes, where the difference in treatment effects ranges from 8 to 12 percentage points in the ITT specifications. The treatment differential is larger in the ToT specifications, ranging between 9 and 17 percentage points, and distinct with roughly 90 percent confidence for day-ahead pricing events.

Within each treatment cell we observe some interesting differences in the response to pricing events based on household schedule. As expected, households that tend to have someone at home during the day are more responsive (as compared to households with someone home after work) to the short-notice TM pricing events. Households with someone home after school appear to be similar to those home during the day for these events, though "after school" may have a different meaning in different households. These results suggest that short notification events inhibit households that are not at home during the day from responding to them.

5.3.4 Central Air Conditioning

The estimated reduction in usage during pricing events and the heightened response for those with feedback may be related to the presence of central air conditioning. One hypothesis is that households responding to pricing events are those with central AC, a high usage activity that can easily account for large usage reductions when turned off (or down) during events. Table (10) shows the percentage of households with central AC and reports treatment effects for households with and without it.

Results are contrary to our initial expectations. Households without central AC are on average more responsive to pricing events. This implies that households exhibiting large responses are doing so by adjusting other, likely smaller, margins of response. This result is once again consistent with the learning hypothesis, though it is by no means dispositive. There may also be a selection effect at work: households with central AC may be different in an observable way that is correlated with treatment response. The latter is consistent with, say, younger households being more interested in energy conservation, but living in dwellings less likely to have central AC.

6 Dynamic Effects

In this section we explore two views of dynamic household usage, one roughly corresponding to the medium-term (weeks to months) and the other examining the short-run (minutes to hours). Insofar as households are conserving outside of the treatment window, this would attenuate our treatment effects by lowering the baseline for within-household comparison. In addition, this broader view of behavior will provide a deeper understanding of household response and has environmental implications. In the medium run, as weeks pass and households are exposed to multiple events they may learn to form habits that cause their usage patterns to decrease in meaningful ways during non-event days. In the short run, there are competing hypotheses. Usage in hours on either side of the event window may increase, relative to normal, which is consistent with households pre-cooling their homes before anticipated price increases or delaying activities (e.g. laundry) until the price returns to its normal level. Alternatively, activities undertaken in response to high prices may spill over into the hours preceding or following the events. If true, this would manifest as lower usage.

Habit formation is responsible for large conservation effects on non-event days over the months of our study. Table (12) reports coefficients from a single regression in which we estimate separate trends for each hour of the peak period over the 62 days of July and August. For example, row 1 displays the trend in kWh for the period noon to 1pm, implying an average daily decrease in usage of 0.20 percent for price-only households during this noontime hour. This corresponds to a 13 percent decrease in usage from July 1 to August 31. At the upper end of the coefficient range conservation effects are equivalent to 22 percent during 5-6pm for the price+IHD group. We also find that the conservation trend is steeper for price+IHD households during all hours. The maximum differential between households with and without feedback implies 6 percent more conservation over these months. The magnitude of this result is very large when measured against the backdrop of other treatments seeking to achieve conservation.

In the short run we find that conservation exhibited during pricing events spills over in to adjacent hours for the price+IHD group. Table (11) reports estimates from the ITT specification, but with the addition of indicator variables for the two hours preceding and following price events for each treatment group. The results clearly do not support the so-called "pre-cooling" hypothesis. Rather, in the hours preceding an event, households exposed to day-ahead notification exhibit no change in usage while those with feedback significantly reduce usage by 9%. We also find that with advanced notification the treatment effects spill over into the two hours following an event. In contrast, we find no significant evidence of load shifting or spillovers in response to TM events, probably because of the limited amount of time to anticipate them. These results imply that households are not shifting load to the hours immediately preceding or following events. If anything, the spillovers that accompany DA events are consistent with the habit formation results discussed above.

The dynamic effects inform the study in three important ways. Evidence of habit formation lends further support to the hypothesis that learning is taking place, and that the IHDs facilitate this learning. Second, the identification of our baseline effects comes from variation within households, so reductions in usage outside of the treatment windows will bias the primary treatment effects towards zero. Further, since habit formation and spillovers are stronger for price+IHD households, the treatment differential will also be attenuated, implying an even larger true effect of information feedback. Finally, these conservation effects imply a favorable result with respect to greenhouse gas emissions.

7 Welfare and Policy Implications

The assumption of full information underpins countless important economic results about market efficiency, including any application of the first and second fundamental theorems of welfare.[30] It is not surprising, then, that the absence of full information may lead to inefficient market outcomes. The main result of this paper points to the importance of information, whereby consumers with additional information about the quantity of consumption make vastly different decisions than their uninformed counterparts. This finding

[30] See Greenwald and Stiglitz (1986).

may extend to other settings in which quantity is also shrouded from the consumer such as cellular phone minutes, Internet or data, water, heating oil and natural gas usage.

The electricity sector exhibits an additional inefficiency: the disconnect between wholesale and retail prices.[31] Time-varying electricity prices that better reflect generation costs would mitigate short-run inefficiencies by aligning social cost with marginal benefit, reduce over-investment in capital in the long run, lessen the harm from market power in the generation sector, and (depending on the location and timing of consumption) potentially reduce the external cost of emissions (Borenstein 2002, Borenstein 2005, Borenstein and Holland 2005, Holland and Mansur 2008).[32] A historical barrier to the implementation of time-varying retail prices was the ability to collect high-frequency usage data and transmit this information to customers (Joskow and Wolfram 2012). The device used in our study removes this barrier, and is precisely the type of technology being considered by regulators to facilitate dynamic retail electricity pricing. Our results highlight the potential of time-variant pricing to reduce usage, especially when coupled with information feedback.

Now that the technology exists to implement time-variant pricing, electric utilities and their regulators are trying to understand whether to deploy HAN technology in the field, and if so, the appropriate payment structure and nature of deployment. In this section, we first estimate the expected net private benefit to consumers of purchasing an IHD, and then shift our focus to estimating net social benefits. The magnitude of net benefits accruing to households provides insight into the likelihood of a successful free-market solution in which households voluntarily purchase HAN technology for their homes.

Assuming that all households are placed on a time-variant tariff (similar to the one we implemented), we now use our treatment effect on treated households to quantify the potential savings to households of owning an IHD. Table 13 shows the change in expenditure from price increases of either \$0.50/kWh or \$1.25/kWh with and without an IHD. The difference between the Price and Price+IHD columns represents the incremental value of an IHD. We vary the number of hours - 20 hours, 40 hours and 60 hours - that households are exposed to price increases. If a household is exposed to 40 hours of \$0.50/kWh or \$1.25/kWh prices respectively, the IHD would facilitate \$3-\$8 in savings for the household. Note that this is an overestimate of the marginal net benefit of utilizing the IHD since we omit from our measure of cost any disutility incurred by households from altering their behavior. Currently, a basic IHD is available at retail stores for just under \$40,[33] though

[31] There are also others, most notably market power and external costs of pollution.

[32] Other long-run benefits include a reduction in ramping costs, as well as potential savings in transmission and distribution infrastructure.

[33] http://www.amazon.com/Black-Decker-EM100B-Energy-Monitor/dp/B001ELJKLE

many more expensive models exist (including the ones used in this study). According to our estimates, at a 5 percent discount rate, the payback period is over 20 years for \$0.50/kWh events, but under 6 years for \$1.25/kWh events. Even if these payback periods were not underestimates, it seems unlikely that the "average" electricity consumer would find the voluntary purchase of an IHD an attractive investment.

Since many of the benefits of energy conservation are external to the consumer, we also examine the net social benefit of deployment, accounting for some of the general equilibrium effects of IHDs and dynamic pricing on the need for excess generation capacity. This estimate will be informative as to the role of regulation and government intervention in HAN deployment. However, the net social welfare gains are more difficult to quantify. Borenstein and Holland (2005) estimate that the industry-wide deadweight loss (excluding external environmental costs) from flat rate electricity pricing is on the order of 5-10 percent of the wholesale electricity market. In 2009, when \$350 billion of electricity was sold in the United States, this translates into \$17-\$35 billion dollars, a portion of which is attributable to the residential sector. The question is, how much of this deadweight loss would be remedied by the presence of both CPP (or a similar dynamic pricing scheme) and feedback technology (IHDs)?

To get a sense for this, we apply the price elasticities estimated in our setting to a previously developed simulation evaluating the long-run welfare gains from switching one-third of residential customers in the U.S. onto RTP (Borenstein 2005). While this is a different flavor of dynamic pricing than the one we implement, a well-executed CPP program targeting the most grid-strained days will likely yield a majority of the benefits that a RTP scheme would. Using the price elasticities in our study, we find that placing one-third of customers on RTP would increase annual total surplus by \$200 to 250 million (in year 2000 dollars).[34] Approximately 35 to 50 percent of these gains would accrue to residential customers remaining on the flat-rate tariff, since they now face lower average electricity prices because of the efficiency gains. To deploy IHDs to one-third of households nationwide would cost approximately \$1.4 billion dollars (35 million households * \$40/IHD), implying a payback period of between 7 and 9 years.[35] This payback period is an upper-bound estimate since the simulations in Borenstein (about which he is perfectly transparent) do not consider market power, ramping costs, and transmission and distribution costs. There may also be economies of scale in the procurement of the feedback technology, reducing the up-front cost of deployment. Finally, since the exhibited responses to feedback are conservationist as opposed to load shifting, there will

[34]This turns out to be true in the simulations when starting from a very low base elasticity (-0.050) or a somewhat higher one (-0.100).

[35]The payback period assumes a discount rate of 3 percent, and annual gains of either \$200 million (9 years) or \$250 million (just over 7 years).

be environmental gains from a decrease in generation-related pollution.

While our estimates are internally valid, our sample is select in that it is comprised of households with broadband that volunteered to participate in this study. Our results must be interpreted in this context, and their external validity depends on the setting being considered. That said, the main qualitative result may be broadly applicable to other locations and settings, most clearly those in which choices ought to incorporate a metered quantity (e.g. cellular phone, water, or cellular data usage).

To help gauge the external validity of our results and the potential of a national deployment of HAN technology, we estimate the ITT and ToT specifications, weighting our observations to reflect the income distribution in the United States. Consistent with our setting, we assume that this deployment will target technologically advanced households, specifically those with broadband Internet. This is quite a large proportion of the U.S. population; over 77 percent had Internet access in 2009 and over 27 percent of the population has broadband.[36] Results are shown in Table (14). They are qualitatively similar to the preferred baseline estimates, and if anything show a stronger treatment effect in IHD households.

8 Conclusion

This paper shows that information feedback about quantity of consumption increases the price elasticity of demand, using the residential electricity sector as the field setting. Random assignment of households to a control group, price treatment and price-plus-information treatment allows us to identify both the price elasticity as well as the marginal contribution of information feedback. Households only experiencing price increases respond to them, reducing usage by 0 to 7 percent. But those exposed to both price increases and information feedback produce a 8 to 22 percent reduction. The incremental contribution of feedback is both economically and statistically meaningful, inducing a three-fold increase in the behavioral response. By industry standards, the magnitude of these effects is remarkable and points to a level of customer engagement not often seen.

The usage reductions extend beyond the pricing event windows, to both non-event hours on event days and non-event days. For households in the price+IHD group, energy conservation spills over into the hours immediately preceding and following events. The trend over July and August suggests that households,

[36]International Telecommunications Union.

particularly in the price+IHD group, are forming conservation habits. The combined effects of spillovers and conservation habit formation cause us to underestimate our primary treatment effects. Further, since these dynamic effects are more pronounced for feedback households, we will also understate the incremental information effect.

The totality of evidence lends support to the hypothesis that information feedback facilitates learning. First, customers who more frequently look at their IHD are significantly more responsive to pricing events. It is natural to expect that frequent experience with these displays may lead customers to learn about electricity usage, particularly the relationship between electricity services and their cost. Second, we find that informed households are achieving conservation in ways not confined to central AC, the most energy-intensive use. In fact, IHD households without central AC are on average more responsive to price changes that those with it. This is consistent with a story by which information feedback allows people to learn about many available margins of adjustment. Lastly, the formation of stronger conservation habits on non-event days for feedback households also suggests that learning is occurring and that it is facilitated by the IHDs.

These results confirm the practical importance of one of economics' most ubiquitous assumptions – that decision-makers have perfect information. Indeed, the absence of perfect information is likely to cause substantial efficiency losses both in this setting and others in which quantity is also infrequently or partially observed by decision-makers. Under conservative assumptions, our results imply that a widespread deployment of information feedback technology would pay for itself in no more than 7-9 years in the residential electricity setting. An additional social benefit of feedback comes in the form of greenhouse gas abatement, with IHDs inducing up to a 6 percent conservation effect in the medium-run. Our study suggests price remains a promising lever by which to influence electricity demand when consumers are well-informed about their own usage.

References

[1] Allcott, H. (2011b). "Rethinking Real-Time Electricity Pricing," *Resource and Energy Economics*, 33(4): 820-842.

[2] Allcott, H. (2011a). "Social Norms and Energy Conservation," *Journal of Public Economics*, 95(9-10): 1082-1095.

[3] Borenstein, S. (2002). "The Trouble with Electricity Markets: Understanding California's Restructuring Disaster," *The Journal of Economic Perspectives*, 16(1): 191-211.

[4] Borenstein, S. (2005). "The Long-run Efficiency of Real-time of Electricity Pricing," *The Energy Journal*, 26(3): 93-116.

[5] Borenstein, S. and S. Holland (2005). "On the Efficiency of Competitive Electricity Markets With Time-Invariant Retail Prices," *Rand Journal of Economics*, 36(3): 469-493.

[6] Chetty, R., A. Looney, and K. Kroft (2009). "Salience and Taxation: Theory and Evidence," *American Economic Review*, 99(4): 1145-1177.

[7] DellaVigna, S. and J. Pollet (2008). "Investor Inattention and Friday Earnings Announcements," *Journal of Finance*, 64: 709-749.

[8] DellaVigna, S. (2008). "Psychology and Economics: Evidence from The Field," *Journal of Economic Literature*, 47: 315-372.

[9] Faruqui, A. and S. Sergici (2010). "Household Response to Dynamic Pricing of Electricity: A Survey of the Experimental Evidence," *Journal of Regulatory Economics*, 38: 193-225.

[10] Faruqui, A., S. Sergici, and L. Akaba (2012). "Dynamic Pricing in a Moderate Climate: New Evidence from Connecticut," Working Paper.

[11] Feng, L. and M. Seasholes (2005). "Do Investor Sophistication and Trading Experience Eliminate Behavioral Biases in Financial Markets?" *Review of Finance*, 9: 305-351.

[12] Greenwald, B. and J. Stiglitz (1986). "Externalities in Economies with Imperfect Information and Incomplete Markets," *Quarterly Journal of Economics*, 101 (2): 229-264.

[13] Grubb, M. and M. Osborne (2012). "Cellular Service Demand: Biased Beliefs, Learning, and Bill Shock," Working Paper.

[14] Harding, M. and D. Rapson (2012). "Does Absolution Promote Sin? Green Marketing and Energy Demand," Working Paper.

[15] Herter Energy Research Solution (2012). "SMUD's Residential Summer Solutions Study," Working Paper.

[16] Hirshleifer, D., S. Lim and S. H. Teoh (2009). Driven to Distraction: Events and Under-reaction to Earnings News, *Journal of Finance*, 64(5): 2289-2325.

[17] Hossain, T. and J. Morgan (2006). ". . . Plus Shipping and Handling: Revenue (Non) Equivalence in Field Experiments on eBay," *B.E. Journals in Economic Analysis and Policy: Advances in Economic Analysis and Policy*, 6(2): 1-27.

[18] Holland, S. and E. Mansur (2008). “Is Real-Time Pricing Green? The Environmental Impacts of Electricity Demand Variance,” *The Review of Economics and Statistics*, 90(3): 550-561.

[19] Ito, K. (2011). “Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing,” Working Paper.

[20] Jacobsen, G., M. Kotchen, and M. Vandenbergh (2010) “The Behavioral Response to Voluntary Provision of an Environmental Public Good: Evidence from Residential Electricity Demand,” *European Economic Review*, 56: 946-960.

[21] Joskow, P and Wolfram, C. (2012) “Dynamic Pricing of Electricity,” forthcoming, *American Economic Review, Papers and Proceedings.*

[22] Lacetera, N., D. Pope, and J. Sydnor (2012). “Heuristic Thinking and Limited Attention in the Car Market,” forthcoming, *American Economic Review.*

[23] List, J., (2003). “Does Market Experience Eliminate Market Anomalies?” *Quarterly Journal of Economics*, 118(1): 41-71.

[24] Reiss, P. and M. White (2005). “Household Electricity Demand Revisited,” *The Review of Economic Studies*, 72: 853-883.

[25] Simon, H. (1955), “A Behavioral Model of Rational Choice,” *Quarterly Journal of Economics*, 69(1): 99-118.

Tables and Figures

Table 1: Treatment Events

Event Date	Desc	Type	Start Hour	High Temp	Mean Temp	Humidity
07/21/11	4 hr $0.50	DA	12	89	82	75
07/22/11	4 hr $1.25	TM	12	103	90	61
08/04/11	2 hr $0.50	DA	15	80	74	68
08/10/11	2 hr $1.25	TM	16	88	80	63
08/17/11	2 hr $1.25	TM	16	86	75	64
08/26/11	4 hr $0.50	DA	12	84	78	69

Table 2: Summary Statistics by Control and Treatment Group

Panel A: Initial Group

	Control		Price			Price+IHD		
	Mean	Obs	Mean	Obs	Difference	Mean	Obs	Difference
Off-peak usage (kWh/h)	1.159	207	1.279	130	0.121	1.203	100	0.044
	(0.687)		(0.737)		(1.524)	(0.646)		(0.542)
Peak usage (kWh/h)	1.422	207	1.529	130	0.107	1.383	100	-0.038
	(1.107)		(1.034)		(0.887)	(0.954)		(-0.298)
TOU Rate (1=yes)	0.184	207	0.200	130	0.016	0.240	100	0.056
	(0.388)		(0.402)		(0.373)	(0.429)		(1.153)
Home ownership (1=yes)	0.768	203	0.798	129	0.030	0.773	97	0.005
	(0.423)		(0.403)		(0.641)	(0.42)		(0.091)
Annual income ($1000)	72.00	203	74.00	129	2.000	71.00	97	-0.001
	(29.00)		(29.00)		(0.690)	(31.00)		(-0.181)
Home size (1000 square feet)	1.529	189	1.880	119	0.351 **	1.451	91	-0.078
	(1.11)		(1.83)		(2.100)	(1.14)		(-0.550)
Age of home (years)	52.423	156	57.619	97	5.195	52.239	71	-0.184
	(30.29)		(31.34)		(1.309)	(26.94)		(-0.044)

Panel B: Final Group

	Control		Price			Price+IHD		
	Mean	Obs	Mean	Obs	Difference	Mean	Obs	Difference
Off-peak usage (kWh/h)	1.161	203	1.294	124	0.121 *	1.202	72	0.044
	(0.69)		(0.73)		(1.52)	(0.62)		(0.542)
Peak usage (kWh/h)	1.432	203	1.551	124	0.107	1.432	72	-0.038
	(1.11)		(1.03)		(0.89)	(0.96)		(-0.298)
TOU Rate (1=yes)	0.182	203	0.202	124	0.016	0.181	72	0.056
	(0.39)		(0.40)		(0.37)	(0.39)		(1.153)
Home ownership (1=yes)	0.774	199	0.821	123	0.030	0.855	69	0.005
	(0.42)		(0.39)		(0.64)	(0.36)		(0.091)
Annual income ($1000)	72.00	199	75.00	123	0.002	76.00	69	-0.001
	(29.00)		(28.00)		(0.69)	(28.00)		(-0.181)
Home size (1000 square feet)	1.541	185	1.908	114	0.351 **	1.611	66	-0.078
	(1.10)		(1.84)		(2.10)	(1.16)		(-0.550)
Age of home (years)	52.221	154	56.574	94	5.195	53.375	56	-0.184
	(30.43)		(31.02)		(1.31)	(28.59)		(-0.044)

*Notes: Means are reported by treatment with standard deviations in parantheses below. Difference displays the difference in means between each treatment and control group with T-stats reported in parantheses below; * significant at the 0.10 level, ** significant at the 0.05 level, *** significant at the 0.01 level.*

Table 3: Group Assignment Balance on Observables, Initial and Compliers

	Initial Group		Compliers	
	Price	Price + IHD	Price	Price + IHD
Mean Peak kWh	0.022	-0.019	0.024	0.082*
	(0.027)	(0.027)	(0.021)	(0.049)
TOU Rate (1=yes)	0.002	0.094	-0.020	-0.294***
	(0.074)	(0.071)	(0.053)	(0.108)
F-Statistic	0.392	0.921	0.701	4.040
P-Value	0.422	0.473	0.246	0.095
N	337	307	130	100

*Results denoted "Initial Group" from a linear probability model regressing observables on the treatment group indicator. Results denoted "Compliers" from a LPM regressing observables on a compliance indicator. P-Value corresponds to probability that coefficients are jointly equal to zero. Control group used as control in each specification. Standard errors in parantheses. * denotes significant at the 0.10 level, ** significant at the 0.05 level, *** significant at the 0.01 level.*

Table 4: Survey Compliance by Treatment Group

	Pre-Survey		Post-Survey	
	Price	Price + IHD	Price	Price + IHD
Mean Peak kWh	0.006	-0.013	-0.072*	-0.045
	(0.009)	(0.019)	(0.043)	(0.049)
Rate RT (1=yes)	0.002	-0.006	-0.016	0.032
	(0.022)	(0.043)	(0.110)	(0.108)
F-Statistic	0.368	0.290	1.996	0.419
P-Value	0.486	0.511	0.098*	0.362
N	130	100	130	100

*Results from a linear probability model regressing observables on pre- and post-survey compliance indicator. Sample in each specification is the initial treatment group. P-Value corresponds to probability that coefficients are jointly equal to zero. Standard errors in parantheses. * denotes significant at the 0.10 level, ** significant at the 0.05 level, *** significant at the 0.01 level.*

Table 5: Intention-to-Treat Effects

	Preferred Specification (Full Sample, Controls)			Treatment Window Only	Full Sample	
	All Events	$0.50	$1.25	All Events	All Events	All Events
Event Type:	Pooled	Day Ahead (DA)	30min (TM)	Pooled	Pooled	Pooled
Price Only	-0.038	-0.071*	0.006	-0.008	0.127*	0.024
	(0.036)	(0.042)	(0.044)	(0.108)	(0.074)	(0.108)
Price + IHD	-0.137***	-0.171***	-0.084	-0.123	0.044	-0.067
	(0.046)	(0.051)	(0.057)	(0.108)	(0.079)	(0.106)
Prob(P = P+I)	0.044**	0.066*	0.130	0.325	0.471	0.429
Hour-by-day FEs	Y	Y	Y	N	N	Y
HH FEs	Y	Y	Y	N	N	N
Number of Events	6	3	3	6	6	6
Number of HHs	437	437	437	437	437	437
R-Square	0.58	0.58	0.58	0.00	0.00	0.05

*Dependent variable: log(kWh). * Significant at the 0.10 level, ** Significant at the 0.05 level, *** Significant at the 0.01 level. Standard errors clustered at the household level.*

Table 6: Treatment Effect on Treated Households

	Preferred Specification (Full Sample, Controls)			Treatment Window Only	Full Sample	
	All Events	\$0.50	\$1.25	All Events	All Events	All Events
Event Type:	Pooled	Day Ahead (DA)	30min (TM)	Pooled	Pooled	Pooled
Price Only	-0.040	-0.074*	0.007	-0.008	0.132*	0.024
	(0.037)	(0.044)	(0.046)	(0.112)	(0.076)	(0.111)
Price + IHD	-0.170***	-0.217***	-0.100	-0.157	0.055	-0.083
	(0.057)	(0.064)	(0.067)	(0.138)	(0.098)	(0.131)
Prob(P = P+I)	0.023**	0.025**	0.115	0.276	0.561	0.411
Hour-by-day FEs	Y	Y	Y	N	N	Y
HH FEs	Y	Y	Y	N	N	N
Number of Events	6	3	3	6	6	6
Number of HHs	437	437	437	437	437	437
R-Square	0.58	0.58	0.58	0.00	0.00	0.05

*Dependent variable: log(kWh). * Significant at the 0.10 level, ** Significant at the 0.05 level, *** Significant at the 0.01 level. Standard errors clustered at the household level.*

Table 7: Heterogeneous Treatments: Awareness of Treatment Events

	% of HHs	All events	DA events	TM events
Price*1[Unaware]	8%	0.052	-0.006	0.119
		(0.097)	(0.141)	(0.108)
Price*1[Aware]	64%	-0.069	-0.098*	-0.026
		(0.042)	(0.051)	(0.051)
Price*1[Missing]	28%	0.020	-0.016	0.065
		(0.052)	(0.055)	(0.073)
Price+IHD*1[Unaware]	3%	-0.071	-0.224***	0.129
		(0.059)	(0.037)	(0.105)
Price+IHD*1[Aware]	60%	-0.192***	-0.221***	-0.143**
		(0.054)	(0.060)	(0.064)
Price+IHD*1[Missing]	37%	-0.019	-0.063	0.045
		(0.078)	(0.082)	(0.099)
P-Value (PIHD*A = P*A)		0.0386**	0.0721*	0.0956*
HH FEs		Yes	Yes	Yes
Hour-by-day FEs		Yes	Yes	Yes
Number of hhs		230	230	230
R-Square		0.571	0.571	0.571

*Heterogeneous treatment effects by survey response. P-Value reports probability of equal treatment effects across groups, conditional on awareness of event. Standard errors clustered by household in parantheses. * denotes significant at the 0.10 level, ** significant at the 0.05 level, *** significant at the 0.01 level.*

Table 8: Heterogeneous Treatments: Frequency of IHD Interaction

	% of HHs	All events	DA events	TM events
Price+IHD*1[0/None]	4%	-0.366***	-0.585***	-0.086
		(0.111)	(0.154)	(0.260)
Price+IHD*1[1-2 times]	10%	-0.008	0.003	-0.022
		(0.082)	(0.084)	(0.100)
Price+IHD*1[3-5 times]	8%	0.065	0.042	0.089
		(0.070)	(0.066)	(0.079)
Price+IHD*1[More than 5 times]	40%	-0.241***	-0.256***	-0.208***
		(0.065)	(0.072)	(0.076)
Price+IHD*1[Missing]	38%	-0.007	-0.031	0.027
		(0.074)	(0.077)	(0.095)
P-Value (PIHD*>5 = PIHD*1-2)		0.022**	0.015**	0.128
P-Value (PIHD*>5 = PIHD*3-5)		0.001***	0.001***	0.005***
HH FEs		Yes	Yes	Yes
Hour-by-day FEs		Yes	Yes	Yes
Number of obs		100	100	100
R-Square		0.571	0.571	0.571

*Heterogeneous treatment effects by survey response. P-Value reports probability of equal treatment effects across frequency of experience with IHDs. Standard errors clustered by household in parantheses. * denotes significant at the 0.10 level, ** significant at the 0.05 level, *** significant at the 0.01 level.*

Table 9: Heterogeneous Treatments: Home During Daytime

			Intention to Treat			Treatment on the Treated		
	% of HHs: Initially Assigned	% of HHs: Compliers Only	All events	DA events	TM events	All events	DA events	TM events
Price*1[Home during day]	51%	52%	-0.036	-0.056	-0.009	-0.032	-0.052	-0.005
			(0.043)	(0.052)	(0.051)	(0.044)	(0.053)	(0.052)
Price*1[Home after school]	22%	20%	-0.042	-0.095	0.028	-0.041	-0.1	0.035
			(0.056)	(0.060)	(0.083)	(0.063)	(0.067)	(0.090)
Price*1[Home after work]	27%	27%	-0.031	-0.068	0.017	-0.027	-0.065	0.022
			(0.067)	(0.080)	(0.080)	(0.069)	(0.082)	(0.082)
Price*1[Missing]	1%	0%	-0.426***	-0.957***	0.636***			
			(0.027)	(0.034)	(0.038)			
IHD+price*1[Home during day]	41%	44%	-0.142**	-0.179**	-0.086	-0.173**	-0.221**	-0.103
			(0.072)	(0.079)	(0.083)	(0.088)	(0.097)	(0.103)
IHD+price*1[Home after school]	21%	22%	-0.124*	-0.13	-0.109	-0.127*	-0.14	-0.105
			(0.071)	(0.091)	(0.082)	(0.076)	(0.102)	(0.082)
IHD+price*1[Home after work]	35%	29%	-0.123*	-0.174**	-0.045	-0.164*	-0.244**	-0.054
			(0.067)	(0.072)	(0.086)	(0.094)	(0.104)	(0.113)
IHD+price*1[Missing]	3%	4%	-0.275	-0.28	-0.25	-0.420***	-0.950***	0.640***
			(0.300)	(0.235)	(0.365)	(0.027)	(0.033)	(0.037)
P-Value (PIHD*HDD = P*HDD)			0.174	0.154	0.386	0.120	0.092	0.350
HH FEs			Yes	Yes	Yes	Yes	Yes	Yes
Hour-by-day FEs			Yes	Yes	Yes	Yes	Yes	Yes
Number of obs			230	230	230	230	230	230
R-Square			0.571	0.571	0.571	0.571	0.571	0.571

*Heterogeneous treatment effects by survey response. P-Value reports probability of equal treatment effects across groups, conditional on being "generally home during the day". Standard errors clustered at the household in parantheses. * denotes significant at the 0.10 level, ** significant at the 0.05 level, *** significant at the 0.01 level.*

Table 10: Heterogeneous Treatments: Central Air Conditioning

	% of HHs	All events	DA events	TM events
Price*1[No Central AC]	55%	-0.090**	-0.085*	-0.089*
		(0.042)	(0.048)	(0.053)
Price*1[Yes Central AC]	42%	0.035	-0.051	0.144**
		(0.049)	(0.064)	(0.057)
Price*1[Missing]	4%	-0.078	-0.071	-0.081
		(0.143)	(0.132)	(0.190)
IHD+P*1[No Central AC]	61%	-0.184***	-0.226***	-0.117*
		(0.054)	(0.059)	(0.071)
IHD+P*1[Yes Central AC]	35%	-0.071	-0.083	-0.052
		(0.073)	(0.082)	(0.079)
IHD+P*1[Missing]	4%	-0.048	-0.17	0.123
		(0.445)	(0.420)	(0.297)
P-Value (P*Yes CAC = PIHD*Yes CAC)		0.191	0.744	0.0259**
P-Value (P*No CAC = PIHD*No CAC)		0.124	0.032**	0.722
P-Value (PIHD*Yes CAC = PIHD*No CAC)		0.180	0.123	0.504
P-Value (P*Yes CAC = P*No CAC)		0.027**	0.630	0.000***
HH FEs		Yes	Yes	Yes
Hour-by-day FEs		Yes	Yes	Yes
Number of hhs		230	230	230
R-Square		0.571	0.571	0.571

*Heterogeneous treatment effects by survey response. P-Value reports probability of equal treatment effects across groups, conditional on the indicated observable. Standard errors clustered by household in parantheses. * denotes significant at the 0.10 level, ** significant at the 0.05 level, *** significant at the 0.01 level.*

Table 11: Load Shifting: Anticipation and Spillovers

	DA events	TM events
Price-Only: 2hrs Pre-Event	0.003	0.053
	(0.039)	(0.038)
Price-Only: 2hrs Post-Event	-0.043	-0.051
	(0.047)	(0.046)
Price+IHD: 2hrs Pre-Event	-0.095**	-0.024
	(0.042)	(0.045)
Price+IHD: 2hrs Post-Event	-0.103*	-0.027
	(0.055)	(0.057)
HH FEs	Yes	Yes
Hour-by-day FEs	Yes	Yes
Number of hhs	437	437
R-Square	0.568	0.568

*Specification is the baseline ITT with additional regressor indicator variables for 2-hrs pre- and post-treatment event. Specification includes treatment indicators, coefficients for which are not reported. Standard errors clustered by household in parantheses. * denotes significant at the 0.10 level, ** significant at the 0.05 level, *** significant at the 0.01 level.*

Table 12: Habit Formation

	Price	Price + IHD
12-1pm Calendar Day Trend	-0.002	-0.0025*
	(0.001)	(0.002)
1-2pm Calendar Day Trend	-0.0019	-0.0022
	(0.001)	(0.001)
2-3pm Calendar Day Trend	-0.0021	-0.0031**
	(0.001)	(0.001)
3-4pm Calendar Day Trend	-0.0022*	-0.0031**
	(0.001)	(0.001)
4-5pm Calendar Day Trend	-0.0026**	-0.0032**
	(0.001)	(0.001)
5-6pm Calendar Day Trend	-0.0027**	-0.0036***
	(0.001)	(0.001)
6-7pm Calendar Day Trend	-0.0032**	-0.0034**
	(0.001)	(0.001)
7-8pm Calendar Day Trend	-0.0030**	-0.0030*
	(0.002)	(0.002)
HH FEs	Yes	
Hour-by-day FEs	Yes	
Number of hhs	437	
R-Square	0.579	

*Results from a single regression specification which interacts a calendar day time trend for each peak hour with initial treatment assignment. The sample is restricted to all non-pricing event weekdays in July and August. Standard errors clustered by household in parantheses. * denotes significant at the 0.10 level, ** significant at the 0.05 level, *** significant at the 0.01 level.*

Table 13: Annual Value of Conservation During Price Events

	DA, $0.50/kWh			TM, $1.25/kWh		
	Price	Price+IHD	Difference	Price	Price+IHD	Difference
20hrs	$1.17	$3.24	$2.07	$0.08	$4.01	$3.93
40hrs	$2.33	$6.48	$4.15	$0.16	$8.02	$7.87
60hrs	$3.50	$9.72	$6.22	$0.24	$12.04	$11.80

Estimated from average peak kWh usage and estimated treatment effects by event type.

Table 14: ITT and ToT Specifications Weighted to Reflect US Income Distribution

Estimator:	ITT			ToT		
Event Type:	All Events	Day Ahead (DA)	30min (TM)	All Events	Day Ahead (DA)	30min (TM)
Price Only	-0.013	-0.031	0.012	-0.013	-0.032	0.012
	(0.048)	(0.058)	(0.061)	(0.050)	(0.061)	(0.062)
Price + IHD	-0.133**	-0.189**	-0.047	-0.178**	-0.267**	-0.059
	(0.061)	(0.083)	(0.109)	(0.086)	(0.116)	(0.137)
Prob(P = P+I)	0.049**	0.070*	0.588	0.042**	0.041**	0.597
Hour-by-day FEs	Y	Y	Y	Y	Y	Y
HH FEs	Y	Y	Y	Y	Y	Y
Number of Events	6	3	3	6	3	3
Number of HHs	429	429	429	429	429	429
R-Square	0.57	0.57	0.57	0.57	0.57	0.57

*Dependent variable: log(kWh). * Significant at the 0.10 level, ** Significant at the 0.05 level, *** Significant at the 0.01 level. Standard errors clustered at the household level.*

Figure 1: In-Home Display (1)

Figure 2: In-Home Display (2)

Figure 3: July 21, 2011: 4hr \$0.50 increase, day-ahead notice

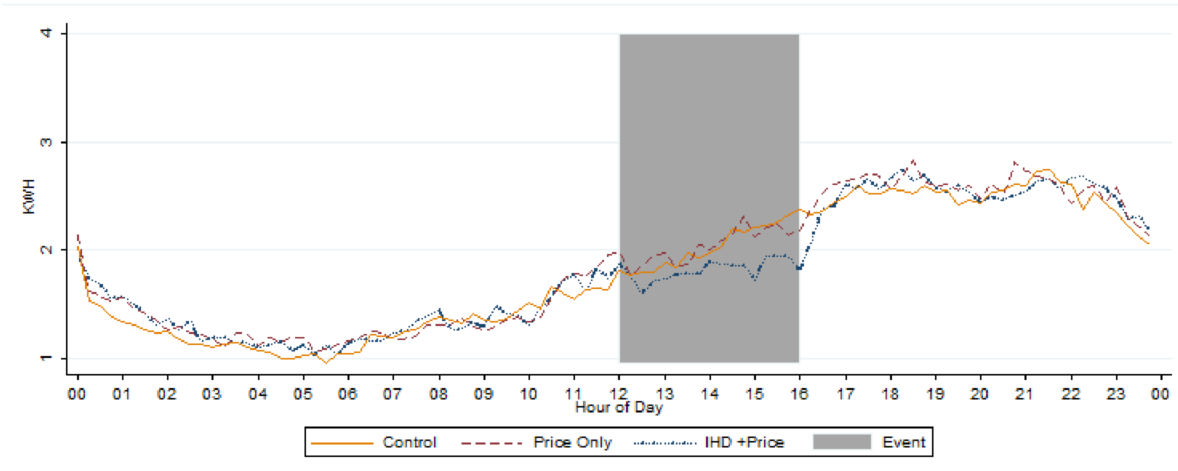

Figure 4: July 22, 2011: 4hr \$1.25 increase, 30-min notice

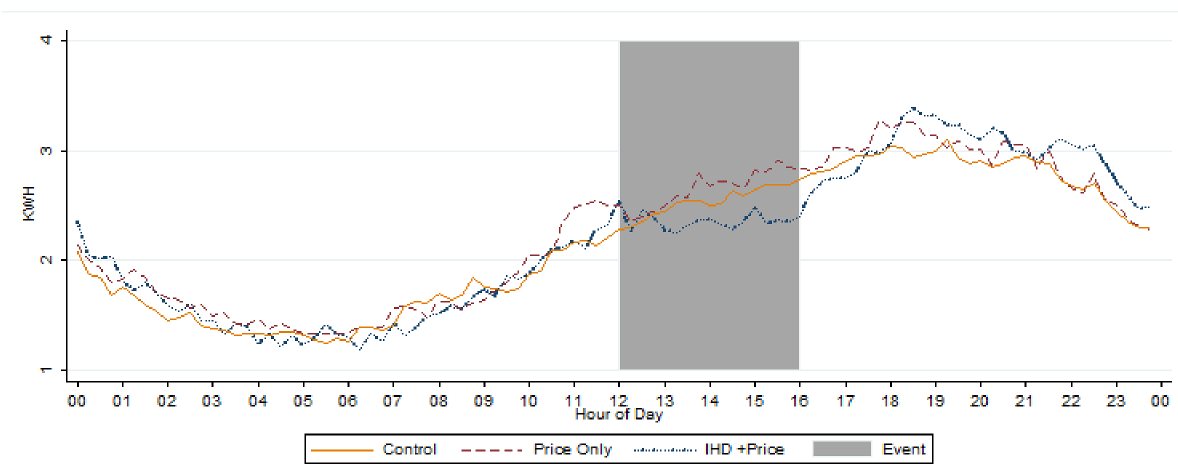

Figure 5: August 4, 2011: 2hr $0.50 increase, day-ahead notice

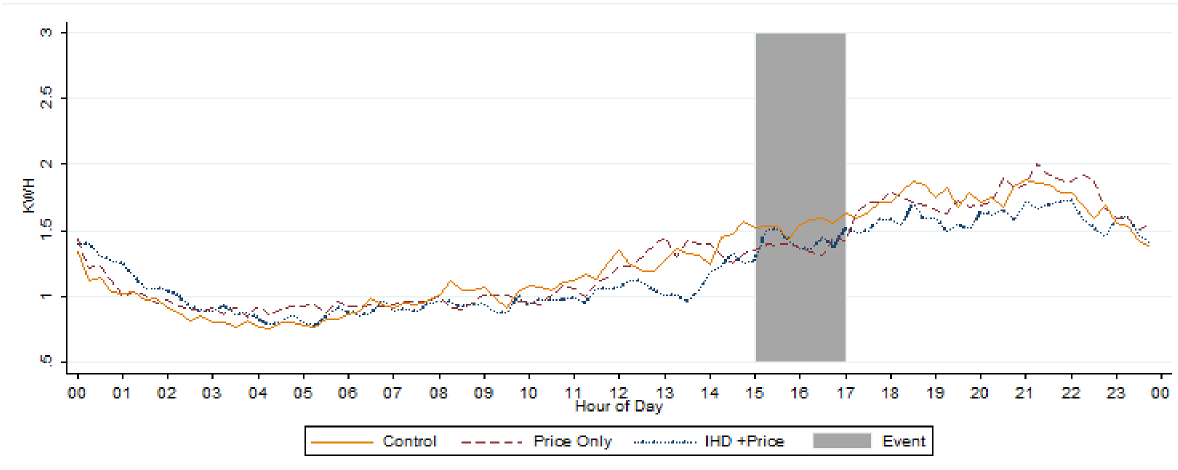

Figure 6: August 10, 2011: 2hr $1.25 increase, 30-min notice

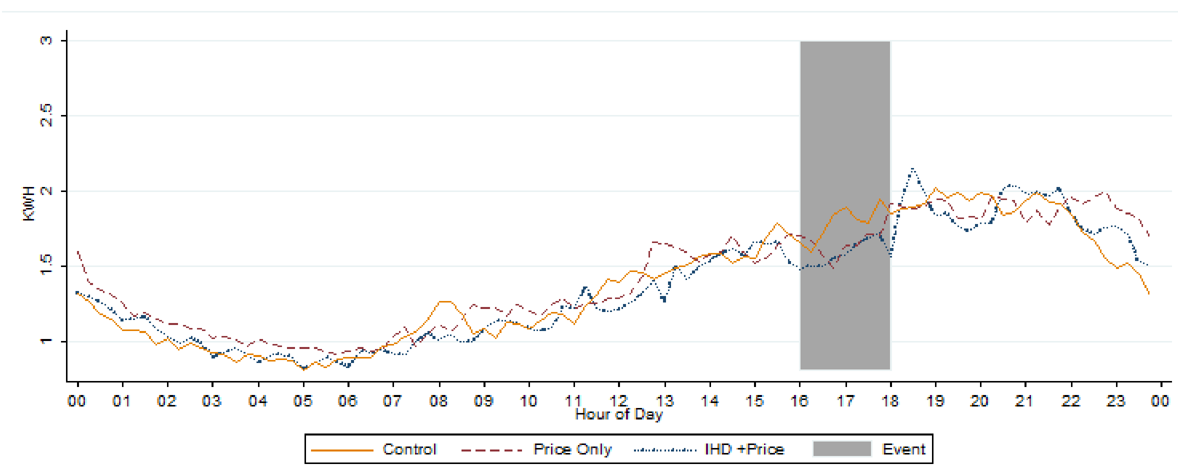

Figure 7: August 17, 2011: 2hr $1.25 increase, 30-min notice

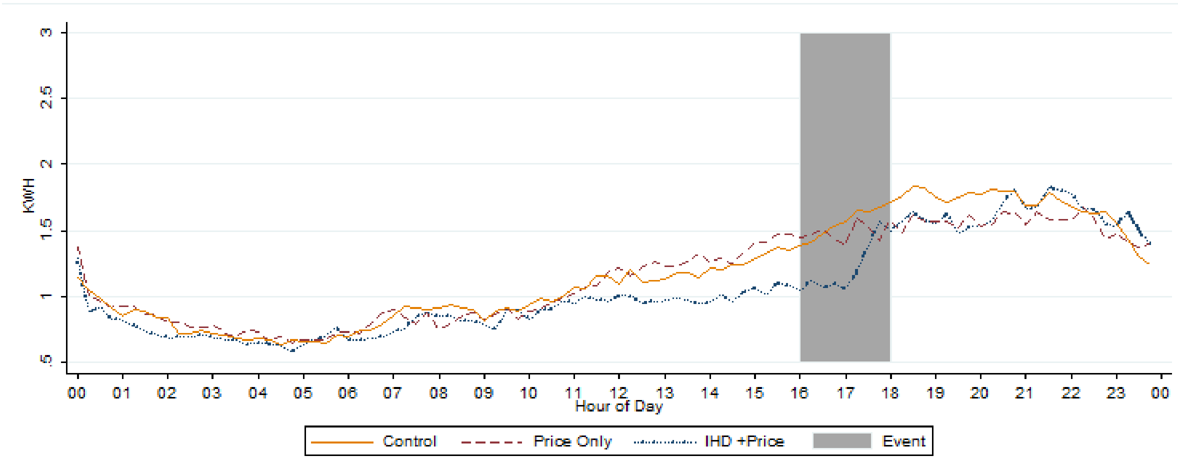

Figure 8: August 26, 2011: 4hr $0.50 increase, day-ahead notice