Do Smart Technologies Deliver? Smart Thermostats and Energy Conservation

Alec Brandon
University of Chicago

Christopher M. Clapp*
University of Chicago

John A. List
University of Chicago

Robert Metcalfe
Boston University

Michael Price
University of Alabama

February 25, 2020†

[Click for Most Recent Version]

Abstract

We estimate the causal impact smart thermostats have on energy use using data from a field experiment in which treated households were randomized into free installation of a smart thermostat. We combine this experimental data with 18 months of high-frequency data on household energy consumption in the form of more than 16 million hourly electricity use records and almost 700 thousand daily observations of natural gas consumption. We model the effect of a smart thermostat on energy consumption using a difference-in-differences instrumental variables (DDIV) specification. In contrast to advertised savings based on engineering models, we find no evidence that smart thermostats have a statistically or economically significant effect on energy use. This result is robust to the inclusion of numerous controls and when the model is estimated on various subsamples (e.g., by hour). We explore potential mechanisms using almost four million observations of system events including user interactions with their smart thermostat. Results indicate that user behavior dampens energy savings and explains the discrepancy between estimates from engineering models and those from the field.

*Corresponding author. Kenneth C. Griffin Department of Economics at the University of Chicago. Address: 1126 East 59th Street; Chicago, IL 60637; USA. Email: cclapp@uchicago.edu. Phone: (773) 795-1480. Web: www.chrisclapp.org.

†Acknowledgements: Koichiro Ito, David Jimenez-Gomez, David Novgorodsky, and Florian Rundhammer provided helpful comments and suggestions. Ariel Listo and Uditi Karna provided excellent research assistance. We would like to thank Richard Caperton, Augustin Fonts, Ilan Frank, Kevin Hamilton, Nicholas Payton, Jim Kapsis, and many others at Opower for sharing data and offering insights. Opower provided the data analyzed in this paper to the authors under a nondisclosure agreement. The authors and Opower structured the agreement in a way that maintains the authors' independence. In particular, the agreement stipulates that Opower has the right to review the publication prior to public release solely for factual accuracy.

JEL Codes: D01, Q41, Q48, Q55.
Keywords: Energy Demand, Smart Grid, Smart Technology, Field Experiment.

1 Introduction

According to the Energy Information Administration (EIA), American households consume an average of 166.3 million British thermal units (BTU) of energy per year (EIA, 2019a), and the average single-family home owner spends just over $2,200 on energy annually (EIA, 2018). In total, the production of that residential energy results in the emission of over one billion metric tons of carbon dioxide into the atmosphere each year (EIA, 2019b).[1] These high private and social costs have led to substantial interest in smart technologies that reduce energy use without reducing consumer utility by increasing efficiency. Given that the largest share (almost 40%) of residential energy use goes to heating and cooling the home (EIA, 2019a), smart thermostats are an increasingly popular example of such a technology.[2]

Smart thermostats allow individuals to program temperature setpoint schedules and adjust settings remotely via a smart phone application. While producers of these devices promise consumers substantial savings on their home heating and cooling bills, projected savings are often based on engineering studies that fail to account for how people actually use their smart thermostat or rely upon non-experimental data that fails to estimate the causal relationship between smart technologies and household energy use (Peffer et al., 2011; Peffer et al., 2013).[3] Thus, the true impact of smart thermostats on energy usage "in the field" is uncertain.

In order to determine the causal impact that smart thermostats have on home energy usage, we examine data from a field experiment conducted by Opower and Honeywell in conjunction with Pacific Gas and Electric (PG&E) – the second largest utility in California. As part of the experiment, the 1,379 households that volunteered to participate in the study were randomized into either a treatment group who received free installation of a Honeywell two-way programmable thermostat linked to an Opower platform or a control group that received neither.[4] We combine this house-

[1]This is roughly 20% of the total annual carbon dioxide pollution due to the production of energy.

[2]Mordor Intelligence estimates that the smart thermostat market was worth $689.8 million in 2018 (Mordor Intelligence, 2019).

[3]For instance, the Nest smart thermostat website advertises a 10 to 12% savings on heating and a 15% savings on cooling costs. The ecobee smart thermostat touts a savings of "up to 23%" on heating and cooling costs. Honeywell has a savings calculator that reports annual savings by location based on estimates of local energy costs. Their website cautions, "The savings in [these] calculations are based on schedules you may set up by programming your thermostat.... The demonstration in this calculation uses typical events people may use to schedule/program their thermostats. Your results may vary depending on your dynamic lifestyle."

[4]In addition to the ability to schedule permanent temperature setpoints and interact with the thermostat remotely, the smart thermostat given to households in our experiment provided households with a social norm framing of their setpoint choices. Although not in the context of smart technologies, extensive work has shown the responsiveness of household energy consumption to social norm framing (e.g., Allcott, 2011; Ferraro and Price, 2013; Ayres et al., 2012; Costa and Kahn, 2013; Allcott and Rogers, 2014; Dolan and Metcalfe, 2015). Given this finding and the Peffer et al. (2013) result that individuals do not use their smart thermostats as intended, this feature should provide the best chance for a smart thermostat to reduce energy consumption. Additionally, some form of framing is an increasingly common feature of more modern smart thermostats.

hold thermostat technology data with high-frequency data on household energy consumption over an 18-month period in the form of more than 16 million hourly electricity use records and almost 700 thousand daily observations of natural gas consumption. We model the effect of thermostat technology on household energy consumption using a difference-in-differences instrumental variables (DDIV) specification where we use assignment to treatment as the instrument for installation of the thermostat. The coefficient of interest measures the differential change in energy use across pre- and post-intervention periods for treated versus control households. Thus model allows us to identify the causal impact of a smart thermostat on energy use.[5]

Across numerous specifications, we find that smart thermostats have neither a statistically nor economically significant effect on energy use. This result is robust to the inclusion of controls for weather conditions and a battery of household, location, and time effects. To investigate whether this aggregate result masks significant, but offsetting, effects, we estimate the model on subsamples by day of the week, hour of the day, and ambient temperature/humidity quintiles. We find no evidence of heterogeneous treatment effects.

In order to explore potential mechanisms that would explain this null result, we rely on almost four million observations of heating, ventilation, and air conditioning (HVAC) system activity and user interactions with their smart thermostat in the form of scheduled temperature setpoints, temporary overrides, and HVAC system events. We first provide descriptive evidence that users do take advantage of the smart features of their devices by showing that they frequently schedule permanent setpoints, the pattern of those setpoints across hours of the day is intuitive, and the temperatures that they set are in line with Environmental Protection Administration (EPA) energy-efficiency guidelines. We then establish that users frequently override scheduled temperature setpoints, and when they do, override settings are less efficient than their previously scheduled counterparts. To more formally test the hypothesis that user behavior explains the discrepancy between the decrease in energy use purported by the engineering studies and our experimental estimates, we categorize smart thermostat households into flexible, energy-efficiency type categories based on their relative permanent setpoint and temporary override behavior. We match these categories to our experimental, energy use data, interact energy-efficiency type with an indicator for treatment, and estimate difference-in-differences intention-to-treat (DDITT) models.[6] Estimates indicate that some high-efficiency type users do realize significant savings by installing a smart thermostat, but that human behavior explains the discrepancy between engineering estimates and our null experimental results. Our findings suggests that engineering models fail to adequately incorporate how people actually use smart technologies, thus severely limiting the usefulness of

[5]Our treatment is likely to result in heterogeneous treatment effects, as evidenced by the fact that not all subjects who were randomized into the treatment group ultimately installed a smart thermostat in their home. Under the assumptions of instrument exogeneity and one-sided noncompliance (although some households in treatment group do not take-up the treatment, no households in the control group install a smart thermostat), our estimates can be interpreted as the average treatment effect on the treated (ATT) of a smart thermostat (Cornelissen et al., 2016). We more formally discuss identification and provide evidence based on an external dataset that one-sided noncompliance is a reasonable assumption in Section 3.2.

[6]We estimate DDITT models instead of DDIV models because our experiment did not (and could not) stratify households by ex-post energy-efficiency type, so we do not have valid instruments for type.

their estimates.

We make several contributions to the literature. First, while there has been considerable research on smart grid investments (Joskow, 2012), much less work has been done exploring the impact of smart technologies on residential energy use. Initial assessments of these technologies have focused on changes in average energy use induced by in-home displays of real-time energy price or quantity information (see, e.g., Faruqui and Sergici, 2010; Jessoe and Rapson, 2014; Alberini et al., 2013). [7]

Second, our results have important policy implications as there are both government and industry funded subsidies of smart technologies. Between 2009 and 2014, the Department of Energy (DOE) invested $7.9 billion in smart technologies under the Smart Grid Investment Grant (SGIG) program (DOE, 2016).[8] Additionally, the joint EPA and DOE ENERGY STAR program certifies the efficacy of smart thermostats. The program partners with 17 utility companies to sponsor rebates for purchases of smart thermostats. In 20 states, over half of all households are eligible for a smart thermostat rebate, and in the most generous case, all of the residents in Nevada are eligible to receive a smart thermostat for free (Bloomberg New Energy Finance, 2019). Given the current information available, energy producers and policymakers alike are subsidizing these devices based on misleading information with funds that would be better spent on more effective policy interventions.

The remainder of this study is organized as follows: In Section 2 we describe the details of the randomized control trial (RCT), the sample of households in the study, and our data. The following section formalizes our empirical specification. Section 4 presents our model estimates, and Section 5 explores the mechanisms that drive our findings. The final section concludes.

2 Experimental Design

2.1 Smart Thermostat

The intervention in our field experiment occurs when a given household's existing (dumb) thermostat is replaced by a smart device.[9] Smart thermostats are designed to increase consumer utility by improving the efficiency of the home's HVAC system and reducing adjustment costs. To these ends, the device in our experiment has two key features common to most smart thermostats. First, the thermostat allows the user to program an extensive schedule of permanent temperature setpoints for each day of the week. Second, the user can either interact with the device directly or remotely

[7] Harding and Lamarche (2016) is a notable exception. The authors consider the effect of technologies that automate temperature setpoint changes to dynamic pricing.

[8] While more than two-thirds of these investments went towards outfitting households with smart meters and communication systems that allow utilities to integrate real-time market conditions into household consumption decisions via dynamic pricing plans or demand response messaging, a complementary set of investments targeted the development and dissemination of technologies such as smart thermostats that allow individuals to remotely communicate with their appliance and HVAC system.

[9] Specifically, Opower/Honeywell installed a Honeywell Z-Wave Touchscreen Thermostat that communicates with a website portal and smartphone app designed and hosted by Opower.

via a web portal or smartphone app. Both lower the cost of adjusting temperature settings.[10]

While the effect of these features on energy usage is theoretically ambiguous depending the schedule the user sets and how she interacts with the device, there are several additional features of the thermostat used in our experiment that encourage users to make changes that reduce energy consumption. First, when choosing setpoints, users receive messages that compare their settings to those of similar households. Analogous to the social comparison module studied in Allcott (2011), the thermostat interface presents: (i) descriptive norms with information on peer setpoint choices and (ii) injunctive norms with efficiency ratings of setpoints. Second, the thermostat app interface is designed to facilitate toggling to a less energy intensive setting when the user leaves home and toggling it back to the previous setting when the user returns. Finally, when a user overrides a permanent setpoint to make a temporary change that is more energy efficient than the scheduled one, she is prompted by a query asking if she wants to make this more energy efficient setting permanent.[11]

2.2 Experiment

Figure 1 illustrates the execution of the field experiment. It describes the assignment of households to treatment and control groups, as well as the subsequent installation decisions of treatment households. Potential subjects were recruited in public places (e.g., malls, markets, and festivals).[12] A total of 1,379 eligible households agreed to participate in the study and were randomized into either a treatment or control group.[13] After group assignment, the experimenter had no further contact with the 690 control households. The 689 households assigned to the treatment group were offered the smart thermostat described in the previous section and installation at no cost. The smart thermostat was successfully installed in 73% of homes in the treatment group. Of the remaining treatment homes, 19% percent declined, and 8% had complications that prevented installation (e.g., compatibility issues).

[10]Appendix Section A provides a more detailed description of the device. Panel (a) of Appendix Figure 16 displays the thermostat and associated applications. Panel (b) shows a screen-shot of scheduling using the smartphone app.

[11]Appendix Figure 17 highlights features of the smart thermostat. Panel (a) illustrates the social norm framing displayed when households choose setpoints. Panel (b) shows how households can remotely toggle the thermostat in response to leaving and returning home via a smartphone or personal computer.

[12]To be eligible, an individual had to own her residence and have central air conditioning, a smart phone, and high-speed Internet. See Appendix Section D.1 for a summary of the eligibility requirements. For more information on canvassing, see Appendix Section D.2 for the original recruitment and enrollment guide.

[13]All household counts in this section are based on the households for which we observe electricity consumption. There are 1,379 unique households in the electricity sample, 1,369 unique households in the natural gas sample, and 1,385 unique households across both samples. Stated another way, we observe 16 households with electricity consumption data, but not natural gas information and another six households that consume natural gas, but for which we have no electricity consumption information.

Figure 1: Sample Randomization

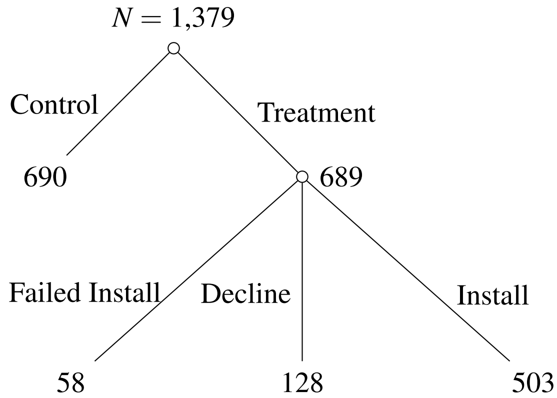

Figure 2 plots the cumulative density function (CDF) of the difference in time between assignment and installation dates that illustrates how long it takes households in the treated group to install the smart thermostat (conditional on eventual installation of the smart thermostat). Most households had the smart thermostat installed shortly after being assigned to the treatment group: 50% of households had their thermostat installed within 5 days, and 95% had it installed within 30 days.

Subjects were recruited in two waves. The first wave of recruitment took place across four counties in Northern California from July through October of 2012. The second wave of subjects were recruited from December of 2012 to February of 2013 in three Central California counties.[14] Figure 3 depicts the locations of homes in the experiment and provides visual evidence that treatment and control groups are spatially balanced across locations.[15]

2.3 Energy Data

All households in the study were equipped with smart meters that enabled PG&E to record household-level data on hourly electricity use and daily natural gas consumption. The quantity of electricity consumed is measured in kilowatt hours (kWh), and the unit of measurement for natural gas is a therm (thm).[16] As we cannot observe temperature setpoints directly for control households with a traditional thermostat, and energy is the policy-relevant good, these measures are the main outcome variables in our analyses. In total, we observe an average of 11,908 hourly electricity use decisions for the 1,379 households in electricity sample and 495 natural gas use decisions for the 1,369 households in the natural gas sample over an 18 month period from July 2012 through December 2013.

[14] Northern California wave subjects were recruited from the greater San Francisco/Sacramento area (Contra Costa, San Joaquin, Solano, and Yolo counties). The Central California wave households are located in and around Fresno and Bakersfield (Fresno, Kern, and Madera counties).

[15] We formally test balance in Section 2.6 and fail to reject the null of spatial balance in the counties where households are located.

[16] A therm is a unit of heat energy equivalent to 100,000 BTUs.

Figure 2: Conditional Distribution of Time from Assignment to Installation

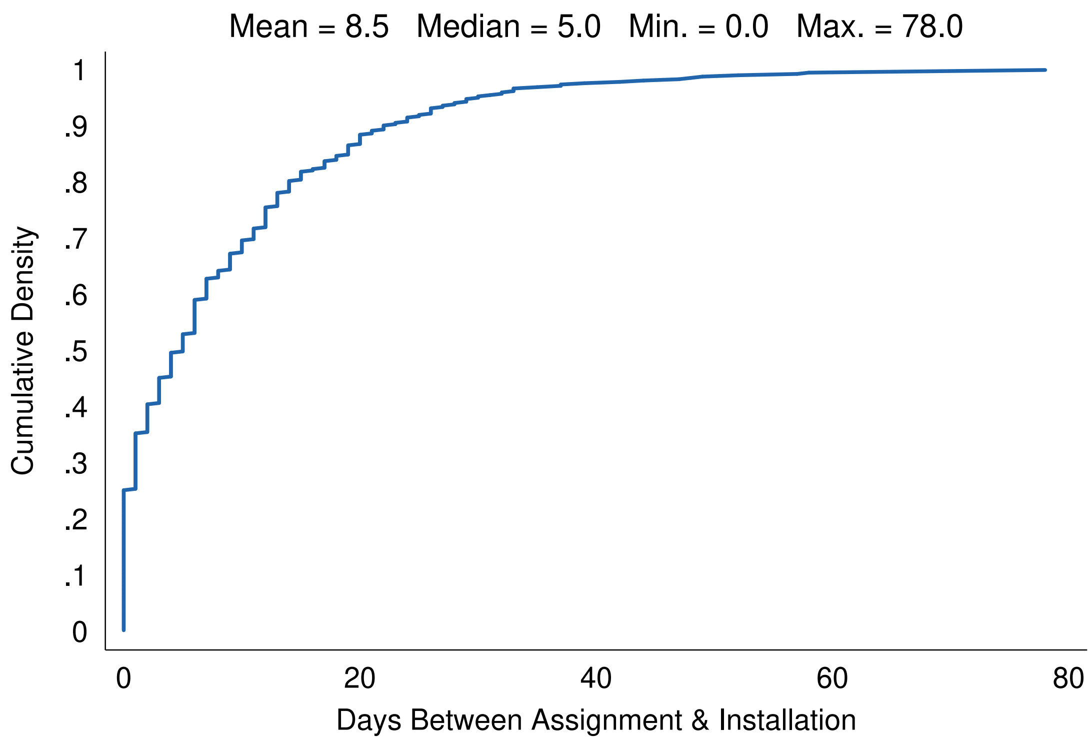

Cumulative density conditional on eventual installation.

2.4 Timing

Figure 4 presents three visual depictions of important timing issues associated with the experiment and data. Panel (a) is a timeline that illustrates the temporal relationship between the two waves of subject recruitment and the the period over which we observe energy data. The red line indicates the range of time during which individuals were recruited in Northern California, and the blue line depicts the Central California wave of recruitment. The black line is the span of time over which we observe energy data.

Figure 3: Locations of Treatment and Control Groups

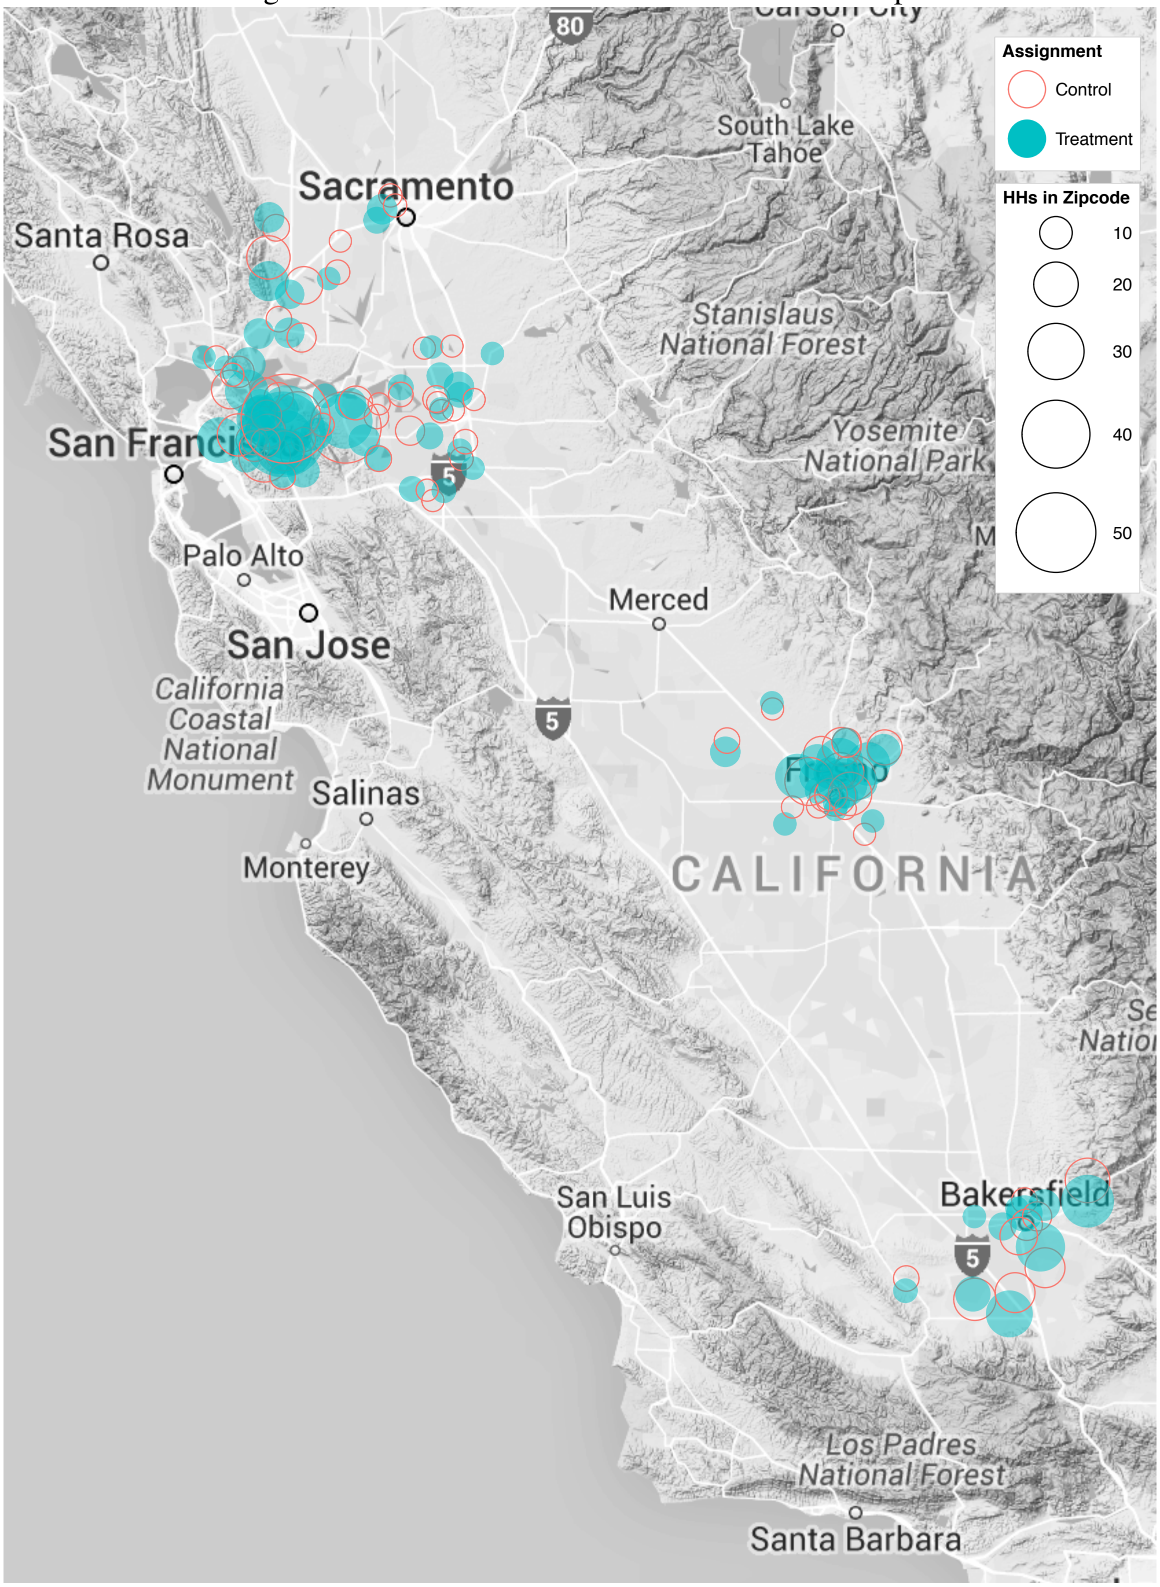

Figure 4: Timeline

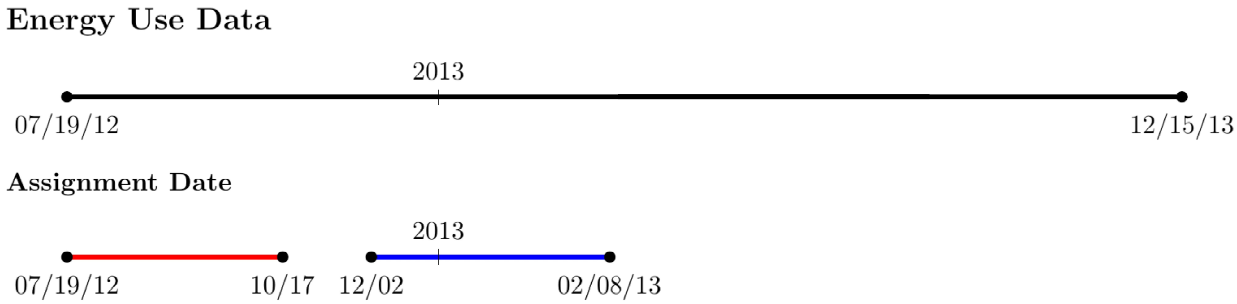

(a) Timeline

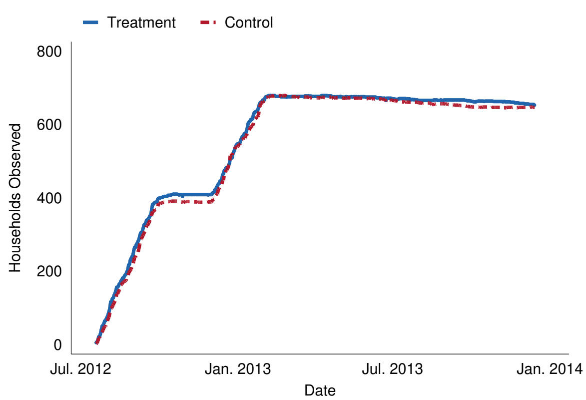

(b) Number of Households Observed by Experimental Status and Date

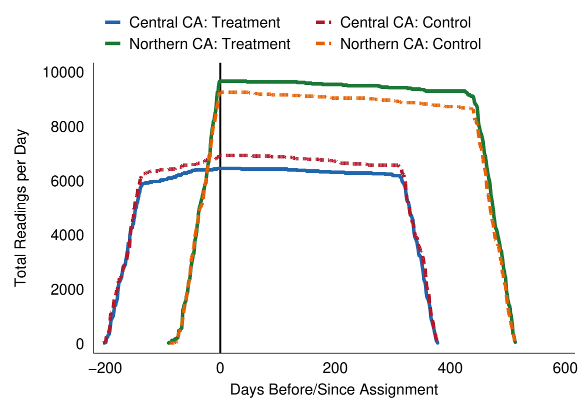

(c) Total Electricity Readings by Experimental Status, Wave, and Event Time

Panel (b) of Figure 4 plots the flow of households into the treatment and control groups over time. Importantly, the subfigure shows that treatment and control households are temporally balanced, as they were assigned at similar rates over time.[17] Unfortunately, due to an error in the energy data collection process, we only observe energy readings starting on the first day of recruitment in Northern California. Panel (c) illustrates the effect of this issue by plotting the number of electricity readings per day for each wave relative to event time (where assignment to the treatment or control group occurs at time zero).[18] The figure shows that we do not observe a substantial pre-period for all households in the Northern California recruitment wave, but we do for the Central California wave. While the model estimates reported in Section 4 are based on the full sample, we also report results separately by wave in Appendix Section C.1 to account for this issue. The subsample estimates are not qualitatively different to those based on the full sample.

[17]We formally test balance in Section 2.6 and fail to reject the null of temporal balance in the month of assignment to experimental group.

[18]Plotting an analogous graph for natural gas readings results in the same the same pattern.

2.5 Additional Data

2.5.1 External Data

We supplement the main experimental dataset with information from several external sources and additional data collected as part of the experiment. First, we compile hourly temperature, humidity, and heat index readings for each county in the study from the National Oceanic and Atmospheric Administration (NOAA).[19] Figure 5 summarizes the outdoor temperatures households face across the different counties in our sample by plotting time series of the minimum and maximum daily ambient temperatures over the sample period. The figure highlights two facts that inform our model specification. First, despite our sample being drawn from a temperate part of the country, there is substantial seasonal (within-county) variation in the NOAA data. Temperatures in the sample range from below freezing to well over 100 degrees Fahrenheit. Summers are hot and require the use of air conditioning to ensure comfortable indoor temperatures. While the rest of the year is more moderate, there are many days cold enough to necessitate home heating. Thus, we estimate separate models of the effect of a smart thermostat on two energy sources: electricity (the energy source used for cooling) and natural gas (the predominate energy source for heating). Additionally, as is born out in Table 1, there is both between-county and daily variation in temperatures. We include hourly outdoor temperature and humidity measures, as well as location and time effects, as controls to address concerns that ambient weather fluctuations affect our results.

[19] We are missing values for 0.09% of the temperature and 0.5% of the humidity observations in the sample. We interpolate these missing values using the predicted values from separate regressions of the given weather variable on location, day, and hour fixed effects. We calculate the heat index from the temperature and humidity readings (see: https://www.wpc.ncep.noaa.gov/html/heatindex_equation.shtml for the formula).

Figure 5: Minimum and Maximum Daily Outdoor Temperatures $(^{\circ}F)$ by County

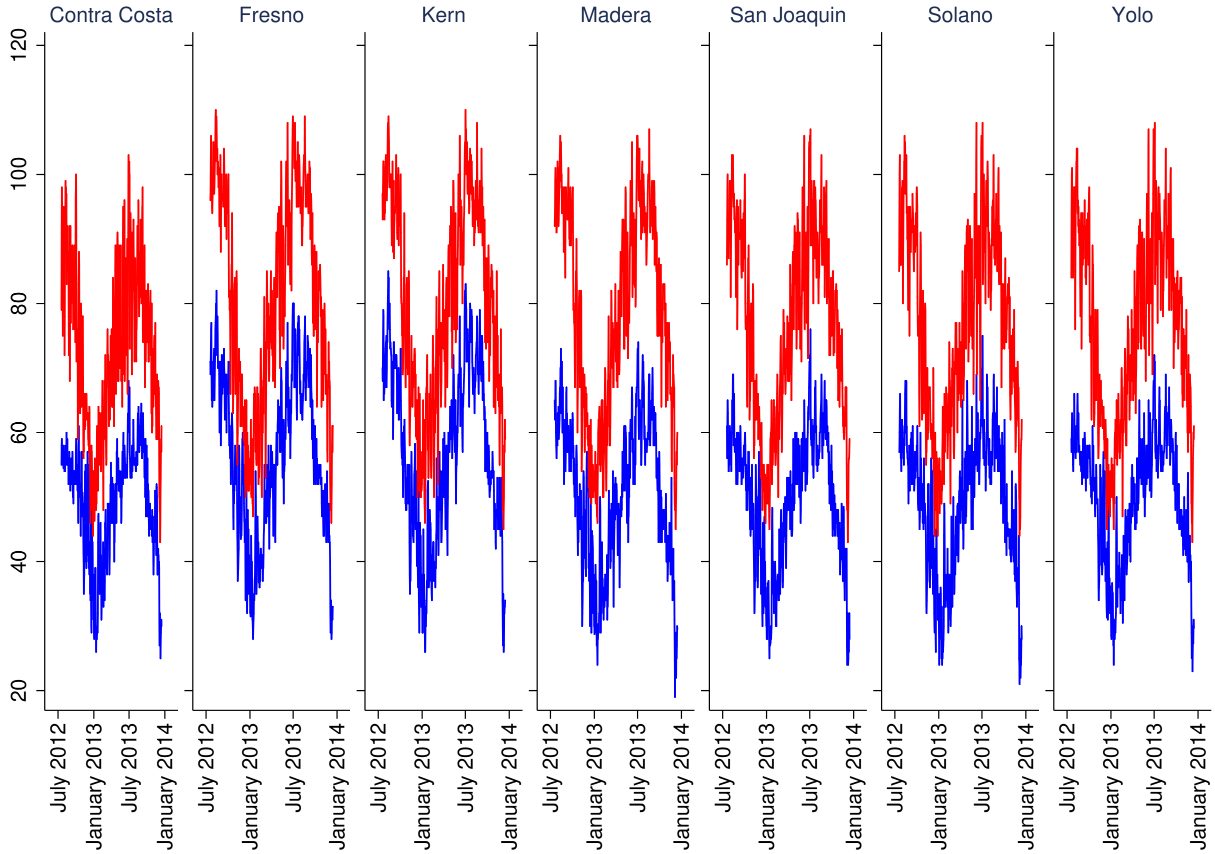

Table 1: Daily Outdoor Temperature $(^{\circ}F)$ Summary Statistics

Variable	Mean	Std. Dev.	Between Std. Dev.	Within Std. Dev.	Min.	Max.
Mean Daily Temp.	63.70	13.06	3.20	12.71	32.63	96.04
Minimum Daily Temp.	51.34	11.55	3.43	11.10	19.00	85.00
Maximum Daily Temp.	77.52	15.23	2.58	15.05	43.00	110.00
N			7			
$N \times T$			3,605			

Second, since electricity is produced from many sources with different production and external costs, we supplement the household electricity use measure with data on the average hourly real-time price of electricity from the California Independent System Operator (CAISO).[20] By combining these data sources, we are able to construct a household-specific measure of the social

[20]The real-time market for electricity in California clears every five minutes. We use this data to calculate the average price each hour. A similar measure is not available for natural gas prices.

marginal cost of producing the electricity the household uses each hour.[21] We use this measure to test whether smart thermostats have a differential effect on usage during peak load times when cost of electricity production to society is the greatest. Figure 6 plots the mean spot price by decile to illustrate the variation in production and social costs that exists in our data.

Figure 6: Average CAISO Spot Price by Decile of Price

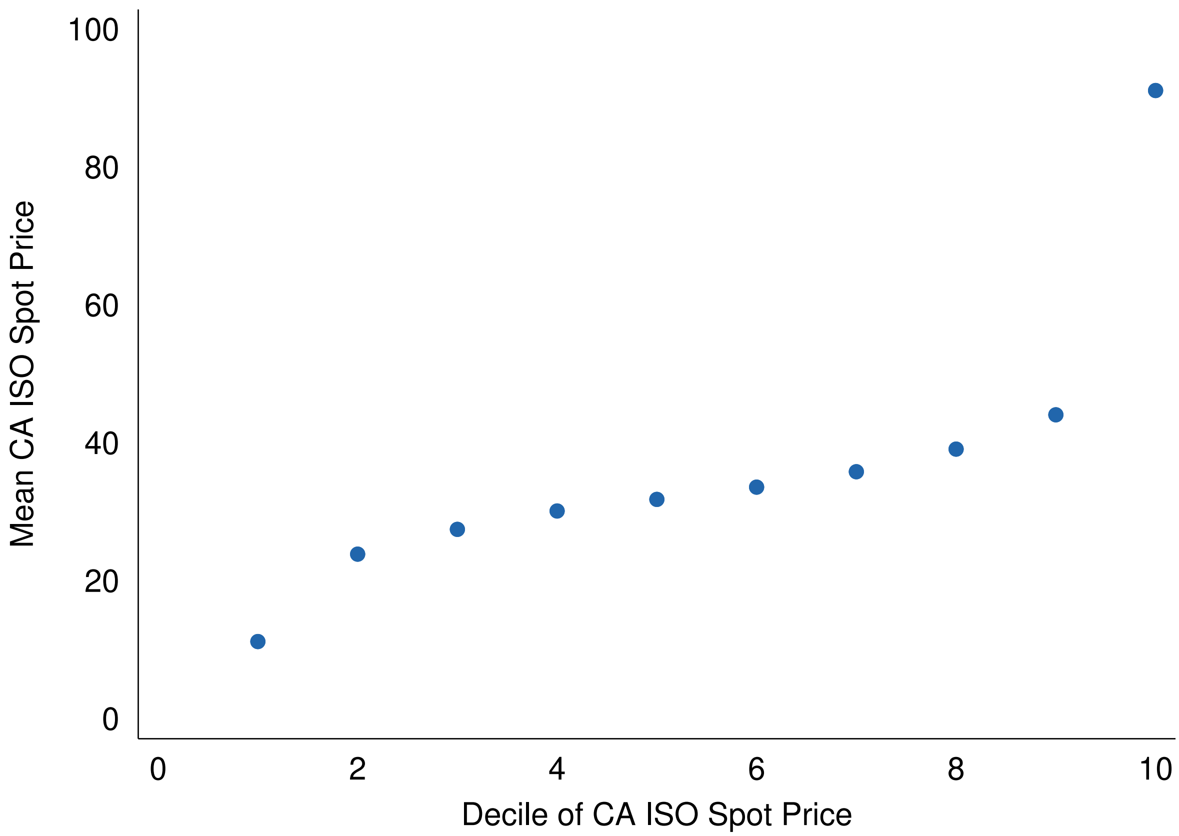

2.5.2 Internal Data

In addition to the external data we collect, we also observe a high-frequency, exact-time log of 3,967,558 HVAC system events, including user interactions with their smart thermostat, from 365 households. The unbalanced panel dataset spans from July 2012 to January 2013, and Figure 7 illustrates the number of households observed by calendar date. Recruitment and installation of

[21] California instituted a cap-and-trade carbon emissions program in the summer of 2012, so the price of electricity on the state's wholesale market reflects both the marginal cost of production and the prevailing market price for emissions as reflected in the price of carbon permits. The extent to which electricity prices reflect true social costs depends on whether regulators in California issued the socially optimal number of permits. If allowances exceed this level, then the prevailing permit price is less than the external cost and electricity prices are a lower bound on social marginal costs.

smart thermostats first began in Northern California in July of 2012, whereas those in Central California began in December of 2012. Since this dataset is truncated in January of 2013, the majority of the cross-sectional units in this dataset are homes from Northern California, while roughly 5% of the households in the data are from Central California.

Figure 7: Number of Households Observed in Events Data by Date

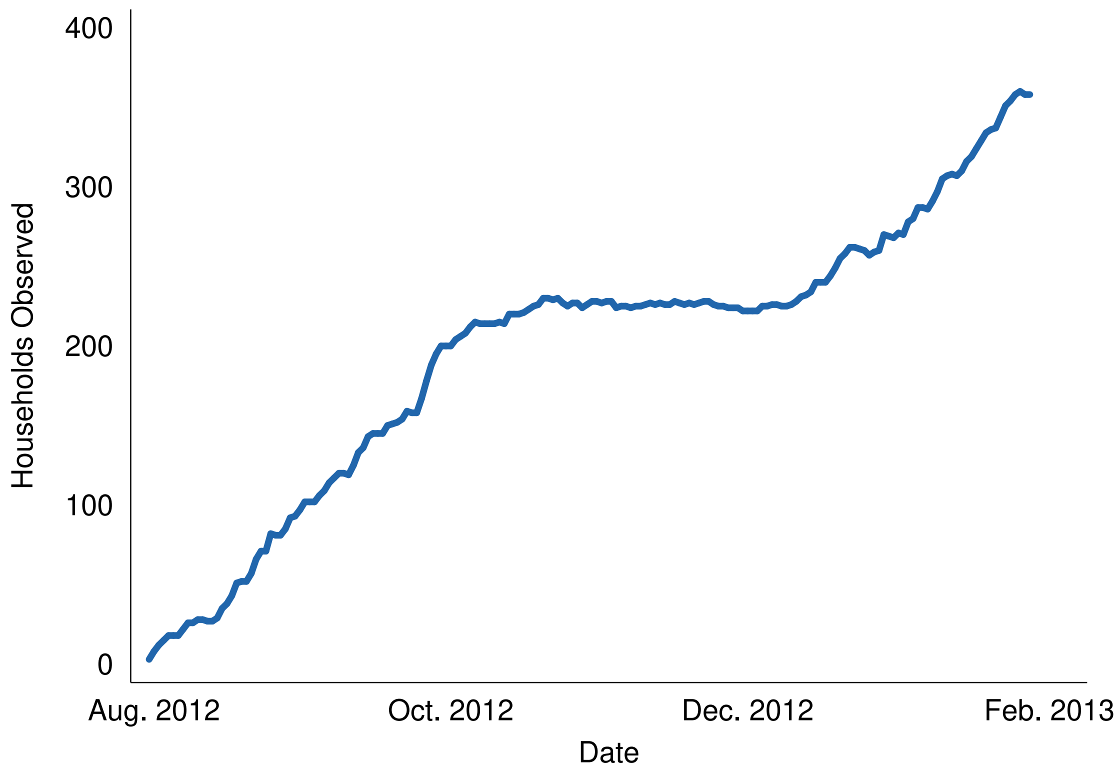

The system events and user interactions we observe include ambient temperature, HVAC state, and heating/cooling setpoints (which we classify into permanent setpoints and temporary overrides).[22] We aggregate these measures to hour-level observations. Table 2 summarizes the data. The table shows that while there are more observations from the Northern California wave of the

[22]Permanent setpoints are thermostat temperature settings previously scheduled to occur automatically at specific times on a periodic basis. Temporary overrides are changes to the current setpoint which result from a concurrent interaction with the thermostat. We do not actually observe whether system temperature changes are permanent setpoints or temporary overrides. Appendix Figure 18 informs our approach to classifying setpoints and overrides. Panel (a) plots the density of the second of the minute at which temperature changes take place. The density is roughly uniform with a probability of about 0.70 across all seconds, save for a large increase in the probability of changes occurring at :00 through :02 (and to a lesser extent :03) seconds of the minute. Since we would expect temporary overrides to occur uniformly across seconds of the minute, we code temperature changes occurring at less than :03 seconds of the minute as permanent setpoints and all other temperature changes as temporary overrides. Panel (b) plots the density of permanent setpoints (as determined by our classification rule) by minute of the hour. Consistent with our priors, users schedule most setpoints on the hour or half hour (and to a lesser extent, at :15 and :45 minutes past the hour). This is both a finding and a confirmation of the validity of our classification rule.

experiment, settings in the two locations are remarkably similar.[23]

Finally, Opower and Honeywell conducted an online survey to collect baseline information on both treatment and control households in the experiment. We do not use these time-invariant household characteristics in our main analysis because they are redundant to household fixed effects, but we use them to test the validity of Opower and Honeywell's randomization process.

2.6 Balance

To test for balance, we estimate a linear probability model with an indicator for assignment to treatment as the dependent variable. Table 3 reports estimates from that model that summarize the results of our balance tests. Column (1) reports estimates based on our full sample of households, and the estimates in Columns (2) and (3) are from models estimated on subsamples by wave. The reported F-statistics test the null hypothesis that all parameters in the given model are jointly equal to zero. Consistent with an appropriate randomization process, we fail to reject the null in all three models and find that control and treatment households are statistically balanced for all outcomes.

The significance of each coefficient estimate provides an additional hypothesis test of balance. The lack of significance for all but two of the reported coefficient estimates indicates that households are balanced across all individual measures, save for mean pre-period electricity use. The significant *Mean (kWh)* estimate in Column (1) indicates that treatment households used 4.5% less electricity per hour in the pre-period on average. The lack of balance is unlikely to be due to poor randomization for multiple reasons. First, energy use was not known at the time of randomization. Second, the comparable *Mean (kWh)* estimates in Columns (2) and (3) indicate that this effect is driven by the Northern California wave of the experiment where pre-period durations are limited for some households. Finally, the *Mean (thm)* estimates from all three models indicate that mean pre-period natural gas use is balanced across experimental groups. Regardless, out of an abundance of caution, we estimate DD models to control for potential imbalance.

2.7 Descriptive Analysis

To illustrate basic patterns in the raw data and the effect of the installation of a smart thermostat on energy use, Figure 8 plots mean energy consumption against event time (days before/after the installation of a smart thermostat) by wave. Panel (a) displays electricity use, and Panel (b) illustrates the patterns in natural gas consumption. The figure does not suggest that smart thermostats have a large effect on energy use, but the raw data is too noisy to be visually conclusive.[24] For this reason, we develop empirical models that allow us to include additional controls that mitigate residual variation in the raw data and formally test for statistically significance.

[23] Average ambient temperatures are higher in Northern than Central California because of seasonal variation. The Northern California panel spans July through January, whereas the Central California panel runs from December through January.

[24] For instance, the seasonal effects of summer for electricity use and winter for natural gas can be seen in the patterns in the data.

Table 2: User Interactions Summary Statistics by Wave

	Central California			Northern California			All		
Variables	Mean	Std.Dev.	Obs.	Mean	Std.Dev.	Obs.	Mean	Std.Dev.	Obs.
Ambient Temp.	66.88	4.38	25,240	69.07	5.20	314,428	68.91	5.17	339,668
Cooling Setpoints	78.20	4.55	2,975	78.84	4.10	52,861	78.80	4.12	55,836
Heating Setpoints	64.65	5.45	5,749	63.89	5.58	71,667	63.95	5.58	77,416
Cooling Overrides	77.55	5.01	1,191	77.49	3.88	14,082	77.50	3.98	15,273
Heating Overrides	68.17	4.59	6,490	67.36	4.15	40,983	67.47	4.22	47,473
N		133			233			365	
$N \times T$		28,365			350,175			378,540	

Table 3: Balance Table

Variable	(1) All Waves Treatment Indicator	(2) Wave 1: N. CA Treatment Indicator	(3) Wave 2: C. CA Treatment Indicator
Household Characteristics			
Family in the Household Indicator	0.026	-0.026	0.085
	(0.053)	(0.071)	(0.080)
Pets in the Household Indicator	0.015	0.020	0.008
	(0.029)	(0.038)	(0.045)
HER Subject Indicator	0.019	0.002	0.045
	(0.031)	(0.040)	(0.049)
HER Recipient Indicator	0.007	-0.026	0.062
	(0.039)	(0.049)	(0.063)
Home Characteristics			
Multi-Family Home Indicator	-0.019	-0.024	0.039
	(0.080)	(0.091)	(0.166)
Year Home Built	0.000	-0.001	0.001
	(0.001)	(0.001)	(0.001)
Size of Home (Sq. Ft.)	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)
Pool Indicator	0.001	0.045	-0.079
	(0.034)	(0.044)	(0.053)
Electric Heat Indicator	0.015	-0.059	0.126
	(0.095)	(0.131)	(0.140)
Pre-Period Energy Use			
Mean (kWh)	-0.045*	-0.057**	-0.003
	(0.024)	(0.028)	(0.048)
Mean (thm)	-0.024	-0.008	-0.046
	(0.031)	(0.048)	(0.040)
N	1,385	821	564
R^2	0.013	0.019	0.021
F	0.731	0.822	0.687

Notes: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$
The table reports linear probability model estimates of the probability of assignment to treatment. Models also include indicators for month and county of recruitment, as well as indicators for missing year built, home size, and heating type data. All omitted coefficient estimates are statistically insignificant. The F-statistic tests the null hypothesis that all parameters are jointly equal to zero. We fail to reject the null in all three models.

Figure 8: Average Energy Use by Experimental Status and Wave

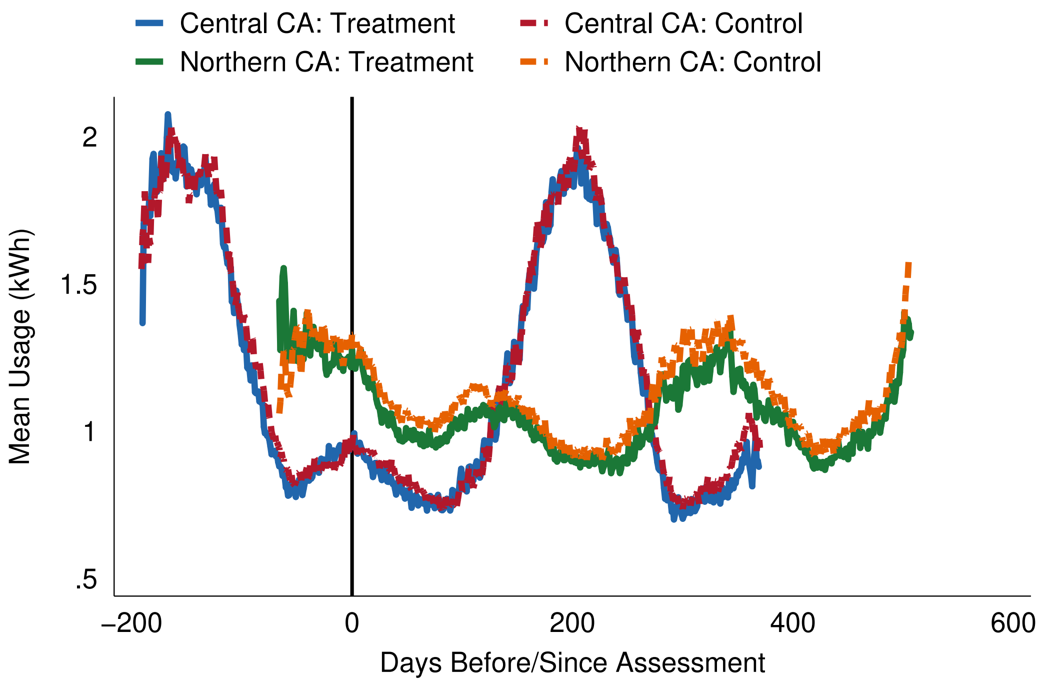

Only means based on 30 or more homes per day are included in the figure.

(a) Electricity

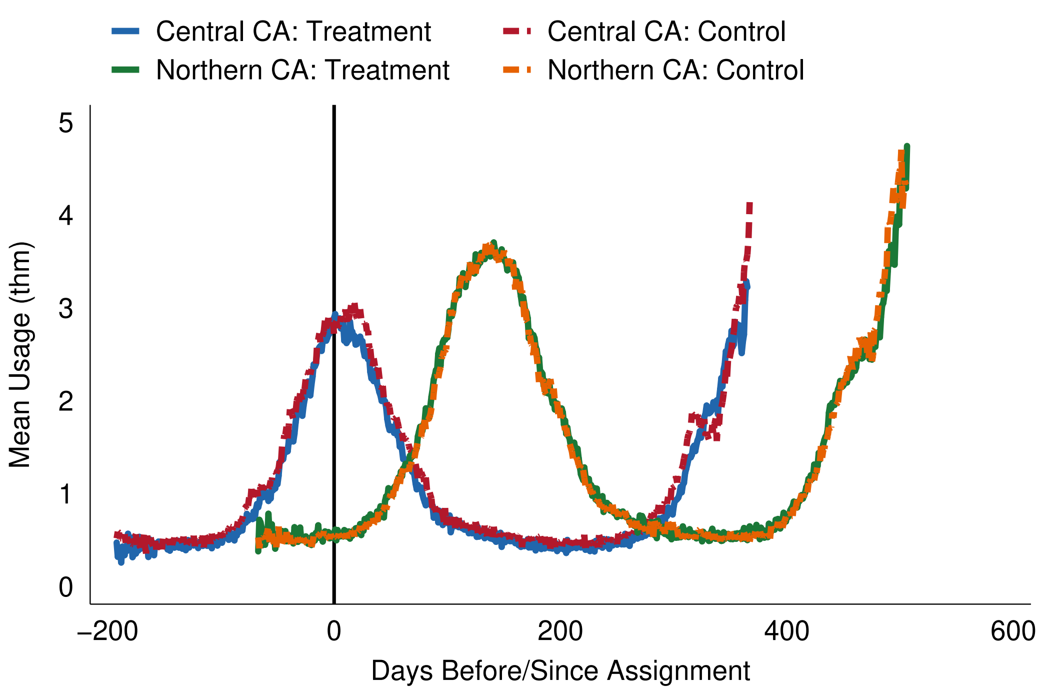

Only means based on 30 or more homes per day are included in the figure.

(b) Natural Gas

3 Empirical Model

Our field experiment randomizes receipt of a smart thermostat among eligible applicants. We observe a long time series of household-level energy use for treatment and control groups before and after experimental assignment. Both motivate our empirical strategy. Given the potential pre-period imbalance in electricity use discussed in Section 2.6, we estimate difference-in-differences (DD) models. Given the potential for self selection after randomization, we augment our DD model with instrumental variables (IV) model techniques. We begin by formalizing our model specification, then discuss identification issues.

3.1 Model Specification

We begin by modeling the effect of a smart thermostat on household i's consumption of energy type $j \in$ {electricity, natural gas} in time period t (e^j_{it}) using a DD model:

$$e^j_{it} = \alpha^j_i + \beta^j_t + \gamma^j S_i P_t + X_{it}\delta^j + u^j_{it}, \tag{1}$$

where S_i is an indicator equal to one if household i installs a smart thermostat, P_t is an indicator for post-assignment status in time period t, X_{it} is a vector of controls, α^j_i is a household fixed effect, β^j_t is a vector of time effects, and u^j_{it} is a household/time varying unobservable.[25] We cluster standard errors at the household level to account for serial correlation (Bertrand et al., 2004) and estimate the model separately for each energy type. When j denotes electricity, energy is measured in kWh and the time period is an hour. If j denotes natural gas, the energy unit is a therm and observations are recorded daily.

Our parameter of interest is γ^j, which measures the differential change in energy use across pre- and post-intervention periods for smart relative to traditional thermostat households. This specification implicitly assumes that smart thermostats have a constant effect for all households. Given that individuals in our treatment sample are each optimizing over their household's expected energy savings and installation costs when deciding whether or not to follow through on installing smart thermostat, our treatment is likely to result in heterogeneous effects and Roy (1951) selection on gains. Consistent with this underlying model of behavior, there is incomplete installation compliance among the treated households in our experiment (see Figure 1). To address concerns of bias from self selection after randomization, we estimate a DDIV model that uses our randomization as an instrument for the installation of a smart thermostat. Formally, we estimate γ^j using two-stage least squares (2SLS) methods with $E\left[Z^j_{it}u^j_{it}\right] = 0$, where $Z^j_{it} = \left(\alpha^j_i, \beta^j_t, T_i P_t, X_{it}\right)'$, and

[25] If the randomization in our experiment is valid, our coefficient of interest is identified regardless of whether or not we include household fixed effects (α^j_i), time effects (β^j_t), or additional controls (X_{it}). Thus, we begin by estimating a basic specification of the model without any additional covariates that replaces α^j_i with $\alpha^j S_i$ and β^j_t with $\beta^j P_t$. Subsequent specifications add controls for the weather (which cannot be randomized a priori), household fixed effects, and various time effects to demonstrate robustness and to improve the efficiency of our estimates. Results are qualitatively similar across all specifications.

T_i is an indicator for household i's treatment status in our experiment.[26]

3.2 Identification

If our instrument is relevant and valid, and there is one-sided noncompliance in our experiment, our DDIV coefficient of interest, γ^j, identifies the ATT of a smart thermostat (Cornelissen et al., 2016).[27] This is the average impact of a smart thermostat on the energy use of households that install one.[28]

We provide evidence that our experimental design satisfies the first two of these requirements (instrument relevance and validity) and that the third (one-sided noncompliance) is a reasonable assumption.[29] First, instrument relevance requires that assignment to treatment affects the probability that a household installs a smart thermostat. We report the first-stage F statistics with all of our results. As one would expect of a field experiment, we always easily reject the null of weak instruments.

Second, the instrument validity assumption in a DDIV model can be thought of as two separate conditions (Hudson et al., 2017). The first is the traditional IV assumption that the instrument is exogenous and the only way assignment to the treatment group affects energy use is through the installation of a smart thermostat. The second is the assumption implicit in all DD analyses that post-period randomization does not affect the pre-period values of outcomes (energy use) or treatment (smart thermostat installation). Both assumptions are satisfied by the nature of our experiment: households are randomly assigned to a treatment or control group. Assignment occurs both (shortly) after the household first interacts with the experimenter and after the household's pre-period energy use decisions have been made.[30]

Finally, if there is two-sided noncompliance in our experiment, our estimates will be confounded by substitution bias (Heckman and Smith, 1995). This is a cause for concern to the extent that "the need for treatment under question is widely acknowledged and there is competition over implementation" (Ito, 2007). This is not the case in our context as smart thermostat technology was in its infancy at the time of our study. Using data from the EIA's Residential Energy Consumption

[26]Equation 1 is the second-stage equation, and the first stage is modeled as

$$S_i P_t = \theta_i^j + \kappa_t^j + \lambda^j T_i P_t + X_{it}\pi^j + w_{it}^j. \tag{2}$$

[27]In our context, one-sided noncompliance means that while some households randomized into treatment do not install a smart thermostat, no households in the control group install one.

[28]Instead, we can replace S_i in Equation 1 with T_i to recover the ITT estimate of γ^j. This is an estimate of the average effect of being randomized into the treatment group in our experiment. We estimate DDITT models in Section 5.5 as we do not observe additional instruments for household energy-efficiency types. Alternatively, we can relax the one-sided noncompliance assumption to one of monotonicity (or uniformity): the experiment makes all households in the treatment group more (or less) likely to get a smart thermostat than they would have been otherwise. Under this alternative assumption, the DDIV specification recovers the LATE estimate of γ^j (Imbens and Angrist, 1994). This is an estimate of the average impact of a smart thermostat on the energy consumption of households that were induced to install one by our experiment.

[29]Thus, our estimates can be interpreted as ATTs.

[30]The analyses in Section 2.6 are consistent with an appropriate randomization process.

Survey (RECS), we find that two to three years after our experiment, only 4.09% of all households in the survey and 10.58% of households observationally similar to those in our study own a smart thermostat.[31] Additionally, while we are unable to directly observe whether any households in the control group upgrade their thermostat, we never observe control households using a smart thermostat on Opower platform. Thus, the available evidence supports the validity of the assumption of one-sided noncompliance in our experimental context.

4 Results

We begin by reporting estimates of the parameters in Equation 1 for electricity and natural gas in the next section. We then re-estimate the model on restricted subsamples of the data to investigate whether our main results mask significant, but offsetting, heterogeneous treatment effects. We estimate the model on subsamples by quintiles of ambient weather conditions, day of the week, and hour of the day in the subsequent section.

4.1 Main Estimates

Table 4 reports estimates of the effect of a smart thermostat on energy use based on the full sample comprised of households recruited during both waves of the experiment.[32] Panel (A) reports estimated effects on electricity usage, and Panel (B) reports analogous estimates based on consumption of natural gas. Column (1) reports estimates of a basic version of the DDIV model neither any fixed effects, time effects nor additional controls.[33] Column (2) reports estimates from a model that adds an indicator for the wave the household was recruited during, as well as linear and quadratic hourly county temperature and humidity readings as controls for recruitment and ambient weather conditions, respectively. Column (3) reports estimates from a model that adds household fixed effects to control for all time-invariant unobserved characteristics of a household and home (e.g., age and square footage of the home, number of family members).[34] Column (4) reports estimates from a model that adds month-of–year (MOY) effects to control for aggregate, time-varying effects such

[31] The RECS is not conducted annually, so we use data from the 2015 survey as it is the closest possible survey iteration subsequent to the time period observed in our data. The previous iteration of the survey in 2009 did not ask questions about smart devices. We define "observationally similar" households by restricting the RECS sample to homes that would pass Opower's initial eligibility screening to join the trial (to the extent possible given the measures available). Specifically, we condition on owner-occupied, single-family homes located in the Pacific Division (state of residence is not observed) that have a functioning central furnace or heat pump, central air conditioning, and an electrical connection. We are not able to condition on whether or not the household has a high-speed Internet connection or whether the occupants plan to move in the next year, as those questions are not part of the RECS survey.

[32] Given the issue with the observation of pre-experiment energy outcomes for households in the first wave of recruitment (illustrated in Panel (c) of Figure 4) that results in short pre-periods for some homes from the Northern California wave of the experiment, Appendix Section C.1 reports analogous estimates separately by wave based on subjects recruited during the Northern California and Central California waves of the experiment, respectively. Estimates based on these samples are qualitatively similar and do not affect the the conclusions drawn from our analysis.

[33] Relative to Equation 1, the model in Column (1) replaces α_i^j with $\alpha^j S_i$, replaces β_t^j with $\beta^j P_t$, and restricts $\delta^j = 0$.

[34] Since the experimental wave indicator is perfectly collinear with recruitment wave, we drop the wave indicator from this and subsequent specifications.

as seasonal variation in weather patterns.[35] Column (5) adds day-of-week effects to control for variation in daily usage patterns due to occupant work and schooling schedules. Finally, Column (6) replaces the time effects with day and hour-of-day effects.[36]

The coefficient estimate of -0.031 reported in Column (1) of Panel (A) indicates that a smart thermostat causes a 0.031 kWh decrease in electricity usage per hour. The cluster-robust estimate of the standard error of 0.036 reported in parentheses indicates that this estimate is statistically insignificant.[37] The estimated effect is equivalent to about two percent of the baseline energy use of 1.293 kWhs per hour (the constant). The natural gas coefficient estimate in Panel (B) of the same column is equivalent to almost 6.5% of baseline energy use, but the coefficient estimate is positive. Both estimates are well short of the savings estimates from engineering studies touted by the thermostat manufacturers. Across all specifications in both panels, the lack of economic or statistical significance indicates that smart thermostats do not reduce energy usage. In fact, for both electricity and natural gas use, the estimates reported in Column (6), are positive, and the natural gas estimate is statistically significant.

4.2 Heterogeneity in Treatment Effects

In order to investigate the possibility of significant, heterogeneous effects that are not apparent in the aggregate, we estimate the model conditional on various sub-sample selection criteria. Since smart thermostats will only have an effect on energy usage when there is a need for the HVAC system to heat or cool the house, moderate ambient temperature observations may attenuate a significant effect. To address this concern, Table 5 reports estimates by ambient temperature quintile based on our preferred specification reported in Column (5) of Table 10. If the effect of a smart thermostat is only apparent when the HVAC system is in use, we would expect to find significant effects in the upper quintiles of temperature for electricity use and in the lower quintiles for natural gas. This is not the case. Only one of the 10 estimates is statistically significant, and the significant effect occurs in the second quintile of temperature for electricity consumption. Given the overall pattern of results, this finding is likely spurious.[38]

Alternatively, since smart thermostats may only have an effect on energy use during the weekdays when individuals have predictable schedules, Table 6 reports estimates by day of the week

[35] Estimates based on models that include and week-of–year (WOY) effects and models that instead include month-by-year or week-by-year effects result in qualitatively similar results.

[36] Estimates based on models that include weather controls, day-of-week effects, and household-by-MOY (or household-by-WOY) effects do not affect our findings. The specification identifies off of hourly (electricity) or daily (gas) variation in usage within a household at a particular time of year. Intuitively, identification comes from the change in consumption in a given month of a the year for a treated home before and after treatment, relative to that same change for a control home. We also estimate models that include ZIP Code-by-MOY and ZIP Code-by-WOY effects that similarly identify off of variation within a neighborhood at a particular time of year. Again, results are qualitatively similar.

[37] Standard errors in parentheses are clustered at the household level. The rk *LM* and Wald *F* statistics are first-stage diagnostic tests of under and weak identification, respectively, in models with non-i.i.d. errors. In all specifications, we reject the nulls of an under or weakly identified model. See Kleibergen and Paap (2006) for details.

[38] In Appendix Section C.2, we report estimates from analogous models that condition on ambient humidity and heat index quintiles. Results are qualitatively similar.

Table 4: ATT Estimates of the Effect of a Smart Thermostat on Energy Use

	(1)	(2)	(3)	(4)	(5)	(6)
	Power Use (kWh or thm)					
Panel A: Electricity (kWh)						
$\hat{\gamma}^{kWh}$	-0.031	-0.031	-0.003	-0.001	-0.001	0.026
	(0.036)	(0.035)	(0.022)	(0.022)	(0.022)	(0.017)
Constant	1.293***	2.873***				
	(0.024)	(0.060)				
N	1,379	1,379	1,379	1,379	1,379	1,379
$N \times T$	16,421,734	16,421,734	16,421,734	16,421,734	16,421,734	16,421,734
R^2	0.004	0.118	0.135	0.091	0.091	0.051
F statistic	67.704	538.083	818.852	749.195	743.903	547.479
rk LM statistic	738.263	749.372	611.958	612.274	612.275	488.960
rk Wald F statistic	790.294	819.435	1,948.381	1,951.624	1,951.629	1,931.185
Panel B: Natural Gas (thm)						
$\hat{\gamma}^{thm}$	0.062	0.065	0.028	0.023	0.023	0.055**
	(0.060)	(0.049)	(0.028)	(0.026)	(0.026)	(0.022)
Constant	0.963***	14.320***				
	(0.028)	(0.227)				
N	1,369	1,369	1,369	1,369	1,369	1,369
$N \times T$	677,304	677,304	677,304	677,304	677,304	677,304
R^2	0.011	0.429	0.482	0.104	0.104	0.015
F statistic	126.946	685.010	910.597	687.556	686.021	87.637
rk LM statistic	733.785	744.065	618.764	619.162	619.163	497.269
rk Wald F statistic	790.386	817.152	1,976.210	1,980.104	1,980.097	1,958.933
Wave Indicator		x				
Weather Controls		x	x	x	x	x
HH Fixed Effects			x	x	x	x
Month-of-Year Effects				x	x	
Day-of-Week Effects					x	
Day Effects						x
Hour-of-Day Effects						x

Note: Standard errors in parentheses are clustered at the household level.

*** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$.

All estimates are based on a sample comprised of both waves of the experiment. Estimates based on the "Northern California" and "Central California" samples do not qualitatively affect our results. See Appendix Section C.1 for full results. Note that the estimates reported in Column (2) are based on a model that includes an indicator for the first wave of experiment (N. CA). This indicator is perfectly co-linear with household fixed effects, so it is dropped from subsequent models. The sample used to produce the estimates in Panel A is based on *hourly* electricity meter readings in kWh, while the sample underlying the estimates in Panel B is based on *daily* natural gas meter readings (thm). Based on the values of the rk *LM* and Wald *F* statistics, we reject the nulls of an under or weakly identified model across all specifications.

and by weekday/weekend. Across all days of the week and when we aggregate to the weekday/weekend level, we find no evidence that smart thermostats reduce energy consumption. Similarly, smart thermostats may only have an effect during the times of day that individuals typically schedule permanent temperature changes (e.g., before leaving for work/school or after returning home). Table 7 reports estimates by hour of the day. We are only able to calculate estimates conditional on the hour of the day for the effects of a smart thermostat on electricity usage, as we observe natural gas use at the daily level. Again, there is scant evidence that smart thermostats have a significant effect on energy use.[39]

[39]We also report estimates from models that condition on both hour of the day and weekday/weekend in Appendix Section 13. Results are qualitatively similar.

Table 5: ATT Estimates of the Effect of a Smart Thermostat on Energy Use by Ambient Temperature Quintile

	(1) Quintile 1	(2) Quintile 2	(3) Quintile 3	(4) Quintile 4	(5) Quintile 5
	Power Use (kWh or thm)				
Panel A: Electricity (kWh)					
$\hat{\gamma}^{kWh}$	-0.036	-0.033*	-0.024	-0.008	0.009
	(0.022)	(0.019)	(0.019)	(0.024)	(0.044)
N	1,376	1,379	1,379	1,379	1,378
$N \times T$	3,345,085	3,541,064	3,239,489	3,102,224	3,193,872
R^2	0.000	0.000	0.000	0.000	0.001
F statistic	1.522	1.610	14.502	16.745	24.597
rk LM statistic	368.164	652.296	681.468	600.120	545.434
rk Wald F statistic	1,379.806	1,920.331	1,966.682	1,879.175	1,769.185
Panel B: Natural Gas (thm)					
$\hat{\gamma}^{thm}$	-0.054	-0.013	0.005	-0.008	0.010
	(0.064)	(0.038)	(0.023)	(0.018)	(0.015)
N	1,364	1,366	1,369	1,368	1,365
$N \times T$	145,525	147,440	120,087	138,512	125,737
R^2	0.001	0.000	0.000	0.000	0.000
F statistic	22.958	0.550	6.339	6.145	0.431
rk LM statistic	360.657	435.244	563.227	699.424	403.356
rk Wald F statistic	1,375.353	1,587.271	1,323.568	1,802.507	1,377.126
Weather Controls					
HH Fixed Effects	x	x	x	x	x
Month-of-Year Effects	x	x	x	x	x
Day-of-Week Effects	x	x	x	x	x

Note: Standard errors in parentheses are clustered at the household level.

*** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$.

All estimates are based on a sample comprised of both waves of the experiment. The sample used to produce the estimates in Panel A is based on *hourly* electricity meter readings in kWh, and temperature quintiles are calculated from the distribution of *hourly* average ambient temperature readings. The sample underlying the estimates in Panel B is based on *daily* natural gas meter readings (thm), and temperature quintiles are calculated using the distribution of *daily* average ambient temperature readings. Based on the values of the rk LM and Wald F statistics, we reject the nulls of an under or weakly identified model across all specifications.

Table 6: ATT Estimates of the Effect of a Smart Thermostat on Energy Use by Day of the Week and Weekday/Weekend

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Weekday	Weekend
	Power Use (kWh or thm)								
Panel A: Electricity (kWh)									
$\hat{\gamma}^{kWh}$	-0.006	-0.014	-0.007	-0.001	0.005	0.005	0.010	-0.002	0.002
	(0.024)	(0.023)	(0.024)	(0.023)	(0.023)	(0.023)	(0.023)	(0.022)	(0.023)
N	1,379	1,379	1,379	1,379	1,379	1,379	1,379	1,379	1,379
$N \times T$	2,338,599	2,331,710	2,331,777	2,331,619	2,362,409	2,362,700	2,362,920	11,720,215	4,701,519
R^2	0.115	0.086	0.076	0.072	0.084	0.080	0.110	0.081	0.112
F statistic	682.877	596.313	582.663	552.920	692.183	555.214	688.254	694.402	732.848
rk LM statistic	610.468	605.939	604.236	600.642	624.896	620.751	616.090	611.775	613.368
rk Wald F statistic	1,946.894	1,941.134	1,936.473	1,933.095	1,972.928	1,966.289	1,952.711	1,951.954	1,950.200
Panel B: Natural Gas (thm)									
$\hat{\gamma}^{thm}$	0.024	0.015	0.026	0.024	0.018	0.017	0.033	0.022	0.028
	(0.029)	(0.029)	(0.029)	(0.029)	(0.029)	(0.030)	(0.031)	(0.026)	(0.027)
N	1,369	1,369	1,369	1,369	1,369	1,369	1,369	1,369	1,369
$N \times T$	94,575	96,480	96,480	96,474	97,760	97,764	97,771	484,958	192,346
R^2	0.135	0.090	0.112	0.104	0.136	0.085	0.098	0.102	0.114
F statistic	646.968	296.353	634.809	519.942	661.750	588.390	592.766	616.326	737.276
rk LM statistic	622.274	611.094	609.913	606.781	630.704	626.904	623.225	617.610	622.893
rk Wald F statistic	1,981.320	1,965.987	1,959.102	1,960.855	1,998.501	1,993.533	1,984.444	1,978.155	1,983.795
Weather Controls	x	x	x	x	x	x	x	x	x
HH Fixed Effects	x	x	x	x	x	x	x	x	x
Month-of-Year Effects	x	x	x	x	x	x	x	x	x

Note: Standard errors in parentheses are clustered at the household level.

*** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$.

All estimates are based on a sample comprised of both waves of the experiment. The sample used to produce the estimates in Panel A is based on *hourly* electricity meter readings in kWh, while the sample underlying the estimates in Panel B is based on *daily* natural gas meter readings (thm). Based on the values of the rk *LM* and Wald *F* statistics, we reject the nulls of an under or weakly identified model across all specifications.

Table 7: ATT Estimates of the Effect of a Smart Thermostat on Electricity Use by Hour of the Day

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	12:00	1:00	2:00	3:00	4:00	5:00	6:00	7:00	8:00	9:00	10:00	11:00
	Power Use (kWh)											
Panel A: AM												
$\hat{\gamma}^{kWh}$	-0.022	-0.012	-0.020	-0.030	-0.015	0.009	0.003	-0.003	0.005	-0.029	-0.041	-0.042
	(0.027)	(0.023)	(0.021)	(0.021)	(0.021)	(0.023)	(0.024)	(0.027)	(0.030)	(0.036)	(0.039)	(0.042)
N	1,379	1,379	1,379	1,379	1,379	1,379	1,379	1,379	1,379	1,379	1,379	1,379
$N \times T$	684,283	684,283	682,930	684,283	684,283	684,283	684,283	684,283	684,284	684,290	684,291	684,295
R^2	0.423	0.442	0.460	0.456	0.454	0.451	0.455	0.436	0.393	0.392	0.410	0.431
F statistic	265.776	246.431	217.654	202.070	184.084	174.976	155.358	175.907	179.396	206.676	254.733	305.851
rk LM statistic	614.284	614.247	614.352	614.225	614.218	614.175	614.184	614.190	614.150	614.089	613.938	613.624
rk Wald F statistic	1,956.322	1,956.072	1,956.121	1,956.061	1,956.057	1,955.827	1,955.958	1,955.923	1,955.732	1,955.709	1,955.444	1,955.012
Panel B: PM												
$\hat{\gamma}^{kWh}$	-0.028	-0.005	0.015	0.021	0.046	0.076*	0.048	0.035	-0.004	-0.032	-0.022	-0.019
	(0.045)	(0.047)	(0.048)	(0.047)	(0.045)	(0.042)	(0.039)	(0.036)	(0.034)	(0.032)	(0.031)	(0.027)
N	1,379	1,379	1,379	1,379	1,379	1,379	1,379	1,379	1,379	1,379	1,379	1,379
$N \times T$	684,296	684,299	684,301	684,304	684,307	684,308	684,308	684,308	684,308	684,308	684,308	684,308
R^2	0.444	0.449	0.448	0.453	0.452	0.442	0.431	0.426	0.425	0.424	0.423	0.421
F statistic	383.020	453.596	527.898	615.775	665.775	701.382	748.405	680.792	617.310	548.500	409.111	308.266
rk LM statistic	612.714	611.982	611.318	610.725	610.204	609.946	609.792	609.666	609.584	609.590	609.614	609.611
rk Wald F statistic	1,952.825	1,950.654	1,949.534	1,948.465	1,947.225	1,946.455	1,946.583	1,946.181	1,946.106	1,946.458	1,947.013	1,947.114
Weather Controls	x	x	x	x	x	x	x	x	x	x	x	x
HH Fixed Effects	x	x	x	x	x	x	x	x	x	x	x	x
Month-of-Year Effects	x	x	x	x	x	x	x	x	x	x	x	x
Day-of-Week Effects	x	x	x	x	x	x	x	x	x	x	x	x

Note: Standard errors in parentheses are clustered at the household level.

*** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$.

All estimates are based on a sample comprised of both waves of the experiment. The sample used to produce the estimates is based on *hourly* electricity meter readings in kWh. Panel A reports estimates from models that are conditional on the given hour in the AM. Panel B reports analogous estimates based on the given hour in the PM. Based on the values of the rk LM and Wald F statistics, we reject the nulls of an under or weakly identified model across all specifications.

5 Potential Mechanisms

In order to better understand the null results from our experiment, we supplement our experimental data with a high-frequency log of user interactions with their smart thermostat. Our goal is to explore potential mechanisms that may explain our robust finding that smart thermostats have a null effect on energy usage.[40] We observe thermostat interactions in the latter data only for those in the treatment group who install a smart thermostat. This limits our ability to use traditional techniques to conduct causal inference, so we present a combination of descriptive and causal evidence to examine whether patterns in the data are consistent with hypothesized mechanisms that would cause a null result in the experiment. We investigate these mechanisms by asking five questions about smart thermostat use:

1. Do users program their smart thermostats?
2. Do users program their smart thermostats for energy savings?
3. Do users deviate from their programed schedules?
4. Do user deviations increase or decrease energy use?
5. Do smart thermostats save any users energy?

If answers to the first or second questions are "no," smart thermostats would have no effect on energy use. If the answers to those questions are instead "yes," but the answers to the third and fourth questions are "yes" and "increase," user override behavior may attenuate or negate the energy savings gained by using the programmable features of the smart thermostat. Finally, if the answer to the last question is "yes," determining which users save energy informs the important dimensions that engineering models fail to account for.[41] We proceed by using the available data to answer each of these questions.

5.1 Do Users Program Their Smart Thermostats?

Peffer et al. (2013) find that programmable thermostats fail to achieve their advertised savings due, in part, to poor usability.[42] If users do not program schedules for their smart thermostats to follow because the interfaces are too complicated or they do not understand how thermostats and/or their HVAC systems work, we would not expect the installation of a smart thermostat to affect energy consumption.

[40] Ge and Ho (2019) use similar high frequency data to analyze the effect that outdoor temperature has on smart thermostat usage.

[41] Spoiler alert: our analysis suggests that the answers are "yes," "yes," "yes," "increase," and "yes."

[42] Programmable thermostats are a precursor technology to smart thermostats. The two types of thermostats share the ability to schedule permanent temperature setpoints in advance, but users cannot interact with programmable thermostats remotely, nor do they offer built-in setpoint framing. Peffer et al. (2013) report that they were so difficult to program that most users disabled their defining feature, and the ENERGY STAR program stopped certifying them in December 2009.

To determine the fraction of households who install the smart thermostat use the programmable features of the device and how long it takes them to begin doing so, Figure 9 plots the CDF of the time between the installation date and the first scheduled setpoint. The figure shows that almost all users who install a smart thermostat program at least one permanent setpoint, and most households do so almost immediately. The median time from installation to the first permanent setpoint is one day.

Figure 9: Distribution of Time from Installation to First Scheduled Setpoint

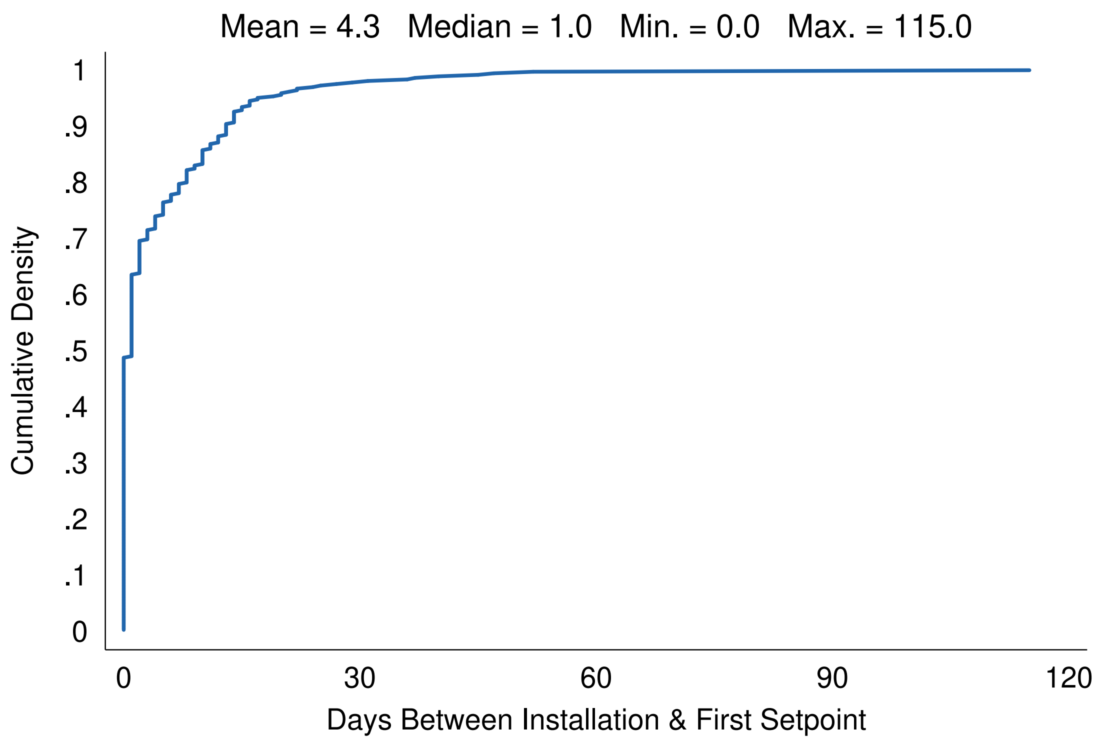

Cumulative density conditional on observing the household in the HVAC events data.

Additionally, users do not just quickly schedule a permanent setpoint, then fail to continue to use the smart features of the device. Individuals who have a smart thermostat installed as part of our experiment set an average of 3.749 (heating and cooling) setpoints per day. Figure 10 plots a measure of the frequency of permanent setpoints by hour of the day (denoted in military time) for both heating (red bars) and cooling (blue bars) setpoints. The figure provides visual evidence that setpoints occur frequently and when we would expect them: in the morning from about 5:00 AM until 10:00 AM when most users wake and leave for work and/or school. Similarly, there is a small increase in frequency of setpoints during the afternoon from 4:00 PM until 7:00 PM when users return home at the end of their days. Consistent with scheduling setpoints when most users go to sleep, we also observe frequent setpoints in the evening from about 10:00 PM until 12:00 AM. Thus, our analysis suggests that users do program their smart thermostats both quickly and frequently.

Figure 10: Average Permanent Setpoints per Household per Day by Time of Day

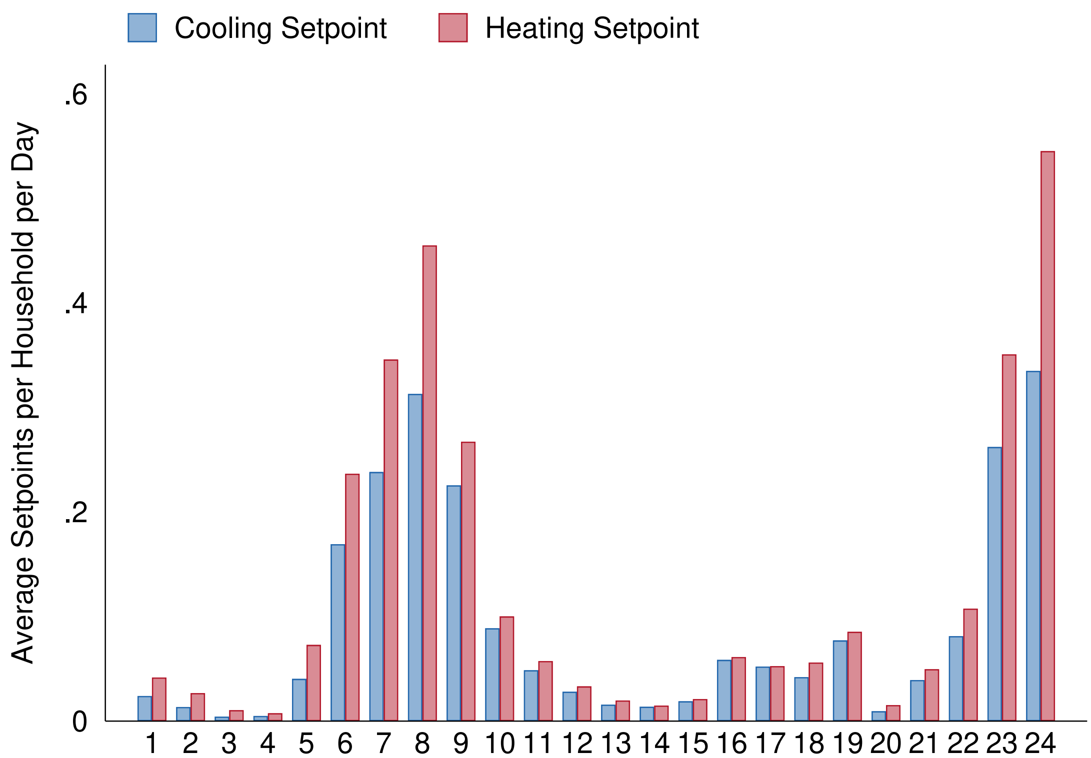

5.2 Do Users Program Their Smart Thermostats for Energy Savings?

The previous analysis is consistent with users taking advantage of their device's scheduling feature, but is inconclusive as to whether or not they are programming setpoints to achieve energy savings. To inform the latter, Figure 11 is a box and whisker plot of heating and cooling setpoints by hour of the day. The dashed lines represent the cooling and heating temperature settings the EPA recommends for energy savings of 78 degrees F for cooling and 68 degrees F for heating.[43] The figure illustrates that median (as well as the 25th and 75th percentiles of) temperatures are in line with the EPA's recommendations. According to Table 2, cooling setpoints average 78.80 degrees F and are higher than heating setpoints, which average 63.95 degrees F. Additionally, the figure illustrates that there is temporal variation in setpoints over the course of the day consistent with individuals adjusting settings when they leave the house: cooling setpoints increase slightly starting at around 9:00 AM and drop back to baseline around 3:00 PM. Heating setpoints follow a similar, but opposite pattern with a more pronounced discrepancy between evening and daytime temperature setpoints. Overall, while the figure illustrates variation in setpoints across households,

[43] Source: https://www.energy.gov/energysaver/thermostats.

our analysis suggests that users program their smart thermostats to save energy.[44]

Figure 11: Box and Whisker Plots of Permanent Setpoints by Time of Day

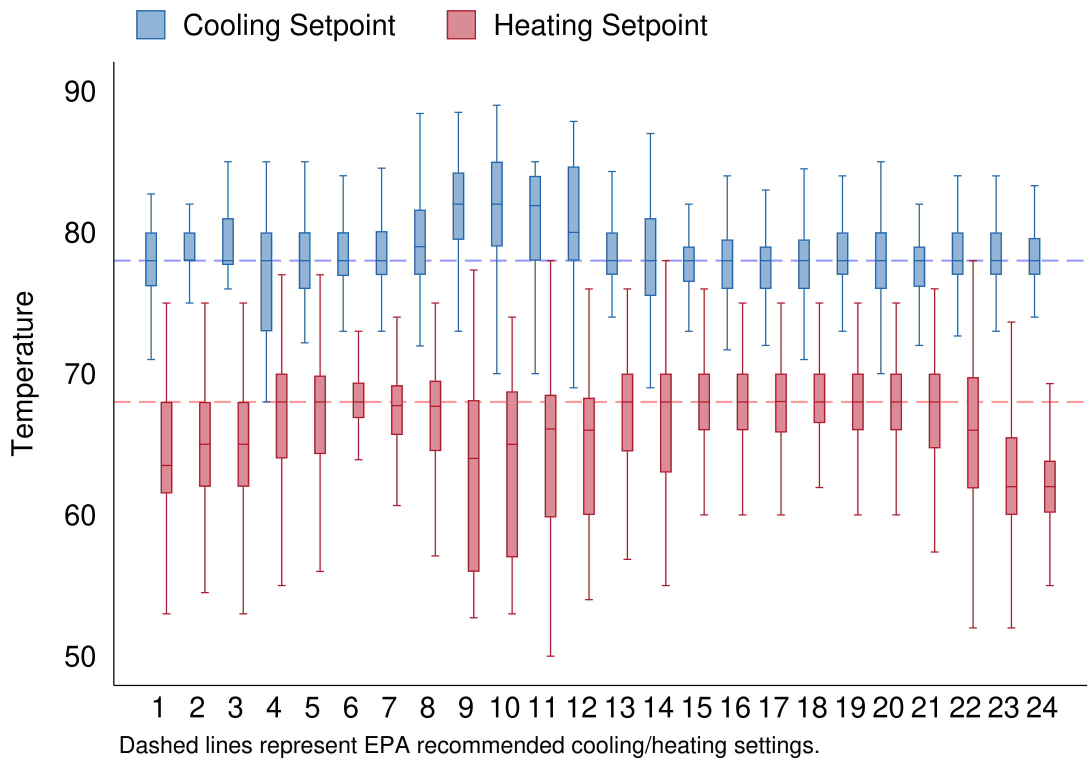

5.3 Do Users Deviate from Their Programed Schedules?

Given that users seem to be programing their smart thermostats and doing so with energy savings in mind, we turn to an alternative explanation for our null findings. The remote features of the thermostat reduce the costs associated with both permanent and temporary setpoint changes. If users program their thermostats to reduce energy usage, but the ability to more easily adjust temperature settings via a computer or smart phone makes individuals more likely to deviate from their schedules, individuals may undo the benefits of their smart thermostat. If so, the effects of the scheduling and override features of smart thermostats have opposing effects on energy use and could result in a net null effect.

To explore this possibility, Figure 12 plots a measure of the frequency of setpoint overrides by time of the day.[45] As we would expect, overrides are more frequent when most individuals

[44]Regarding the variation in setpoints, Table 2 reports standard deviations of 4.12 degrees for cooling and 5.58 degrees for heating setpoints.

[45]The figure is the analog to Figure 10 for temporary overrides, save for our definition of "per day." While users program both heating and cooling setpoints every day, we typically only observe heating (cooling) overrides on heating (cooling) degree days. Given that we predominantly observe the HVAC system events data during the fall and winter,

are likely to be awake, from about 6:00 AM to 11:00 PM. Heating overrides peak in the morning and early eventing, while cooling overrides rise throughout the day until about 6:00 PM. More importantly given our focus, the figure illustrates that users often override their permanent schedule both when heating and cooling their homes. Compared to the previously noted 3.749 setpoints per day, users in our data temporarily override their permanent setpoints an average of 1.699 times per day. The hourly measures are also substantial relative to the number of permanent setpoints reported in Figure 10.

Figure 12: Average Temporary Overrides per Household per Day by Time of Day

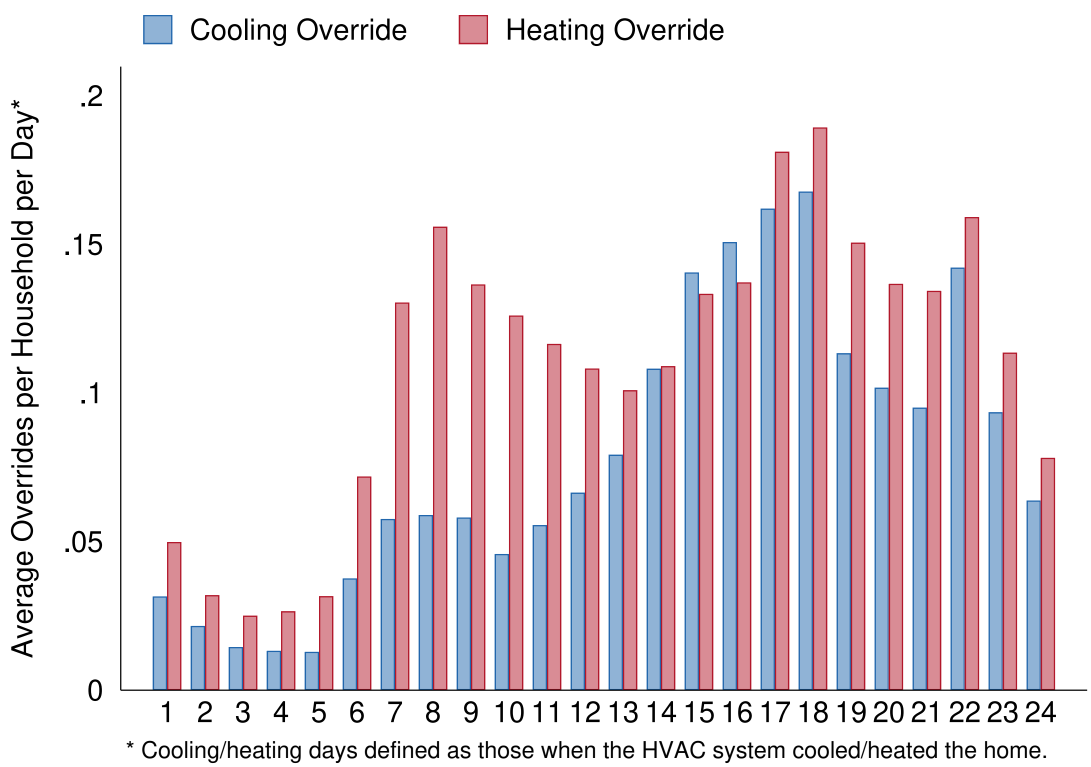

5.4 Do User Deviations Increase or Decrease Energy Use?

Evidence that smart thermostat users frequently override their setpoints offers a potential explanation for our null findings, but it is not conclusive of one. The features of the smart thermostat that lower adjustment costs both make it easier to override in ways that increase energy use (e.g., users no longer have to get off the couch or out of bed and walk to the thermostat when they are uncomfortable) and to override to decrease energy use (e.g., by toggling the HVAC system off when leaving home). To determine which effect dominates, Figure 13 plots kernel densities of the

failure to address this issue results in heating and cooling override measures that are of different magnitudes. To account for this artifact in the data, we adjust the numerator of our measure to days on which the HVAC system heated or cooled the home to standardize the scales of the heating and cooling override measures.

difference between the override temperature a user sets and the permanent setpoint, conditional on a temperature override, by temperature setting (cooling or heating). The figure illustrates that when users override their permanently scheduled setpoints, they generally do so in ways that use more energy: when cooling, they set temperatures colder and when heating, they set it warmer.[46] Taken together with the previous figure, our analysis suggests that individuals undo the benefits of their preset smart thermostat schedule when they are uncomfortable in the moment. This suggests a potential explanation for our null experimental findings.

Figure 13: Density of Difference between Temporary Override and Permanent Setpoint Temperatures by Heating/Cooling

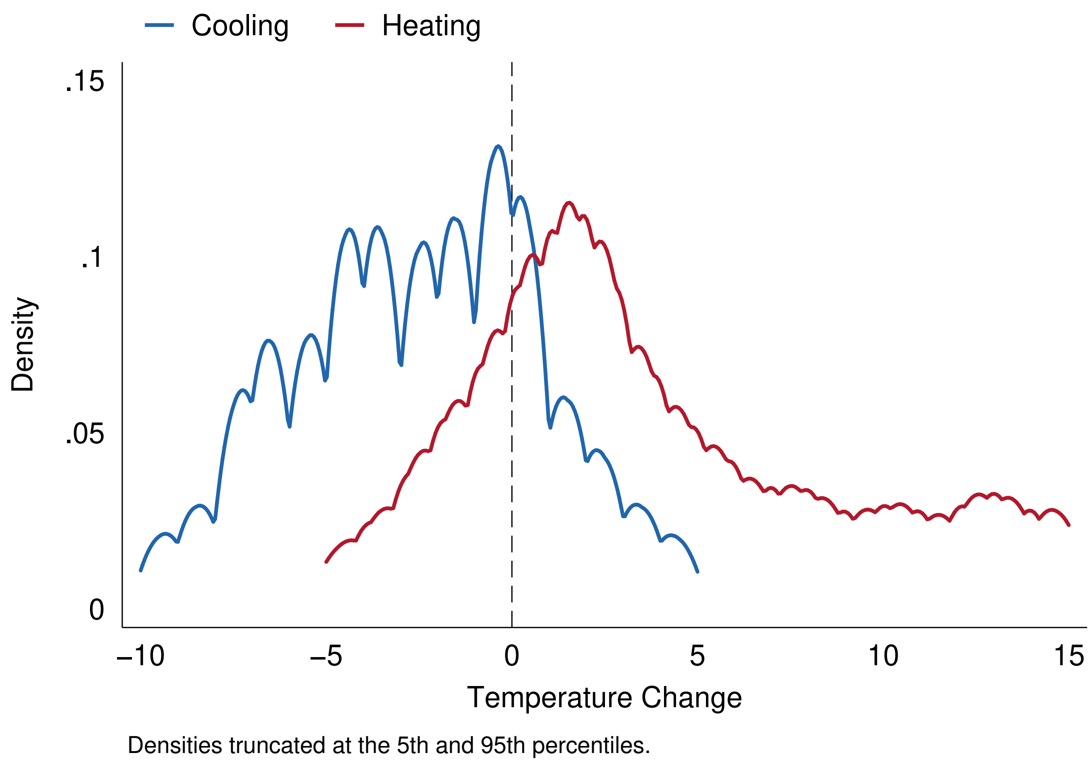

5.5 Do Smart Thermostats Save Anyone Energy?

To provide more definitive evidence that lowered adjustment costs undo the benefits of the other features of the smart thermostat and, more generally, to determine whether user behavior explains

[46]There is a non-trivial mass at large override-setpoint temperature differences (e.g., greater than 10 degrees F). This is primarily driven by a small number of households that program setpoints (~55 degrees F) that essentially turn off the HVAC system in the morning and override those setpoints at varying times in the afternoon/evening every day. This is consistent with using the programmable features of the smart thermostat based on a consistent daily departure time and a variable return time. Additionally, we note that the figure plots override-setpoint temperature differences, not override-ambient temperature differences. The ambient temperature may not actually be as low as the setpoint, so the actual temperature change caused by the override may not be so extreme.

the discrepancy between our estimates and those from engineering studies, we combine measures calculated from the HVAC events data with our experimental data and estimate heterogeneous treatment effects by multiple definitions of user energy-efficiency type. Since engineering estimates contend that smart thermostats reduce energy use absent human intervention, our goal for this analysis is to estimate whether any users "act like engineers" and are able to reduce energy use by installing a smart thermostat. Determining which user energy-efficiency types enjoy savings informs what relevant dimensions of user behavior are missing from the engineering studies.

We begin by using the HVAC events data to classify households who installed a smart thermostat based on how diligently they use their device to achieve energy savings. We do so by defining three energy-efficiency types: high-efficiency (H), low-efficiency (L), and unknown types (?). Figure 14 illustrates how this classification builds on our existing experimental design. The unknown type is necessary because we do not observe all households who install a smart thermostat in the HVAC events data. The high and low types are based on the distributions of two measures of energy-efficiency: the average number of permanent setpoints and temporary overrides observed per hour. For both metrics, we specify models based on various cutpoints between high and low types. As an example, we define high-efficiency type households based on the permanent setpoint measure as those above the median and low-efficiency types as those below the median. In contrast, for the other metric, we define high-efficiency types as those below the median number of average overrides per hour and low-efficiency types as those above the median. Figure 15 plots the CDFs of both measures of energy-efficient behavior.

Figure 14: Modified Sample Randomization with Energy-Efficiency Types

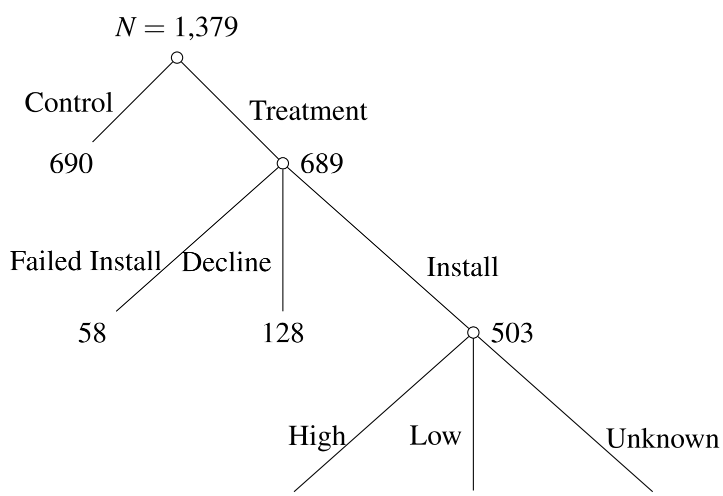

Figure 15: Distributions of Permanent Setpoints and Temporary Overrides

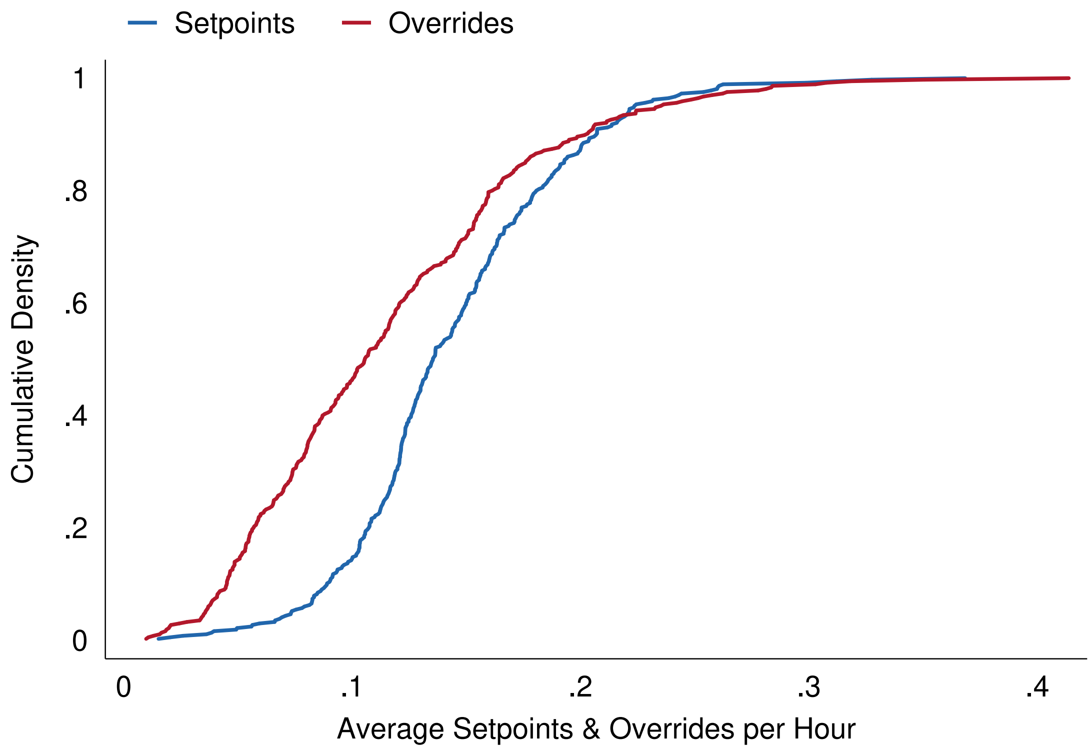

Given these classifications, we interact indicators for type with treatment and estimate a DDITT model.[47] Letting $k \in$ {H, L, ?} index the three energy-efficiency types, we estimate

$$e_{it}^{j} = \alpha_i^j + \beta_t^j + \sum_k \gamma_k^j T_i R_i^k P_t + X_{it}\beta_X^j + u_{it}^j, \tag{3}$$

where R_i^k is an indicator for household i being of energy-efficiency type k and all other indexes, variables, and parameters are defined as in Equations 1 and 2. The parameters of interest in this model are the γ_k^j which are the the ITT effects of a smart thermostat on the consumption of energy j for households of type k.

Table 2 indicates that we have the best coverage of households recruited during the Northern California wave of the experiment in the HVAC system events data, and Figure 7 illustrates that the majority of the events are observed during the fall and winter when natural gas is the predominant type of energy consumed. For these reasons, we estimate our model for natural gas use based on the the Northern California wave subsample because doing allows us to most accurately classify household energy-efficiency type.

Table 8 reports estimates of the γ_k^{thm} parameters based on this subsample. Panel (A) reports estimated effects from a model based on the permanent setpoint energy-efficiency type classifi-

[47] We are not able to estimate an analogous DDIV model because we do not have valid instruments for types.

cation, and Panel (B) reports analogous estimates based on the temporary override type definition. Column (1) reports estimates from a baseline DDITT model that does not differentiate by energy-efficiency type. Consistent with our DDIV model estimates, the effects are not statistically significant. Columns (2) through (6) report estimates based on varying definitions of the high- vs. low-type percentile cutpoint.[48] The estimates in Panel (A) in these columns indicate that households above the 10th percentile of average permanent setpoints per hour enjoy statistically significant savings, with those above the 90th percentile seeing the greatest reduction in their natural gas use. In contrast, low-efficiency types who program relatively few setpoints never reduce their energy consumption after installing a smart thermostat.

Similarly, the estimates reported in Column (6) of Panel (B) indicate that high-efficiency types above the 75th and 90th percentiles of fewest temporary overrides enjoy significant savings, while low types below the 90th percentile see a significant increase in their energy use. The overall pattern is not as consistent as in Panel (A), but the significant, negative estimate for the low-efficiency types below the 10th percentile in Column (2) (those who most often temporarily override their permanent setpoints) suggests that both those who most frequently and most infrequently override their setpoints see significant energy savings. This is consistent with the two opposing types of overrides that result from smart thermostats reducing adjustment costs.

Overall, these results support our descriptive analysis findings, inform the mechanisms that drive our null experimental estimates, and confirm that estimates of the effect of installing a smart thermostat based on engineering models fail to adequately account for how individuals use the device. Consumer and policymaker decisions based on these estimates are destined to fall short of their expected result.

[48] For instance, the estimates reported in Column (4) of Panel (A) define high-types as those with greater than the median number of setpoints per hour and low-types as those below the median.

Table 8: ITT Estimates of the Effect of a Smart Thermostat on Natural Gas Use by Energy-Efficiency Type

	(1)	(2)	(3)	(4)	(5)	(6)
		High/Low-Type Percentile Cutpoint				
	Baseline	10	25	50	75	90
			Power Use (thm)			
Panel A: Permanent Setpoint Type Classification						
$\hat{\gamma}^{thm}$	0.047					
	(0.037)					
$\hat{\gamma}^{thm}_{H}$		-0.070*	-0.104**	-0.138***	-0.173***	-0.291***
		(0.042)	(0.043)	(0.046)	(0.057)	(0.073)
$\hat{\gamma}^{thm}_{L}$		-0.049	0.104	0.018	-0.018	-0.041
		(0.143)	(0.100)	(0.065)	(0.051)	(0.044)
$\hat{\gamma}^{thm}_{?}$		0.056	0.056	0.056	0.056	0.056
		(0.160)	(0.160)	(0.160)	(0.160)	(0.160)
N	805	805	805	805	805	805
$N \times T$	398,243	398,243	398,243	398,243	398,243	398,243
R^2	0.650	0.650	0.650	0.650	0.650	0.650
F statistic	22.503	16.711	17.077	17.147	17.201	17.831
Panel B: Temporary Override Type Classification						
$\hat{\gamma}^{thm}$	0.047					
	(0.037)					
$\hat{\gamma}^{thm}_{H}$		0.067	-0.060	-0.063	-0.077*	-0.089**
		(0.144)	(0.078)	(0.057)	(0.046)	(0.041)
$\hat{\gamma}^{thm}_{L}$		-0.085**	-0.072	-0.078	-0.032	0.423***
		(0.042)	(0.046)	(0.050)	(0.071)	(0.080)
$\hat{\gamma}^{thm}_{?}$		0.056	0.056	0.056	0.056	0.056
		(0.160)	(0.160)	(0.160)	(0.160)	(0.160)
N	805	805	805	805	805	805
$N \times T$	398,243	398,243	398,243	398,243	398,243	398,243
R^2	0.650	0.650	0.650	0.650	0.650	0.650
F statistic	22.503	16.784	16.698	16.726	16.868	22.463
Weather Controls	x	x	x	x	x	x
HH Fixed Effects	x	x	x	x	x	x
Day Effects	x	x	x	x	x	x
Hour-of-Day Effects	x	x	x	x	x	x

Note: Standard errors in parentheses are clustered at the household level.
*** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$.
All estimates are based on a sample comprised of the "Northern California" wave of the experiment. The sample underlying the estimates in both panels is based on *daily* natural gas meter readings (thm).

6 Conclusion

Our work informs the efficacy of a popular technology designed to conserve energy by exploring how smart technologies affect energy use—both through actual measurement and by investigating the mechanisms that prevent the realization of advertised energy savings. We provide evidence from a field experiment wherein residential households are randomized into either a treatment group that receives a smart thermostat or a control group. The smart thermostats given to the treatment group allow households to set more advanced schedules and adjust temperature settings remotely via a smart phone app. In addition, the smart thermostats provide households in the treatment group with information designed to promote energy-efficient setpoints.

In contrast to the commonly held prior that smart thermostats are an effective way to reduce residential energy use, we find no evidence that the installation of a smart thermosat reduces household energy consumption. This null result is robust to numerous specifications. We believe that the discord between the results of our field experiment and the extant belief stems from the source of the latter: engineering studies that do not adequately account for how individuals use their smart devices. We augment our experimental analysis with data on user interactions with their smart thermostat and find evidence that supports this belief.

There are many ways to extend our research. One avenue would be to better understand how different smart technology features, that often have opposing theoretical energy impacts, affect actual usage. Another would be to understand why smart thermostats are so popular given their costs and trivial energy-efficiency benefits. This avenue speaks to the energy efficiency gap literature as outlined by Allcott and Greenstone (2012). A further avenue would be to explore the impact that such technologies have on the price elasticity of energy demand (some preliminary evidence from Herter (2007) suggests that they do). If technology can enable people to better optimize their energy consumption, then price might become even more salient and therefore make people more marginal.

In summary, cooling and heating homes, powering transportation, and producing the wealth of goods and services enjoyed in modern economies are all heavily reliant on energy. Given that most of the world relies on non-renewable resources to produce energy, this reliance does not appear to be ending any time soon (Covert et al., 2016), and the negative externalities of energy production, one of the greatest policy challenges of this century centers on energy use. Without efforts to promote energy conservation and associated reductions in greenhouse gas emissions, future generations will face a lower quality of life due to a degraded environment. We believe this paper is one small step towards ensuring that decision makers focus their energies on the smartest policies possible.[49]

[49]Puns intended.

References

Alberini, Anna, Silvia Banfi, and Celine Ramseier, "Energy Efficiency Investments in the Home: Swiss Homeowners and Expectations about Future Energy Prices," *The Energy Journal*, 2013, *34* (1), 49–86. 4

Allcott, Hunt, "Social norms and energy conservation," *Journal of Public Economics*, 2011, *95* (9), 1082 – 1095. Special Issue: The Role of Firms in Tax Systems. 2, 5

_ **and Michael Greenstone**, "Is There an Energy Efficiency Gap?," *Journal of Economic Perspectives*, February 2012, *26* (1), 3–28. 37

_ **and Todd Rogers**, "The Short-Run and Long-Run Effects of Behavioral Interventions: Experimental Evidence from Energy Conservation," *American Economic Review*, October 2014, *104* (10), 3003–37. 2

Ayres, Ian, Sophie Raseman, and Alice Shih, "Evidence from Two Large Field Experiments that Peer Comparison Feedback Can Reduce Residential Energy Usage," *The Journal of Law, Economics, and Organization*, 08 2012, *29* (5), 992–1022. 2

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How Much Should We Trust Differences-In-Differences Estimates?," *The Quarterly Journal of Economics*, 2004, *119* (1), 249. 18

Cornelissen, Thomas, Christian Dustmann, Anna Raute, and Uta Schönberg, "From LATE to MTE: Alternative Methods for the Evaluation of Policy Interventions," *Labour Economics*, 2016, *41*, 47 – 60. SOLE/EALE conference issue 2015. 3, 19

Costa, Dora L. and Matthew E. Kahn, "Energy Conservation "Nudges" and Environmentalist Ideology: Evidence From a Randomized Residential Electricity Field Experiment," *Journal of the European Economic Association*, 2013, *11* (3), 680–702. 2

Covert, Thomas, Michael Greenstone, and Christopher R. Knittel, "Will We Ever Stop Using Fossil Fuels?," *Journal of Economic Perspectives*, February 2016, *30* (1), 117–38. 37

Department of Energy, "Smart Grid Investment Grant Program Final Report," Technical Report, U.S. Department of Energy, Office of Electricity Delivery & Energy Reliability December 2016. 4

Dolan, Paul and Robert Metcalfe, "Neighbors, Knowledge, and Nuggets: Two Natural Field Experiments on the Role of Incentives on Energy Conservation," *Becker Friedman Institute for Research in Economics Working Paper*, 2015, (2589269). 2

Energy Information Administration, "Energy Consumption and Expenditures Tables," Technical Report, U.S. Department of Energy, 2015 Residential Energy Consumption Survey May 2018. 2

—, "Annual Energy Outlook 2019 with Projections to 2050," Technical Report, U.S. Department of Energy, Office of Energy Analysis January 2019. 2

—, "Monthly Energy Review," Technical Report, U.S. Department of Energy, Office of Energy Statistics August 2019. 2

Faruqui, Ahmad and Sanem Sergici, "Household response to dynamic pricing of electricity: a survey of 15 experiments," *Journal of Regulatory Economics*, Oct 2010, *38* (2), 193–225. 4

Ferraro, Paul J. and Michael K. Price, "Using Nonpecuniary Strategies to Influence Behavior: Evidence from a Large-Scale Field Experiment," *The Review of Economics and Statistics*, 2013, *95* (1), 64–73. 2

Ge, Qi and Benjamin Ho, "Energy Use and Temperature Habituation: Evidence from High Frequency Thermostat Usage Data," *Economic Inquiry*, 2019, *57* (2), 1196–1214. 27

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. 4

Heckman, James J. and Jeffrey A. Smith, "Assessing the Case for Social Experiments," *Journal of Economic Perspectives*, June 1995, *9* (2), 85–110. 19

Herter, Karen, "Residential implementation of critical-peak pricing of electricity," *Energy Policy*, 2007, *35* (4), 2121 – 2130. 37

Hudson, Sally, Peter Hull, and Jack Liebersohn, "Interpreting Instrumented Difference-In-Differences," Technical Report 2017. 19

Imbens, Guido W. and Joshua D. Angrist, "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 1994, *62* (2), 467–475. 19

Ito, Seiro, "A Practical Guide to the Program Evaluation Methods," *Health Service and Poverty: Making Health Services More Accessible to the Poor*, 2007, pp. 143–203. 19

Jessoe, Katrina and David Rapson, "Knowledge Is (Less) Power: Experimental Evidence from Residential Energy Use," *American Economic Review*, April 2014, *104* (4), 1417–38. 4

Joskow, Paul L., "Creating a Smarter U.S. Electricity Grid," *Journal of Economic Perspectives*, February 2012, *26* (1), 29–48. 4

Kleibergen, Frank and Richard Paap, "Generalized reduced rank tests using the singular value decomposition," *Journal of Econometrics*, 2006, *133* (1), 97 – 126. 21

Peffer, Therese, Daniel Perry, Marco Pritoni, Cecilia Aragon, and Alan Meier, "Facilitating energy savings with programmable thermostats: evaluation and guidelines for the thermostat user interface," *Ergonomics*, 2013, *56* (3), 463–479. PMID: 23005033. 2, 27

_ **, Marco Pritoni, Alan Meier, Cecilia Aragon, and Daniel Perry**, “How people use thermostats in homes: A review,” *Building and Environment*, 2011, *46* (12), 2529 – 2541. 2

Roy, A. D., “Some Thoughts on the Distribution of Earnings,” *Oxford Economic Papers*, 1951, *3* (2), 135–146. 18

A Smart Thermostat

Figure 16: Smart Thermostat Overview

(a) Interfaces: The left panel shows the web portal, the middle panel shows the smartphone app, and the right panel shows the thermostat.

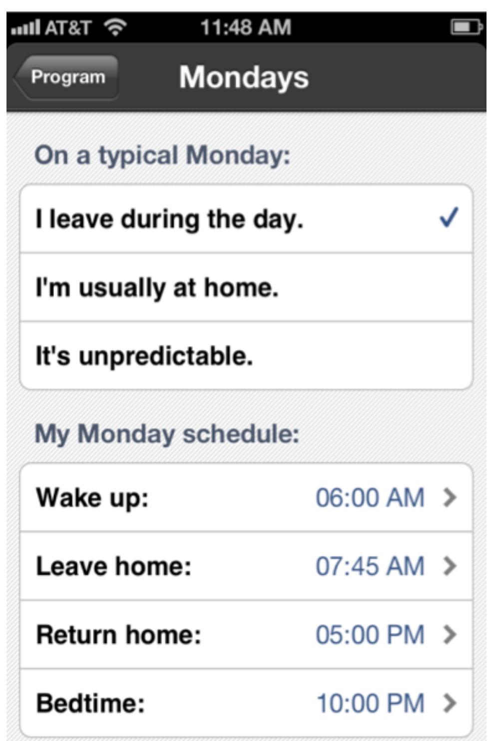

(b) Permanent Setpoint Scheduling: Screenshot of the smartphone app scheduling interface.

Figure 17: Smart Thermostat Features

(a) Setpoint Choice Messaging: Screenshots of smartphone app that shows the messaging associated with different temperature set points.

(b) Temporary Overrides: Screenshots of the smartphone app that facilitates changes to the temperature setpoint. The left panel shows the interface after the user indicates she is not home. The right panel shows the same interface when the user indicates she is at home.

B HVAC System Events Data

Figure 18: Timing of HVAC System Events

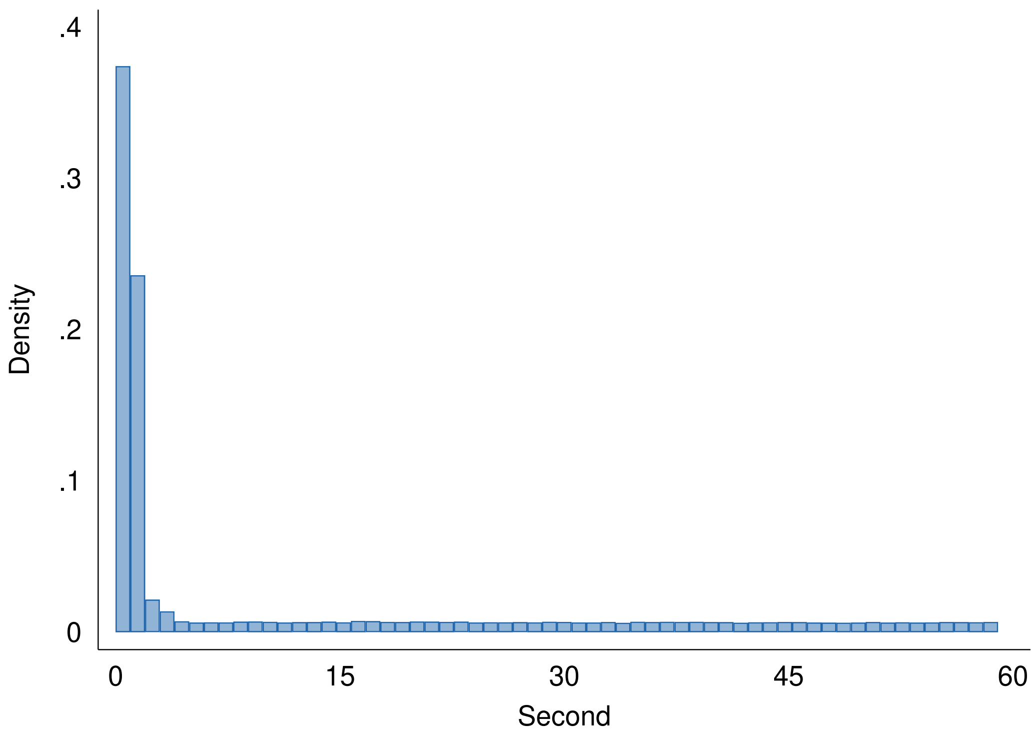

(a) Density of Temperature Changes by Second of the Minute

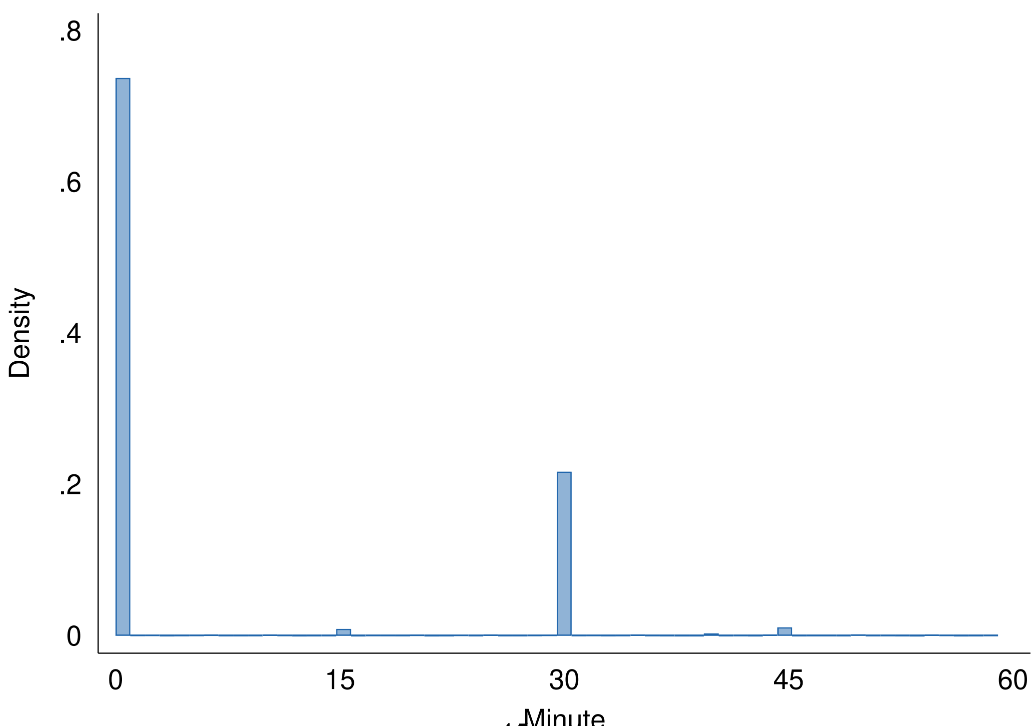

(b) Density of Permanent Setpoints by Minute of the Hour

C Additional Heterogeneous Treatment Effects Estimates

C.1 Northern and Central California Wave Estimates

Tables 9 and 10 report estimates based on a samples comprised of the "Northern California" and "Central California" waves of the experiment. Results are not qualitatively different from those based on the full sample.

Table 9: ATT Estimates of the Effect of a Smart Thermostat on Energy Use Based on the Northern California Wave of the Experiment

	(1)	(2)	(3)	(4)	(5)	(6)
	Power Use (kWh or thm)					
Panel A: Electricity (kWh)						
$\hat{\gamma}^{kWh}$	-0.055	-0.061	-0.016	-0.016	-0.016	-0.003
	(0.058)	(0.058)	(0.046)	(0.046)	(0.046)	(0.041)
Constant	1.294***	2.553***				
	(0.035)	(0.072)				
N	815	815	815	815	815	815
$N \times T$	9,729,849	9,729,849	9,729,849	9,729,849	9,729,849	9,729,849
R^2	0.003	0.051	0.062	0.052	0.052	0.017
F statistic	44.591	270.070	343.596	353.483	350.375	171.299
rk LM statistic	391.219	391.264	313.225	313.190	313.190	269.656
rk Wald F statistic	379.956	380.003	670.871	670.765	670.766	639.637
Panel B: Natural Gas (thm)						
$\hat{\gamma}^{thm}$	-0.009	0.009	0.085	0.075	0.075	0.069
	(0.061)	(0.063)	(0.068)	(0.066)	(0.066)	(0.055)
Constant	0.523***	17.739***				
	(0.020)	(0.304)				
N	805	805	805	805	805	805
$N \times T$	398,243	398,243	398,243	398,243	398,243	398,243
R^2	0.021	0.439	0.504	0.111	0.111	0.003
F statistic	801.768	568.771	674.486	519.934	520.789	22.446
rk LM statistic	386.783	386.896	313.868	313.885	313.886	270.288
rk Wald F statistic	377.042	377.090	672.580	672.617	672.609	641.179
Weather Controls		x	x	x	x	x
HH Fixed Effects			x	x	x	x
Month-of-Year Effects				x	x	
Day-of-Week Effects					x	
Day Effects						x
Hour-of-Day Effects						x

Note: Standard errors in parentheses are clustered at the household level.

*** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$.

All estimates are based on a sample comprised of the "Northern California" wave of the experiment. The sample used to produce the estimates in Panel A is based on *hourly* electricity meter readings in kWh, while the sample underlying the estimates in Panel B is based on *daily* natural gas meter readings (thm). Based on the values of the rk LM and Wald F statistics, we reject the nulls of an under or weakly identified model across all specifications.

Table 10: ATT Estimates of the Effect of a Smart Thermostat on Energy Use Based on the Central California Wave of the Experiment

	(1)	(2)	(3)	(4)	(5)	(6)
	Power Use (kWh or thm)					
Panel A: Electricity (kWh)						
$\hat{\gamma}^{kWh}$	0.009	0.006	0.002	0.002	0.002	-0.001
	(0.029)	(0.028)	(0.025)	(0.025)	(0.025)	(0.023)
Constant	1.292***	3.105***				
	(0.030)	(0.090)				
N	564	564	564	564	564	564
$N \times T$	6,691,885	6,691,885	6,691,885	6,691,885	6,691,885	6,691,885
R^2	0.003	0.204	0.229	0.110	0.109	0.047
F statistic	49.321	411.636	539.983	392.831	389.611	249.936
rk LM statistic	394.996	395.009	384.992	384.985	384.985	374.160
rk Wald F statistic	677.494	677.449	1,352.535	1,352.620	1,352.619	1,365.852
Panel B: Natural Gas (thm)						
$\hat{\gamma}^{thm}$	-0.003	0.007	0.001	0.001	0.001	-0.021
	(0.044)	(0.031)	(0.027)	(0.026)	(0.026)	(0.026)
Constant	1.101***	10.130***				
	(0.034)	(0.233)				
N	564	564	564	564	564	564
$N \times T$	279,061	279,061	279,061	279,061	279,061	279,061
R^2	0.001	0.439	0.496	0.096	0.096	0.003
F statistic	3.488	357.120	408.612	280.326	281.833	15.312
rk LM statistic	393.909	393.941	390.416	390.404	390.404	379.295
rk Wald F statistic	675.636	675.284	1,376.620	1,376.557	1,376.527	1,388.599
Weather Controls		x	x	x	x	x
HH Fixed Effects			x	x	x	x
Month-of-Year Effects				x	x	
Day-of-Week Effects					x	
Day Effects						x
Hour-of-Day Effects						x

Note: Standard errors in parentheses are clustered at the household level.

*** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$.

All estimates are based on a sample comprised of the "Central California" wave of the experiment. The sample used to produce the estimates in Panel A is based on *hourly* electricity meter readings in kWh, while the sample underlying the estimates in Panel B is based on *daily* natural gas meter readings (thm). Based on the values of the rk LM and Wald F statistics, we reject the nulls of an under or weakly identified model across all specifications.

C.2 Additional Ambient Weather Estimates

Table 11 reports estimates by ambient humidity quintile based on our preferred specification reported in Column (5) of Table 4. Table 12 reports analogous estimates by ambient heat index quintile. Across all specifications, we find little evidence that smart thermostats reduce energy consumption. The only exception that is robust to both types of energy consumption occurs when there is high humidity (Column (5) of Table 11).

Table 11: ATT Estimates of the Effect of a Smart Thermostat on Energy Use by Ambient Humidity Quintile

	(1) Quintile 1	(2) Quintile 2	(3) Quintile 3	(4) Quintile 4	(5) Quintile 5
	Power Use (kWh or thm)				
Panel A: Electricity (kWh)					
$\hat{\gamma}^{kWh}$	0.050	-0.010	-0.021	-0.041**	-0.066***
	(0.048)	(0.024)	(0.019)	(0.018)	(0.020)
N	1,379	1,379	1,379	1,379	1,379
$N \times T$	3,313,684	3,333,963	3,255,920	3,239,969	3,278,198
R^2	0.002	0.000	0.000	0.000	0.000
F statistic	45.607	3.514	8.612	4.219	7.804
rk *LM* statistic	521.960	564.647	595.843	638.333	623.192
rk Wald F statistic	1,763.238	1,860.182	1,910.165	1,944.091	1,612.296
Panel B: Natural Gas (thm)					
$\hat{\gamma}^{thm}$	0.004	-0.010	-0.005	0.047	-0.022
	(0.017)	(0.025)	(0.036)	(0.044)	(0.067)
N	1,367	1,369	1,369	1,369	1,367
$N \times T$	141,016	133,650	132,648	153,013	116,975
R^2	0.000	0.000	0.000	0.000	0.002
F statistic	0.930	0.188	0.149	14.963	65.458
rk *LM* statistic	380.444	564.518	647.032	611.390	550.812
rk Wald F statistic	1,356.189	1,740.682	1,908.480	1,522.235	1,306.659
Weather Controls					
HH Fixed Effects	x	x	x	x	x
Month-of-Year Effects	x	x	x	x	x
Day-of-Week Effects	x	x	x	x	x

Note: Standard errors in parentheses are clustered at the household level.
*** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$.
All estimates are based on a sample comprised of both waves of the experiment. The sample used to produce the estimates in Panel A is based on *hourly* electricity meter readings in kWh, and humidity quintiles are calculated from the distribution of *hourly* average ambient relative humidity readings. The sample underlying the estimates in Panel B is based on *daily* natural gas meter readings (thm), and humidity quintiles are calculated using the distribution of *daily* average ambient relative humidity readings. Based on the values of the rk *LM* and Wald F statistics, we reject the nulls of an under or weakly identified model across all specifications.

Table 12: ATT Estimates of the Effect of a Smart Thermostat on Energy Use by Ambient Heat Index Quintile

	(1)	(2)	(3)	(4)	(5)
	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
	Power Use (kWh or thm)				
Panel A: Electricity (kWh)					
$\hat{\gamma}^{kWh}$	-0.036	-0.030	-0.026	-0.009	0.009
	(0.022)	(0.019)	(0.019)	(0.024)	(0.043)
N	1,376	1,379	1,379	1,379	1,378
$N \times T$	3,296,464	3,272,861	3,296,156	3,273,130	3,283,123
R^2	0.000	0.000	0.000	0.000	0.001
F statistic	1.491	1.632	13.865	17.538	24.681
rk *LM* statistic	367.624	636.517	691.267	604.526	546.840
rk Wald F statistic	1,381.488	1,927.034	1,955.091	1,883.345	1,770.575
Panel B: Natural Gas (thm)					
$\hat{\gamma}^{thm}$	-0.060	-0.004	-0.004	-0.003	0.009
	(0.066)	(0.044)	(0.024)	(0.018)	(0.015)
N	1,364	1,366	1,369	1,367	1,365
$N \times T$	135,502	136,401	134,876	135,317	135,204
R^2	0.001	0.000	0.000	0.000	0.000
F statistic	18.708	6.519	10.808	12.692	0.289
rk *LM* statistic	351.296	404.357	586.160	702.841	413.818
rk Wald F statistic	1,364.503	1,468.623	1,403.564	1,797.169	1,406.956
Weather Controls					
HH Fixed Effects	x	x	x	x	x
Month-of-Year Effects	x	x	x	x	x
Day-of-Week Effects	x	x	x	x	x

Note: Standard errors in parentheses are clustered at the household level.
*** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$.
All estimates are based on a sample comprised of both waves of the experiment. The sample used to produce the estimates in Panel A is based on hourly electricity meter readings in kWh, and heat index quintiles are calculated from the distribution of hourly average ambient heat index readings. The heat index is calculated using temperature and humidity readings. See https://www.wpc.ncep.noaa.gov/html/heatindex_equation.shtml for the exact formula. The sample underlying the estimates in Panel B is based on daily natural gas meter readings (thm), and heat index quintiles are calculated using the distribution of daily average ambient heat index readings. Based on the values of the rk *LM* and Wald F statistics, we reject the nulls of an under or weakly identified model across all specifications.

C.3 Additional Hour of the Day Estimates

Table 13: ATT Estimates of the Effect of a Smart Thermostat on Energy Use by Hour of the Day and Weekday/Weekend

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	12:00	1:00	2:00	3:00	4:00	5:00	6:00	7:00	8:00	9:00	10:00	11:00
	Power Use (kWh)											
Panel A: Weekdays Only												
$\hat{\gamma}^{kWh}$ (AM)	-0.026	-0.015	-0.025	-0.036*	-0.017	0.011	0.000	-0.010	0.003	-0.032	-0.039	-0.042
	(0.028)	(0.024)	(0.022)	(0.021)	(0.021)	(0.023)	(0.025)	(0.027)	(0.030)	(0.036)	(0.040)	(0.042)
$\hat{\gamma}^{kWh}$ (PM)	-0.027	-0.009	0.010	0.016	0.051	0.085*	0.053	0.036	-0.007	-0.035	-0.022	-0.027
	(0.045)	(0.048)	(0.049)	(0.050)	(0.047)	(0.044)	(0.041)	(0.038)	(0.035)	(0.033)	(0.032)	(0.028)
N	1,379	1,379	1,379	1,379	1,379	1,379	1,379	1,379	1,379	1,379	1,379	1,379
$N \times T$	~488,000	~488,000	~488,000	~488,000	~488,000	~488,000	~488,000	~488,000	~488,000	~488,000	~488,000	~488,000
Panel B: Weekend Days Only												
$\hat{\gamma}^{kWh}$ (AM)	-0.010	-0.006	-0.007	-0.017	-0.009	0.004	0.011	0.013	0.011	-0.022	-0.047	-0.041
	(0.028)	(0.024)	(0.022)	(0.022)	(0.021)	(0.023)	(0.024)	(0.027)	(0.032)	(0.038)	(0.042)	(0.044)
$\hat{\gamma}^{kWh}$ (PM)	-0.032	0.005	0.027	0.031	0.034	0.053	0.037	0.032	0.002	-0.026	-0.023	0.000
	(0.048)	(0.050)	(0.050)	(0.048)	(0.045)	(0.043)	(0.040)	(0.037)	(0.036)	(0.034)	(0.031)	(0.029)
N	1,379	1,379	1,379	1,379	1,379	1,379	1,379	1,379	1,379	1,379	1,379	1,379
$N \times T$	~196,000	~196,000	~196,000	~196,000	~196,000	~196,000	~196,000	~196,000	~196,000	~196,000	~196,000	~196,000
Weather Controls	x	x	x	x	x	x	x	x	x	x	x	x
HH Fixed Effects	x	x	x	x	x	x	x	x	x	x	x	x
Month-of-Year Effects	x	x	x	x	x	x	x	x	x	x	x	x
Day-of-Week Effects	x	x	x	x	x	x	x	x	x	x	x	x

Note: Standard errors in parentheses are clustered at the household level.
*** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$.
All estimates are based on a sample comprised of both waves of the experiment. The sample used to produce all estimates is based on *hourly* electricity meter readings in kWh. All panels report estimates of the coefficient of interest from separately estimated models that are conditional on readings corresponding to the given hour in the AM or PM. The results in Panel A are based on a sample comprised of all days of the week. Panel B (C) reports analogous estimates based on weekdays (weekend days) only. Across all specifications, the min(rk *LM* statistic)=609.463 and the min(rk Wald *F* statistic)=1,942.447. Based on the values of the rk *LM* and Wald *F* statistics, we reject the nulls of an under or weakly identified model across all specifications. Full regression diagnostics are available from the authors by request.

D Recruitment and Enrollment

D.1 Subject Eligibility

Table 14: Subject Eligibility Summary

	Eligible	Not Eligible
Rent or own?	Own	Rent
Home Type	House or Condo	Apartment or Other
Phone	iPhone or Android	Blackberry or Other
# of Thermostats	1	≥ 2
A/C	Central Air	Box Unit, Fans, Other
Heating	Air Vents	Baseboard or Other
High-speed Internet?	Yes	No
Plan to move in next year?	No	Yes

D.2 Trial Recruitment and Enrollment Guide

FOR ONLINE PUBLICATION: ONLINE APPENDIX B

UTILITY Smart Thermostat Trial

Recruitment and Enrollment Guide

November 19, 2012

Introduction

Experimental Trial Information

UTILITY is running an experimental thermostat trial with Opower and Honeywell, offering eligible customers a free remote-controlled thermostat solution (a thermostat controlled by a smartphone and web application). The goal of the experiment is to test the energy savings and customer experience of the thermostat solution. Customers gain a thermostat and app that helps them save energy, by creating a customized, energy efficient schedule that fits their lifestyle.

For this trial, 1 in 2 qualifying customers will receive the thermostat solution. Customers who meet the eligibility qualifications must complete the online enrollment process to determine if they will receive a thermostat or not. At the end of the online enrollment process the system will randomly flip a coin to determine which customer will receive the remote-controlled thermostat and which will not. All customers who enroll for a chance to participate are benefiting the trial (even those who do not receive a thermostat), and it is important that all qualified customers complete the full enrollment process.

Customers should be encouraged to enroll for a chance to receive this exciting solution, which allows them to control their thermostat on-the-go. UTILITY, Opower, and Honeywell are grateful for the time each customer takes to enroll online for a chance to participate, and all customers should be thanked for their time regardless of the outcome.

Customers should be encouraged to answer all qualification and enrollment questions honestly. If a customer provides inaccurate information during enrollment it negatively impacts the trial and the customer will ultimately be turned down for the trial.

Talking Points for Recruitment Events

Initial Communication

Initial communication should be a call to action, provide quick benefits (FREE remote-controlled thermostat), provide a fun atmosphere and garner attention.

- Do you own an iPhone or an Android? If so, would you be interested in a free thermostat controlled by your smartphone?
- How would you like to gain better control of your energy use at home? You can control your thermostat at home from right here! Want to know how?
- Sign-up for a free remote-controlled thermostat, a $500 dollar value and take control of your energy consumption and improve the comfort of your home.
- I know you're in a hurry but this opportunity will allow you to take control of your energy use and you'll always come home to a house at the perfect temperature.
- Save energy while you're away and stay comfortable while you're at home, all by using your smartphone or the web.
- How would you like to control your heating/cooling by your iPhone or Android and

through the internet from anywhere in the world?

After Initial Communication

After initial communication, you should be focused on getting the customer more excited about the offering by providing key information and benefits unique to the opportunity.

- We are conducting a trial on behalf of UTILITY that allows you to interact with your heating & cooling system using your smartphone or the web. That means you can control your home's comfort at your fingertips from wherever you are. All you need is your smartphone of the web. Are you ready to take control?
- Did you know that a typical family spends almost half (49%) of its energy cost on heating and cooling? (*Source:* Energy Star)-- How would you like to have the opportunity to be selected for a special trial UTILITY is conducting to provide a limited number of customers a thermostat controlled by your smartphone? That's right you can control the comfort of your home at anytime or any place using your smartphone or the web.
- How would you like to be one of the lucky UTILITY customers who receives a free thermostat controlled on-the-go from your smartphone or the web? This is over a $500 value completely free with professional installation and a 1-year warrantee. UTILITY is conducting this trial to allow customers a unique way to reduce energy use and save money. The process for signing up only takes a few minutes of your time. Let's see if you qualify.
- Check out this free thermostat controlled by your smartphone. You'll have complete control over your comfort, and you can see how your temperature settings stack up against other participants in the trial.

Overcoming Initial Objections

Objection: "I don't have time"

- You'll never come home to a cold house again and sign-up only takes a few minutes.

Objection: "I still don't have time"

- Okay; here's how you can see if you qualify and sign-up from home (postcard)

Objection: "I don't want to give out my personal information"

- You're information is completely confidential and will be only used to determine if you qualify for the free thermostat.

Objection: "I'm not interested"

- Here is a free pen, compliments of UTILITY. Have a great day!

Initial Eligibility Screening

	Eligible	**Not eligible**

Do you rent or own your home?	Own	Rent
What kind of home do you own?	Single family, Townhome, Condo	- Apartment - Other
What kind of phone do you have?	- iPhone - Android	- Blackberry - Other
How many thermostats do you have in your home?	One (1)	Two (2) or more
How do you cool your home?	Central air	- Window box unit - Fans - Other
What is the main way you heat your home?	Air vents	- Baseboard - Other - None
Are your heating and air conditioning systems functional and have you used them the last 6 months?	Yes	No
Do you have high-speed internet access (Cable, DSL, satellite, Broadband)?	Yes	No
Do you have an available *ethernet* port on your internet router?	Yes	No
Do you plan to move to a new home in the next 12 months?	**No**	Yes
Will other adults in your household object to enrolling in this program?	**No**	Yes

Customer Does NOT Pass Initial Eligibility Screening

- Thank you for your interest, but unfortunately you don't meet the eligibility requirements for this trial. However, UTILITY is developing a number of residential energy efficiency programs that you may qualify for. Please fill out this post card in to enable them to contact you in the future for other offerings. Thank you and please accept this free pen, compliments of UTILITY. We appreciate your time!

- If you do know someone else who may be interested, please let them know about this free trial and they can sign-up right away. (Staffer hands the customer a post card.)

Customer Passes Initial Screening

- Great! You've pre-qualified to participate in the selection process, which only takes a few minutes. Would you like to learn how the thermostat and app works? (demo)

- Let's get you signed-up and see if you are selected to join the UTILITY Smart Thermostat Trial, with a free remote- controlled thermostat and professional installation. The sign-up process just takes a few minutes and we can help you complete it here.

- You'll need your UTILITY account number for enrollment. You can use my phone to retrieve your utility account number from UTILITY. You will also be asked to provide the last four digits of the Social Security Number of the UTILITY account holder—this may be you or a housemate. Staffer provides customer phone & contact number (1-888-743-0011).

Customer is Selected to Join the Trial

Encourage customers to take the first available appointment. Explain that technicians are only in the area for a limited amount of time.

- Congratulations! You've been selected to participate in the UTILITY Smart Thermostat Trial. A customer service representative will contact you with further information about your free installation. You will receive an email reminder with the date and time of your installation appointment, but you may want to write it down now, so you don't forget.
- Tell your friends and family to see if they are eligible and sign-up online! (postcard)
- Here is a free lens cleaner or smartphone holder for your smartphone, compliments of UTILITY. We appreciate your time!
- You will be contacted within a few days to confirm your eligibility and appointment time. (Honeywell CSR will conduct a follow-up call to confirm appointment time & answer any additional questions)

Customer is NOT Selected for the Trial

Thank you for your interest in the Smart Thermostat Trial. Unfortunately, this is currently a trial so participation cannot be granted for everyone.

- In the event the trial is extended, would you like to leave your contact information, which will only be used to contact you regarding other opportunities to participate in UTILITY residential trials or programs?
- Please accept this free pen, compliments of UTILITY. Have a great day.
- Tell your friends and family to see if they are eligible and sign-up! (postcard).
- Here is a free lens cleaner or smartphone holder for your smartphone, compliments of UTILITY. We appreciate your time!

How Online Enrollment Works

If a customer passes the initial qualification screening, direct them to the Opower Web application to enroll online. Eligible customers have a 1 in 2 chance of being selected to receive a thermostat.

Enroll online at: https://thermostat.opower.com/

The customer begins by clicking "See if your household qualifies."

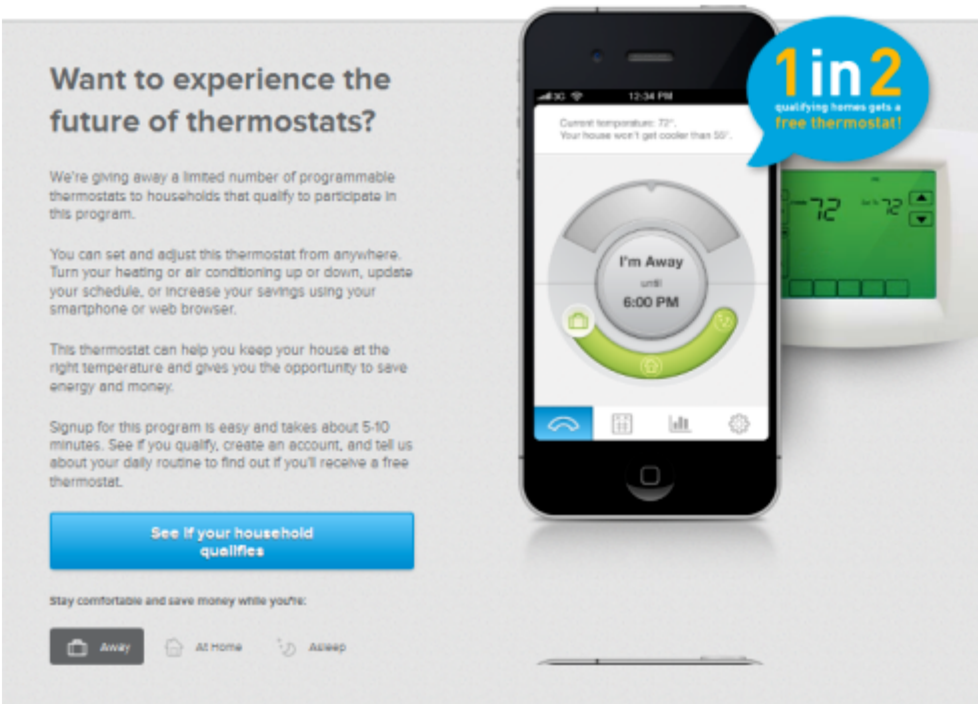

Verifying if the Household Qualifies

In order to verify that they can participate in the program, customers must answer a series of questions about their home.

On the first verification screen, they are asked to provide the following information:

- Zip code: Qualified zip codes are those within the greater Fresno and Bakersfield areas, see list provided by Honeywell.
- Whether they rent or own: Customers must own their own home.
- What kind of home they live in: Customers can select any option except "other."
- Whether they plan on moving in the next year: Customers must plan on remaining in the same home.
- What kind of phone they have: Customers must have an iPhone or Android phone if the utility program requires a smartphone.

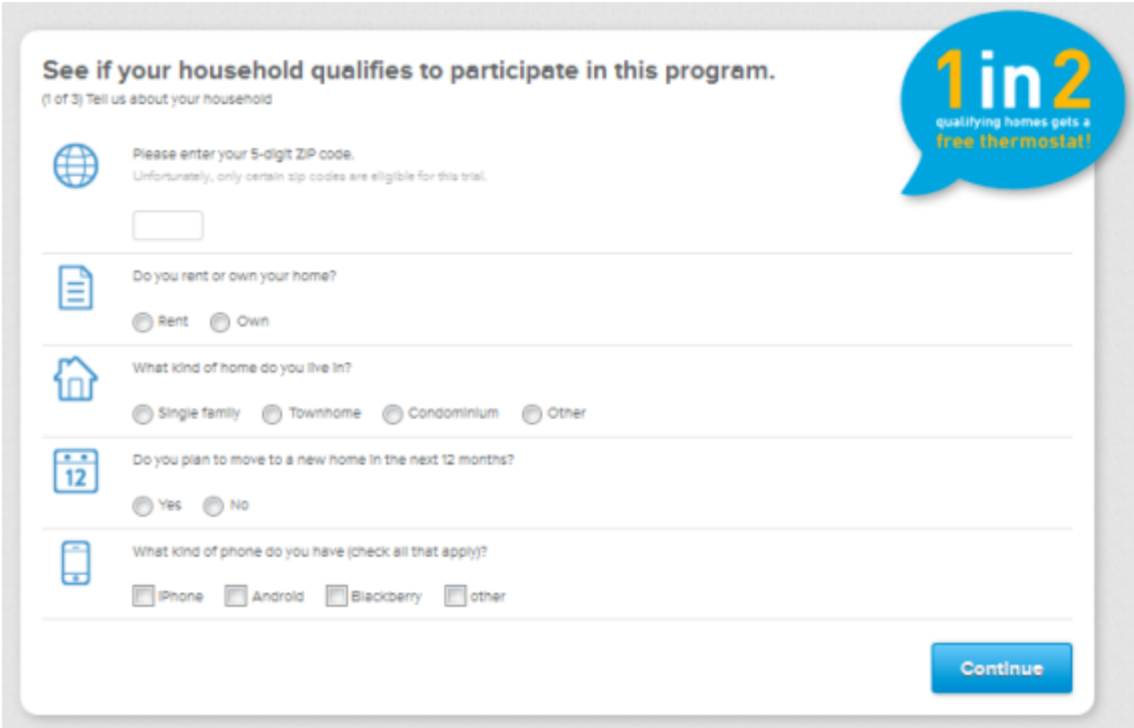

If a customer qualifies based on the answers to the questions above, they are asked to provide the following additional information:

- Number of thermostats: Customers can have only one thermostat.
- Primary cooling system: Customers must have central air.
- Main way they heat their home: Customers must have a gas furnace.
- If their air conditioning and heat are currently working: Customers must have an operational air conditioner and heater that they have used in the last 6 months.

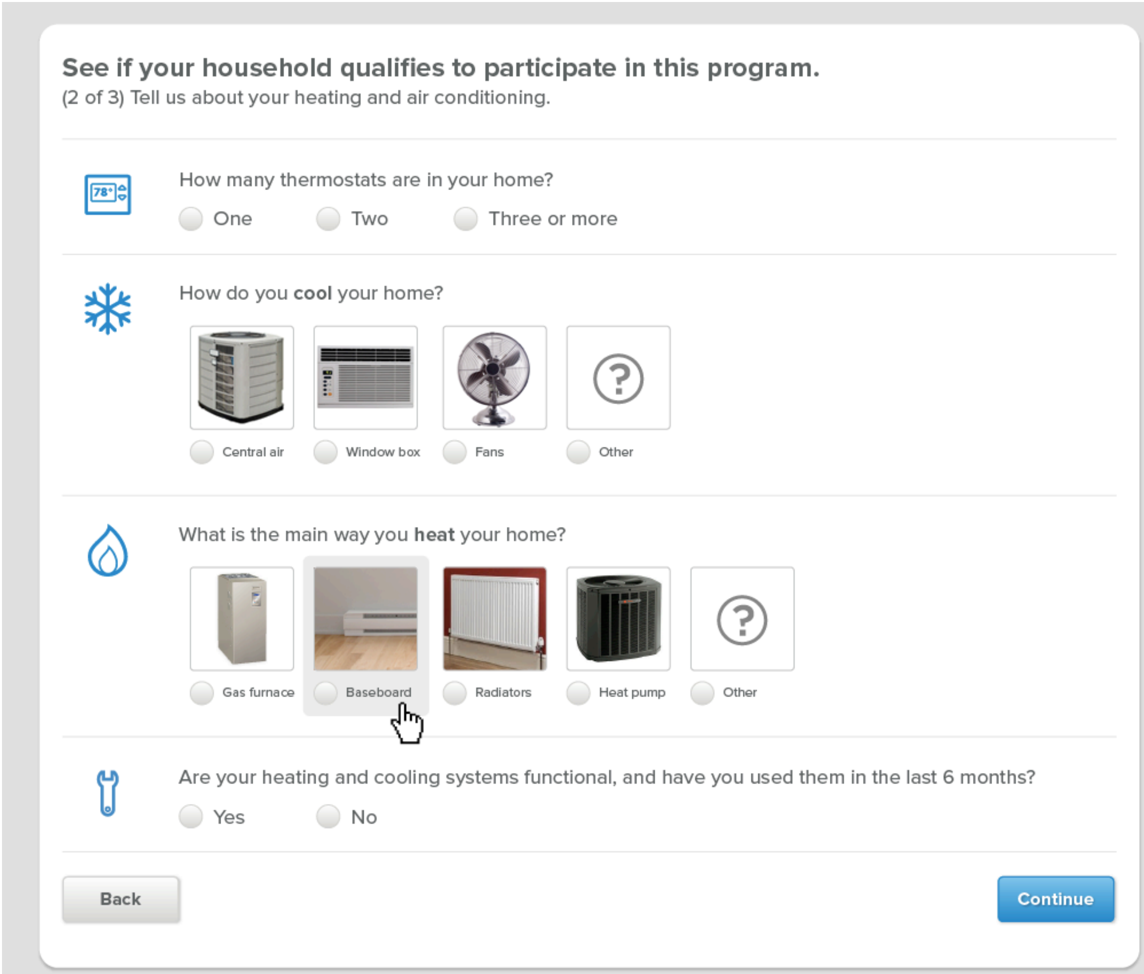

Finally the customer is asked, if they:

- Have high-speed Internet access: Customers must have high-speed access.
- Have an available Ethernet port on their router: Customers must have an available port.
- Are in agreement with the terms and conditions of the program: Customers must agree to the terms. Terms vary by utility.

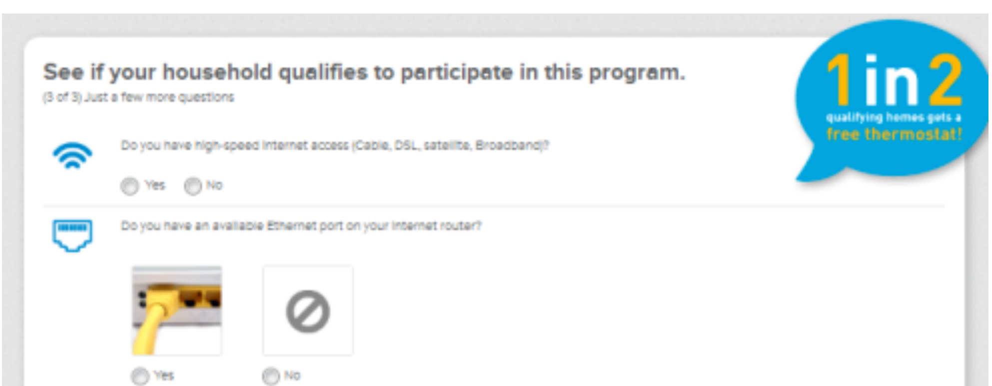

When they complete the final verification screen, they are told if they are eligible to receive an account. They must meet all of the qualifications to be considered for the program.

If a customer answers any of the qualification questions with a response that makes them ineligible, they are excluded from the program.

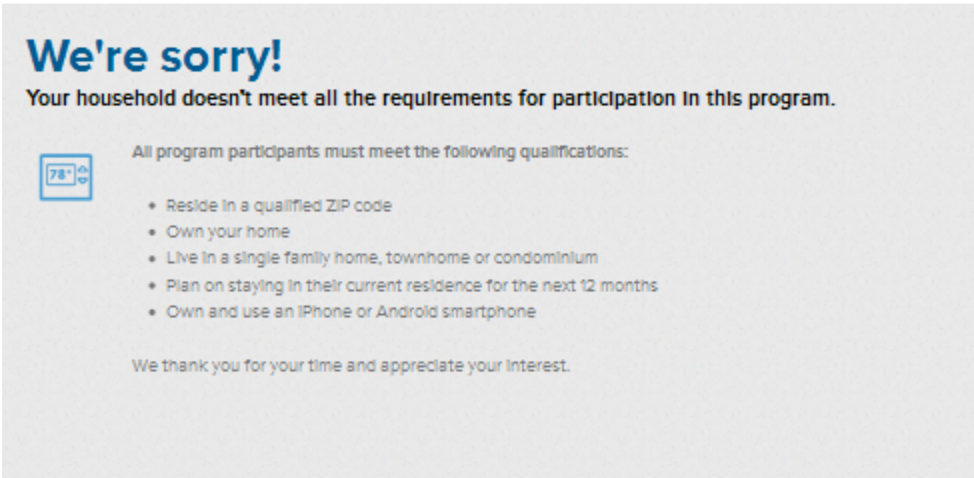

Creating an Account

Customers who are eligible for the program are required to enter the following information to create an account:

- The email address they will use to access the Web application. Basic validation is performed to verify that the email address is well-formed.
- A unique password. The password must be at least eight characters long. Passwords must not be or contain the customer's name or email address.
- Customers enter the same password again and are prompted to correct the password if it is not identical in the two password fields.
- The full name of the utility account holder exactly as it appears on the utility bill. The customer enrolling in the program must enter the name of the utility account holder as it appears on the utility bill, even if they are not the account holder.
- The utility account number exactly as it appears on the utility bill. This includes spaces or any other characters included in the data.

Customers are prompted to agree to the Opower Terms of Use.

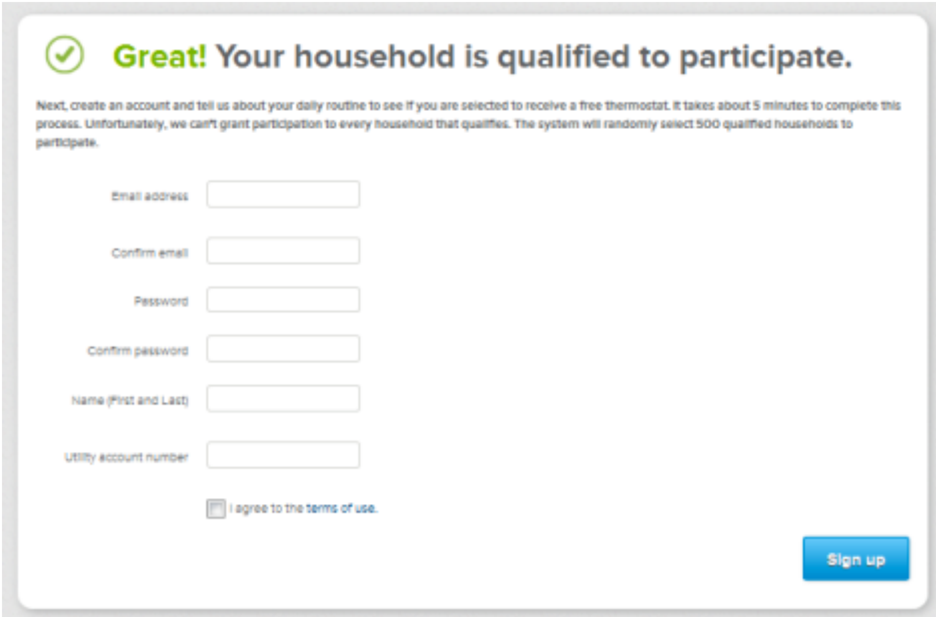

Customers submit their account information, and then a new page prompts the customer to check their email.

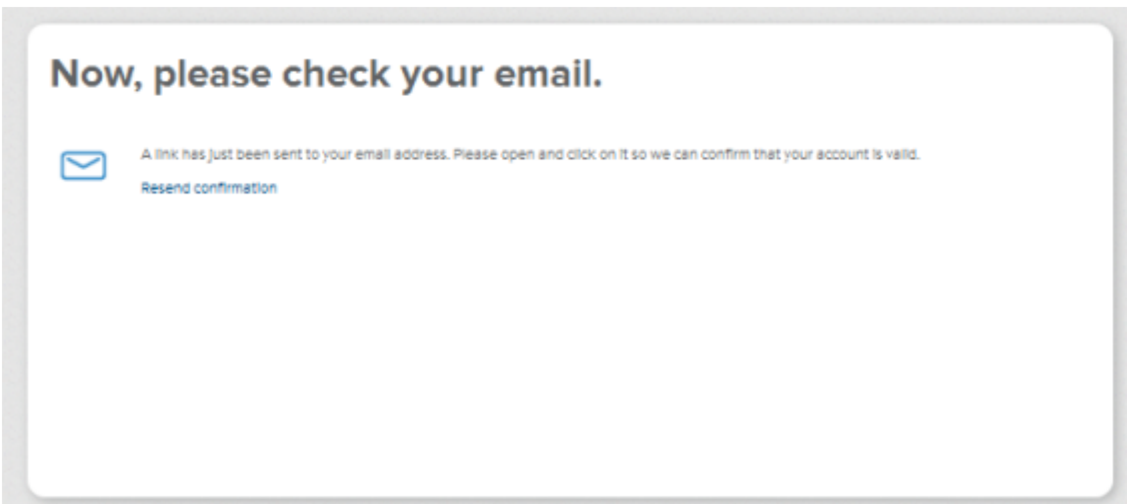

Customers should receive an email message at the address they specified. If the customer does not receive the email, they have the option to "Resend confirmation" in the Web application. The email is titled "Your Thermostat," and it will arrive from an @opower.com email address. The customer may need to check their junk/spam folder for the email.

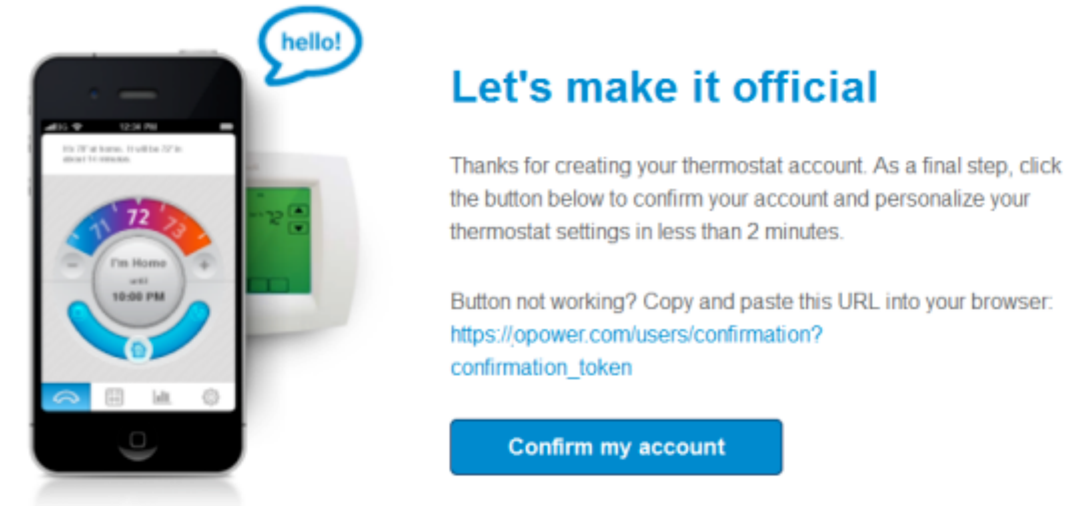

The customer must click "Confirm my account" to complete their registration and verify their email address. If nothing happens when the button is clicked, the customer can copy and paste the customer-specific URL provided in the email to their Internet browser to confirm the account.

Thermostat Registration

Once the customer has confirmed their account, they are provided with more information about the program and asked to describe their daily routine.

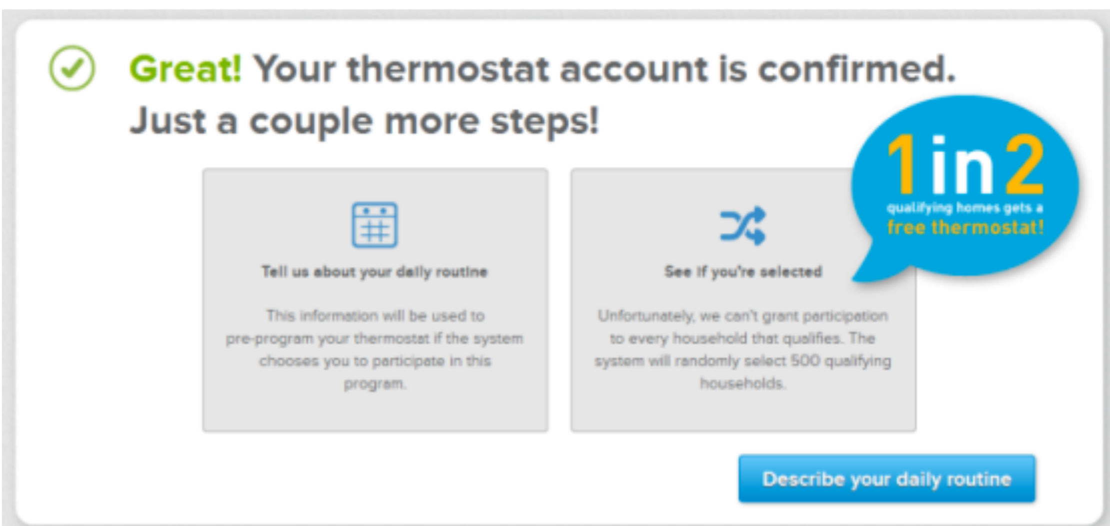

Qualifying Questions

The customer begins to program their thermostat by providing the following information:

- Whether multiple people live in their home. Opower tailors the language in the application to the number of people in the household.

- Whether they have pets. If the customer has pets, the default away temperature of the home is adjusted to a safe temperature for household pets. For homes with pets, the default away temperature is 82 instead of 85 for cooling and 60 instead of 55 for heating.

- Their mobile phone number. Customers are sent a text message to this number with a link to the Opower mobile application..

Setting an Initial Schedule

After completing the qualification questions, the customer is prompted to create a personalized schedule. By default, customers set a schedule for all weekdays and then Saturday and Sunday.

For all weekdays, Saturday, and Sunday, the customer has the following options:

- They can set a schedule for when they typically wake, leave the home, return home, and go to sleep.

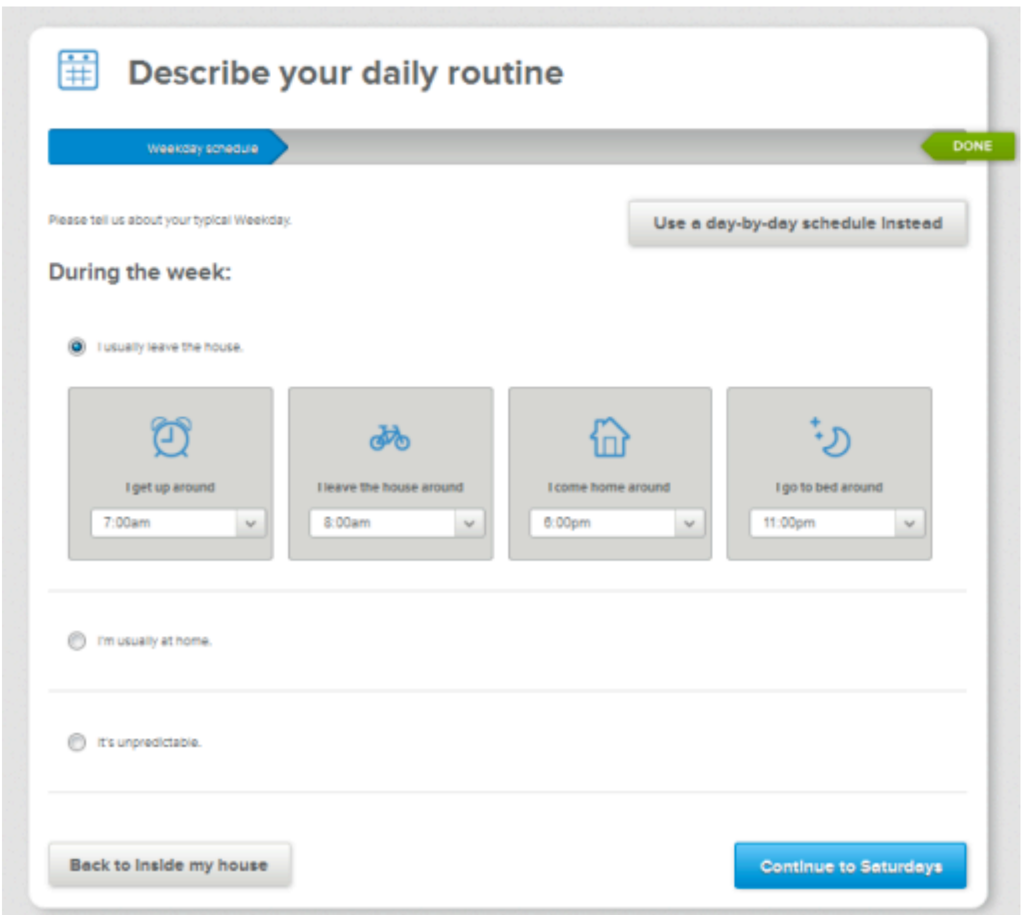

- They can indicate they are home all day and set the time for when they usually wake and go to sleep.

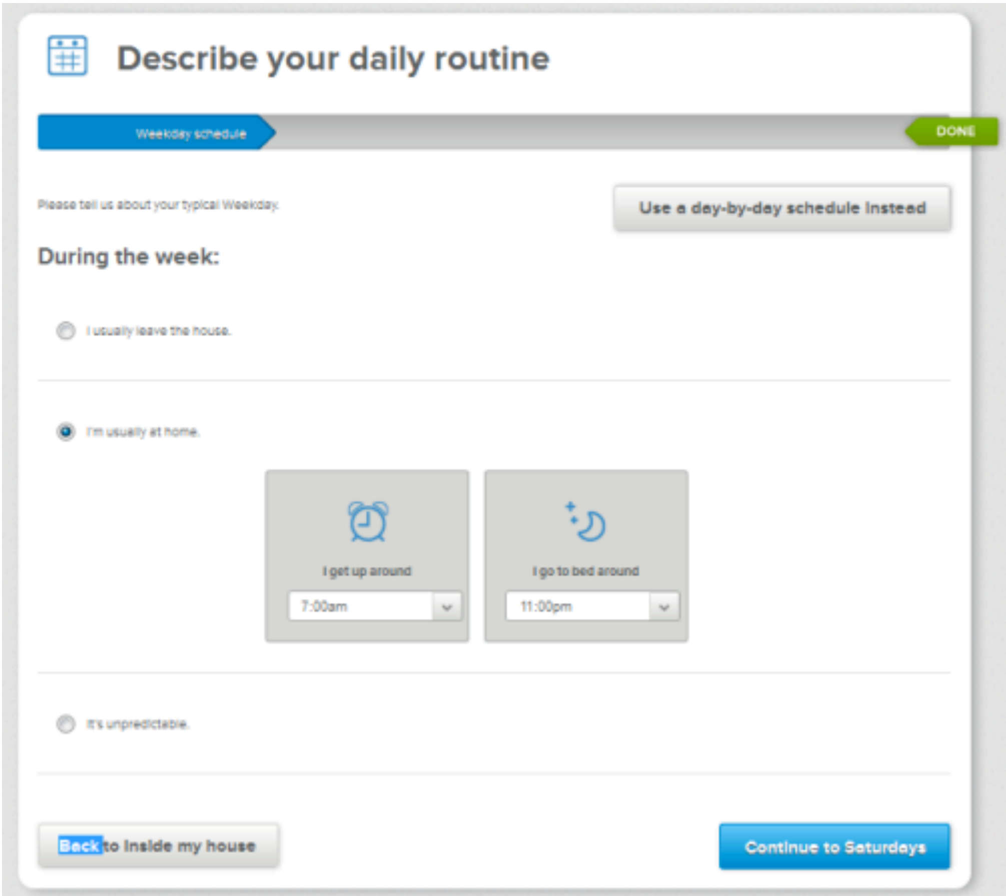

- They can indicate their schedule is unpredictable. In this case, they are still asked when they typically wake and go to sleep.

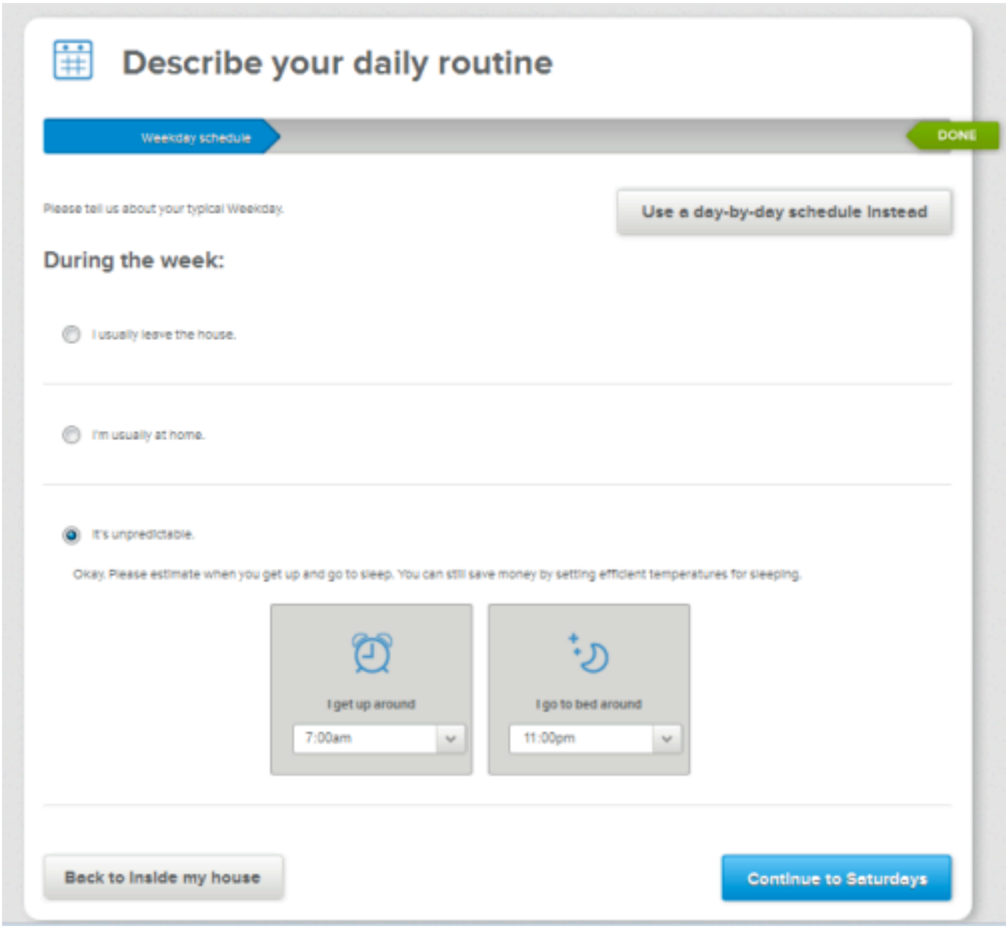

Instead of setting the same schedule for all weekdays, a customer can also create a day-by-day schedule for each weekday separately. The same schedule options are available on a daily basis.

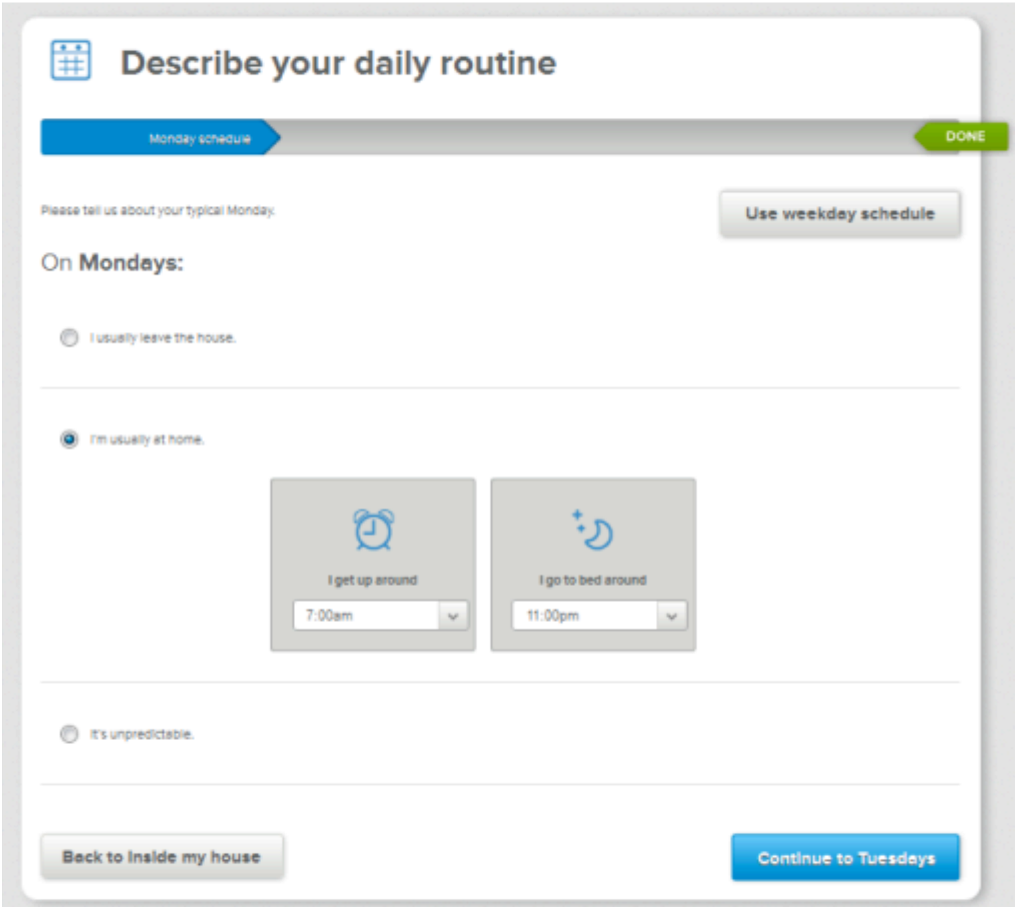

Setting Initial Temperatures

Customers are prompted to set their home and sleep temperatures for heating and cooling. The default temperatures for these settings are based on the suggested Energy Star settings (ENERGY STAR® Program Requirements for Residential Climate Controls, Version 1.0 Partner Commitments, DRAFT 2).

On the heating page, customers are asked how warm they would like their home to be when they are home and asleep.

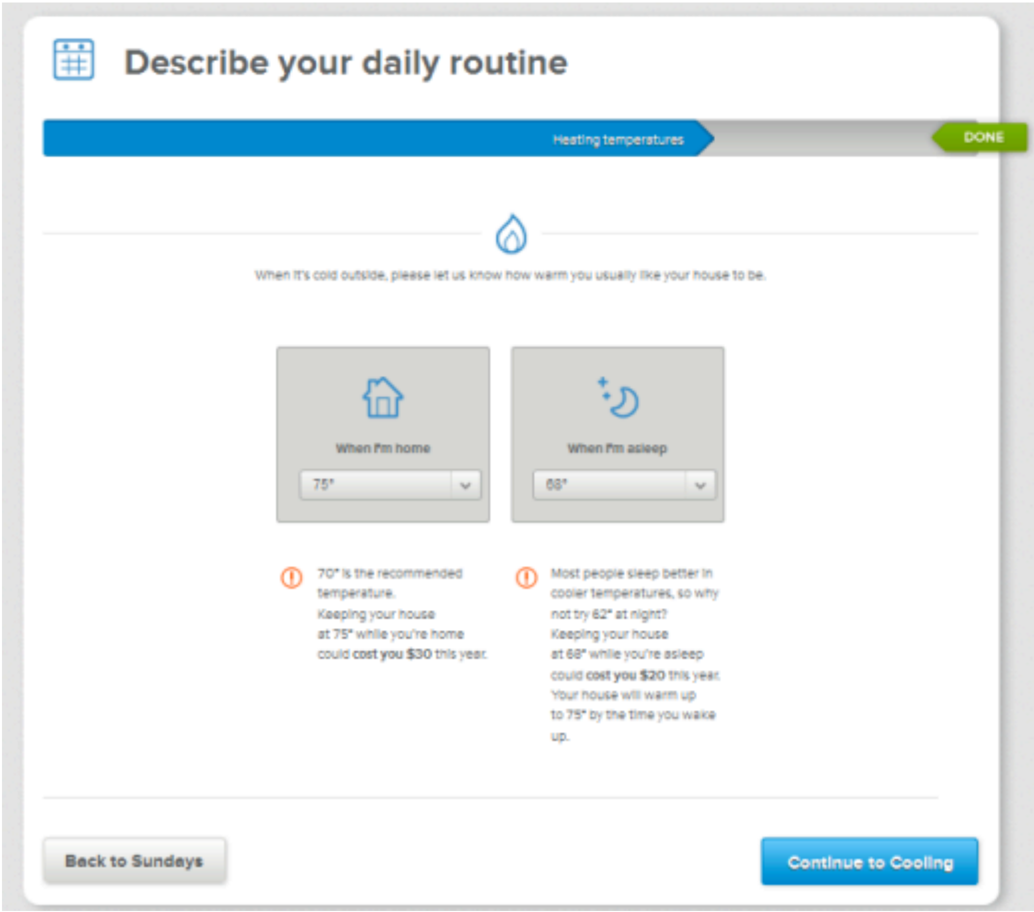

If the home temperature is greater than the recommended setting (less efficient), an insight appears to tell them how much money they will spend during the winter keeping the home at this higher temperature. If the away temperature is higher than the recommended setting, they are prompted to try setting the temperature lower since the house will warm up to a comfortable setting before they wake up.

On the cooling page, customers are asked how cool they would like their home to be when they are home and asleep.

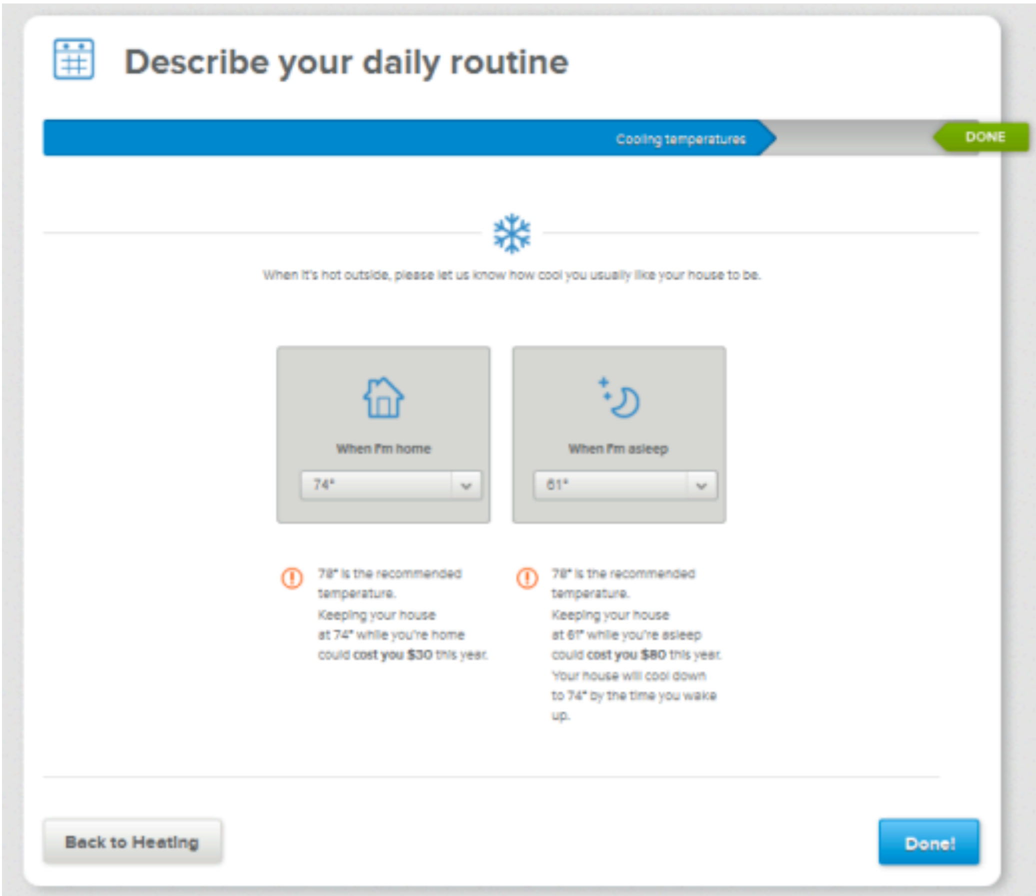

If the home temperature is less than the recommended setting (less efficient), an insight appears to tell them how much money they'll spend during the summer keeping the home at this lower temperature. If the away temperature is lower than the recommended setting, they are prompted to try setting the temperature higher since the house will cool down to a comfortable setting before they wake up.

Installation

After submitting their temperature settings, the customer is randomly selected to be part of the test or control group.

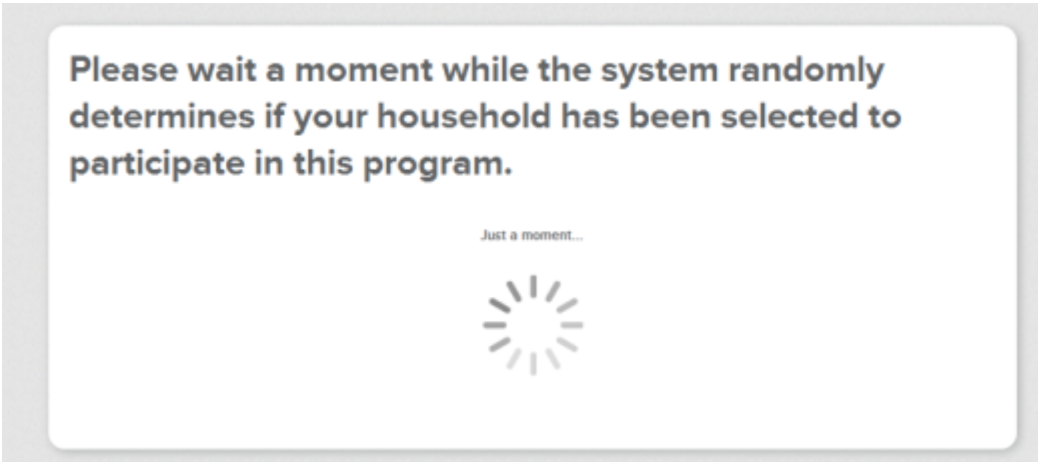

If they are part of the control group, they will not receive a thermostat. Customers in the control group may opt to sign up for a waiting list and may receive a thermostat if the program is expanded.

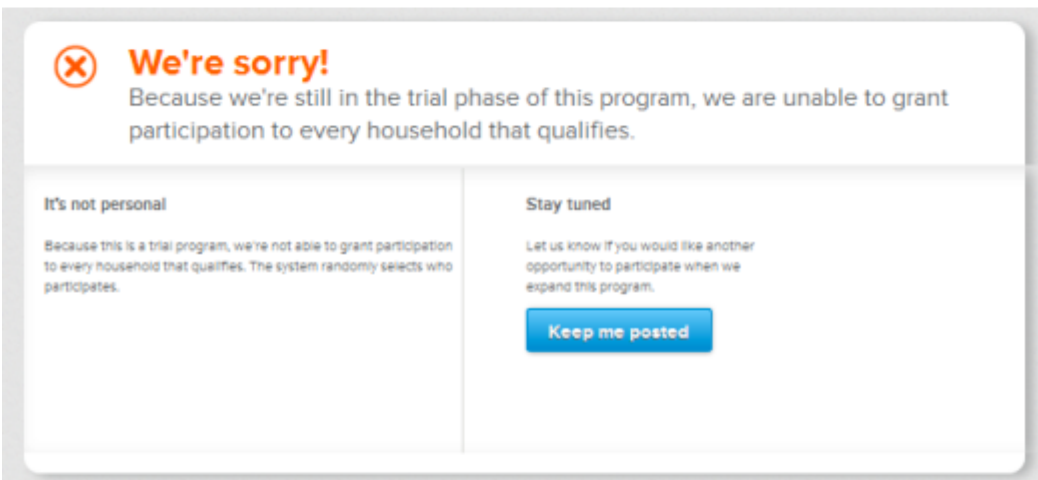

If they are randomly selected into the test group, they will receive a thermostat and become part of the program. Customers participating in the test group can schedule an appointment to have their thermostat installed.

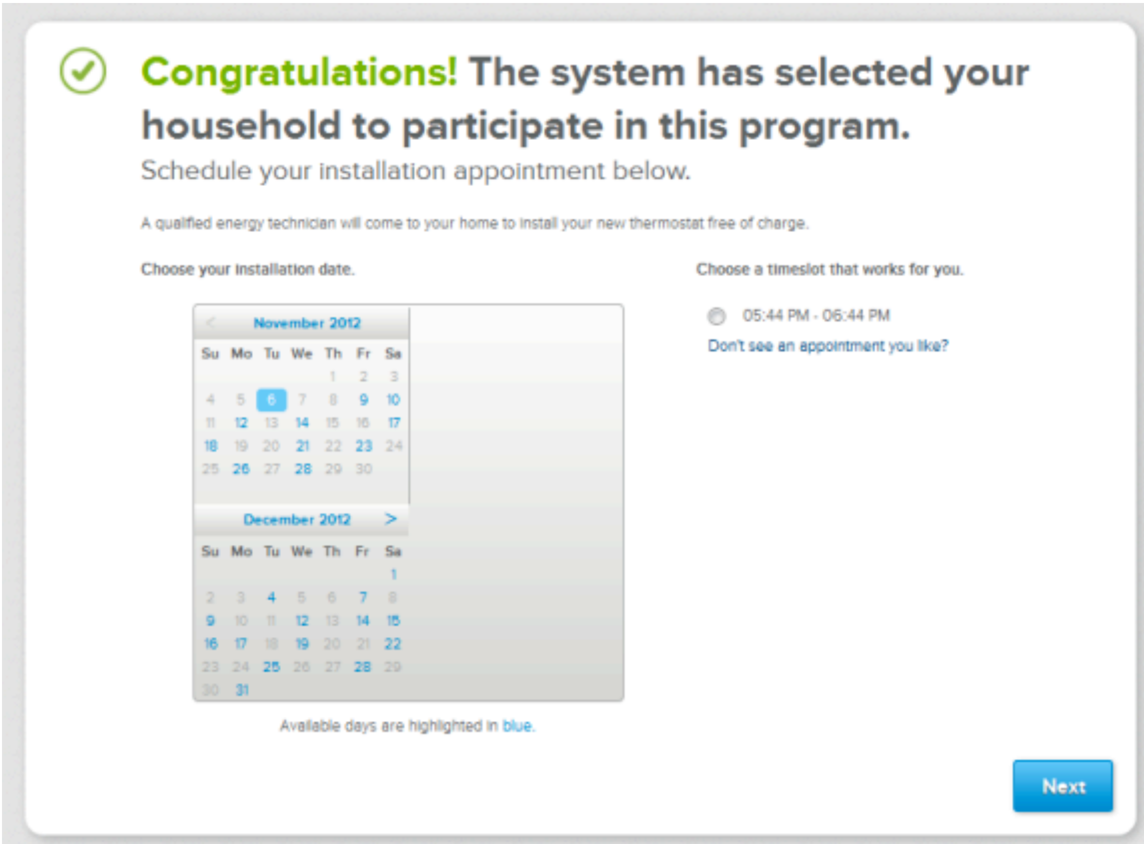

If none of the times available on the screen are convenient for the customer, they can click "Don't see an appointment you like?" to see a phone number they can call to schedule the appointment (1-888-660-5028).

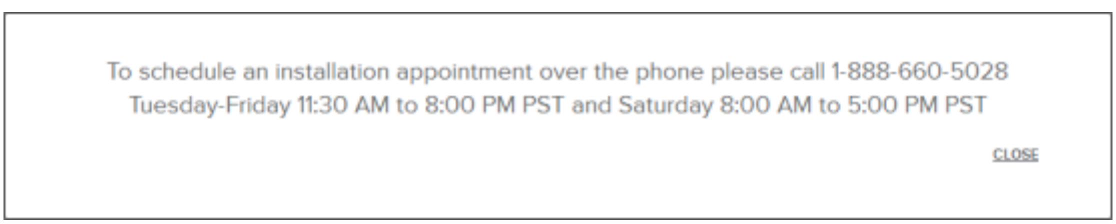

Once they have selected the date and time for their appointment, they will see a confirmation screen. This includes information on how to reschedule the appointment and where to download the mobile application.

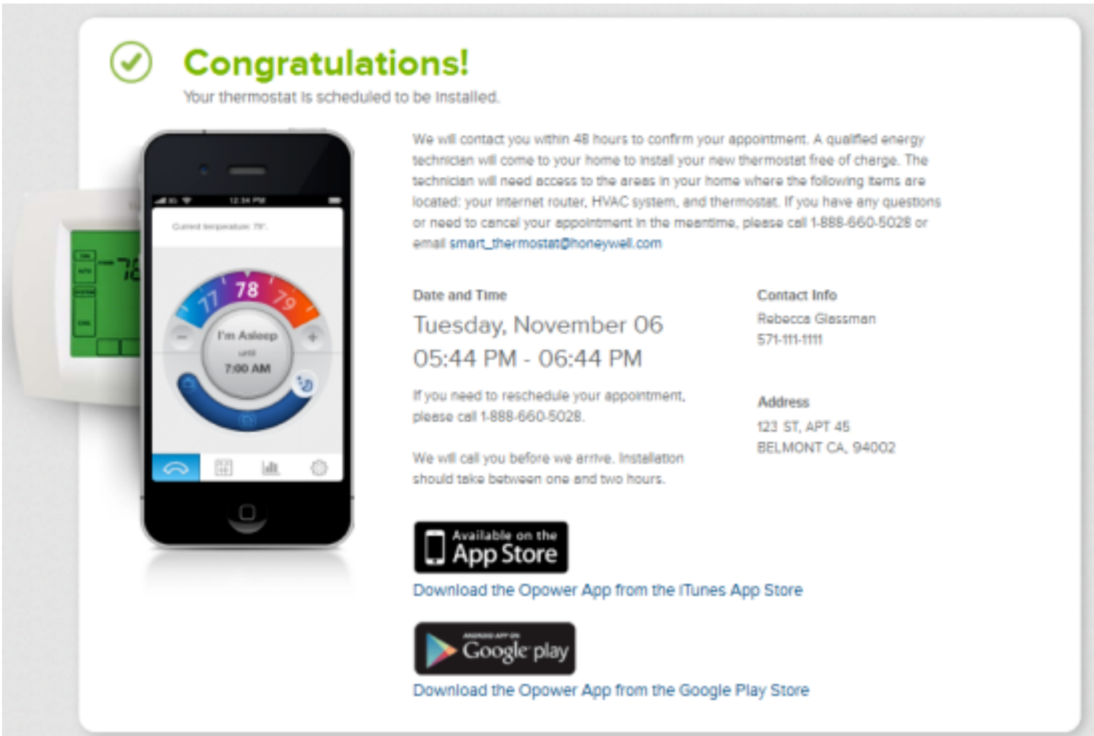

The customer will also receive an email confirmation for their appointment and a reminder to install the mobile application in advance of the appointment.

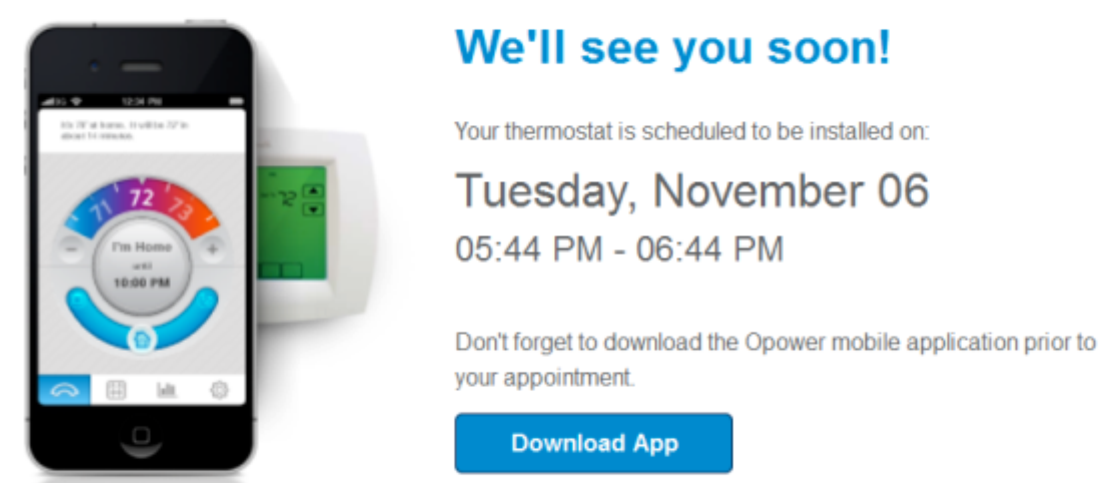

Mobile Application Tour

The mobile application tour can be launched at anytime, using the Opower mobile app on the iPod Touches, and later on the customer's smartphone. Click on the Settings tab, click "Launch tour," slide through the tour pages, and click "Done" to exit. The tour provides an overview of some of the main application functions and customer messages.

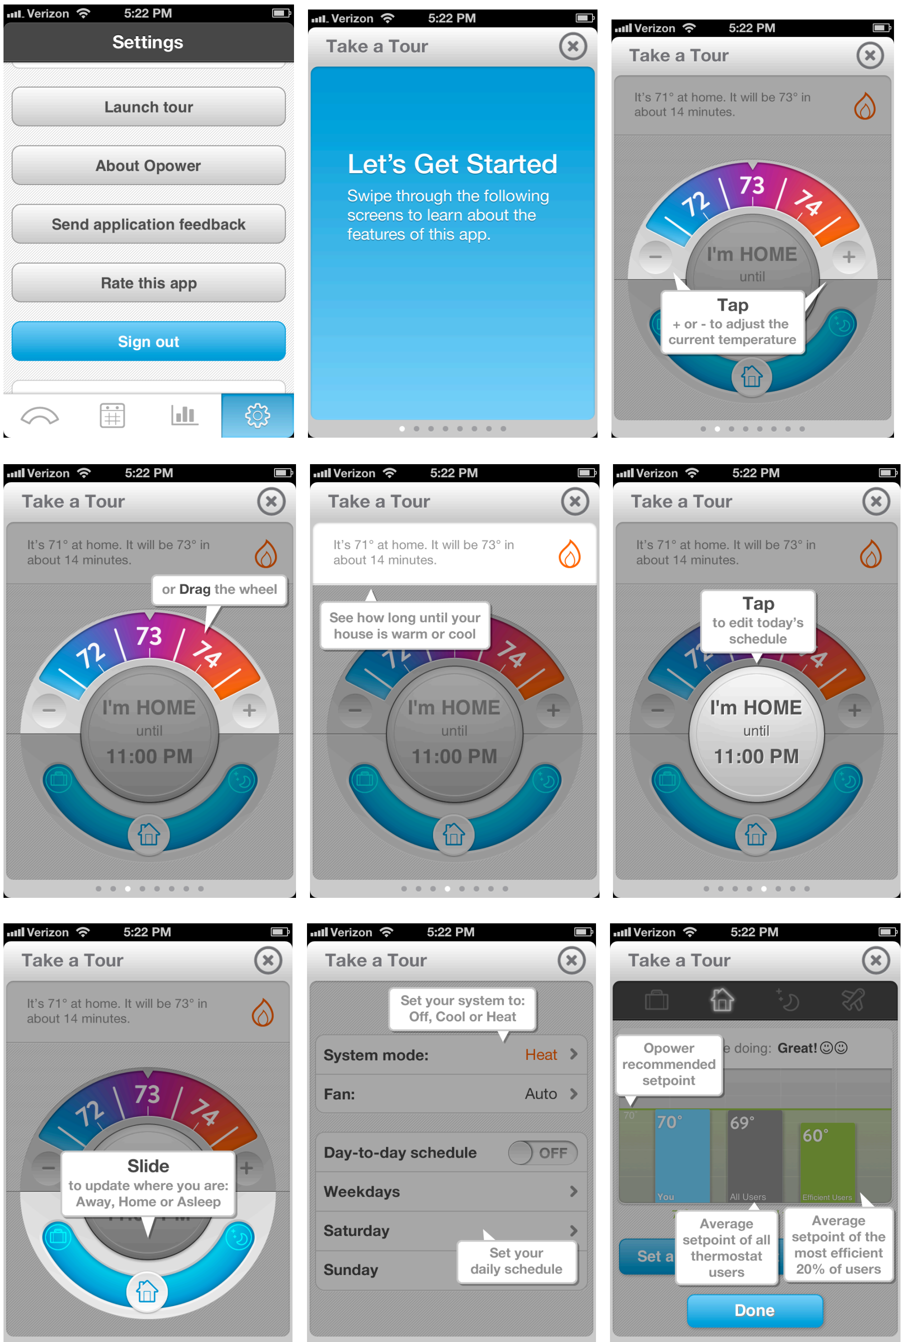

Answering Customer FAQs

This section will help you answer customer questions about the program, mobile and Web applications, and thermostat. A full set of customer FAQs can be found at https://thermostat.opower.com/faq.

What is this thermostat program?

Opower and Honeywell have partnered to create a smart thermostat solution, which allows utility customers to program and monitor heating and cooling energy usage, not just from the thermostat itself, but also via Internet-connected devices like smartphones. This solution also gives you the ability to create optimal thermostat schedules that fit your lifestyle and provides customized recommendations to help you trim your energy bills.

How can I save?

A programmable thermostat can help reduce your heating and cooling costs. You can save all year long if you ensure your thermostat is set at the optimum program settings that match your lifestyle. You can manipulate your temperature setting and conserve energy, even while you are away, through the use of the Internet or your smartphone. Setting your programmable thermostat to the highest comfortable temperature in the summer and lowest comfortable temperature in the winter can help you reduce your energy bill.

What are the estimated savings based on?

The estimated costs and savings calculations are based on average heating and air conditioning usage and utility billing rates in your area. These are only estimations and are not a guarantee of savings from your utility company.

What other benefits does this program provide?

This thermostat program also benefits the community by helping to educate customers about energy use and energy efficiency goals. The energy customers save will not only help the environment, but also help reduce the need for new power plants and the occurrence of power outages.

Are there any safety or privacy concerns I should be aware of related to this thermostat program?

The Honeywell VisionPro thermostat used for this program was rigorously tested prior to being installed in customers' homes. These devices go through numerous quality control checks by multiple parties, to ensure they meet a high level of customer safety, reliability, and satisfaction.

It is also our top priority to protect our customers' information. We apply the same privacy protection standards to all data collected by the company from customers. We treat each customer's personal information and data as confidential, consistent with all regulatory requirements, including those established by the Public Utilities Commission. Therefore, be assured that your information is kept private.

Can I get this device for my other properties and/or business?

The smart thermostat program is only available for residential use at this time. Only a single thermostat is available for each program participant.

How many devices can I access the applications from?

Only a single wall-mounted thermostat is available for each program participant. You can install and access the mobile application from as many smartphones as you would like, but the application must be registered with the same username and password. Similarly, you can use the

Web application from any supported web browser on any computer. If more than one member of your household uses the application at the same time, the changes are preserved for the last person who saves their changes.

Can people see if I am home or not?

No. We apply the same privacy protection to this data as other all other data collected by the company for customers. The only way someone can see your status and schedule is if you give them your login credentials to the web or smartphone application.

If I work from home or have a severe illness for which I have special temperature needs, can I still benefit from this program?

You will always have control of your thermostat, so you can set safe and comfortable temperatures that are suitable for your lifestyle. An easy way to save energy is to lower your heating temperatures and raise your cooling temperatures when you are away. Depending on your personal needs, you may also be able to use more efficient temperatures while you are asleep.

How safe is the program? Can anyone hack into the system?

It is our top priority to protect our customers' information. Our system employs industry-standard defense mechanisms against brute-force attacks, code injection, and other malicious activity. We apply the same privacy protection standards to all data collected by the company from customers. We treat each customer's personal information and data as confidential, consistent with all regulatory requirements, including those established by the Public Utilities Commission. Therefore, be assured that your information is kept private.

What smartphones support the mobile application?

The mobile application is currently supported on the Apple iPhone 3GS or later, running IOS 4.3 or later, and Android phones running 2.2 or above. To locate your operating system on your iPhone, open the *Settings* app, click on "About," and see what "Version" your iPhone is running (needs to be 4.3 or above). To locate your operating system on your Android, open the *Settings* app, click on "About phone," and see what "Android version" your phone is running (needs to be 2.2 or above).

How do I make a one-time change to my schedule?

You can use the "Thermostat" page of the mobile application or the " My Thermostat" page of the Web application to manually change your temperature, change your current state (away, home, asleep), or set a new time to come home, wake, or go to sleep. On the thermostat on the wall, you can also manually change your temperature.

How can I change my email address and/or password?

Open the Web application, and then select "My account" to change your password or email address.

I now have three ways to change my thermostat. How are they different?

You can use your thermostat to manually change temperatures, turn on and off your heating and AC, and control your fan. The Web application has the same functionality as the thermostat and also allows you to register for an account, set a vacation schedule, and change your account settings, primary schedule, default temperatures, state (home, away, asleep), and schedule for today. The mobile application has all of the functionality of the thermostat and Web application, plus it allows you to compare your temperature settings, set a passcode, and set and receive notifications.

Which browsers are supported for the Web application?

The current major release and previous major release of the four desktop browsers with the largest market share are supported. Currently, this means Internet Explorer, Safari, Mozilla Firefox, and Google Chrome are supported.

Will my house really be comfortable enough when I get home?

Yes. You just set the time you will return home and your thermostat does the rest. Your home will be heated or cooled for you before you return home after being away or on vacation. Your smart thermostat learns the amount of time it takes to heat or cool your house before you arrive, based on the actual temperature in your home and past usage.

Can I enroll in the program using my smartphone?

You can only enroll in the program using the Web application. If you are selected for the program, you will receive information about how to install the mobile application.

