Knowledge is (Less) Power: Experimental Evidence from Residential Energy Use[1]

Katrina Jessoe[2] and David Rapson[3]

December 18, 2013

Abstract

Imperfect information about product attributes inhibits efficiency in many choice settings, but can be overcome by providing simple, low-cost information. We use a randomized control trial to test the effect of high-frequency information about residential electricity usage on the price elasticity of demand. Informed households are three standard deviations more responsive to temporary price increases, an effect that is not attributable to price salience. Conservation extends beyond pricing events in the short- and medium-run, providing evidence of habit formation and implying that the intervention leads to greenhouse gas abatement. Survey evidence suggests that information facilitates learning.

JEL: C93, D83, L94, Q41

Keywords: Information; Price Elasticity; Randomized Controlled Trial; Energy Demand

---

[1]We would like to thank two anonymous referees, Hunt Allcott, Jim Bushnell, Colin Cameron, Scott Carrell, Colin Carter, Michael Carter, Meredith Fowlie, Koichiro Ito, Stephen Holland, Hilary Hoynes, Michael Price, Nancy Rose, Burkhard Schipper and Aaron Smith. This paper benefited from helpful comments by seminar participants at the CU Boulder, NBER, POWER, 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. Rapson thanks the Energy Institute at Haas for support under a research contract from the California Energy Commission. All errors are our own.

[2]Department of Agricultural Economics, University of California, Davis, One Shields Ave, Davis, CA 95616; Phone: (530) 752-6977; Email: `kkjessoe@ucdavis.edu`

[3]Department of Economics, University of California, Davis, One Shields Ave, Davis, CA 95616; Phone: (530) 752-5368; Email: `dsrapson@ucdavis.edu`

"*One should hardly have to tell academicians that information is a valuable resource: knowledge* is *power. And yet it occupies a slum dwelling in the town of economics. Mostly it is ignored...*"
- George Stigler (1961)

A key assumption underpinning central theorems in economics is that agents are fully informed. Yet information is rarely free to decision-makers. Information costs may take many forms - time, cognitive effort, monetary expense, technological hurdles - and are pervasive in choice settings. A growing body of literature explores the importance of some of these information costs, demonstrating that decisions are dramatically altered when information is conveyed in a simple way. Retirement plan decisions change when employees are provided benefits information (Duflo and Saez 2003); the take-up rates of government transfer programs increase among eligibles when information is provided in a simple format (Bhargava and Manoli 2013); consumers restrict their use of cell phone minutes when informed of approaching a higher price tier (Grubb and Osborne 2012); and fishing markets operate more efficiently with the adoption of mobile phones, a new information technology (Jensen 2007). Even President Obama has acknowledged that decisions can be improved through the dissemination of relevant information in some important choice settings, as witnessed by his "College Scorecard" which provides students with comprehensive access to information to facilitate college choice. This paper explores how an economic treatment can interact with improvements in information to enhance welfare.

The full information assumption implies that utility-maximizing consumers can identify the tradeoff between product attributes and price (e.g. Rosen 1974). However, basic information may often be unclear. The recent literature has focused on whether agents perfectly know and comprehend the *price* of a good (i.e. price "salience"). In many settings they do not. Consumers treat taxes featured and not featured in the posted price differently, becoming more price elastic when taxes are salient (Chetty, Looney, and Kroft 2009); they are inattentive to opaque "add-on" costs such as shipping and handling fees (Hossain and Morgan 2006); and drivers become less price elastic when road tolls become less salient due to automated payment technology (Finkelstein 2009). Another setting has garnered less attention: that in which there is uncertainty about *non-price* attributes (exceptions include Jin and Leslie 2003, Gabaix and Laibson 2006). A common form of this uncertainty arises in household choice settings, where we consume *services*, not inputs directly. Durable goods that require water or energy inputs fall into this category, and markets have a mixed record of providing information about the relationship between input and output quantities. Advancements in car dashboard displays have increased drivers' knowledge of the gasoline required to travel a mile (Stillwater and Kurani 2012). But information about the household production function is lacking in other markets, leaving individuals uncertain as to how common actions like watering the lawn or cooling a house by one degree translate into water and electricity usage. This is also observed in caloric intake from food (Bollinger, Leslie, and Sorensen 2011, Wisdom, Downs, and Loewenstein 2010) and carbon emissions from driving. Generally, this variety of poor information will lead to inefficiency since households must know input requirements (not just price of inputs) to equate marginal benefit with marginal cost.

This paper assesses the importance of providing information about non-price attributes, specifically input quantities, in a familiar setting: residential electricity demand. Electricity customers traditionally exhibit low price elasticity (Reiss and White 2005, Allcott 2011, Ito 2011), but this may be due to features of the setting that inhibit full information. Quantity consumed tends to be shrouded because of coarse and infrequent billing making it difficult to know both electricity usage at any moment in time and the input requirements of each appliance. Further, since electricity comprises only a modest share of household budgets, it may be rational for households not to invest the time and effort to resolve this uncertainty. Therefore, it is important to distinguish whether residential electricity consumers are actually price inelastic or just appear to be inelastic because they lack complete information.

To address this question, we conduct a randomized controlled trial (RCT) (a "framed field experiment" in the terminology of Harrison and List 2004) which exposes all treatment households to exogenous price changes. A random subset of these households is also exposed to real-time feedback on quantity of electricity consumed via an in-home display (IHD). The price treatments expose customers to two- or four-hour long pricing events during which the price of electricity increased by 200 to 600 percent, thereby allowing us to isolate the effect of price changes alone (for households without IHDs) on usage.[1] The IHD enables customers to inform themselves about electricity usage and price at almost zero marginal cost, allowing them to learn about the energy requirements associated with various forms of household production. For example, households can view their electricity load both before and after turning off a light or running the air conditioner. Since all treatment households were informed of upcoming price changes via email, voicemail or text message, the main incremental information provided by the IHD (relative to the price-only treatment) was the real-time quantity of electricity being used.

Our central result shows that providing residential electricity customers with real-time information about energy usage increases their price elasticity of demand. Conditional on changes in price, households exposed to information feedback consistently exhibit price elasticities that are roughly three standard deviations larger than those without feedback. Households in the price-only group reduce their usage by between 0 and 7 percent on average during pricing events (depending on the amount of advance notification they received), relative to control. In contrast, those exposed to the same price changes but who also have IHDs, exhibit much larger usage reductions of 8 to 22 percent. The treatment gradient attributable to information is not due to price salience, as confirmed by an analysis of event notification receipts (the emails/texts/phone calls sent to households in advance of events). Instead, our empirical evidence suggests that experience with IHDs facilitates consumer learning, improving households decision making when confronted with high prices.

The treatment effects also spill over into non-event hours, both within an event day and on non-event days. In the long run, an evaluation of trends in usage over the days of the summer reveals that households in both the price and price-plus-information groups are forming conservation habits even when events are not occurring. The combined effects imply a social benefit in the form of potential long-run capital efficiencies, as well as greenhouse gas abatement on the order of 1-2 percent of emissions from the residential electricity sector.

This paper contributes to an ongoing literature that explores the effectiveness of dynamic pricing

[1] These so-called "critical peak prices" are a variant of dynamic pricing that is intended to transmit wholesale market price incentives to the retail sector. By design, the magnitude of our price treatments is consistent with peak fluctuations in the wholesale electricity market.

programs in the electricity sector, an industry with over $370 billion per year in retail sales. One of the industry's most difficult challenges relates to the mismatch between wholesale and retail prices within and across days, despite recent advances in metering technologies (Borenstein 2002, Borenstein 2005, Borenstein and Holland 2005, Joskow and Wolfram 2012). Increases in the retail price elasticity would lead to short-run and long-run efficiency gains from market power mitigation and improvements in capital efficiencies. Our results suggest that when coupled with real-time feedback, price acts as an effective lever to achieve these gains.

Our main result is timely due to a recent tendency for electricity and water regulators to eschew market-based approaches in favor of non-market instruments. Electric utilities have made social pressure a cornerstone of recent energy conservation efforts (e.g. Opower) and federal authorities rely heavily on efficiency standards. To respond to droughts, water authorities tend to use non-pecuniary approaches (Olmstead and Stavins 2009). And while these programs have been shown to achieve moderate conservation (Allcott 2011b, Ferraro and Price 2011), policymakers should view price as an effective tool to achieve their objective if consumers are informed.

The paper proceeds as follows: section I explains the experimental setting and research design, and section II describes the data; empirical methods and results are presented and discussed in section III; section IV describes conservation implications; and section V concludes.

I Research Design

In partnership with a regulated electric utility, The United Illuminating Company (UI), we designed and implemented a framed field experiment that introduced short-term price increases and real-time information to a sample of residential electricity customers in the Bridgeport and New Haven areas of Connecticut. The treatment events occurred in July and August of 2011, the months during which peak electricity demand strains the capacity of the grid.

To be eligible for participation in the pilot a customer needed to reside in a townhouse or single family home, have a broadband internet connection, and sign and return an end-user agreement indemnifying UI against litigation risk.[2] As an additional participation incentive, we 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. To recruit households into the pilot, UI emailed 60,000 customers that had enrolled in paperless billing, indicating the likely presence of Internet in their home. We estimate that approximately 7,000 households opened the emails.[3] Of the 7,000 or so households that received the invitation to participate, we recruited 1,152 households (approximately 1 in 6) to participate in the project.

A subset of these households, 437, form the sample for this study. We randomly assigned households to one of three groups: control, price ("price-only"), and price-plus-information ("price+IHD").

Control Group: A total of 207 households were assigned to the control group. These households (and all others in the pilot) received a mailing that notified them they were in the pilot, informed them of their group assignment, and contained an energy conservation pamphlet documenting "101 Ways to Conserve Electricity".

Price-Only Treatment: The 130 households in this group experienced pricing events that varied in the magnitude of the price increase and the timing of event notification. There were two event types. The first, "DA", provided *day-ahead* notification that the per-kWh price of electricity would be increased by $0.50 (or a roughly 250 percent increase over the standard rate). These events mimic a pricing policy that a utility might use to transmit electricity conservation incentives when high temperatures are expected in the following afternoon. The second type of event, "TM", sent notification *thirty minutes* before a $1.25 increase in the per-kWh price of electricity. A utility may 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).[4]

[2]Some aspects of device-to-utility communication were configured to occur wirelessly via the Internet. It was thus important for experimental validity 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.

[3]It is well known that only a small fraction of households receiving these sorts of marketing emails actually open them, let alone read them. While we are unable to report the exact open rate (the fraction of delivered email messages that are actually opened), 12 percent is a rule-of-thumb open rate used by industry experts. For example, see http://www.mailermailer.com/resources/metrics/2012/open-rates.rwp.

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

Ex-ante we cannot predict to 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. All events occurred during peak hours, but there was variation in the length and exact timing of events. Table 1 lists each event, including the start time, event duration, and measured temperature. Households received notification of pricing events by email, phone call and/or text message, depending on their stated preference.

Since billed electricity rates are determined through a periodic, external rate-case process, we transmitted the experimental price incentives via 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 period, and that amount was subtracted from the household's off-bill account balance. At the conclusion of the experiment, any balance remaining in the account was the customer's to keep. This arrangement achieved the intended marginal incentives while also shielding participants from bill increases attributable to pricing events.

The implementation of price changes may raise concerns about construct validity. However, the central result of the paper relates to the differential price response between households with and without information feedback. As such, any concerns about construct validity would apply equally to these two groups, so the treatment differential is unaffected by the implementation of the price changes.

Price+IHD Treatment: The 100 households randomly assigned to this treatment group experienced the pricing program described above, and also received real-time information about their electricity use. 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 (see Online Appendix figures A.1 and A.2). Households received the IHDs and professional installation free of charge. The latter ensured that displays were set up and activated, and maximized the likelihood that participants understood

on the 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.

how to use them. Two separate vendors supplied the IHDs; both provided the same information in a similar format: real-time usage, electricity price, estimated monthly usage and bill-to-date. Allowing all subsequent results to differ by IHD vendor produces identical results.

The main difference between this treatment and the price-only treatment is that customers are able to view in real time the quantity of power being consumed, the price of electricity, and their estimated monthly bill-to-date. The device removes the information acquisition costs involved in informing oneself about how electricity-consuming actions translate into electricity usage and expenditure. In particular, 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 potentially enabling them to more fully optimize in response to price changes. Price+IHD households were also informed of price changes via phone, email and text, so while the IHD displayed price in real-time, the additional information conveyed by this feature was likely minimal. We test the hypothesis that IHDs made price changes more salient in section 4.3.1.

II Data

We use high-frequency meter data on household electricity usage as the primary data source. Advanced meters were installed in all participating households, enabling electricity usage of all participants (including those that dropped out of the study) to be collected at 15-minute intervals. The utility also collected data which confirmed the receipt of event notifications. Lastly, we rely on data collected from two household 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 demographic and housing unit characteristics, appliance ownership, conservation-related actions, tendency to be home during the day, and the frequency with which households checked their IHDs.

Technical issues associated with back-end system software precluded UI from retrieving meter data for some households until after the experiment had begun. Roughly one-quarter of participants are absent from the billing data for at least one event. We are not concerned about the effect of these omissions due to both the exogenous nature of the software issue and the fact that results are quantitatively identical irrespective of the inclusion or exclusion of these households. To make maximum use of the variation in our data, we use the unbalanced panel (as defined by presence in the data for *at least one* treatment event) in what follows. However, in the Online Appendix we provide results using the balanced sample as well.

A Randomization and Compliance

Before exploring the impact of information on the price elasticity of demand, we provide evidence to support the integrity of the randomization. Table 2 presents descriptive statistics of observable characteristics 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"). A comparison across control and treatment groups indicates statistical balance in observables. Usage and rate class, characteristics provided in the UI billing data and thus observed for all households, are similar across groups.[5] The availability of demographic data is subject to survey compliance, so the observed statistical difference in the size of the home across groups may be due to either sampling variation or survey non-compliance. Still, other household characteristics do not differ across groups, contributing to the case for balance. The table also highlights that households in our sample are wealthier (and likely different across variables correlated with wealth) than both the national average and the population served by the utility.

To further test the integrity of the randomization we estimate a linear probability model regressing each treatment indicator on mean off-peak electricity usage and rate class (flat rate versus time-of-use rate).[6] 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

[5] Where data were available, we restricted comparison of usage data to the days preceding the first event (i.e. July 1-July 21). We found no systematic differences in mean peak and off-peak usage across treatment and control groups. Technical difficulties prevented UI from accessing the meter data from some households until after the first pricing event. As such this restricted sample is comprised of 99, 81 and 184 households in the price only, price+IHD, and control groups respectively.

[6] 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.

price+IHD group in column 2. Neither variable is statistically significant in explaining assignment to the price and price+IHD treatments. These results provide further evidence that households were randomly assigned to treatments.

Some attrition occurred in each group. In total, thirty-eight households (approximately 9 percent) did not complete the pilot. These households either moved, requested to be removed from the study or failed to arrange an installation appointment for the IHD. Of the 100 households assigned to the price+IHD treatment, 28 did not complete the study, which we attribute to time and scheduling costs of IHD installation. In contrast, compliance was high in the non-technology groups. Only 4 of the 207 control households and 6 of the 130 price-only households were non-compliers.

A concern raised by this asymmetry is that success or failure to schedule an installation appointment is systematically related to the desire or ability to respond to treatment. 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. In this scenario, our estimated treatment effect may partly reflect that households best able to respond to price are more likely to be in the price+IHD group.

We test for asymmetric attrition by estimating a linear probability model in which we regress an indicator set equal to 1 for compliers and 0 otherwise on off-peak electricity usage and rate class. Results are presented in columns 3 (price-only) and 4 (price+IHD) of Table 3. The significant coefficient on rate class for price+IHD households suggests that some selective attrition may be occurring. Household fixed effects will strip out any time-invariant unobservables (such as rate class), and in our analysis we use intent-to-treat (ITT) and treatment-on-the-treated (ToT) estimators to account for asymmetries in non-compliance.

B Raw Data and Mean Differences

Figures 1-6 plot raw mean hourly usage by treatment group on each of the six event days. In these figures, mean 15-minute interval consumption is averaged across all households in each group. The shaded area marks the period during which a pricing event occurred. In Figure 1, for example, a

DA pricing event was held between noon and 4pm. The commonalities between these events are evident: households exposed to information feedback exhibit visibly lower usage during treatment events than price-only or control households. In the hours preceding an event, hourly electricity usage in the three groups is approximately the same, though this appears to change over time.[7] 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 in each hour of the pricing event. Interestingly, we do not observe visible differences between price-only and control households, despite large price increases in the former. The figures also reveal that the treatment effect differs from event to event, and in some cases is less visually discernible. For example, it is difficult to see differences between group outcomes in Figure 3. This may be partly due to a smaller treatment effect on that day, or that there appear to be similar treatment effects in the Price-only and Price+IHD groups. Clear temporal patterns also emerge from the plots, which are predictable features of electricity demand given weather and lifestyle patterns.

A simple comparison of means in the raw data (Table 4) reflects what is visible in most of the usage plots. For each event type, the control group exhibits higher mean usage than either treatment group. The magnitude of the raw treatment effects varies by treatment and event type. Price+IHD households exhibit between a 12-18 percent usage reduction (to all events), as compared to a 0-7 percent reduction by price-only households. Response to DA events is higher than to TM events, providing evidence that advance notice facilitates a stronger response, even under a weaker price incentive.

III Results: Information and Price Elasticity

We begin the regression analysis by estimating a simple difference-in-differences model,

$$q_{it} = \sum_{g \in \{P, P+I\}} \beta_g D_{it}^g + \gamma_g + \delta_e + \mu_{it} \tag{1}$$

[7]There is some visual evidence of potential habit formation and load-shifting affecting non-event-hour usage. In section 5, we test for the presence of these behaviors.

in which the dependent variable, q_{it}, is the natural log of energy usage by household i in 15 minute interval t. 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 household i in interval t. This specification includes a pricing event indicator, δ_e, set equal to 1 during pricing events, and separate treatment group dummies, γ_g, for the price only and price+IHD groups. We also report results from specifications that use additional controls, including hour-by-calendar-date dummies, household fixed effects and a combination of the two.

A Intent-to-Treat

Panel A of Table 5 reports results from the ITT estimator in which D_{it}^{P} and D_{it}^{P+I} in Equation 1 denote initial assignment to the price and price+IHD treatments, irrespective of compliance status. The coefficients, $\hat{\beta}_P$ and $\hat{\beta}_{PI}$, in this specification are consistent estimates of the average percentage change in electricity usage from assignment to treatment during pricing events.

Column 1 displays results from the simple difference-in-differences model; column 2 includes hour-by-calendar-date fixed effects and column 3 includes household fixed effects. In column 4, our preferred specification, we include both household fixed effects and hour-by-calendar-date dummies, allowing us to account for both time-invariant differences across households and aggregate hourly patterns in weather and lifestyle. Here, the coefficients of interest are identified from variation within households over time, controlling for aggregate hourly shocks separately for each calendar day. Finally, in columns 5 and 6 we present the effects of DA and TM events separately, using the preferred controls.

This table makes clear that households with real time information feedback are significantly more responsive to price changes than those without. When DA and TM events are pooled together (columns 1-4), households with an IHD reduce usage by 11 to 14 percent and are over three standard deviations more responsive to pricing events as those without IHDs (column 4). A comparison of the treatment effect across groups highlights that the differential is robust to the controls included, both in terms of magnitude and statistical significance. That the inclusion of household and time controls does not meaningfully alter the magnitude of treatment effects provides further evidence for the integrity of the randomization. As shown in Table A.1 of the Appendix, the treatment differential is similar when we use the "balanced" sample, suggesting that the estimated information gradient is not sensitive to our choice of sample. These results provide strong evidence that the cumulative effect of real-time information feedback in this setting is to increase the price elasticity of demand. In the absence of information, the price effects are weak. Price-only households on average decrease usage by 2 to 6 percent during events, an effect which is not statistically significant. These results confirm the pattern illustrated in the daily plots which suggests little, if any, response to events by uninformed households.

We break out the events separately by type (DA or TM) in columns 5 and 6, and continue to find that households with IHDs are significantly more price elastic. The percentage point treatment differential between households in the two treatment groups is robust to event type and significant with 85 percent confidence. However, individuals are more responsive, both economically and statistically, to pricing events that occur with more 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). Households assigned to the price group reduce usage by 7 percent during DA events where this response is weakly significant. In contrast, they are largely unresponsive to TM events and if anything increase usage during these events. We find similar patterns across event type for households assigned to the price+IHD treatment. With advanced notification, these households reduce usage by 17 percent as compared to an 8 percent (but not statistically significant) response with thirty minute notification. Even with strong financial incentives, with only 30 minutes of warning individuals may not have the ability to respond to a price change.

B Treatment on Treated

The treatment effect on treated households (ToT) is the causal effect of the price and price+IHD treatments on compliers. The ToT specification uses initial treatment assignment as an instrument for receipt of treatment, and is estimated using two-stage least squares. Randomization and the strong rate of compliance ensure strength and validity of the instruments: compliance was 98 percent, 95 percent, and 72 percent in control, price-only, and price+IHD, respectively.

Estimates of ToT specifications are presented in Panel B of Table 5, where the specification estimated in each column corresponds to those described for Panel A. Again, we find that information feedback meaningfully impacts the usage response to pricing events. Compared to the ITT estimates, the magnitude of the treatment effect and treatment differential are slightly larger. In our preferred specification (column 4), households with IHDs are 13 percentage points (3-4 standard deviations) more responsive to pricing events than those without, where this differential is present with 95 percent confidence. Households in both treatment groups are more responsive to advance notification events, and we continue to observe a treatment differential with over 85 percent confidence when estimating event types separately.

Many readers may wish to see these results presented as elasticities, which is of course mathematically possible. The implied arc elasticities are low at less than -0.12. However, it is unclear as to what information is actually conveyed by these statistics due to the magnitude of the price increases: 200 percent during CPP events and 600 percent during DR events. 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 price changes may induce the same absolute behavioral response.

C Secondary Results and Robustness

Potential explanations for the strong information effect are many. The hypothesis that we believe is most consistent with observed behavior is that experience with the display facilitates learning about energy choices, and in particular the mapping of these choices to expenditure. This mechanism is consistent with a framework in which baseline information about quantity is imperfect, and whereby IHDs serve to inform consumers and (in doing so) influence their subsequent choices. In this section, we first present some evidence that attempts to rule out the leading alternative hypothesis – that IHDs generate more awareness of price, and consumers who possess them become more price responsive for this reason. We then provide some evidence consistent with the hypothesis that consumers learn through experience with the IHDs, and that this plays an important role in the treatment differential. In the Appendix, we also explore the practical mechanisms by which households respond to price changes and find that, while different household characteristics are correlated with customer response, they are not driving the information gradient.

In some of the analyses that follow, we use data collected from the pre- and post-surveys, thereby limiting the sample to households that completed them. As such the estimates may be vulnerable to differences in survey non-compliance across groups that are also correlated with unobservables; we discuss this in the Appendix. A second caveat in the specifications is that we interact treatment indicators with survey responses that are themselves outcome variables. The results that follow should be viewed as cross-tabulations of the treatment effect with various survey responses.

C.1 Notification and Awareness of Price Events

A plausible alternative to the "learning" hypothesis is that IHDs facilitate a heightened awareness of price events. Our experimental design sought to control for this by making electricity price changes salient. This was achieved by having the utility send all customers in the price and price+IHD groups notification in the form of a combination of a text message, email and/or phone call in advance of each event. The messages alerted households of the event and informed them what the price would be. Still, we would like to test that there is no additional awareness effect from the IHDs that is driving the main treatment differential. One indicator that customers were aware of pricing events would be if they actually received notification of them. If for some reason, notifications were not received, it becomes more plausible that the IHD may provide incremental awareness of pricing events (making it more difficult to rule out the price salience mechanism).

For each pricing event, the utility collected data confirming the receipt of event notification and documenting the extent to which notification information was conveyed if notification was not confirmed. To acknowledge notification of an event, a household needed to a push a button on the telephone at the end of of an automated phone call (informing customers of an event) or click a button confirming receipt of a notification email. These events, as well as households who had notification texted to them, were defined as confirmed. The remaining events were classified as not confirmed or intermediate, where the latter category includes instances when someone answered the phone but hung up before manually confirming receipt or when an automated message was left on a voicemail. We classify this intermediate group as confirmed. As shown in Table A.2 in the Appendix, awareness of the events was high, suggesting that the treatment events were front-of-mind for a majority of households.

Were the differential response to exist because IHDs increase awareness of price events, then we would expect to see that, conditional on confirmation, households in the two groups would respond equivalently. To test this hypothesis, we estimate

$$q_{it} = \sum_{g \in \{P, P+I\}} \sum_{A \in \{0,1\}} \beta_g D_{it}^g * \mathbb{1}\{A_{it} = A\} + \gamma_i + \sigma_h + \mu_{it} \tag{2}$$

in which we interact the treatment indicator D_{it}^g with A_{it}, a variable set equal to 1 if receipt of notification of the event occurring in interval t was confirmed by household i, and control for household and hour-by-calendar-day fixed effects (γ_i and σ_h, respectively). Results are presented in Table 6.[8] It is clear that event notifications are important, as households confirming receipt of them exhibit a larger response overall and for the DA events. We also see that confirmation of event notification alone cannot explain the treatment differential between those with and without IHDs. Conditional on confirmation of event notification, we reject the null that the coefficient estimates are equal with 95 percent confidence overall, and with 90 and 85 percent confidence when estimating DA and TM events separately. At the same time, conditional on no confirmation we find a statistically indistinguishable response across the price and price+IHD groups. This suggests that IHDs do not appear to be either informing households of events or enabling unaware households to respond.

[8]In an alternative specification, "intermediate" confirmation and full confirmation indicators were interacted separately with the treatment dummy. We found that the treatment effect for price+IHD households was similar between the intermediate and confirmed designations. This motivated the specifications in which the intermediate confirmations are defined as fully confirmed.

C.2 Experience

We have shown that awareness does not account for the information gradient in price response, and in the Appendix discuss alternate hypotheses which also fail to explain the gradient. We now hypothesize that IHDs facilitate learning about the electricity usage associated with the portfolio of household production alternatives. To the extent that customers observe and interact with the IHD, they are gaining information about features of the mapping between electricity-consuming actions and the electricity inputs that they require. For example, if a consumer views the display before and after turning on her air conditioner (or light), she will notice that it consumes a large (or small) amount of electricity. This knowledge, if accumulated for multiple appliances, better equips households to respond to price changes.

To test if frequent experience with the IHD increases price responsiveness, we use responses from the post-survey that asked, "How many times per week did you look at the IHD in the first month that it was activated?". 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. Table 7 shows that most households frequently experimented with their IHD in the first month that they received it. Conditional on answering this question, only 6 percent of respondents did not look at it and approximately 65 percent of respondents looked at it more than 5 times.

In columns 2-4, we present results from the estimation of Equation 2, where the sample is restricted to control households and households with an IHD, and 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.[9] This evidence suggests that more frequent experience with the IHDs facilitates learning about the quantity of electricity consumed by energy consuming durables. Still, it is possible that the households inherently more responsive to price also happen to look at the IHDs more frequently. If this is case then unobservables may bias the estimates.

[9]Households reporting never to look at the displays have large measured responses as well. However, these are few in number and there are potentially reasonable explanations (e.g. principal-agent problems, or conservation responses that do not require knowledge of tradeoffs on the intensive margin, like leaving the home).

IV Conservation Implications

Despite confining the pricing events to small periods in time, the treatments produce behavioral responses which extend beyond events. These occur in both the short-run (hours adjacent to the price event) and the medium-run (on non-event days over the course of the summer). In the short run, activities undertaken in response to high prices spill over into the hours preceding and following the events, resulting in lower usage outside of the treatment window. This result is highlighted in Table 8, which reports estimates from the preferred ITT specification but with the addition of indicator variables for the two hours preceding and following price events for each treatment group. In the hours preceding and following an event, price households exposed to DA notification exhibit no change in usage, while those with feedback significantly reduce usage by 10 percent. In contrast, we find no significant evidence of spillovers (or load shifting) in response to TM events, likely because households have limited time to prepare for them. Our results provide no evidence of load shifting, suggesting that the conservation experienced during events is not offset by increases in usage in adjacent non-event hours.

In the medium-term (weeks to months) as households are exposed to multiple events their usage decreases in meaningful ways on *non-event* days. This effect translates into measurable conservation on non-event days over the months of the study. To explore these effects, we restrict the sample to include only non-event days for the balanced panel of households and estimate the following specification:

$$q_{it} = \sum_{g} \beta_g D_{it}^g + \sum_{g} \sum_{hod} \lambda_{g,hod} * D_i^g * d + \gamma_i + \sigma_h + \mu_{it} \tag{3}$$

The new term is $\sum_{g} \sum_{hod} \lambda_{g,hod} * D_i^g * d$, where d is a running variable counting days from July 1 to August 31 2011 (ie. 1 to 62), D_i^g is a treatment group indicator, and *hod* is a binary variable indicating each hour of the day between noon and 8pm, allowing us to estimate separate trends, $\lambda_{g,hod}$, for each hour of the peak period by treatment group g.

Table 9 reports the $\lambda_{g,hod}$ coefficients from Equation 3, with each row reflecting the estimated trend associated with each hour of day. For example, row 1 displays the trend in kWh for the period noon to 1pm, implying an average daily decrease (gradient) in usage of 0.23 percent for price-only households during this noontime hour. This corresponds to a 14 percent decrease in usage on August 31 relative to July 1. At the upper end of the coefficient range, conservation effects are equivalent to 24 and 21 percent during 7-8pm for the price-only and price+IHD groups, respectively. We find that the conservation trend is steeper for price+IHD households in early hours, but late in the peak period, price-only households exhibit a larger change in usage.

Taken together, three observations emerge from these non-event period results. First, our main estimates of the treatment affect are attenuated. Given that identification comes from within household variation, reductions in usage outside the event window will lower the baseline against which the event-period effects are compared.[10] Further, since spillovers are stronger for price+IHD households, the treatment *differential* will also be attenuated, implying an even larger gradient of information feedback on price elasticity. Second, given the magnitude of the main treatment differential during the price events, it is noteworthy that habit formation across the two groups over non-event days is quite similar. This contrast in results may suggest that recurring price events induce a cumulative response in the medium-run irrespective of the presence of information feedback, but that the information enables a more sophisticated and precise response to short-run incentives.

The third observation is that the incremental contribution of the price treatments extends beyond load reductions during critical periods, conferring environmental benefits in the form of greenhouse gas abatement. Both in the short-run and medium-run we find that the cumulative effect of treatment results in electricity conservation. The magnitude of conservation associated with behavior during the actual pricing events is likely to be negligible, given their short duration and infrequency. However, habits that form in the medium-run may lead to measurable changes, amounting

[10]To test if the primary treatment effects are attenuated, we restrict the sample to include only (i) event days and (ii) non-event days that precede the first pricing event, and estimate our preferred specification. Results are shown in Table A.3. Compared to the results shown in Table 5, the treatement effect increases (in absolute terms) for both price only and price+IHD households. This suggests that conservation on non-event days attenuates the primary treatment effects reported earlier.

to approximately a 1-2 percent reduction in residential electricity sector carbon dioxide emissions.[11]

V Conclusion

This paper reinforces the importance of information when non-price attributes of the choice setting are imperfectly known to consumers. Using residential electricity markets as the empirical setting, our experimental results show that information feedback about electricity usage increases the price elasticity of demand. Evidence suggests that feedback helps consumers to learn about the input requirements of their energy-consuming home durables. The incremental contribution of this information is both economically and statistically meaningful: a three standard deviation increase in the behavioral response.

Our results contribute to ongoing energy policy discussions about the ability of price to influence consumer demand. The ever-decreasing cost of information technology indicates that a major obstacle to price-based solutions – poorly informed consumers – is becoming surmountable. This implies that when combined with feedback, price may be an efficient tool to address the challenges that characterize the electricity industry. For example, dynamic retail pricing can be used to reduce long-term capital investments, align marginal costs and marginal benefits in the short-run and mitigate market power (when consumers are price elastic). Further, evidence that price events cause overall conservation will lead to greenhouse gas abatement equivalent to 1-2 percent of emissions from the residential electricity sector.

More broadly, these results confirm the practical importance of one of economics' most ubiquitous assumptions – that decision-makers have perfect information. Other important choice settings share features that make these findings broadly relevant. Providing consumers with information about the quantity of frequently-consumed goods such as water, calories or greenhouse gas emissions may

[11] Many assumptions are required to reach this estimate. Mainly, we assume that the conservation attributable to habits at the end of our treatment period (August 31) persists through a full summer season and is generalizable across all regions in the continental US. We apply this electricity reduction during the hours of 4-8pm, and use regional estimates of the hourly marginal load and hourly CO_2 emissions of the marginal generator from FERC Form 714 and the analysis by Graff Zivin, Kotchen, and Mansur (2013), respectively. Details of these calculations are available from the authors upon request.

improve private efficiency. However, the direction of social benefits may be setting-specific. While in our setting information makes consumers more price elastic, in other settings consumers may realize that they consume too little or are too price responsive. If there are externalities in these markets, then this response will make private decisions more efficient but may increase social costs.

References

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

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

[3] Bhargava, Saurabh and Dayanand Manoli (2013). "Why are Benefits Left on the Table? Assessing the Role of Information, Complexity, and Stigma on Take-up with an IRS Field Experiment," Working Paper.

[4] Bollinger, Brian, Phillip Leslie, and Alan Sorensen, (2011). "Calorie Posting in Chain Restaurants," *American Economic Journal: Economic Policy*, 3(1): 91-128.

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

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

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

[8] Chetty, Raj, Adam Looney, and Kory Kroft (2009). "Salience and Taxation: Theory and Evidence," *American Economic Review*, 99(4): 1145-1177.

[9] Duflo, Esther and Emmanuel Saez (2003). "The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment," *The Quarterly Journal of Economics*, 118(3): 815-842.

[10] Ferraro, Paul J. and Michael K. Price, (2013). "Using Non-pecuniary Strategies to Influence Behavior: Evidence from a Large Scale Field Experiment," *The Review of Economics and Statistics*, 95(1): 64-73.

[11] Finkelstein, Amy (2009). "EZ-Tax: Tax Salience and Tax Rates," *The Quarterly Journal of Economics*, 124(3): 969-1010.

[12] Gabaix, Xavier and David Laibson (2006). "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets," *The Quarterly Journal of Economics*, 121 (2): 505-540.

[13] Graff Zivin, Joshua, Matthew J. Kotchen, and ErinT. Mansur (2013). "Spatial and Temporal Heterogeneity of Marginal Emissions: Implications for Electric Cars and Other Electricity-Shifting Policies," NBER Working Paper No. 18462.

[14] Grubb, Michael D. and Matthew Osborne (forthcoming). "Cellular Service Demand: Biased Beliefs, Learning, and Bill Shock," Working Paper.

[15] Harrison, Glenn W. and John A. List (2004). "Field Experiments," *Journal of Economic Literature, 42(4), 100955.*

[16] Hossain, Tanjim and John 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.

[17] Ito, Koichiro (forthcoming). "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing," *American Economic Review.*

[18] Jensen, Robert (2007). "The Digital Provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector," *Quarterly Journal of Economics*, 122(3): 879-924.

[19] Jin, Ginger Zhe and Phillip Leslie (2003). "The Effect of Information on Product Quality: Evidence from Restaurant Hygiene Grade Cards," *The Quarterly Journal of Economics*, 118(2): 409-451.

[20] Joskow, Paul L. and Catherine D. Wolfram (2012) "Dynamic Pricing of Electricity," *American Economic Review Papers and Proceedings*, 102(May): 381385.

[21] Stillwater, Tai and Kenneth S. Kurani (2012) "Goal Setting, Framing, and Anchoring Responses to Ecodriving Feedback," Working Paper.

[22] Olmstead, Sheila M. and Robert N. Stavins (2009). "Comparing Price and Non-price Approaches to Water Conservation," *Water Resources Research*, 45, W04301, doi:10.1029/2008WR007227.

[23] Reiss, Peter C. and Matthew W. White (2005). "Household Electricity Demand Revisited," *The Review of Economic Studies*, 72: 853-883.

[24] Rosen, Sherwin (1974). "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition," *The Journal of Political Economy*, 82(1): 34-55.

[25] Wisdom, Jessica, Julie S. Downs and George Loewenstein, (2010). "Promoting Healthy Choices: Information versus Convenience," *American Economic Journal: Applied Economics*, 3(2): 164-178.

Figures and Tables

Figure 1: July 21, 2011: 4hr $0.50 increase, day-ahead notice

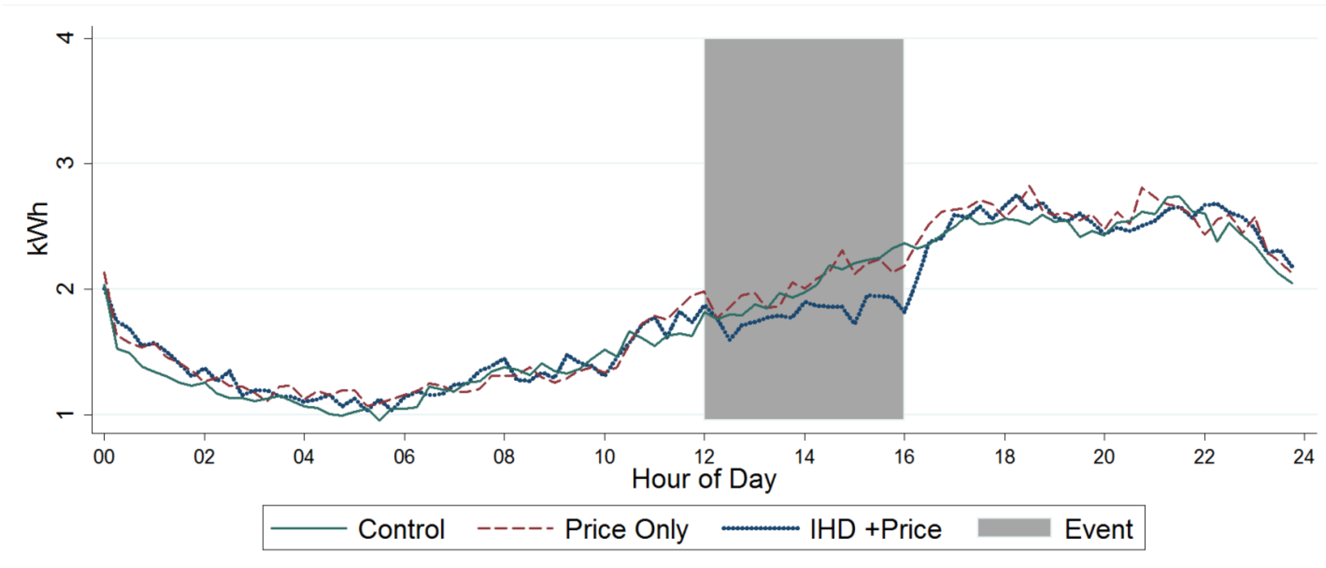

Figure 2: July 22, 2011: 4hr $1.25 increase, 30-min notice

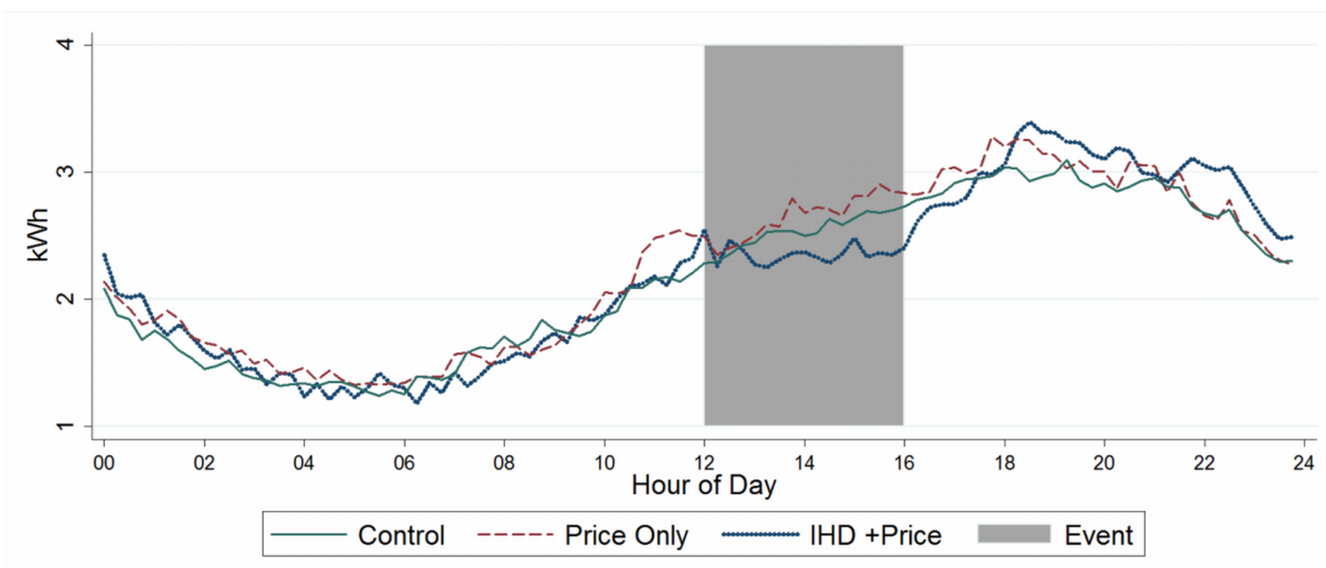

Figure 3: August 4, 2011: 2hr $0.50 increase, day-ahead notice

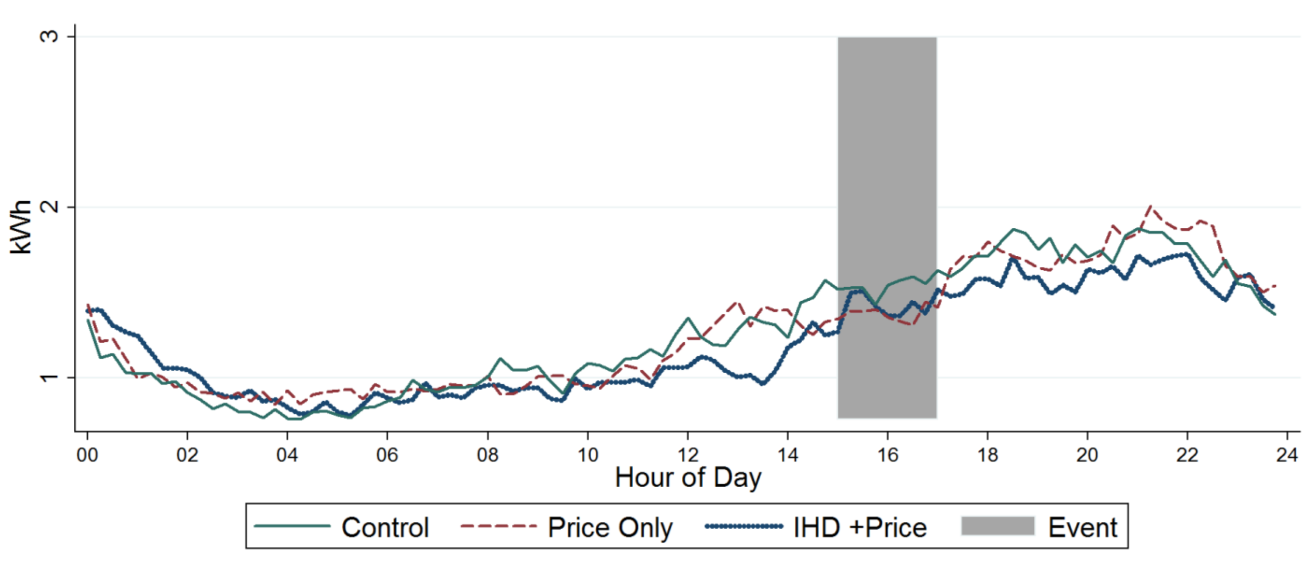

Figure 4: August 10, 2011: 2hr $1.25 increase, 30-min notice

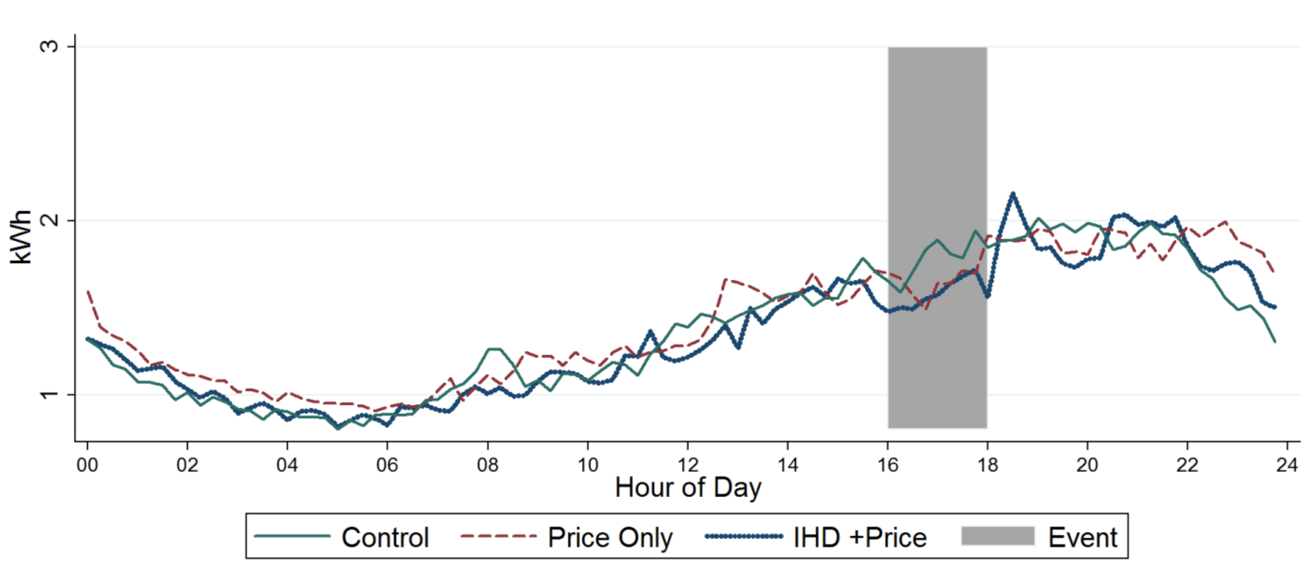

Figure 5: August 17, 2011: 2hr $1.25 increase, 30-min notice

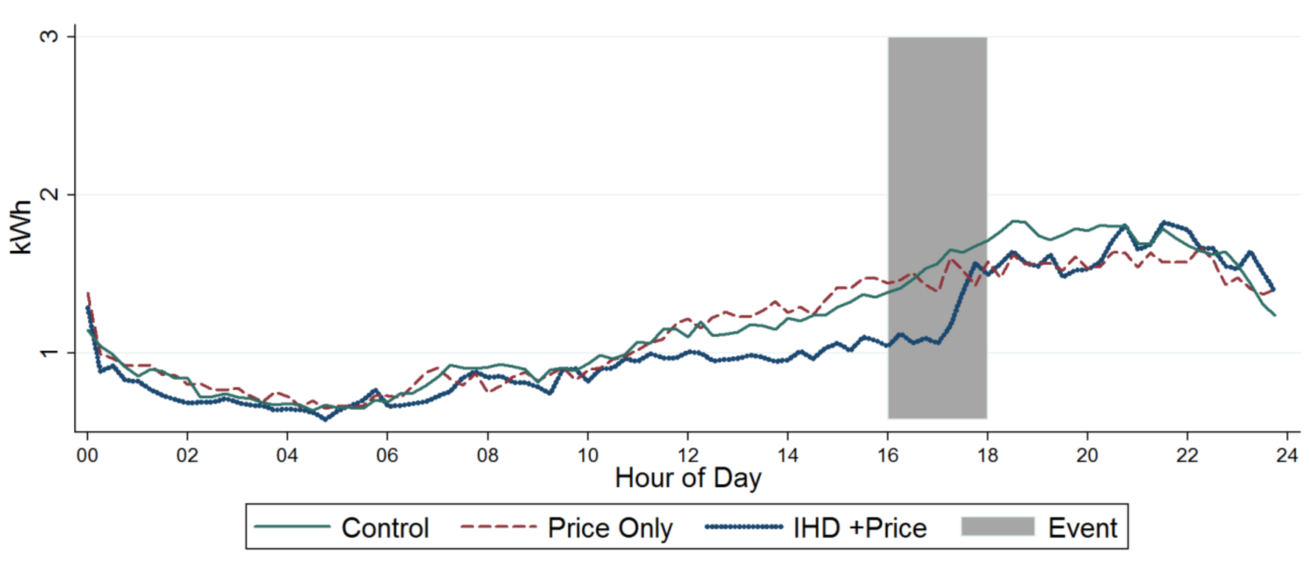

Figure 6: August 26, 2011: 4hr $0.50 increase, day-ahead notice

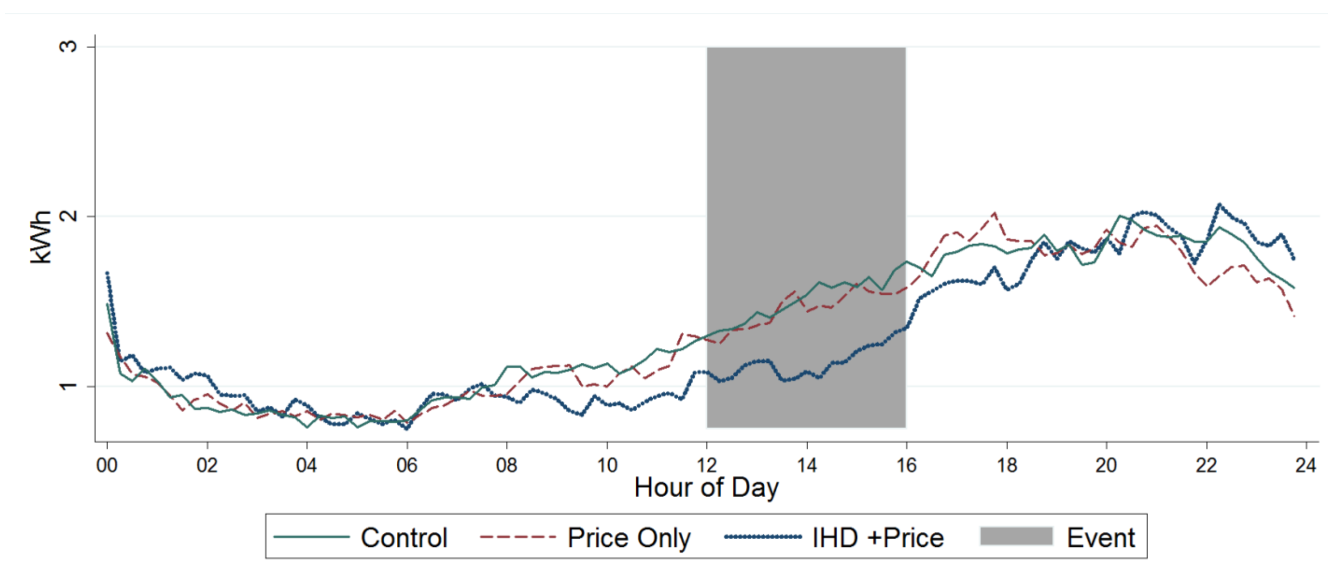

Table 1: Treatment Events

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)	1.230	207	1.282	130	0.052	1.225	100	-0.005
	(0.738)		(0.739)		(0.629)	(0.658)		(-0.058)
Peak usage (kWh/h)	1.519	207	1.533	130	0.014	1.413	100	-0.106
	(1.197)		(1.036)		(0.109)	(0.984)		(-0.772)
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)	1.232	203	1.297	124	0.065	1.229	72	-0.003
	(0.74)		(0.73)		(0.77)	(0.63)		(-0.033)
Peak usage (kWh/h)	1.529	203	1.556	124	0.026	1.468	72	-0.061
	(1.20)		(1.04)		(0.20)	(0.99)		(-0.389)
TOU Rate (1=yes)	0.182	203	0.202	124	0.019	0.181	72	-0.002
	(0.39)		(0.40)		(0.43)	(0.39)		(-0.032)
Home ownership (1=yes)	0.774	199	0.821	123	0.047	0.855	69	0.081
	(0.42)		(0.39)		(1.01)	(0.36)		(1.439)
Annual income ($1000)	72.00	199	75.00	123	0.003	76.00	69	0.004
	(29.00)		(28.00)		(0.93)	(28.00)		(1.006)
Home size (1000 square feet)	1.541	185	1.908	114	0.367 **	1.611	66	0.069
	(1.10)		(1.84)		(2.16)	(1.16)		(0.433)
Age of home (years)	52.221	154	56.574	94	4.354	53.375	56	1.154
	(30.43)		(31.02)		(1.09)	(28.59)		(0.247)

*Notes: Means are reported by treatment group, with standard deviations in parentheses below. "Difference" displays the difference in means between each treatment group and control, with t-statistics reported in parentheses below. *,**,*** denote significant at the 0.10, 0.05, and 0.01 level.*

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

	Initial Group		Compliers	
	Price	Price + IHD	Price	Price + IHD
Mean Off Peak kWh	0.021	-0.019	0.030	0.060
	(0.040)	(0.040)	(0.029)	(0.071)
TOU Rate (1=yes)	0.010	0.088	-0.018	-0.263**
	(0.074)	(0.071)	(0.053)	(0.109)
F-Statistic	0.206	0.775	0.579	2.915
P-Value	0.814	0.462	0.562	0.059
N	337	307	130	100

*Notes: Results denoted "Initial Group" are from a linear probability model regressing observables on the treatment group indicator. Results denoted "Compliers" are 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 parentheses. *,**,*** denote significant at the 0.10, 0.05, and 0.01 level.*

Table 4: Mean kWh Differences (wrt Control) by Treatment Group

		Mean kWh During Events			Difference in Mean kWh wrt Control	
Event Type	Variable	Control	Price	Price + IHD	Price	Price + IHD
Sample: Unbalanced Panel						
DA	Mean	1.65	1.59	1.35	-0.06	-0.30 *
	Std Dev	(1.51)	(1.25)	(1.22)		
	Obs	207	130	100		
TM	Mean	2.07	1.99	1.79	-0.07	-0.28
	Std Dev	(1.77)	(1.54)	(1.42)		
	Obs	186	128	87		
Sample: Balanced Panel						
DA	Mean	1.79	1.67	1.54	-0.13	-0.25
	Std Dev	(1.56)	(1.13)	(1.24)		
	Obs	172	90	77		
TM	Mean	2.17	2.17	1.92	0.00	-0.25
	Std Dev	(1.79)	(1.39)	(1.44)		
	Obs	172	90	77		

*Notes: This table presents raw household-level means of hourly kWh during each 15-minute interval, and their difference (for treatment groups with respect to control) during price event periods of each type (DA and TM). The "Unbalanced Panel" and "Balanced Panel" samples are comprised of households assigned to a treatment for which we observe usage data for "at least one pricing event" and "all pricing events", respectively. Standard deviations in parentheses. *,**,*** denote significant at the 0.10, 0.05, and 0.01 level.*

Table 5: Treatment Effects (Unbalanced Panel)

Event Type:	All	All	All	All	Day Ahead (DA)	30min (TM)
Column:	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: ITT Unbalanced Panel						
Price Only	-0.031	-0.054	-0.027	-0.038	-0.071*	0.006
	(0.036)	(0.036)	(0.036)	(0.036)	(0.042)	(0.044)
Price + IHD	-0.116**	-0.137***	-0.123***	-0.137***	-0.171***	-0.084
	(0.048)	(0.048)	(0.047)	(0.046)	(0.051)	(0.057)
Prob(P = P+I)	0.096*	0.098*	0.051*	0.044**	0.066*	0.130
R-Square	0.001	0.054	0.536	0.583	0.583	0.583
Panel B: ToT Unbalanced Panel						
Price Only	-0.032	-0.056	-0.028	-0.040	-0.074*	0.007
	(0.037)	(0.037)	(0.037)	(0.037)	(0.044)	(0.046)
Price + IHD	-0.143**	-0.170***	-0.153***	-0.170***	-0.217***	-0.100
	(0.058)	(0.058)	(0.057)	(0.057)	(0.064)	(0.067)
Prob(P = P+I)	0.061*	0.052*	0.030**	0.023**	0.025**	0.115
R-Square	0.001	0.054	0.536	0.583	0.583	0.583
HH FEs	N	N	Y	Y	Y	Y
Hour-by-day FEs	N	Y	N	Y	Y	Y
Number of Events	6	6	6	6	3	3
Number of HHs	437	437	437	437	437	401

*Notes: The dependent variable is ln(kwh) in 15-minute intervals. ITT results are reported from an OLS regression on usage on intial assignment to treatment. ToT results are reported from a 2SLS regression where initial treatment assignment is used as an instrument for receipt of treatment. The sample is comprised of households for which we observe usage data for AT LEAST ONE pricing event (i.e. the unbalanced panel). All specifications include a treatment group indicator and an event window indicator (except where subsumed by time or household fixed effects). In columns 1-4 the treatment window indicator is set equal to 1 if any event (DA or TM) is occurring. Column 2 includes hour-by-day fixed effects; column 3 includes household fixed effects and column 4 includes both. Columns 5 and 6 present results separately from DA and TM events, respectively. Standard errors in parentheses are clustered at the household level. *,**,*** indicates significance at 0.10, 0.05, and 0.01.*

Table 6: Notification Confirmation and Treatment Effects

Event Type:	All events	DA events	TM events
Column:	(1)	(2)	(3)
Price*1[Not Confirmed]	-0.007	-0.043	0.038
	(0.048)	(0.066)	(0.057)
Price+IHD*1[Not Confirmed]	-0.050	-0.104	0.037
	(0.080)	(0.087)	(0.110)
Price*1[Confirmed]	-0.049	-0.080*	-0.005
	(0.040)	(0.046)	(0.051)
Price+IHD*1[Confirmed]	-0.162***	-0.192***	-0.113*
	(0.052)	(0.057)	(0.062)
P-Value (PIHD*NC = P*NC)	0.628	0.557	0.991
P-Value (PIHD*C = P*C)	0.047**	0.073*	0.120
HH FEs	Yes	Yes	Yes
Hour-by-day FEs	Yes	Yes	Yes
Number of hhs	437	437	401
R-Square	0.583	0.583	0.583

*Notes: Results are reported from an OLS regression where the dependent variable is ln(kwh) and the treatment indicator is interacted wtih notification confirmation category. The sample is comprised of households assigned to a treatment for which we observe usage data for AT LEAST ONE pricing event. P-Value reports probability of equal effects across groups. Standard errors in parentheses are clustered by household. *,**,*** indicates significance at 0.10, 0.05, and 0.01.*

Table 7: Frequency of IHD Interaction

	% of HHs	All events	DA events	TM events
Price+IHD*1[0/None]	4%	-0.453**	-0.690***	-0.161
		(0.196)	(0.181)	(0.338)
Price+IHD*1[1-2 times]	10%	-0.013	-0.028	0.007
		(0.139)	(0.137)	(0.160)
Price+IHD*1[3-5 times]	8%	0.02	-0.02	0.06
		(0.083)	(0.083)	(0.091)
Price+IHD*1[More than 5 times]	40%	-0.248***	-0.279***	-0.204**
		(0.077)	(0.085)	(0.086)
Price+IHD*1[Missing]	38%	-0.023	-0.065	0.037
		(0.096)	(0.095)	(0.119)
P-Value (PIHD*>5 = PIHD*1-2)		0.123	0.102	0.225
P-Value (PIHD*>5 = PIHD*3-5)		0.011**	0.017**	0.020**
HH FEs		Yes	Yes	Yes
Hour-by-day FEs		Yes	Yes	Yes
Number of hhs		307	307	273
R-Square		0.526	0.526	0.526

*Notes: Treatment effect and survey-reported initial weekly frequency of IHD interaction. P-Value reports probability of equal treatment effects across frequency of experience with IHDs. Standard errors clustered by household in parentheses. *,**,*** indicates significance at 0.10, 0.05, and 0.01.*

Table 8: Load Shifting: Anticipation and Spillovers

	DA events	TM events
Price-Only: 2hrs Pre-Event	0.002	0.053
	(0.038)	(0.038)
Price-Only: 2hrs Post-Event	-0.043	-0.051
	(0.046)	(0.045)
Price+IHD: 2hrs Pre-Event	-0.097**	-0.024
	(0.042)	(0.045)
Price+IHD: 2hrs Post-Event	-0.103*	-0.027
	(0.054)	(0.056)
HH FEs	Yes	Yes
Hour-by-day FEs	Yes	Yes
Number of hhs	437	401
R-Square	0.579	0.579

*Notes: The specification is the baseline ITT with additional regressor indicator variables for 2-hrs pre- and post-treatment event. The sample includes households that were present for at least one treatment event (what we are calling the unbalanced panel). The specification includes treatment indicators, for which coefficients are not reported. Standard errors clustered by household in parentheses. *,**,*** indicates significance at 0.10, 0.05, and 0.01.*

Table 9: Habit Formation

	Price	Price + IHD
12-1pm Calendar Day Trend	-0.0023	-0.0030**
	(0.0016)	(0.0015)
1-2pm Calendar Day Trend	-0.0024	-0.0027*
	(0.0015)	(0.0014)
2-3pm Calendar Day Trend	-0.0025*	-0.0032**
	(0.0014)	(0.0013)
3-4pm Calendar Day Trend	-0.0027*	-0.0031**
	(0.0014)	(0.0013)
4-5pm Calendar Day Trend	-0.0033**	-0.0034***
	(0.0014)	(0.0013)
5-6pm Calendar Day Trend	-0.0032**	-0.0033**
	(0.0014)	(0.0013)
6-7pm Calendar Day Trend	-0.0038**	-0.0032**
	(0.0015)	(0.0014)
7-8pm Calendar Day Trend	-0.0037**	-0.0029**
	(0.0017)	(0.0015)
HH FEs	Yes	
Hour-by-day FEs	Yes	
Number of hhs	339	
R-Square	0.556	

*Notes: 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, and includes only households that were present for all treatment events (what we are calling the balanced panel). Standard errors clustered by household in parantheses. *,**,*** indicates significance at 0.10, 0.05, and 0.01.*