DEFAULT EFFECTS AND FOLLOW-ON BEHAVIOR: EVIDENCE FROM AN ELECTRICITY PRICING PROGRAM

ONLINE APPENDIX

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

June 6, 2017

*We received many helpful comments from seminar participants at Arizona State University, Cornell, Toulouse School of Economics, UC Berkeley and University of Oxford. The authors gratefully acknowledge contributions from and discussions with Hunt Allcott, Stefano Dellavigna, Steven George, Nick Kuminoff, Brigitte Madrian, 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 Load Shape Balance across Treatment Groups

Table 1 in the main text discusses balance in covariates between control and treatment groups. Because we analyze consumption across hours of the day, we are also concerned about balance in hourly consumption profiles. Figure A1 plots each treatment group's hourly electricity consumption overlaid with control group consumption, obtained from a regression of electricity consumption on a set of indicator variables for each hour. The left 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 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 for customers of 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 Figure A2 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.

2 Alternative Specifications

Tables A1-A3 report results similar to those in Tables 3, 4 and 8 in the text using the log of hourly consumption as the dependent variable. The results are very consistent across specifications: the ITT estimate is about twice as large in the CPP opt-out

Table A6 reports specifications similar to those in Table 8, where the interaction term is for households that were structural winners. The coefficients on the interaction terms are all either negative or insignificant, suggesting that structural winners if anything reacted more to changes in prices.

3 Assumptions Underlying the LATE Estimates

We take a standard approach to identifying the local average treatment effects using our two encouragement instruments (i.e., the opt-in offer and the opt-out offer, respectively). This section explains how we leverage this research design to estimate local average treatment effects in different sub-groups of our study sample.

Let $D_i = 1$ if the individual participates in the dynamic pricing program. Let $D_i = 0$ if the individual remains in the standard pricing regime. Let $Z_i = 1$ if the individual was assigned to the opt-in encouragement treatment, let $Z_i = 2$ if the individual was assigned to the opt-out; otherwise $Z_i = 0$.

Conceptually, we define four sub-populations:

- Never takers (NT): Do not opt in if $Z_i = 1$. Opt out if $Z_i = 2$.
- Complacents (C): Do not opt in if $Z_i = 1$. Do not opt out if $Z_i = 2$.
- Always takers (AT): Opt in if $Z_i = 1$. Do not opt out if $Z_i = 2$.
- Defiers (D): Opt in if $Z_i = 1$. Opt out if $Z_i = 2$.

To identify the LATE separately for the opt-in and opt-out interventions, respectively, we make the following assumptions:

- **Unconfoundedness**: We assume that the assignment of the encouragement intervention Z_i is independent of/orthogonal to other determinants of energy consumption. This assumption is satisfied (in expectation) by our experimental research design.
- **Stable unit treatment values:** Electricity consumption at household i is affected by the participation status of household i but not the participation decisions of other households.
- **Exclusion restriction:** Our encouragement intervention affects energy consumption only indirectly through the effect on program participation.
- **Monotonicity**: Our encouragement intervention weakly increases (and never decreases) the likelihood of participation in the time varying rate. This implies that there are no defiers.

Let π^{NT}, π^{C}, and π^{AT},denote the population proportions of never takers, complacents, and always takers, respectively. Let $Y_i(D_i = 1)$ and $Y_i(D_i = 0)$ define the potential electricity consumption outcomes associated with consumer i conditioning on participation in the dynamic pricing program. Given the exclusion restriction, these potential outcomes need not condition on the encouragement intervention.

With the opt-in design, the average electricity consumption among households assigned to the control group ($Z_i = 0$) is:

$$E[Y_i|Z_i = 0] = \pi^{NT}E[Y_i(0)|NT] + \pi^{C}E[Y_i(0)|C] + \pi^{AT}E[Y_i(0)|AT].$$

The average consumption among households assigned to the opt-in encouragement:

$$E[Y_i|Z_i = 1] = \pi^{NT}E[Y_i(0)|NT] + \pi^{C}E[Y_i(0)|C] + \pi^{AT}E[Y_i(1)|AT].$$

Mechanically, it is straightforward to construct an estimate of the effect of the pricing program on average consumption among always takers by taking the difference in these two expectations and dividing by π^{AT}

$$LATE^{AT} = \frac{E[Y_i|Z_i = 0] - E[Y_i|Z_i = 1]}{\pi^{AT}} = E[Y_i(0)|AT] - E[Y_i(1)|AT],$$

where $\pi^{AT}$is estimated by the share of participants in the encouraged group. We take a similar approach using the opt-out design to construct an estimate of the local average treatment effect in the combined AT and C groups:

To isolate the average treatment effect in the complacent population, we compare outcomes across the two groups assigned to $Z_i = 1$ and $Z_i = 2$, respectively. Taking the difference across these two groups and dividing by π^{C}yields:

$$LATE^{C} = \frac{E[Y_i|Z_i = 1] - E[Y_i|Z_i = 2]}{\pi^{C}} = E[Y_i(0)|C] - E[Y_i(1)|C].$$

The estimate of π^{C} is obtained by taking the difference in program participation across the opt-in and opt-out treatments.

If our encouragement intervention affects electricity consumption directly, this will violate the exclusion restriction and confound our ability to identify these local average treatment effects. The exclusion restriction would be violated, for example, if the act of encouraging customers to opt into the dynamic rate plan directly impacts consumption by increasing the salience of energy use. In this scenario, potential outcomes should be represented by $Y_i(D_i, Z_i)$. Taking the opt-in design as an example, the local average treatment effect among takers is now more accurately estimated as:

$$LATE^{AT} = \frac{E[Y_i|Z_i = 0] - E[Y_i|Z_i = 1]}{\pi^{AT}} - \frac{\Delta_C}{\pi^{AT}} - \frac{\Delta_{NT}}{\pi^{AT}},$$

where $\Delta_C = E[Y_i(0, 0)|C] - E[Y_i(0, 1)|C]$ and $\Delta_{NT} = E[Y_i(0, 0)|NT] - E[Y_i(0, 1)|NT]$. If these encouragement-induced changes in electricity consumption among non-participants are not equal to zero, they will bias our LATE estimates.

We cannot estimate these Δ terms directly. We can, however, compare consumption patterns at households that did not participate in the dynamic pricing program across encouraged and unencouraged groups. These differences are difficult to interpret as they compare electricity consumption across different subsets of the non-participant population. But they do provide some sense of how large the bias from violating the exclusion restriction might be.

We re-estimate Equation (1) using only those households who did not participate in dynamic pricing. Table A4 summarizes these comparisons. For the opt-in experiments, these results represent the difference in average consumption among households assigned to the control group and the average consumption among all non-participants who received the opt-in offer (i.e., complacents and never takers). For the opt-out experiments, we compare consumption across all households assigned to the control group and the never-takers in the encouraged group.

Some of these differences are statistically different from zero. For example, we estimate a significant difference of -0.025 across encouraged and unencouraged non-participants in the opt-in TOU experiment. It seems likely that some of this difference is driven by differences in composition- we are comparing consumption across all households in the control group with consumption of never-takers and complacents in the encouraged group. However, if we interpret this difference as entirely caused by the opt-out intervention, this would imply that our local average treatment effect overstates the true effect by $\frac{0.025}{0.19} = 0.13$

4 Modeling Attrition out of the Program

As reported in Table 2, approximately 6-7% of the customers on the dynamic pricing programs opted to leave the program at some point during the two-year study. Figure A3 reports Kaplan-Meier survival estimates for each of the four treatment groups. The vertical orange lines indicate critical event days and the vertical blue line indicates the date on which the second summer reminder letter was sent out to all study participants letting them know that the rate would start again. We see some attrition from all four groups before the event days started, slightly more attrition from the CPP groups throughout the first summer, and then a relatively big drop after the reminder.

To gain more insight into attrition timing, we model the propensity for customers to leave the dynamic pricing programs once enrolled using an accelerated failure time (AFT) model. We elected to use an AFT model instead of a proportional hazard model as it better accommodates the impact of specific events, such as the critical peak pricing days. In the AFT, the exponential of the estimated coefficient on a variable indicates the "acceleration factor" in the influence on that variable on the survival time. The results of the hazard analysis are presented in Table A5.

One might expect that customers who actively opted in to the new rates would be less likely to later change their minds and opt-out. In fact, the attrition rates are similar across opt-in and opt-out for the TOU rates, and the opt-in customers were even quicker to get out of the rates than the opt-out in the CPP case. In particular, the CPP opt-out group had a survival time (i.e., time remaining in treatment before dropping out) that was 40 percent higher than (calculated as exp(0.339)) the opt-in group, which is the omitted category, although the difference is not stastically significant. This could reflect the fact that opt-in customers are self-selected to have low switching costs.

As for other customer-level impacts, there is some evidence that low-income customers were less likely to drop out of the study quickly relative to non-EAPR customers. Structural winners tended to remain in the study longer than those that were not structural winners, although the difference is not stasticially significant. Customers with “Your Account” were no more likely to drop out of the study more quickly.

In the case of effects over time, the second summer reminder had a strong effect that accelerated the rate of drop-outs (reduced the survival time) across all the treatment groups. The occurrence of CPP event days enters the model in the following way: there is an indicator variable included for CPP event days for the two CPP treatment groups (“CPP event date”). In addition, the variable “CPP event date count in each summer” is a variable that increases by one each occurrence of a CPP event date within each summer. So, it is equal to 1 on the first occurrence of a CPP event both in the first and second summer, and is equal to 2 for the second occurrence of a CPP event within each summer, etc. The results for CPP event days indicate that for the opt-in CPP treatment group, the experience of CPP event days reduced the survival time in the study by slightly less than the reminder. However, this effect was attenuated over the course of more CPP events within each summer. For the CPP opt-out treatment group, however, the effect of experiencing a CPP event at all is close to zero (the sum of the coefficient and the interaction), and the effect of CPP events appears to increases the rate of drop-outs slightly over multiple events. Finally, we tested whether there was any disproportional additional effect of experiencing a string of consecutive (two or three in a row) events. There does not appear to be a discernible effect of experiencing multiple CPP event beyond the baseline CPP event effect for either CPP treatment group.

The bottom rows of the table list the number of participants and the number of dropouts for each treatment group. As emphasized in the main text, we find the attrition results suggestive but are hesitant to put too much emphasis on them given the relatively small number of dropouts.

5 Cost-Benefit Analysis

This section describes the cost-benefit calculations reported in Section 6. Many of the assumptions used in our calculations are summarized in **?**, a consulting report that provided, among other things, a cost-benefit calculation of several components of the SMUD program. Other assumptions are based on personal communications with SMUD employees and their consultants.

5.1 Benefits

At a high level, reduced demand during CPP and TOU peak hours avoids two types of expenses – the energy associated with generating electricity during these hours and the expected cost of adding new capacity to meet peak demand, where the expectation is taken over the probability that demand in a particular hour would drive capacity expansion decisions. The components of the benefit calculations are summarized in Figure A4.

Consider the first row, reflecting capacity benefits. The first box represents assumptions on the cost of adding a new peaking plant. Our calculations are based on proprietary information provided by SMUD and summarized in **?**. As reported by Potter et al., the costs "range from roughly $50 to $80/kW-year in the first few forecast years and increase to around $125/kW-year by the end of the forecast period" (p. 112, **?**). These costs are slightly lower than other estimates of generation capacity costs from Northern California. For example, the California Public Utilities Commission (CPUC) publishes capacity values for assessing the cost effectiveness of demand response programs. The "Generation Capacity Values" range from $174 to $209/kW-year for 2012-14, considerably higher than the numbers SMUD uses. Notably, SMUD did not include the capacity costs associated with the transmission and distribution system. According to the CPUC model, those can account for approximately 25% of the capacity benefits of a peak demand reduction program, so SMUD's decision likely understates the benefits of the program. The values represented by the second box, "# of Enrolled Customers on Time-Variant Pricing Plans," reflect participation rates, summarized in Table 2, multiplied by 600,000, an estimate of the number of customers SMUD will have in 2018. We assumed a customer attrition rate of approximately 7% per year. As shown in Table 2, attrition rates over the 16 months the program operated were approximately 5.5 to 7 percent. We converted these to annual attrition rates and then added 2% to account for customers moving out of SMUD's service territory, assuming that customers who moved within the service territory would remain on the rate.

The values represented by the third box, "Average Reduction by Enrolled Customer by Hour and Month" are the LATE coefficients summarized in Table 4. **?** estimated separate LATE effects for each hour of the program and provide suggestive evidence that customers reduce more when day are hotter. Hotter days also have higher "Capacity Risk Allocation" values, so this likely explains why the numbers in **?** are slightly higher than ours.

The "Capacity Risk Allocation by Hour and Month" figures are based on proprietary values provided by SMUD. They are based on a simulation model which estimates the probability that demand exceeds supply on SMUD's system across any of the hours on representative weekend days and weekdays for each month of the year (called the "loss of load probability.") These values are then normalized to sum to one across all hours of the year. We use the sum of the normalized values in hours targeted by the CPP and TOU rates. Finally, following **?**, we assume a 7.1% nominal discount rate and a 4.5% real discount rate.

In row two, reflecting the calculations to arrive at the avoided energy benefits, two of the three boxes are the same as in row 1. "Avoided Energy Costs by Hour and Month" are based on an avoided cost of $0.04 per kilowatt-hour, which is approximately the midpoint of the range provided in **?**. As mentioned in the text, these numbers do not reflect the environmental externalities associated with electricity generation.

5.2 Costs

Table 6 summarizes one-time fixed costs, one-time variable costs and recurring fixed and variable costs. One-time fixed costs do not vary with enrollment and include items such as IT costs to adjust the billing system and initial market research costs. One-time variable costs primarily include the customer acquisition costs, including the in-home devices offered to customers as part of the recruitment. Note that **?** model opt-in programs that do not include outbound calls to enroll customers, while we include the costs of the calls, as well as the customers recruited through them. Our objectives are different from theirs, as they were modeling a hypothetical program that SMUD might run in the future, while we are modeling the program that was actually run. Recurring annual fixed and variable costs include personnel costs required to administer the program and costs associate with customer support and equipment monitoring. They go down slightly over time with attrition from the program.

References