Extending the Impact of Smart Thermostats: measured impacts from automated set point schedule tune-ups

Michael Blasnik, Nest Labs

ABSTRACT

Internet-connected thermostats are becoming recognized as a potentially significant technology for promoting energy efficiency and behavior change. The technology can provide energy savings in several ways -- by making it easier for consumers to create efficient heating and cooling setpoints, by employing algorithms to optimize control of HVAC equipment, and by providing feedback to consumers about their energy usage and settings. One feature enabled by connected thermostats is to offer customers a "tune-up" of their scheduled set points to more efficient values. This paper describes a version of this feature that, on an opt in basis, makes small daily adjustments to customer schedule set points over a short time and then returns to its normal mode with the new schedule.

This paper provides measured results from multiple large-scale deployments of this tune-up algorithm across hundreds of thousands of homes in multiple climates. Key findings include: customer opt-in rates of more than 70%; scheduled thermostat set point shifts of about 0.5°F in the summer and 0.7°F in the winter; and HVAC runtime reductions of about 3%-7% on average with results varying by climate. Early indications about persistence suggest that changes to the schedules degrade modestly (~25%) into the second year but that many customers re-enroll and achieve further savings. A new variation designed to reduce summer peak demand has been tested and delivered estimated peak load shifts of 0.1 – 0.2 kW per thermostat. Many of the deployments included true experimental designs with randomized control groups.

Introduction

"Smart" internet-connected thermostats, on the market for fewer than 10 years, have become recognized as a potentially significant technology for achieving residential energy savings. Smart thermostats may provide energy savings in a variety of ways:

- by making it easier for consumers to use efficient heating and cooling set points via schedules, learning algorithms, occupancy detection, and other means;
- by employing algorithms to optimize control of HVAC equipment such as managing and minimizing the use of auxiliary heat on heat pumps
- by providing feedback to consumers about their energy usage and settings to encourage greater efficiency

Several independent studies and pilot projects have found significant energy savings – on the order of 10% of HVAC usage -- from smart thermostat installations (Aarish and Jones, 2016; Apex Analytics LLC 2014; Parker, Sutherland and Chasar 2016) but other studies have found a wider range of savings – especially across brands/models (for example, Apex Analytics LLC 2016). The number of studies that found significant energy savings prompted the US EPA to create an Energy Star standard for certifying smart thermostats (EPA 2016). The standard is unique in that it includes criteria based on a statistical analysis of data collected by the devices in a random sampling of customers.

In addition to directly providing energy savings from installation, smart thermostat may also be able to enhance energy efficiency and the utility grid in other ways, such as:

- using the data collected by the thermostat to diagnose HVAC system faults or efficiency problems
- using thermostat data to identify homes that need efficiency improvements and/or rate home performance
- tracking the impact of efficiency improvements in homes
- providing demand response and other load shifting that can benefit the grid and also optimize usage patterns for time-of-use rates

The potential for direct and indirect energy savings and utility system benefits from smart thermostats has led to widespread interest in the efficiency field (Energetics and VEIC 2016; Robinson et al. 2016). Many utility companies are currently offering rebates and incentives for their customers to purchase smart thermostats as part of their efficiency and demand response programs. As of late 2017, more than half of US consumers were eligible to receive a rebate for a smart thermostat (Nest 2017).

Thermostat Optimization / Schedule Tune-ups

One feature enabled by smart thermostats is to "tune-up" thermostat settings to more efficient values for customers. In 2013 Nest Labs created a feature called Seasonal Savings (Nest 2018). Seasonal Savings is designed to help customers fine tune their heating or cooling set point schedules to save energy while maintaining comfort. Customers are offered the service, on an opt-in basis, directly on the thermostat and through the app. Enrolled customers see small adjustments (<0.2°F) made to their scheduled set points each day over a few weeks – nudging them to greater efficiency. The algorithm focuses on different set points each week and assesses the customers' response to these changes to create a more efficient schedule going forward.

The Seasonal Savings algorithm is only deployed when an energy partner (typically a utility company) arranges with Nest to provide it to their customers. The target group of potential participants is commonly identified from the existing population of Nest thermostats located in the partner's service territory.

In summer 2017, Nest created a modified version of Seasonal Savings for a few deployments, tentatively named Peak Aware Seasonal Savings. As the name implies, this algorithm is designed to make schedule adjustments that also reduce utility peak demand while providing energy savings to customers. The algorithm was customized for the peak periods defined by each of four energy partners where it was deployed.

Participation

The millions of existing Nest customers have created the opportunity for individual partner deployments of tens of thousands of customers and in some cases more than 100,000 customers. The easy enrollment process – requiring just a single click on the device or app results in high participation rates – much larger than the typical sub-10% opt-in rates for standard enrollment-based programs that require great customer effort (see Henschel and Corsetti 2014).

When Seasonal Savings is deployed to a thermostat it must meet certain criteria to qualify. It checks to see that a thermostat is currently on-line, controlling an appropriate HVAC system (cooling in summer or heating in winter), in the proper operating mode (heating or cooling), and has a schedule with potential for efficient adjustments. The most common disqualification reasons are systems that aren't in the proper mode (more common in mild climates) or are off line. If a system isn't qualified, the algorithm waits and keeps checking for the next few weeks to see if it becomes qualified.

Once a system is qualified, the customer is offered the opportunity to enroll directly on the thermostat and in the phone app. Enrollment requires just a single click. While the algorithm is running over the next few weeks the customer can always simply adjust the thermostat as they normally do any time they choose and they can also opt out of the process early.

Table 1 shows participation data for Seasonal Savings deployments for the summer of 2017 and the winter of 2017/18. A few deployments are shown separately to illustrate the typical range of results with abbreviated names indicating geography (e.g., MW=Midwest).

Table 1. Seasonal Savings Participation: Summer 2017

	Targeted	Qualified		Opted In		
Utility /Deployment	Thermostats	Thermostats	% of target	Thermostats	% qual	% target
Summer 2017						
MW-1	86,447	71,910	83%	53,345	74%	62%
MW-2	13,475	10,601	79%	7,669	72%	57%
NE-1	16,395	12,086	74%	8,510	70%	52%
CA-1	24,997	16,663	67%	11,653	70%	47%
CA-2 peak aware	63,525	44,854	71%	28,508	64%	45%
Other: 14 partners	139,309	94,939	68%	68,652	72%	49%
Peak Aware: (3)	18,709	14,474	77%	10,290	71%	55%
Total Summer	362,857	265,527	73%	188,627	71%	52%
Winter 2017/18						
MW-1	113,617	99,225	87%	72,275	73%	64%
MW-2	34,106	29,416	86%	20,226	69%	59%
NW-1	16,301	14,361	88%	10,749	75%	66%
CA-1	47,000	36,138	77%	26,742	74%	57%
CA-5	183,136	115,911	63%	79,879	69%	44%
Other (9 partners)	100,832	80,147	79%	57,596	72%	57%
EU (7 partners)	50,408	42,848	85%	33,557	78%	67%
Total Winter	545,400	418,046	77%	301,024	72%	55%

The table shows that Seasonal Savings was targeted at more than 350 thousand thermostats in the summer of 2017 and that 73% of these thermostats qualified to participate and 71% of those that qualified actually opted to enroll. Summer qualification rates ranged from a low of 56% for a mild climate California utility (where many customers did not use their air conditioners) to a high of 89% for a hot climate utility that only targeted customers who had participated in a demand response program (these were smaller deployments included within the "Other" line item).

The summer opt-in rate among qualified thermostats averaged 71% and ranged from 64% to 78%, but the majority of deployments were in the fairly narrow range of 69% to 73%.

The winter qualification rates are a little larger than the summer but the opt-in rates are comparable. The lowest winter qualification rate was 63% for a mild region of California where many customers rarely heat. In colder climates, winter qualification rates were near 90%. Winter Seasonal Savings was also deployed to seven energy partners in Europe where opt-in rates were generally a little higher than in the US. All five of the partners with opt-in rates greater than 80% were in Europe.

Evaluation Methods

The impacts of Seasonal Savings can be evaluated using multiple approaches. Unlike savings from the initial installation, seasonal savings can be evaluated using thermostat data because the data are available from before and after deployment. The analysis can assess scheduled set points, actual set points (after any manual or other adjustments) and HVAC runtime. Energy impacts can be calculated from runtime savings based on estimated HVAC input rates.

The ideal evaluation design employs a large scale randomized control trial (RCT) where a group of customers from the target population are randomly assigned to be a control group. An RCT removes many potential sources of bias from an evaluation. Because Seasonal Savings is "opt-in" -- customers must choose to participate -- there is a potential for self-selection bias if comparing the opt-in group to the control group. But this potential bias can be removed by analyzing the entire targeted participant group (including those that did not qualify or did not opt-in) compared to the entire control group. This evaluation approach is called an intent-to-treat design or a randomized encouragement design (RED).

An RED eliminates self-selection bias while directly estimating the impact of being in the target participant group (not the impact of actually enrolling). Note that throughout this paper the term “target group” is used to refer to all thermostats in the targeted population regardless of qualification or opt-in status. To estimate the savings per opt-in customer, the RED results can be adjusted for the opt-in rate. For example, if the RED analysis found 2% savings from being in the target group and there was a 50% opt-in rate then the estimated savings per opt-in customer would be 4% (2% / 50% = 4%). Instrumental variable regression can also make this adjustment.

In some cases, the RED design is considered infeasible and quasi-experimental methods are used such as creating a control group by matching on key covariates such as geography and prior set points and runtime. This approach can work well especially in cases where there are many nearby thermostats not in the target population.

When evaluating a large-scale deployment, energy savings can be evaluated by analyzing the HVAC runtime from before and after treatment for the treatment and control groups. One potential drawback of this approach is that runtime data are noisy which requires large samples to accurately assess the modest savings expected from Seasonal Savings. In cases with small samples or little post deployment runtime data (e.g., mild climates or late deployments) the results can be too noisy to be useful. But there is an alternative two step approach that can be used instead that involves first analyzing the impact of seasonal savings on thermostat set points and then analyzing the impact of set points on HVAC runtime (potentially over a much larger population of thermostats/time). The results from this two-step approach have been consistent with results from a direct runtime analysis for large scale RED studies.

Impacts on Thermostat Set Points

Data from a winter 2016/17 deployment in the Northwest is used throughout the remainder of this paper to illustrate the impact analysis.

Figure 1 shows the daily average scheduled heating set points for the target group (whether opt-in or not) and the control group for the example deployment. The target group included 9,745 thermostats and the control group had 6,418 thermostats. A total of 6,746 thermostats actually opted in. The vertical line marks the deployment date (8-Dec-2016). The saw-tooth pattern is due to weekday/weekend differences. The impact of the algorithm is clear as it nudged the schedule to more efficient values over a three-week period.

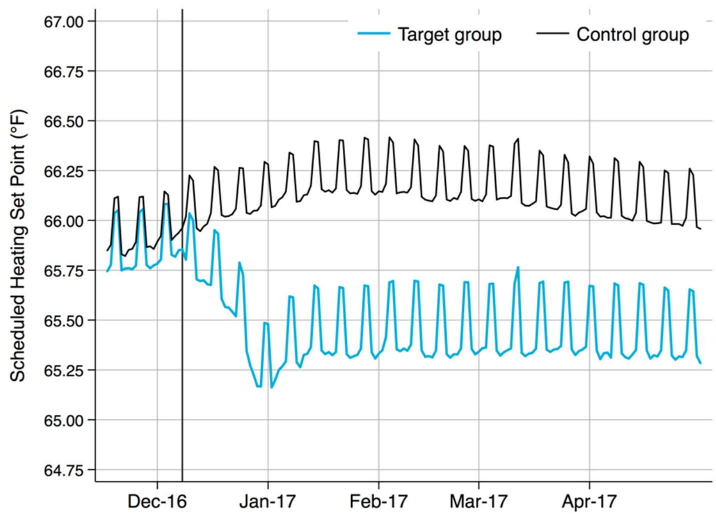

Figure 1. Daily average scheduled heating set points

Figure 2 plots the net difference in set point temperatures (target minus control) for the scheduled set points (i.e., Figure 1 data) and for the actual set point temperatures.

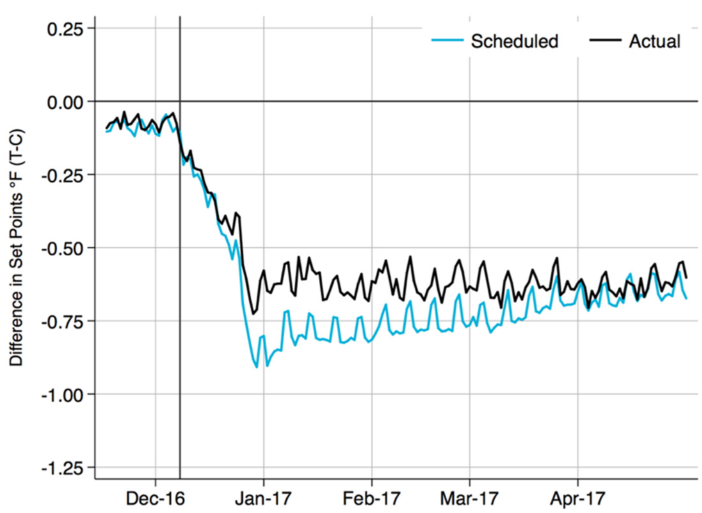

Figure 2. Net difference in daily average set points (Target minus Control)

Actual set points can differ from the scheduled set points due to manual adjustments, the use of Away/Eco mode, and other algorithms running on the thermostat.

Figures 1 and 2 provide a good illustration of how Seasonal Savings affects set points. The net change in set points can be quantified using a regression analysis of daily average set point that includes thermostat and date fixed effects and a simple post treatment group indicator variable. This specification estimates the net impact while accounting for imbalances in the data across dates or thermostats.

For this specific deployment, the regression analysis found a net reduction in executed set points of 0.52°F per targeted participant and 0.69°F (±0.05°F) per opt-in participant (based on a 74% opt-in rate for customers in the analysis).

Across 10 US deployments analyzed from the winter of 2016/17, the average net change in actual set points was 0.65°F per opt-in participant (ranging from 0.44°F to 0.92°F). For the 2017 cooling season, the net change in actual set points averaged 0.35°F and ranged from 0.30°F to 0.41°F. The changes in set points are typically smaller for cooling than heating due to the common preference of maintaining cool temperatures when sleeping – which allows for more efficient set points at night in the winter but not the summer.

The energy savings expected from changes in set points can be statistically estimated using a regression model that fits HVAC runtime as a function of degree days and set points, with thermostat fixed effects. For the example analysis, the regression model estimated that heating usage changes by 6.4% per degree F change in actual set point, resulting in estimated overall heating savings of 4.4% (0.69°F * 6.4%/°F).

The estimated impact of set points on heating runtime varies substantially across climates -- ranging from about 5%/°F in cold climates such as Minnesota to about 10%/°F in mild climates like much of California. This pattern makes sense because heat loss is driven by temperature difference and a 1° change is a smaller fraction of the temperature difference in a cold climate than in a warm climate. For cooling, the impact of set points on runtime are mostly in a narrower range of about 8%-12%/°F due to the smaller range of outdoor temperatures in the cooling season across most climates.

Impacts on HVAC Runtime

The analysis of set points and the estimated impacts on HVAC runtime provide an indirect means to evaluate the energy savings from Seasonal Savings. But the HVAC runtime recorded by the thermostats can also be directly analyzed. Figure 3 shows the daily heating run time for the target group and matched controls over the season for the example deployment.

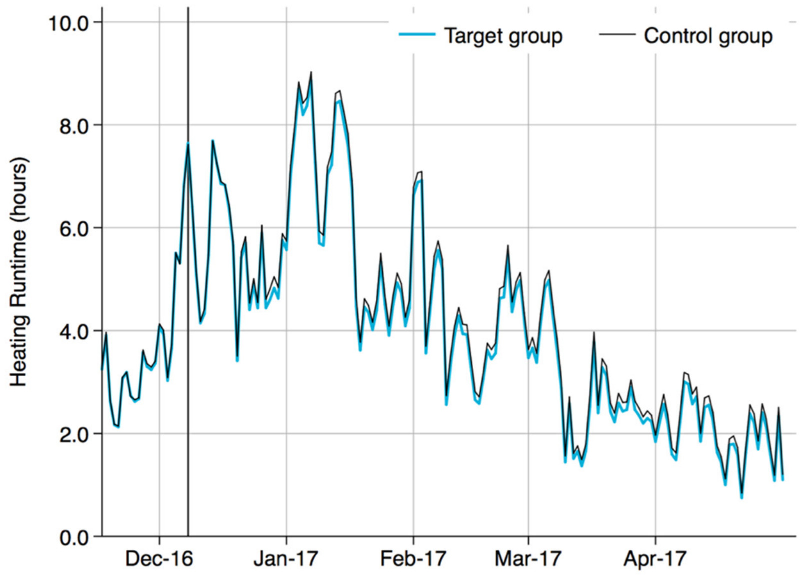

Figure 3. Daily heating runtime for target treatment group and control group

The runtime is nearly identical for the two groups prior to deployment and then the target group runtime becomes slightly, but noticeably, lower than the controls. Figure 4 plots the difference in average daily runtime between the two groups aggregated by week.

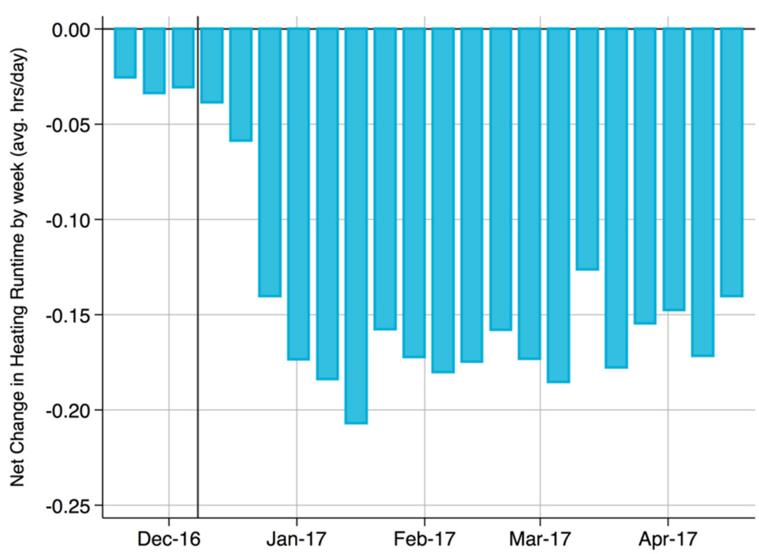

Figure 4. Net difference in daily heating runtime by week (target minus control)

The plot shows that the difference in runtime between the groups is quite small prior to the deployment and grows sharply over the next few weeks as the algorithm adjusts the scheduled set points, leveling out at about 0.15 hours/day. Comparing this value to the average hourly runtimes shown in Figure 3 indicates about a 3%-4% reduction in average runtime for the full target population.

The runtime impacts can be quantified using regression analysis similar to methods commonly employed to analyze utility meter data. The primary modeling approach fit daily heating runtime as a function of heating degree days (HDD base 60°F); and indicator variables for the target group and the post-treatment period and interaction terms and included thermostat fixed effects. The dataset included more than 3 million daily observations from 9,144 target participants and 6,024 control customers.

The net heating savings from the analysis was 4.8% ±1.0% (95% confidence interval, standard errors clustered by thermostat) -- nearly identical to the 4.4% savings estimated from the set points analysis. Alternative model specifications were also tested (fixed effects without degree days, adding date fixed effects, a post only model using pre-deployment runtime as an explanatory variable) and all methods resulted in savings estimates between 4% and 5%.

Energy Savings

The runtime analysis provides estimates of the percent savings and hours of runtime saved. But the evaluation goal typically includes measuring energy savings and so the runtime must be converted into kWh and therms, which requires estimating the HVAC input rates. This estimate can come directly from the partner utility based on existing research (e.g., end use load studies, HVAC equipment rebate/replacement program data). If no other estimate is available, Nest can provide an estimate based on climate and other factors that is generally similar to estimates from these other sources.

The input rates of air conditioners (and heat pumps) will vary some with outdoor temperature (approximately 1%/°F for air conditioners). This non-linearity can be included in the analysis but the net impact is minimal and tends to slightly boost estimated cooling savings due to the greater set point changes achieved in the middle of the day.

For this sample winter deployment, the estimated input rates resulted in savings of 20 therms of natural gas and 16 kWh for each gas heating system and 126 kWh for each heat pump. The savings only include the specific post-deployment period (from 8-Dec-2016 through May 2017) and include no savings for future years (i.e. zero persistence).

Results From Multiple Evaluations

Table 2 summarizes key evaluation results from large scale (> 5,000 opt ins) deployments in the winter of 2016/17 and the summer of 2017. The energy savings listed are the results from the runtime regression analysis.

Table 2. Seasonal Savings Impact Summary

Utility /Deployment	Set Point Change	Energy Savings	# opted in
Cooling 2017			
MW-1	+0.36	3.7% ±1.3%	53,345
MW-2	+0.32	3.7% ±2.5%	7,669
CA-1	+0.33	2.9% ±0.9%	11,653
CA-2 peak aware	+0.35	4.2% ±1.5%	28,508
NE-1	+0.37	6.5% ±1.9%	8,510
Heating 2016/17			

Utility /Deployment	Set Point Change	Energy Savings	# opted in
CA-3	-0.62	8.2% ±3.8%	50,904
MW-2	-0.74	3.5% ±0.7%	10,491
NW-1	-0.69	4.8% ±1.0%	6,746
EU-1 (UK)	-0.54	5.0% ±0.9%	7,817

The table shows that set points are changed by similar amounts across deployments – especially in the summer. The lowest heating season set point reduction was for a UK partner, where set points were often already fairly efficient. The percent savings vary more widely but the uncertainty in the estimates may account for much of this -- especially in milder climates like CA-3 where the 8.2% heating savings was equal to just 5 therms of natural gas.

Independent Review/ Verification

Several of the Seasonal Savings evaluations conducted by Nest have been reviewed and independently re-analyzed by third party consultants hired by utility partners. These reviews have been generally consistent with the Nest results. In some cases there have been small differences due to different modeling choices – some resulting in larger savings and some in smaller savings.

One potential option that has been proposed for verifying the impacts without relying on thermostat data is to perform a utility billing data analysis. But there are significant obstacles for this approach. First, there is the logistical issue that Nest doesn't have customer utility account numbers (or even addresses in many cases). Second, Nest has strict customer privacy policies that would prevent sharing data without a specific customer sign off. One approach to address this issue has been to email customers and ask for utility account information and authorization. But response rates from these efforts have been low (~5%), resulting in small samples that may also suffer from response bias. This email approach was used for the example winter deployment described previously and the billing analysis found savings virtually identical to the thermostat data analysis: 21±18 therms, but with a much wider confidence interval and significant potential for bias.

Other utility partners have suggested requiring customers to provide utility account data as part of enrollment. But this approach would reduce participation rates dramatically – the 70%+ opt-in rates come from the ease of enrollment. Any additional steps would create a hurdle to participation that renders the approach not worth pursuing.

Savings Persistence and Repeated Participation

To date there have not been a sufficient number of large scale RCT deployments from prior years to adequately assess the net persistence of impacts or the impacts of customer re-enrolling in follow up years. The winter of 2017/18 results may shed some light on these issues. The limited analyses performed thus far suggest that schedules may degrade by about 25% when tracked into the next year and this degradation may continue across the 2nd season. But customer who enroll again in the second year end up with even more efficient schedules which should lead to even greater net savings.

"Peak Aware" Seasonal Savings

In the summer of 2017, Nest launched a "peak aware" version of Seasonal Savings designed to help utilities reduce peak demand while at the same time savings customers on their energy bills. Unlike a demand response program, which typically provides financial incentives to customers and can use substantial pre-cooling before the peak and may push some comfort limits during peak, the peak aware version of Seasonal Savings is much less aggressive at shifting load. Instead, schedules are shifted to reduce loads during peak periods without making set points less efficient immediately prior to peak (no explicit pre-conditioning) and without pushing set points up as high as demand response might during peak to maintain comfort.

Figure 5 shows the pre/post changes in scheduled weekday set points by hour of day for the target and control groups for a deployment in California. The algorithm shifted the set point up by about 1°F during the day but then returned it to the pre-deployment value just prior to the peak – creating a sort of pre-cooling. During peak, the set points were shifted up again. The algorithm attempts to shift set points up by more during the peak than the morning, but the peak coincides with people coming home from work and any manual adjustments customers make are learned by the algorithm to maintain comfort.

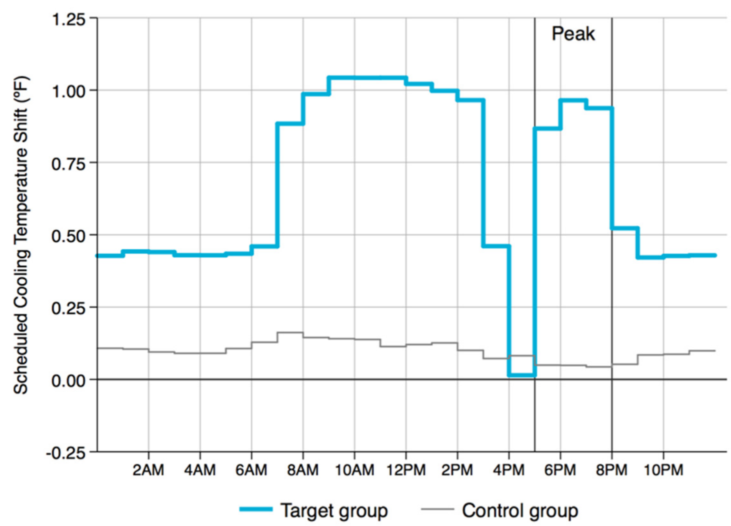

Figure 5. Pre/Post change in scheduled cooling set points by hour of day

Regression analysis of thermostat runtime data was used to quantify the load shifts employing models similar to the standard energy savings analysis applied to each hour of the day. The runtime impacts were converted to watts using the estimated HVAC input rate.

Measuring peak demand impacts can be complicated because there are different definitions of peak demand. One definition is simply the average load shift on weekdays during the cooling season. But it may be of more interest to a utility to focus on hotter days when peak demand reductions are usually more valuable. Figure 6 shows the net load shift on average cooling weekdays post deployment (22-Aug to 30-Sep-2017 in this example).

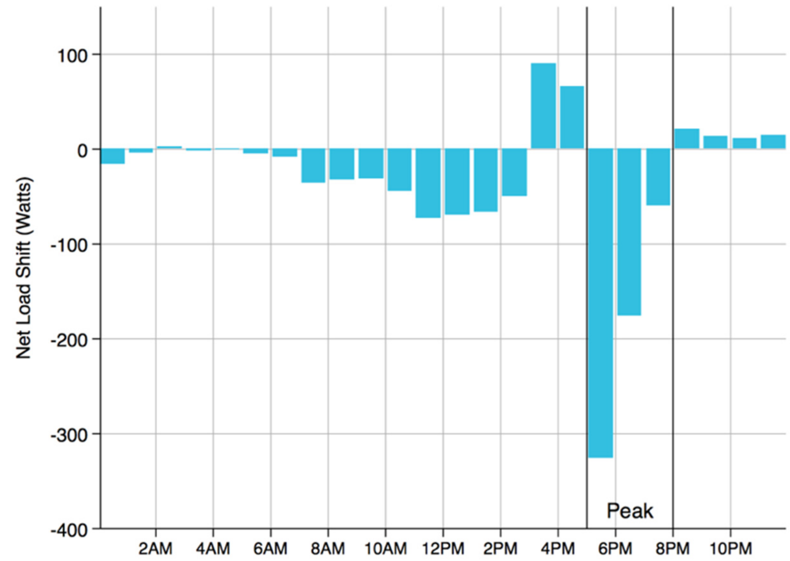

Figure 6. Net Load shift by hour of day – cooling season weekdays

The plot shows a small amount of pre-cooling followed by a peak reduction of 326 watts per opt-in participant during the first hour of peak (a 27% reduction) which then tapered off over the next two hours of peak. The load shift averaged 187 watts (16%) across the three hours. The analysis was repeated for just warmer weekdays – defined as days where the average outdoor temperature exceeded 75°F -- and the resulting first hour load shift was 463 W (24%) and the average peak period shift was 245 W (13%). The same analysis applied to a deployment in the Midwest found an average 135 W (21%) load shift across a four hour peak and 213 W (14%) on warmer days.

The peak-aware version of Seasonal Savings produces load shifts that are smaller and shorter-lived compared to standard demand response programs. But it has the advantage of quick and easy large-scale deployment with high participation rates and no need for individual customer incentives. Work is continuing on the algorithms to provide even greater impacts.

Conclusions

Nest's Seasonal Savings algorithm provides an example of how connected smart thermostats create new opportunities for energy savings and load shifts beyond what is achieved from the initial installation. Easy enrollment has led to opt-in rates that typically exceed 70% of qualified thermostats (>50% of the entire thermostat population) -- resulting in hundreds of thousands of participants. The approach of making many tiny shifts in scheduled set points over time has resulted in significant energy savings while maintaining customer comfort. These results show some of the promise that connected technology and well-designed algorithms can provide and suggest many future opportunities.

References

Apex Analytics, LLC . 2016. *Energy Trust of Oregon Smart Thermostat Pilot Evaluation.* Energy Trust of Oregon. https://www.energytrust.org/wp-content/uploads/2016/12/Smart_Thermostat_Pilot_Evaluation-Final_wSR.pdf

Apex Analytics, LLC and Demand Side Analytics, 2017. *Energy Trust of Oregon Nest Thermostat Seasonal Savings Pilot Evaluation.* Energy Trust of Oregon. https://www.energytrust.org/wp-content/uploads/2017/12/Energy-Trust-of-Oregon-Nest-Seasonal-Savers-Pilot-Evaluation-FINAL-wSR.pdf

Energetics Inc. and VEIC (Vermont Energy Investment Corporation). 2016. *Overview of Existing and Future Residential Use Cases for Connected Thermostats*. U.S. Department of Energy Office of Energy Efficiency and Renewable Energy

EPA (Environmental Protection Agency). 2016. "Connected Thermostat Specification V1.0" Accessed March 2018. https://www.energystar.gov/products/spec/connected_thermostats_specification_v1_0_pd

Henschel, R. and N. Corsetti. 2014. "A Behavior Program for Everyone: Moving Beyond the Opt-out Home Energy Report" . In *Proceedings of the 2014 ACEEE Summer Study on Energy Efficiency in Buildings* 7:124-132. Washington, DC: ACEEE https://aceee.org/files/proceedings/2014/data/papers/7-848.pdf

Nest. 2017. https://nest.com/blog/2017/08/30/e-is-for-everyone/

Nest. 2018. https://nest.com/support/article/What-is-Seasonal-Savings

Robinson, J., R Narayanamurthy, B. Clarin, C. Lee, and P. Bansal. 2016. "National Study of Potential of Smart Thermostats for Energy Efficiency and Demand Response". In *Proceedings of the 2016 ACEEE Summer Study on Energy Efficiency in Buildings* 2:1–11. Washington, DC: ACEEE https://aceee.org/files/proceedings/2016/data/papers/2_1172.pdf