

Lawrence Berkeley National Laboratory

LBL Publications

Title

Default Effects and Follow-On Behaviour: Evidence from An Electricity Pricing Program

Permalink

<https://escholarship.org/uc/item/3wv2r6t2>

Journal

Review of Economic Studies, 88(6)

ISSN

0034-6527

Authors

Fowlie, M

Wolfram, C

Baylis, P

et al.

Publication Date

2021-11-01

DOI

10.1093/restud/rdab018

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial-NoDerivatives License, available at <https://creativecommons.org/licenses/by-nc-nd/4.0/>

Peer reviewed



Electricity Markets & Policy
Energy Analysis & Environmental Impacts Division
Lawrence Berkeley National Laboratory

Default Effects and Follow-On Behaviour: Evidence From an Electricity Pricing Program

Meredith Fowlie, Catherine Wolfram, Patrick Baylis, C Anna Spurlock,
Annika Todd-Blick, Peter Cappers

April 2021

This is a pre-print version of a journal article published in *The Review of Economic Studies*. DOI: <https://doi.org/10.1093/restud/rdab018>



This work was supported by the U.S. Department of Energy's Office of Electricity Delivery and Energy Reliability under Lawrence Berkeley National Laboratory Contract No. DE-AC02-05CH11231.

DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, or The Regents of the University of California.

Ernest Orlando Lawrence Berkeley National Laboratory is an equal opportunity employer.

COPYRIGHT NOTICE

This manuscript has been authored by an author at Lawrence Berkeley National Laboratory under Contract No. DE-AC02-05CH11231 with the U.S. Department of Energy. The U.S. Government retains, and the publisher, by accepting the article for publication, acknowledges, that the U.S. Government retains a non-exclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for U.S. Government purposes.



Energy Institute WP 280R

**Default Effects and Follow-On Behavior: Evidence
From An Electricity Pricing Program**

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

August 2020

**Revised version published in
Review of Economic Studies, March 2018**

Energy Institute at Haas working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to review by any editorial board. The Energy Institute acknowledges the generous support it has received from the organizations and individuals listed at <https://haas.berkeley.edu/energy-institute/about/funders/>.

© 2020 by Meredith Fowlie, Catherine Wolfram, C. Anna Spurlock, Annika Todd, Patrick Baylis, and Peter Cappers. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit is given to the source.

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

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

August 21, 2020

Abstract

We study default effects in the context of a residential electricity-pricing program. We analyze a large-scale randomized controlled trial in which one treatment group was given the option to opt-in to time-varying pricing while another was defaulted into the program but allowed to opt-out. We provide dramatic evidence of a default effect on program participation, consistent with previous research. A novel feature of our study is that we also observe how the default manipulation impacts customers' subsequent electricity consumption. Passive consumers who did not opt-out but would not have opted in — comprising more than 70 percent of the sample — nonetheless reduce consumption in response to higher prices. Observation of this follow-on behavior enables us to assess competing explanations for the default effect. We draw conclusions about the likely welfare effects of defaulting customers onto time-varying pricing.

*Fowlie: UC Berkeley and NBER, fowlie@berkeley.edu. Wolfram: UC Berkeley and NBER, cwolfram@berkeley.edu. Spurlock: Lawrence Berkeley National Laboratory, caspurlock@lbl.gov. Todd: Lawrence Berkeley National Laboratory, ATodd@lbl.gov. Baylis: University of British Columbia; patrick.baylis@ubc.ca. Cappers: Lawrence Berkeley National Laboratory, pacappers@lbl.gov. 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, Leslie Martin, 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 Office of Electricity Delivery and Energy Reliability, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof. Corresponding author: Catherine Wolfram, Haas School of Business, UC Berkeley, Berkeley, CA 94720-1900, (510)-642-2588.

1 Introduction

When confronted by a choice with a default option, decision-makers are often predisposed to accept the default. Prior work in psychology and economics has documented this “default effect” for a range of decisions that would seem to merit deliberate choices, including retirement plans (Madrian and Shea 2001), health insurance (Handel 2013), and organ donations (Johnson and Goldstein 2003). This phenomenon is of general interest because it provides businesses and public policy makers with a relatively easy and non-intrusive way to influence choices.

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 harder to discern. One reason is that these impacts are a function of both the initial choice subject to the default manipulation and any “follow-on” behaviors that can depend on the initial choice. For example, consumers who are defaulted onto a health insurance plan with high co-payments may invest less in preventative health compared to those who actively chose such a plan. Similarly, consumers who are forced to actively choose particular privacy settings on a social media platform may subsequently share less information than consumers who are defaulted into a data-sharing regime. Given that many default manipulations aim to induce changes in some form of follow-on behavior, it is important to account for both direct and follow-on impacts of default manipulations on economic outcomes.

This study analyzes the use of a default manipulation in a new choice setting: time-varying electricity pricing. Electricity customers were randomized into two different types of treatment groups. In one type, customers were invited to opt in to a new time-varying pricing plan. In another set of treatment groups, customers were informed that they would be defaulted onto the new pricing programs unless they opted out. The field experiment was run by the Sacramento Municipal Utility District (SMUD) in 2011-2013. We observe both the initial pricing plan choice and follow-on electricity use. We are able to isolate impacts on the follow-on behavior of those who actively opted in (referred to here as “active joiners”), from those who enrolled in the new pricing structure because of the default (referred to here as “passive consumers”).

It is important to understand how default manipulations can affect consumers' response to time-varying electricity pricing because a significant increase in customer participation could generate substantive efficiency gains. Benefits include lower electricity system operating costs, lower renewable energy integration costs, and a more resilient electricity grid. Importantly, the scale of these benefits increase with the number of customers confronted by, and responding to, time-varying prices, and are therefore critically contingent on both the enrollment rate and follow-on behavior once enrolled.

The vast majority (over 94 percent in 2018) of U.S. residential customers face time-invariant prices for electricity (EIA 2018). Recent investments in smart grid infrastructure, including smart meters, make it technologically feasible to enroll many customers in time-varying pricing programs. As of 2019, almost 100 million smart meters had been deployed to over half of US households (Cooper and Shuster 2019).¹ The large discrepancy between the share of customers for whom it is technologically feasible to face time-varying pricing and the share who actually do suggests that proactive approaches to increasing active participation in time-varying pricing will be required to fully leverage its potential.

We show that making time-varying pricing the default choice can significantly increase participation — over 90 percent of the customers stayed with time-varying pricing when defaulted onto it. In contrast, only 20 percent actively opted in. In this setting, the economic importance of the default effect depends critically on whether the households susceptible to the default effect, i.e., the passive consumers who neither opt in nor opt out, follow on to actively reduce their peak consumption in response to the time-varying electricity prices. If passive customers do not adjust consumption, then there is little point in defaulting them into this pricing regime. We obtain detailed measurements of electricity consumption in the periods prior to and following the experimental intervention. We show that passive customers, who comprise more than 70 percent of the sample, do reduce consumption when prices increase during peak times. Although the average demand response among passive customers is approximately half as large as the average

1. The deployment of smart grid technology was dramatically accelerated under the American Recovery and Reinvestment Act of 2009.

response among customers who actively opted in, higher participation rates in the opt-out group mean that the average effect of the opt-out offer on peak demand is significantly larger than the average effect of the opt-in offer.

These findings notwithstanding, policy makers may be reluctant to authorize the use of default provisions until they understand the consumer welfare implications. For example, if the default effect is driven by high switching costs, some customers could be considerably worse off under a new pricing plan. Because alternative explanations for the default effect can have very different welfare implications, it is important to investigate the underlying mechanisms. We assess the extent to which alternative explanations for the default effect are consistent with the participation choices and detailed electricity consumption patterns we observe. We document a striking lack of correlation between households' participation choices and the savings they stand to gain from participation, even in the presence of a program enrollment deadline. This is hard to fully explain if switching costs, discounting, or present-biased preferences drive the default effect. An alternative model in which consumers are inattentive to the participation decision performs much better in explaining observed choices. We offer further evidence to show that inattention is plausibly rational in this setting.

Previous work has analyzed follow-on behavior (though not explicitly labeled it as such) in the context of household savings, where individuals were originally subject to a default for their retirement savings plan (Chetty et al. 2014; Choukhmane 2018). Our paper adds to this literature by exploring a situation where the variation in the initial default is randomly assigned, and the subsequent impacts on follow-on behavior can be cleanly identified.

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. Section 6 investigates alternative explanations for the default effect in light of the empirical evidence we document. In section 7, we summarize the implications of our findings for consumer welfare and present calculations on the net benefits of the time-varying pricing programs from

the utility’s perspective. Section 8 concludes.

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

A rich literature considers default effects in a range of settings, including participation in retirement savings plans (Samuelson and Zeckhauser 1988; 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, Jun, and MacInnis 2000), and email marketing (Johnson, Bellman, and Lohse 2002). Thaler and Sunstein (2009) motivate the main thesis of their book *Nudge* with an introductory example on the default effect, suggesting that, “[a]s we will show, setting default options, and other similar seemingly trivial menu-changing strategies, can have huge effects on outcomes, from increasing savings to improving health care to providing organs for lifesaving transplant operations” (p. 8).

In many of the contexts where default provisions are used to influence choice outcomes, follow-on behavior plays a critical role in determining economic impacts. We make a distinction between two types of follow-on behavior. First, individuals may choose to subsequently modify the option they chose by default. For example, a consumer who accepts a particular health insurance plan as a default option might subsequently adjust this choice by changing to a different plan. Second, there may be important choices or actions that are contingent on — but distinct from — the initial choice. Building on the health insurance plan example, participating in a plan with a high co-pay could impact subsequent choices about whether or not to go to the doctor, lifestyle choices that can affect health outcomes, or choice of medical procedures.

To date, the literature on default effects has emphasized the initial choice and placed less emphasis on subsequent decisions that can be significantly — albeit indirectly — impacted by default manipulations. Analyses of retirement savings decisions have considered the first type of follow-on behavior: modifications to the original choice. For example, Brown, Farrell, and

Weisbenner (2016) present survey evidence suggesting that employees who were irreversibly defaulted into a defined benefit retirement plan are more likely to later express a desire to enroll in a different plan. Sitzia, Zheng, and Zizzo (2015) consider the effects of defaults using a choice experiment on electricity tariffs, but because these decisions are hypothetical it is not possible to observe follow-on behavior. Other work includes information about follow-on choices, but does not model the impact of the default setting on those choices. For example, Ketcham, Kuminoff, and Powers (2016) include information about Medicaid recipients' prescription drug spending in their welfare calculation, but do not model how plan choice impacts drug expenditures. 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 decisions.

Our paper is most closely related to Chetty et al. (2014) and Choukhmane (2018). Chetty et al. (2014) analyze Danish policies to encourage retirement savings and differentiate “active savers,” who respond to tax incentives and/or mandatory savings policies by adjusting their investments, and “passive savers,” who do not. Choukhmane (2018) investigates retirement savings behavior in the U.S. and the U.K. and finds that individuals who are not enrolled by default make future adjustments to retirement savings that eventually bring them in line with those who were enrolled by default. Our choice setting is similar in that consumers must first navigate, either actively or passively, an initial participation offer which will then impact follow-on choices. One difference is that the follow-on behaviors and outcomes analyzed in the retirement savings literature are related by the budget constraint: whether or not passive savers respond explicitly, their spending and/or saving behavior must adjust in some way to the change to retirement savings induced by the default. By contrast, because electricity accounts for a small share of total consumption, and because the pricing plans we study affect a small share of electricity consumption, our passive consumers could have been defaulted onto the time-varying pricing and then completely ignored it. That they do not exposes a difference in the two decision settings.² Nonetheless, we

2. There are other relevant differences between Chetty et al. (2014) and Choukhmane (2018) and our work. Neither paper addresses the difference between opting out of a mandatory plan, which would be analogous to our active

also note a remarkable similarity between these two very different settings of retirement savings and electricity consumption: passive consumers comprise roughly 60-70 percent of the population. Understanding how passive consumers who are nudged onto a program by default respond to the program once enrolled is relevant to a range of settings outside of electricity consumption.³

Our empirical results on both the initial choice and follow-on behavior also shed light on the underlying mechanisms that can give rise to default effects. Recent papers have investigated the welfare effects of nudges in a variety of contexts, including retirement savings plan default provisions (Carroll et al. 2009; Bernheim, Fradkin, and Popov 2015), health insurance plan choices (Handel 2013; Handel and Kolstad 2015; Ketcham, Kuminoff, and Powers 2016), and home energy conservation reports (Allcott and Kessler 2019). These papers augment the more standard utility maximization framework to accommodate features of consumer behavior (such as inattention) that could rationalize a default effect. Bernheim, Fradkin, and Popov (2015) and Blumenstock, Callen, and Ghani (2018) go one step further and mediate between several different explanations for the default effect. Our work extends this line of inquiry to a context where inattention to one choice leads to economically significant efficiency gains in a subsequent set of consumer choices.

3 Empirical Setting and Experimental Design

Economists have noted for some time that efficient pricing of electricity should reflect changing electricity market conditions (e.g., Boiteux 1964b, 1964a). 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 week during our study. The red line depicts hourly electricity demand, which

refusers group, and actively taking advantage of a non-mandatory plan, analogous to our active joiners group. Additionally, both papers take advantage of quasi-experimental variation across similar but not identical settings, while our experimental setting allows to randomly allocate customers into plans that are identical except for their enrollment mechanism.

3. In addition to health insurance and privacy on social media platforms, consider, for example, on line marketing: customers are often passively enrolled into e-mail campaigns, but the degree to which they respond to subsequent appeals is of central interest to the marketer. Other examples, such as employee incentive programs for fitness or volunteering, or car purchasing and post-purchase driving behavior, indicate that this type of two-stage decision-making is widespread.

cycles predictably over the course of a day, varying by a factor of 1.5 on some days to almost 3 on others 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 in 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 94 percent in 2018) of U.S. residential customers pay time-invariant prices for electricity (EIA 2018). 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-varying rates, utilities could construct 44 percent fewer peaking plants.

This suggests that these inefficiencies can be mitigated – or eliminated – with the introduction of time-varying retail electricity pricing. Residential customers have an important role to play in electricity demand response, particularly in areas of the country where peak residential demand (driven by air conditioning in many parts of the U.S.) coincides with the system peak. When residential customers have been exposed to time-varying prices, prior analyses suggest they are willing and able to adjust consumption in response (see, for example, EPRI 2012).⁴

To reap benefits from time-varying pricing, however, utilities need to enroll more customers in time-varying pricing programs and these customers need to respond to the prices. In what

4. In a 2012 meta-analysis, authors identified what they deemed to be the best seven U.S. residential pricing studies up to that time (EPRI 2012). These studies document peak demand response to time-varying pricing in the range of 13-33%, depending on the existence of automated control technology (e.g., programmable communicating thermostat). These estimates imply an elasticity of substitution in the range of 0.07 - 0.24 and an own-price elasticity in the range of -0.07 - -0.3. Note that the experimental nature of our study allows us to assess many dimensions of customers' responses to time-varying pricing, including spillovers within and across days. Some previous evaluations of time-varying pricing have relied on within-customers comparisons, which assume there are no spillovers of this sort.

follows, we describe a large-scale field experiment designed to evaluate a novel approach to increasing participation among residential electricity customers.⁵

3.1 The Experiment

The experiment we analyze was implemented as part of the Smart Grid Investment Grant (SGIG) program, which received \$3.4 billion in funds from the American Recovery and Reinvestment Act of 2009. The goal of this program was to invest in the expansion of the smart grid in the U.S., and thereby create jobs and accelerate the modernization of the nation’s electric system (Department of Energy 2012). One of the objectives articulated in the Funding Opportunity Announcement (DE-FOA-0000058) under the heading of Consumer Behavior Studies (CBS) was to document the impacts and benefits of time-varying rate programs and associated enabling control and information technologies.

The Sacramento Municipal Utility District (SMUD), a municipal utility that serves approximately 530,000 residential households in and around Sacramento, California, implemented one of the 11 consumer behavior studies that were funded under the SGIG program.⁶ They were awarded a \$127 million grant overall, which comprised part of a \$308 million smart grid project. SMUD viewed the opportunity to study the impact of time-varying rates within their own service territory as a major benefit to participating in the program (Jimenez, Potter, and George 2013). SMUD had some demand response programs in place prior to the SGIG program (e.g., an air conditioner direct control program and some rates that varied by time-of-day), but these programs had not been broadly emphasized or marketed for a long time. Historic adoption of their “legacy” Time-

5. A much smaller-scale experiment was conducted in Los Alamos. Results of this study are summarized in a recent working paper (Wang and Ida 2017). Residential customers were recruited to participate in a demand response experiment. Of these, 365 were given the option to opt in to a time-varying rate and 183 customers were defaulted onto the new rate. Whereas opt-in rates typically fall within the range of 2-10%, 64% of customers opted into the time varying rate. Presumably, this is because the study sample is comprised of only those customers who actively select into a field experiment. Interpretation of the estimated demand response is further complicated by the fact that program participants were insured against losses (i.e., they could only gain from participating in the experiment).

6. The other ten studies are described in Cappers and Sheer (2016). Most evaluated other aspects of time-varying pricing, such as the impact of providing customers with “shadow” bills, which documented how much they would have paid under standard pricing. Only one of the other studies compared opt-in and opt-out recruitment approaches (Lakeland Electric) but the data the utility provided did not contain enough detail to perform a comparable analysis.

of-Use (TOU) rates had been extremely low. From SMUD’s perspective, the SGIG program was an opportunity to maximize the benefits of their smart-grid technology investments, and to test time-varying rates that were designed to meet their evolving load management needs (Jimenez, Potter, and George 2013).

The study sample was drawn from SMUD’s population of residential customers. To define the experimental population, several selection criteria were applied. Households were excluded: if their smart meter had not provided a year’s worth of data by June 2012; if 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 population.⁷

Households in the experimental population were randomly assigned to one of ten groups, five of which are the focus of this paper.⁸ 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 TOU and a Critical Peak Pricing (CPP) program. 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 by default, but had the opportunity to leave the program if they wished. All encouraged households (opt-in and opt-out) were also offered enabling technology – an in-home display that provided real-time information on consumption and the current price.

Figure 2 summarizes the standard, TOU, and CPP rate structures that are evaluated in this study. All SMUD customers faced an increasing block pricing structure. This means that the price paid for the first block or “tier” of electricity consumed during a billing period was lower

7. SMUD reports no statistically significant differences between the households in the study sample and the larger residential customer base. We did not have access to these sample comparisons, and we do not know which variables were analyzed. Most residential customers had smart meters in time for the experiment, though many were excluded because their meters had not reported a full year of data by June 2012.

8. The other five groups were: defaulted to another time-varying rate that did not have a corresponding opt-in group treatment (i.e., Critical Peak Pricing (CPP) plus TOU rate); encouraged to opt in to CPP or TOU without the enabling technology described below; or were part of a recruit-and-deny randomized controlled trial for TOU rates.

than the price paid for the higher tier. During the time period of our study, customers on the 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 within a monthly billing period. Under the TOU program, customers faced the same monthly fixed charge of \$10. These customers paid a higher rate, \$0.2700 per kWh, for electricity consumed during the “peak period” from 4PM to 7PM on non-holiday weekdays. They paid a lower rate (relative to the standard rate structure), in all other “off-peak” hours, \$0.0846 per kWh for the first 700 kWh 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 a significantly higher rate, \$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.

[FIGURE 2 HERE]

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 in the Energy Assistance Program Rate (EAPR) were eligible to participate in the study. No matter the pricing plan, EAPR customers received about a 30 percent discount on their rates. Both the TOU and CPP rates were designed to be approximately revenue neutral to the utility if customers selected their rate plan randomly and did not adjust their consumption (Jimenez, Potter, and George 2013).

To summarize, the five randomized groups we study include: the CPP opt-in group, which was encouraged to enroll in the CPP program; the CPP opt-out group, which was notified of enrollment and encouraged to stay in the CPP program; the TOU opt-in group, which was encouraged to enroll in TOU program; the TOU opt-out group, which was notified of enrollment and encouraged to stay in TOU program; and the control group, which was not encouraged to participate in a time-varying rate, nor even told about the program at all, and remained on SMUD’s standard

rates.⁹

3.2 Encouragement Messages

Materials and messages encouraging participation were virtually identical across the opt-in and opt-out treatment groups. The encouragement effort for opt-in households consisted of two separate mailed packets. The first was sent in either October 2011, to about 20 percent of the encouraged households, or November 2011, to the remaining 80 percent. The second was sent in January 2012. Each packet included 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. The recruitment materials listed generic benefits of participating in rate programs, including saving money, taking control, and helping the environment. 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.

Recruitment activities and program enrollment are summarized in Figure 3. About half of the customers 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).

[FIGURE 3 HERE]

The opt-out groups were mailed one packet containing a letter, brochure, and business reply card. These materials were designed to look as similar as possible to the materials received by members of the opt-in groups. Packet mailings were followed within two weeks by a reminder post card. About 10 percent of the packets were sent on March 12, 2012 and the remaining 90 percent were sent on April 5, 2012.

9. Sample sizes for control and treatment groups were determined using a set of power calculations designed to account for different enrollment probabilities between groups, required Type I and Type II error rates, minimum detectable effect size, cost of treatment, and the comparison of all treatment groups to a single control group. In general, the opt-in treatment groups were larger because a smaller proportion of these customers were expected to enroll than in the opt-out groups. Additional detail on these calculations is available in the appendix to Jimenez, Potter, and George (2013).

The TOU opt-in group received slightly different encouragement messages from the other groups because they were part of a recruit-and-delay randomized controlled trial (which we are not incorporating into this analysis). 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 experiment. 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 in 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 active joiners in the CPP opt-in group could be somewhat different from the active joiners in the TOU opt-in group, as the TOU active joiners had to be willing to accept some probability that their enrollment would be delayed. Thus, while the CPP opt-in group can be directly compared to the CPP opt-out group, comparisons between the TOU opt-out and opt-in groups are drawn with the caveat that these two groups were encouraged and recruited slightly differently.

4 Data and Methodology

4.1 Data Description

Our analysis uses household-specific data, electricity consumption data, and weather data. The household-specific data include experimental cell assignment, dates of enrollment, disenrollment, and account closure information for households who moved. Finally, for some households, we have responses to two large-scale surveys: a demographic survey and a customer satisfaction survey.

We also have data on households' energy consumption, as well as their associated expenditures. Specifically, we have 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. Electricity consumption is measured in kilowatt hours (kWh). We collect energy consumption data for all

households in the experimental sample, including the control group, for the duration of the study period. Households that moved are one exception. These households were not tracked to their new location, 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. Hourly energy consumption and billing data are quite complete. Less than one percent of these data are missing. The frequency of missing data does not differ systematically across treatment groups.

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 do not vary across households, only over time.

4.2 Validation of Randomization

Table 1 provides summary statistics by experimental group. The top three rows summarize information on daily consumption, the ratio of peak to off-peak energy consumption, and billing from the pre-treatment summer (June to September 2011). Sample households consume slightly less electricity than the average U.S. household – approximately 27 kWh per day during the four summer months compared to almost 31 kWh per day across the U.S. in 2011. The ratio of peak to off-peak usage provides one indication of a customer’s exposure to the higher peak prices under CPP or TOU, and bill amounts reflect the average monthly bill in the pre-treatment summer. Bills in our sample are very close to the national average, reflecting that SMUD customers pay higher prices than the average U.S. residential customer. For all three variables, we also report t-statistics on the test that the mean for each treatment group equals the mean for the control group.¹⁰ The t-statistic exceeds one for only one of these comparisons, suggesting that the randomization yielded groups with very similar means across these three variables.¹¹

10. We also run t-tests comparing the opt-in to opt-out treatment groups, and find no statistically significant differences.

11. Given that we will be analyzing consumption across hours of the day, we are particularly concerned about balance in consumption profiles. In addition to the ratio of peak to off-peak usage, the Appendix provides a breakdown

The “structural winner” variables measure the share of households that would pay less on either the CPP or TOU pricing policy, assuming no change in their consumption (following industry convention, we refer to households who would pay less as “structural winners” and those who would pay more as “structural losers”). Approximately half of all customers are estimated to be structural winners, based on consumption data collected before the intervention.¹²

[TABLE 1 HERE]

4.3 Methodology

4.3.1 Estimating average impacts for encouraged groups

We estimate a difference-in-differences (DID) specification using data from the pre-treatment and treatment periods to identify the average intent-to-treat (ITT) effect, i.e., the average effect for each encouraged group. Equation (1) serves as our baseline estimating equation, where y_{it} measures hourly electricity consumption for household i in hour t . All specifications described below are estimated separately for the opt-in and opt-out groups, unless otherwise noted. Z_{it} is an indicator variable equal to one starting on June 1, 2012 if household i was encouraged to be in the treatment group, and zero otherwise. γ_i is a household fixed effect that captures systematic differences in consumption across households, and τ_t is an hour-of-sample fixed effect.

$$y_{it} = \alpha + \beta_{ITT}Z_{it} + \gamma_i + \tau_t + \varepsilon_{it} \quad (1)$$

We estimate four sets of regression equations. Each set uses data from the control group and one of the four treatment groups. The coefficient of interest is β_{ITT} , which captures the average of consumption across all 24 hours of the day (Figures 1.1, 1.2). Again, all four treatment groups look very similar to the control group.

12. For several of our analyses, including identifying structural winners under the CPP program, we need to simulate 12 CPP days in the pre-treatment period. We do this by choosing the 12 hottest non-holiday summer weekdays. To ensure that our estimates of structural winners do not result from idiosyncratic variation on these 12 days, we also estimate specifications where we randomly select 12 of the 24 hottest non-holiday summer weekdays and recompute our estimate of pre-period CPP bills. We repeat this exercise 10,000 times and then average over the estimated pre-period CPP bills to obtain an alternative measure of structural winners. The correlation between the two measures is 0.97.

difference in hourly electricity consumption across treated and control groups, controlling for any pre-treatment differences by group.¹³ Within each set, we estimate the model separately using data from event day peak hours (4pm to 7pm on the twelve CPP days in each summer) and non-event day peak hours (4pm to 7pm on non-event, non-holiday weekdays during the summer).¹⁴

4.3.2 Estimating average impacts for treated households

We estimate a DID instrumental variables (IV) specification using data from the pre-treatment and treatment periods to identify a local average treatment effect (LATE). Specifically, we estimate Equation (2), where y_{it} , γ_i , and τ_t are defined as in Equation (1). $Treat_{it}$ is an indicator variable equal to one starting on June 1st, 2012 if household i was actually enrolled in the time-varying pricing program, zero otherwise (estimated separately for the opt-in and opt-out groups). We instrument for $Treat_{it}$ using the randomized encouragement to the corresponding treatment Z_{it} , which is defined as in Equation (1).

$$y_{it} = \alpha + \beta_{LATE}Treat_{it} + \gamma_i + \tau_t + \varepsilon_{it} \quad (2)$$

The β_{LATE} coefficient captures the average reduction in household electricity consumption among customers enrolled in the time-varying pricing program. To interpret β_{LATE} as a causal effect, we must invoke an exclusion restriction, which requires that the encouragement (i.e., the offer to opt in or the default assignment into treatment with the ability to opt out) affects electricity consumption only indirectly via an effect on participation. We also invoke a monotonicity assumption which requires that our encouragement weakly increases (versus reduces) the participation probability for all households.¹⁵ In Section 1.3, we conduct a partial test of these identi-

13. We present specifications with the dependent variable measured in levels because the cost savings from time-varying pricing are a function of kWh reduced, not the percent reduction. Our results are not sensitive to alternative functional forms, and the Appendix presents specifications in logs (Tables 1.3 to 1.5).

14. Note that customers under the TOU pricing plan face the same prices on event and non-event days. We estimate separate impacts for comparison to CPP.

15. This is equivalent to “frame monotonicity” as defined by Goldin and Reck (2019), and means, for instance, that consumers would not always select against the default and opt in to time-varying pricing if its not the default and

fying assumptions using a separate treatment group that was encouraged to enroll but not given the opportunity to participate in time-varying pricing during our study. Examining the response of households in this treatment allows us to place an upper bound on the degree of possible bias resulting from a violation of the exclusion restriction via an encouragement effect. Even under conservative assumptions regarding this possibility, our findings are qualitatively unchanged.

4.3.3 Estimating average impacts for passive consumers

Conceptually, our sample of residential customers can be divided into three groups (see Figure 4). Active leavers are households who opt out of an opt-out program and do not enroll in an opt-in program. Passive consumers are households who do not actively enroll in an opt-in program, but who also do not actively drop out of an opt-out program. Active joiners are households who actively enroll in an opt-in program and remain in an opt-out program. Note that a comparison of average electricity consumption across the opt-in and opt-out groups (the top two rows in Figure 4) estimates the average effect of being assigned to the opt-in versus opt-out groups. Scaling this difference by our estimate of the population share of passive consumers yields an unbiased estimate of the average effect of time-varying rates on electricity consumption among passive consumers.¹⁶

[FIGURE 4 HERE]

We estimate the DID IV specification using data from 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 1, 2012 if a household was encouraged into the opt-out treatment group only. This IV specification isolates the average causal effect of these pricing programs on electricity consumption among passive consumers. To interpret our estimates in this way, we again invoke the exclusion restriction which requires

opt out of it when it is the default.

16. Our approach to isolating the response of the passive consumers is very similar to Kowalski (2016), although our setting is considerably more straightforward since we randomized the assignment of both the opt-in and the opt-out treatments.

that the encouragement (the offer to opt in or the default assignment with the ability to opt out) does not directly affect electricity consumption among active joiners, active leavers, or passive consumers. As Figure 4 makes clear, we are also assuming that active joiners who actively enroll in the pricing programs under the opt-in treatment do not respond differently to time-varying pricing, on average, as compared to active joiners who are defaulted onto the programs through the opt-out treatment. Section 1.3 contains a detailed discussion of these exclusion restrictions. We note that, to the extent actively encouraging households to opt in leads to a larger demand response, our estimates of the active joiners' reductions in response to prices will be overstated and our estimates of the demand response among passives will be understated.

5 Main Results

5.1 Default Effects on Program Adoption

Table 2 summarizes customer acceptance of time-varying pricing in the opt-in and opt-out groups, respectively. The columns titled "Initial" summarize customer participation at the beginning of June 2012 (the month the new rates went into effect). The columns titled "End line" summarize participation at the end of the second summer (September 2013). In both sets of results, the first column reflects the share of customers on the time-varying rate while the second column reports the number of customers on the rate.

[TABLE 2 HERE]

The initial participation results provide striking evidence of the default effect. For both the CPP and TOU rates, approximately 20 percent of those assigned to the opt-in encouragement elected to opt in. 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.¹⁷

17. It is worth noting that SMUD was more successful than expected at recruiting customers onto time-varying rates. 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

To interpret the “End line” columns, it is important to understand how we are describing the eligible population. If customers moved, they were no longer eligible for the time-varying rates, even if they moved within SMUD’s service territory. Also, new occupants were not included in the pilot program. The numbers in Table 2 report rates and enrollees after dropping movers. For instance, the number of customers on CPP from the opt-in group fell from 1568 to 1169 because 399 households (approximately 25 percent) moved between June 2012 and September 2013. SMUD reports move rates of approximately 20 percent per year across their entire residential population, so a move rate of 25 percent over a 16-month period that includes the summer, when moves are most likely, is reasonable. Across the four treatment groups, the move rates are very similar, ranging from 23 percent in the CPP opt-out group to 26 percent in the TOU opt-in group.¹⁸

5.2 Choice Modification

We observe modifications to consumers’ participation choices after the program started, although program rules constrained the set of possible changes. Customers in the opt-in group were not allowed to enroll after June 1, 2012; customers in the opt-out group who had already opted-out were not allowed to change their minds and opt back in. However, customers in both groups who had initially chosen to participate in the time-varying rate program could revert to the standard rate at any time.

The final column of Table 2 reports the difference between initial and end line participation rates, 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 7 percent of the original participants. Participation in both of the opt-out groups fell by more percentage points (6.6 in the case of CPP opt out, 96.0 – 89.4, and 5.3 in the case of TOU opt out), but again reflected 7 or fewer percent of the original participants.

Although only a small share of households dropped out of these programs, we conducted

satisfied with the utility and trust its recommendations, they may have been more likely to accept the default. SMUD anticipated that approximately 50 percent of the customers would remain on the rate with opt-out.

18. Moving rates are not statistically significantly different from one another (z-statistic on the largest difference equals 1.3).

a hazard analysis of attrition, described in Appendix 1.4. Comparisons of attrition rates across the opt-in and opt-out groups are under-powered, but some suggestive patterns emerge. 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 statistically significant effect on drop-outs.

5.3 Follow-on Behavior

5.3.1 Average impacts for encouraged households

Table 3 summarizes the estimation results for the DID specification in Equation (1) that uses data from the pre-treatment and treatment periods to identify an ITT effect. The first two columns use data from peak hours on “critical 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 the hottest non-holiday weekdays during the summer of 2011.¹⁹ 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 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, we conclude that providing households the opportunity to opt-in to the CPP treatment leads to an average reduction in electricity consumption of 0.129 kWh during peak hours of event days (averaged across all household that received the opt-in offer). The estimate for the opt-out group is considerably larger at 0.305 kWh across all households defaulted onto the CPP rate.

The coefficients in the last two columns show that CPP customers *reduced* their consumption during peak hours on *non-event* days (by 0.029 kWh per household in the opt-in group and 0.094

19. We have also estimated specifications based on random samples of 12 days within the hottest 24 days. Our results are not sensitive to this choice.

kWh per household in the opt-out group). Recall that CPP customers faced rates that are slightly lower than the standard rates on these non-event days. These kWh reductions are considerably smaller compared to event days for the CPP households, but still statistically significant.

Why might consumers respond to a decrease in electricity price with a decrease in consumption? 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, who faced higher prices during peak hours for all weekdays (not just event days), the results show that households reduced their daily peak consumption by 0.091 kWh on average in the opt-in treatment, and 0.130 kWh on average in the opt-out treatment on days that were called as event days for CPP customers (i.e., relatively hotter days). On all other peak days average reductions are estimated to be 0.054 kWh per household in the opt-in treatment, and 0.100 kWh per hour in the opt-out treatment. Given that non-event-day consumption is lower, the results are approximately the same in percentage terms (3.6-5.2% for the TOU opt-in group and 5.9 - 7.3% for the TOU opt-out group – see Table 1.3).

Finally, we regenerate the results reported in Table 3 using only the post-intervention data. In other words, we do not use the pre-period data, and we simply compare treated households' consumption to the control households' during event and non-event peak hours. This exercise yield qualitatively similar results, which are summarized in Table 1.6. The average reductions for the opt-out group are nearly 3 times larger than the average reductions for the opt-in group for CPP and 2 times larger for TOU. The coefficient estimates do differ slightly from those reported in Table 3 since there were some statistically insignificant pre-period differences by group.

5.3.2 Average impacts for treated households

Table 4 reports on the instrumental variables specifications that correspond to Equation (2). Similar to Table 3, the columns on the left of the table report estimates using data from CPP event

hours and the columns on the right report results estimated using data from non-event-day peak hours. The top of the table corresponds to CPP customers while the bottom corresponds to customers participating in TOU programs.

Estimates in the first two columns suggest that the active joiners in the opt-in CPP group reduced consumption during event-day peaks by almost twice as much as the larger group of active joiners plus passive consumers participating in the CPP program in the opt-out group (0.658 compared to 0.330 kWh per household). The magnitude of the reduction for the opt-in group (0.658 kWh) is large and suggests consumers did more than simply turn off a few light bulbs. Given that electricity rates increased by approximately 350 percent during critical peak events, this reduction off a mean of almost 2.5 kW is consistent with a price elasticity of approximately -0.075. This is comparable to other short-run demand elasticities estimated for electricity consumption, though typically those estimates are based on demand reductions over longer time periods (EPRI 2012).

In the fourth and fifth columns of Table 4, we see again that households in both the opt-in and opt-out CPP treatments significantly reduced their consumption on non-event peak days. Passive consumers' average reductions on non-event days comprise a larger share of the average critical peak reductions than is true for active joiners. This is consistent with the latter group fine-tuning their demand to changing conditions, whereas passive consumers may rely to a larger extent on modifications that do not require sustained attention (such as reprogramming a thermostat to reduce cooling load during peak hours on all days).

In the case of the TOU treatments, the LATE estimates indicate that active joiners 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 active joiners plus passive consumers in the TOU opt-out group (0.480 relative to 0.136 kWh per household), and almost three times as much (0.287 relative to 0.105 kWh per household) during non-event regular peak days.²⁰

[TABLE 4 HERE]

20. In joint specifications, we can reject that the coefficient estimates are equal across the opt-in and opt-out groups in all cases except for the CPP treatment on non-event days ($p=0.249$).

The results in the third and sixth columns isolate the effect of time-varying rates on electricity consumption among the passive households. Comparing the results in the first column (active joiners), to the results in the third column (passive consumers), suggests that the average response among active joiners to the CPP rate was about 2.7 times larger than the response among passive consumers during event hours. Passive consumers were more similar to active joiners during non-event peak hours, reducing by only half as much.²¹ Differences between active joiners and passive consumers are more pronounced with the TOU rates. Given that there are so many more passive consumers exposed to the rates under an opt-out experimental design, the aggregate savings from an opt-out design is significantly higher than from an opt-in design (as is made evident in Table 3).

Tables 3 and 4 have averaged treatment effects across all peak hours. Figure 5 illustrates these effects graphically, disaggregating by hour. The figure depicts hour-by-hour LATE estimates for event days across the four treatment groups relative to the control group. We also test for changes in consumption during non-peak hours. One might expect that some consumers would increase consumption in the hours leading up to the peak period (cooling the house when prices are relatively low, for example). However, we find that consumers are reducing consumption in the hours before the peak period, statistically significantly so for the active joiners in both the CPP and TOU groups.

[FIGURE 5 HERE]

Finally, we estimate the impacts on household electricity expenditures with an alternative form of Equation (2) that features total bill amount as the dependent variable. Table 5 summarizes these estimation results. The coefficient estimate in the first column of the top panel suggests that bills for customers who opted in to the CPP rate plan fell by 5.7% on average, with a mean reduction of \$6.52 on an average summer bill of \$114. Bills for the typical participant in the opt-out group fell by less — around \$4.50 for the group overall and slightly less for the passive

21. Note that the coefficient estimates for the opt-out group in Table 4 are equal to the weighted sum of the coefficients for the active joiners (e.g., -0.658 for CPP event hours) and the passive consumers (-0.242), with weights set equal to the share of active joiners relative to total opt-out enrollees and one minus this number from Table 2.

consumers. This is consistent with the results presented in Table 4, which shows how passive households reduced consumption by less during critical peak periods.²²

[TABLE 5 HERE]

5.3.3 Impacts over time

Since our study period includes two years of post-intervention data, we can 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 passive 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 6 summarizes the estimation results. For the CPP treatments, the interaction term is positive for the active joiners in the opt-in group (columns 1 and 4) and negative for the passive consumers (columns 3 and 6). Three out of four of the coefficients are statistically significant.²³ This pattern suggests that demand response is attenuating over time among active joiners. In contrast, the average demand response is increasing over time among passive consumers. This could be due to a growing number of passive consumers responding over time, or an escalating demand response as passive customers gain experience with the program.²⁴

[TABLE 6 HERE]

We also investigate whether customers who experienced higher than normal bills once the program took effect had different treatment effects. We construct a binary ‘bill shock’ indicator which equals one if a participating customer received a bill in the first year of the program that

22. Bill reductions should not be interpreted as a measure of consumer welfare impacts; customers may have made adjustments that were costly from a monetary or welfare perspective. We return to this point below.

23. The results for the TOU treatment are less pronounced, although columns 1 and 4 suggest that the active joiners are responding less over time. Since we had attrition in the set of participating customers over time, the results could also reflect changes in the types of customers who are still treated in the second summer. Table 1.2 estimates these specifications on a balanced panel, i.e., on customers who did not change their enrollment status during the treatment period. We find that the results are qualitatively the same and slightly larger in magnitude overall.

24. Table 1.1 explores additional heterogeneity in impacts by customer type. We show, for instance, that structural winners are, if anything, more responsive to time-varying rates while low-income customers are less responsive.

was 20% greater than the bills received in the pre-program summer. Presumably, this group of ‘shocked’ customers is largely comprised of structural losers who did not initially adjust consumption. Mechanically, demand reductions among shocked customers are smaller on average during the first year of the program as compared to unshocked customers in that first year. But notably, we find the demand reduction in the second year among shocked customers to be significantly larger (see Table 1.17). While we find no evidence that those who were structural losers dropped out of the program at a faster rate (see Section 1.4), bill shocks may have caught the attention of customers, explaining the larger than average demand reductions among these customers in the second year.

6 Explanations for the Default Effect

We next investigate the underlying mechanisms that could be generating the default effect in our setting. We assess these explanations in light of three key empirical facts. First, we have shown how switching the default choice significantly impacts the rate of participation in time-varying electricity pricing programs. Second, whereas the impacts of time-varying pricing on *aggregate* energy consumption and expenditures are economically significant, we have shown that the *household-level* impacts are quite small. Finally, we will document a striking lack of correlation between a household’s likely gains from program participation and its program enrollment decision, even in the presence of an enrollment deadline. Taken together, we will argue that these empirical facts are more consistent with a model that uses inattention to generate a default effect, versus high switching costs or present-biased preferences.

6.1 Program Benefits Are Poor Predictors of Participation Choices

To model the relationship between program benefits and consumers’ participation choices, we construct measures of *household-level* gains from participation. Let \bar{X}_i denote the optimal vector of electricity consumption (i.e., peak versus off-peak consumption) under the standard price

schedule \bar{P} for a representative household i . Let \tilde{X}_i denote the optimal vector of electricity consumption under the time-varying price schedule \tilde{P} . If we assume that utility is quasi-linear in electricity consumption and consumption of other goods, a monetary measure of the annual benefits from switching from \bar{P} to \tilde{P} can be summarized by: $\max\{\bar{P}'\tilde{X}_i - \tilde{P}'\tilde{X}_i, \bar{P}'\tilde{X}_i - \tilde{P}'\tilde{X}_i - A_i\}$. The first argument measures the change in electricity expenditures holding consumption patterns constant. We refer to this subsequently as ‘structural gains,’ recognizing that these gains will be negative if expenditures increase on the time-varying rate. The second argument measures the change in expenditures if the household re-optimizes consumption net of any adjustment costs, A_i .²⁵

If we assume that the household will only choose to re-optimize if the benefits from adjustments exceed the costs, then the structural gains provide a lower bound on household-level benefits. We estimate structural gains for each household under both types of time-varying pricing programs (CPP and TOU) using hourly data on household-level electricity consumption from the pre-treatment period (2011).²⁶ Using these monthly benefits estimates, and assuming a discount rate of 5%, we construct household-specific estimates of the net present value of structural gains from participating in a time-varying pricing program. The bottom panel of Figure 6 summarizes the distributions of these values by group.²⁷ These structural gains are not large; we estimate that 96% and 93% of households would have experienced monthly bill differences of less than \$10 under CPP and TOU pricing, respectively. This is consistent with SMUD’s goal to limit impacts on monthly bills for most customers (Jimenez, Potter, and George 2013).

The top panel of Figure 6 shows a striking lack of correlation between the structural gains

25. A_i captures any costs of re-optimization in response to the change in price schedule. This can include both the utility impacts of changes in energy consumption patterns (e.g., tolerating warmer indoor temperatures on hot days) or any adjustment costs (e.g., the effort required to reprogram a thermostat).

26. We use the control group to assess the extent to which a customer’s structural gains in 2011 are correlated with structural gains in subsequent years. For the TOU rate, the correlation between pre-period and treatment period structural gains is 0.82 (correlation with the 2012 summer) and 0.79 (correlation with 2012-2013 summers). For the CPP rate, these correlations are 0.73 and 0.72. These strong correlations indicate that structural gains are persistent over time and support our use of pre-period data to estimate consumers’ expected structural gains under time-varying pricing across all treatment groups.

27. Under CPP pricing, the average customer has structural gains of \$0.33 (in net present value) and 51% of households are structural winners. Under TOU pricing, the average customer has structural losses of \$10.54, and 34% of households are structural winners.

from participation and the participation rate. A significant share of the structural losers participate in the new rates while some of the largest structural winners don't participate.²⁸ Notably, 48% of the households opting out of the CPP program and 60% of the households opting out of the TOU program actively switched *away* from a pricing regime under which our lower bound estimates suggest they should expect to benefit. Formally, regressions of program participation on structural gains identify small and inconsistent effects across treatment groups, ranging from a 0.03% increase ($p < 0.01$) for each additional dollar of structural gains in the TOU opt-in group to a 0.06% decrease ($p < 0.1$) for the CPP opt-out group (see Table 1.13). Given the limited range of structural gains, these differences represent minor shifts in the likelihood of participation, as demonstrated in the figure. In what follows, we show that this lack of empirical correlation between structural gains and participation choice is inconsistent with some standard explanations of default effects.

[FIGURE 6 HERE]

6.2 Switching Cost Model

We now turn to our consideration of underlying causes of the default effect, beginning with the most standard explanation: switching costs. A simple model elucidates the mechanism and provides a framework for evaluating this explanation empirically.

We assume that the benefits from participation B_i are distributed in the population according to some distribution $f()$. As Figure 6 shows, these benefits can be negative if the household would fare worse on the time-varying program. Switching away from the default choice incurs a cost of s . In our context, this could reflect the cost of calling the utility or visiting the website to switch away from the default. If households make fully informed decisions, customers defaulted onto the standard rate will actively opt in to the time-varying rate if $B_i > s$. Customers defaulted on to the time-varying program will actively opt out if the cost of switching away from the default,

28. For example, in Figure 6, 5.6% of households are associated with structural losses that exceed \$50 in net present value on the CPP rate. These losses notwithstanding, 21% of these customers participated in the new rate.

s , is less than the future cost (or negative benefit) of remaining in the pricing program: $B_i < -s$.

This simple choice model will generate a significant default effect if $F(s) - F(-s)$ is large. This will be the case if switching costs are large relative to discounted participation benefits. The model predicts that program participation will be positively correlated with discounted benefits.

Figure 6 provides graphical evidence that is inconsistent with this prediction. To demonstrate more formally how this model fails to rationalize the participation choices we observe, we implement this switching cost model empirically. In the opt-in treatment, for example, we assume that a household will actively opt in if:

$$B_i - s_i + \epsilon_i > 0. \tag{3}$$

We use the household-specific structural gains to proxy for household-specific benefits B_i . The s_i is a household-specific switching cost to be estimated.

To identify the parameters of the switching cost distribution that best rationalize observed participation choices, we must invoke some additional assumptions. We assume that the error term ϵ_i is a mean zero, type I extreme-value random variable. And we assume that households correctly anticipate how they would benefit under the new program (B_i). Section 1.5.2 describes this econometric exercise in detail. Overall, these estimates are implausibly large given that switching away from the default option required only a phone call, a text, or an email. Instructions were clearly displayed on all marketing materials.

Why are these cost estimates so large? Intuitively, switching costs are identified relative to benefits which the model assumes are fully accounted for by households. If, in fact, these benefits are not fully accounted for (or ignored) by households, the model will be mis-specified in a way that inflates switching cost estimates. Not only are these cost estimates too large (in absolute value), but a model that assumes participation choices are driven by comparisons of expected benefits against a reasonable switching cost predicts a strong correlation between expected benefits and participation. As noted above, we do not see this correlation in the data.

6.3 Present-Biased Preferences Model

The significant default effect we document could also reflect present-biased preferences and procrastination. In our context, it seems quite plausible that households might have intended to opt-out or opt-in, but did not get around to doing so. Were this the case, the participation choice should more accurately be represented as a choice between switching today, planning to switch later, or never switching at all. In addition to the exponential discount rate δ , the household may also exhibit a present-bias, parameterized by β , which additionally discounts all future periods by a constant amount. This type of discounting is also referred to in the literature as hyperbolic discounting (e.g., Frederick, Loewenstein, and O’Donoghue 2002; DellaVigna 2009).

We outline a simple model with present-biased preferences in Appendix A.5 and demonstrate how a key testable prediction of the model is that households that face a deadline and have higher structural gains will be more likely to switch, while households without a deadline may never actively make a choice to switch or not, and therefore their participation status will be uncorrelated with structural gains. In our setting only the opt-in treatment groups faced a deadline (they had to join the program by June 1st, 2012 or they would be prevented from joining for the duration of the program), while the opt-out treatment group could return to the standard rate at any point in the program and therefore did not face a specific deadline. If procrastination and present-biased preferences explain the default effect, we should see a positive correlation between structural gains and participation status in the opt-in arm, but not necessarily in the opt-out arm.²⁹

Figure 6 indicates that, for the TOU group, there is a slightly higher probability that the participation decision is correlated with structural gains for the opt-in group compared to the opt-out group, consistent with present-biased preferences in the presence of a deadline for the opt-in group. However, the degree to which the correlation differs between the TOU opt-in and opt-out

29. (Gottlieb and Smetters 2019) document the role of forgetfulness in explaining why consumers miss deadline in a different context. These authors study consumers who fail to make premium payments before a scheduled deadline which results in policy termination. In this insurance setting, the insurance company stands to benefit if consumers miss the deadline. In our setting, the utility has a strong incentive to remind consumers about the approaching participation deadline. Figure 3 shows how customers in the opt-in group received frequent reminder phone calls, door hangers, and mailings right up to the participation deadline. We therefore assume that forgetting is less likely in this setting.

groups is, while statistically significant, very small. The difference is not statistically different from zero for the CPP treatments. This suggests that, while present-biased preferences may be one factor contributing to the default effect, they cannot explain all the variation observed in the data, so there must be additional explanations at play.

6.4 Inattention

We now introduce the possibility that customers are inattentive to benefits when making their participation decisions. In our context, customers must exert significant effort to collect the information they would need to fully understand how their households' energy consumption patterns would determine expenditures under the new, time-varying rate.³⁰ Inattention to this information could be rational if the impacts of switching from the standard rate are small relative to the effort costs required to make an informed decision (Sallee 2014).

In this augmented framework, the participation choice is modeled in two steps. First, the customer decides whether to exert the effort required to collect the information she would need to make an informed decision. Specifically, we assume there is some basic information about participation benefits b_i that she can easily assess without effort. The customer could make her decision on this basis. Or, if she exerts more effort, she can collect additional information in order to refine her estimate of benefits to $B_i = b_i + \alpha_i$, her true benefits.³¹ If α_i is pivotal the household's participation decision will differ across the informed and uninformed states.

A rational decision-maker will exert effort if the expected returns exceed the effort costs. Consider, for example, opt-in households that are defaulted onto the standard rate. For customers who would opt in on the basis of b_i , the expected returns to exerting additional effort are: $(1 - F(s))E[s - B_i]$, assuming that customers know the true distribution of B_i . For those who would

30. In fact, there is considerable evidence to suggest that consumers poorly understand the relationship between energy consumption and energy bills. For example, Attari et al. (2010) show that when asked to guess, customers underestimate the electricity used by high energy activities, such as clothes dryers, by more than an order of magnitude and overestimate electricity used by other energy services such as lighting. See also Myers, Puller, and West (2019) and Todd-Blick et al. (2020).

31. We are assuming that customers who invest effort to learn their benefits learn their true benefits. Under a more complicated model, customers could invest to obtain a better, though still imperfect, estimate.

not opt in based on uninformed priors, the expected returns on effort are $F(s)E[B_i - s]$. Note that for both of these expressions, the first term is the probability that a decision-maker will change their program participation decision given the additional information α_i and the second is the expected value of the decision change.

We first demonstrate how this model can rationalize the patterns of participation we observe. We consider a scenario in which all households make uninformed program participation decisions. We assume that prior beliefs b_i are distributed normally in the population and are uncorrelated with structural gains. Given the complexity involved in mapping electricity consumption to time-varying rate schedules described above, this lack of correlation seems plausible and consistent with previous work documenting inattention to electricity consumption. If we further assume that switching costs are constant across households, we can identify the mean and variance of the distribution of uninformed priors that best rationalizes observed participation choices over a range of switching costs. The two participation shares we observe in the opt-in and opt-out groups, respectively, allow us to identify the two parameters of the distribution of prior beliefs.³²

The “Uninformed prior” rows of Table 7 report the estimated means and standard deviation of the distribution of prior beliefs about participation benefits. For an assumed switching cost of \$10, the average benefit prior is \$3.53 (standard deviation is \$7.73). As customers could reasonably expect to reduce expenditures by a few dollars, these estimates seem reasonable.

Having estimated the distribution of prior beliefs about benefits, we can now ask whether, conditional on this distribution, inattention to the enrollment decision could be rational. More precisely, we simulate the participation choices that households would make based on uninformed priors and contrast these with informed choices. The difference in benefits net of switching costs across these two scenarios can be interpreted as an estimate of the return on effort. If these returns look small relative to the effort cost required to collect information, inattention to this decision could be rational.

To calibrate *informed* estimates of participation benefits, we construct household-specific es-

32. More specifically, the identifying conditions set $F(-s) = 0.96$ and $1 - F(s) = 0.20$.

estimates of the present discounted gains associated with program participation. We consider two heuristics in particular, which we use to bound the expected returns. Under the first, the “No adjustment” heuristic, we assume that *informed* household decision-makers do not account for the possibility that they might adjust consumption in response to the time-varying rate. In other words, we use our estimates of household-specific structural gains to proxy for informed expectations about participation benefits. Under the “With adjustment” heuristic, we assume that informed household decision-makers anticipate that they will adjust their consumption patterns in response to the time-varying rate. To estimate the additional value obtained via demand response, we use the average additional monthly bill savings expected for participating customers in the opt-out group (\$4.50 per month in the CPP group, see Table 5). Subtracting the average monthly structural gains under CPP pricing (\$0.045) we estimate average benefits of \$4.45 per month in addition to structural gains.

The second panel in Table 7 summarizes simulated choices under the two heuristics for the CPP experiment. We compare bill impacts associated with the informed choice (net of assumed switching costs) against bill impacts associated with the uninformed prior (net of assumed switching costs). This yields a distribution of estimated returns to paying attention. The distribution under the “No adjustment” heuristic provides a lower bound on the costs of inattention. Under the “With adjustment” heuristic, which assumes consumers account for demand response, this difference provides an upper bound because it reflects bill savings but does not account for the dis-utility associated with re-optimization of energy consumption.

[TABLE 7 HERE]

From the perspective of a household faced with this program participation choice, we estimate that the average returns on effort are low. Under the first heuristic, the average discounted returns on effort are in the range of \$3 to \$14 over the full two-year pricing program. Estimates are lower in the opt-out case because these consumers are, on average, nudged in the right direction. Upper bound estimates under the second heuristic are somewhat higher (\$24 to \$31). Given the time and cognitive effort required to gather and process information about how one’s

electricity expenditures might change under time-varying pricing, we speculate that the effort costs of making an informed decision could easily exceed the returns on this effort for a majority of households.

To summarize, the fact that benefits are largely uncorrelated with program participation (see the top of Figure 6) suggests the default effect is unlikely to be driven purely by switching costs or discount rates in our setting. To rationalize the participation patterns we observe, we need an explanation that somehow breaks the link between participation benefits and enrollment choices. Present-biased preferences could offer such an explanation; people who stand to benefit from switching away from the default could procrastinate indefinitely. However, the fact that the relationship between structural gains and participation remains weak, even in the presence of a participation deadline, suggests that present-biased preferences are not a sufficient explanation. Also, crucially, present-biased preferences do not explain active decisions to switch away from dynamic pricing even if lower bound estimates suggest they would gain.³³ Inattention is another mechanism that breaks the link between structural gains and switching. In our context, an electricity consumer would have to make a substantive effort to understand how she would benefit from participating in the time-varying pricing regimes. We offer evidence to suggest that, given relatively low returns on this effort, it could be rational for most consumers to make uninformed or inattentive participation decisions. Rational inattention also explains why some customers might switch away from dynamic pricing even if it is likely to provide benefits as these customers' uninformed priors may suggest that they will not benefit.

7 Implications for Welfare and Cost Effectiveness

We have argued that inattention offers a plausible explanation for the participation choices we observe across experimental treatment groups and the observed lack of correlation between structural gains and participation choice. If rationally inattentive households are nudged into a pricing

33. Note that households that actively switched in to dynamic pricing when our estimates of structural gains suggest they would lose money are easier to understand since our estimate is a lower bound on the gains they expect. These household may know that they will adjust their consumption in response to the electricity prices.

regime that they are more or less indifferent about, this default manipulation offers a powerful means of unlocking an economically significant (in aggregate) and social welfare improving demand response. In this section, we further investigate the implications of this default effect from both the household and utility perspective.

7.1 Consumer Welfare

In Section 5, we find that passive consumers mount an economically significant demand response to time varying prices. These results present something of a puzzle. If passive consumers are largely inattentive to the pricing program participation decision, why are they subsequently attentive to electricity consumption choices?

We posit that inattention is less likely to be rational once customers are nudged into the program. Responding to the new electricity pricing regime once enrolled required less effort as compared to understanding the implications of the initial participation choice. First, enrolled customers were provided with highly salient information (and frequent reminders) about how electricity prices vary across off peak, peak, and critical-peak hours. Second, whereas the initial time-varying pricing program participation decision was an entirely new and unfamiliar choice, all consumers have prior experience with electricity price changes. Although the new pricing plans featured a different kind of inter-temporal price variation (i.e., variation across days and hours versus across billing cycles), decisions about electricity consumption are similar to choices and trade-offs evaluated in the past.

To generate further insights into the likely welfare implications of this default effect, an end-line survey was sent to all households enrolled on the CPP and TOU pricing plans and a subset of the control group after the pricing pilot had concluded. Among participants in time-varying pricing programs, survey responses in the opt-out group were lower (26%; N=566) as compared to the opt-in group (36%; N=183). Although survey respondents are not a random subset of the study sample, their responses shed light on consumers' motivations and sentiments about the pricing programs.

Overall, survey responses indicate a positive customer experience with time-varying electricity pricing. In both the opt-in and opt-out groups, fewer than 7% disagreed 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.” 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 respondents are “very” or “somewhat” satisfied with the standard rate.

Our results are consistent with a scenario in which consumers are nudged onto an unfamiliar electricity rate structure that offers the average consumer small but positive gains. Over time, as consumers gain experience with the new pricing regime many come to prefer it. Although the evidence we document is consistent with a model of rational inattention, we cannot rule out an alternative or additional “endorsement” explanation. In this unfamiliar choice context, households may have viewed the default option as the choice endorsed by their electricity provider. Disentangling endorsement from rational inattention could be important with respect to external validity; the default effect could be less strong with a less trusted electricity supplier. However, either explanation is consistent with positive welfare gains in this particular context.

7.2 Cost-Effectiveness

Thus far, we have analyzed the empirical evidence from the perspective of the household. We now evaluate program outcomes from the perspective of the utility. More precisely, we investigate whether the default-induced demand response confers benefits to the utility that offset the additional program costs.

Table 8 compares the costs of enrolling participants and implementing the program against the benefits (i.e., costs avoided when peak consumption is reduced). Our analysis in Table 8 assumes each pricing program was scaled to SMUD’s entire residential customer base and run

for 10 years.³⁴

[TABLE 8 HERE]

The two columns on the left summarize the two main benefits of the program. Reduced demand during CPP and TOU peak hours avoids two types of expenses: the costs incurred to supply sufficient electricity to meet peak demand during these hours, and the expected cost of new investments in peaking plants needed to meet demand in peak hours. To estimate avoided capacity investment costs, the expectation is taken over the probability that demand in CPP or TOU hours would drive capacity expansion decisions. Notably, the avoided energy costs are considerably smaller than the avoided capacity costs, particularly for the CPP programs. This reflects the fact that a small number of peak hours drives costly generating capacity expansions. Reducing demand in peak periods avoids the need to construct and maintain these “peaker” plants.³⁵

We break the program costs into three components: (1) one-time fixed costs, which include items such as IT costs to adjust the billing system and initial program design costs, (2) one-time per-household costs which primarily include the customer acquisition costs, including the in-home devices offered to customers as part of the recruitment, and (3) recurring annual fixed and variable costs, which include personnel costs required to administer the program. The one-time variable cost of recruiting customers is lower under the opt-out programs than under the opt-in.

Net benefits are reported in the final column of Table 8. We estimate that both opt-out programs would be cost-effective with net benefits to the utility in excess of \$55 and \$23 million for the CPP and TOU programs, respectively. The CPP opt-in program is estimated to be marginally cost-effective. The TOU opt-in program, which led to much smaller demand reductions than the CPP program, is projected to incur costs in excess of savings. In other words, in the case of

34. Some of these program benefits and costs are summarized in Potter, George, and Jimenez (2014), a consulting report prepared to help SMUD decide whether to expand the pilot. We obtained additional information from personal communications with SMUD and their consultants. Section 1.3 summarizes underlying assumptions, and explains why some of the assumptions pertaining to program benefits are likely conservative.

35. As we explain in Section 1.6, the calculations reflected in Table 8 may understate the capacity benefits, for example because they do not measure reductions in transmission- and distribution-level investments. Because the numbers in Table 8 reflect private benefits to the utility, they do not incorporate the value of avoided pollution. Given that the avoided energy savings benefits are low relative to the avoided capacity, we suspect that avoided pollution would not change the overall cost-benefit calculus by much.

the TOU program, the default manipulation turns a cost-ineffective program into a cost-effective endeavour from the utility's perspective.

8 Conclusion

The default effect is one of the most powerful and consistent behavioral phenomena in economics, with examples documented across many settings, including health care, personal finance and internet marketing. This paper studies this phenomenon in a new context – time-varying pricing programs for electricity. Residential customers served by a large municipal utility in the Sacramento area were randomly allocated 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 but allowed to opt out; and (3) a control group. We document 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.

This empirical setting offers several innovations relative to the existing literature on default effects. In addition to observing the initial decision that was directly manipulated by the default effect, we also collect detailed data on follow-on behavior. We distinguish between follow-on behavior that modifies the original choice, such as opting out of the time-varying pricing program once it has begun, and behavior that is conditional on, but distinct from, the original choice. In our case, the latter involves adjusting electricity consumption in response to time-varying electric prices. We argue that this conditional behavior can be equally, if not more, important than the original choice for two reasons: (1) it is observation of the follow-on behavior that enables us to assess competing explanations for the default effect and in turn draw qualitative conclusions about the welfare implications of defaulting people onto these time-varying pricing programs; and (2) societal and grid benefits from such a program are critically contingent on whether consumers join the program, and if they change their electricity consumption in response to the

program conditional on joining. We find that consumers do adjust electricity consumption in response to the time-varying prices, even if they did not actively select them.


An additional innovation of this work results from analyzing systematic differences in the initial participation choice and follow-on behavior across different groups. Notably, we find that customers who should expect to have lower bills on the program without changing their behavior (so-called “structural winners”) were no more likely to enroll in the program. This observation underpins our assessment of competing explanations for the default effect. This pattern is inconsistent with explanations for the default effect that assume consumers perform well-informed, cost-benefit calculations before making their choice. We find that expected gains from making a fully attentive choice are small, on average, relative to the effort that is presumably required to make those calculations. We show how a choice model that accommodates inattention can rationalize the default effects we observe. Because Sacramento is a particularly well-regarded utility with high customer satisfaction, consumers in our study might assume the utility had their interests in mind when choosing the default. This could reduce the likelihood that consumers will attend to this choice. However, even without these potential endorsement effects, we show that the costs of inattention are low in this setting, suggesting that inattention to the participation choice is plausibly rational.

Once defaulted, passive consumers (i.e., consumers who would not have actively enrolled in the pricing program but did not opt out) are sufficiently attentive to time-varying pricing to mount an economically significant response on average. Moreover, as passive customers gain experience with the new pricing regime, their average demand response increases. We see convergence between active joiners and passive consumers in the second year of the program, which we take as evidence that nudged consumers acclimated to the new pricing regimes. We expect that future work can similarly use follow-on behavior to draw inferences about default effects.

In sum, we find that placing households onto time-varying pricing by default can lead to significantly more customers on time-varying pricing and, more importantly, significantly higher aggregate responses to price changes. Inattention to this choice appears rational from the per-


spective of a household; the welfare implications of the default effect are small for an individual customer. Aggregated across households, the social benefits associated with higher participation in demand response are substantial.

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




Note: This figure shows fluctuations of hourly electricity demand and wholesale spot prices over a week in June, 2011. Wholesale spot prices reported by the California independent system operator (CAISO).

Figure 2: Electricity rate structures



Notes: This figure shows SMUD electricity rate structures in place during the treatment period. 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 called per summer on weekdays during the hours of 4 PM and 7 PM. 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 efforts



Notes: This figure shows pre-period encouragement efforts and enrollment proportion. For opt-out groups, vertical lines indicate dates on which packets were mailed out to the households. For opt-in groups, 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, and the final solid vertical line depicts distribution of door hangers on March 1st, 2012. Gray vertical lines between April 4th and June 1st, 2012 indicate the phone bank campaign, when calls went out on almost a daily basis. The solid decreasing (increasing) lines in each figure represent the proportion of households in the opt-out (opt-in) group that remained enrolled (chose to enroll) in treatment over the course of the recruitment efforts.

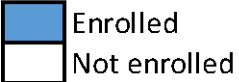
Table 1: Comparison of means by treatment assignment

	Control group	Treatment groups			
		CPP		TOU	
		Opt-in	Opt-out	Opt-in	Opt-out
Daily usage (kWh)	26.63	26.81 (-0.82)	26.92 (-0.45)	26.49 (0.83)	26.38 (0.71)
Peak to off-peak ratio	1.77	1.77 (0.02)	1.78 (-0.50)	1.78 (-0.57)	1.78 (-0.37)
Bill amount (\$)	109.10	109.44 (-0.34)	109.12 (-0.01)	108.20 (1.08)	107.86 (0.69)
Structural winner (CPP)	0.51	0.51 (-0.51)	0.52 (-0.39)	0.51 (-0.11)	0.50 (0.70)
Structural winner (TOU)	0.34	0.34 (-0.13)	0.35 (-0.15)	0.34 (0.41)	0.33 (1.14)
Households	45,839	9,190	846	12,735	2,407

Notes: This table compares pre-period usage statistics across control and treatment groups. Cells contain group means and t-statistics (in parentheses) obtained from a two-sample t-test comparing means in the control group to means in the given treatment group. Daily usage is the average per-customer electricity usage over the pre-period summer. Peak to off-peak ratio is the average hourly consumption during peak periods (4-7pm on weekdays) divided by the hourly kWh used during non-peak times over the pre-period summer. Bill amounts reflect monthly bills over the pre-period summer. Structural winner is an indicator variable for whether the household would have experienced reduced bills in the pre-period summer had they been enrolled in either the CPP or TOU pricing plans.

Figure 4: Identification of active joiners, passive consumers, and active leavers

	Active leavers	Active joiners	Passive consumers
Opt-out experimental design	Not enrolled	Enrolled	Enrolled
Opt-in experimental design	Not enrolled	Enrolled	Not enrolled
Control group	Not enrolled	Not enrolled	Not enrolled



Notes: This figure describes enrollment choice of different customer types by experimental group. Rows indicate the three groups into which customers in our sample were randomly assigned: opt-out, opt-in, and control. Columns signify types of customers (active leavers, active joiners, and passive consumers). Shading indicates that the customer type enrolls in time-varying pricing program under the associated experimental group.

Table 2: Participation rates

	Initial		Endline		Attrition
	Proportion	Count	Proportion	Count	Change
CPP opt-in (AJ)	0.201	1,568	0.189	1,169	0.057
CPP opt-out (AJ + PC)	0.960	701	0.894	537	0.070
TOU opt-in (AJ)	0.193	2,088	0.181	1,551	0.062
TOU opt-out (AJ + PC)	0.979	2,019	0.926	1,507	0.055

Notes: This table describes participation by experimental group. AJ stands for active joiners, PC stands for passive consumers. Proportions are the count of enrolled customers divided by the count of total customers in each group at a given point in time. Counts include only customers who have not moved away by a given point in time. Initial participation reflects the beginning of the treatment period (June 1st, 2012), while endline participation reflects the end of the treatment period (September 30th, 2013). An enrolled customer is one who entered the program (either by opting in or by being defaulted in) and did not opt-out before the given date. Attrition is the percentage change between initial and end-line participation proportions (where a value of 0.057 signifies a percent change of 5.7 percent).

Table 3: Average effects for encouraged groups

	Critical event		Non-event peak	
	Opt-in	Opt-out	Opt-in	Opt-out
Encouragement (CPP)	-0.129*** (0.010)	-0.305*** (0.037)	-0.029*** (0.006)	-0.094*** (0.020)
Mean usage (kW)	2.49	2.5	1.8	1.8
Customers	55,028	46,684	55,028	46,684
Customer-hours	4,832,874	4,104,263	31,198,201	26,495,612
Encouragement (TOU)	-0.091*** (0.008)	-0.130*** (0.019)	-0.054*** (0.006)	-0.100*** (0.013)
Mean usage (kW)	2.49	2.49	1.79	1.79
Customers	58,573	48,245	58,573	48,245
Customer-hours	5,141,976	4,240,163	33,195,961	27,374,276

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Notes: This table presents estimates of the impact of encouragement assignment on average hourly electricity usage in kilowatts, irrespective of enrollment status. To estimate the critical event hour effects, data include 4-7pm during simulated CPP events in 2011 (hottest 12 non-holiday weekdays) and 4-7pm during actual 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 and hour-of-sample fixed effects. Standards errors are clustered by customer.


Table 4: Average effects for treated households

	Critical event hours			Non-event day peak hours		
	Opt-in (AJ)	Opt-out (AJ+PC)	Passive (PC)	Opt-in (AJ)	Opt-out (AJ+PC)	Passive (PC)
Treatment (CPP)	-0.658*** (0.051)	-0.330*** (0.040)	-0.242*** (0.053)	-0.146*** (0.031)	-0.101*** (0.022)	-0.089*** (0.028)
Mean usage (kW)	2.49	2.50	2.44	1.80	1.80	1.79
Customers	55,028	46,684	10,036	55,028	46,684	10,036
Customer-hours	4,832,874	4,104,263	880,075	31,198,201	26,495,612	5,679,023
Treatment (TOU)	-0.480*** (0.044)	-0.136*** (0.020)	-0.051* (0.027)	-0.287*** (0.029)	-0.105*** (0.014)	-0.059*** (0.018)
Mean usage (kW)	2.49	2.49	2.43	1.79	1.79	1.75
Customers	58,573	48,245	15,142	58,573	48,245	15,142
Customer-hours	5,141,976	4,240,163	1,325,077	33,195,961	27,374,276	8,555,447

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Notes: This table presents estimates of the impact of enrollment on average hourly electricity usage in kilowatts. AJ stands for active joiners, PC stands for passive consumers. The sample for critical event hours includes hours between 4pm and 7pm during simulated CPP events in 2011 (hottest 12 non-holiday weekdays between June and September) and actual CPP events in 2012-2013. Sample for non-event day peak hours include hours between 4pm and 7pm of non-holiday, non-CPP event weekdays during the 2011-2013 summers (June to September). Opt-in and opt-out effects are estimated by comparing the opt-in and opt-out experimental groups, respectively, to the control group. Passive consumer effects are estimated 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 into treatment used as an instrument for treatment enrollment. All regressions include customer and hour-of-sample fixed effects. Standard errors are clustered at the customer level.

Figure 5: Event day effects for treated households by hour



Notes: This figure depicts hourly impacts of enrollment on electricity usage in kilowatts during event days. The sample includes simulated CPP events in 2011 (hottest 12 non-holiday weekdays between June and September) and actual CPP events in 2012-2013. Opt-in and opt-out effects are estimated by comparing the opt-in and opt-out experimental groups, respectively, to the control group. Passive consumer effects are estimated 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 into treatment used as an instrument for treatment enrollment. Dashed lines indicate the 95 percent confidence interval of the estimates with standard errors clustered by customer. The vertical bars indicate the peak period, between 4pm and 7pm.

Table 5: Average bill impacts for treated households

	Opt-in (AJ)	Opt-out (AJ+PC)	Passive (PC)
Treatment (CPP)	-6.515*** (2.358)	-4.499*** (1.428)	-3.121** (1.485)
Mean bill (\$)	114	114	114
Customers	55,029	46,685	10,036
Customer-months	552,087	468,843	100,552
Treatment (TOU)	-2.816 (2.196)	-1.985** (0.872)	-1.423 (0.935)
Mean bill (\$)	114	114	113
Customers	58,574	48,246	15,142
Customer-months	587,406	484,364	151,392

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Notes: This table documents the impact of treatment enrollment on monthly bills. The sample is composed of summer months. AJ stands for active joiners, PC stands for passive consumers. Opt-in and opt-out effects are estimated by comparing the opt-in and opt-out experimental groups, respectively, to the control group. Passive consumer effects are estimated 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 into treatment used as an instrument for treatment enrollment. All regressions include customer and month-of-sample fixed effects. Standard errors are clustered at the customer level.


Table 6: Usage impacts vary by year of program

	Critical event hours			Non-event day peak hours		
	Opt-in (AJ)	Opt-out (AJ+PC)	Passive (PC)	Opt-in (AJ)	Opt-out (AJ+PC)	Passive (PC)
Treatment (CPP)	-0.714*** (0.054)	-0.298*** (0.043)	-0.186*** (0.056)	-0.161*** (0.031)	-0.079*** (0.022)	-0.057** (0.029)
× Year 2	0.126** (0.054)	-0.069* (0.037)	-0.124** (0.049)	0.036 (0.035)	-0.051** (0.023)	-0.075** (0.030)
Treatment (TOU)	-0.545*** (0.046)	-0.156*** (0.021)	-0.058** (0.028)	-0.310*** (0.029)	-0.112*** (0.014)	-0.062*** (0.018)
× Year 2	0.146*** (0.049)	0.044** (0.020)	0.017 (0.027)	0.056* (0.033)	0.018 (0.013)	0.007 (0.017)

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Notes: This table presents estimates of the treatment effects separately for each year of the program. Year 2 refers to the the second year of treatment period, 2013. For the first, second, fourth and fifth columns regressors are instrumented with indicators for encouragement group and its interaction with the indicator variable for structural winners. The sample for these four columns is composed of the control group and given treatment group. For the third and sixth columns, the instruments are enrollment into opt-out group and its interaction with the indicator variable for Year 2 and the sample includes only opt-in and opt-out treatment groups. Event hours include 4 to 7 PM on simulated critical peak event days in 2011 and actual event days in 2012 and 2013. Non-event peak day hours include all peak hours excluding critical event hours. All models include customer and hour of sample fixed effects. Standard errors are clustered at the customer level.

Figure 6: Program participation by structural gains



Notes: Figure documents the distribution of structural gains and the relationship between structural gains and program participation on June 1, 2012 (the start of the experiment) for both CPP and TOU time-varying pricing treatments. Bottom panels: distributions of structural gains. Top panels: points are the proportion of households participating at the start of the program in each bin, lines and confidence intervals are the fitted values and confidence intervals for the prediction of a regression of participation on structural gains.

Table 7: Returns on attention (CPP)

	Opt-in	Opt-out
Participation	20%	96%
Passive customers	76%	
Returns on attention (\$)		
<i>Switching cost c = \$10</i>		
Uninformed prior	$\mu = 3.53, \sigma = 7.73$	
No adjustment	13.51	3.76
With adjustment	31.03	2.70
<i>Switching cost c = \$20</i>		
Uninformed prior	$\mu = 7.05, \sigma = 15.45$	
No adjustment	11.57	2.92
With adjustment	24.37	3.00

Notes: This table summarizes our estimates of the distribution of uninformed priors and the associated returns on attention for customers in the CPP group. Participation indicates the observed initial participation percentages in the time-varying pricing program. In the second panel, uninformed priors are the means and standard deviations of the normal distribution that rationalizes the observed participation in the opt-in and opt-out groups, given the assumed switching costs. We simulate returns on attention, or the average value of becoming informed about their structural gains (and possibly changing their enrollment choice) across all customers. The “No adjustment” assumes that customers can become informed about their structural gains under time-varying pricing. “With adjustment” assumes that customers both become informed about their structural gains and anticipate their own changes in energy consumption in response to time-varying pricing. Under both heuristics, we assume a discount rate of 5%.

Table 8: Cost-effectiveness

	Benefits		Costs			10-year NPV	Benefits - Costs
	Avoided Capacity	Avoided Energy	One-time Fixed Costs	One-time Variable Costs	Recurring Annual Total Costs		
CPP opt-in	44.0	0.9	1.4	31.0	0.9	36.5	8.4
CPP opt-out	92.1	2.1	1.4	21.0	3.1	38.8	55.4
TOU opt-in	27.0	5.0	0.8	30.0	0.5	32.5	-0.5
TOU opt-out	41.8	7.3	0.8	18.5	1.3	26.1	23.0

Notes: This table presents the estimates of the cost-effectiveness for each treatment group. All figures are in millions of dollars and assume the program is scaled to SMUD’s whole residential customer base and run for 10 years. See Appendix 6 for details.

References

- Abadie, Alberto, and Sebastien Gay. 2006. "The impact of presumed consent legislation on cadaveric organ donation: A cross-country study." *Journal of Health Economics* 25 (4): 599–620.
- Allcott, Hunt, and Judd B. Kessler. 2019. "The Welfare Effects of Nudges: A Case Study of Energy Use Social Comparisons." *American Economic Journal: Applied Economics* 11 (1): 236–276.
- Attari, Shahzeen Z., Michael L. DeKay, Cliff I. Davidson, and Wändi Bruine de Bruin. 2010. "Public perceptions of energy consumption and savings." *Proceedings of the National Academy of Sciences* 107 (37): 16054–16059. eprint: <https://www.pnas.org/content/107/37/16054.full.pdf>.
- Bernheim, B. Douglas, Andrey Fradkin, and Igor Popov. 2015. "The Welfare Economics of Default Options in 401 (k) Plans." *The American Economic Review* 105 (9): 2798–2837.
- Blumenstock, Joshua, Michal Callen, and Tarek Ghani. 2018. "Why Do Defaults Affect Behavior? Experimental Evidence from Afghanistan." *American Economic Review* 108:2868–2901.
- Boiteux, Marcel. 1964a. "Marginal Cost Pricing." In *Marginal Cost Pricing in Practice*, edited by J. R. Nelson, 51–58. Englewood, NJ: Prentice-Hall.
- . 1964b. "The Choice of Plant and Equipment for the Production of Electricity." In *Marginal Cost Pricing in Practice*, edited by J. R. Nelson, 199–214. Englewood, NJ: Prentice-Hall.
- Borenstein, Severin, and Stephen Holland. 2005. "On the Efficiency of Competitive Markets with Time-Invariant Retail Prices." *RAND Journal of Economics* 36 (3): 469–93.
- Brown, Jeffrey R., Anne M. Farrell, and Scott J. Weisbenner. 2016. "Decision-making Approaches and the Propensity to Default: Evidence and Implications." *Journal of Financial Economics* 121 (3): 477–495.
- Cappers, Peter, and Steven Sheer. 2016. *American Recovery and Reinvestment Act of 2009: Final Report on Customer Acceptance, Retention, and Response to Time-Based Rates from Consumer Behavior Studies*. LBNL-1007279. LBNL.

- Carroll, Gabriel D., James J. Choi, David Laibson, Brigitte C. Madrian, and Andrew Metrick. 2009. “Optimal Defaults and Active Decisions.” *Quarterly Journal of Economics* (November): 1639–1674.
- Chetty, Raj, John N. Friedman, Søren Leth-Petersen, Torben Heien Nielsen, and Tore Olsen. 2014. “Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from Denmark.” *The Quarterly Journal of Economics* 129 (3): 1141–1219.
- Choi, James J., David Laibson, Brigitte C. Madrian, and Andrew Metrick. 2002. “Defined contribution pensions: Plan rules, participant choices, and the path of least resistance.” In *Tax Policy and the Economy, Volume 16*, edited by James M. Poterba, 67–114. Cambridge, MA: MIT Press.
- . 2004. “For Better or for Worse: Default Effects and 401 (K) Savings Behavior.” In *Perspectives on the Economics of Aging*, 81–126. University of Chicago Press.
- Choukhmane, Taha. 2018. “Default Options and Retirement Saving Dynamics.”
- Cooper, Adam, and Mike Shuster. 2019. *Electric Company Smart Meter Deployments: Foundation for a Smart Grid*. Technical report. The Edison Foundation, Institute for Electric Innovation.
- DellaVigna, Stefano. 2009. “Psychology and Economics: Evidence from the Field.” *Journal of Economic Literature* 47:315–372.
- Department of Energy. 2012. *Smart Grid Investment Grant Program*. Staff Report. United States Department of Energy, July.
- EIA. 2018. *Annual Electric Power Industry Report*. Data retrieved from <https://www.eia.gov/electricity/data/eia861/>.
- EPRI. 2012. *Understanding Electric Utility Consumers - Summary Report: What We Know and What We Need to Know*. Final Report. Electric Power Research Institution (EPRI), October.
- Frederick, Shane, George Loewenstein, and Ted O’Donoghue. 2002. “Time discounting and time preference: A critical review.” *Journal of Economic Literature* 40 (2): 351–401.

- Goldin, Jacob, and Daniel Reck. 2019. "Revealed Preference Analysis with Framing Effects." *Journal of Political Economy*: forthcoming.
- Gottlieb, Daniel, and Kent Smetters. 2019. *Lapse-Based Insurance*.
- Handel, Benjamin R. 2013. "Adverse selection and inertia in health insurance markets: When nudging hurts." *The American Economic Review* 103 (7): 2643–2682.
- Handel, Benjamin R., and Jonathan T. Kolstad. 2015. "Health insurance for "humans": Information frictions, plan choice, and consumer welfare." *The American Economic Review* 105 (8): 2449–2500.
- Jimenez, Lupe, Jennifer Potter, and Stephen George. 2013. *SmartPricing Options Interim Evaluation*. https://www.smartgrid.gov/files/MASTER_SMUD_CBS_Interim_Evaluation_Final_SUBMITTED_TO_TAG_20131023.pdf, accessed February 2020. U.S. Department of Energy and Lawrence Berkeley National Laboratory.
- Johnson, Eric J., Steven Bellman, and Gerald L. Lohse. 2002. "Defaults, Framing and Privacy: Why Opting In-Opting Out." *Marketing Letters* 13, no. 1 (February): 5–15.
- Johnson, Eric J., and Daniel Goldstein. 2003. "Do Defaults Save Lives?" *Science* 302, no. 5649 (November): 1338–1339.
- Johnson, Eric J., John Hershey, Jacqueline Meszaros, and Howard Kunreuther. 1993. "Framing, probability distortions, and insurance decisions." *Journal of Risk and Uncertainty* 7:35–51.
- Ketcham, Jonathan D, Nicolai V Kuminoff, and Christopher A Powers. 2016. "Estimating the Heterogeneous Welfare Effects of Choice Architecture: An Application to the Medicare Prescription Drug Insurance Market." *NBER Working Paper*.
- Kowalski, Amanda E. 2016. "Doing More When You're Running LATE: Applying Marginal Treatment Effect Methods to Examine Treatment Effect Heterogeneity in Experiments." *NBER Working Paper*.

- Madrian, Brigitte C., and Dennis F. Shea. 2001. "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior." *The Quarterly Journal of Economics* 116 (4): 1149–1187.
- Myers, Erica, Steven L Puller, and Jeremy D West. 2019. *Effects of Mandatory Energy Efficiency Disclosure in Housing Markets*. Working Paper, Working Paper Series 26436. National Bureau of Economic Research, November.
- Park, C. Whan, Sung Youl Jun, and Deborah 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.
- Potter, Jennifer M., Stephen S. George, and Lupe R. Jimenez. 2014. *SmartPricing Options Final Evaluation*. Technical report.
- Sallee, James M. 2014. "Rational Inattention and Energy Efficiency." *The Journal of Law and Economics* 57 (3): 781–820.
- Samuelson, William, and Richard Zeckhauser. 1988. "Status quo bias in decision making." *Journal of Risk and Uncertainty* 1 (1): 7–59.
- Sitzia, Stefania, Jiwei Zheng, and Daniel John Zizzo. 2015. "Inattentive consumers in markets for services." *Theory and Decision* 79 (2): 307–332.
- Thaler, Richard H., and Cass R. Sunstein. 2009. *Nudge: Improving decisions about health, wealth, and happiness*. Penguin.
- Todd-Blick, Annika, C. Anna Spurlock, Ling Jin, Peter Cappers, Sam Borgeson, Daniel Fredman, and Jarett Zuboy. 2020. "Winners are not keepers: Characterizing household engagement, gains, and energy patterns in demand response using machine learning in the United States." *Energy Research and Social Science* 70 (101595).
- Wang, Wenjie, and Takanori Ida. 2017. "Default Effect Versus Active Decision: Evidence from a Field Experiment in Los Alamos." *Kyoto University Discussion Paper No. E-14-010*.

1 Appendix

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

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

Contents

- 1.1 Load Shape Balance across Treatment Groups 1
- 1.2 Alternative Specifications 2
- 1.3 Assumptions Underlying the LATE Estimates 2
- 1.4 Modeling Attrition out of the Program 5
- 1.5 Evidence on Alternative Mechanisms 6
 - 1.5.1 Heterogeneous switching costs 7
 - 1.5.2 Present-biased preferences 8
- 1.6 Cost-Benefit Analysis 10
 - 1.6.1 Benefits 10
 - 1.6.2 Costs 11
- 1.7 Customer losses from default assignment 11

1.1 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 1.1 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 1.2 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.

1.2 Alternative Specifications

Table A1 reports heterogeneity in impacts of time-varying rates by customer type, including for structural winners, low-income customers and customers that had signed up for SMUD’s online portal (“Low income” is a dummy variable indicating enrollment in the low-income rate and “My Account” is a dummy variable indicating whether or not the household had signed up to use SMUD’s online portal prior to our experiment). Tables 1.3, 1.4, and 1.5 report results similar to those in Tables 3, 4 and 6 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 treatment compared to the CPP opt-in.

Table 1.2 replicates Table 6 using data only from the balanced panel, i.e., customers who did not change their treatment status (either by opting in, opting out, or moving) during the treatment period. Table 1.6 reports results similar to those in Table 4 in the text using only post-treatment period data.

1.3 Assumptions Underlying the LATE Estimates

This section explains how we leverage our research design to estimate local average treatment effects in different sub-groups of our study sample. We use the randomly assigned encouragements (i.e., the opt-in offer and the opt-out offer, respectively) as instruments.

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:

- Active leavers (AL): Do not opt in if $Z_i = 1$. Opt out if $Z_i = 2$.
- Passive consumers (PC): Do not opt in if $Z_i = 1$. Do not opt out if $Z_i = 2$.
- Active joiners (AJ): 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 observable and unobservable 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 pricing program participation.

- **Monotonicity:** Our encouragement intervention weakly increases (and never decreases) the likelihood of participation in the pricing program. This implies that there are no defiers.

Let π^{AL} , π^{PC} , and π^{AJ} , denote the population proportions of active leavers, passive consumers, and active joiners, 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:

$$E[Y_i|Z_i = 0] = \pi^{AL} E[Y_i(0)|AL] + \pi^{PC} E[Y_i(0)|PC] + \pi^{AJ} E[Y_i(0)|AJ].$$

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

$$E[Y_i|Z_i = 1] = \pi^{AL} E[Y_i(0)|AL] + \pi^{PC} E[Y_i(0)|PC] + \pi^{AJ} E[Y_i(1)|AJ].$$

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

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

where π^{AJ} 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 AJ 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^{PC} yields:

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

The estimate of π^{PC} 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 encouragement (i.e., the dynamic price offers) increased the salience of energy use in a way that impacts energy consumption. In this scenario, potential outcomes are more accurately represented by $Y_i(D_i, Z_i)$. Taking the opt-in design as an example, the local average treatment effect among active joiners is now more accurately estimated as:

$$LATE^{AJ} = \frac{E[Y_i(0,0)|AJ] - E[Y_i(0,1)|AJ]}{\pi^{AJ}} - \frac{\Delta_{PC}}{\pi^{AJ}} - \frac{\Delta_{AL}}{\pi^{AJ}},$$

where $\Delta_{PC} = E[Y_i(0,0)|PC] - E[Y_i(0,1)|PC]$ and $\Delta_{AL} = E[Y_i(0,0)|AL] - E[Y_i(0,1)|AL]$. 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. However, we can estimate bounds on the bias from these terms by comparing consumption patterns at households that did not participate in the dynamic pricing program across encouraged and unencouraged groups. Differences in electricity consumption among non-participants across experimental groups are difficult to interpret as they compare electricity consumption across different subsets of the consumer 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 1.7 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., passive consumers and active leavers). For the opt-out experiments, we compare consumption across all households assigned to the control group and the always leavers in the encouraged group.

Some of these differences are statistically different from zero. For example, we estimate a statistically significant difference of -0.025 across encouraged non-participants and unencouraged 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 always leavers and passive consumers 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 estimate of energy reduction by the always taker group overstates the true effect by $\frac{0.025}{0.19} = 0.13$.

Our estimates of average treatment effects for passive consumers (see Tables 4 to 6 and Figure 5) assume that active joiners who actively enroll in the pricing programs under the opt-in treatment do not behave differently than active joiners who are defaulted onto the programs through the opt-out treatment. In other words, we are assuming that $E[Y_i(1, 1)|AJ] = E[Y_i(1, 2)|AJ]$. Again, we cannot verify this assumption directly, but we can use the recruit-and-delay treatment group who were encouraged to opt-in to the TOU program but only placed on the pricing schedule in 2014 after our sampling frame (i.e., delayed), rather than in 2012. (This group is introduced in footnote 8 in the main text.) This allows us, in principle, to estimate $E[Y_i(0, 0)] - E[Y_i(0, 1)|AJ]$, and provides insight on customers who actively enrolled in the program (i.e., are active joiners) but did not immediately face time-varying prices.

We re-estimate Equation (2) using this group and find that the customers who opted-in but for whom time-based pricing was delayed reduced their usage by a statistically significant 0.09 kwh during event hours and 0.08 during non-event peak hours on average during the 2012 and 2013 summers, despite experiencing an identical price schedule to the control group (see Table 1.8). If we assume that this difference is driven entirely by the recruitment encouragement, then our exclusion restriction is violated. In this case, this means that we are underestimating the magnitude of the average reduction for passive consumers. To estimate the extent of this underestimation, we return to our estimate of $LATE^{PC}$, but now allow $E[Y_i(0, 0)] \neq E[Y_i(0, 1)|AJ]$. This yields:

$$\widehat{LATE}^{PC} = \underbrace{E[Y_i(1)|PC] - E[Y_i(2)|PC]}_{LATE^{PC}} + \underbrace{\frac{\pi^{AJ}(E[Y_i(0, 0)] - E[Y_i(0, 1)|AJ])}{\pi^{PC}}}_{\text{bias}}$$

To compute the size of the bias, we use take the treatment effect in the recruit-and-deny group (row 1, columns (1) and (2) in Table 1.8) as $E[Y_i(0, 0)] - E[Y_i(0, 1)|AJ]$ and substitute in

the endline participation rates in the TOU group for π^{AJ} and π^{PC} , which are 0.163 and 0.707, respectively. The final row in Table 1.8 estimates the implied bias from this calculation for event and non-event hours.

This gives a bias of -0.023 during event hours and -0.019 for non-event peak hours. This represents an upper bound on the degree to which our estimates in columns (3) and (6) in Table 4 could understate event and non-event peak hour reductions by passive consumers.

However, another likely explanation for our finding of a recruitment or encouragement effect is that some customers were not fully informed about the delay in their start date. In order to investigate this possibility, we conduct two tests. First, we examine whether the treatment effect declined more from the first to the second summer of treatment, which would indicate that households who were previously misinformed about their 2014 start date became aware by the second summer that they were not yet on time-varying prices and reduced their energy saving behaviors accordingly. We find that this is the case: Table 1.8, column (4) shows that there is about a 59% (0.065/0.11) reduction in savings between years 1 and 2, larger than the 18% (0.056/0.31) reduction in the comparable non-delayed group seen in Table 6, third panel, fourth column. Our second test of treatment start date confusion examines whether the TOU opt-in and opt-out groups reduced usage in the two months prior to the treatment period start date on June 1, 2012. Table 1.9 documents this test, which demonstrates that both opt-in and opt-out households did reduce usage relative to the control group even before the treatment began. The opt-in households reduced slightly more, although the point estimate for the opt-in effect is not statistically significant at conventional levels. Finally, in the survey, customers in the recruit and delay group were much more likely to think (wrongly) that they were on the time-of-use rate compared to the control group.

Overall, these results suggest the estimates in columns (1) and (2) of Table 1.8 are likely an overestimate of the extent to which actively enrolling influenced energy consumption. If we use the effect estimated for the second treatment year only in columns (3) and (4), the bias is roughly halved. In any case, and as discussed in the main text, this assumption only affects our estimates of the passive consumers' behavior, and, to the extent actively enrolling leads to reductions, our main estimates understate the true complacent response.

1.4 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 1.3 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 1.12.

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 statistically 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 statistically 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 increase 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.

1.5 Evidence on Alternative Mechanisms

In this section we provide a more detailed discussion of the extent to which the participation choices we observe are consistent with alternative explanations for the default effect. In Section 1.5.1, we estimate a discrete choice model which rationalizes the default effect with switching costs or discount rates that can vary across households. Section 1.5.2 extends the model to account for time inconsistent preferences and documents empirical evidence which suggests that that this explanation cannot explain the significant default effect we document.

1.5.1 Heterogeneous switching costs

The baseline model assumes that well-informed and risk neutral consumers faced with the choice of opting into the time-varying rate will do so if the difference in utility net of switching costs is positive. For example, a consumer assigned to the opt-in group will switch if the following inequality holds:

$$B_i - s_i + \epsilon_i > 0.$$

We use household-specific estimates of discounted structural gains, a lower bound on participation benefits, to proxy for B_i . The s_i is a customer-specific switching cost. Our empirical objective is to estimate the parameters of the distribution of switching costs that are most consistent with the participation choices we observe.

Once the pilot began, non-participants could not switch into the time-varying pricing program, but participants could switch out. To the extent that customers value this option, estimated switching costs are net of option value for opt-in customers defaulted onto the standard rate. To accommodate the possibility that consumers account for this option value, we estimate cost distributions separately for the opt-in and opt-out group. If consumers see value in preserving the participation option, we should expect smaller cost estimates in the opt-in group. We will allow this cost parameter to be distributed differently across the opt-in and opt-out treatments because we cannot separate switching costs and option value which varies with the default assignment.

To identify the parameters of these switching cost distributions, we must invoke a series of assumptions. We assume that the error term ϵ_i is a mean zero, type I extreme-value random variable. We begin by assuming that switching costs are distributed normally and independent of consumer-specific benefits. We assume a discount rate (we will estimate the cost parameters under a range of discount rate assumptions). And finally, to pin down the level of the cost parameters, we implement a standard normalization that fixes the coefficient on discounted benefits to equal one. Thus, our switching cost estimates are identified relative to our measure of benefits. If households don't fully attend to benefits in their participation decision, our model specification will over-weight benefits, and this will inflate our switching cost estimates.

Table 1.14 summarizes the estimated parameters of the switching costs distribution across different choice contexts under a range of assumed discount rates. The first two rows report the estimated mean and variance of the implied switching cost distributions in the opt-in and opt-out groups, respectively. To put these estimates in perspective, recall that moving away from the default required a single phone call, text, or email. Switching instructions were clearly displayed on all marketing materials. Given the limited effort actually required to switch away from the default option, the as-if cost distributions we estimate are implausibly large in absolute value. In the case of the opt-out group, switching costs have the wrong sign. Taken literally, this would imply that the act of switching generates utility improvements. This result follows from the fact that a key assumption implicit in this participation choice model finds no empirical support in our data. Whereas the model predicts/assumes a strong positive correlation between

gains and participation, in the CPP opt-out enrollment group we observe a negative relationship. To rationalize this negative correlation, the model estimates negative average switching costs (implying a love of switching). This is the only way this rational choice model can rationalize the participation choices we observe.

In the context of behavioral welfare analysis, it can be insightful to eliminate from consideration those choices that are most susceptible to misunderstanding or characterization failure (see, for example, Bernheim and Rangel (2009)). In our setting, it is difficult to identify those choices that are clearly mistakes. But we do have a way to isolate those households who have, in the past, paid closer attention to their electricity consumption. In the years prior to the pricing pilot, customers had the option of signing up for a “My Account” program which allows them to access detailed information about their electricity consumption. Tables 1.10 and 1.11 show how the customers who have historically engaged with these pre-existing information programs are more likely to take an active choice and either opt-in or opt-out. We re-estimate Equation (??) using only the choices made by these ‘attentive’ households. Even among these informed customers, implied switching costs (reported in Table 1.14) seem implausible given that the act of switching required a simple email, call, or text.

As we note in the paper, another way to generate a large default bias using this simple choice model is to assume a significant degree of myopia on the part of households. In our setting, consumers incur switching costs immediately, whereas benefits (in the form of lower electricity costs) accrue over the two subsequent treatment years. An alternative way to specify the model is to assume a value for switching costs and search for the discount rate values that best rationalizes observed choices. We continue to assume that that households use an exponential discount function. All of our aforementioned identifying assumptions about program benefits, the statistical properties of the error term are maintained. The key difference is that we assume switching cost values and identify the discount rates that best rationalize participation choices conditional on the maintained assumptions.

Estimated discount rates are reported in Table 1.15. These large discount rates are implausibly large in most cases. This is, again, an artifact of assuming structural gains are a significant determinant of participation decisions, and forcing the coefficient on these benefits to be one (consistent with full attention) when observed choices suggest otherwise.

1.5.2 Present-biased preferences

An alternative way to explain the significant default effect is to appeal to present-biased preferences, which could induce a procrastination effect, where households that might have intended to opt-out or opt-in did not get around to doing so. In our choice context, one can assume a household is choosing between switching today, planning to switch later, or never switching at all. In addition to the exponential discount rate δ , the household may also exhibit a present-bias, parameterized by β , which additionally discounts all future periods by a constant amount. This type of discounting is also referred to in the literature as hyperbolic discounting (e.g., Frederick, Loewenstein, and O’Donoghue 2002; DellaVigna 2009).

Our setting is slightly different from the more familiar present-bias behavioral contexts where benefits to a choice start accruing once the choice is made. In our situation as customers were

presented with their choices several months ahead of the beginning of the dynamic pricing program, and any benefits they would accrue from their choices did not begin until a specific start date (June 1, 2012) no matter whether they made their decision right away or waited. Customers who were invited to opt-in faced a deadline of June 1, 2012 and if they had not opted in by that date, they were ineligible to opt in for the rest of the program. In this case, both procrastination and simple discounting in this simple model predict that customers would wait until the deadline to incur the switching cost. We modify the notation slightly so as to explicitly represent the discounting of per-period benefits b_i . Specifically, assuming the offer was made in period 0 and customers could opt in during periods 1 or 2 and the program began in period 3, opt-in customers compared:

Value of opting in during period 1: $\beta\delta^2b_i - s_i$

Value, from the perspective of period 1, of opting in during period 2: $\beta\delta^2b_i - \beta\delta s_i$

Value of not opting in: 0

Where s_i is the cost of switching away from the default choice and b_i is the benefit of having made that switch. The household will choose to delay the decision until later, although this could reflect present-bias ($0 < \beta < 1$), but is also true for $\beta = 1$ and simply reflects the household's desire to put off incurring the switching cost. This pattern will continue up until there is some binding deadline (period 2 in the example), at which point the household no longer has the option of delaying their decision and must now choose between the following options:

Value of opting in during period 2: $\beta\delta b_i - s_i$

Value of not opting in: 0

At this point, households will switch if $\beta\delta b_i - s_i > 0$.

Customers who could opt out did not face a deadline and could continue to opt out as the program ran. At that point, they faced a more commonly modeled present-bias situation where their choice to opt-out would result in an immediate reversion back to the standard rate, so for this group, the fact that many people who would have been better off opting out (structural losers) did not could reflect procrastination. Specifically, during the program, opt-out customers compared:

Value of opting out today: $\beta\delta b_i - s_i$

Value, from the perspective of today, of opting out tomorrow: $\beta\delta^2b_i - \beta\delta s_i$

Value of not switching: 0

As long as $\beta\delta^2b_i - \beta\delta s_i > 0 \Rightarrow \delta b_i - s_i > 0$ and $\beta_1 < 1$, the household will choose to delay the decision until later.

This model results would result in two testable hypotheses. First, opt-in households will wait until the last possible moment to switch while opt-out customers may continue to procrastinate. Second, households that face a deadline and have higher structural gains are more likely to switch, while households not faced with a deadline may never actively make a choice to switch or not, and therefore their choice to do so may not be correlated with structural gains. In our context,

however, the recruitment strategy used by the utility limits our ability to examine the first of these predictions, as the weeks leading up to the deadline involved heavy phone-banking of the opt-in group (and not the defaulted group) trying to increase their participation in the program. These phone calls could be interpreted as a reduction in s_i . This makes it hard to assess the extent to which there is any "bunching" of opt-in households joining right around the deadline. However, the second testable prediction is still something we can examine in the data.

Specifically, we consider the extent to which opt-in households, who faced a deadline, are more likely to switch if they were structural winners, while opt-out households, who did not face a deadline, may never actively make a choice to switch or not, and therefore their choice to do so may not be correlated with structural gains. The results of this assessment are discussed in the body of the paper.

1.6 Cost-Benefit Analysis

This section describes the cost-benefit calculations reported in Section 7. Many of the assumptions used in our calculations are summarized in Potter, George, and Jimenez (2014), 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.

1.6.1 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 1.4.

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 Potter, George, and Jimenez (2014). 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, Potter, George, and Jimenez 2014). 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. Potter, George, and Jimenez (2014) 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 Potter, George, and Jimenez (2014) 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 Potter, George, and Jimenez (2014), we assume a 7.1% nominal discount rate and a 4.5% real discount rate.

1.6.2 Costs


Table 8 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 Potter, George, and Jimenez (2014) 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.

1.7 Customer losses from default assignment

To better understand customer incentives, here we offer a separate discussion in which we examine the number of and degree to which customers were defaulted into a program that was costly to them. To simplify the exposition, here we take observed usage during the experimental period as given and describe how the population of customers would have faired under the non-default rate (the time-varying rate and the flat rate for the opt-in and opt-out groups, respectively).


Table 1.18 summarizes these estimates. Under CPP, a higher proportion of customers lost money due to the default in the opt-in group than the opt-out group, while the opposite was the case for the TOU group. Average losses were higher on average than losses from the opt-out in both the CPP and TOU groups.

Figure 1.1: Pre-treatment electricity usage



Notes: Figure depicts average pre-treatment weekday electricity usage in kW. Panels plot average treatment group hourly electricity consumption overlaid with control group consumption, with coefficients and standard errors clustered by household obtained from a regression of electricity consumption on a set of indicator variables for each hour. Dashed lines indicate 95% confidence intervals.

Figure 1.2: Difference between treatment and control groups' electricity consumption prior to treatment



Notes: Figure depicts average difference in pre-treatment weekday electricity usage in kW between treatment and control groups. Lines represent regression coefficients from interactions between hourly indicator variables and a treatment indicator. Dashed lines indicate 95% confidence intervals, clustered by household. Vertical bars indicate peak period, between 4pm and 7pm.

Table 1.1: Usage impacts vary by customer characteristics

	Critical event hours			Non-event day peak hours		
	Opt-in (AJ)	Opt-out (AJ+PC)	Passive (PC)	Opt-in (AJ)	Opt-out (AJ+PC)	Passive (PC)
<i>Structural winner</i>						
Treatment (CPP)	-0.675*** (0.071)	-0.350*** (0.054)	-0.183*** (0.057)	-0.063 (0.042)	-0.058** (0.027)	-0.036 (0.028)
× Structural winner	0.036 (0.100)	0.039 (0.079)	-0.172 (0.121)	-0.172*** (0.063)	-0.086** (0.043)	-0.153** (0.067)
Treatment (TOU)	-0.414*** (0.050)	-0.100*** (0.023)	-0.022 (0.030)	-0.252*** (0.033)	-0.085*** (0.016)	-0.044** (0.021)
× Structural winner	-0.190* (0.098)	-0.108** (0.047)	-0.087 (0.062)	-0.099 (0.065)	-0.061* (0.032)	-0.048 (0.042)
<i>Low income</i>						
Treatment (CPP)	-0.815*** (0.066)	-0.370*** (0.047)	-0.267*** (0.060)	-0.181*** (0.040)	-0.096*** (0.025)	-0.075** (0.032)
× Low income	0.543*** (0.098)	0.176** (0.089)	0.104 (0.125)	0.122** (0.062)	-0.023 (0.051)	-0.076 (0.072)
Treatment (TOU)	-0.547*** (0.056)	-0.148*** (0.024)	-0.061** (0.031)	-0.321*** (0.037)	-0.111*** (0.017)	-0.063*** (0.021)
× Low income	0.227*** (0.086)	0.055 (0.043)	0.051 (0.061)	0.117** (0.057)	0.026 (0.030)	0.020 (0.042)
<i>My Account</i>						
Treatment (CPP)	-0.600*** (0.080)	-0.225*** (0.045)	-0.151*** (0.056)	-0.152*** (0.049)	-0.077*** (0.026)	-0.063** (0.032)
× My Account	-0.108 (0.104)	-0.251*** (0.085)	-0.238** (0.117)	0.012 (0.063)	-0.057 (0.046)	-0.067 (0.062)
Treatment (TOU)	-0.336*** (0.070)	-0.080*** (0.024)	-0.032 (0.030)	-0.204*** (0.046)	-0.065*** (0.017)	-0.039* (0.021)
× My Account	-0.274*** (0.089)	-0.143*** (0.043)	-0.055 (0.059)	-0.157*** (0.059)	-0.099*** (0.030)	-0.058 (0.040)

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Notes: Table estimates treatment impacts separately for structural winners, low income customers and My Account holders. Structural winners are customers predicted to experience savings under the time-varying rate assuming their energy consumption during the treatment period is identical to their consumption in the pre-period. Low income is an indicator variable for customers enrolled in the low income rate. My Account indicates whether if the customer has enrolled in the online My Account program. For columns 1, 2, 4, and 5, regressors are instrumented with indicators for encouragement group and its interaction with the indicator variable for structural winners. Sample for columns 1, 2, 4, and 5 is composed of the control group and given treatment group. For columns 3 and 6, the instruments are enrollment into opt-out group and its interaction with the indicator variable for structural winners and sample includes only opt-in and opt-out treatment groups. Event hours include simulated critical peak events in 2011 and actual events in 2012 and 2013. Non-event peak day hours include all peak hours excluding critical event hours. All models include customer and hour of sample fixed effects, plus an interaction between the post-treatment period and given dimension of heterogeneity. Standard errors clustered by customer in parentheses.

Table 1.2: Usage impacts vary by year of program (balanced panel)

	Critical event hours			Non-event day peak hours		
	Opt-in (AJ)	Opt-out (AJ+PC)	Passive (PC)	Opt-in (AJ)	Opt-out (AJ+PC)	Passive (PC)
Treatment (CPP)	-0.775*** (0.062)	-0.346*** (0.048)	-0.238*** (0.061)	-0.163*** (0.036)	-0.092*** (0.025)	-0.075** (0.032)
× Year 2	0.144** (0.056)	-0.051 (0.037)	-0.100** (0.048)	0.038 (0.037)	-0.043* (0.022)	-0.063** (0.029)
Treatment (TOU)	-0.605*** (0.053)	-0.168*** (0.023)	-0.069** (0.030)	-0.351*** (0.034)	-0.118*** (0.015)	-0.065*** (0.020)
× Year 2	0.161*** (0.051)	0.053*** (0.020)	0.028 (0.026)	0.061* (0.033)	0.020 (0.012)	0.010 (0.016)

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Notes: Replicates Table 6 with a balanced panel, i.e., including only households who did not change their enrollment status (by opting in, opting out, or moving) during the treatment period.

Table 1.3: Intent to treat effects (logged outcome)

	Critical event		Non-event peak	
	Opt-in	Opt-out	Opt-in	Opt-out
Encouragement (CPP)	-0.083*** (0.007)	-0.173*** (0.022)	-0.021*** (0.005)	-0.055*** (0.014)
Mean usage (kW)	2.5	2.5	1.8	1.8
Customers	55,024	46,680	55,028	46,684
Customer-hours	4,824,157	4,097,167	31,141,456	26,448,932
Encouragement (TOU)	-0.052*** (0.005)	-0.073*** (0.012)	-0.036*** (0.004)	-0.059*** (0.010)
Mean usage (kW)	2.49	2.5	1.79	1.79
Customers	58,569	48,241	58,573	48,245
Customer-hours	5,133,166	4,232,869	33,137,047	27,326,082

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Notes: Replicates Table 3 with $\log(\text{Usage})$ as outcome variable, coefficients are proportion change in consumption.

Table 1.4: Average treatment effects (logged outcome)

	Critical event hours			Non-event day peak hours		
	Opt-in (AJ)	Opt-out (AJ+PC)	Passive (PC)	Opt-in (AJ)	Opt-out (AJ+PC)	Passive (PC)
Treatment (CPP)	-0.424*** (0.032)	-0.187*** (0.024)	-0.124*** (0.031)	-0.106*** (0.024)	-0.059*** (0.015)	-0.046** (0.020)
Mean usage (kW)	2.50	2.50	2.44	1.80	1.80	1.79
Customers	55,024	46,680	10,036	55,028	46,684	10,036
Customer-hours	4,824,157	4,097,167	878,222	31,141,456	26,448,932	5,667,680
Treatment (TOU)	-0.275*** (0.028)	-0.077*** (0.012)	-0.028* (0.016)	-0.190*** (0.022)	-0.062*** (0.010)	-0.030** (0.013)
Mean usage (kW)	2.49	2.50	2.44	1.79	1.79	1.75
Customers	58,569	48,241	15,142	58,573	48,245	15,142
Customer-hours	5,133,166	4,232,869	1,322,933	33,137,047	27,326,082	8,540,421

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Notes: Replicates Table 4 with $\log(\text{Usage})$ as outcome variable, coefficients are proportion change in consumption.

Table 1.5: Usage impacts vary by customer observables and year of program (logged outcome)

	Critical event hours			Non-event day peak hours		
	Opt-in (AJ)	Opt-out (AJ+PC)	Passive (PC)	Opt-in (AJ)	Opt-out (AJ+PC)	Passive (PC)
<i>Structural winner</i>						
Treatment (CPP)	-0.425*** (0.043)	-0.227*** (0.033)	-0.124*** (0.038)	-0.043 (0.033)	-0.044** (0.020)	-0.029 (0.023)
× Structural winner	0.002 (0.063)	0.078* (0.046)	0.002 (0.066)	-0.130*** (0.048)	-0.030 (0.030)	-0.049 (0.044)
Treatment (TOU)	-0.250*** (0.035)	-0.066*** (0.015)	-0.020 (0.020)	-0.163*** (0.028)	-0.055*** (0.012)	-0.027* (0.016)
× Structural winner	-0.074 (0.057)	-0.031 (0.026)	-0.019 (0.035)	-0.076* (0.045)	-0.020 (0.021)	-0.008 (0.028)
<i>Year 2</i>						
Treatment (CPP)	-0.453*** (0.034)	-0.169*** (0.026)	-0.092*** (0.033)	-0.114*** (0.024)	-0.044*** (0.016)	-0.026 (0.021)
× Year 2	0.065** (0.033)	-0.041* (0.022)	-0.071** (0.029)	0.019 (0.025)	-0.034** (0.017)	-0.049** (0.022)
Treatment (TOU)	-0.305*** (0.029)	-0.083*** (0.013)	-0.027 (0.018)	-0.206*** (0.022)	-0.065*** (0.010)	-0.030** (0.013)
× Year 2	0.066** (0.031)	0.013 (0.013)	-0.001 (0.017)	0.040* (0.024)	0.008 (0.010)	-0.001 (0.013)
<i>Low income</i>						
Treatment (CPP)	-0.504*** (0.040)	-0.219*** (0.028)	-0.152*** (0.035)	-0.122*** (0.030)	-0.059*** (0.017)	-0.043** (0.022)
× Low income	0.275*** (0.065)	0.136*** (0.051)	0.129* (0.072)	0.056 (0.049)	-0.003 (0.036)	-0.017 (0.052)
Treatment (TOU)	-0.306*** (0.035)	-0.084*** (0.014)	-0.035* (0.018)	-0.210*** (0.027)	-0.067*** (0.012)	-0.035** (0.015)
× Low income	0.102* (0.056)	0.032 (0.028)	0.039 (0.040)	0.069 (0.044)	0.021 (0.023)	0.025 (0.033)
<i>My Account</i>						
Treatment (CPP)	-0.386*** (0.052)	-0.139*** (0.028)	-0.089** (0.035)	-0.124*** (0.039)	-0.053*** (0.019)	-0.039 (0.024)
× My Account	-0.071 (0.065)	-0.117** (0.049)	-0.090 (0.067)	0.032 (0.049)	-0.015 (0.031)	-0.019 (0.042)
Treatment (TOU)	-0.200*** (0.044)	-0.050*** (0.015)	-0.021 (0.019)	-0.130*** (0.034)	-0.043*** (0.012)	-0.026* (0.015)
× My Account	-0.143** (0.056)	-0.070*** (0.026)	-0.019 (0.036)	-0.113** (0.044)	-0.047** (0.022)	-0.011 (0.029)

* p < 0.1, ** p < 0.05, *** p < 0.01. Standard errors in parentheses.

Notes: Replicates Tables 6 and 1.1 with log(Usage) as outcome variable, coefficients are proportion change in consumption.

Table 1.6: Intent to treat effects (post-treatment period only)

	Critical event		Non-event peak	
	Opt-in	Opt-out	Opt-in	Opt-out
Encouragement (CPP)	-0.100*** (0.023)	-0.291*** (0.061)	-0.012 (0.017)	-0.073 (0.047)
Mean usage (kW)	2.51	2.52	1.82	1.82
Customers	46,024	39,086	47,155	40,054
Customer-hours	2,855,231	2,426,418	18,751,449	15,935,568
Encouragement (TOU)	-0.089*** (0.017)	-0.193*** (0.035)	-0.061*** (0.013)	-0.143*** (0.026)
Mean usage (kW)	2.51	2.52	1.81	1.82
Customers	48,971	40,383	50,188	41,387
Customer-hours	3,037,095	2,506,122	19,947,727	16,460,956

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Notes: Replicates Table 3 using only post-treatment period data. The estimating equation is identical, except that customer-specific fixed effects are no longer included due to the exclusion of pre-treatment period data.

Table 1.7: Average effects on non-participating households

	Critical event		Non-event peak	
	Opt-in	Opt-out	Opt-in	Opt-out
Encouragement (CPP)	-0.012 (0.010)	0.027 (0.109)	-0.003 (0.007)	-0.016 (0.095)
Bound of bias	-0.06	0.00	-0.02	-0.00
Mean usage (kW)	2.52	2.52	1.80	1.79
Customers	53,381	45,867	53,381	45,867
Customer-hours	4,675,263	4,031,723	30,179,735	26,026,802
Encouragement (TOU)	-0.025*** (0.009)	0.035 (0.094)	-0.012** (0.006)	-0.089 (0.085)
Bound of bias	-0.13	0.01	-0.06	-0.01
Mean usage (kW)	2.52	2.52	1.79	1.79
Customers	56,378	45,881	56,378	45,881
Customer-hours	4,934,493	4,033,157	31,853,310	26,036,009

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Notes: Table estimates effect of encouragement on usage of households who did not enroll in treatment. Table specification similar to Table 3, but sample includes control customers and encouraged customers who did not enroll in the treatment by not opting in or opting out, depending on whether they were in the opt-in or opt-out treatments, respectively. Bound of bias rows calculate the potential bias $\frac{(1-P)}{P}\beta$ (where P is the proportion enrollment for that group) in Table 4 as a result of the estimated encouragement effects on non-enrolling customers under the assumption that selection does not bias the given estimates.

Table 1.8: Average effects on recruit-and-delay group

	Critical event	Non-event peak	Critical event	Non-event peak
Treatment (TOU)	-0.098** (0.041)	-0.083*** (0.028)	-0.137*** (0.044)	-0.111*** (0.028)
× 2013			0.085* (0.047)	0.065** (0.031)
Bound of bias	0.023	0.019	0.012	0.011
Mean usage (kW)	2.50	1.80	2.50	1.80
Customers	58,532	58,532	58,532	58,532
Customer-hours	5,140,696	33,188,035	5,140,696	33,188,035

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Notes: Table estimates impact of treatment on usage for recruit-and-delay households (RITTD). Dependent variable is usage in kwh. Sampling frame is summer weekday CPP event hours and non-event peak hours from 2011-2013 and includes control group and TOU opt-in recruit-and-delay households. Regressions include household and hour of sample fixed effects, standard errors clustered by household.

Table 1.9: April-May LATE impacts

Treatment (TOU)	-0.060 (0.039)	-0.039** (0.016)
Mean usage (kW)	1.05	1.06
Customers	52,153	42,991
Customer-hours	6,748,730	5,564,183

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Notes: Table estimates effect of treatment on pre-treatment period usage. Dependent variable is usage in kwh. Sampling frame is April and May weekday peak hours in 2012 and includes control group and the given treatment group. Regressions include hour of sample fixed effects, standard errors clustered by household.

Table 1.10: Household characteristics by customer type (means)

	AJ	PC	AL
<i>CPP customers</i>			
Daily usage	26.75	26.97	26.90
Peak to off-peak	1.77	1.78	1.79
Bill amount	106.18	109.84	112.50
Structural winner (CPP)	0.50	0.52	0.49
Structural winner (TOU)	0.35	0.34	0.33
My Account	0.54	0.42	0.52
My Account logins	9.16	6.65	11.81
Paperless	0.24	0.19	0.18
Low income	0.29	0.19	0.15
<i>TOU customers</i>			
Daily usage	26.94	26.25	27.36
Peak to off-peak	1.74	1.78	1.73
Bill amount	107.32	107.99	114.93
Structural winner (CPP)	0.53	0.50	0.57
Structural winner (TOU)	0.35	0.33	0.39
My Account	0.53	0.39	0.48
My Account logins	8.25	5.91	10.79
Paperless	0.24	0.18	0.22
Low income	0.29	0.18	0.11

Notes: Mean customer characteristics for the three customer types: active joiners (AJ), passive consumers (PC), and active leavers (AL). Means for active joiners (μ^{AJ}) and active leavers (μ^{AL}) are computed as the average value for customers who enrolled in the opt-in groups or disenrolled in the opt-out groups, respectively. Means for passive consumers (μ^{PC}) are computed using the following formula: $\mu^{OO} = p^{AJ}\mu^{AJ} + p^{PC}\mu^{PC}$, where proportions p^{AJ} and p^{PC} are the relative proportions of active joiners and passive consumers who enroll in the opt-out group, which we compute from the difference in enrollments between opt-in and opt-out groups.

Table 1.11: Household characteristics by customer type (differences)

	AJ - PC	AJ - AL	PC - AL
<i>CPP customers</i>			
Daily usage	-0.22 (0.81)	-0.15 (0.95)	0.07 (0.98)
Peak to off-peak	-0.01 (0.62)	-0.02 (0.76)	-0.00 (0.99)
Bill amount	-3.66 (0.42)	-6.32 (0.59)	-2.66 (0.83)
Structural winner (CPP)	-0.02 (0.34)	0.01 (0.87)	0.03 (0.51)
Structural winner (TOU)	0.01 (0.73)	0.03 (0.54)	0.02 (0.70)
My Account	0.13 (0.00)	0.02 (0.64)	-0.11 (0.03)
My Account logins	2.51 (0.00)	-2.65 (0.30)	-5.17 (0.04)
Paperless	0.05 (0.02)	0.06 (0.12)	0.01 (0.80)
Low income	0.10 (0.00)	0.13 (0.00)	0.04 (0.32)
<i>TOU customers</i>			
Daily usage	0.69 (0.22)	-0.42 (0.70)	-1.11 (0.33)
Peak to off-peak	-0.04 (0.03)	0.01 (0.79)	0.05 (0.12)
Bill amount	-0.67 (0.81)	-7.62 (0.19)	-6.94 (0.25)
Structural winner (CPP)	0.03 (0.05)	-0.04 (0.14)	-0.08 (0.01)
Structural winner (TOU)	0.02 (0.17)	-0.04 (0.13)	-0.07 (0.02)
My Account	0.14 (0.00)	0.05 (0.07)	-0.08 (0.01)
My Account logins	2.34 (0.00)	-2.54 (0.05)	-4.89 (0.00)
Paperless	0.06 (0.00)	0.02 (0.52)	-0.04 (0.11)
Low income	0.12 (0.00)	0.18 (0.00)	0.07 (0.00)

Notes: Differences in means between the three customer types: active joiners (AJ), passive consumers (PC), and active leavers (AL). Means are computed following the notes in Table 1.10. The first column compares active joiners to passive consumers, the second compares active joiners to active leavers, and the third compares passive consumers to active leavers. Each cell gives the difference in means with p-value in parentheses for the two given customer types across the given household characteristic. P-values are computed using the following variance formula: $V(\mu^{OO}) = V(p^{AJ}\mu^{AJ} + p^{PC}\mu^{PC})$.


Table 1.12: Hazard Analysis - Accelerated Failure Time (AFT) Weibull Model

	Estimate	s.e.
<i>Model Estimates</i>		
TOU opt-in	0.340	(0.213)
TOU opt-out	0.561**	(0.226)
CPP opt-out	0.339	(0.301)
Low Income (EAPR)	0.514**	(0.202)
Structural winner	0.278	(0.172)
Your account	-0.0438	(0.162)
Second summer reminder date	-4.252***	(0.563)
CPP event date	-3.088***	(0.609)
CPP opt-out \times CPP event date	2.406	(1.618)
CPP event date count in each summer	0.208**	(0.106)
CPP opt-out \times CPP event date count in each summer	-0.352*	(0.210)
Final event in a string of consecutive event dates	-0.0937	(0.673)
Constant	9.869***	(0.292)
ln(p)	-0.328***	(0.0575)
Observations	2,690,168	
<i>Drop out counts</i>		
	Number of households	Number of drop outs
TOU opt-in	2110	92
TOU opt-out	2019	77
CPP opt-in	1585	101
CPP opt-out	701	35

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.


Notes: Top panel in table estimates predictors of time in treatment using an Accelerated Failure Time (AFT) specification, assuming a Weibull distribution parameterized by p . An estimate greater than zero indicates time in the program is extended (reduction in drop-out rate), while a number smaller than zero indicates that the time in the program is reduced (increase in drop-out rate). The omitted category is the CPP opt-in group. Bottom panel counts enrolled households and drop outs by treatment group.

Figure 1.3: Kaplan-Meier Survival Analysis



Notes: Kaplan-Meier survival estimates for each of the four treatment groups. Declining solid line is the proportion of households enrolled at the beginning of the treatment period who remain enrolled over time. 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.

Figure 1.4: Measuring Benefits of Time-Varying Pricing



Notes: Schematic of estimated net benefits of time-varying pricing programs used in Table 8. Source is Potter et. al. (2014), Figure 10-1.

Table 1.13: Structural gains and participation

	CPP	TOU
Opt-in	0.1773 ^{***} (0.0039)	0.1734 ^{***} (0.0033)
Opt-out	0.7448 ^{***} (0.0134)	0.7443 ^{***} (0.0083)
Structural gains X Opt-in	-0.0001 (0.0001)	0.0003 ^{***} (0.0001)
Structural gains X Opt-out	-0.0006 [*] (0.0003)	-0.0003 (0.0002)

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Notes: Table estimates regressions of initial program participation on structural gains for each treatment group. Coefficients are equivalent to the slopes and intercepts of the fitted lines in the top panel of Figure 6.

Table 1.14: Switching costs (random utility model)

	Opt-in	Opt-out
All customers		
CPP	\$1,303 (\$232)	\$-1,003 (\$-231)
TOU	\$1,522 (\$2)	\$2,363 (\$601)
Attentive customers		
CPP	\$8,422 (\$2,089)	\$-145 (\$-0)
TOU	\$1,576 (\$20)	\$278 (\$33)

Notes: Table documents estimated switching costs using a random utility model. Values given in dollars, means are listed with standard deviations in parenthesis. We model the choice to switch as a logit, where the outcome is whether a customer switched away from their default choice (flat rate for the opt-in groups, time-varying rate for the opt-out groups) and the covariates are a household-specific random intercept that represents the switching costs and the NPV of the net structural gains from switching. We scale the estimated mean and standard deviation of the coefficient on switching cost by fixing the coefficient of the NPV of the net change in bill to be one. The first two rows include all customers who never move during the treatment period, while the second two rows include only customers who access their My Account web portal at some point during the study. All estimates assume a discount rate of 5%.

Table 1.15: Discount rates (random utility model)

	Opt-in	Opt-out
All customers		
CPP	7,260% (40,546%)	
TOU	8,293% (5,107%)	22,494% (1,188%)
Attentive customers		
CPP	72,151% (15,248%)	
TOU	8,522% (19,722%)	1,727% (850%)

Notes: Table documents estimated discount rate using a random utility model. Values are in percentages, means are listed with standard deviations in parenthesis. Specifically, we estimate a logit where the choice variable is defined as whether or not a household chooses to switch away from their default choice and the covariates are a fixed intercept that represents the switching cost and the estimated structural gains for each household. We allow the coefficient on structural gains to vary by household and normalize the coefficients by fixing the switching cost to be \$10. This gives us a estimated distribution of the value of summer savings for the sample. From this distribution, we back out the mean and standard deviation of the discount rate as the value that rationalizes the distribution of the value of summer savings we observe.

Table 1.16: Returns on attention (TOU)

	Opt-in	Opt-out
Participation	20%	96%
Passive customers	76%	
Returns on attention (\$)		
<i>Switching cost c = \$10</i>		
Uninformed prior	$\mu = 4.02, \sigma = 6.90$	
No adjustment	13.35	9.51
With adjustment	20.84	3.48
<i>Switching cost c = \$20</i>		
Uninformed prior	$\mu = 8.04, \sigma = 13.79$	
No adjustment	12.33	6.26
With adjustment	16.79	2.49

Notes: This table summarizes our estimates of the distribution of uninformed priors and the associated returns on attention for customers in the TOU group. Participation indicates the observed initial participation percentages in the time-varying pricing program. In the second panel, uninformed priors are the means and standard deviations of the normal distribution that rationalizes the observed participation in the opt-in and opt-out groups, given the assumed switching costs. We simulate returns on attention, or the average value of becoming informed about their structural gains (and possibly changing their enrollment choice) across all customers. The “No adjustment” assumes that customers can become informed about their structural gains under time-varying pricing. “With adjustment” assumes that customers both become informed about their structural gains and anticipate their own changes in energy consumption in response to time-varying pricing. Under both heuristics, we assume a discount rate of 5%.

Table 1.17: Household effects on bill shocked customers by year

	Critical event hours		Non-event day peak hours	
	Opt-in (AJ)	Opt-out (AJ+PC)	Opt-in (AJ)	Opt-out (AJ+PC)
Treatment (CPP)	-0.883*** (0.067)	-0.496*** (0.062)	-0.193*** (0.037)	-0.109*** (0.030)
× 2013	0.197*** (0.064)	0.009 (0.048)	0.083** (0.042)	-0.047 (0.030)
× Shocked	0.505*** (0.099)	0.368*** (0.077)	0.156*** (0.057)	0.034 (0.040)
× Shocked × 2013	-0.240** (0.113)	-0.148** (0.072)	-0.143* (0.074)	0.008 (0.044)
Customers	46,607	39,585	46,607	39,585
Customer-hours	4,528,890	3,847,946	29,258,359	24,859,574
Treatment (TOU)	-0.594*** (0.056)	-0.248*** (0.030)	-0.334*** (0.035)	-0.172*** (0.019)
× 2013	0.160*** (0.057)	0.055** (0.027)	0.105*** (0.038)	0.019 (0.017)
× Shocked	0.173** (0.086)	0.156*** (0.039)	0.101* (0.052)	0.098*** (0.025)
× Shocked × 2013	-0.073 (0.102)	-0.019 (0.040)	-0.139** (0.067)	0.0001 (0.025)
Customers	49,597	40,899	49,597	40,899
Customer-hours	4,818,062	3,974,957	31,128,325	25,681,202

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Notes: Adds interaction between bill shock indicator and 2013 indicator.

A bill shocked customer is who receives a 2012 bill with a per-day average that is 20% greater than the bills they received in the pre-period summer.

Table 1.18: Losses from default enrollment

	Opt-in		Opt-out	
	Prop. losing	Average loss	Prop. losing	Average loss
CPP	0.66	-30.23	0.29	-16.98
TOU	0.45	-34.61	0.55	-25.64

Notes: Table summarizes customer losses from being defaulted into either flat pricing (opt-in groups) or time-varying pricing (opt-out groups). Prop. losing is the proportion of total customers who would have paid less under the non-default option, average loss is the average NPV of the financial loss incurred by those customers.

References

- DellaVigna, Stefano. 2009. "Psychology and Economics: Evidence from the Field." *Journal of Economic Literature* 47:315–372.
- Frederick, Shane, George Loewenstein, and Ted O'Donoghue. 2002. "Time discounting and time preference: A critical review." *Journal of Economic Literature* 40 (2): 351–401.
- Potter, Jennifer M., Stephen S. George, and Lupe R. Jimenez. 2014. *SmartPricing Options Final Evaluation*. Technical report.