Default Effects, Follow-on Behavior and Welfare in Residential Electricity Pricing Programs

Peter Cappers, Meredith Fowlie, C. Anna Spurlock, Annika Todd, Catherine Wolfram, and Patrick Baylis*

July 19, 2015

Preliminary draft. Please do not cite.

Abstract

We study default effects in a large randomized controlled trial, in which one treatment group was allowed to opt-in to time-based pricing for electricity while another was allowed to opt-out. We provide dramatic evidence of default effects – a significantly higher fraction of households enrolled in the time-based pricing plan in the opt-out enrollment group compared with the opt-in group, even though deviating from the default simply involved making a phone call or clicking through to a website. A distinguishing feature of our empirical setting is that we observe follow-on behavior subsequent to the default manipulation. This, in conjunction with randomization of the default enrollment mechanism, allows us to separately identify the subsequent response of "complacent" households (i.e., those who only enroll in time-based pricing if assigned to the opt-out treatment). We find that the complacent households do reduce energy use during higher priced peak periods, though significantly less on average compared to customers who actively opt in. However, because significantly more customers face time-based pricing in the opt-out group relative to the opt-in group, the pricing incentives produced much larger aggregate demand reductions during peak periods in the opt-out group. Finally, because welfare implications hinge on the mechanisms that give rise to the default effect, we examine the extent to which our results, together with ancillary evidence, lend support to alternative hypotheses including transaction costs, inattention, and explanations that assume preferences are constructed versus revealed.

*We received many helpful comments from seminar participants at Cornell and UC Berkeley. The authors gratefully acknowledge contributions from and discussions with Hunt Allcott, Stefano Dellavigna, Steven George, Jennifer Potter, Lupe Strickland, Michael Sullivan and Nate Toyama. We also thank Severin Borenstein, Lucas Davis, and Michael Greenstone for helping to make this project possible through their initial involvement with the Smart Grid Investment Grant program. This material is based upon work supported by the Department of Energy under Award Number OE0000214.The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

1 Introduction

Economists and psychologists have found that, when confronted by a choice with a default option, people are often predisposed to accept the default. Existing work has documented this "default effect" for a range of decisions that would seem to merit deliberate choices, including health insurance, retirement plans, charitable giving, and social media privacy settings. This phenomenon is of general interest because it provides businesses and public policymakers with a relatively easy and non-intrusive way to influence choice outcomes.

Although the effect of default options on decision-making has been clearly demonstrated in the literature, the broader economic implications of these default effects have been much harder to discern. One reason is that the economic impacts of a default effect can manifest through direct and indirect channels. To comprehensively assess these impacts, one must consider not only the initial choice that is subject to the default manipulation, but also any "follow-on" behaviors that depend on the initial choice. Welfare analysis is particularly complicated because it hinges on the cognitive mechanisms behind default effects which are not well understood. Moreover, several of the proposed explanations for default effects involve non-standard preferences that render standard welfare analysis inapplicable.

This study analyzes the implications of manipulating the default option in a new choice setting: time-based electricity pricing. In addition to documenting an economically significant default effect, we collect detailed data on subsequent choices and behaviors. In our setting, there are at least two important reasons to be concerned with follow-on behavior. First, this is a context in which influencing an intitial decision (i.e., customer acceptance of a time-varying electricity price) is not an end in itself but a means to an end (i.e., reduced electricity consumption at peak times). The economic importance of the default effect depends critically on whether the desired electricity consumption impacts manifest among households susceptible to default effects. Second, this is a context in which the welfare implications of the default effect are of great interest. The follow-on behavior we observe sheds light on the extent to which the default choice is compatible with the normative preferences of the customers who accept the default option, thus helping us to understand the mechanisms that give rise to the default effect in this setting.

Increasing the level of active participation in residential demand response programs can significantly increase the efficiency of electricity market outcomes. Economists have noted for some time that efficient pricing of electricity should reflect changing electricity market conditions (e.g., Boiteux, 1964a,b).

Electricity demand, marginal system operating costs, and firms' abilities to exercise market power vary significantly and systematically over hours of the day and seasons of the year. Figure 1 demonstrates the extent of this variation for a representative week during our study. The red line depicts hourly electricity demand, which cycles predictably over the course of a day, varying by a factor of 1.5 to almost 3 from the middle of the night to the peak hours in the late afternoon. The blue line depicts hourly wholesale prices, which fall below $60/MWh for most hours, but spike to over $1,000/MWh at critical peak times.

[FIGURE 1 HERE]

Although wholesale electricity prices can vary significantly across hours, at least partially reflecting variations in marginal costs, retail prices do not generally reflect these dynamic market conditions. The vast majority (over 95 percent in 2012) of U.S. residential customers pay time-invariant prices for electricity (FERC, 2014). If customers are not exposed to prices that reflect variable marginal operating costs, economic theory suggests that consumers will under-consume in periods of low marginal costs and over-consume in periods of high marginal costs. This further implies over-investment in capacity to meet excessive peak demand. For example, Borenstein and Holland (2005) simulate that by shifting a fraction of customers to time-based rates, utilities could construct 44 percent fewer peaking plants.

In principle, these inefficiencies can be mitigated - or eliminated - with the introduction of time-varying retail electricity pricing. Residential customers are a crucial component of a time-based pricing strategy since their highest demand (e.g., driven by air conditioning in many parts of the U.S.) drives system peaks. When residential customers have been exposed to time-based prices, existing analyses suggest they are willing and able to respond to them accordingly. EPRI (2012) identified what the authors' deemed to be the best seven U.S. residential pricing studies up to that time, finding that customers could reduce peak demand in response to a CPP rate design by 13-33%, depending on the existence of automated control technology (e.g., programmable communicating thermostat).[1]

The deployment of smart grid technology was dramatically accelerated under the American Recovery and Reinvestment Act of 2009. As of 2014, more than 50 million smart meters had been deployed to over 40 percent of US households (IEI, 2014). In principle, this technology investment could facilitate widespread

[1]These estimates imply an elasticity of substitution in the range of 0.07 - 0.24 and an own-price elasticity in the range of -0.3 - -0.07. Note that the experimental nature of our study allows us to assess many dimensions of customers' responses to time-based pricing, including spillovers within and across days. Some previous evaluations of time-based pricing have relied on within-customers comparisons, which assume there are no spillovers of this sort.

adoption of time varying pricing programs. However, participation in these programs has historically been very low (Joskow and Wolfram, 2012). Proactive approaches to increasing active participation in these programs will be required to fully leverage demand response potential.

This paper explores an innovative approach to increasing participation - and demand response - in a retail time-varying electricity pricing program. The analysis is based on a field experiment run by the Sacramento Municipal Utility District (SMUD) in 2011-2012. In one set of treatment groups, customers were invited to opt-in to a new time-based pricing structure. In another set of interventions, customers were informed that they would be defaulted onto the new pricing program unless they opted out. We show that making time-based pricing the default choice can significantly increase participation – over 90 percent of the customers stayed with time-based pricing when defaulted onto it, while only about 20 percent actively opted in.

As noted above the economic importance of this default effect depends critically on whether the households susceptible to default effects actively reduced their peak consumption in response to the time varying electricity prices. Our experimental design allows us to characterize the electricity demand response among the "complacent" households who defaulted into the program but likely would not have selected it if asked to opt in. We show that these customers, 75 percent of the sample, do reduce consumption when peak-period prices increase, though by about half as much as customers who actively opted in.

In addition to estimating the direct and indirect effects of the default manipulation, we are also interested in understanding what drives the default effect because different explanations can have very different welfare, and hence public policy, implications. On possible interpretation of our findings is that perfectly rational consumers face very high transaction costs from switching away from the default. Under this scenario, consumers can suffer welfare losses when they are defaulted into a time-based pricing program. Under an alternative scenario, switching inattentive or uninformed customers onto the new rate allows them to experience time varying prices and remain in the program upon learning they prefer it. In this case, defaulting customers into the new rate is welfare improving. We explore ancillary data to provide insight on the relevance of alternative explanations for the default effect. While not dispositive, the bulk of the evidence points to customers with limited attention (either through rational inattention or lack of awareness), and non-standard choice heuristics.

The paper proceeds as follows. Section 2 situates our paper relative to the existing work on the default effect. Section 3 describes the experiment. Section 4 describes the data and our empirical approach.

Section 5 presents our main results on the default effect and follow-on behavior. We are able to isolate the follow-on behavior of those who actively opted in (referred to here as “always takers”), from those who only ended up on the new pricing structure because of the default (referred to here as “complacents”). In Section 6, we consider similarities and systematic differences between complacent households and other households, both in terms of pre-determined characteristics and subsequent follow-on behavior. This provides suggestive evidence on both the mechanisms behind the default effect and the welfare implications of the default choice architecture in this setting. Section 7 concludes.

2 Default Effects, Choice Modification, and Follow-on Behavior

A rich literature documents and explores various aspects of default effects in a range of settings, including 401(k) participation (Madrian and Shea, 2001; Choi et al., 2002, 2004), organ donation (Johnson and Goldstein, 2003; Abadie and Gay, 2006), car insurance (Johnson et al., 1993), car purchase options (Park et al., 2000), and email marketing Johnson et al. (2002). This literature offers a range of possible explanations for default effects. In instances where the choice is relatively simple and not particularly important, default effects may stem merely from rational inattention (Bellman et al., 2001; Sims, 2005). When confronting a decision that is more complicated or stressful, choosing not to choose (and thus accepting the default) can allow the decision-maker to avoid incurring the costs of gathering information or evaluating difficult tradeoffs (Kressel and Chapman, 2007; Pichert and Katsikopoulos, 2007). If the consumer has limited personal experience with the choice context, the default option can be appealing, particularly if it is perceived to be the prescribed or recommended option (Beshears et al., 2009).

The manipulation of default options to influence behavior is attractive insofar as it guides behavior without constraining individual choice. However, the welfare implications of switching the default choice are not a priori positive. Welfare impacts depend critically on whether the default choice is well-suited to those who are susceptible to default effects. In some cases, the welfare impact or value generated by a default effect is completely determined by the initial choice. For example, the social impact of a default effect on the choice to become an organ donor is largely - if not entirely- determined by the increase in organ donor consent rates. In other contexts, subsequent ‘follow-on’ behavior plays a critical role in determining the outcomes that matter.

We distinguish between two types of follow-on behavior, both of which can play a role in determining welfare impacts. First, individuals may choose to subsequently modify the choice that is subject to the default effect. For example, a consumer who accepts a particular 401(k) plan as a default option might subsequently adjust the parameters of this choice by changing the savings rate, changing the asset allocation, or dropping off the plan altogether. Second, there may be important choices or actions that are contingent on - but distinct from- the initial choice. Returning to the health insurance example, the relevant follow-on behavior could include subsequent choices about whether or not to go to the doctor, how frequently, what procedures to have, etc. Another example in the area of privacy settings is when a consumer agrees to a default privacy policy for a service like Facebook, and the relevant follow-on behaviors are the choices made about whether to post personal photos on Facebook or whether to include particular information about one's background or location.

To date, the literature on default effects has emphasized the initial choice itself with less emphasis on the implications for subsequent choices and behaviors. In particular, we are not aware of studies that consider the contingent behaviors that can be indirectly influenced by default effects.[2] Our study provides an unusual opportunity to analyze not only the direct effect of a default manipulation on an initial choice, but also the ways in which the default effect operates through the initial choice to affect subsequent consumer choices and behaviors. The follow-on behavior we observe among complacent customers clearly shows that the default effect on program participation results in valuable reductions in peak electricity consumption. Observed patterns of behavior also sheds light on the underlying mechanisms that give rise to the default effect in this setting.

3 Experimental Design

3.1 SMUD Rates and Random Assignment

SMUD serves approximately 530,000 residential households. Households were excluded from our experiment if they did not have interval meters to capture hourly electricity usage installed prior to June 2011, they were participating in SMUD's Air Conditioning Load Management program, Summer Solutions study, PV solar programs, budget billing programs, or medical assistance programs, or if they had master-metered accounts. After these exclusions, approximately 174,000 households remained eligible for the experimental

[2]Analyses of 401(k) investment decisions have analyzed the first type of follow-on behavior – modifications to the original choice. For example, Carroll et al. (2009) analyze savings outcomes over time as a function of different default options at the initial plan participation decision.

population.

Households in the experimental population were randomly assigned into ten groups, five of which are the focus of this paper. Households in four of these five groups were encouraged to participate in a new pricing program; the fifth group received no encouragement and serves as the control group. There were two pricing treatments: a time-of-use (TOU) and a Critical Peak Pricing (CPP) pricing program.

Figure 2 summarizes the three standard, TOU, and CPP rate structures. During the time period of our study, customers on SMUD's standard rate plan (i.e., customers in the control group) paid a $10 monthly fixed charge plus $0.0938 per kWh for the first 700 kWh of consumption and $0.1765 per kWh for consumption above 700 kWh. Under the TOU program, customers paid $0.2700 per kWh for electricity consumed from 4PM to 7PM on non-holiday weekdays plus the same monthly fixed charge and $0.0846 per kWh for the first 700 kWh in all other hours and $0.1660 for consumption above 700 kWh. (On-peak consumption did not count towards the 700 kWh total.) Customers on the CPP plan paid $0.7500 per kWh for consumption between 4PM and 7PM on twelve "event days" over the course of the summer. Customers were alerted about event days at least one day in advance. Consumption outside of the CPP event window was charged at a rate of $0.0851 per kWh up to 700 kWh and $0.1665 per kWh beyond. Both the CPP and TOU rates were only in effect between June 1 and September 30 for the two summers in the study (2012 and 2013). Low-income customers enrolled on the Energy Assistance Program Rate (EAPR) received about a 30 percent discount on their electricity rates.

There were also two forms of encouragement: opt-in, where households were encouraged to enroll in the rate program; and opt-out, where households were notified that they were enrolled and were encouraged to stay in the rate program, but had the opportunity to leave the program if they wished. All encouraged groups were also offered enabling technology – an in-home display that provided real-time information on consumption and the current price.

The five randomized groups we focus on are as follows: the CPP opt-in group was encouraged to enroll in the CPP program; the CPP opt-out group was notified of enrollment and encouraged to stay in the CPP program; the TOU opt-in group was encouraged to enroll in TOU program; the TOU opt-out was notified of enrollment and encouraged to stay in TOU program; and the unencouraged, or control, group was not encouraged to participate in a rate program.

3.2 Encouragement Messages

The encouragement effort for the opt-in groups was significant relative to the opt-out groups. The opt-out groups were mailed one packet containing a letter, brochure, and business reply card. About 10 percent of the packets were sent on March 12th, 2012 and the remaining 90 percent were sent on April 5th, 2012. Recruitment activities targeted at the opt in groups are summarized in Figure 3. The encouragement effort for opt-in households consisted of two separate mailed packets, the first sent in either October 2011, to about 20 percent of the encouraged households, or November 2011, to the remaining 80 percent, and the second sent in January of 2012. The packet included a a letter, a brochure, and a postage-paid business reply card that the household could mail back to SMUD indicating their choice to either join the program or not. These materials were designed to look as similar as possible to the materials received by members of the opt-out groups. Each packet mailing was followed up within two weeks by a reminder post card. Then, in March of 2012, door hangers were placed on the doorknobs of encouraged households. Finally, an extensive phone bank campaign was carried out throughout April and May of 2012, with calls going out almost daily. For the opt-in groups, about a quarter of the households enrolled using the business reply cards. Additionally, about half of those enrolled following the packet and door hanger recruitment phase, while the second half were successfully enrolled over the timeframe of the phone campaign (though about 22 percent of these still indicated their desire to enroll by way of the business reply cards). The recrutiment materials listed generic benefits of participating in rate programs, including saving money, taking control, and helping the environment.

The TOU opt-in group received slightly different encouragement messages from the other groups because they were also part of a recruit-and-delay randomized controlled trial (which we are not incorporating into this paper). In the first packet mailed in late 2011, the households were given the same information as other groups regarding the starting date of the pricing expermiment. However, in the packet mailed in January 2012, there was text that informed them that if they decided to opt-in to the rate program, they would be randomly assigned to a start date of either 2012 or 2014. The other three groups were told that their participation date would start in 2012 if they decided to opt-in or not opt-out throughout all communications they received. This means that the set of always takers in the CPP opt-in group is slightly differnt from the always-takers in the TOU group, as the TOU always takers had to be willing while the CPP opt-in group can be directly compared to the CPP opt-out group, there is a caveat to comparisons between the

TOU opt-out and opt-in groups given the slight different wording in the recruitment materials.

4 Data and Methodology

4.1 Data Description

We use three primary sources of data: household-specific data, energy consumption and expenditure data, and weather data. The household-specific data includes experimental cell assignment, and dates of enrollment, disenrollment, or account closure due to moving. In addition, we observe whether households were on SMUD's Energy Assistance Program Rate (EAPR) for low-income customers, as well as whether or not they had set up a "My Account" online to interface with their SMUD account, and the number of times they had signed in to their My Account page. Finally, for some households, we have responses to two large-scale surveys administered to customers on the new rate programs as well as a sample of control households, including a demographic survey and a customer satisfaction survey.

We also have data on households' energy consumption and expenditures. These include data on hourly energy consumption for each household starting on June 1, 2011 and continuing through October 31, 2013, the end of the pilot period. The unit of measurement is kilowatt hours (kWh). We collect energy consumption data for households whether or not they ended up enrolled on the treatment pricing, and whether or not they opted out at any point in the pilot period. If a household moved they were not tracked to their new location, even if it was within SMUD's service territory, so data for these households ends when they moved from their initial location.

In addition to the hourly energy consumption data, billing data were also obtained for all households in the experiment. These data include the total energy (kWh) charged in each bill, as well as the total dollar amount of the bill.

Coverage of the hourly energy consumption and billing data was quite complete. While there were a handful of missing observations (less than one percent) they do not differ systematically across treatment groups, nor across those who did and did not opt in or opt out of treatment.

The final type of data we use are hourly weather data, including dry and wet bulb temperature as well as humidity. There is only one weather station in close proximity to all participants in the SMUD service area, so the weather data does not vary across households, only over time.

4.2 Validation of Randomization

Table 1 provides summary statistics by experimental group. The top four rows summarize information on billing as well as daily, peak, and the ratio of peak to off-peak energy consumption from the pre-treatment summer (June to September 2011). SMUD households consume slightly more electricity than the average U.S. household – approximately 27 kWh per day during the four summer months compared to XX kWh per day across the U.S. The ratio of peak to off-peak usage is also slightly higher than the U.S. average, reflecting high air conditioning penetration and usage in Sacremento.[3]

The bottom three rows of Table 1 summarize the household-level covariates that we were able to observe for every household in the experiment. My Account is a dummy indicating whether or not the household had signed up to use SMUD's online portal and My Account logins is the average number of log-ins across enrolled customers. EAPR is a dummy variable indicating enrollment in the low-income rate. All means are statistically indistinguishable across the various treatment groups as compared to the control group.

[TABLE 1 HERE]

Figure 4 summarizes average consumption across pre-treatment, summer weekdays (i.e., weekdays between June 1 and September 30, 2011). The left-hand side of the figure compares customers who were offered the opportunity to opt-in to either the CPP or TOU treatment to control customers, while the right hand side compares customers who were defaulted on to either the CPP or TOU plan to the same control customers. The graph highlights the variation in electricity consumption over the day, from a low below .75 kWh in the middle of the night to a peak nearly three times as high at 5PM. This consumption profile is typical across electricity consumers around the country, although SMUD customers' peak consumption tends to be slightly later than other utilities.

The graph also highlights that we cannot reject that both sets of treated households had statistically identical consumption profiles to the control households. The graphs in the bottom row of Figure 4 show the differences between treated and control, highlighting that these are well within the 95 percent confidence intervals for all hours. The standard errors for the CPP opt-out group are notably larger since that group had one tenth as many households.

[3] 9X% of the survey respondents reported having air conditioing, compared to XX% in the U.S.

[FIGURE 4 HERE]

4.3 Methodology

4.3.1 Estimating ITT for experimental treatment groups

We use a difference-in-differences (DID) specification that uses data from the pre-treatment and treatment periods to identify an intent to treat (ITT) effect. Specifically, we estimate versions of equation 1, where y_{it} captures hourly electricity consumption for household i in hour t; Z_{it} is an indicator variable equal to one starting on June 1st, 2012 if household i was encouraged to be in one of the treatment groups, zero otherwise (estimated separately for the opt-in and opt-out groups); γ_i is a household fixed effect and τ_i is an hour-of-sample fixed effect.

$$y_{it} = \alpha + \beta Z_{it} + \gamma_i + \tau_t + \varepsilon_{it} \tag{1}$$

We estimate equation 1during both event day peak hours (4pm to 7pm on the twelve CPP days in each summer) and during non-event day peak hours (4pm to 7pm on every other non-holiday weekday during the summer). The coefficient of interest isβ, which captures the average difference between treated and control groups, controlling for any pre-treatment differences by group.

4.3.2 Estimating LATE for experimental treatment groups

We use a DID instrumental variables (IV) specification with data from the pre-treatment and treatment periods to identify a Local Average Treatment Effect (LATE) effect. Specifically, we estimate versions of equation 2, where y_{it}, γ_i, and τ_i are defined as in equation 1. $Treat_{it}$ is an indicatory variable equal to one starting on June 1st, 2012 if household i was actually enrolled in treatment, zero otherwise (estimated separately for the opt-in and opt-out groups). We instrument for $Treat_{it}$ with randomized encouragement to treatment Z_{it}.

$$y_{it} = \alpha + \beta Treat_{it} + \gamma_i + \tau_t + \varepsilon_{it} \tag{2}$$

The β coefficient captures the Local Average Treatment Effect (LATE). In this specification, the LATE measures the average reduction in peak period electricity consumption (either during event days or during peak hours on non-event days) per household among customers on the experimental rate. To interpret β as a causal effect, we invoke an exclusion restriction, which implies that the encouragement (the offer to opt in or default assignment into treatment with the opt to opt out) affects electricity consumption only indirectly via an effect on participation.[4]

4.3.3 Estimating LATE for Complacents

Our experimental design allows us to disentangle the effect of the TOU and CPP rate programs on an interesting subset of households, who we label, "complacents." If we regard the opt-out program as identical to the opt-in program, where the only difference is that opt-out is a stronger encouragement mechanism than opt-in, then we can define the following household types: never-takers, which are households who would drop-out of an opt-out program and would not enroll in an opt-in program; the complacents, which are households that do not actively enroll in an opt-in program, but who also do not actively drop out of an opt-out program; and the always-takers, households who would actively enroll in an opt-in program and would remain in an opt-out program (see Figure 5).

[FIGURE 5 HERE]

We use a DID IV specification with data from the pre-treatment and treatment periods for the opt-in and opt-out groups, as shown in equation 2, where all variables are defined as above, except now $Treat_{it}$ is instrumented for with an indicator variable equal to one for observations starting on June 1st, 2012 if a household was encouraged into the opt-out treatment group only.

This IV specification provides an intuitive way to isolate the LATE of the rate program on the complacents. To estimate this, we assume that (a) the always-takers participating in an opt-in program are affected by the program in the same way as they would if they were participating in an opt-out program; and (b) the never-takers who were given an opt-out encouragement and the never-takers who were given an opt-in encouragement both exhibit zero response.

[4] We have run a simple analysis to test the validity of the exclusion restriction assumption. An explanation of this and results from this test are shown in Appendix XXX.

5 Main Results

5.1 Default Effects in Program Adoption

Table 2 summarizes take up in the opt-in and opt-out groups. The columns titled “Initial participation” report information at the beginning of June 2012, the month the time-varying rates went into effect, and the columns titled, “Endline participation” report information as of the end of the second summer (September 30, 2013). In both sets of columns, the first number reflects the share of people on the time-varying rate while the second column reports the eligible population.

[TABLE 2 HERE]

The Initial participation results in Table 2 provide striking evidence of the default effect. For both the CPP and TOU rates, approximately 20 percent of those approached agreed to try the new rates on an opt-in basis while fewer than 5 percent opted out when defaulted onto the new rate structure, leaving over 95 percent of the customers on the new rates in the default treatment.

One question is whether these differences are idiosyncratic to SMUD. In terms of opt-in programs, SMUD was more successful than expected at recruiting customers onto the rate. The company’s expectations, and the basis for our ex ante statistical power calculations, were that between ten and fifteen percent of customers would opt-in. On the other hand, given that SMUD customers are generally satisfied with the utility and trust its recommendations, they may have been more likely to accept the default. We anticipated that approximately 50 percent of the customers would remain on the rate with opt-out.

To interpret the “Endline Participation” columns, it is important to understand how we are describing the eligible population. If people moved, they were removed from the pilot program, meaning they were no longer eligible for the time-based rate structure, even if they moved within SMUD’s service territory. Also, any new occupant was not allowed into the pilot program. The numbers in Table 2 report rates and populations after dropping movers. For instance, the eligible population of the CPP opt-in group fell from 1589 to 1169 because 420 households (approximately 26 percent) moved between June 2012 and September 2013. SMUD reports move rates of approximately 20 percent per year among their residential population, so a move rate of 26 percent over a 16-month period that includes the summer, when moves are most likely, is reasonable. Also, across the four columns, the move rates are very similar, ranging from 23.5 percent in the CPP opt-out group to 26.7 percent in the TOU opt-in. Even for these endpoints, the rates are not statistically significantly different from one another (t-statistic on the difference equals 1.7).

5.2 Choice Modification

We also observe some changes in participation after the program started. While customers in the opt-in group were not allowed to enroll after June 1, 2012 and customers in the opt-out group who had already opted-out were not allowed to change their minds and enroll, we do observe customers in both groups who actively requested to revert to the standard rate.

The final column of Table 2 reports the "Attrition rate," which reflects the difference between initial and endline participation, divided by the initial participation rate. Participation in both of the opt-in groups fell by fewer than 1.5 percentage points, reflecting fewer than 10 percent of the original participants. Participation in both of the opt-out groups fell by more percentage points (6.8 in the case of CPP opt out, 96.2 − 89.4, and 5.3 in the case of TOU opt out), but again reflected fewer than 10 percent of the original participants.

Though the small share of households that dropped out makes tests comparing attrition rates between the opt-in and opt-out groups relatively low powered, the Appendix reports results from a hazard analysis of drop outs. Several interesting patterns emerge from that analsysis. First, although the rates of attrition over the entire study were similar, the opt-in participants (both TOU and CPP) dropped out sooner than opt-out. For households in the opt-out groups, the reminder sent to participants before the second summer had a strong effect on drop-outs.

In sum, sections 5.1 and 5.2 provide strong evidence of a default effect and relatively little evidence of subsequent re-optimization.

5.3 Follow-on Behavior

5.3.1 Intent to Treat (ITT) Effects

Table 3 reports a difference-in-differences (DID) specification of equation 1 that uses data from the pre-treatment and treatment periods to identify an intent to treat (ITT) effect. The left two columns of Table 3 use only data from "event" days. In the post-treatment period, these correspond to days when a CPP event was called. In the pre-treatment period, these correspond to weekdays during the summer of 2011 where maximum daily dry bulb temperature exceeded 96 degrees Fahrenheit, the average temperature on event days during the treatment period. The right two columns use data from all other summer weekdays. In all cases the analysis is limited to the peak periods of the relevant set of days (4PM to 7PM).

[TABLE 3 HERE]

If we interpret the coefficients in Table 3 as estimates of the causal impact of encouragement to join the time-varying rates on electricity consumption, we conclude that providing households the opportunity to opt-in to the CPP treatment leads to a reduction in consumption on event days of 0.130 kWh per household that received the offer . The estimate for the opt-out group is considerably larger at 0.312 kWh per CPP household that received the offer. .

The coefficients in the non-event day columns suggest that CPP customers reduced their consumption during peak hours even on days when their rates were not adjusted (by 0.0276 kWh per household in the opt-in group and 0.0923 kWh per household in the opt-in group). In fact, the CPP customers faced slightly lower rates than the control group. The kWh reductions are considerably smaller as compared to event days for the CPP households, but still statistically significant. This is consistent with habit formation, learned preferences, (e.g. if households learn that they can comfortably open windows instead of turning on the air conditioning), or a fixed adjustment cost (e.g., if customers set programmable thermostats to run air conditioning less between 4 and 7 PM on all days, even when they only face higher prices on a subset of those days).

In the case of the TOU group, which were charged their peak prices during all weekdays and not just event days, the results show that average households reduced their daily peak consumption by 0.089 kWh per household in the opt-in treatment, and 0.133 kWh per household in the opt-out treatment on days that were called as event days for CPP customers (i.e., relatively hotter days), and by 0.054 kWh per household in the opt-in treatment, and 0.101 kWh per hour in the opt-out treatment on all other peak days. Given that non-event-day consumption is considerably lower, the results are approximately the same in percentage terms (3.1-3.6% for the opt-in group and 5.5 - 5.8% for the opt-out group).

The fact that the estimates between the opt-in and opt-out experimental groups differ by so much in both the CPP and TOU context suggests that the complacents (included in the opt-out group but not the opt-in group) are reducing consumption, under the assumption that the always takers in the opt-out group are responding in a similar manner to what they would do if they had actively opted in to the new rate. Both coefficients are highly statistically significant, and we can reject that each is equal to the other with a high t-statistic. Future work will develop formal tests of their equality

Results from a specification similar to Table 3 that does not use the pre-period data and simply compares treated households' consumption to the control households' during event and non-event peaks yield qualitatively similar results, suggesting that the average reductions for the opt-out group were nearly 3 times larger than the average reductions for the opt-in group for CPP and 1.5 to 2 times larger for TOU. The coefficient estimates do differ from those reported in Table 3 since there were some pre-period differences by group, even if those differences are not statistically significant.

Figure 6 depicts hour-by-hour results for the CPP event days. The results suggest that consumers are reducing consumption in the hours before the peak period, statistically significantly so in the 3PM and 2PM hours for the opt-in group. There is no evidence of negative spillovers (i.e., offsetting increases in consumption) to other hours.

[FIGURE 6 HERE]

5.3.2 Local Average Treatment Effects (LATE)

Table 4 reports coefficients from several versions of the instrumental variables specification described in equations 1 and 2 above. Similar to Table 3 , the columns on the left of the table report results during the CPP event hours and the columns on the right report results during non-event-day peak hours. The top of the table is based on customers in the CPP tretaments while the bottom is for customers in TOU treatments.

To interpret the coefficients in Table 4 as causal local average treatment effects, we invoke an exclusion restriction, which implies that the encouragement (the offer to opt-in or opt-out) affects electricity consumption only indirectly via an effect on participation. The results in columns (1) and (2) suggest that the always-takers in the opt-in CPP group reduced consumption during event-day peaks by almost twice as much as the combination of always takers and complacents in the CPP opt-out group (0.665 compared to 0.338 kWh per household). The magnitude of the reduction for the opt-in group (665 watts per hour) is quite large and suggests consumers did more than simply turn off a few light bulbs. Given that their rates increased by almost 100 percent, though, this reduction off a mean of almost 2,500 watts is consistent with a price elasticity of approximately -0.25, which is on the high side of other short-run demand elasticities estimated for electricity consumption, though typically those estimated demand reductions over longer time periods (EPRI 2012). In columns (4) and (5), we see again that households in both the opt-in and opt-out CPP treatments were reducing their consumption on non-event peak days significantly.

In the case of the TOU treatments, the LATE estimates indicate that always-takers reduced consumption during daily peaks that were called as event days for the CPP treatment by about three times as much as the combination of always-takers and complacents in the TOU opt-out group (0.470 relative to 0.140 kWh per household), and almost three times as much (0.284 relative to 0.106 kWh per household) during non-event regular peak days.

[TABLE 4 HERE]

Comparing the results in columns (1), which desribes the always-takers, to columns (3), which describes the complacents, suggests that CPP complacents responded about 2.5 times less than always takers during event hours. Complacents were more similar to always takers during non-event peak hours, reducing by only 50% less.[5] This is consistent with complacents being more likely than always takers to undertake a single action to reduce their consumption during all peak hours – both event and non-event, such as programming their thermostats to keep their homes slightly warmer during peak hours. Always takers, by contrast, appear more likely to fine-tune their behavior on event days. Countering this interpretation, always takers on TOU also reduce their consumption more on CPP event-days, eventhough their rates to not change. However, the differences between always takers on event and non-event days is more pronounced for customers on the CPP rate.

Figure 7 is analogous to Figure 6 , and shows the LATE estimates for event days across the four treatment groups relative to the control group.

[FIGURE 7 HERE]

While these results underline the fact that complacents responded less than always-takers, they did respond to these rates. Therefore, given that there are so many more of them exposed to the rates under an opt-out experimental design, the aggregate savings from an opt-out designis significantly higher than from an opt-in design (as is made evident in Table 3 and Figure 6).

[5] Note that the coefficient estimates for the opt-out group in Table 4 are basically equal to the weighted sum of the coefficients for the always takers (e.g., -0.665 for CPP event hours) and the complacents (-0.250), with weights set equal to the share of always takers relative to total opt-out enrollees and one minus this number from Table 2 .

5.4 Bill Impacts and Cost-Effectiveness

This subsection summarizes the impact of the pricing plans on customer bills and SMUD's net revenues. Table 5 reports IV versions of equation 2 that use customer-by-month observations and total bill amount as the dependent variable. The coefficient estimate in the top panel, column (1) suggests that bills for customers who opted in to the CPP rate plan fell by more than 10% on average, with a mean reduction of \$8.70 on an average summer bill of nearly \$80. Bills for the typical participant in the opt-out group fell by much less – around \$3.80 for the group overall and only slightly less for the complacents. This is consistent with the results presented in Table 4 suggesting that the opt-out group overall and complacents in particular reduced consumption by less during critical peak periods. Of course, customers may have made adjustments that were either costly from a monetary or welfare perspective, so bill reductions should not be interpreted as necessary welfare gains. We provide more insight on consumer welfare in the next section.

Table 6 analyzes the pricing programs from the perspective of the utility, comparing the costs of enrolling participants in the programs versus the benefits in terms of the avoided costs associated with meeting customers demand on event and non-event days. The results in this table are based on information in Potter et al. (2014). The costs, summarized in the first row, include the recruitment costs, plus the cost of the in-home devices multiplied the probability that a participant requested the device. The benefits reflect both the avoided energy costs and the expectation of avoided future generation capacity. All costs are marginal, in the sense that they do not account for the fixed cost SMUD incurred to set up the program. The final two rows, reflecting the net benefits and benefit cost ratio, suggest that all the programs generate more benefits than costs, excpet for the TOU opt-in. For both CPP and TOU, the opt-out treatments are more cost-effective because the recruitment costs are so much lower. The benefit-to-cost ratio for both opt-out programs exceed 2. Note that in addition to the possibility of other sources of utility losses or gains, a full welfare analysis would recognize avoided pollution costs. Changes in energy consumption are a small share of the benefits to the utility – they are dominated by the capacity benefits. This suggests that accounting for pollution externalities associated with energy generation would be low. This is consistent with the fact the utilities must spend capital to build plants that are used for a very small number of peak hours a year. Reducing demand in a small number of peak hours avoids the need to build these plants.

6 What Explains the Default Effect?

In addition to assessing the benefits and costs of the default manipulation from the perspective of the utility, we are also interested in the larger welfare implications of the default effect. If consumers who are susceptible to the default effect are in fact well suited to (or prefer) time-varying electricity pricing, the default manipulation can leave everyone better off, implying an unambiguous welfare improvement. At the other extreme, if all complacent customers are unhappy with the new default choice, but face switching costs or other barriers that prevent them from opting out, the default effect could reduce overall welfare.

From a social welfare perspective, understanding *why* people are predisposed to choose the default option is important. Prior studies have identified several potential explanations for default effects. But, the critical task of evaluating the welfare effects of default options has been almost entirely ignored, in large part because it is difficult to identify precisely which mechanisms are at work. Although a full welfare analysis is beyond the scope of this paper, we are uniquely positioned to investigate alternative explanations for the default effect in the choice context we study. We use detailed information about the determinants of the initial choice, together with rich data on follow-on behavior, to shed light on what factors give rise to the default effect and what this could imply for consumer welfare.

There are a number of alternative explanations for the default effect that could apply in this setting. Some of these presume stable, well-defined preferences. In one scenario, consumers have a clear understanding of their preferences for one choice over another, but choose the default to avoid incurring switching costs that offset the gains from switching to the preferred choice. In another scenario, consumers must incur a cost to collect the information required to make a choice that is most consistent with their well-defined preferences. If the effort associated with collecting and processing this information exceeds the expected benefits from switching, it can be quite rational to avoid exerting this effort, and instead choose the path of least resistance (i.e. the default choice). Under either of these scenarios, the key challenge is to identify the distribution of switching and information costs, and incorporate these in welfare calculations.

Under an alternative perspective, preferences are constructed – versus uncovered – by consumers as they weigh and experience alternative options. When faced with an unfamiliar choice, individuals may not have well defined preferences for one option over another. In this setting, observed choices reveal not only the agent's valuation of the alternatives, but also the processing strategies used to construct the prefered choice. This perspective introduces some additional heuristic explanations for default effects.

For one, people may interpret the default choice as an informative suggestion or endorsement helping to guide an otherwise uninformed choice. Or, the default choice can serve as an anchor or point of reference. If preferences are formed as customers experience the new pricing structure, welfare analysis becomes more complicated. Standard approaches that seek to rationalize default effects using switching costs and information costs may overestimate the role of these costs.

We cannot definitively distinguish between the alternative explanations for the default effect we document. In this section, we describe heterogenous patterns in default procilivty and systematic differences in follow-on behavior which provide suggestive evidence on both the mechanisms behind the default effect and the associated welfare implications.

6.1 Heterogeneity in Default Sensitivity

Table 7 summarizes some important household-characteristics (including prior program participation, income, and projected bill savings) by sub-group. It is straightforward to summarize these variables for never-takers (i.e., households assigned to the opt-out group who actively opt-out) and always-takers (i.e., households assigned to the opt-in group who actively opt-in). To impute the summary statistics for complacents, we leverage the random assignment across opt-in and opt-out groups which implies that the share of always-takers, never-takers, and compliers will be the same in expectation across the two groups. [6] The three columns on the right of Table 7 summarize statistical significance levels for each pairwise comparison. The top of the table applies to the CPP treatments and the bottom to TOU.

"My Account" and "My Account logins" reflect actions that customers could take to monitor their consumption in the pre-treatment period. Customers who have historically engaged with existing utility programs, either enrolled in the online My Account service, or, conditional on enrollment, frequently accessed the service, are more likely to take an active choice and either opt-in or opt-out. This is true for both CPP and TOU treatments. In both cases, the differences between compliers and always takers as well as between compliers and never takers are statistically significant for My Account. The difference between compliers and always takers is statistically significant for number of logins. If we interpret these variables as proxies for attentiveness, we find that complacent households have historically been significantly less attentive to their electricity consumption. This could reflect that members of the complacent group incur

[6] Specifically, we calculate the mean of each variable for the complacents as follows: μ_C=(μ_Opt-out - p_AT*μ_AT)/(p_Opt-out - p_AT) where μ_Opt-out and p_Opt-out are the means and proportions for all participants in the opt-out group and μ_AT and p_AT are the means and proportions for all participants in the opt-in group.

higher costs to engage and monitor their use in general. The lack of engagement with the existing programs could also raise the costs of making an active choice about enrolling in time-varying pricing.

"CPP Savings" and "TOU Savings" summarize projected summer bill savings under the time varying rate (relative to the standard rate) based on household consumption in the pre-treatment summer (2011). A significant fraction of consumers in our study would pay lower electricity bills in the summer if they moved onto a time-varying rate, even if they made no change in their consumption patterns. Following industry practice, we use the term "structural winner" to refer to these customers, realizing that "structural losers" could also "win" by switching to a time-varying rate if the welfare cost of adjusting consumption patterns is more than offset by the associated bill savings.[7]

If customers were making a well-informed choice, we might expect to find that our projected savings variable is a good predictor of participation in the time-varying rates. However, for CPP, projected gains and losses are not a good predictor of the decision to opt in or the decision to opt out, even among the apparently informed customers who are monitoring their accounts. In fact, households that actively opt into the CPP rate have lower projected savings on average than households that opt out, although the difference is not statistically significant.[8] Customers appear slightly better at predicting whether they will benefit under the new pricing with TOU. Always takers would lose significantly less than compliers, although never takers appear poised to lose the least under TOU.

"Low Income" summarizes participation in the utility's low-income electricity pricing program. We find that low-income consumers are much more likely to opt in and a little less likely to opt-out, though the second difference is not statistically significant for CPP. Because households must actively sign up for this low-income rate, these households may be more attentive than other low income households in the sample.

[TABLE 7 HERE]

In sum, we find systematic differences in the extent to which customers have historically been engaged in monitoring consumption, with complacent households sigificantly less engaged than other households

[7]The "CPP/TOU Savings" variables are based on consumption in one summer; year-to-year variation in weather and other factors that determine demand will generate year-to-year variability in these structural gains and losses. However, customers who consume a relatively large share of their electricity in off-peak hours can expect to gain with some certainty. We find that over half of the control group households would see structural gains in both the pre and post-treatment periods. See Appendix D.

[8]Under the assumption that customers enrolled in My Account are more informed about their usage, we also calculated projected savings on CPP conditional on participation in My Account. Even among these more attentive customers, program participants are not associated with higher projected savings.

in the sample. This is consistent with the default effect reflecting inattention (rational or otherwise). We also find that structural gains are not a good predictor of program participation, even among engaged households. The average projected gain or loss is quite small (average gains among winners, and average loss among losers, are on the order of $15 over an entire summer). Given that gathering information about consumption patterns and alternative rate structures to make an informed decision requires effort, inattention to these savings could be rational.

6.2 Heterogeneity in Follow-on Behavior

Table 8 tests for systematic heterogeneity in the electricity consumption response to time-varying prices with respect to our proxy for attentiveness (My Account participation). Specifically, we estimate a more flexible specification of Equation 2 that includes an interaction between the participation indicator and the My Account indicator. Note that the direct effect of My Account is absorbed by the customer fixed effect.

The coefficients on the interaction terms are negative in all 12 cases and statistically significant in 6 of those 12. In other words, customers who had signed up for My Account prior to the study reduced consumption by significantly more on average during both event and non-event peak hours. The most striking differences are found among complacents. The coefficient on the interaction term is negative and larger than the coefficient on the treatment variable alone, though only statistically significant during event hours (column (3)). We note that the responses of complacents enrolled in My Account appear more similar to always takers than for complacents who have not activated My Account.

For TOU, customers enrolled in My Account are more responsive across the board, and provide significantly higher peak demand reductions than customers who are not enrolled in My Account (except in the case of complacents). The effects are large for complacents, even proportionately larger than for always takers, but the point estiamtes are small, so they are not statistically significant.

[TABLE 8 HERE]

Table 9 tests for systematic variation in price responsiveness across income groups. We separate households that are participating in the utility's low-income electricity pricing program from the larger sample. Low income households are of particularly of interest to regulators; financial impacts of the rates may lead to larger welfare impacts. The households in our study that are participating in this low-income program not only have relatively low incomes, but they also pay relatively low rates.

The results presented in Table 9 indicate that always takers on the low-income rate are significantly less responsive during event and non-event hours for both the CPP and the TOU treatments. This indicates that low-income customers that actively opted in did not provide as much peak savings. Among compliers, the average demand response among low-income customers is also smaller during critical events, although the differences are not statistically significant. Overall, these results are not consistent with the story that low income customers were disproportionately impacted by the default effect. Recall that low income consumers were more likely to opt-in to the time-varying pricing programs. The demand response of low income customers who were susceptible to the default effect is statistically indistinguishable from the other complacent households.

[TABLE 9 HERE]

6.3 Persistence in Follow-on Behavior

Our study period includes two years of post-intervention data. This allows us to analyze how electricity demand response to the time-varying rates evolves over time. In particular, we can test for differences in this evolution across customers who actively opted-in and the complacent households who were nudged in by the opt-out encouragement.

We modify Equation 2 to include an interaction between the treatment indicator and an indicator for the second summer. Table 10 reports the estimation results. For the CPP treatments, the interaction term is positive for the always takers in the opt-in group (columns (1) and (4)) and negative for the complacents (columns (3) and (6)). Three out of four of the coefficients are statistically significant.[9] This pattern suggests that demand response is attenuating over time among always takers. I contrast, the average demand response is increasing over time among compliers. This could be due to a growing number of complacents responding over time, or an increasing demand response from those complacent customers who had been actively responding in the first summer.

On average, the extent to which complacent households rely on fixed adjustments (such sa re-programming their thermostat) versus variable adjustments (fine-tuning thermostats during critical event days) seems to be increasing over time. In the second summer, the demand reduction among complacents on the CPP

[9]The results are not as pronounced for the TOU treatment, although column (1) suggests that the always takers in the opt-in group are responding less over time.

rate increases by a factor of two during the *non-critical* peak hours when prices have fallen relative to the standard rate.

Overall, these results are consistent with the complacents gradually learning about and acclimating to the time-varying rate, and less consistent with a scenario in which complacents knew they would dislike the rate but elect to remain on account of high switching costs.

[TABLE 10 HERE]

6.4 Survey Results

Another source of evidence on households' tastes and decision process is a set of follow-up surveys that SMUD conducted after the pricing program ended. The survey was sent to all households enrolled on the CPP and TOU pricing plans and a subset of the control group. While the survey respondents are by no means a random subset of the larger sample, the responses can provide some insight into consumers' motivations and sentiments about the pricing programs. One notable fact is that the opt-out participants were less likely to respond to the survey – 26% for opt-out (N=566) versus 36% for opt-in (N=183), consistent with the general finding thus far that complacents tend to be less engaged and less responsive. Also, only 60% of the respodnents from the opt-out groups affirmed that they understood they were paying time-varying rates, compared to around 85% of the respondents from the opt-in group.

Survey responses generally suggest that customers are not averse to the new pricing plans. In both the opt-in and opt-out groups, fewer than 7% disagree with the statement, "I want to stay on my pricing plan." More of the opt-in customers strongly agree with that statement and more of the opt-out customers express, "no opinion," perhaps indicative of their complacency. Similarly, across both groups, almost 90% of respondents are either "Very satisfied" or "Somewhat satisfied" with their current pricing plan, with no statistically significant differences across those two categories by group. In contrast, only 80% of the control group are "very" or "somewhat" satisfied with the standard rate.

The results in this section suggest that customers who are more engaged with utility programs are more likely to make an active choice and either opt in to or opt out of the dynamic pricing programs, while customers who were expected to have lower bills on the program without changing their behavior (so-called "structural winners") were no more likely to enroll in the program, even if they were engaged in utility programs. We find these patterns inconsistent with explanations for the default effect that rely on consumers performing well-informed, cost-benefit calculations before making their choice and more consistent with other explanations, such as inattention and possibly some form of constructed preferences.

Once on dynamic pricing, consumers who were more attentive are also more likely to respond to the prices, although we still see significant reductions by the less attentive consumers in both the always taker and complacent populations. We also see convergence between always takers and complacents in the second summer, which we take as evidence that nudged consumers acclimated to the new pricing regimes. Finally, at least among consumers who responded to the survey, there seems to be general acceptance of dynamic pricing. In sum, we see these results as consistent with a scenario where consumers are nudged onto the rates, perhaps because they are not paying attention, and once on the rates, they learn to adjust to them and some even prefer them to standard rates.

7 Conclusion

The default effect is arguably the most powerful and consistent behavioral outcome in economics, with examples identified across many settings, including health care, personal finance and internet privacy settings. This paper studies the default effect in a new context – time-varying pricing programs for electricity. In an experiment with SMUD, a large municipal utility in the Sacramento area, we randomly allocated residential customers to one of three groups: (1) a treatment group in which they were offered the chance to opt in to a time-varying pricing program, (2) a treatment group that was defaulted on to time-varying pricing unless they opted out, and (3) a control group. We see stark evidence of a default effect, with only about 20% of customers opting into the new pricing programs and over 90% staying on the programs when it was the default option. This holds for both critical peak pricing and time-of-use programs.

Our study offers several innovations relative to the existing literature on default effects. First, in addition to observing the initial decision, where the consumer was influenced by the default effect, we also track follow-on behavior. We distinguish between follow-on behavior that modifies the orginal choice, such as opting out of the dynamic pricing program once it's begun, and behavior that is conditional on, but distinct from, the original choice. In our case, this last type of behavior involves taking steps to respond to the dynamic changes in electricy prices. We argue that this conditional behavior can be equally, if not more, important than the original choice. In our context, the dynamic pricing programs will have no impact on SMUD's costs or social welfare if consumers do not ultimately respond to the pricing programs.

This type of follow-on behavior is also important in other contexts. For example, decisions about what type of health care plan to enroll in will likely impact decisions about how much health care to consume. To our knowledge, ours is the first study to identify and study follow-on behavior.

We find that consumers do respond to the time-varying prices, even if they did not actively select them. In particular, the complacents in our study (i.e., consumers who would not have actively enrolled in the pricing program but did not opt out) reduced their consumption during critical peak pricing periods by about 10%, when the price of electricity increased by nearly a factor of 10. Always takers, who actively selected the rates, reduced consumption by more than 25%, although over time, the always takers respond by less and complacents by more.

Our second innovation is to analyze the initial decisions and follow-on behavior across different groups in our study in order to draw inferences about the likely welfare impacts of the default effect in our context. While our conclusions about the welfare effects are speculative, our findings cast doubt on explanations for the default effect based on high transaction costs. We find more support for explanations under which consumers are not paying attention to the initial choice, but come to understand it and like it.

Policy makers in the electricity sphere are actively contemplating the costs and benefits of assigning customers to time-varying pricing. Our results suggest that placing households onto time-varying pricing by default can lead to significantly more customers on time-varying pricing and, more importantly, significantly higher responses to price changes. Going forward, policy makers are focused on increasing the share of electricity generated from renewables, like wind and solar, driven by concerns over the environmental costs of fossil-fuel-fired electricity generation. Numerous studies have shown that increasing the share of renewable capacity on an electricity system can accentuate the volatility of wholesale prices, which can range from negative, when there is high wind and little demand, to sudden price spikes as the sun goes down and solar plants stop generating while consumers turn on lights. In this context, time-varying pricing can be even more important to reflect costs, and simulations have shown that providing consumers with the opportunity to respond to wholesale price changes may be one of the most cost-effective ways to smooth this price volatility. While the relatively simple time-based price structures that we consider in this study may not capture much of the price volatility caused by renewables, they are a step in that direction.

References

Abadie, A. and S. Gay (2006). The impact of presumed consent legislation on cadaveric organ donation: A cross-country study. *Journal of health economics 25*(4), 599–620.

Bellman, S., E. J. Johnson, and G. L. Lohse (2001). To opt-in or opt-out?: It depends on the question. *Communications of the ACM 44*(2), 25–27.

Beshears, J., J. J. Choi, D. Laibson, and B. C. Madrian (2009). The importance of default options for retirement saving outcomes: Evidence from the United States. In *Social security policy in a changing environment*, pp. 167–195. University of Chicago Press.

Carroll, G. D., J. J. Choi, D. Laibson, B. C. Madrian, and A. Metrick (2009, November). Optimal Defaults and Active Decisions. *The Quarterly Journal of Economics*, 1639–1674.

Choi, J. J., D. Laibson, B. C. Madrian, and A. Metrick (2002). Defined contribution pensions: Plan rules, participant choices, and the path of least resistance. In James M. Poterba (Ed.), *Tax Policy and the Economy, Volume 16*, pp. 67–114. Cambridge, MA: MIT Press.

Choi, J. J., D. Laibson, B. C. Madrian, and A. Metrick (2004). For better or for worse: Default effects and 401 (k) savings behavior. In *Perspectives on the Economics of Aging*, pp. 81–126. University of Chicago Press.

Johnson, E. J., S. Bellman, and G. L. Lohse (2002, February). Defaults, Framing and Privacy: Why Opting In-Opting Out. *Marketing Letters 13*(1), 5–15.

Johnson, E. J. and D. Goldstein (2003, November). Do Defaults Save Lives? *Science 302*(5649), 1338–1339.

Johnson, E. J., J. Hershey, J. Meszaros, and H. Kunreuther (1993). Framing, probability distortions, and insurance decisions. *Journal of Risk and Uncertainty 7*, 35–51.

Kressel, L. M. and G. B. Chapman (2007). The default effect in end-of-life medical treatment preferences. *Medical Decision Making 27*(3), 299–310.

Madrian, B. C. and D. F. Shea (2001). The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior. *The Quarterly Journal of Economics 116*(4), 1149–1187.

Park, C. W., S. Y. Jun, and D. J. MacInnis (2000). Choosing what I want versus rejecting what I do not want: An application of decision framing to product option choice decisions. *Journal of Marketing Research 37*(2), 187–202.

Pichert, D. and K. V. Katsikopoulos (2007). Green defaults: Information presentation and pro-environmental behaviour. *Journal of Environmental Psychology 28*(1), 63–73.

Potter, George, and Jimenez (2014).

Sims (2005).

A Exclusion Restriction

[TABLE ?? HERE]

B Variable Correlations

[TABLE ?? HERE]

C Hazard Analysis

[HAZARD ANALYSIS RESULTS HERE]

D Structural gains as a predictor of participation

Figure 1: Hourly electricity demand (SMUD) and wholesale electricity price (CAISO)

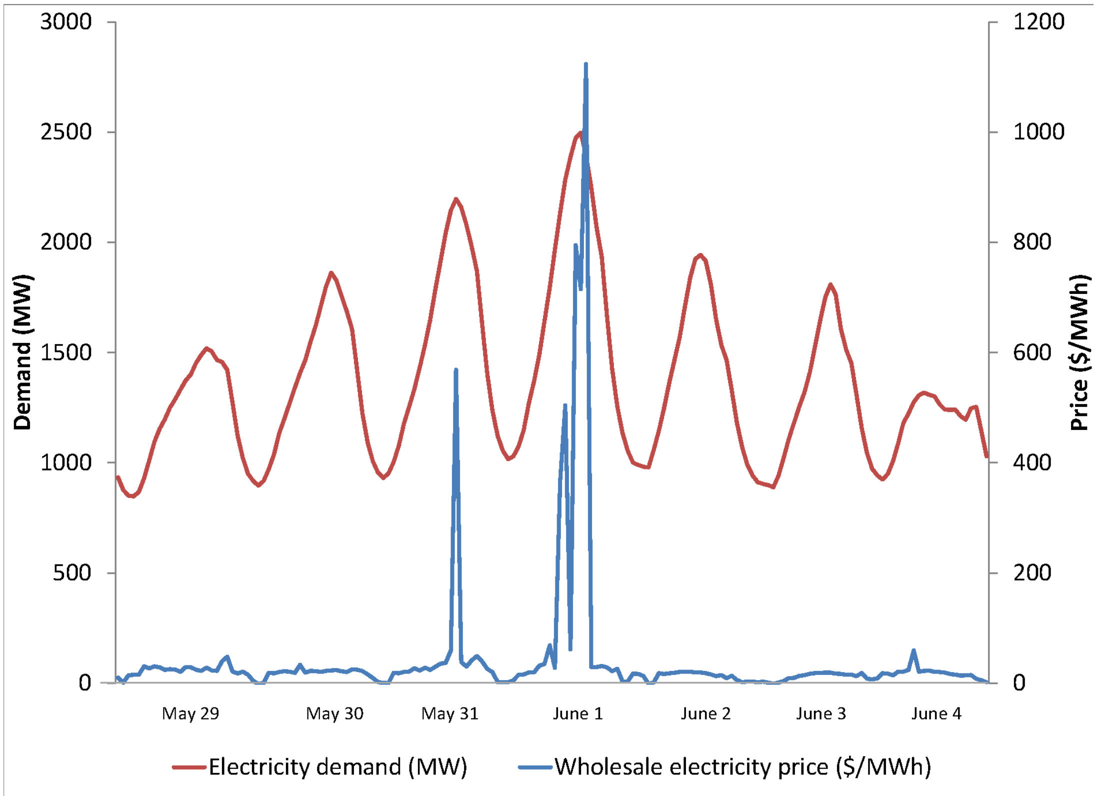

Note: Hourly electricity demand in red, wholesale spot price for electricity for CAISO in blue.

Figure 2: rates

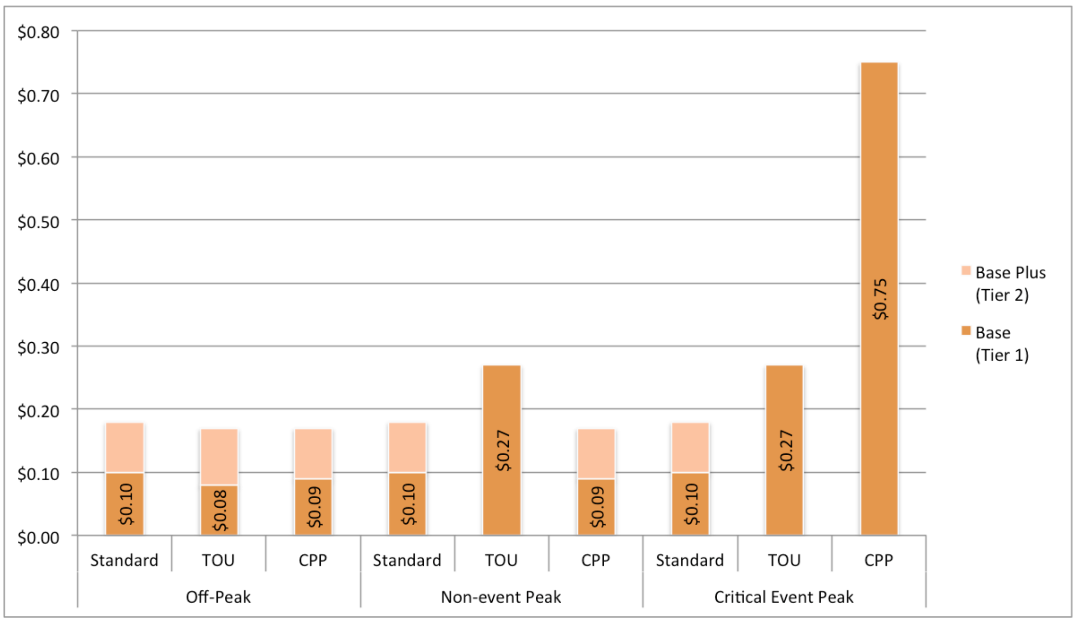

Note: On the base rate, customers are charged $0.1016 for the first 700 kWh in the billing period, with additional usage billed at $0.1830. Participants on the TOU rate were charged an on-peak price of $0.27/kWh between the hours of 4 PM and 7 PM on weekdays, excluding holidays. For all other hours, participants were charged $0.0846/kWh for the first 700 kWh in each billing period, with any additional usage billed at $0.1660/kWh. On the CPP rate, participants were charged a price of $0.75/kWh during CPP event hours. There were 12 CPP events caller per summer on weekdays during the hours of 4 PM and 7 PM on weekdays. For all other hours, participants were charged $0.0851/kWh for the first 700 kWh in each billing period, with any additional usage billed at $0.1665/kWh.

Figure 3: Encouragement

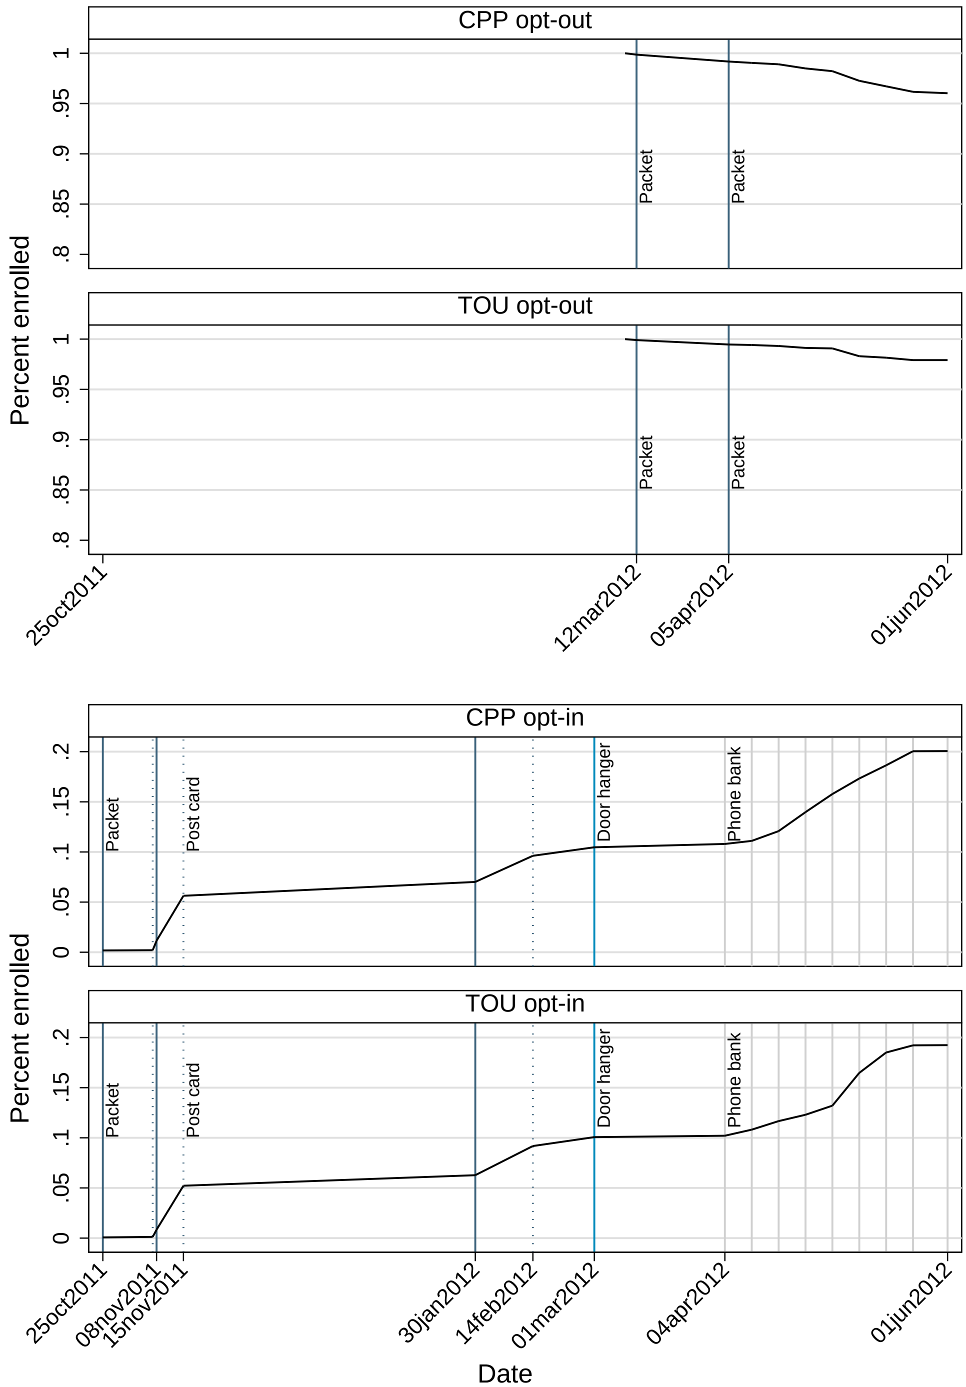

Note: This figure depicts the encouragement efforts experienced by the various treatment groups. In the case of the opt out groups, the two vertical lines indicate dates on which packets were mailed out to the households. In the case of the opt in households, the first three solid vertical lines are dates on which packets were mailed out, the three dotted vertical lines indicate dates on which follow-up post cards were mailed out, then on March 1st, 2012 door hangers were distributed. Finally, between April 4th and June 1st 2012 there was a phone bank campaign, with calls going out almost daily. This period is indicated by the grey vertical lines during that period. The figure shows the percent of those randomized into encouragement that were enrolled in treatment over the course of the recruitment efforts.

Table 1: Comparison of means across treatment groups

	Controls	CPP		TOU	
	(1)	(2) Opt-in	(3) Opt-out	(4) Opt-in	(5) Opt-out
Average kWh per day (usage)	26.62	26.80	26.92	26.48	26.37
		(0.198)	(0.584)	(0.168)	(0.351)
Peak to off peak ratio	1.791	1.790	1.796	1.795	1.798
		(0.006)	(0.020)	(0.006)	(0.012)
Average monthly bill ($)	109.2	109.5	109.2	108.3	107.9
		(0.967)	(2.939)	(0.844)	(1.769)
My Account	0.425	0.430	0.442	0.432	0.419
		(0.006)	(0.017)	(0.005)	(0.010)
My Account log-ins	6.714	7.095	7.139	6.818	6.349
		(0.468)	(1.337)	(0.410)	(0.867)
Low income	0.194	0.196	0.210	0.200	0.200
		(0.005)	(0.014)	(0.004)	(0.008)
Households	45,839	9,190	846	12,735	2,407

Note: Cells contain the group means. Standard errors of the difference in means between the treated and control groups are in parentheses. Average kWh during peak hours reflects consumption on weekdays between June 2011 and September 2011 from 4-7pm. Peak-to-off-peak ratio is the by-customer hourly kWh used during peak times divided by the hourly kWh used during non-peak times, averaged across customers. Households are eligible for the low income rate if their income does not exceed 200 percent of the federal poverty level.

Figure 4: Pre-treatment hourly energy usage comparisons

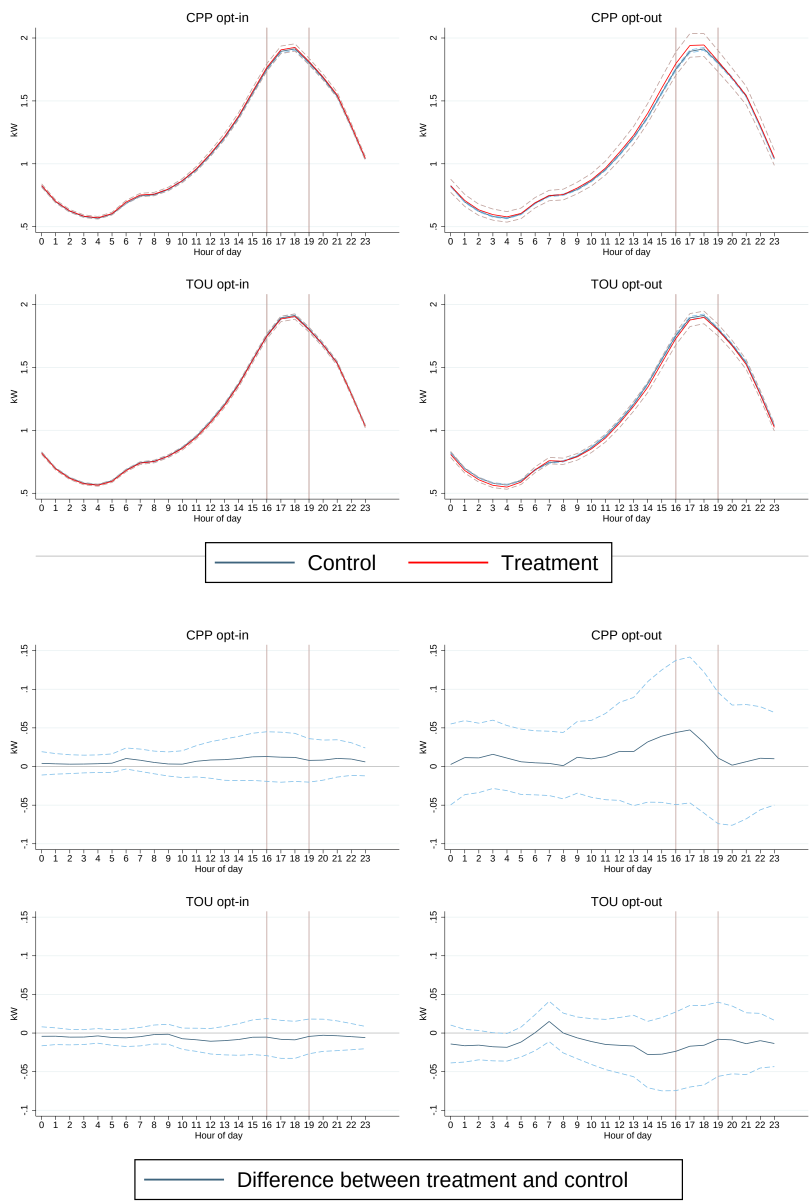

Note: Hourly average energy use in the pre-treatment period for each of the four treatment groups (CPP opt-in, CPP opt-out, TOU opt-in and TOU opt-out) is shown along side that of the control group in the top four panels. The bottom four panels show the difference in hourly energy use between each treatment group and the control group. In all cases the dashed lines indicate the 95 percent confidence interval.

Figure 5: Identification of always takers, complacents, and never takers

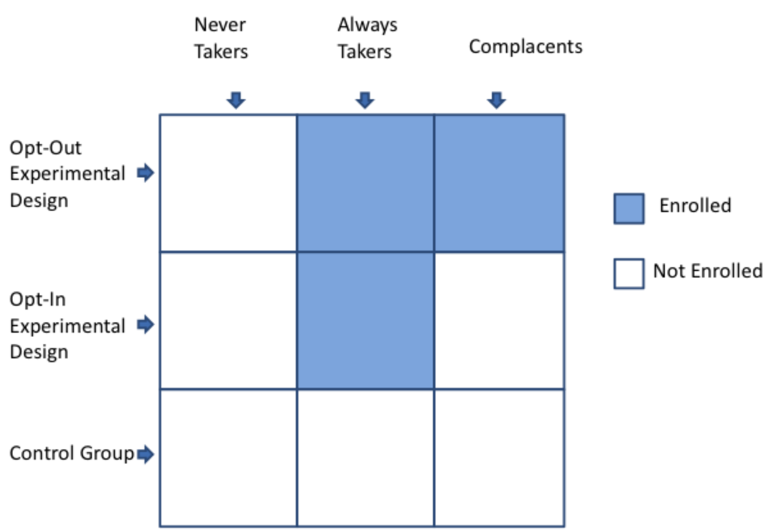

Note: Rows signify the three groups into which customers in our sample were randomly assigned: opt-out, opt-in, and control. Columns signify types of customers. Shading indicates that the customer type enrolls in time-based pricing program under the associated experimental group.

Table 2: Participation rates

	(1) Initial participation		(2) End-line participation		(3) Attrition rate
CPP opt-in	0.203	(1,589)	0.189	(1,169)	0.068
CPP opt-out	0.962	(703)	0.894	(538)	0.071
TOU opt-in	0.195	(2,115)	0.181	(1,551)	0.072
TOU opt-out	0.979	(2,021)	0.926	(1,508)	0.055

Note: The participation rates are shown for each treatment group, with the resulting count of households enrolled in treatment for each group shown in parentheses. Initial participation reflects all those enrolled as of June 1st, 2012. End-line participation reflects all those enrolled as of October 1st, 2013. Participation is counted if the customer entered the program and did not opt out before the relevant date. Customers who moved away are removed from both the rate (the denominator) and the participation pool (the numerator) on the date they move, and so the participation rates do not reflect churn from customer relocation. The attrition rate is the percentage change between initial and end-line participation.

Table 3: Intent to Treat Effects

	Critical event hours		Non-event day peak hours	
	(1)	(2)	(3)	(4)
	Opt-in	Opt-out	Opt-in	Opt-out
Encouragement (CPP)	-0.130^{***}	-0.312^{***}	-0.027^{***}	-0.092^{***}
	(0.010)	(0.036)	(0.006)	(0.020)
Mean usage (kW)	2.44	2.44	1.76	1.76
Customers	55,028	46,684	55,028	46,684
Customer-hours	6,149,966	5,221,696	29,881,109	25,378,179
Encouragement (TOU)	-0.089^{***}	-0.133^{***}	-0.054^{***}	-0.101^{***}
	(0.008)	(0.019)	(0.006)	(0.013)
Mean usage (kW)	2.43	2.44	1.75	1.75
Customers	58,573	48,245	58,573	48,245
Customer-hours	6,543,842	5,394,991	31,794,095	26,219,448

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, standard errors clustered by customer.

Note: The dependent variable is hourly electricity usage in kilowatts. To estimate the critical event hour effects, data include 4-7pm during simulated CPP events in 2011 (weekdays with average dry bulb temperature above 96 degrees Fahrenheit) and 4-7pm during real CPP events in 2012-2013. To estimate the peak period non-event hour effects, data include 4-7pm on all non-holiday weekdays during the 2011, 2012 and 2013 summers, excluding simulated CPP event days in 2011 and excluding actual CPP event days in 2012 and 2013. Intent to treat effects are identified by comparing the opt-in and opt-out experimental groups to the control group. Intent to treat effects are estimated using ordinary least squares. All regressions include customer-hour of day and hour of sample fixed effects.

Figure 6: Event day intent to treat effects by hour

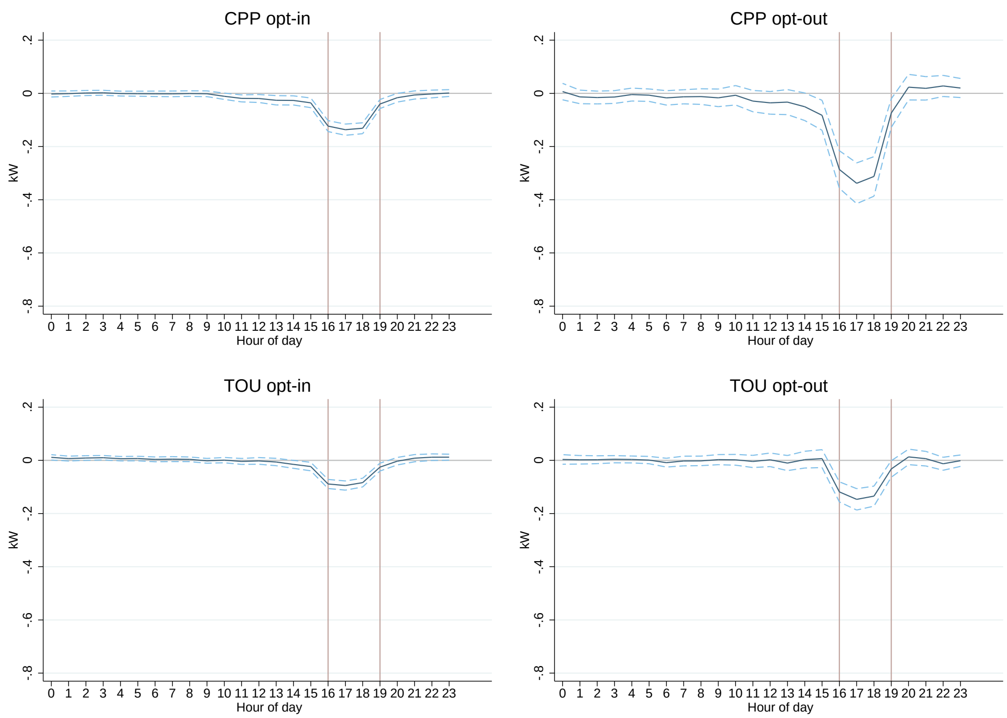

Note: The dependent variable is hourly electricity usage in kilowatts. To estimate the hourly effects on critical event days, data include all hours during simulated CPP events in 2011 (weekdays with average dry bulb temperature above 96°F) and all hours during real CPP events in 2012-2013. Intent to treat effects are identified by comparing the opt-in and opt-out experimental groups to the control group. Intent to treat effects are estimated using ordinary least squares. All regressions include customer-hour of day and hour of sample fixed effects. The hourly effects are shown for each treatment group, with the dashed lines indicating the 95 percent confidence interval of the estimates generating using standard errors clustered by customer. The vertical lines indicate the peak period.

Table 4: Average treatment effects by group

	Critical event hours			Non-event day peak hours		
	(1)	(2)	(3)	(4)	(5)	(6)
	Opt-in (AT)	Opt-out (AT+C)	Complacents (C)	Opt-in (AT)	Opt-out (AT+C)	Complacents (C)
Treat (CPP)	-0.665^{***}	-0.338^{***}	-0.250^{***}	-0.136^{***}	-0.099^{***}	-0.089^{**}
	(0.049)	(0.039)	(0.051)	(0.031)	(0.021)	(0.028)
Mean usage (kW)	2.44	2.44	2.40	1.76	1.76	1.75
Customers	55,028	46,684	10,036	55,028	46,684	10,036
Customer-hours	6,149,966	5,221,696	1,120,210	29,881,109	25,378,179	5,438,888
Treat (TOU)	-0.470^{***}	-0.140^{***}	-0.058^{*}	-0.284^{***}	-0.106^{***}	-0.062^{***}
	(0.042)	(0.020)	(0.026)	(0.029)	(0.014)	(0.018)
Mean usage (kW)	2.43	2.44	2.39	1.75	1.75	1.71
Customers	58,573	48,245	15,142	58,573	48,245	15,142
Customer-hours	6,543,842	5,394,991	1,687,381	31,794,095	26,219,448	8,193,143

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, standard errors clustered by customer.

Note: The dependent variable is hourly electricity usage in kilowatts. AT stands for always takers, AT+C stands for always takers and complacents. To estimate the critical event hour effects, data include 4-7pm during simulated CPP events in 2011 (weekdays with average dry bulb temperature above 96 degrees Fahrenheit) and 4-7pm during real CPP events in 2012-2013. To estimate the non-event day peak hour effects data include 4-7pm on all non-holiday weekdays during the 2011, 2012 and 2013 summers, excluding simulated CPP event days in 2011 and excluding actual CPP event days in 2012 and 2013. Complacent effect estimates are identified by comparing the opt-out experimental group to the opt-in experimental group, while the always taker effects are identified by comparing the opt-in experimental group to the control group. Treatment effects are estimated using two-stage least squares, with randomized encouragement to treatment used as an instrument for whether or not treatment was experienced. All regressions include customer-hour of day and hour of sample fixed effects.

Figure 7: Event day average treatment effects by hour

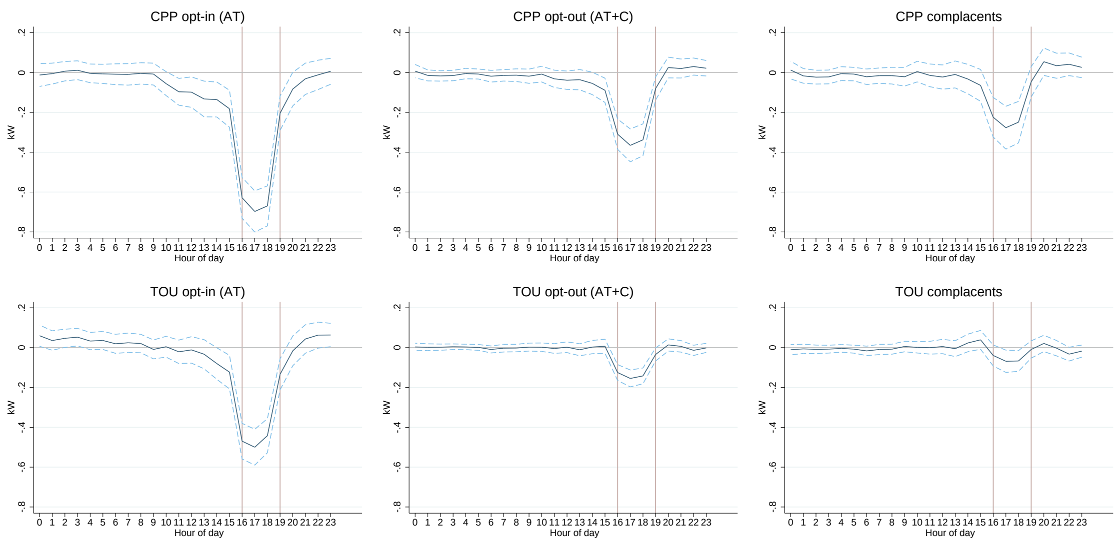

Note: The dependent variable is hourly electricity usage in kilowatts. To estimate the hourly effects on critical event days, data include all hours during simulated CPP events in 2011 (weekdays with average dry bulb temperature above 96°Fahrenheit) and all hours during real CPP events in 2012-2013. Opt-in and opt-out effects are identified by comparing the opt-in and opt-out experimental groups, respectively, to the control group. Complacent effect estimates are identified by comparing the opt-out experimental group to the opt-in experimental group. Treatment effects are estimated using two-stage least squares, with randomized encouragement to treatment used as an instrument for whether or not treatment was experienced. All regressions include customer-hour of day and hour of sample fixed effects. The hourly effects are shown for each group, with the dashed lines indicating the 95 percent confidence interval of the estimates generating using standard errors clustered by customer. The vertical lines indicate the peak period.

Table 5: Billing individual effects

	(1) Opt-in (AT)	(2) Opt-out (AT+C)	(3) Complacents (C)
CPP enrollment	-6.429^{**}	-4.174^{**}	-3.570^{*}
	(2.232)	(1.362)	(1.795)
Mean bill ($)	113.71	113.71	113.10
Customer-months	556,989	472,991	101,426
TOU enrollment	−2.947	-1.965^{*}	−1.723
	(2.091)	(0.825)	(1.105)
Mean bill ($)	113.55	113.59	112.32
Customer-months	592,641	488,681	152,768

Note: Dependent variable is monthly bill. Sample composed of summer months. AT stands for always takers, AT+C stands for always takers and complacents. Complacent effect estimates are identified by comparing the opt-out experimental group to the opt-in experimental group, while the always taker effects are identified by comparing the opt-in experimental group to the control group. Treatment effects are estimated using two-stage least squares, with randomized encouragement to treatment used as an instrument for whether or not treatment was experienced. All regressions include customer and month fixed effects.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, standard errors clustered by customer.

Table 6: Cost-effectiveness

	CPP		TOU	
	Opt-in	Opt-out	Opt-in	Opt-out
Marginal costs	219	63	206	46
Marginal benefits	355	153	188	100
Net Benefits	136	90	-18	54
Benefit/Cost Ratio	1.62	2.43	0.91	2.17

Note: Data from Nexant SMUD SmartPricing Options Pilot Evaluation Table 10-4. Marginal costs and benefits represent cost of adding a new customer given that the program already exists. Marginal costs include recruitment, equipment, and annual recurring costs. Marginal benefits include avoided capacity additions and avoided energy costs.

Table 7: Heterogeneity in Participation: Always Takers, Complacents, and Never Takers

	μ^{NT}	σ^{NT}	μ^{AT}	σ^{AT}	μ^{CM}	σ^{CM}	t/z^{NT-CM}	t/z^{AT-CM}	t/z^{AT-NT}
CPP: My Account	0.57	0.50	0.54	0.50	0.40	0.49	3.17	6.03	-0.54
CPP: My Account logins	20.86	42.69	8.98	22.26	5.94	0.55	0.92	3.05	-1.00
CPP: Projected savings	2.17	31.47	-1.71	19.01	-1.49	17.85	1.11	-0.25	-1.41
CPP: Low income	0.14	0.35	0.28	0.45	0.20	0.40	-1.51	3.82	3.11
TOU: My Account	0.46	0.50	0.53	0.50	0.39	0.49	2.23	9.09	2.03
TOU: My Account logins	5.90	5.70	8.32	25.49	5.87	6.90	0.02	2.32	0.21
TOU: Projected savings	-5.69	18.57	-6.46	20.49	-8.25	20.72	1.96	2.70	-0.41
TOU: Low income	0.12	0.33	0.29	0.46	0.19	0.39	-2.55	7.60	5.64

Note: NT indicates never takers, AT indicates always takers, and CM indicates complacents. μ are means, σ are standard deviations, t are the results from a two-tailed t-test for My Account logins and CPP savings, and z are the results from a two proportion z-test for My Account and Low income.

Table 8: Heterogeneity: My Account

	Critical event hours			Non-event day peak hours		
	(1)	(2)	(3)	(4)	(5)	(6)
	Opt-in (AT)	Opt-out (AT+C)	Complacents (C)	Opt-in (AT)	Opt-out (AT+C)	Complacents (C)
Treat (CPP)	-0.605^{***}	-0.236^{***}	-0.163^{**}	-0.133^{**}	-0.075^{**}	-0.063^{*}
	(0.076)	(0.043)	(0.054)	(0.049)	(0.025)	(0.031)
Treat $\times$ My Account	-0.113	-0.245^{**}	-0.227^{*}	-0.006	-0.059	-0.069
	(0.099)	(0.083)	(0.114)	(0.063)	(0.045)	(0.061)
Customers	55,028	46,684	10,036	55,028	46,684	10,036
Customer-hours	6,149,966	5,221,696	1,120,210	29,881,109	25,378,179	5,438,888
Treat (TOU)	-0.342^{***}	-0.082^{***}	-0.033	-0.197^{***}	-0.066^{***}	-0.041^{*}
	(0.066)	(0.023)	(0.029)	(0.045)	(0.016)	(0.021)
Treat $\times$ My Account	-0.244^{**}	-0.149^{***}	-0.072	-0.166^{**}	-0.101^{***}	-0.057
	(0.085)	(0.042)	(0.057)	(0.058)	(0.029)	(0.040)
Customers	58,573	48,245	15,142	58,573	48,245	15,142
Customer-hours	6,543,842	5,394,991	1,687,381	31,794,095	26,219,448	8,193,143

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, standard errors clustered by customer.

Note: Dependent variable is hourly energy usage in kW. Coefficient estimates produced using TSLS. For columns 1, 2, 4, and 5, enrollment in program and its interaction are instrumented with encouragement group and its interaction and sample includes encouragement and control groups. For columns 3 and 6, enrollment in program is instrumented with enrollment into opt-out group and sample includes only opt-in and opt-out groups. Event hours include simulated CPP events in 2011 and real CPP events in 2012-2013. Non-event hours include all peak hours except CPP event hours. All models include customer and hour of sample fixed effects, plus an interaction between the post-treatment period and dummies for My Account. Standard errors clustered by customer in parentheses.

Table 9: Heterogeneity: Low income

	Critical event hours			Non-event day peak hours		
	(1)	(2)	(3)	(4)	(5)	(6)
	Opt-in (AT)	Opt-out (AT+C)	Complacents (C)	Opt-in (AT)	Opt-out (AT+C)	Complacents (C)
Treat (CPP)	-0.818^{***}	-0.378^{***}	-0.275^{***}	-0.173^{***}	-0.095^{***}	-0.077^{*}
	(0.062)	(0.046)	(0.058)	(0.039)	(0.025)	(0.031)
Treat × Low income	0.527^{***}	0.176^{*}	0.109	0.126^{*}	−0.017	−0.062
	(0.093)	(0.085)	(0.119)	(0.061)	(0.050)	(0.070)
Customers	55,028	46,684	10,036	55,028	46,684	10,036
Customer-hours	6,149,966	5,221,696	1,120,210	29,881,109	25,378,179	5,438,888
Treat (TOU)	-0.537^{***}	-0.154^{***}	-0.070^{*}	-0.326^{***}	-0.113^{***}	-0.066^{**}
	(0.053)	(0.023)	(0.030)	(0.037)	(0.016)	(0.021)
Treat × Low income	0.228^{**}	0.063	0.059	0.146^{*}	0.032	0.023
	(0.082)	(0.041)	(0.058)	(0.057)	(0.030)	(0.042)
Customers	58,573	48,245	15,142	58,573	48,245	15,142
Customer-hours	6,543,842	5,394,991	1,687,381	31,794,095	26,219,448	8,193,143

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, standard errors clustered by customer.

Note: Dependent variable is hourly energy usage in kW. Coefficient estimates produced using TSLS. For columns 1, 2, 4, and 5, enrollment in program and its interaction are instrumented with encouragement group and its interaction and sample includes encouragement group and control group. For columns 3 and 6, enrollment in program is instrumented with enrollment into opt-out group and sample includes only opt-in and opt-out groups. Event hours include simulated CPP events in 2011 and real CPP events in 2012-2013. Non-event hours include all peak hours except CPP event hours. All models include customer and hour of sample fixed effects, plus an interaction between the post-treatment period and a dummy for low income. Standard errors clustered by customer in parentheses.

Table 10: Heterogeneity: Program year

	Critical event hours			Non-event day peak hours		
	(1)	(2)	(3)	(4)	(5)	(6)
	Opt-in (AT)	Opt-out (AT+C)	Complacents (C)	Opt-in (AT)	Opt-out (AT+C)	Complacents (C)
Treat (CPP)	-0.719^{***}	-0.304^{***}	-0.192^{***}	-0.152^{***}	-0.078^{***}	-0.058^{*}
	(0.052)	(0.042)	(0.055)	(0.031)	(0.022)	(0.029)
Treat $\times$ Year 2	0.120^{*}	-0.073^{*}	-0.127^{**}	0.037	-0.051^{*}	-0.074^{*}
	(0.054)	(0.037)	(0.049)	(0.035)	(0.022)	(0.030)
Customers	55,028	46,684	10,036	55,028	46,684	10,036
Customer-hours	6,149,966	5,221,696	1,120,210	29,881,109	25,378,179	5,438,888
Treat (TOU)	-0.534^{***}	-0.160^{***}	-0.065^{*}	-0.307^{***}	-0.114^{***}	-0.065^{***}
	(0.044)	(0.021)	(0.027)	(0.029)	(0.014)	(0.018)
Treat $\times$ Year 2	0.143^{**}	0.043^{*}	0.016	0.056	0.018	0.007
	(0.049)	(0.020)	(0.027)	(0.032)	(0.013)	(0.017)
Customers	58,573	48,245	15,142	58,573	48,245	15,142
Customer-hours	6,543,842	5,394,991	1,687,381	31,794,095	26,219,448	8,193,143

Note: Coefficient estimates produced using TSLS, enrollment in program and its interaction with a dummy for 2013 are instrumented with encouragement into program and its interaction with with a 2013 dummy. Event hours include simulated CPP events in 2011 and real CPP events in 2012-2013. Non-event hours include all peak hours except CPP event hours. All models include customer and hour of sample fixed effects. Standard errors clustered by customer in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, standard errors clustered by customer.