WHITE PAPER

Understanding Energy Efficiency Benefits from Smart Thermostats in Southern California

PREPARED FOR ENERGYHUB BY BEN HO, PHD – MAY 2014

CUSTOMERS WHO INSTALLED A SMART THERMOSTAT CONNECTED TO WEB AND MOBILE SOFWARE SAW A SIGNICIANT REDUCTION IN WHOLE-HOME ENERGY USE.

Key Findings

- In a study of 89 California households that adopted a smart thermostat between 2012 and 2013, we find that smart thermostats reduce overall electricity bills by 6% in summer months, for an average savings of $15.60 per month. Savings were largest in August, where estimated savings are as high as 17%.
- In particular, we find that savings came from lower usage in afternoons (smart thermostat customers increased usage somewhat between 1 a.m. and 5 a.m.) and early in the week (Sunday through Tuesday).
- Our model controls for outdoor weather changes, household characteristics, and seasonal effects.
- We find a learning effect as well. Users initially see an increase in their electricity usage as they use the new device to make regular manual adjustments to the settings. However, once consumers adapt the thermostat's programming to their particular usage needs, we see lasting and substantial savings. Those who make the most use of the program settings of the thermostats saw an additional 10% reduction in their electricity usage.
- Nearly all the savings come from the high energy use households. Those with below average energy usage saw statistically insignificant savings. The largest savings observed were for households in the highest quartile of energy use who saw their energy use decline by approximately 8%.

Executive Summary

> **The Southern California households included in our study saw a 6% reduction in whole-home electricity use.**

Programmable thermostats have repeatedly failed to live up to their energy efficiency potential, resulting in the end of many utility thermostat rebate programs and the suspension of the ENERGY STAR programmable thermostat certification program. On the other hand, our recent study of southern California households showed a 6% reduction in whole-home electricity use for customers that installed a smart thermostat connected to web and mobile software. The study results provide not only a framework for determining the sources of energy efficiency but also direction for helping more smart thermostat owners achieve positive energy efficiency results.

EnergyHub set out to quantify the energy efficiency benefits of residential internet-connected "smart" thermostats distributed outside of utility programs. We also wanted to determine the sources of energy efficiency and what separates customers that achieved efficiency gains from those that did not. Finally, we were motivated to evaluate the impact of smart thermostat installations on residential peak load.

From our study comparing whole home summer electricity usage before and after residential customers installed a smart thermostat, we learned that the primary determinant in the resulting change in electricity consumption was how much customers "fiddled" with their temperature settings and that the bulk of energy savings came from high use households. These results yield new insight into behavioral techniques that could encourage more efficient usage patterns.

Using a model that controlled for outdoor weather changes, household characteristics, and seasonal effects, we also found a significant learning effect. Customers that installed smart thermostats initially saw an increase in their electricity usage as they used the new device to make regular manual adjustments to the settings. However, once the adapted the thermostat's programming to their particular usage needs, we saw lasting and substantial savings. Those who made the most use of the programming settings of the thermostats saw the largest reductions in their electricity consumption. From an implementation perspective, these results provide additional support for utilities and regulators employing or planning to employ rebates and subsidies for customer-selected smart thermostats deployed through a variety of channels. From a research perspective, this study provides a novel framework for determining the sources of energy efficiency from programmable thermostats and insights on where further efficiency opportunities exist.

Introductions

Smart thermostats that give users programmatic control over their home heating and cooling, as well as remote access and easy programming via web and mobile interfaces, promise to be a potent tool in both saving households money and also conserving energy. Heating and cooling alone account for about half of a typical American household's annual energy bill (of approximately $2,000 per year) and much of it is potentially wasted on hours when families are either not at home, or when thermostat temperatures are set unnecessarily high or low.

Here, we analyze the impact of one such smart thermostat implementation using utility-provided interval electricity meter data to see what kind of impact these thermostats have on energy consumption on an hour-by-hour basis. This study quantifies the amount of energy savings, the patterns of energy savings by time of day and month of year, the largest beneficiaries of energy savings, and the usage patterns that correlate with energy savings.

We use data from the Summer of 2012 and the Summer of 2013 from 89 households that installed a Wi-Fi smart thermostat connected to web and mobile applications provided by EnergyHub during this time period, and we build a regression model that controls for local weather changes, household characteristics, and seasonal energy use variation. To assess whether the savings can actually be attributed to the thermostat itself, as opposed to other behaviors that may be correlated with the thermostat purchase, we estimate how much of the energy savings can be attributed to the actual thermostat settings, and we find evidence that most of the savings can indeed be attributed to the thermostat settings.

We find that smart thermostats reduce overall electricity consumption by 6% in Summer months, for an average bill savings of $15.60 per month. Savings were largest in August, where estimated savings are as high as 17%. In particular, we find that savings came from lower usage in afternoons (smart thermostat customers increased usage somewhat between 1am to 5am) and early in the week (Sunday through Tuesday). We find a learning effect as well. Users initially see an increase in their consumption as they use the new device to make regular manual adjustments to the settings. However, once consumers adapt the thermostats (e.g., programming to their particular usage needs), persistent and substantial savings follow. Those who make the most use of the program settings of the thermostats see 30% reductions in electricity bill. Nearly all the savings come from high electricity use households. Those with below average electricity usage saw statistically insignificant savings. With the largest savings observed for households in the highest quartile of electricity use who saw their usage decline by approximately 8%.

More broadly, smart thermostats represent a first step in the forthcoming "Internet of Things." By understanding how these devices can be used to save money and conserve energy as well as reshape the environment we live in terms of ambient household temperature, this study gives us a glimpse of things to come.

Data

The data consists of hourly meter readings, thermostat data, and outdoor weather conditions for 89 California households for the months of June until September for both 2012 and 2013. The average price of electricity in California at this time was approximately 17 cents per kWh, but because of a tiered pricing plan, households in our sample faced a marginal price of electricity of closer to 30 cents per kWh. There were 12 utility-initiated demand response days during the period of study, on which customers received a monetary incentive for reducing their electricity usage, but excluding those days from the dataset had no effect on model estimates.

Summary Statistics

Variables	Mean	Standard Deviation
Hourly Meter Reading (Wh)	1,292	1,398
Target Temperature (degrees F)	75.2	12.0
Outdoor Temperature (degrees F)	75.6	11.2

The smart thermostats in question are Wi-Fi enabled programmable thermostats, capable of either four or seven unique temperature set points per day. The thermostat can be easily programmed via its companion web and mobile applications, which can also be used to make remote adjustments to the thermostat settings when the user is not at home. These thermostats report a significant amount of data related to their operation to their remote management platform (approximately 50,000 data points per thermostat per month), enabling much of the savings decomposition described below.

Main Findings

Figure 1 shows both the actual average hourly electricity meter readings for households in our sample between 2012 and 2013 as well as the regression models of estimated electricity usage for households before they purchased a thermostat compared to households who did not.

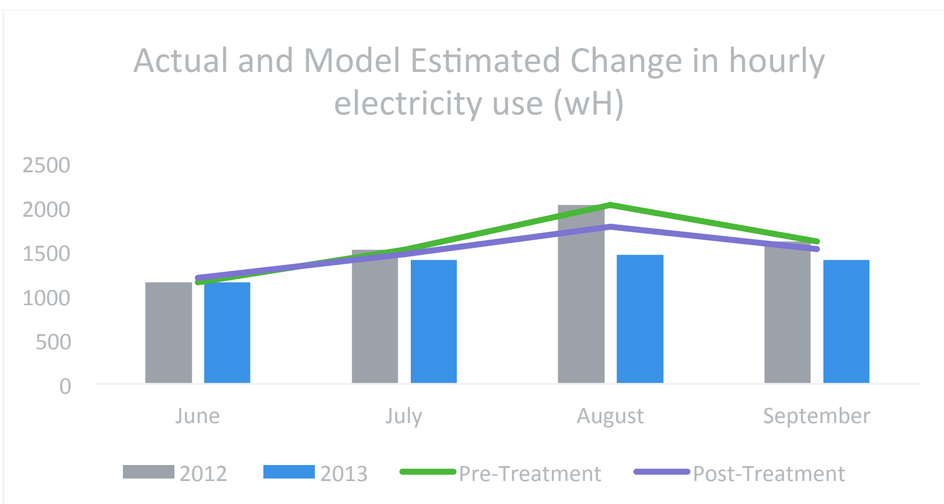

Figure 1 Actual average hourly meter readings per household (Bars) and the model estimated electricity usage by household (lines), holding temperature fixed at 2012 levels.

At 30 cents per kWh, the 6% reduction in energy use amounts to savings of $15.60 per month.

While the average household in the sample experiences clear savings between 2012 and 2013, one might worry that these savings were due to cooler temperatures in 2013. Therefore, we develop a model of household energy consumption by conducting a differences-in-differences estimate using an ordinary least squared (OLS) regression using time and household fixed effects, to calculate the treatment effect of installing a smart thermostat on electricity usage. Regressions were clustered by household, and all reported differences are significant at the 95% level (see Appendix).

Essentially the model compares households who had the thermostat on a given day in 2013 with the same household on a similar day in 2012 (for example a typical Tuesday in August 2013 with a typical Tuesday in August 2012) when it didn't have the thermostat. While the raw difference in electricity usage per hour between 2013 and 2012 was 200 Wh per hour, the model estimates that the installation of a smart thermostat accounted for 72.5 Wh of the difference. By comparison, for each degree Fahrenheit decrease in the outdoor temperature, electricity usage decreased by 42 Wh per hour.

Given that the average hourly electricity consumption level was 1,292 Wh, the estimated reduction in electricity use after the thermostat was installed was 6%. This translates to 52 kWh per month. At 30 cents per kWh (a typical marginal price of electricity for households from this part of California in August 2013), that corresponds to savings of $15.60 per month.

Locus of Savings

In this section, we use the same specification that controlled for time, household, and outdoor temperature, but we break down the savings by time of day, day of week, and month of year to see when and where savings were concentrated.

Month to Month Comparisons in Wh/hour (negative numbers indicate savings)

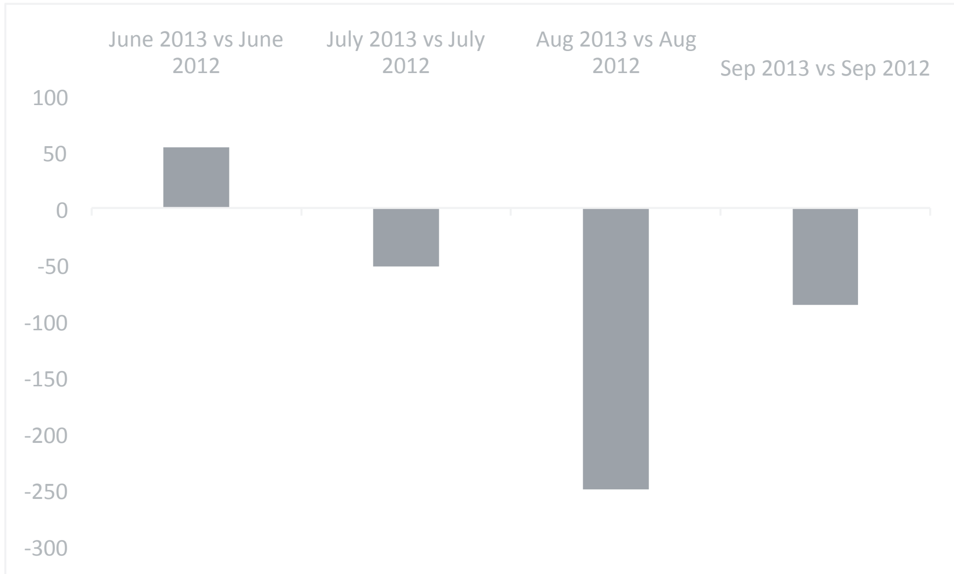

Not surprisingly, the savings were concentrated in August, which had the hottest temperatures in our data. Average hourly electricity usage in August for both years was 1,475 Wh, so 250 Wh represents a 17% reduction in households' monthly electricity consumption.

Time of Day Comparisons in Wh/hour (Negative numbers indicate savings)

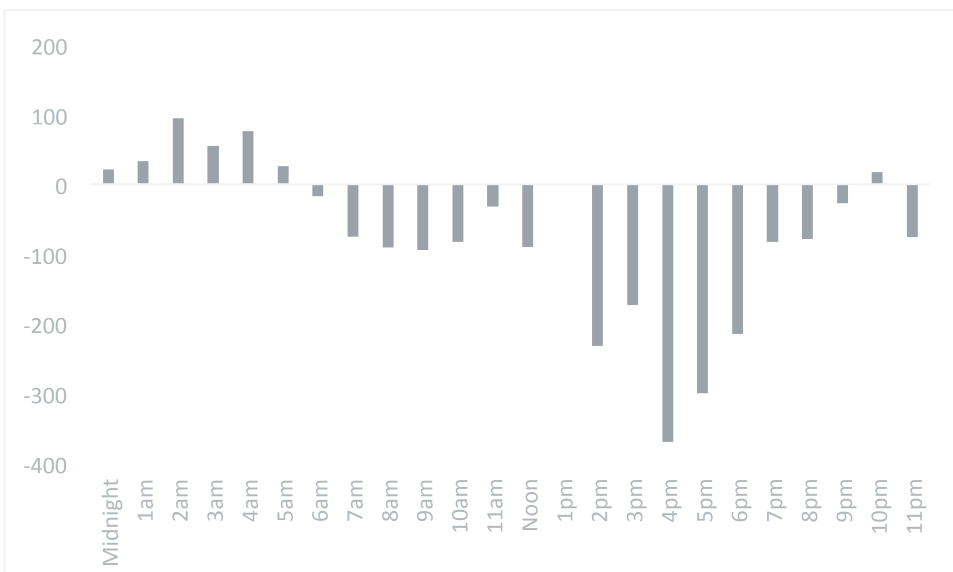

Not surprisingly, the savings were concentrated in the afternoon when electricity usage is the highest but people are not yet at home. A smart thermostat is ideal at this time, as it can turn off heating and cooling when people are not at home. Interestingly, electricity usage went up late at night, perhaps because users became more comfortable leaving the air conditioner on overnight knowing that they could schedule it to turn it off automatically.

Day of Week Comparisons in Wh/hour (negative numbers indicate savings)

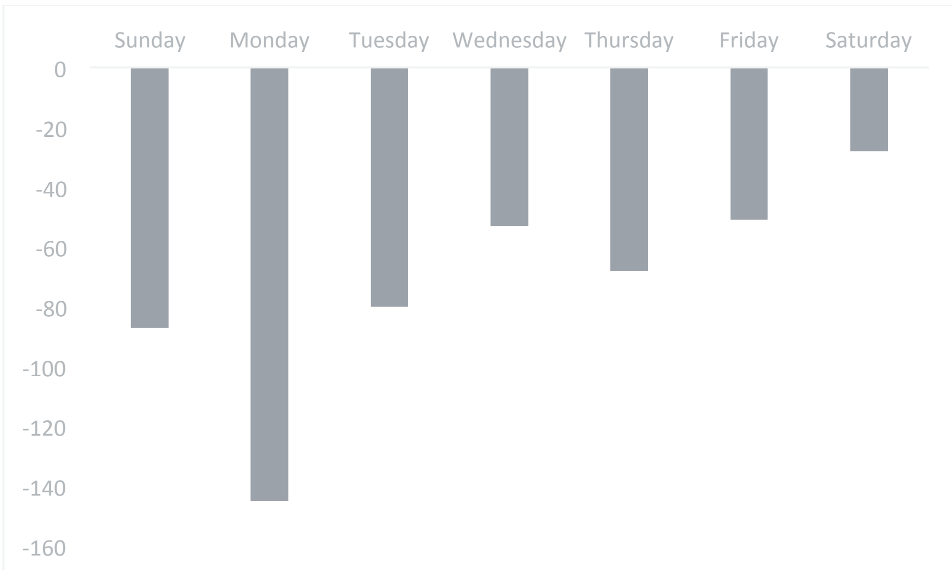

Here, we see the lowest savings occur on Saturdays, which makes sense because many people keep their homes at the “comfort” set point for the entire day. The greatest advantage of the thermostats comes when people are away at work. The highest savings on Monday may be explained by the fact that people, who once forgot to turn off their thermostats on Monday for the workweek, now used the device to automatically do so.

Savings by household size in Wh per hour (negative numbers indicate savings)

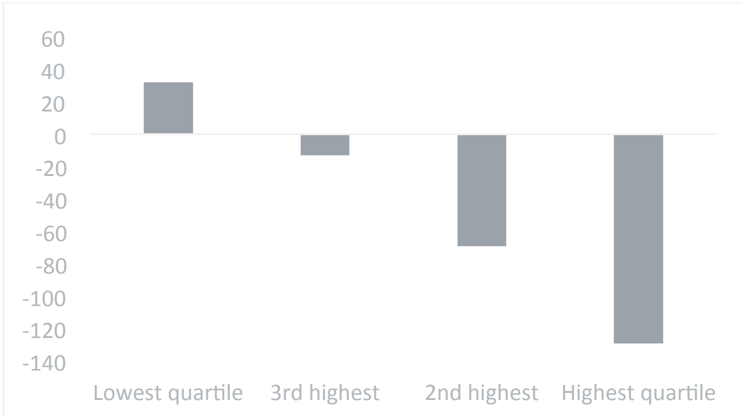

The savings seem primarily located amongst households in the highest energy usage quartile. However, we should note that given the relatively small number of households in each quartile, these estimates are only marginally significant.

We also compared the savings of the 11% of households in the sample who have on-site generation installed (primarily in the form of rooftop solar panels) to households who do not. Households that had on-site generation saw a savings of 380 Wh compared to the 30 Wh of savings for other households. We also compared the estimates of models that accounted for the net savings of on-site generation households during warmer weather, but we found statistically similar estimates.

Robustness of Estimates

The estimates above include controls for

- Prevailing weather
- Month of year, time of day, day of week
- Date of thermostat adoption
- Household characteristics

The estimates are also robust to a number of specifications that allow outdoor temperature to affect electricity usage non-linearly, i.e., differently from month to month or year to year.

However, there are two important caveats to this analysis that we will explore in the following section.

The first is the problem of omitted variable bias. All we have shown is that households that have adopted the smart thermostat have lower energy use (controlling for weather, time of day, and household characteristics). However, our results would also be consistent with the idea that some other behavior that is temporally correlated with the purchase of a smart thermostat—e.g. the adoption of energy efficient light bulbs or appliances—may instead be driving these results. Without having measures of these other behaviors as controls, we can only speculate.

The second is that all the savings we observe are year to year. We do not observe within year savings, i.e., if a thermostat was installed in July, we do not observe a reduction from June to August (analysis was done primarily using a differences in differences design with a polynomial flexible time trend). In fact, if anything, energy usage seems to go up in the short run. One possible explanation is simply that we have relatively few households that installed the thermostat in the summer months for which we have data; most thermostats were installed after September 2012 and before May 2013. However, the difference in short run versus long run still bears further examination.

We resolve both of these caveats in the following section where we decompose the savings using the household's thermostat usage patterns, which we observe once the smart thermostats have been installed. This exercise allows us to say that not only are the energy savings correlated with the purchase of the thermostat, they are also correlated with how the thermostat is used. This limits the number of other confounds that could be biasing our results.

The decomposition exercise also was informative about which patterns of smart thermostat use yield the most savings, and why electricity use seems to go up in the short run after thermostat purchase but go down over the course of the first year.

Decomposing the Savings

To decompose the calculated savings, we took our preferred regression model, and constructed interaction term between the treatment variable, and variables that measured different usage patterns for the smart thermostats. The usage patterns we identify are based on the temperature set point setting on the thermostat, combined with data about the external temperature.

Specifically

- The mean level of the set point.
- The difference between the level of the set point and the outdoor temperature.
- The standard deviation of the set point (a measure of the number of adjustments made to the set point both automatic and user controlled).
- Set point unpredictability (changes that can't be predicted by time of day or day of week).

Using the same regression analysis as before (see Appendix), we decompose the overall savings (of approximately 72 Wh/hour) to see how much of the savings can be explained by our measures of thermostat usage.

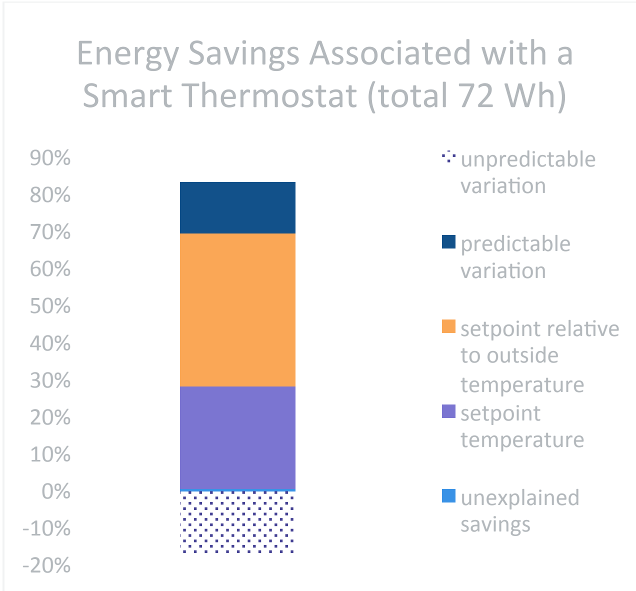

Figure 2 We Decompose the total energy savings (85 Wh per hour) associated with owning a smart thermostat by how the thermostat was used

The colored bars of the decomposition show how the total energy savings can be explained by the four usage patterns. The dotted area shows the effect of unpredictable variation in usage, which actually yields "negative savings," that is periods of time with high unpredictable variation increases electricity usage relative to when the household did not have a smart thermostat. Looking at the decomposition we can conclude a few things:

- It is reasonable to assume that most of the energy savings correlated with installation of the thermostat was due to their thermostat usage. The choice of temperature set point explains nearly all of the savings; households that use the thermostat to choose higher average set points are the ones who experience the most energy savings.
- Unpredictable variation of set point settings leads to increased energy use: households that had more unpredictable adjustments (as owners fiddled with the settings) saw their energy use increase.
- Households that had more predictable variation (following a program according to hour of day) saw their energy use decrease.

This decomposition tells us that most of the savings are correlated not just with ownership of a thermostat but also in how that thermostat was used. It also explains the difference between the short run effect and the long run.

When the device is first purchased, the amount of unexplained variation is high. People are fiddling with their device, learning their own preferences as well its optimal use. During this period, electricity usage in that household goes up. However, over time, once users become accustomed to the device, set points start to follow a predictable pattern. Households who take the most advantage of the automated programming features also benefit with the greatest savings.

Using estimates from this treatment effect decomposition model, we can simulate the estimated savings for a household in the top 10 percentile of predictable variation. In other words, we can simulate a typical household, typical in every way except that their predictable set point variation is greater than 90% of the households in our sample. Here, the model predicts approximately an additional 100 Wh per hour of savings, or 10% of typical hourly usage (see Appendix for details).

This study takes a look at the impact of residential smart thermostat installation using rich high frequency data to explore the potential energy and cost savings these devices can offer. Even in jut the first year of usage, $16 per month is substantial, with added benefits amortizing over time, both in terms of cost savings and in reduced environmental impact. As with any retrospective study, there is the potential issue that we have identified a correlation rather than a causation. However, the richness of our dataset allows us to say that if there was some other factor responsible for these energy savings besides the thermostat, it would have to have been not only temporally correlated (as in started at the same time) with when the thermostat was installed but also correlated with how its owners were setting the thermostat set points. As more data becomes available, our analysis points the way for future research that can refine these estimates and better disentangle the mechanisms that make smart thermostats so effective.

Appendix

	(1)	(2)	(3)	(4)
Energy Savings	-72.49***	2.526	2.013	5.257
	(4.114)	(6.261)	(6.270)	(6.270)
… explained by set point		-1.470***	-1.372***	-1.470***
		(0.0867)	(0.106)	(0.106)
… explained by temp difference		-6.732***	-6.841***	-7.044***
		(0.207)	(0.218)	(0.219)
… explained by overall set point variation			-1.258	-7.447***
			(0.787)	(0.970)
… explained by hourly set point variation				8.156***
				(0.745)
Outdoor temperature	41.80***	38.15***	38.08***	37.97***
	(0.281)	(0.305)	(0.308)	(0.308)
month controls	yes	yes	yes	yes
day of week controls	yes	yes	yes	yes
time of day controls	yes	yes	yes	yes
household controls	yes	yes	yes	yes
Observations	416,154	416,154	416,154	416,154
R-squared	0.362	0.364	0.365	0.365

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

Above is the regression model used for this data. Each column represents a separate differences-in-differences estimate using an ordinary least square regression of meter reading on whether the household had a thermostat at the time (the first row), the impact of the thermostat based on how the set points were set (rows 2 through 6), controlling for outdoor temperature, month controls, day of week controls, time of day controls, household fixed effects.

Other specifications tried included clustering standard errors at the household level, where our results were still significant for the main specification (column 1). We also tried controlling for time trends using a flexible polynomial time trend variable, but we found that such specifications yielded inconsistent estimates (see discussion the main text).

Millions of connected devices are already changing the way your customers manage their energy. Power that relationship with EnergyHub.

Contact EnergyHub today for a free evaluation:

info@energyhub.com
T: +1 718 522-7051
Or learn more at:
energyhub.com

232 3rd Street Suite C201 • Brooklyn, NY 11215