

Take the Load Off: Effort and Technology as Determinants of Electricity Demand Response*

Megan R. Bailey[†] David P. Brown[‡] Blake Shaffer[§]
Frank A. Wolak[¶]

July 2025

Abstract

As electricity systems transition to greater shares of renewable energy, flexible demand becomes vital for grid stability. We test whether smart technologies can unlock this flexibility or if human behavior remains the key barrier. Our field experiment randomizes peak event timing across households enrolled in one of three demand response programs: fully automated, app-enabled, and manual. Households with automated systems reduced consumption five times more than those requiring any action. Simply providing smart technology made no difference when effort was still required. Our results suggest that time and attention—not technology—are the main constraints to realizing greater demand-side flexibility.

*We thank Meredith Fowlie, Koichiro Ito, Katrina Jessoe, Ryan Kellogg, Robert Metcalfe, Mar Reguant, and attendees at the 2022 Electricity Camp in the Rockies, 2022 Berkeley Energy Camp, 2022 Canadian Resource and Environmental Economics Annual conference, 2023 ASSA sessions, the University of Alberta, 2023 Berkeley POWER conference, 2023 NBER Spring Environment and Energy Economics Meeting, and 2023 AERE Summer Conference for useful feedback. We are grateful to our partner electric utility company for sponsoring and managing the field experiment on which this study is based. This work is supported by excellent research assistance from Mallika Sharma. We acknowledge funding from our utility partner, the University of Calgary, the Canada First Research Excellence Fund as part of the University of Alberta’s Future Energy Systems research initiative, the University of Calgary’s Global Research Initiative, and the Social Sciences and Humanities Research Council Insight Development Grant. This experiment is registered in the AEA RCT Registry under ID #AEARCTR-0008680 and the UCalgary Conjoint Faculties Ethics Board (REB21-0161).

[†]Haskayne School of Business and the School of Public Policy, University of Calgary
megan.bailey@ucalgary.ca

[‡]Department of Economics, University of Alberta dpbrown@ualberta.ca

[§]Department of Economics, University of Calgary blake.shaffer@ucalgary.ca

[¶]Program on Energy and Sustainable Development and Department of Economics, Stanford University and NBER wolak@zia.stanford.edu

1 Introduction

“While the electricity system of the past involved forecasting demand and dispatching supply, going forward grid operators will increasingly find themselves forecasting supply and dispatching demand.”

–Jeff Dagle, Pacific Northwest National Laboratory

The energy transition is rapidly changing electricity systems around the world as supply decarbonizes and new sources of demand are electrified (e.g. heating, transportation). However, the growth of variable renewable generation sources, such as wind and solar, has raised reliability challenges for power grids, where supply and demand must balance at every instant. These market dynamics highlight the potential for flexible demand to play a critical role in facilitating the transition to a low-carbon electricity supply. Despite its promise, questions remain as to how to elicit such flexibility from notoriously inelastic electricity consumers.

One solution, long advocated by economists, is to send dynamic price signals to customers to reflect the time-varying nature of marginal electricity system costs (Boiteux, 1960; Kahn, 1970; Joskow and Wolfram, 2012). However, when facing time-varying pricing, customers face a tradeoff between the financial benefits, discomfort from adjusting consumption, and the cost of taking action which includes the time, effort, attention, and cognitive burden of interpreting and responding to price signals, making such first-best policies unrealistic (Schneider and Sunstein, 2017). In addition to the onus of having to gather (and understand) their electricity price schedule and usage, consumers face an even more fundamental obstacle: The reward for adjusting consumption may simply not be worth their opportunity cost of time (Becker, 1965).

Many utility demand response programs amount to consumers “picking up pennies” in a series of irregular and relatively low stakes opportunities where consumers earn rewards or savings for reducing consumption during “peak events” (Harding and Sexton, 2017). Consumers can respond to these events, but doing so requires allocating time to actively participate and exert effort. Consumers may be understandably reluctant to allocate their scarce time when the individual private gains are small—despite potentially large social gains in the aggregate. This “rational inattention” (Sallee, 2014) can lead to low price responsiveness and is reflected in inertia in consumer decisions.

In this paper, we examine the role of effort and attention as barriers to flexible residential electricity demand. In particular, we investigate the role that emerging

technologies and the requirement to respond actively versus passively to events can play in reducing these barriers. Partnering with a large electric utility in Canada, we recruited approximately 1,000 households to a demand response experiment that lasted 17 months. Participants received individually randomized notifications of “peak events”, roughly 3 to 4 per month, whereby they received financial rewards for reducing electricity consumption during the 3-hour window of each event.

Participants were assigned to one of the three demand response programs that differed in the provision of enabling technology and whether active versus passive responses to events were required. Households in the most basic **Manual program** received information on their real-time consumption as well as notifications of peak events via the Utility’s App and had to respond *actively* by manually adjusting their consumption, as the name suggests. Households in the **Tech program** received the same information but also had app-enabled load (electricity demand) controllers installed on their baseboard thermostats, hot water heaters, and/or electric vehicle (EV) chargers, allowing them to respond *actively* to events but with the ease of doing so remotely via an App. Finally, households in the **Central program** had the same information provision and load control technology installed as those in the Tech program, but with the key difference that the Utility automatically reduced their consumption in response to an event, i.e. their response was *passive*. Even if they exert no effort, or pay no attention to the event, households in this program will have their consumption reduced and can earn rewards. Households could override the automatic reductions, but, in contrast to the Tech program, they needed to actively make an effort to *not* reduce consumption by pushing a button on their App.¹

Our treatment of interest is the effect of peak events on electricity consumption and, importantly, how that differs across participants in different programs. Accordingly, we randomize our treatment events, i.e. peak event notifications, at the level of household-day. That is, each participant receives a unique randomized schedule of peak events over the 17-month period of the study.² This granularity of randomization differs from typical “between-subject” designs, where randomization occurs at the level of group allocation, with participants subsequently receiving common treatments. Our design incorporates what is often called a “panel experiment” or

¹There are a number of recently developed programs that include utility-managed load-control for various appliances including hot water heaters ([Wattersaver, 2023](#)), thermostats ([PG&E, 2023](#)), electric vehicles ([DTE Energy, 2022](#)), and solar-plus-storage systems ([Spector, 2020](#)).

²The schedule is blind to the participant prior to the notification that occurs 21 and 2 hours preceding an event.

“within-subject” design (Charness et al., 2012; Bojinov et al., 2021; List, 2025). Within-subject designs are far less common than between-subject designs, in part because of the complexity involved in implementing an experiment with personalized randomization.

There are several key benefits of leveraging within-subject identification. First, we are able to estimate the average treatment effect from peak events, separately by program, for the population of households that accepted the demand response program offers. We can directly estimate this effect due to 100% compliance of receipt of treatment events, i.e. peak event notifications, by participants in the study. Second, the within-subject design allows us to estimate household-specific treatment effects, providing estimates for the full distribution of responses to peak events for each demand response program. We analyze these household-level estimates to better understand the mechanisms driving our results.³

Our results provide three key insights. First, we find that people do respond to financial incentives, but the requirement to respond actively is a major barrier to performance. Using both within- and between-subject variation, we find participants in the Central program reduced consumption, on average, by 26.3% during events, as compared to 4.8% and 5.3% in the active-response-requiring Tech and Manual programs, respectively. Having to take an action to reduce consumption, even one as small as pushing a button on an app to respond to an event, results in one-fifth of the effect. Second, technology alone is not sufficient. Households in the Tech program, despite receiving remote control capabilities, performed similarly to those in the Manual program. Simply providing consumers with smart home technology to monitor and control their energy use is insufficient to drive significant consumption changes. Third, the magnitude of the Central program’s response reinforces the considerable potential of demand flexibility as a resource for electricity grids.

We find that the distributions of household treatment effects, estimated solely based on within-subject variation, differ significantly across the programs. Participants in the Central program are symmetrically distributed around a mean reduction of 24%. However, those in the Tech and Manual programs have mean reductions of approximately 5%, consisting of a large mass around 0% and a small subset of high performers. Our within-subject design also allows us to examine the consistency of individual household treatment effects by exploring the variance of responsiveness

³Another benefit of the within-subject design is that the variation provided by the individualized randomization schedules likely increases the statistical precision of our estimates.

across events within each household. We find that the standard errors are similar across the three programs, with slightly smaller values for the Central program. This suggests that households’ responses to events are broadly consistent. That is, the Central program households are consistently large responders, while the Tech and Manual households are in large part consistently low responders. This finding has important policy implications as the Central program can be viewed as a reliable source of demand-side flexibility.

We explore potential mechanisms driving our results by combining the household-level estimated treatment effects with detailed data on participants’ interactions with the utility’s App. We provide evidence that differences in performance are driven by differences in time allocation and attention. We find that when households do not interact with their App on peak event days, the average household-level treatment effect is roughly 2% for the Tech and Manual programs and 22% for the Central program. When households do interact with their App, the reductions increase to 6% for the Tech and Manual programs, and 25% for the Central program. These results emphasize that while all programs increased responsiveness when households interacted with the App, the Central program’s “headstart” from not having an active response requirement is associated with the largest effects.

The App interaction data provides further insights into the intensity to which households allocated time to monitor and provide demand flexibility and how this correlates to their performance in responding to peak events. We find that the small subsample of “high achievers” in the Tech and Manual programs interacted with their App very frequently (61% and 51% of days on average), which suggests that their performance relied on considerable time devoted to their electricity consumption decisions. In contrast, the high achievers in the Central program interacted with the App significantly less (31% of days on average).

We also investigate how price affects the degree of responsiveness by randomizing peak event incentives into regular and “high” reward events, the financial rewards for consumption reductions being roughly double in the latter. We find no statistically significant evidence that higher prices motivate greater consumption reductions.⁴ This further suggests a story of responsiveness in which the costs of actively responding act as the key barrier, rather than a smooth response to price signals.

Finally, we use data from an exit survey to understand the relationship between

⁴This finding is similar to that of [Prest \(2020\)](#) who estimates the price-responsiveness of households facing time-of-use prices in Ireland (without automated technologies) and find that households are responsive to TOU price signals, but the magnitude of the price does not matter.

our estimated household-level treatment effects and measures of individuals’ opportunity cost of time and willingness to exert effort. Households with lower income and/or a higher perceived personal net benefit of providing demand flexibility were associated with larger responses to peak events. This analysis provides further support that time, effort, and attention are key factors limiting demand response. The Central demand response program is able to overcome these barriers through the use of enabling technologies and utility-controlled/automated responses to events.

Our results point to the need to take an active action as a key barrier to responsiveness. This is consistent with a notion of adjustment costs (or “ordeals”) resulting in a strong default effect. Default effects have been considered in a range of settings, including retirement savings (Chetty et al., 2014; Bernheim et al., 2015), organ donations (Abadie and Gay, 2006), subscription decisions (Einav et al., 2025), and health insurance (Brot-Goldberg et al., 2023; Shepard and Wagner, 2025). In this literature, the work most closely related to ours is Fowlie et al. (2021) who look at opt-in vs opt-out default effects at the extensive margin of selecting time-varying retail electricity pricing plans. Our paper complements this work by focusing on default effects at the intensive margin reflecting the decision to alter consumption decisions in response to financial incentives. Our key contribution beyond the existing literature is the finding of significantly greater responsiveness when consumption reductions are made the default, or passive, action in response to demand response events. Requiring customers to take action—even with the provision of technology that makes the associated cost as minimal as remote control with a mobile phone app—is no match for the power of demand response that is managed on the consumer’s behalf. This speaks to the importance of recognizing the opportunity cost of time/cost of effort in settings where inattention is high and individual event rewards are relatively small, despite potentially large social gains in aggregate.

There is an emerging literature that explores automation of key household appliances for consumers to reduce time and effort-related costs. There is evidence that automation of smart thermostats and electric vehicles (EVs) can assist in facilitating short-run demand responsiveness when combined with pricing (Harding and Lamarche, 2016; Bollinger and Hartmann, 2020; Burkhardt et al., 2023; La Nauze et al., 2024; Bailey et al., 2025a,b). Our paper builds on this work by investigating the role of prices, technology, and effort barriers, leveraging automation in a broader range of smart home technologies. In particular, we are able to isolate the relative impact of automation compared to a program that has the exact same technology provided

without automation. Our findings suggest that there is a threshold effect where even a seemingly small amount of effort required to respond serves as a sufficient barrier to limit demand response.

Recent work finds that consumers may override important settings with demand response enabling technologies, reducing the anticipated benefits (Brandon et al., 2022). Consistent with this work, we find our Tech program performs no better than the Manual program on average. That is, given the ability to remotely control large appliances as well as automate some aspects of their electricity usage (e.g. thermostat settings), they fare no better than consumers who require a more manual action. The key takeaway is that to reach their full potential, enabling technologies must account for how consumers actually engage with technology—including their reluctance—as a barrier to responsiveness.

Finally, our work builds on the expanding experimental literature that relies on within-subject variation (Bojinov et al., 2021; List, 2025). In addition to improving statistical power, this framework allows us to identify household-specific treatment effects. This provides insights into the heterogeneity of responses and allows us to better understand the potential channels through which time, effort, and attention relate to the degree of price-responsiveness provided across demand response programs, further quantifying the role of technology and automation.

The rest of the paper proceeds as follows. Section 2 describes our experimental design and data. We start the analysis with descriptive analysis in Section 3, followed by our formal estimation framework in Section 4 and estimation results in Section 5. Section 6 presents empirical evidence to understand the mechanisms behind our results. Section 7 concludes.

2 Experimental Design and Data

2.1 Overview

We partnered with a large Canadian electric utility (hereafter referred to as the “Utility”) to create three demand response programs that vary in terms of the enabling technology provided to customers, as well as household versus utility-initiated electricity demand changes on specific devices. We are primarily interested in how customers in each program adjust their electricity consumption in response to “peak events”, times during which the Utility asks consumers to reduce consumption and rewards them financially for doing so.

A unique feature of our study is the fact that we randomized the timing of peak events at the household level. We leverage the randomization of these events (as well as the richness of our data) to estimate household-specific treatment effects. These features allow us to look at the distribution of effects within each program to better understand what is driving average program-specific results.

2.2 Treatment Events

Customers in each of the demand response programs received notifications of peak events through an electricity consumption management phone App offered by the Utility. Peak events had the possibility of occurring at one of two time periods: morning (7am to 10am) or evening (5pm to 8pm). The schedule and timing (morning or evening) of events were unknown ex-ante to the customer, each receiving a unique, randomized schedule of events over the course of the experiment. Consequently, households could not predict the day or event time when they would receive a peak event. Households received two event notifications – 21 and 2 hours before the event – that included an offer for households to receive financial rewards for reducing electricity consumption during the peak event period, relative to their household-specific baseline.⁵ Further, because events were randomized in time for each household, they are not correlated with other drivers of household electricity consumption.

“Event types” were also randomized and were one of two pricing levels, “normal” and “high”, with rewards increasing in the latter for large reductions. High peak events were only possible during evening periods. During normal events, households could receive \$1 for a 10% reduction, \$2 for a 30% reduction, or \$3 for a 50% reduction. During high peak events, households could receive \$1 for 10%, \$3 for 30%, or \$6 for 50% reductions.⁶ By randomizing the pricing levels, we are able to estimate the effect of greater price incentives on household consumption decisions. The incentive amounts translate to payments ranging from approximately \$1.11 to \$2.22 per kWh of electricity reduced, for the average household.⁷ These incentives are in the range

⁵Baselines were calculated based on a household’s average consumption during the relevant event time window over the last five weekdays prior to the event, excluding days where events occurred. Customers did not know how the baseline was calculated to avoid the potential for customers to “game” their baseline 21 hours was selected as the longest notification period to avoid consumers having knowledge of a pending event while their baseline was still being set.

⁶Report dollars are in Canadian dollars. CAD\$1 equals approximately USD\$0.75 or EUR€0.68 as of December 2023.

⁷The average household consumes 1.8 kWh in each hour in our sample. A 10%, 30%, and 50% reduction translates to a 0.54, 1.62, and 2.7 kWh reduction over the three-hour event, respectively. Consequently, for a 50% reduction during a normal peak event, we compensated households $\frac{\$3}{2.7} =$

of wholesale price caps used in a number of jurisdictions in North America.⁸

Events randomly occurred on weekdays, excluding holidays. Households typically experienced two “normal” and one “high” event per month. This schedule was altered in the summer months of July and August when the likelihood of peak events is lower in Canada. During these months, households experienced no “high peak” events. Events started on February 22, 2022, and continued until June 30, 2023 resulting in 30,502 household-event days in our experiment.

Event notifications provided information on the time of the event and the financial incentives for the different demand reduction levels. Once consumers received the 21-hour notifications, they could also see event details in the App itself. See Appendix B.1 for examples of the notification and in-App event messages. Households’ rewards for consumption reductions during events were displayed in the App at a two- to three-day lag. The App also gave households a summary of their total rewards to date. See Appendix B.2 for details on each program’s in-App experience.

2.3 Demand Response Programs

Our initial sample consists of all households that first downloaded the utility’s App. The App provides customers with household-level hourly consumption posted at a one-day lag. From the pool of approximately 9,000 households that joined the App, we identified 1,661 eligible households based on factors such as being located near a major metropolitan city and household characteristics.⁹ Eligible households were randomized into one of three demand response programs or two never-treated groups. Email invitations were sent to eligible households randomized into the demand response programs, with several reminders. See Appendix A.1 for a complete description of the recruitment and assignment process.

Table 1 summarizes our demand response programs and never-treated groups. Households in the **Manual** program earned financial rewards for demand reductions during events but did not have any load controller equipment installed by the Utility

\$1.11 per kWh. For a 50% reduction during a high peak event, compensation was $\frac{\$6}{2.7} = \2.22 per kWh. The other percent reductions lie between these two cases.

⁸Examples include the wholesale price cap of CAD\$1.00/kWh in Alberta (Brown and Olmstead, 2017), USD\$3.50/kWh in the Mid Continent Independent System Operator that operates in the Midwest United States (IRC, 2017), and USD\$5/kWh in Texas (Smith, 2022).

⁹Eligible households were those that were in or near a major metropolitan city for which it was feasible for utility-partnered electricians to install load control equipment; homeowners; not in condos or apartments; had at least one month of historical electricity consumption data as of September 2021; and had at least one controllable electrical device (level 2 electric vehicle charger, electric baseboard heater as the primary heat source, and/or an electric hot water tank).

Table 1. Summary of Household Programs

Programs	DR Control	Load Controller	Price Incentive	Real-Time Usage Info
Central	Utility	✓	✓	✓
Tech	Household	✓	✓	✓
Manual	Household		✓	✓
Info	Household			✓
Control	Household			

Notes. DR Control represents whether demand response to events is controlled entirely by the household or by the Utility for the load-controlled devices. Load Controller denotes whether the household has load controller equipment installed. Price Incentive reflects if households receive peak events and rewards for reduced demand during events. Real-Time Usage Info denotes whether households receive real-time household-level consumption information. ✓ indicates categories that are applicable to each program.

to manage their consumption via the App. They had to respond to events manually.¹⁰ Households were sent a device that allowed them to monitor their real-time electricity usage information in the Utility’s App.

The **Tech** program differs from the Manual program in that the Utility installed load controller equipment on one or more of the household’s electric hot water heaters, baseboard thermostats, and level 2 EV chargers, to enable remote electricity consumption (“load”) reductions. This equipment allows households to see device-specific electricity consumption and turn on and off devices remotely via the App. Critically, while the Tech program households are equipped with load control technology to ease their effort in responding to events, they still must take an active action to do so via their phone’s App. See Appendix B.2 for detail on the experience that participants had with the App, by program. Both the Manual and Tech programs allow us to test the efficacy of an active-response-requiring approach to demand-side flexibility—one without and one with enabling smart technology, respectively.

The **Central** program received the same equipment installed in their homes as the Tech program. However, during an event, the passive setting for Central program participants was for the Utility to manage their load-controlled devices by reducing electricity consumption. That is, without any active response, the Central program

¹⁰We can observe if a household independently installs its own equipment and links it to the Utility’s App. Only 3 (out of 242) Manual households installed equipment independently. These 3 all installed smart thermostats that allow the monitoring and remote control of the household’s electric baseboard heaters.

participants would reduce consumption via demand management initiated by the Utility. For example, during events, their thermostat temperature is automatically reduced, EV charging is delayed, and/or the water heater is turned off. Central program households needed to actively choose *not* to respond to an event (i.e., select out of utility management) by pushing a button on their App. The Central program allows us to examine the efficacy of centralized demand management and, as compared to the Tech program, the difference between having to take active action during an event versus passively responding by conceding control to a third party.¹¹ It is important to note that the Central, Manual, and Tech programs had symmetric information about their real-time usage which could be observed within the App.

Finally, we have two groups of households that serve as *never-treated* baselines throughout the study. One is a **Control** group that receives no intervention or messaging regarding the experiment after joining the utility’s App. Another is an **Info** group that is identical to the Control group, but these households have access to real-time consumption information for their home via the App after they accepted and installed a device provided by the Utility. Both of these never-treated groups do not receive peak events or financial incentives. We passively monitor their consumption.

While households outside of the Control group needed to accept being in each program, the final sample of participating households are representative of those that utility companies are likely to be interested in until such time that demand response programs serve as the default.

2.4 Acceptance

Table 2 summarizes the number of households invited and the acceptance rates for each program offer. Acceptance rates among all programs were high. In particular, the acceptance rate for the Central offer was 42%, and only marginally statistically different than the acceptance rate for the Tech offer (48%). A difference in means test between these two values yielded a p-value of 0.072. Compared to Tech and Central, the Manual program had a statistically significantly higher acceptance rate of 59%, followed by the Info-only group at 68%.¹² Finally, Control had 100% acceptance

¹¹Customers in the Central program received advanced notifications of events just as other program participants, and the language of the notifications reminded them that the Utility would manage their devices. See Appendix B.1 for event notifications for each program. They had the option to opt-out of utility management for all devices ahead of an event or for individual devices during an event. These options are described in Appendix B.2.

¹²Like the Manual program, the info group required actively accepting the offer to join the experiment and installing a device that facilitates the monitoring and reporting of real-time consumption.

because their participation was not subject to an offer. The acceptance rates of the Central and Tech programs were lower than the others due in part to the need for load controllers to be successfully installed in households that accepted these offers.¹³

Table 2. Program Acceptance by Program

	Central	Tech	Manual	Info	Control
Invited	423	382	409	259	188
Accepted	177	184	242	177	188
Pct. Accepted	(42%)	(48%)	(59%)	(68%)	(100%)

Notes. “Invited” reflects the number of households invited to participate in the experiment, by program. “Accepted” is the number of households that accepted our offer and made it through the equipment installation process (as applicable, by program). “Pct. Accepted” displays acceptance rates relative to the number of households invited.

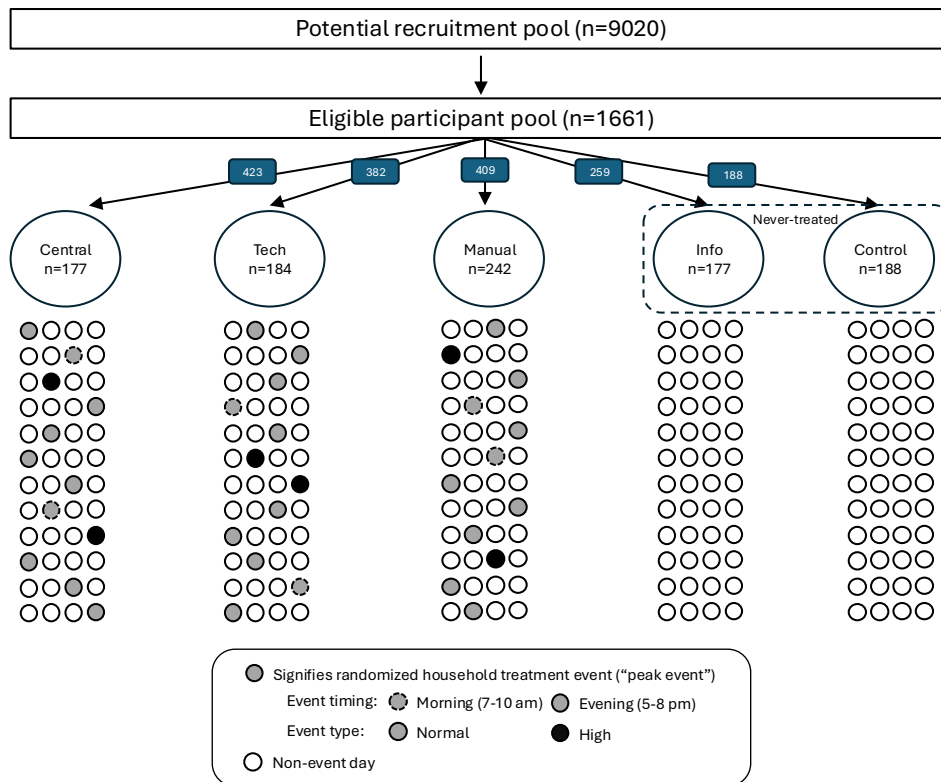
We take the similarity among final acceptance rates between the Central and Tech programs as the first set of evidence that we can compare our estimated treatment effects between these programs. While the Manual program had a higher final acceptance rate, concerns that the Manual program participants systematically differ from those of the other two programs are mitigated based on a comparison of observable characteristics across programs in our final sample, as are concerns about differences in the composition of households invited to each program. See Appendices A.2 and A.3 for a detailed discussion of balance using pre-treatment data.

Figure 1 provides a summary of our experimental design, starting with our initial recruitment and eligible participant pools, and subsequent allocation and acceptance into our demand response programs and never-treated groups. Next, we illustrate our within-subject design with a column of dots for each household (limited to four per program here, solely for exposition), with each dot representing a day in the experiment and shaded dots indicating randomly allocated peak events. Event variants, such as normal vs high reward days and morning vs evening events, are indicated based on the colour and linetype of the shaded dots. This figure highlights how the randomly allocated peak events occur on different days for the treated households, although the distribution of events (and variants) is similar across households. We leverage this within individual randomization in our identification strategy outlined

¹³We observed unsuccessful installation at households that initially accepted these offers due to, for example, households never responding to subsequent inquiries to receive and install equipment or households not being in compliance with local electrical codes.

in Section 4 to identify the average treatment effect of peak events on the population of households that accepted participation in our demand response programs. We further evaluate if responses vary by peak event type (i.e., morning, normal evening, high evening) by program.

Figure 1. Experimental Design Summary



Notes: Dots represent individual household-days. Treatment events are indicated by filled dots, which are randomized at the household level, i.e. running vertically in this diagram. Dashed contours on treatment events indicate morning events, solid contours indicate evening events. Grey fill indicates “normal” event types and darker fill indicates “high type” (i.e. larger incentive) events.

2.5 Data Description

For all households in our experiment, we track hourly household-level consumption (in kWh) from October 1, 2020 until June 30, 2023. We also have information on a number of household characteristics, such as household appliances, that were provided through survey responses as a necessary condition to enter the first phase of

our recruitment process. In addition, the Utility provided supplementary household information, including the type of household (e.g., single-family/duplex, row home) and an approximate geographical location. We are also provided time-stamped information on household interactions with the Utility’s App at the daily level.

We complement the detailed household-level data with demographic information from the 2016 Canadian Census ([Statistics Canada, 2021](#)). We are provided a household’s Census Dissemination Area (CDA) identifier; the CDA is the most granular geographical unit for which all Census information is provided publicly. We collect hourly weather information to control for environmental factors that impact electricity consumption, including temperature and humidity at three stations that are geographically representative of the households located in our study.¹⁴ These data were accessed at Environment and Climate Change Canada.

In the last month of the experiment (mid-June 2023), we conducted an additional survey that contained questions on participants’ experience with their respective demand response program. We provide a subset of questions from the survey in [Appendix D.1](#). We use these survey responses in our analysis in [Section 6](#).

3 Descriptive Results


We begin our analysis with descriptive evidence that participating households reduce their electricity consumption during peak events and show how this response differs across demand response programs. [Figure 2](#) provides average hourly household-level consumption for the Central, Tech, and Manual programs for the entire sample period during non-event (solid lines) and event (dashed lines) days. The shaded regions reflect the relevant event hours for each event type.

[Figure 2a](#) demonstrates that the Central program had a large reduction in average consumption during events regardless of the event type. After each event, we observe a large spike in consumption. This “snap-back” is consistent with the devices turning on immediately after the event (e.g., to reheat the water tank and/or home, or restart EV charging).¹⁵ Comparing High Evening to Evening event consumption, we see no

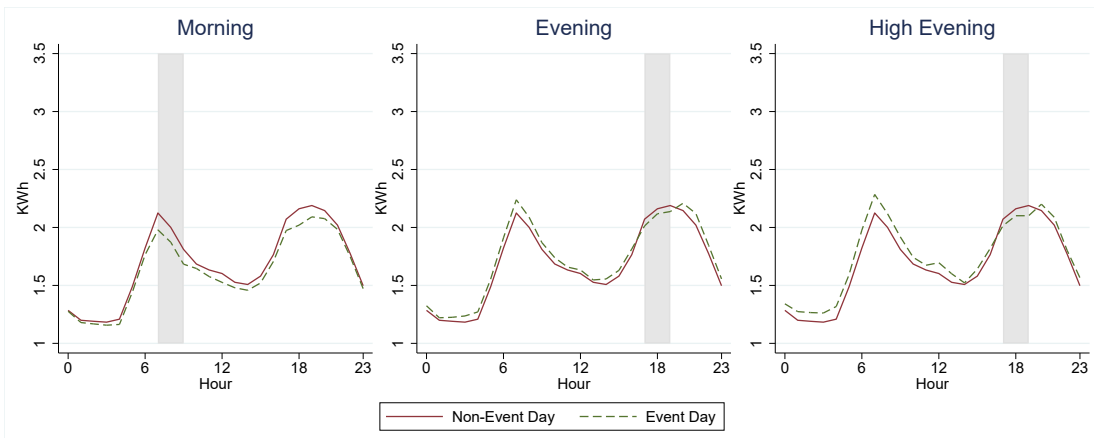
¹⁴We match the households in our sample with their closest weather station.

¹⁵Because we are interested in event-induced demand flexibility among programs, not energy efficiency, we are unconcerned with demand being shifted to another time. However, based on conversations with the Utility, the new problematic peaks in demand created by the observed snap-back could be mitigated by the Utility staggering the beginning and/or end of the load-controlled event across households or only partially adjusting the demand levels on controllable devices. Managing the “snap-back” or “shadow peak” from demand response is an important area for further research.

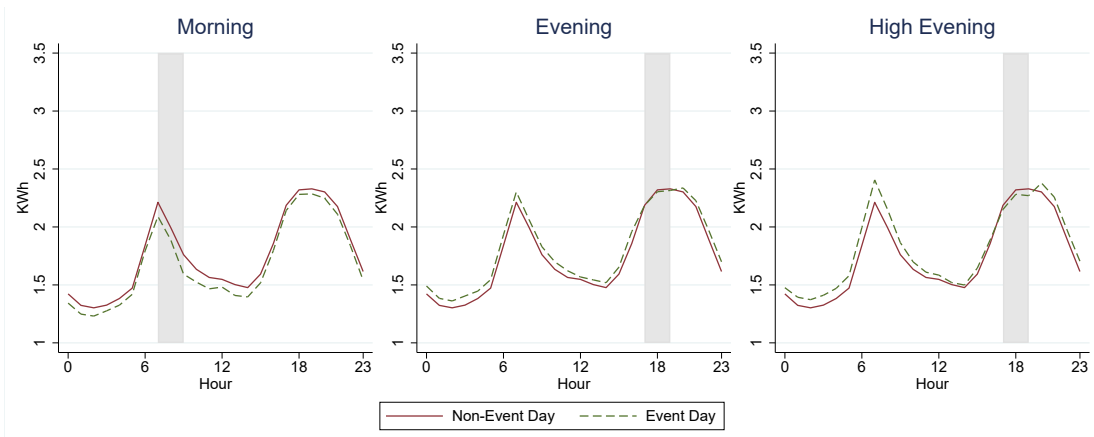
Figure 2. Average Household Consumption



(a) Central Program



(b) Tech Program



(c) Manual Program

Notes: Figure plots mean household consumption by hour and demand response program on weekdays, on event days and non-event days over the period February 1, 2022 - June 30, 2023. Event days are separated by type: Morning, Evening, and High-Evening. The shaded area represents the relevant 3-hour event period.

discernible difference in response to this higher reward.

Figures 2b and 2c demonstrate that the Tech and Manual programs show negligible changes in average consumption patterns during events. This limited observable response for the Tech program arises despite the fact that this program has access to the same equipment as the Central households. However, unlike the Central program where changes are automated, the Tech households must actively engage with the App or device to turn off the same appliances during an event.

Taken together, these descriptive results suggest that the Central program has a considerably larger response to each event type. Further, we see no visual evidence of greater performance for greater financial rewards. Rather, the largest difference appears to be whether the household is in a passive versus active response-required program. In the sections that follow, we undertake a formal empirical analysis to quantify these effects and control for potentially confounding factors.

4 Empirical Framework

4.1 Program-Level Regressions

We estimate the average treatment effect of peak events by comparing electricity consumption for households facing a randomized peak event to those who, contemporaneously, are not, using the following model at the household i and hour t level:

$$\ln(c_{it}) = \sum_{j \in \{C, T, M\}} \beta_j \text{Program}_{ji} \cdot E_{it} + \alpha_i + \tau_t + \delta X_{it} + \varepsilon_{it} \quad (1)$$

in which $\ln(c_{it})$ is the log of household electricity consumption, E_{it} is the household-specific event indicator that equals one if the household is (randomly) assigned an event in hour t , and Program_{ji} is a categorical variable for which demand response program the household i is enrolled in (e.g., Central [C], Tech [T], or Manual [M]). A key advantage of this model is that it allows us to readily test for differences in responsiveness to events across the three demand response programs. We use the log of household electricity consumption on the left-hand side to account for the right-skewed nature of consumption.¹⁶

We include α_i , household fixed effects, which control for time-invariant household characteristics. We also include τ_t , an hour-of-sample fixed effect, which controls for

¹⁶Our results are robust to functional form; we observe similar results with a linear-linear specification.

time-varying factors that impact consumption. Household electricity consumption and consumer responses to events may vary with local weather conditions (especially due to thermostat settings). To control for this, we include X_t , a vector of hourly weather controls that include the relative humidity and cooling degrees and heating degrees above and below 65° F (18.33° C). Since these may vary in weather conditions in a nonlinear way, we include a flexible functional form with a polynomial up to the third degree for each weather-related covariate. ε_{it} is the error term. We cluster standard errors at the household level.

We also consider a version of this regression specification where the event indicator variable, E_{it} , is adjusted to be a categorical variable for the three potential event-types: Morning, Evening, and High-Evening. This analysis allows us to evaluate if households’ responses to events differ by time-of-day and the financial reward for responding, in the case of Evening compared to High Evening events.

Our parameters of interest are β_j for $j \in \{C, T, M\}$, which measure the change in household-level electricity consumption during peak events for each of the Central, Tech, and Manual demand response programs. Because of our log-linear specification, we transform our estimates to report the percentage change in hourly consumption during an event via $100 \times (\exp(\hat{\beta}_j) - 1)$.

Our empirical framework relies on three identifying assumptions to recover the effect of events on household-level consumption. First, events are not correlated with other drivers of household electricity consumption. This is met via our randomization of events. We include weather controls to ensure that our estimated treatment effects can be interpreted as weather-agnostic.

Second, our analysis falls within the literature on experiments that leverage “within-subject” variation, which have been referred to as “panel experiment” designs (Charness et al., 2012; Bojinov et al., 2021; List, 2025). Panel experiments involve the treatment of interest (e.g., peak events) that vary in time. This differs from designs that use a one-time treatment and between-subject variation to identify treatment effects. A key identification assumption needed for estimating event-level treatment effects in the presence of within-subject variation is that our random treatment events do not “carryover” to a persistent change in consumption decisions in similar hours on non-event days.¹⁷ This could occur if, for example, experiencing an event led to a household persistently scheduling consumption reductions during the event hours on all days going forward. Or, conversely, a savvy participant might suspect a financial

¹⁷List (2025) refers to this as the *causal transience* assumption.

benefit of “gaming” the baseline by purposely increasing consumption during event hours on non-event days.¹⁸ We test the validity of the “no carryover” assumption using a DID regression to evaluate whether households in each demand response program adjusted their consumption during the event windows on non-event days in the post-treatment period, relative to the never-treated households. Details of the empirical approach are provided in Appendix C.3.

Third, as noted above, for each program, Equation (1) compares event time consumption to non-event time consumption of households in the same demand response program, other demand response programs, and never-treated households. This relies on the assumption that the households in the never-treated groups and other demand response programs provide a valid counterfactual on non-event days. While observed characteristics are similar across the demand response programs (see Appendix A.3), one may be concerned that the programs are differentially selected, and this impacts our ability to compare event-time to non-event-time consumption decisions across programs, even after including our various control variables.

To address this potential concern of between program comparison validity, we perform a robustness check whereby we vary the set of comparators. We start with the most restrictive set: limiting the analysis to be solely within each demand response program and excluding the never-treated households. This restrictive specification ensures we are only comparing behavior across participating households within the same program and guards against concerns of comparing consumption between participating and never-treated households. We run three separate regressions, one for each program, to separately estimate β_C , β_T , and β_M . The treatment effect is identified within each program by comparing consumption when households receive peak events versus not. Next, we augment the comparison group to include never-treated households, i.e. those in the Control and Info groups, as non-treated comparators, but again run three separate regressions to estimate treatment effects for each demand response program separately. Finally, we pool all participants and never-treated households to run a regression comparing those receiving a peak event to all households not receiving a peak event at the same time (i.e., Equation (1)). To the extent that our results are robust to these alternative specifications, this will alleviate concerns that potential differential selection into demand response programs impacts our key conclusions.

¹⁸We mitigated this latter effect by providing no information on how the baseline consumption that was used to determine the household’s rewards was calculated.

Our empirical framework provides estimated average treatment effects to peak events by program for households that accepted the demand response program offers. In our setting, the treatment of interest is randomized household-specific peak events that all households in demand response programs received. As a result, there was no non-compliance to our treatment of interest. One might consider estimating an Intention-to-Treat (ITT) measure for the impact of program invitation on event-time consumption (i.e., including households that did not accept the invitation to participate). However, households that did not join a demand response program did not subsequently receive the (randomized) peak events. Because of this, non-participating households will have (noisy) null event responses. Appendix C.4 presents an analysis that assigns “synthetic events” to households that did not join our demand response programs. As expected, this attenuates our main estimated average treatment effects downward. However, our key conclusions that compare estimates across demand response programs persist.¹⁹

Finally, as shown in Figure 2, there is the potential for spillover or “snap back” effects after events. Appendix C.2 extends our main specification to control for potential pre-event and post-event spillover effects. For expositional purposes, and because our results are robust to the inclusion of these covariates, we focus on the specification that estimates event-time-only effects.

4.2 Household-Level Regressions

A unique feature of our setting is our ability to estimate household-level treatment effects solely from *within-subject* variation by leveraging the randomized event timing. This allows us to examine heterogeneity and distribution of event responsiveness across households, and to explore factors associated with household-level responsiveness that speak to the role of attention and effort/time allocated to responding to events.

Separately for each household i in the demand response programs, we estimate a household-specific treatment effect using the following model:

¹⁹The key conclusion to be drawn from our analysis is the relative magnitudes of the average responses to peak events across demand response programs. Although the precise magnitude of the estimate varies by the estimation methodology and subsample considered, our key findings that Central systematically results in a statistically and economically significantly larger (by at least two to three times) reduction in demand compared to Tech and Manual, and that these two programs have similar magnitudes to each other, persists throughout.

$$\ln(c_{it}) = \gamma_i + \beta_i E_{it} + T_t + \delta_i X_{it} + \eta_{it} \quad (2)$$

where, analogous to above, c_{it} is consumption, E_{it} equals 1 when household i has an event and zero otherwise, and X_t includes the same set of temperature controls as the specification in Equation (1). In this specification, T_t is a set of time fixed effects that includes day-of-week, hour-of-day, and year-month to capture time-varying factors that impact consumption.²⁰ η_{it} is the heteroskedastic-robust error term.

The regression in Equation (2) estimates a separate $\hat{\beta}_i$ for each household in our demand response programs. The household-level regression compares consumption during event hours to non-event hours within the same household, conditional on time-based fixed effects and weather variables. Similar to the approach in our program-level estimation, we also consider a specification that allows for differential responses by event type.

The ability to estimate household-level treatment effects provides us with the opportunity to understand the potential mechanisms driving our results. In particular, we have data on when a household has interacted with the App on a given day. App interactions are indicative of the time and attention that households expended to respond to events.²¹ We leverage this to estimate separate event treatment effects, by household, for when households do and do not interact with the App.

We run a specification of Equation (2) that interacts E_{it} with an indicator variable App Interact $_{it}$ that equals one if the household has interacted with the Utility’s App on the relevant day and zero otherwise. A key benefit of this approach is that it allows us to quantify how a specific household’s estimated treatment effect varies by whether or not they interacted with the App on an event day. This helps overcome the sample selection challenge that would arise by running an analogous regression using all households within a given demand response program. With such a regression, it would not be possible to disentangle whether the different treatment effects arise because

²⁰We cannot include an hour-of-sample fixed effect in the household-level regression because it would absorb the variation we are using for identification in this specification. We estimated our program-level specification in Equation (1) using this less granular set of fixed effects. The results closely reflect the estimates reported below. Finally, we verified that our household-level results are robust to alternative calendar fixed effects (e.g., year-month-hour, year-month-day-of-week).

²¹Recall that households can use the App to monitor their hourly household consumption and observe the timing and rewards for upcoming events 21 hours in advance. Households with installed devices can also monitor their device-level consumption in the App. Tech households can adjust their connected devices by pushing a button in the App (e.g., to turn off their use during events). Finally, Central households can adjust their connected devices and opt out of centralized load management during an event.

the household interacted with the App on an event day or whether the households that interact with the App are unique in the way they respond to events.

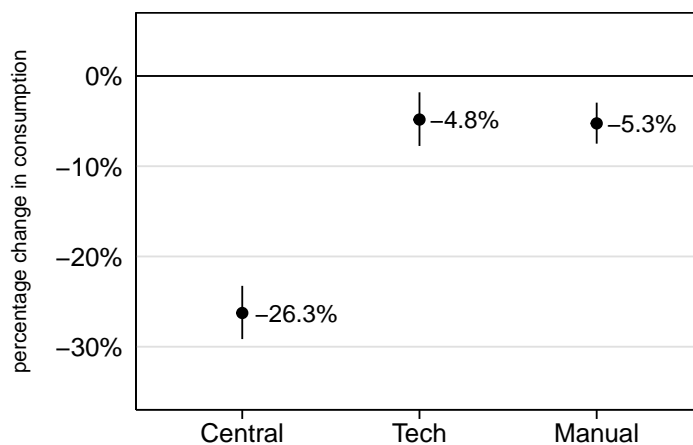
5 Empirical Results

This section presents the results of our demand response program-level econometric analyses. In particular, we provide the average treatment effect of events for each program, across all events and then separated by event type.

5.1 Program-Level Treatment Effects

Figure 3 provides the estimated average response to events by program as a percentage change in household-level consumption using the specification in Equation (1) including the full set of comparator households. We observe an average 26% reduction in consumption during events for the Central program. In contrast, the Tech and Manual programs reduced demand by approximately 5% on average during events. Both of these effects are significantly different from zero. Even though the Tech program had the same equipment as the Central program, it demonstrated a significantly lower response to events. Additionally, the average response for the Tech and Manual program are not significantly different.

Figure 3. Average Estimated Treatment Effects of Participants by Program



Notes: The reported results are program-specific marginal effects calculated from estimating $\hat{\beta}_j$ in (1) for $j \in \{C, M, T\}$. We present the marginal effects to be a percentage change in consumption using the transformation $100 \times (\exp(\hat{\beta}_j) - 1)$. Vertical lines indicate 95% confidence intervals. Standard errors are clustered at the household level.

Recall that the Central program has the ability to opt-out of events using the App. We observe an opt-out rate of only 4% at the event-connected device level from the participants in the Central program. When taken together with the results for the Tech and Manual programs, this low opt-out rate suggests that the large reductions for the Central program are primarily attributable to consumers allowing utility management of their devices during events.²²

These results align with the descriptive data presented in Section 3 that suggest that households in the Tech program did not use the load controller equipment to the same extent as the Central program. These results indicate that the installation of technology in the Tech program that enables remote control of household devices is likely to be insufficient to resolve the time allocation/effort barrier that limits the provision of demand flexibility. In contrast, centrally managed demand response that flips the default from an active to a passive response, thus requiring minimal-to-no effort, results in large demand reductions.

As described in Section 4.1, we undertake additional analyses to evaluate the validity of our identification strategy. Appendix Table C1 provides the results of our regression analysis when we vary the comparison group during non-event hours to estimate the event treatment effects. These results demonstrate that our estimated treatment effects are highly robust to varying the comparison groups included in the regressions. Appendix Table C5 provides the results for our test of the validity of our no carryover assumption. In this analysis, we find no evidence of changes in behavior during the event window on non-event days for any of our demand response programs.


5.2 Program-Level Treatment Effects by Event Type

In addition to randomized event timing, we also randomize event types, varying both the time of the event and the financial reward for reductions. This allows us to estimate how consumers respond to different price incentives and event times.

Figure 4 presents the estimated response to events allowing for differential responses by event type. For the Central program, we see a large demand reduction for all event types, with an approximate 27% reduction during morning events, 25% during evening events, and a 27% average reduction during high evening events. This indicates that the Central households allowed central management of demand during both morning and evening times. It also indicates that they were not distinctly more

²²Notably, when consumers opted out of central management, they generally did so for their thermostats: 90% of device opt-outs occurred by households adjusting thermostats during events.

Figure 4. Average Treatment Effect of Participants by Program and Event Type



Notes: The reported results are program- and event type-specific marginal effects calculated from estimating $\hat{\beta}_j$ in (1) for $j \in \{C, M, T\}$, adjusted to allow for event-type interactions with the program indicator variables D_i . We present the marginal effects to be a percentage change in consumption using the transformation $100 \times (\exp(\hat{\beta}_j) - 1)$. Vertical lines indicate 95% confidence intervals. Standard errors are clustered at the household level.

responsive to the greater incentives offered during the High Evening events.

During the Evening and High Evening events, the Tech program reduced its demand by approximately 6%, while the Manual program had a 4% estimated reduction in demand during these event types. These effects are statistically different from zero. The Evening and High Evening Tech and Manual program effects are not significantly different from each other, when compared within each event type.

The Tech program had a response to Morning events that are not statistically different from zero. This differs (statistically significantly) from the Manual program, which had an average estimated reduction of 8% during the Morning events. This is a counter-intuitive result, as the Tech program had all the same information, incentives, and abilities as the Manual program in making electricity consumption reductions during events, with the added ability to remotely control thermostats, EV chargers, and hot water heaters on which they have load controllers installed.

For all three programs, the change in consumption during High Evening events does not statistically significantly differ from their responses to regular Evening events. This suggests that the increased financial incentives does not motivate participants to undertake additional effort to make greater reductions in usage. This result suggests the barrier to demand responsiveness may have less to do with the scale of financial rewards and more to do with the hurdle of the opportunity costs of time/effort.²³

²³That said, participants were only eligible to receive \$1 more for achieving a 30% reduction in

Similar to the discussion in the previous section, Appendix Table C2 provides the results of our regression analysis when we vary the comparison group used to estimate event treatment effects. These results continue to demonstrate that our estimated treatment effects are largely robust to varying the comparison groups included in the regressions. The estimated response to morning events for the Tech and Manual programs varies with the relevant comparison group, with a smaller response for the Manual and a larger response for the Tech than our main specification. However, our key conclusions persist. The Central program is considerably more responsive to all event types. Further, no program shows a distinct response to the elevated incentives during the high evening events.

Finally, Appendix Table C6 provides the results for our test of the validity of our no carryover assumption, allowing for differential estimates for the morning and evening event windows. In this analysis, we find no evidence of changes in consumption decisions during either the morning or evening event windows on non-event days in the post-treatment period for any of our demand response programs.

6 Time, Effort, and Attention

Our results have so far focused on average estimated treatment effects by demand response programs. Overall, our main results suggest that active versus passive response and the associated differences in effort and attention play a key role in eliciting demand response. Participants in the Tech and Manual programs, both requiring active responses to financial incentives, reduced electricity much less on average during events, with similar behavior across the two programs.

A natural and important economic question is: What drives these differences across the three programs? In particular, what behavior underlies the differences between the responses to events of the Central compared to the Tech programs? In this section, we estimate household-level treatment effects and investigate the extent to which there is heterogeneity in event responsiveness. We then look at potential drivers of this heterogeneity. We use App interaction data during events to analyze the relationship between the frequency of engagement with the App and household-specific estimated treatment effects. Additionally, we use survey data on household


electricity use during a High Evening vs. Evening event. The reward for achieving a 50% reduction was doubled (\$6 vs. \$3). It is possible that a larger scaling of incentives could induce a greater response. However, given our rewards fall in the range of wholesale price caps observed in practice, it is unlikely that incentives provided in a real-world setting would be considerably larger than the amounts we provided.

income and the stated value of event participation to look at the relationship between the household’s opportunity cost of time/willingness to exert effort and its estimated response to events.

6.1 Household-Level Treatment Effects

Figure 5 presents histograms of household-level treatment effects by program, estimated as per Equation (2). We observe that the Tech and Manual programs’ treatment effects are tightly distributed near zero. On average, households in the Tech and Manual programs reduce their consumption during events by 4.6% and 4%, respectively. This corresponds closely to the estimated program-level treatment effect in the previous section. Only 20% and 18% of the household-level estimated treatment effects are negative and significant at the 5% level for the Tech and Manual programs, respectively.

Figure 5. Household-Level Estimated Treatment Effect Histograms by Program



Notes. The reported results summarize the distribution of estimated household-level treatment effects obtained from estimating specification (2). We present the marginal effects to be a percentage change in consumption using the transformation $100 \times (\exp(\hat{\beta}_i) - 1)$.

Despite the fact that the Tech program has enabling devices, the *distribution* of their estimated treatment effects closely resembles that of the Manual households. In

both the Tech and Manual programs, we observe a long left tail, indicating that there is a subset of households exhibiting large responsiveness. This is striking, as even though the Tech and Manual programs have similar mean responses when looking across all households, we might expect that the Tech program would contain a subset of households with higher household-specific event responses given their ability to remotely control their hot water heaters, EV chargers, and/or thermostats during events. There is only modest evidence to support this; the Manual households' estimated effects are more tightly distributed near zero.

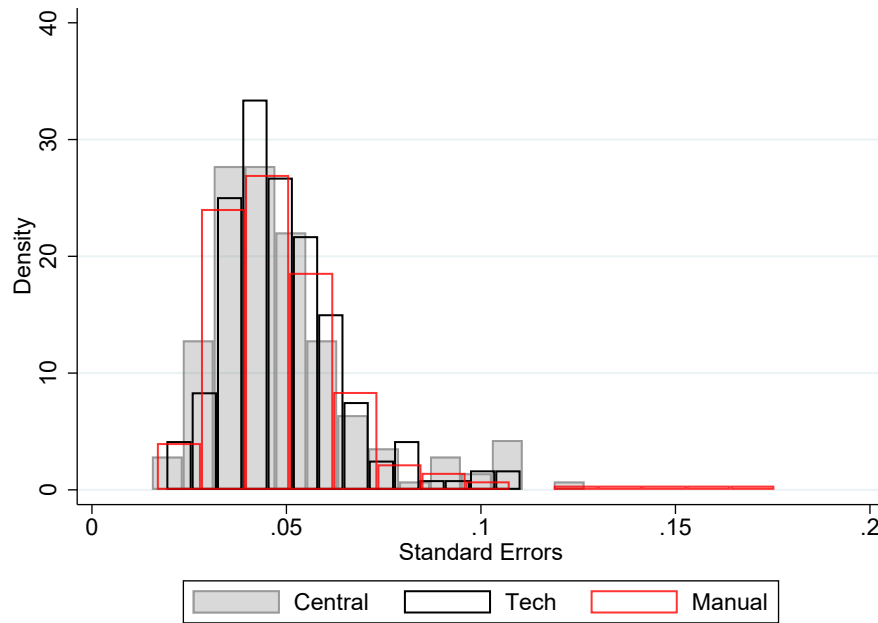
Figure 5 also re-confirms the larger response from households in the Central program. The average household-level treatment effect is a reduction of 24% during events for households in the Central program, again closely reflecting the estimated effects in the program-level regressions above. There is a broad range of household estimates in the Central program, consistent with the fact that there is variation in controllable household devices. However, only the largest household estimates in the Manual and Tech programs achieve reductions that are in the range that is typical in Central. In contrast to these programs, 80% of the Central household-level estimated treatment effects are negative and significant.²⁴

Figure 6 plots the distribution of the standard errors of the estimated household-level treatment effects. These results demonstrate that while the Central program has a slightly lower standard error on average, the distributions are similar across all three programs. The similarity in the precision of the estimates, combined with the average treatment effects noted above, suggest that the Central households are consistently large responders to events, while the Manual and Tech households are consistently low responders.²⁵ These findings provide important policy implications because they suggest that in addition to having larger estimated reductions in consumption during events, the Central program provides a reliable source of demand-side flexibility. Appendix C.6 provides additional empirical evidence demonstrating the consistency in the responses to events by demand response program throughout our experiment.

²⁴Appendix C.5 estimates household-level treatment effects by program and event type. We continue to find no evidence of a larger response to the elevated incentives during high peak events.

²⁵There are several large standard error estimates for the Manual program. This small subset of households has an average consumption reduction of 8.4% during events, more than double the average household treatment effect for the Manual program of approximately 4%. These large standard errors suggest that these households are relatively inconsistently large responders to events.

Figure 6. Standard Errors of the Household-Level Estimated Treatment Effects by Program



Notes. The reported results plot the standard errors of the estimated household-level treatment effects from specification (2).

6.2 App Interactions

In this section, we begin by allowing our household-level treatment effects to vary by whether or not the household interacted with the App. We then break down how the household’s performance during peak events varies by the intensity and types of App interactions.

6.2.1 Estimated Treatment Effects by App Interaction

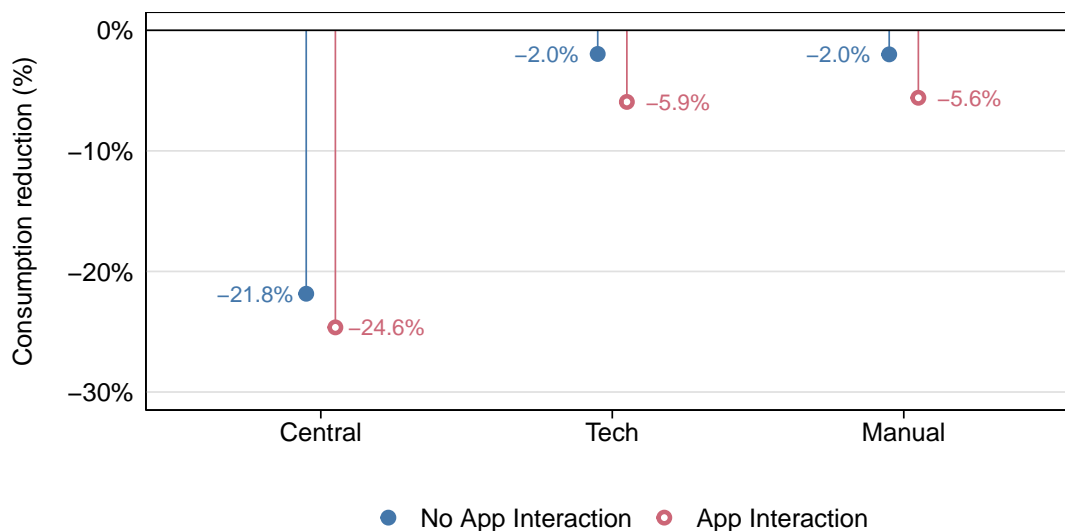
While households could be aware of an upcoming event without opening the App through push notifications sent to their phones, App interactions are a strong indication that the household is aware of the demand response event and serves as a proxy for time spent/effort to respond.²⁶ Despite the lower average response to events for the Tech and Manual households, Figure 5 demonstrates there is a small subset of high performers in the Tech and Manual programs that have large negative estimated

²⁶In addition to actively opening the App on their phone, the App opens and logs an interaction if a user clicks on the event push notification sent to their phone.

treatment effects. In this section, we leverage App interaction data to evaluate if these high-performers differ in their use of the App during events.

The App data tell us when users are interacting with the App on a given day, as well as more details on which features of the App they are accessing. For the Central and Tech households, interacting with the App allows them to control connected devices. For the Central program, the App can be used to opt out of automatic load control before or during events. In all programs, the App allows the household to observe the details of upcoming events 21 hours in advance, detailed information about household consumption in real-time, and performance in previous events.

Figure 7. Household-Level Average Treatment Effects by App Interaction



Notes. The reported results are average household-level treatment effects by whether or not the household interacted with the App on an event day. This represents the specification in equation (2), adjusted to interact the event indicator (E_{it}) with an App Indicator $_{it}$ variable that equals 1 when the household interacts with the App on the event day and zero otherwise. All specifications include fixed effects at the year-month, day-of-week, and hourly levels.

Figure 7 reports the average estimated household-specific treatment effects, allowing for heterogeneous treatment effects by whether or not the household interacted with the App on a given day. For the Central program, the average estimated household-level treatment effect is approximately -22% when the household does not interact with the App, increasing in magnitude to -25% when they do interact with the App during an event day. In contrast, for the Manual and Tech programs, the average household-level treatment effect is -2% when they do not interact with the App, increasing in magnitude to roughly -6% when they interact with the App.

The results shown in Figure 7 provide two key findings. First, there is a positive relationship between App interactions and demand reductions. Households who were more attentive to the App on event days achieved higher reductions. This differential effect is slightly larger for the Tech and Manual programs, where the change in demand reductions between interacting and not interacting with the App was approximately 5% as compared to about 3% for the Central program. Second, there is a massive difference in the “no interaction” estimates between Central program participants and those in the other programs (about -24% vs. -3%). This indicates that the Central program received a roughly 21 percentage-point “headstart” over the other programs by not having to take action during the events to respond to them. This “headstart” is key to the overall finding of greater demand response by the Central program participants, all of whom had a default setting to passively respond to events.

6.2.2 Intensity and Type of App Interactions

As noted above, we also have detailed data on the types of pages/tabs the households interact with in the App (i.e., General Interactions, Energy Usage Dial, Devices Tab, and Advisor Tab). This provides insight into the types of actions households may have taken to respond to events. Table 3 summarizes the average daily App interaction frequency for each demand response program separated by the performance quartile. The performance quartiles are determined by ranking household-level treatment effects over all demand response programs. This separates households into categories of whether or not they are high or low-performing households during our experiment. We summarize the count of the households and the percentage of households enrolled in each demand response program that fall within a specific performance quartile.

Table 3 demonstrates that the majority of Central households (63%) fall within the top quartile of performance. In contrast, only 11% and 8% of Tech and Manual households are in the top quartile. This is consistent with our results above that the Central program has a significantly greater demand response than the Tech and Manual programs, which show similar results.

Within each program, we see a reduction in the frequency of daily General Interactions as we move down the performance quartiles. This suggests that higher-performing households are more attentive and/or spend more time to monitor and react to events. In fact, across all columns, the frequency of App interactions in the top quartile is significantly greater (at the 5% level) than the frequency in the second-highest quartile.

Table 3. Average Daily App Interaction Frequency by Program and Performance Quartiles

Program	Performance Quartile	Household Count	General Interactions	Energy Usage Dial	Devices Tab	Advisor Tab
Central	1	112 (63%)	0.32	0.30	0.18	0.26
	2	42 (24%)	0.20	0.18	0.12	0.16
	3	16 (9%)	0.20	0.19	0.14	0.15
	4	7 (4%)	0.15	0.15	0.11	0.11
Tech	1	20 (11%)	0.54	0.49	0.32	0.42
	2	45 (25%)	0.30	0.27	0.17	0.23
	3	52 (28%)	0.23	0.22	0.11	0.17
	4	67 (36%)	0.13	0.12	0.07	0.09
Manual	1	19 (8%)	0.60	0.59	0.06	0.50
	2	64 (26%)	0.18	0.18	0.03	0.15
	3	83 (34%)	0.15	0.15	0.03	0.12
	4	76 (31%)	0.14	0.14	0.02	0.11

Notes. The reported results provide the daily frequency of App Interactions by program and performance quartile. Performance quartiles are determined using the household-specific estimated treatment effects obtained from estimating specification (2). Household Count represents the number of households that fall within each performance quartile. The percentages report the percentage of households within a program that falls within each quartile. General Interactions reflect any interactions with the App. Energy Usage Dial displays the energy dial in the App that provides data on real-time usage. The Devices Tab displays a household’s connected devices and allows households to adjust the use of the installed devices. The Advisor Tab reports information on upcoming events and historical performance on past events.

We see a considerably higher frequency of App interactions among the Tech and Manual households that achieve the top quartile of performance, interacting with the App on 54% and 60% of days on average, respectively. In contrast, Central households in the top quartile only interact with the App on 32% of the days, suggesting that achieving this high threshold of performance required less time and attention. The top quartile of the Central program interacts significantly less with the App than the top quartile of the Tech and Manual programs (at the 5% level). This is true for each App interaction column in the table.

Looking across the App categories, households interacted with the Energy Usage Dial the most, followed by the Advisor Tab. This suggests that when households used the App, they often monitored their real-time consumption and the details of upcoming events and/or their past performance in the Advisor Tab. The top quartile performers in the Tech and Manual programs stand out in both of these categories, with the frequency of their interactions with both the Energy Usage Dial and Advisor

Tab being roughly 2 to 4 times those of the bottom 3 quartiles.

Tech households in the highest-performing quartile interacted with the Devices Tab in the App at a significantly higher frequency than all other households. This suggests that a (small) subset of households in this program were using the installed devices to achieve larger demand reductions. However, looking at the lower performance quartiles in the Tech and Central programs, households in the Tech program interacted with the Devices Tab about as much as those in the Central program.²⁷

Overall, these results suggest that households in the Tech and Manual programs had to allocate a higher degree of their time/effort to achieve relatively high demand reductions and the associated rewards. In contrast, the Central program participants had to invest less time, effort, and attention in responding to events than others to generate the same scale of demand reductions.

6.3 End-of-Sample Survey

In this section, we use data collected from a survey at the end of our experiment (in June 2023) to further understand the potential mechanisms driving our results. Households who participated in the Central, Tech, and Manual programs were sent a survey and asked questions to better understand their opportunity cost of time/effort and stated net benefit to respond to events. We correlate households' responses to these survey questions with their estimated household-level treatment effects.

There were two key survey questions to assess households' willingness to exert effort/opportunity costs of time.²⁸ First, we asked participants their annual household income level and gave them several income categories (Less than \$50k, \$50-99k, \$100-149k, \$150-200k, and over \$200k). In the context of [Becker's \(1965\)](#) household production model, a household has a finite amount of time to allocate to tasks. The time spent allocating to a specific task, which in our setting includes the time/effort to respond to peak events, is weighted with the opportunity cost of undertaking other activities, such as working and earning income (e.g., the wage rate in the model). In this framing, a household's income serves as a key component of its opportunity cost of allocating effort to reduce consumption during peak events.

²⁷In addition to the installed devices from our experiment, a subset of households had other devices linked to the App. These were primarily smart plugs linked to lighting in the home. This helps explain why we observe interactions with the Devices Tab for the Manual program. Only 3 Manual households installed and linked devices to the App during our experiment, all three were thermostats for electric baseboard heaters.

²⁸Appendix [D.1](#) provides a more detailed summary of the exit survey.

Second, we asked a stated-preference question to assess the extent to which participants felt that participating in events was worth their time: “For the events you noticed, how often was it worth your time to participate by attempting to reduce your electricity consumption?” Respondents choose one of the following: (1) Never, (2) Sometimes, (3) About half the time, (4) Most of the time, or (5) Always. This question encompasses an arguably broader stated-preference assessment of a household’s perceived net benefit of allocating effort, time, and attention to responding to peak events.

We observe a high response rate to the survey, with 75%, 71%, and 69% of households in the Central, Tech, and Manual programs completing the survey. Respondents were paid \$20 upon completion of the survey. The results presented below focus on the households that completed the survey.²⁹

Using the survey responses and household-level estimated treatment effects, we estimate the following equation:

$$Y_i = \beta_0 + \beta_1 I_i + \beta_2 Z_i + \beta_3 G_i + \gamma X_i + \epsilon_i \quad (3)$$

in which Y_i is the estimated treatment effect for household i (estimates from Section 6.1), I_i is a household’s reported income, and Z_i is the household’s response to the Worth Time question described above.³⁰ We include several control variables to control for a household’s appliances and allocation to different demand response programs. G_i is an indicator variable for each demand response program and X_i is a vector of variables controlling for whether the household has an electric hot water heater, a categorical variable for whether a household has an electric vehicle and their corresponding charger type (No, Level 1, Level 2), a categorical variable for air conditioning in the home (No, Window Unit, Central Air), a categorical variable for electric baseboard heating (No, 1 – 3 Units, 4 or more units), and a house/duplex dummy variable (1 for house/duplex, 0 for row home). We report results with heteroskedastic robust standard errors.

²⁹Appendix D.2 compares the characteristics of households that responded to the survey with those that did not using pre-treatment data. While the households are similar on a number of characteristics, responders consumed less electricity on average. To the extent that consumption is correlated with factors that relate to the opportunity cost of time, this could downward weight households with a higher opportunity cost of time in the end of experiment survey.

³⁰We also considered a version of Equation (3) where I_i and Z_i are included separately in separate regressions. While the effects for each covariate are slightly larger, the conclusions reported below persist. This suggests that while these variables are correlated, they are capturing different features of a household’s willingness to allocate time, effort, and attention to responding to peak events.

Table 4. Household Treatment Effect on Worth Time and Income

	Coefficient	Std. Error	P-Value
<u>Worth Time</u>			
Sometimes	-3.49	2.13	0.10
Half of the Time	-6.30	2.39	0.01
Most of the Time	-9.03	2.35	0.00
Always	-15.05	2.86	0.00
<u>Income</u>			
50 - 99k	5.53	4.56	0.23
100 - 149k	7.32	4.55	0.11
150 - 200k	9.89	4.55	0.03
>200 k	9.74	4.50	0.03
<u>Program Indicators</u>			
Central	-17.34	1.74	0.00
Tech	-1.37	1.57	0.38
<u>Controls</u>			
Electric Hot Water Heater	-6.33	1.51	0.00
<u>Electric Vehicle</u>			
Yes, Level 1	-1.07	2.48	0.67
Yes, Level 2	-3.65	2.12	0.09
<u>Air Conditioning</u>			
Yes, Window Unit	1.69	1.74	0.33
Yes, Central Air	0.93	1.85	0.61
<u>Baseboard Heating</u>			
Yes, 1 - 3 Units	3.01	2.23	0.18
Yes, 4 or more	1.75	1.76	0.32
Home/Duplex	3.70	2.25	0.10

Notes. The reported results present the estimates for Equation (3).

The regression results are reported in Table 4. Focusing on the impact of household income, we see that higher income is associated with larger (positive) treatment effects, or smaller electricity consumption reductions during events. The magnitude of the coefficients are increasing in income, with less than \$50,000 per year being the excluded category. Coefficients for income brackets of \$150,000 per year or more are

statistically different than the excluded lowest income bracket.

For a given level of income, demand response program, and stock of household appliances, households may have heterogeneous preferences in their willingness to allocate time and effort to provide demand flexibility. Table 4 finds that all of the Worth Time variable categories are negative and significantly different than the excluded “Never” category (with “Sometimes” being marginally significant). Additionally, the coefficients are more negative as we move down the spectrum of the Worth Time measure. These results suggest that the participants’ perceived net benefits of allocating time to participate in events are correlated with greater reductions during events.

While these results focus on a stated-preference survey, they provide further insights to explain why we observe large differences in responses to peak events across our demand response programs. Factors that lead to a higher costs of allocating effort and attention are associated with lower estimated household-level treatment effects. The Central demand response program, through the use of enabling technologies and utility-controlled/automated default responses to events, helps overcome these time allocation and effort barriers.

7 Conclusion

The flexibility of electricity demand is becoming more valuable as electricity supply evolves to include a growing share of variable renewable sources. Moreover, there is a growing expectation that emerging technologies such as smart thermostats and electric vehicles will provide greater opportunities for flexible electricity demand. However, limited consumer responsiveness to dynamic electricity prices has long posed a problem for flexible demand to be meaningful.

Suspecting that inattention to dynamic pricing is rational—that consumers’ rewards for paying attention to dynamic electricity prices and learning about how to respond to them are not worth the associated costs—we run a large-scale field experiment that tests the efficacy of utility-managed (“passive”) electricity demand management of consumer responses to dynamic prices. Passive demand management has the potential to take the burden of responding actively to price signals off of consumers’ shoulders while allowing them (as well as other consumers and grid operators, in critical conditions) to reap the rewards of adjusting the timing of consumption with changing electricity system conditions.

We find that customers participating in an automated/utility-managed demand

management program, the Central program, reduced consumption by 26% on average during critical “peak events”. In contrast, participants in the Tech program, who had the same smart technology as those in the Central program to remotely control baseboard thermostats, hot water heaters, and electric vehicle chargers, but had to actively initiate reductions themselves, only reduced consumption by 5% during events. This difference indicates that automated/passive electricity demand management has large potential to help consumers overcome barriers to respond to electricity prices. We find that the take-up rates between the Central and Tech programs are not appreciably different, suggesting that automated electricity management is not as unpalatable as one might expect.

Somewhat surprisingly, we find that participants in the Tech program reduce consumption during events no more than those in our Manual program, who do not have smart, remote device adjustment capability. This indicates that smart home energy technology was not sufficient on its own to induce demand flexibility; overcoming the key barrier of effort/attention requires switching the default response, as per the Central program. Financial incentives motivate consumers to take some action, but consumers still face barriers that an automated/passive demand response program can resolve.

Interestingly, consumers in our context were not motivated to reduce consumption more when rewards were increased during periodic “high peak events”. It is possible that consumers need more than what we offered them to induce greater responsiveness and overcome the barriers that the Central program resolves. However, since our rewards were in line with peak electricity system prices, larger offers would not likely be economically efficient.

To understand the mechanisms behind our results, we use evidence from data on participant interaction with the experiment electricity management phone App. Across all programs, App interaction is correlated with larger household-level treatment effects. Average household-level treatment effects when households do not interact with the App are about 2% for the Tech and Manual programs and 22% for the Central program. When households *do* interact with their App, these numbers increase to about 6% for the Tech and Manual programs, and 25% for the Central program. This highlights the Central program’s “headstart”, whereby its participants achieve consumption reduction even in the absence of App interaction. We find that “high achievers” in the Tech and Manual programs, who drive the average consumption reductions during events for these programs, interact with their App on 60%

and 54% of days on average during the experiment, whereas the high achievers in the Central program interacted with the App on average significantly less (32% of days on average). This suggests that high achievers in the Tech and Manual program devoted a lot of attention to their electricity consumption and time/effort in reducing it during events. Taken together, this evidence points to attention and effort (in the form of app interaction) being an important component of responsiveness to events, and that the Central program relieved participants of needing to devote such time to electricity management to achieve large consumption reductions during events.

Given our results, we surmise that programs and policies that relieve consumers of cognitive, time, and other burdens that contribute to rational inattention will have large potential to lead to enhanced demand-side flexibility. In the case of residential electricity, we see utility-initiated/automated demand response as one such program, having the potential to both save money for consumers and facilitate flexible demand to meet emerging needs of electricity markets.

References

- Abadie, Alberto and Sebastien Gay**, “The impact of presumed consent legislation on cadaveric organ donation: a cross-country study,” *Journal of health economics*, 2006, *25* (4), 599–620.
- Bailey, Megan, David Brown, Erica Myers, Blake Shaffer, and Frank Wolak**, “Electric Vehicles and the Energy Transition: Unintended Consequences of Time-of-Use Pricing,” *American Economic Review: Insights*, 2025.
- Bailey, Megan R, David P Brown, Blake C Shaffer, and Frank A Wolak**, “Show Me the Money! Incentives and Nudges to Shift Electric Vehicle Charge Timing,” *American Economic Journal: Economic Policy*, 2025.
- Becker, Gary S**, “A Theory of the Allocation of Time,” *The Economic Journal*, 1965, *75* (299), 493–517.
- Bernheim, B Douglas, Andrey Fradkin, and Igor Popov**, “The welfare economics of default options in 401 (k) plans,” *American Economic Review*, 2015, *105* (9), 2798–2837.
- Boiteux, Marcel**, “Peak-Load Pricing,” *The Journal of Business*, 1960, *33* (2), 157–179.
- Bojinov, Iavor, Ashesh Rambachan, and Neil Shephard**, “Panel Experiments and Dynamic Causal Effects: A Finite Population Perspective,” *Quantitative Economics*, 2021, *12* (4), 1171–1196.
- Bollinger, Bryan K and Wesley R Hartmann**, “Information vs. Automation and Implications for Dynamic Pricing,” *Management Science*, 2020, *66* (1), 290–314.
- Brandon, Alec, Christopher M Clapp, John A List, Robert D Metcalfe, and Michael Price**, “The Human Perils of Scaling Smart Technologies: Evidence from Field Experiments,” Technical Report, National Bureau of Economic Research 2022.
- Brot-Goldberg, Zarek, Timothy Layton, Boris Vabson, and Adelina Yanyue Wang**, “The behavioral foundations of default effects: theory and evidence from Medicare Part D,” *American Economic Review*, 2023, *113* (10), 2718–2758.
- Brown, David P and Derek EH Olmstead**, “Measuring Market Power and the Efficiency of Alberta’s Restructured Electricity Market: An Energy-Only Market Design,” *Canadian Journal of Economics*, 2017, *50* (3), 838–870.
- Burkhardt, Jesse, Kenneth T Gillingham, and Praveen K Kopalle**, “Field experimental evidence on the effect of pricing on residential electricity conservation,” *Management Science*, 2023, *69* (12), 7784–7798.

- Charness, Gary, Uri Gneezy, and Michael A Kuhn**, “Experimental methods: Between-subject and within-subject design,” *Journal of economic behavior & organization*, 2012, 81 (1), 1–8.
- Chetty, Raj, John N Friedman, Søren Leth-Petersen, Torben Heien Nielsen, and Tore Olsen**, “Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from Denmark,” *The Quarterly Journal of Economics*, 2014, 129 (3), 1141–1219.
- DTE Energy**, “DTE Energy Smart Charge Program,” 2022.
- Einav, Liran, Ben Klopach, and Neale Mahoney**, “Selling subscriptions,” *American Economic Review*, 2025, 115 (5), 1650–1671.
- Fowlie, Meredith, Catherine Wolfram, Patrick Baylis, C Anna Spurlock, Annika Todd-Blick, and Peter Cappers**, “Default Effects and Follow-on Behaviour: Evidence From an Electricity Pricing Program,” *The Review of Economic Studies*, 2021, 88 (6), 2886–2934.
- Harding, Matthew and Carlos Lamarche**, “Empowering Consumers Through Data and Smart Technology: Experimental Evidence on the Consequences of Time-of-Use Electricity Pricing Policies,” *Journal of Policy Analysis and Management*, 2016, 35 (4), 906–931.
- **and Steven Sexton**, “Household response to time-varying electricity prices,” *Annual Review of Resource Economics*, 2017, 9 (1), 337–359.
- IRC**, “Market Design Executive Summary: 2017 IRC Markets Committee,” Technical Report, ISO/RTO Council 2017.
- Joskow, Paul L and Catherine D Wolfram**, “Dynamic Pricing of Electricity,” *American Economic Review*, 2012, 102 (3), 381–385.
- Kahn, Alfred E.**, *The Economics of Regulation: Principles and Institutions: Volume 1: Economic Principles*, John Wiley, 1970.
- List, John A.**, “The Experimentalist Looks Within: Toward an Understanding of Within-Subject Experimental Designs,” Technical Report, National Bureau of Economic Research 2025.
- Nauze, Andrea La, Lana Friesen, Kai Li Lim, Flavio Menezes, Lionel Page, Thara Philip, and Jake Whitehead**, “Can Electric Vehicles Aid the Renewable Transition? Evidence from a Field Experiment Incentivising Midday Charging,” Technical Report, CESifo Working Paper 2024.
- PG&E**, “SmartAC Switches,” 2023. https://www.pge.com/en_US/residential/save-energy-money/savings-solutions-and-rebates/smart-ac/smartac-switch.page. Accessed: 2-27-2023.

- Prest, Brian C**, “Peaking interest: How awareness drives the effectiveness of time-of-use electricity pricing,” *Journal of the Association of Environmental and Resource Economists*, 2020, 7 (1), 103–143.
- Sallee, James M**, “Rational Inattention and Energy Efficiency,” *The Journal of Law and Economics*, 2014, 57 (3), 781–820.
- Schneider, Ian and Cass R Sunstein**, “Behavioral Considerations for Effective Time-Varying Electricity Prices,” *Behavioural Public Policy*, 2017, 1 (2), 219–251.
- Shepard, Mark and Myles Wagner**, “Do ordeals work for selection markets? Evidence from health insurance auto-enrollment,” *American Economic Review*, 2025, 115 (3), 772–822.
- Smith, Jordan**, “Texas Regulators Lower Electric Price Cap to \$5,000 per Megawatt Hour,” 2022. <https://www.chooseenergy.com/news/article/texas-regulators-lower-electric-price-cap/>.
- Spector, Julia**, “Portland General Electric Jumps Into Virtual Power Plant Business,” 2020. <https://www.greentechmedia.com/articles/read/portland-general-electric-launches-525-home-battery-fleet-virtual-power-plant>.
- Statistics Canada**, “Census of Population. Data Products, 2016 Census,” 2021. <https://www12.statcan.gc.ca/census-recensement/2016/dp-pd/index-eng.cfm>, Accessed: 2021-07-15.
- Wattersaver**, “WatterSaver Program,” 2023. <https://www.watter-saver.com/>. Accessed: 2-27-2023.

A Supplementary Experimental Framework Material

A.1 Recruitment and Assignment

The study sample was drawn from the population of residential customers in the Utility’s service territory in and near a large metropolitan city in Canada. We employed a two-step recruitment strategy. In Phase 1, starting in August 2021, the Utility invited households to join an App operated by a third-party company in partnership with the Utility. The App provides households with household-level hourly consumption posted at a one-day lag. The App can be coupled with other devices to provide more detailed information on real-time usage and device control. Households were recruited to the App using several marketing strategies, including advertisements on the Utility’s website, social media posts, the Utility’s newsletter, website notifications when users logged into their Utility accounts, and emails sent to approximately 306,000 residential households.

The recruitment onto the Utility’s App provided us with a pool of 9,020 households to draw from. When households signed up to join the App, they were required to answer a six-question survey. The survey asked households about their motivation for joining the App and whether the household rents or owns their home. It asked about devices eligible for load control in our experiment, including whether the household has an electric hot water tank, an electric vehicle (EV), and electric baseboard heaters as the primary heat source. Households with EVs were asked what type of charger (level 1 or 2) they use. It also asked whether households have air conditioning, a major source of demand flexibility.

We applied several selection criteria to this pool of households. Customers had to be in or near a large metropolitan city in the province for which it was feasible for Utility-partnered electricians to install load control equipment, as needed. Only homeowners were permitted to participate. Condos and apartments were removed, leaving primarily single-family homes, duplexes, and row homes as eligible. Households must have at least one month of historical consumption data as of September 2021, and the customers must have at least one controllable electric device. Recall, the set of controllable electric devices includes a level 2 electric vehicle charger, electric baseboard heaters used as the primary heat source, and an electric hot water heater tank. This left us with a sample of 1,661 potential households that we used for our randomized assignment to experimental programs.

In Phase 2 of recruitment, we randomized the eligible households into our treat-

ment programs and never-treated groups.³¹ Starting in October 2021, we sent program-specific emails to households inviting them to join a new “Trial” program. These emails provided details about the specific experiences households would face in the program to which they were being invited, including a summary of the expected rewards they could earn over the course of the Trial, equipment they would receive and its estimated value, and future peak events. Households were also randomly offered a small sign-on incentive of the amounts \$10 or \$20, or no incentive. All households faced a yes/no decision regarding accepting our program-specific offer. The never-treated Control group that received no equipment, price incentives, or real-time usage information (recall Table 1) received no further communication beyond joining the App in the first phase of recruitment.

Households had to accept the invitation to join the relevant experimental program actively. After selecting to join, households were mailed a device called the “Hub” that facilitates monitoring real-time energy usage via the App. Installers contacted households in the Central and Tech programs to install the load controller equipment.

This two-phase recruitment process occurred over the months of August 2021 - February 2022. The second phase of recruitment occurred in five waves starting in October 2021. As additional households joined the App, we collected the survey responses, identified eligible households, randomized households into programs, and sent the second-phase recruitment emails. This process was used to facilitate the time required to schedule and install the load controllers, as well as to achieve the targeted sample size.

Finally, during the invitations to join each program, we randomized the upfront incentive. While we find a higher rate of initial acceptance with higher upfront incentive payments, the differences are small and not significantly different.³²

³¹Specifically, we used a randomization procedure designed to balance important observable characteristics over programs and groups. We first used the machine learning algorithm “kmeans” to group households based on observable characteristics. These included cumulative household electricity consumption (in kWh) and load factor by season (Fall, Spring, Winter, and Summer), variables that indicate if a household has an electric vehicle, electric baseboard heating, or air conditioning, and census data on median household income. Load factor is the average electricity consumption divided by maximum consumption over a specific time period; it is a way to capture the relative utilization rate of consumption at the household level. We then randomized program assignments so households within a cluster were balanced across programs.

³²Households that received a \$0, \$10, and \$20 upfront incentive accepted the initial invitation with a 63%, 67%, and 68% probability, respectively.

A.2 Comparison of Household Characteristics Upon Randomization

We evaluate if there are differences in pre-treatment characteristics across our various programs to assess the quality of our randomization. Table A1 provides summary statistics by program for a number of variables, including those used in the clustering procedure during randomization (recall the discussion in Footnote 31). The sample presented in this Table represents all 1,661 households invited to participate in the experiment. For all variables, we report the p-values from a one-way ANOVA test to evaluate if there are statistical differences in means across the programs.³³

Table A1 shows that we do not find significant differences in key characteristics pre-treatment across our programs. These results indicate that our randomization approach effectively achieved balance on observables pre-treatment. In addition, Table A1 demonstrates that the majority of households in our sample have electric hot water heating and use baseboard heating as the primary heat source. In contrast, electric vehicles are less common, representing approximately 30% of households. The majority of households are single-family homes or duplexes, with the remainder being primarily row homes. The households consume considerably more electricity during the winter, with the lowest consumption arising in summer. This demonstrates the potential for larger opportunities for load shifting during these months.

A.3 Comparison of Household Characteristics After Acceptance

We compare the pre-treatment means in observable characteristics by program, including only the households that accepted our invitation to join each program. Large differences in observable characteristics would raise questions about the comparability of our estimated treatment effects from the main specifications.

Table A2 shows observables across programs for the final set of households included in each program. We observe limited differences in these characteristics across programs. The exceptions are that we find a statistically significant difference in the proportion of households that live in single-family homes/duplexes. There is a larger proportion of households in this building type in the Manual program than in other programs, in particular. We also observe a difference across programs in the proportion of households that have EVs, but this difference is only marginally statistically

³³The seasonal cumulative consumption and load factor data only contain households with a full year’s worth of historical consumption. We computed analogous statistics for the entire sample of households using only data from September 2021, the month in which all households have complete pre-treatment consumption data. We find no evidence of statistically significant differences in means across the programs using this data.

Table A1. Comparison of Means by Programs - Initial Randomization

	Central	Tech	Manual	Info	Control	ANOVA (p-value)
Cumul. kWh						
Winter	5,279 (2,694)	5,268 (3,032)	5,442 (3,076)	4,859 (2,748)	5,265 (2,950)	0.27
Spring	3,760 (1,924)	3,773 (2,112)	3,818 (1,911)	3,503 (2,116)	3,712 (1,974)	0.48
Summer	2,845 (1,742)	2,836 (1,872)	2,708 (1,539)	2,614 (1,861)	2,729 (1,710)	0.54
Fall	3,633 (1,663)	3,670 (1,945)	3,700 (1,974)	3,458 (1,796)	3,623 (1,860)	0.66
Load Factor						
Winter	24.66 (8.20)	24.98 (8.15)	25.41 (8.80)	24.73 (8.29)	24.67 (8.63)	0.81
Spring	19.52 (7.25)	20.12 (6.97)	20.01 (6.70)	19.28 (7.73)	19.91 (7.41)	0.65
Summer	16.82 (7.89)	16.55 (6.30)	16.73 (5.93)	16.12 (8.11)	16.32 (8.29)	0.82
Fall	18.56 (5.89)	18.90 (6.23)	19.34 (6.00)	18.42 (6.48)	19.06 (6.50)	0.42
Electric Vehicle	0.27 (0.44)	0.27 (0.45)	0.27 (0.45)	0.33 (0.47)	0.27 (0.45)	0.41
Baseboard Heating	0.61 (0.49)	0.64 (0.48)	0.61 (0.49)	0.63 (0.48)	0.63 (0.48)	0.95
Air Conditioning	0.52 (0.50)	0.51 (0.50)	0.50 (0.50)	0.51 (0.50)	0.54 (0.50)	0.95
Electric Hot Water	0.70 (0.46)	0.66 (0.47)	0.70 (0.46)	0.66 (0.47)	0.72 (0.45)	0.38
House/Duplex	0.77 (0.42)	0.76 (0.43)	0.81 (0.39)	0.78 (0.42)	0.84 (0.37)	0.17
Median Income	86,376 (19,503)	88,291 (22,227)	85,931 (19,255)	87,470 (21,574)	85,948 (21,541)	0.48
Households	423	382	409	259	188	

Notes. This table compares pre-treatment average values across the five different programs. Parentheses contain the standard deviations. Cumul. kWh and Load Factor represent the cumulative household-level consumption and load factor by season. The seasonal cumulative consumption and load factor data only contain households with a full year's worth of historical consumption. Electric Vehicle, Baseboard Heating, Air Conditioning, and Electric Hot Water are indicator variables denoting the presence of each device. House/Duplex is an indicator variable that equals one if the home type is a single-family home or duplex and zero otherwise. Median Income reports the median household-level income of the Census Dissemination Area where the household is located. ANOVA reports the p-value from one-way ANOVA tests for differences in means across programs. Statistical significance: * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

significant. Overall, these results suggest that the balance on observables that arose due to the initial randomization largely remains in the final sample.

Table A2. Comparison of Means by Program - Final Accepted Households

	Central	Tech	Manual	Info	Control	ANOVA (p-value)
Cumul. kWh						
Winter	5,507 (2,706)	5,302 (2,737)	5,422 (3,240)	5,037 (2,768)	5,265 (2,950)	0.71
Spring	3,900 (1,934)	3,739 (1,791)	3,797 (1,939)	3,642 (2,159)	3,712 (1,974)	0.85
Summer	2,851 (1,869)	2,672 (1,759)	2,766 (1,547)	2,702 (1,849)	2,729 (1,710)	0.93
Fall	3,754 (1,733)	3,550 (1,659)	3,677 (1,992)	3,547 (1,788)	3,623 (1,860)	0.86
Load Factor						
Winter	24.62 (8.68)	25.56 (8.21)	24.93 (9.04)	24.93 (7.76)	24.67 (8.63)	0.90
Spring	19.33 (7.43)	20.48 (6.30)	19.80 (6.45)	19.59 (7.05)	19.91 (7.41)	0.72
Summer	16.33 (8.54)	16.87 (6.02)	16.95 (5.95)	16.61 (7.80)	16.32 (8.29)	0.91
Fall	18.17 (6.27)	18.97 (5.78)	19.11 (6.29)	18.53 (6.11)	19.06 (6.50)	0.65
Electric Vehicle	0.25 (0.43)	0.21 (0.41)	0.30 (0.46)	0.34 (0.47)	0.27 (0.45)	0.07*
Baseboard Heating	0.68 (0.47)	0.70 (0.46)	0.60 (0.49)	0.59 (0.49)	0.63 (0.48)	0.12
Air Conditioning	0.46 (0.50)	0.46 (0.50)	0.51 (0.50)	0.51 (0.50)	0.54 (0.50)	0.41
Electric Hot Water	0.75 (0.43)	0.74 (0.44)	0.68 (0.47)	0.65 (0.48)	0.72 (0.45)	0.16
House/Duplex	0.82 (0.39)	0.77 (0.42)	0.89 (0.32)	0.84 (0.37)	0.84 (0.37)	0.02**
Median Income	84,978 (19,647)	88,274 (20,432)	86,718 (19,494)	89,504 (21,079)	85,948 (21,541)	0.23
Households	177	184	242	177	188	

Notes. This table compares pre-treatment average values across the five different programs for households that were in our final programs. Parentheses contain the standard deviations. Cumul. kWh and Load Factor represents the cumulative household-level consumption and load factor by season. Electric Vehicle, Baseboard Heating, Air Conditioning, and Electric Hot Water are indicator variables denoting the presence of each device. House/Duplex is a indicator variable if the home type is a single-family home or duplex. Median Income reports the median household-level income of the Census Dissemination Area where the household is located. ANOVA reports the p-value from one-way ANOVA tests for differences in means across programs. Statistical Significance * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

B Treatment Details

B.1 Program-Specific Event Notifications

Each treatment program experienced event notifications tailored to their treatment. Each program received a notification 21 and 2 hours before an event. All participants were shown a short notification according to their device and in-app notification settings. If participants touched and pressed the notification, they were shown the long notification specific to their program, featured below, with event incentive details.

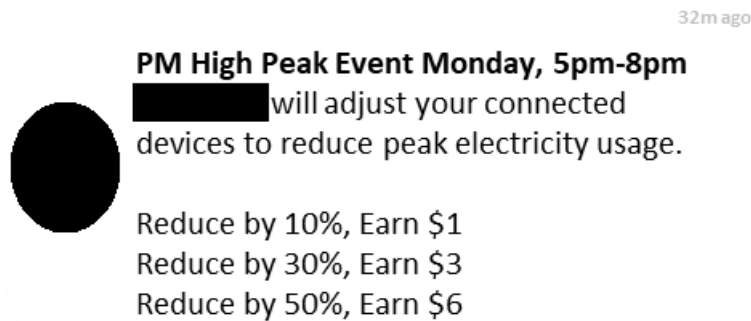


Figure B1. Long Notification for Central program

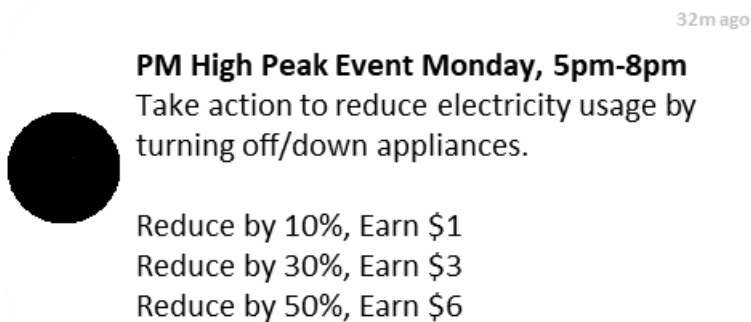


Figure B2. Long Notification for Tech program

32m ago

PM High Peak Event Monday, 5pm-8pm

Take action to reduce electricity usage by turning off/down appliances.



Reduce by 10%, Earn \$1

Reduce by 30%, Earn \$3

Reduce by 50%, Earn \$6

Figure B3. Long Notification for Manual program

Note that all program participants in the three programs were able to locate event details in the “Advisor” tab of the App, a centralized location for information from the App company, once they received an event notification. The “Learn More” button at the bottom right of this information card took participants to the “FAQs” section of the program-specific experiment website.

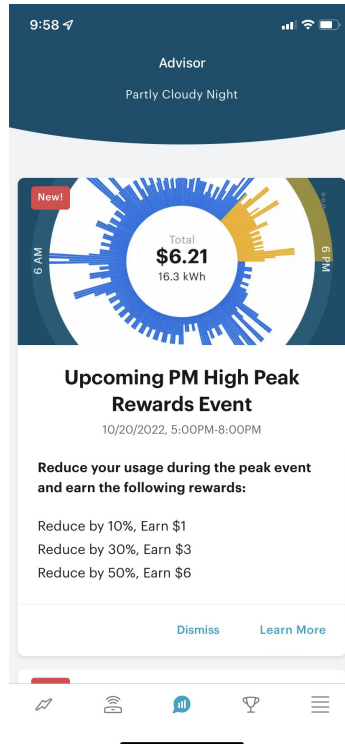


Figure B4. Event info in App

B.2 Treatment Program-Specific App Functionality

Each program in our experiment had an App experience and functionality that differed according to their program assignment. We detail that here and walk through how participants in each program could have responded to peak events, given the options in the App.

B.2.1 Central Program

The Central program participants receive 21-hr and 2-hr notifications regarding upcoming events, as described in Appendix B.1. These notifications allow them to see the timing of the event and the magnitude of rewards for electricity consumption reductions. They also remind participants that their devices with load controllers would be altered by the Utility to reduce consumption, unless they opted-out of the event.

There are several ways that Central program participants can opt-out of events. Before an event starts, they can push an “Opt-out” button in the “My Devices” tab of the App (Figure B5). (This tab is a central App location that allows App users to remotely control devices that have load controllers and see the individual electricity consumption of those devices.) This button removes the participant from the event globally by removing all of their load-controlled devices from the event.

If they do not opt-out in this way, they see a series of screens in the “My Devices” tab. These indicate the progression of the event to the participant and signal when their devices’ electricity consumption is being controlled by the utility, via the icons above the text “You are opted in”, “Event”, and “Complete” (Figure B6).

During an event, participants can cancel Utility device control in a device-specific way. For EV chargers and hot water heaters, they can remotely opt-out their device from being controlled, or they can physically turn off the load controller at the device itself. For thermostats, participants can opt-out of load control by adjusting them physically or remotely through the App, during an event.

Note that the Central program has remote and manual control of all devices with load controllers, just like the Tech program. Central program households can also change anything else in the house to alter their electricity consumption during events.

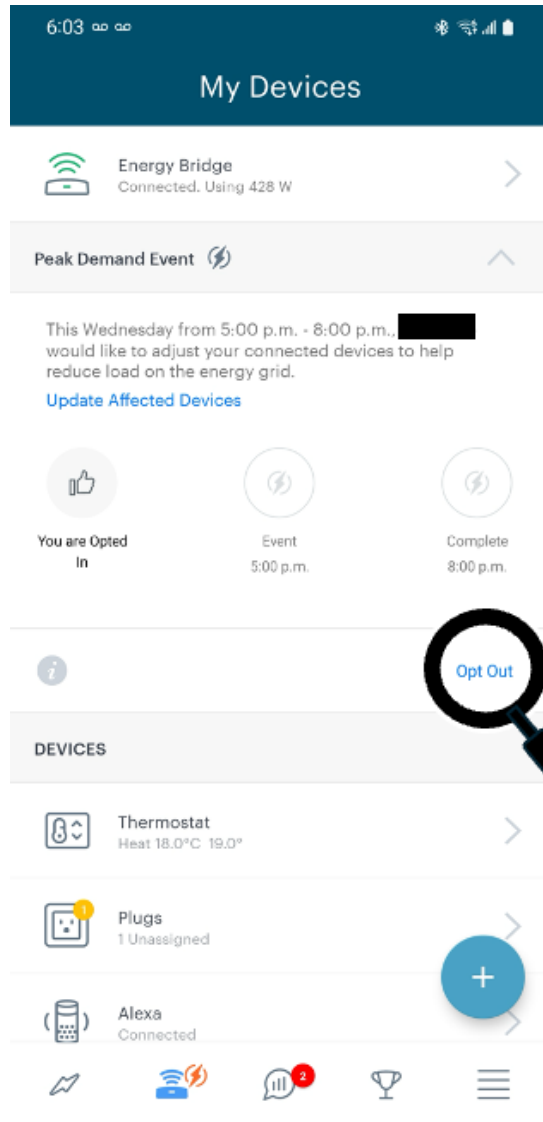


Figure B5. Central Program Opt-Out Functionality

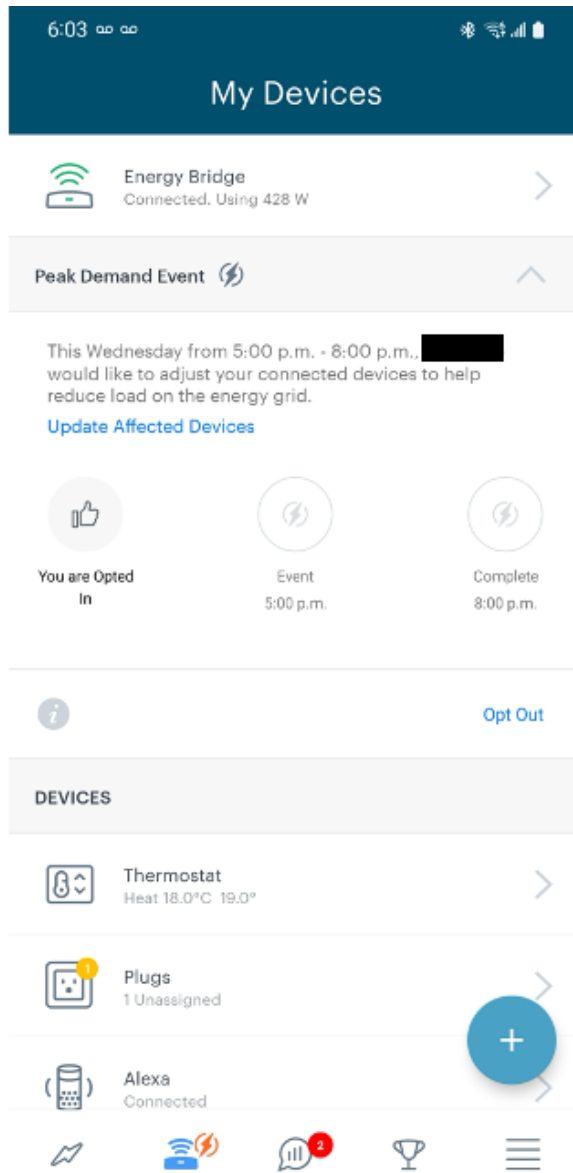


Figure B6. Central Program Event Experience


B.2.2 Tech Program

The Tech program participants receive 21-hr and 2-hr notifications regarding upcoming events, as described in Appendix B.1. These notifications allow them to see the timing of the event and the magnitude of rewards for electricity consumption reductions. They also remind participants that they need to “take action” to make consumption changes to receive the rewards offered.

The Tech program can remotely control any device that has an installed load controller through the App. For EV chargers and hot water heaters, they can turn them off via two clicks from the My Devices section of the App. (See Figure B7 below for the instructions sent to participants that explain these actions.) Tech program participants cannot make a schedule to turn off these devices before events start and must turn them off before or during events to reduce consumption this way. (They must also remember to turn them on unless they set up a turn-on schedule.)

All smart devices, including smart plugs and load controllers can be set up and controlled through the “My Devices” page in the [redacted] app.

Select this icon  at the bottom of your app to go to the “My Devices” page.



On the “My Devices” page your plugs should appear here.
Select “Plugs”

Turning Plugs On or Off



By pressing the icon you can turn a plug on or off.



Orange is on

White is off

Figure B7. Controller Guide for Tech Program

For thermostats, the Tech program can set up schedule for their thermostat set-

point before events, using the App. They can also adjust their thermostats remotely during events with the App.

B.2.3 Manual Program

The Manual program participants receive 21-hr and 2-hr notifications regarding upcoming events, as described in Appendix B.1. These notifications allow them to see the timing of the event and the magnitude of rewards for electricity consumption reductions. They also remind participants that they need to “take action” to make consumption changes to receive the rewards offered.

Manual program participants do not load controllers given to them as part of this experiment or Utility control of any devices. They therefore only observe these notifications as well their aggregate, real-time household consumption through the App. If Manual program participants install their own smart home devices, they may be able to link them to the smart electricity consumption technology ecosystem used in this experiment. If so, they may have the capabilities of the Tech program to observe the real-time consumption of those devices/devices individually and adjust them remotely through the App. (Only three households in the Manual program installed their own smart thermostats over our sample period.)

B.2.4 Central, Tech, and Manual Programs

After each event, all three of the Central, Tech, and Manual programs receive a result on their performance, as depicted below. This appears in the “Advisor” tab of the App, a central location for information from the App company. This result card reminds participants of the event type (reward magnitudes being “high” or not) and the day and time of the event. It shows the incremental reward the participant earned from the event as well as their cumulative rewards throughout the entire experiment, including the reward from the prior event. The text below the reward for the last event is variable and depends on whether a participant met one of the reward tiers. The rewards screen with one of these text options is shown below in Figure B8.

From this rewards screen, participants can select “Event History” and see their recent history of event rewards, as shown in Figure B9.³⁴

³⁴Figure B9 was created for illustrative purposes using a series of simulated events. As a result, the event times differ from the event times considered in our study (i.e., 7:00 AM - 10:00 AM and 5:00 PM - 8:00 PM).

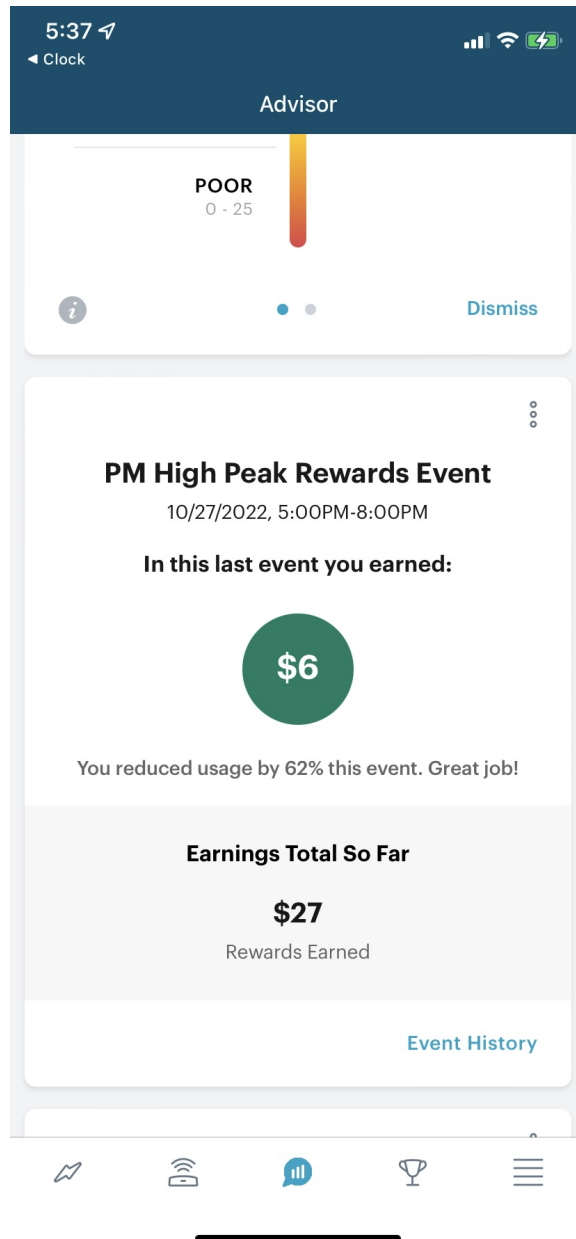


Figure B8. Rewards Screen




Figure B9. Event History

C Extensions and Robustness

C.1 Comparison Group in Program-Level Regression

In our main specification in Equation (1), our analysis includes all three demand response groups and never-treated households. As discussed in Section 4.1, this approach compares event time consumption to non-event time consumption of households in the same demand response program, other demand response programs, and never-treated households. Tables C1 and C2 present our estimated treatment effects of participants by program to events and separated by event type, respectively, allowing for regressions only including households in the same demand response program (Column (1)), same demand response program and the never-treated (Column (2)), and the results from our main specification in Column (3) for comparison purposes. Our results are consistent across all three specifications.

Table C1. Treatment Effects of Participants by Program

Program	(1)	(2)	(3)
Central	-0.3151*** (0.0206)	-0.3007*** (0.0204)	-0.3047*** (0.0204)
Tech	-0.0661*** (0.0132)	-0.0475*** (0.0155)	-0.0495*** (0.0159)
Manual	-0.0507*** (0.0092)	-0.0459*** (0.0117)	-0.0540*** (0.0122)
Comparisons			
Own Program	Y	Y	Y
Other Treated			Y
Never Treated		Y	Y

Notes. The reported results are program-specific treatment effect coefficients. Standard errors are reported in the parentheses and clustered at the household level. Column (1) reports the regression results including only within demand response program comparisons, column (2) includes both within demand response program and the never-treated (Info and Control) groups, and column (3) reports results when all programs/groups are included. All specifications include fixed effects at the household and hour-of-sample levels. Statistical Significance * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

Table C2. Treatment Effects of Participants by Program and Event-Type

Program	(1)	(2)	(3)
Central			
Morning	-0.3198*** (0.0221)	-0.3160*** (0.0224)	-0.3133*** (0.0223)
Evening	-0.3040*** (0.0218)	-0.2850*** (0.0226)	-0.2916*** (0.0227)
High Evening	-0.3291*** (0.0237)	-0.3105*** (0.0241)	-0.3177*** (0.0243)
Tech			
Morning	-0.0495*** (0.0128)	-0.0273 (0.0188)	-0.0158 (0.0199)
Evening	-0.0713*** (0.0155)	-0.0545*** (0.0181)	-0.0623*** (0.0185)
High Evening	-0.0786*** (0.0166)	-0.0604*** (0.0199)	-0.0675*** (0.0205)
Manual			
Morning	-0.0488*** (0.0097)	-0.0689*** (0.0150)	-0.0812*** (0.0159)
Evening	-0.0488*** (0.0108)	-0.0332** (0.0143)	-0.0396*** (0.0149)
High Evening	-0.0567*** (0.0126)	-0.0403** (0.0159)	-0.0467*** (0.0166)
Comparisons			
Own Program	Y	Y	Y
Other Treated			Y
Never Treated		Y	Y

Notes. The reported results are program-specific treatment effect coefficients by event type. Standard errors are reported in the parentheses and clustered at the household level. Column (1) reports the regression results including only within demand response program comparisons, column (2) includes both within demand response program and the never-treated (Info and Control) groups, and column (3) reports results when all programs/groups are included. Each regression is adjusted to include an event-type-specific categorical variable. All specifications include fixed effects at the household and hour-of-sample levels. Statistical Significance * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

C.2 Pre-Event and Post-Event Spillover Effects

In this section, we expand our main specification in Equation (1) to include indicator variables for two hours prior to and after peak events. More specifically, we include additional terms $\sum_{j \in \{C, T, M\}} \text{Program}_{jt} \cdot \text{Pre-Event}_{it}$ and $\sum_{j \in \{C, T, M\}} \text{Program}_{jt} \cdot \text{Post-Event}_{it}$ in Equation (1), where Pre-Event_{it} (Post-Event_{it}) is an indicator variable that equals 1 for the two hours prior to (after) a peak event and zero otherwise. This allows us to estimate if there are treatment effects running up to or after peak events. In addition, this helps alleviate concerns that our main estimate of interest, treatment effects during peak event hours, are robust to controlling for such spillover effects. We consider a two hour window to capture the anticipated “snap back” effects illustrated in Figure 2. However, our results are robust to the consideration of one and three hour pre- and post-event windows.

Table C3 presents estimates for the event coefficients by program. Comparing our results to those in our main analysis (column (3) in Table C1), in column (2) we see that our main treatment effect estimates are systematically unchanged by the inclusion of these additional covariates. We continue to observe large estimated effects for Central, and similar smaller effects for Tech and Manual. Looking at column (1), we find no evidence of a treatment effect prior to events. In the Central program, there was no use of load control (e.g., to pre-cool) prior to events. In column (3), we find a large and significant positive effect for the Central program. There are no significant effects post-event for Tech or Manual. These results are consistent with the “snap back” effect observed for the Central program in Figure 2.

Table C4 extends the spillover effects analysis by permitting differential effects by event type. Comparing our results to those in our main analysis (column (3) in Table C2), in column (2) we see that the conclusions drawn from our main treatment effect estimates are again unchanged by the inclusion of these additional covariates. Looking at column (1), there is no evidence of pre-event treatment effects across all programs and event types. In contrast, for the post-event effects, there is a large and significant positive effect for the Central program for all event types. Tech continues to have no evidence of any post-event effects. For the Manual program, we find some evidence that the estimated treatment effect during morning events persisted after the events ended. This could potentially be explained by households responding to events and leaving the household devices/appliances off after events throughout the morning.

Table C3. Treatment Effects of Participants by Program - Spillover Effects

Program	Pre-Event	Event	Post-Event
Central	0.0056 (0.0123)	-0.3036*** (0.0204)	0.2063*** (0.0195)
Tech	0.0163 (0.0118)	-0.0496*** (0.0160)	0.0034 (0.0121)
Manual	0.0030 (0.0102)	-0.0542*** (0.0123)	-0.0031 (0.0101)

Notes. The reported results are program-specific treatment effect coefficients in (1), with the addition of demand response group-specific Pre-Event and Post-Event covariates. Pre-Event (Post-Event) is an indicator that equals 1 for two hours prior to (after) peak events, and zero otherwise. Standard errors are reported in the parentheses and clustered at the household level. All specifications include fixed effects at the household and hour-of-sample levels. Statistical Significance * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

Table C4. Treatment Effects of Participants by Program and Event Type-Spillover Effects

Program	Pre-Event	Event	Post-Event
Central			
Morning	0.0166 (0.0196)	-0.3122*** (0.0223)	0.2104*** (0.0238)
Evening	0.0024 (0.0173)	-0.2905*** (0.0228)	0.2061*** (0.0238)
High Evening	-0.0047 (0.0139)	-0.3166*** (0.0243)	0.2157*** (0.0336)
Tech			
Morning	0.0284 (0.0219)	-0.0159 (0.0199)	0.0237 (0.0188)
Evening	0.0084 (0.0153)	-0.0625*** (0.0185)	-0.0034 (0.0177)
High Evening	0.0070 (0.0147)	-0.0677*** (0.0206)	-0.0052 (0.0116)
Manual			
Morning	-0.0155 (0.0165)	-0.0815*** (0.0159)	-0.0433* (0.0162)
Evening	0.0161 (0.0151)	-0.0399*** (0.0149)	0.0078 (0.0145)
High Evening	-0.0137 (0.0125)	-0.0470*** (0.0166)	0.0174 (0.0117)

Notes. The reported results are program-specific treatment effect coefficients in (1), with the inclusion of event-type specific estimates and Pre-Event and Post-Event covariates for each event type. Pre-Event (Post-Event) is an indicator that equals 1 for two hours prior to (after) peak events, and zero otherwise. Standard errors are reported in the parentheses and clustered at the household level. All specifications include fixed effects at the household and hour-of-sample levels. Statistical Significance * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

C.3 No Carryover Assumption

As discussed in Section 4.1, an assumption in our identification strategy is that our randomized events do not impact (or “carryover” to have a treatment effect on) persistent changes in consumption in the event hours on non-event days. We test the validity of this assumption by estimating a DID regression. Separately for each demand response program, we run the following regression that excludes event days in the post-treatment period and includes the never-treated households:

$$\ln(c_{it}) = \beta D_i \cdot \text{EventWindow}_{it} + \alpha_i + \tau_t + \delta X_{it} + \varepsilon_{it} \quad (4)$$

where EventWindow_{it} equals 1 in the post-treatment period for hours where morning or evening events occur and 0 otherwise. D_i equals 1 if the household is in a demand-response group (i.e., C , T , or M) and 0 otherwise. All other features of the regression analysis are identical to those in Equation (1). This analysis evaluates whether households in each demand response program adjusted their consumption during the event windows on non-event days in the post-treatment period, relative to the never-treated households. If households treated with events did not systematically alter their behavior on non-event days in response to being exposed to events, β should be statistically indistinguishable from zero. In addition, we consider a specification that estimates separate effects for the morning and evening event windows on non-event days.

Table C5 presents the results of our no carryover assumption DID test, described in Section 4.1. Table C6 presents the results when we allow for differential effects across the morning and evening event windows on non-event days. In both specifications, we find no evidence of changes to non-event day consumption during the event windows.

Table C5. Carry Over DID Estimates by Program

	Central	Tech	Manual
Event Window	0.0241 (0.0169)	0.0302 (0.0181)	-0.0014 (0.0155)

Notes. The reported results are the program-specific Event Window coefficients from equation (4). For each demand response program, the sample includes households from their own treatment program and the never-treated groups. Standard errors are reported in the parentheses and clustered at the household level. All specifications include fixed effects at the household and hour-of-sample levels. Statistical Significance * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

Table C6. Carry Over DID Estimates by Program and Event Type

	Central	Tech	Manual
Morning Event Window	0.0188 (0.0206)	0.0353 (0.0235)	-0.0226 (0.0201)
Evening Event Window	0.0295 (0.0212)	0.0251 (0.0210)	0.0196 (0.0185)

Notes. The reported results are the program-specific Event Window coefficients from equation (4), allowing for differential effects during the morning and evening event windows. For each demand response program, the sample includes households from their own treatment program and the never-treated groups. Standard errors are reported in the parentheses and clustered at the household level. All specifications include fixed effects at the household and hour-of-sample levels. Statistical Significance * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

C.4 Intention-to-Treat

In our main analysis, we estimate the average treatment effects to peak events by program for households that accepted the demand response program offers. In this section, we estimate an estimator similar to an intention-to-treat (ITT) estimate.

Our treatment (events) are randomized within demand response programs for households that participated in our experiment. Therefore, unlike standard ITT specifications, we cannot have a binary treatment indicator that turns on for all households assigned to a specific demand response program, regardless of whether they accept that program. Because we have periodic events that are randomly assigned to those households that chose to enroll in a demand response program, we create an analogous environment in our setting for all households invited to each program. To do this, we assign households that were randomized to receive the Central, Tech, and Manual program offers but did not accept the offer a distribution of randomized “synthetic” events that is the same as the households that participated. For these households, this creates a new variable \hat{E}_{it} that includes the synthetic events for households that were invited to the Central, Tech, or Manual programs but did not accept our offer, and the observed events for households in our experiment (i.e., E_{it} in Equation (1)).

Using the full sample of households (including those that did and did not accept demand response program offers), we estimate the following equation separately for each demand response program and including the never-treated households that were

not subject to events.³⁵

$$\ln(c_{it}) = \beta_{ITT} \text{Assigned}_i \cdot \widehat{E}_{it} + \alpha_i + \tau_t + \delta X_t + \varepsilon_{it} \quad (5)$$

where Assigned_i is an indicator variable that equals one if the household was assigned to receive a program offer (i.e., the Central, Tech, or Manual offer) and zero otherwise. \widehat{E}_{it} is the household-specific event indicator that equals one if household i experiences an event in hour t (or is assigned a synthetic event, for households that did not accept the program offer) and zero otherwise. Similar to our main analysis, we include fixed effects at the household and hour-of-sample levels; X_t includes the same vector of hourly weather controls. We cluster standard errors at the household level.

We also consider a modified version of equation (5) that permits event type-specific treatment effects. More specifically, \widehat{E}_{it} is now a household-specific set of indicator variables that denote whether the household is was assigned a morning, evening, or high evening event. For households that did not enter into our demand response programs, the (synthetic) allocation of the three event types is randomized with the same frequency distribution as for households that did accept the program offers and receive the event treatments. All other details of the regression specification remain the same.

Table C7. ITT Estimates by Program

	Central	Tech	Manual
Assigned \times Event	-0.1479*** (0.0140)	-0.0377*** (0.0102)	-0.0287*** (0.0089)

Notes. The reported results are program-specific ITT coefficients from equation (5). For each demand response program, the sample includes households from their own treatment program and the never-treated groups. Standard errors are reported in the parentheses and clustered at the household level. All specifications include fixed effects at the household and hour-of-sample levels. Statistical Significance * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

Table C7 provides the ITT estimates. The coefficients imply that offering households the opportunity to join the Central, Tech, and Manual demand response programs results in an average reduction in consumption of -13.7%, -3.7%, and -2.8%

³⁵We run two additional specifications of this analysis where we only include households within each demand response program, and all demand response programs and the never-treated simultaneously. Our ITT results are robust to these alternative specifications.

during events, respectively, from each program of invited households.³⁶ As expected, the results are attenuated downward compared to our main estimates as households that did not join the demand response programs were unaware of the peak events. These percent changes are what an electric utility company can expect “at the end of the day” in terms of event-time consumption reductions after both making offers such as these and conducting event-time alerts. That is, these estimates include both the extensive margin of program offer acceptance and intensive margin of household electricity consumption reductions from those who accepted the offer.

These ITT results, like our main event-level treatment effect results, indicate centralized electricity consumption reductions result in an expected average reduction in demand during events that is several orders of magnitude larger than non-centralized programs, without or without the same technology. The ITT estimates here show uniquely that the relative advantage of centralized demand is still realized, even when one considers the differential take-up rates across programs. In addition, the Tech and Manual programs have ITT estimates that are similar in magnitude, consistent with the findings in our main treatment effect estimates.

Table C8 reports the ITT estimates by event type. The findings parallel the qualitative conclusions from our main average treatment effects on the households that participated in our experiment, with smaller estimated effects due to the ITT including the extensive acceptance margin, as described above. The estimated ITT coefficients represent an average reductions in consumption during events ranging from 12% to 16% for the Central program, 3% to 4% for the Tech program, and 2% to 4% for the Manual program, respectively. We continue to find modest differences across event types for each demand response program. The precision of estimates and significance declines for the Manual program; the high peak event estimate is no longer significant. Further, for each program, we do not observe any evidence of a significantly higher response to increased financial incentives during High Evening events.

C.5 Household Treatment Effects by Event Type

Figure B10 presents the distributions of our household-level treatment effects by program and event type. More specifically, we estimate the specification detailed in (2) but permit different treatment effects by event type: morning, evening, and high

³⁶These calculations reflect the relevant coefficient estimate $\hat{\beta}_{ITT}$, transformed to percentage change in consumption as follows: $100 * (\exp(\hat{\beta}_{ITT}) - 1)$.

Table C8. ITT Estimates by Program and Event Type

	Central	Tech	Manual
Assigned \times Morning Event	-0.1714*** (0.0155)	-0.0294** (0.0127)	-0.0433*** (0.0115)
Assigned \times Evening Event	-0.1331*** (0.0152)	-0.0390*** (0.0125)	-0.0234** (0.0110)
Assigned \times High Evening Event	-0.1454*** (0.0167)	-0.0458*** (0.0136)	-0.0198 (0.0124)

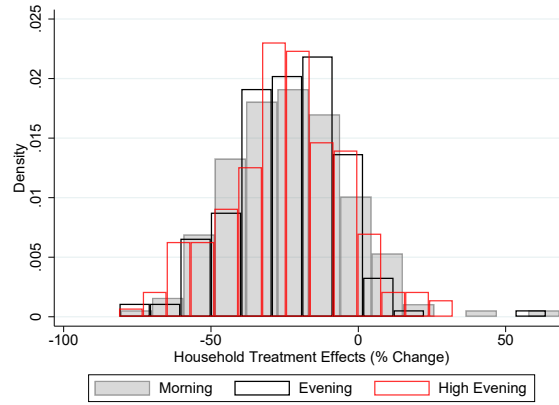
Notes. The reported results are program-specific ITT coefficients from equation (5), allowing for event type-specific estimates. For each demand response program, the sample includes households from their own treatment program and the never-treated. Standard errors are reported in the parentheses and clustered at the household level. All specifications include fixed effects at the household and hour-of-sample levels. Statistical Significance * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

evening.


For the Central program, looking across all households, the average percentage reduction in consumption ranges from 23% during morning events to 24% in evening and high evening events. For Tech, this ranges from 1% during morning events to 5% during evening and high evening events. Finally, for Manual, the average reduction is 3% during morning events and 4% during evening and high evenings. While the precise values vary slightly, the average household-level treatment effects are in a similar range to those found in the program-level regressions summarized in Section 5.2. Importantly, we do not find any evidence of a larger average reduction in consumption during high peak evening events compared to evening events.

A benefit of estimating household-level treatments effects is that we can observe the full distributions for evening and high evening events. Figure B10 shows that the Central program's distributions are quite similar by event type. For the Tech and Manual programs, there is some evidence to suggest that the response to evening events are more tightly distributed around zero compared to the high evening events. However, this does not translate into differences in average treatment effects when looking across all households, as documented above.

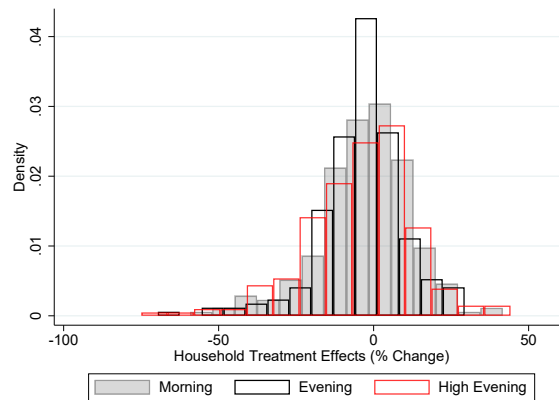
Figure B10. Household-Level Estimated Treatment Effect Distributions by Program and Event-Type



(a) Central Program



(b) Tech Program



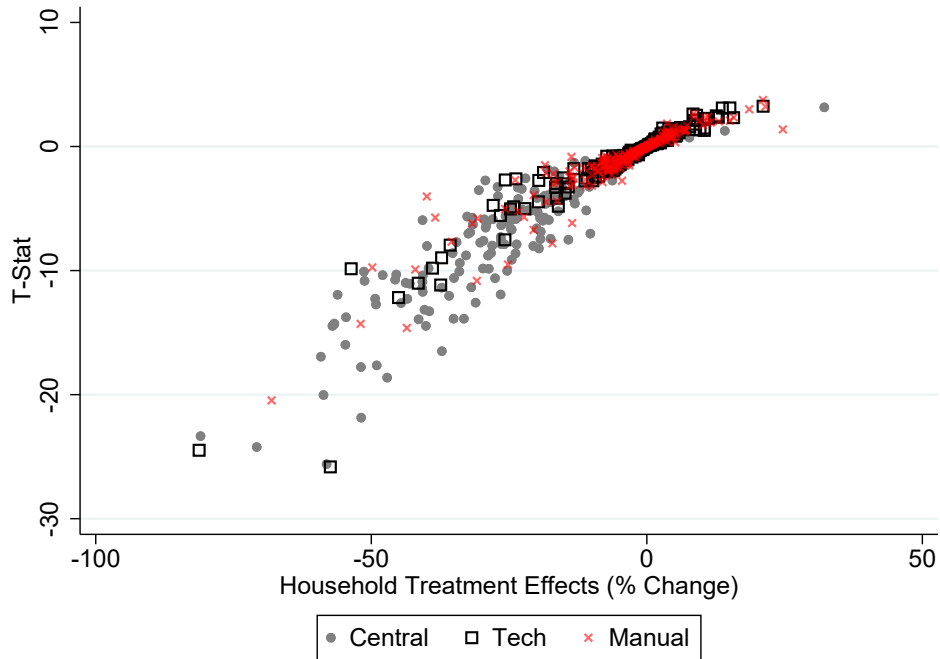
(c) Manual Program

Notes: The reported results summarize the distribution of estimated household-level treatment effects obtained from estimating specification (2), permitting different treatment effects by event type: morning, evening, and high evening. We present the marginal effects to be a percentage change in consumption using the transformation $100 \times (\exp(\hat{\beta}_i) - 1)$.

C.6 Consistency of Demand Response

In this section, we expand upon the household-level results to investigate the consistency of responses to events by program. This builds off of the discussion in Section 6.1 and Figure 6 in particular.

Figure B11. Household-Level Estimated Treatment Effects and T-Statistics by Program



Notes. The reported results plot the estimated household-level treatment effects and corresponding t-statistic obtained from estimating specification (2). We present the marginal effects to be a percentage change in consumption using the transformation $100 \times (\exp(\hat{\beta}_i) - 1)$.

Figure B11 plots the estimated household event responses and the corresponding t-statistics to better understand the significance of the estimates. The majority of the Manual and Tech values are closely distributed near zero in both the estimated event response and t-statistics. Only 20% and 18% of the household-level estimated treatment effects are negative and significant at the 5% level for the Tech and Manual programs, respectively.³⁷ This suggests that only a small subset of high-performing households in these programs were consistent responders to peak events over our 17-

³⁷There is a small subset of households with positive estimated treatment effects. These estimated effects are systematically statistically insignificant, with only 3% being positive and statistically significant.

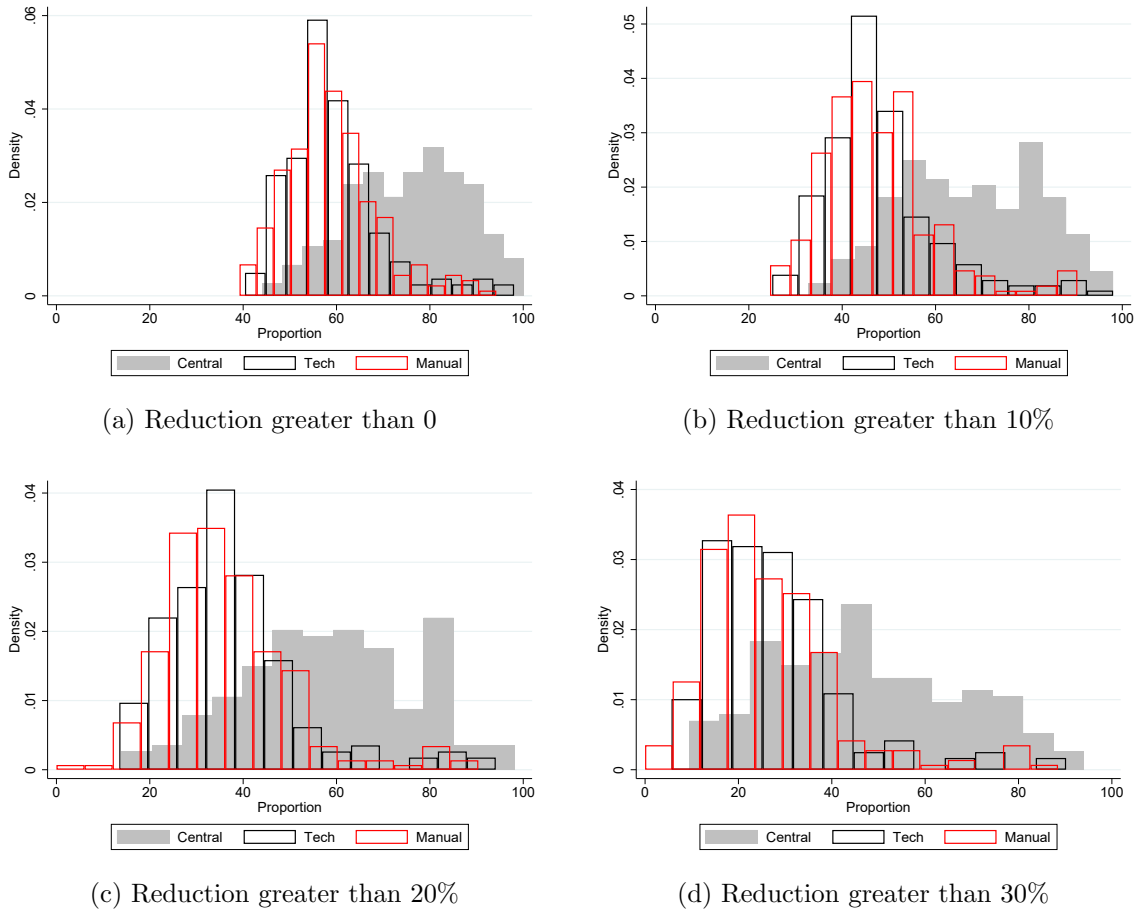
month treatment period. Figure B11 reinforces our finding that the Central program has large and precisely estimated responses to peak events. 80% of the Central household-level estimated treatment effects are negative and significant, indicating that the majority of Central households provide large and consistent reductions in consumption during peak events.

In addition to having estimates for household-level treatment effects, we observe each household's consumption relative to their baseline for *each* event. As described in Section 2.2, the baselines were calculated based on a household's average consumption during the relevant event time window over the last five weekdays prior to the event, excluding event days. For each household, we calculate the proportion of events where their observed percentage reduction in consumption relative to their baseline was greater than 0%, 10%, 20%, and 30%. It is important to recognize that this measure will have some noise. This is particularly relevant for the comparison of consumption relative to the baseline being greater than 0%. Even if a household did not adjust its behavior to respond to the event, natural variation would induce consumption to fall below the baseline for a certain proportion of events. If this randomness is normally distributed around the baseline, the proportion would be 50%. While this makes interpreting the results for this measure more nuanced, as we will describe below, it is still valuable because we can compare the relative performance across programs.

Figure B12 presents, at the household-level, the proportion of events where the observed percentage reduction in consumption relative to the baseline was greater than 0%, 10%, 20%, and 30% by program. Focusing initially on the results for greater than 0% threshold, the Tech and Manual groups have similar distributions and are slightly rightward shifted above 50%. As noted above, noise in this measure could result in values near 50% being viewed as having a minimal reduction relative to the baseline for the greater than 0% threshold. The fact that the Tech and Manual programs have distributions situated just to the right of 50%, with a small subset of households exceeding 70%, is consistent with our household-level regression results. In contrast, the majority of Central households have proportions in excess of 70% of event days being greater than 0%. This continues to suggest that Central households consistently reduce their consumption during events.

Figure B12 shows that as the threshold increases to greater than 10%, 20%, and 30% reductions relative to the baseline, the proportion of events where households achieve reductions above this threshold declines across all programs. However, the proportion of events where Tech and Manual achieve reductions above this threshold

Figure B12. Household-Level Proportion of Responses with Reductions Relative to Baseline Greater Than 0%, 10%, 20%, and 30%



Notes: Figure plots at the household-level the proportion of events where their observed percentage reduction in consumption relative to their baseline was greater than 0%, 10%, 20%, and 30% by program over the period February 1, 2022 - June 30, 2023.

declines at a much faster rate than Central. For example, the majority of Central households achieve reductions greater than 20% in at least 50% of events. In contrast, this number is in the range of 20% - 30% of events for the majority of Tech and Manual households, with a small subset of higher performers who consistently reduce their consumption above this threshold in the majority of events.

Taken together, these results suggest that the Central program was able to achieve a high and consistent demand reductions for the majority of events over the 17-months they were exposed to peak events. This differs from the Tech and Manual programs that achieved more modest demand reductions for a smaller proportion of events.

D End of Experiment Survey

D.1 Survey Details

The following is text from a voluntary, end-of-experiment survey sent out to participants in the Central, Tech, and Manual programs via email, in mid-June, 2023.

Survey instructions:

“This short survey is designed to hear about your experience in the Peak Rewards Trial through [APP NAME]. All homes had a different experience, and we want to hear about yours.

We appreciate the time and thought you put into this survey.

Properly completed surveys will be rewarded with \$20 on bill credit as a token of our gratitude.”

Survey questions used in our analysis:

“What is your approximate household income?”

- Less than \$50k per year
- \$50-99k per year
- \$100-149k per year
- \$150-200k per year
- Over \$200k per year
- Don’t know/Rather not say”

“For the events you noticed, how often was it worth your time to participate by attempting to reduce your electricity consumption?”

- Never
- Sometimes
- About half the time
- Most of the time
- Always
- Don’t know/Not Applicable”

D.2 Survey Response

In this section, we compare the observable characteristics of participants who filled out the end-of-experiment survey to those who did not. The Table below recreates our balance Table A2 (using pre-treatment data), but separates the results by whether or not the household responded to the exit survey. The p-value corresponds to a difference in means test.

Table C9 demonstrates that the non-respondents had larger cumulative consumption during the pre-treatment period. Non-respondents also were more likely to have an electric vehicle. All other characteristics are similar across the two groups.

Table C9. Balance by Exit Survey Response (Pre-Treatment Data)

	Yes	No	p-value
Cumul. kWh			
Winter	5,229 (2,810)	5,892 (3,199)	0.04**
Spring	3,675 (1,788)	4,167 (2,104)	0.02**
Summer	2,614 (1,492)	3,155 (2,142)	0.01***
Fall	3,528 (1,721)	4,012 (2,016)	0.02**
Load Factor			
Winter	24.96 (8.61)	25.21 (8.89)	0.78
Spring	19.82 (6.64)	19.98 (6.93)	0.83
Summer	16.56 (6.87)	17.21 (6.76)	0.37
Fall	18.66 (5.91)	19.12 (6.72)	0.51
Electric Vehicle	0.24 (0.42)	0.32 (0.47)	0.05**
BaseBoard Heating	0.66 (0.47)	0.64 (0.48)	0.71
Air Conditioning	0.47 (0.50)	0.51 (0.50)	0.32
Electric Hot Water	0.72 (0.45)	0.70 (0.46)	0.56
House Duplex	0.83 (0.37)	0.82 (0.38)	0.71
Median Income	86,377 (19,853)	87,434 (19,835)	0.55
Observations	429	174	

Notes. This table compares pre-treatment average values by whether or not the household participated in the exit survey. Parentheses contain the standard deviations. Cumul. kWh and Load Factor represents the cumulative household-level consumption and load factor by season. Electric Vehicle, Baseboard Heating, Air Conditioning, and Electric Hot Water are indicator variables denoting the presence of each device. House/Duplex is a indicator variable if the home type is a single-family home or duplex. Median Income reports the median household-level income of the Census Dissemination Area where the household is located. p-value from a difference in means test across the two groups. Statistical Significance * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.