The R.E.D. Carpet of Thermostat Optimization

Jesse Smith, Demand Side Analytics, Atlanta, GA

ABSTRACT

In 2019 thermostat manufacturer ecobee developed its eco+ optimization platform to improve the energy performance of residential HVAC systems with minimal effort for the user. The software prompts customers to select their personalized savings preferences which guides the performance of each algorithm. The platform consists of three categories of algorithms: Demand Response (DR), Time-of-Use (TOU), and Energy Efficiency (EE). During summer 2019 and summer 2020, the optimization platform was deployed across six climate regions of North America to a large pilot group of thermostats to demonstrate the new capabilities. The pilot included a population of approximately 240,000 thermostats using a Randomized Encouragement Design (RED). Devices were stratified by climate zone then randomly assigned to either an experimental group or a control group. The experimental group was invited (encouraged) to participate in the pilot and the control group was not. The RED provides a robust experimental design against which to measure the impacts of the platform because the control group experiences the same weather and other external factors as the experimental group. Difference-in-differences modeling is used to compare the HVAC runtime characteristics of the experimental group to the control group after the rollout and produce estimates of the impact of the optimization offer. These impacts are divided by the acceptance rate to estimate per-device impacts for users opting into the platform. The study findings illustrate how technology can adapt to customer preferences and deliver energy and cost savings with limited effort required on the part of the consumer.

Introduction

The rollout of eco+ by ecobee was intentionally constructed to facilitate measurement of impacts through a Randomized Encouragement Design. Approximately 50% of thermostats were randomly assigned to a control group and not offered the eco+ platform to serve as the counterfactual, or baseline, against which energy and demand impacts in the experimental group are measured. In this design, the control group is made up of ecobee smart thermostats without eco+ optimization and the experimental group is made up of ecobee thermostat users offered eco+ optimization. The RED framework intentionally sacrifices aggregate energy and demand impacts to facilitate a rigorous evaluation methodology. A control group “buffer” group was also created by random assignment in case ecobee owners in the control group learned of eco+ and asked to be included in the offering.

The eco+ platform consists of three categories of algorithms. By opting in once, consumers receive three types of automated optimization.

- **Demand Response (DR)** – Presented to users as Community Energy Savings (CES), this feature shifts cooling loads away from peak hours when the electrical grid is most constrained.
- **Time-of-Use Optimization (TOU)** – For ecobee owners whose retail electricity rate varies by hour of the day, the TOU algorithm shifts energy use from high-price hours to lower-price hours.
- **Energy Efficiency (EE)** – Features like Enhanced Smart Home & Away, Schedule Assistant, and Adjusting for Humidity help ecobee owners lower their overall heating and cooling energy consumption.

Table 1 shows the count of thermostats across the 11 regions and three experimental cells. Region 1 is Canada. Regions 2 through 6 correspond to five US Department of Energy Building America Climate Zones overlaid on a map in Figure 1. Regions 7 through 11 are specific electric utility service territories with high prevalence of time-varying pricing. These utility service territories were intentionally over-sampled to bolster the sample size for the eco+ TOU optimization algorithm analysis.

Table 1: Thermostat Count by Region and Experimental Cell

Region	Experimental	Control	Buffer	Total
01 Canada	10,062	10,026	1,001	21,089
02 Cold/Very Cold	30,001	30,000	3,000	63,001
03 Hot-Dry/Mixed-Dry	5,579	5,570	557	11,706
04 Hot Humid	15,000	15,000	1,500	31,500
05 Mixed Humid	30,000	30,000	3,000	63,000
06 Marine	5,069	5,085	510	10,664
07 Canada TOU (Hydro One)	1,927	1,932	195	4,054
08 Cold TOU (Fort Collins)	140	139	13	292
09 Dry TOU (Pacific Gas & Electric)	8,156	8,150	815	17,121
10 Dry TOU (Sacramento Municipal Utility District)	2,800	2,800	280	5,880
11 Marine TOU (PG&E)	9,473	9,461	945	19,879
Total	**118,207**	**118,163**	**11,816**	**248,186**

The deployment of eco+ to users was no different in Regions 7-11 than in Regions 1-6. All recipients were given the option to select a time-varying electric tariff and enable the TOU optimization feature. In fact, several thousand devices in Regions 1-6 did utilize the TOU algorithm. However, for analysis purposes, we treated Regions 1-6 as EE/DR cells and Regions 7-11 as TOU cells. Mechanically, this means the DR and EE results presented for the Hot-Dry/Mixed-Dry climate zone are based on findings from Region 3 and do not include devices from Regions 9 and 10 – even though they are in the same parent climate zone.

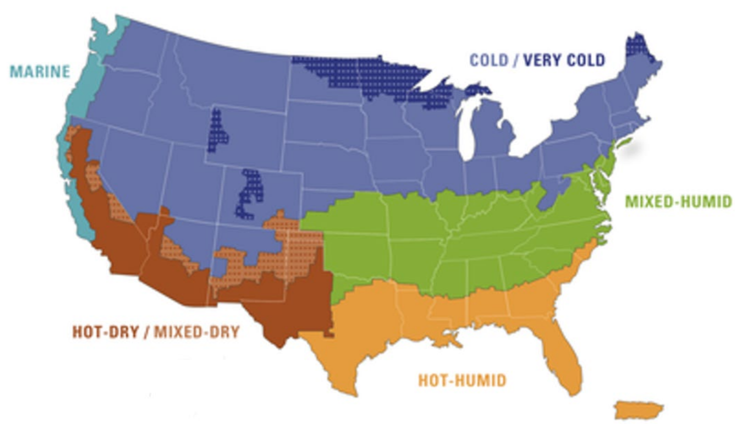

Figure 1: United States climate zone map

The eco+ RED only included devices who were not enrolled in a utility DR program. All DR events were initiated by ecobee for demonstration purposes. The demand response event testing strategy also included a time zone component and most regions included multiple time zones. For example, Region 5 (Mixed Humid) spans both the eastern and central time zones. On some Mixed Humid DR event days, both time zones were dispatched and on other days, DR was initiated for a single time zone within the region. The DR initiation decision was primarily based on temperature triggers.

Pre-Treatment Equivalence

Since the RED framework underpins the entire analysis, a key upfront research question is “was the randomization sound?” If successfully implemented, randomization should result in no meaningful differences between the experimental and control groups. The only difference between the groups should be that the experimental group is offered the treatment while the control group is not. A primary goal of randomization is to avoid systematic differences for key inputs that impact measurement. Specifically, if pre-treatment differences are detected for the dependent variable being measured (in our case cooling runtime), then the actual impact due to treatment is the measured difference during the treatment period minus the known preexisting difference. To be included in the RED, thermostats had to be connected prior to June 1, 2019. Figure 2 compares the distribution of average hourly runtime for treatment and control customers in each region. This shows that on average across pre-treatment days, the distribution of average runtime is nearly the same for treatment and control customers within each region. Figure 2 also reveals drastic differences in average cooling usage across the six study regions, which are largely a function of weather.

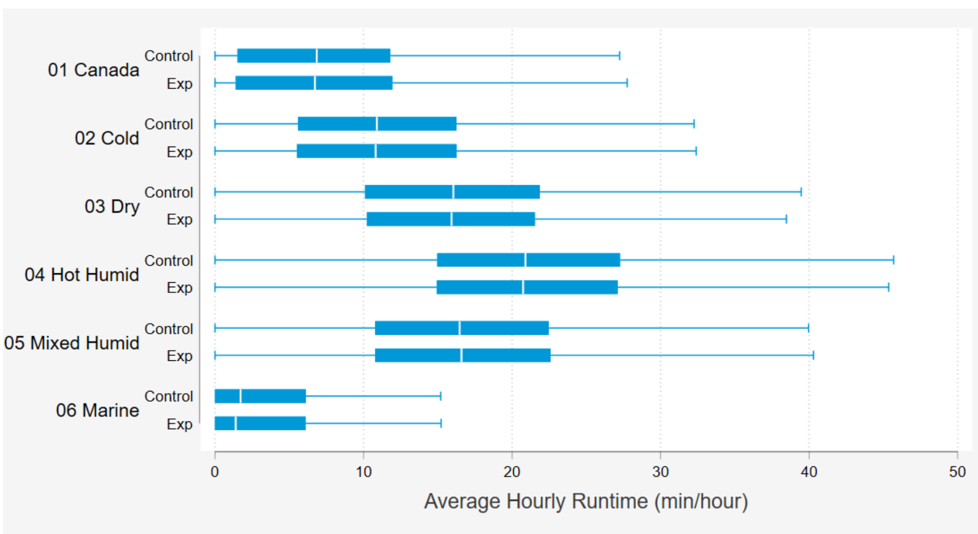

Figure 2: Runtime Distribution Comparison

Acceptance Rates

Regression analysis of the RED produces estimates of the average impact of the eco+ *offer*. These Intention to Treat (ITT) impacts are free of selection bias or endogeneity concerns because of the random assignment. Which devices opt-in to receive the eco+ platform and which ones do not may be correlated with all types of unobservable characteristics, but we can be confident that the full experimental and control groups are balanced with respect to these characteristics. The conversion from ITT impacts to Local Average Treatment Effect (LATE) impacts is performed by dividing the ITT impacts by the proportion of devices in the experimental group that accepted the treatment. For example, if the ITT impacts were 50 kWh per device and 60% of the devices in the experimental group opted-in to the treatment, the per-device LATE impact would be:

$$LATE = \frac{ITT}{\%\,Treated} = \frac{50\,kWh}{0.6} = 83.33\,kWh$$

Determining the percentage of devices “treated” amongst the experimental group is thus an important consideration for the analysis. There are some nuances to the ultimate definition of “treated” that are specific to the DR, EE, and TOU optimization strategies that are discussed in the respective sections. However, there are common aspects across the three strategies which we discuss here.

The first filter in the analysis affects the experimental, control, and control buffer cells. For eco+ to work, the ecobee thermostat needs to be online and connected to WiFi. Around 8% of the devices in the RED population had no runtime data during the experiment. It was determined through discussions with ecobee that many of these devices had been offline for months prior to June 1, 2019 – either because they were no longer installed or being operated offline like a programmable thermostat.

Table 2 shows the thermostat counts by experimental cell. Offline thermostats occur in virtually identical rates for each group as this occurrence is uncorrelated with the rollout. The revised study population becomes 228,492 devices and subsequent calculations of acceptance use the online count as the denominator.

Table 2: Online Rates by Study Cell

Study Cell	Offline	Online	Total	% Online
Control Buffer	931	10,885	11,816	92.12%
Control	9,454	108,709	118,163	92.00%
Experimental	9,309	108,898	118,207	92.12%
Total	**19,694**	**228,492**	**248,186**	**92.06%**

For any one of the 108,898 online experimental thermostats to receive the eco+ optimization features, a series of additional criteria needed to be satisfied.

- **The invitation went out as intended** – A thermostat may have been offline when the invitation was issued, may not be running a cooling schedule, or may have some other technical issue that prevented the invitation from being issued.
- **The device has features data** – Only devices that engaged with the invitation have observations in the features data set.
- **Terms Accepted** – These devices accepted the eco+ terms and conditions.
- **Comfort Setting Greater Than 1** – Users who accept the eco+ terms can select a slider level from 1 to 5 where 5 is the most aggressive implementation and 1 is the least aggressive.

Table 3 shows the number of devices that made it through each of the acceptance stages in the summer 2019 analysis.

Table 3: Eco+ Acceptance Hierarchy

Stage	Stage Description	Device Count	Percentage
A	Randomized	118,207	N/A
B	Online	108,898	100.0%
C	Invited	104,080	95.6%
D	Has Features	81,303	74.7%
E	Terms Accepted	62,748	57.6%
F	Comfort Setting > 1	59,699	54.8%

Across the entire population of online thermostats randomized into the experimental group, approximately 54.8% accepted the eco+ offer. The default slider level of 4 was most common, followed by 5 and then then 3.

Connected Load

Translating HVAC runtime, or changes in runtime, to energy usage and demand requires an assumption about the size of the air conditioning (AC) units connected to ecobee thermostats. Because we did not know the size of ecobee customers' AC units, we collected AC unit size and efficiency measurements/assumptions from various technical resource manuals (TRMs), appliance saturation studies, and evaluation reports. Table 4 shows the results of our research and final assumptions by climate zone.

Table 4: Connected load assumption by climate zone

Climate Zone	Tons	SEER	kW per Device
01 Canada	2.15	10.5	2.45
02 Cold	2.75	10.5	3.10
03 Dry	3.25	10.5	3.48
04 Hot Humid	3.25	10.5	3.60
05 Mixed Humid	2.75	10.5	3.04
06 Marine	2.60	10.5	2.93

Demand Response Analysis

Demand Response event impacts were modeled using a difference-in-differences regression analysis. Regressions were run for each region and time zone separately by hour on non-holiday weekdays. The following components were included in the selected regression model:

- **$\mathbf{Runtime}_{t,d,h}$**: The hourly runtime for thermostat t, on date d, in hour h. Ranges from 0 to 1, where zero is no cooling runtime and 1 means the air conditioner operated for all 60 minutes of the hour.
- **$\mathbf{B}_t$**: The thermostat-level fixed effect.
- **$\mathbf{Post}_{t,d}$**: Indicator equal to 1 on or after the first eco+ invitation for a study region. Zero otherwise.
- **$\mathbf{CDH60t}_{t,d,h}$**: Cooling degree hours, base 60 degrees (F). Equal to the maximum of outdoor temperature minus 60, and zero.
- **$\mathbf{Relative\ Humidity}_{t,d,h}$**: Relative outdoor humidity for thermostat t, on date d, in hour h. Ecobee stores RH values on a scale from 0 to 100.
- **$\mathbf{Mean15}_{t,d,h}$**: Represents the average outdoor temperature from midnight until 3:00 PM.

- **RT_MidDay$_{t,d}$:** Captures the average runtime for thermostat t, on date d, from 10:00 AM to 1:00 PM. This term serves to "calibrate" the regression to any day-specific trends and improve the accuracy of the estimates during the event hours.
- **Treatpost$_{t,d}$**: Equal to 1 for the experimental group in the post-period. Zero otherwise.
- **Date*Treatpost**: Interaction between the treatpost variable and each DR event day. The β_d coefficients are our parameters of interest and capture the average runtime impact in the ITT group, net of any differences between the ITT and control groups during the pre-treatment period.

$$Runtime_{t,d,h} = \beta_0 + \beta_1 Post + \beta_2 CDH60 + \beta_3 Relative\ Humidity + \beta_4 Mean15 + \beta_5 RT\_MidDay + \beta_d (Date * Treatpost)$$

Figure 3 illustrates the DR algorithm on a sample event day called from 3pm to 6pm. The blue line indicates the control group's average cooling runtime and the dashed gray line shows the runtime for the experimental group. Orange bars show the difference between these curves, which are the basis for the ITT impacts. The green line shows the average pre-cooling time for the experimental group and the purple dashed line shows the average DR setback time during the event.

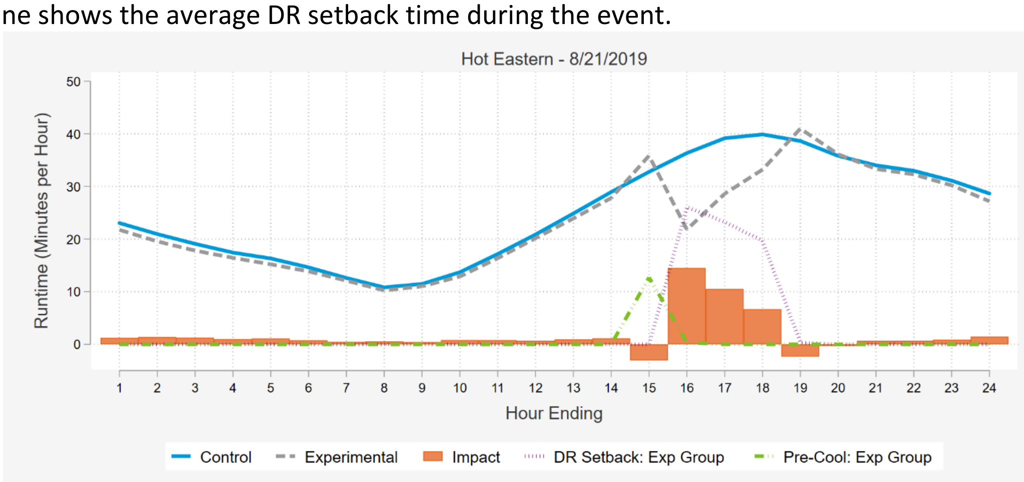

Figure 3: Example DR event

Runtime impacts were modeled via regression and scaled by the percent treated to estimate LATE impacts, or the average impact among devices who received the DR algorithm. Approximately 45% of the experimental group received the eco+ DR algorithm on this event day, so the LATE impacts are roughly 2.2 times the ITT impacts. Figure 4 shows the modeled impacts on the example event day on both an ITT and LATE basis.

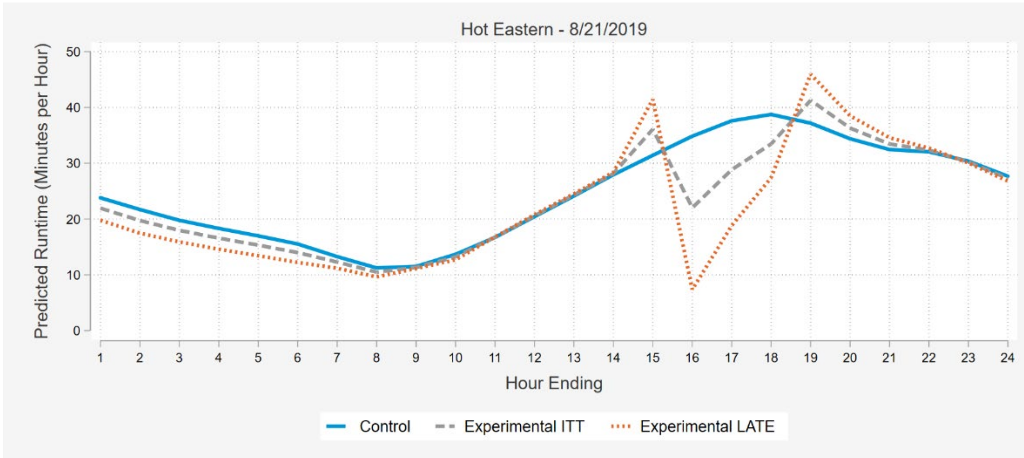

Figure 4: Modeled runtime impacts

There was a total of 55 demand response events in the summer 2019 analysis. Impacts are reported separately for each region and time zone by hour of event. Demand response events ranged from two to four hours in length and were called at various times on hot weekday afternoons. Event days were chosen based on market research of existing DR programs and temperature triggers were used to select the days on which to dispatch DR. Impacts varied by region, time zone, date, and event hour. We estimate average DR savings of 0.91 kW per opt-in thermostat across all event hours. Figure 5 shows the average impacts by event hour and region and the participation rate over the course of events. All summer 2019 events have an hour 1 and hour 2, so average impacts from these hours are weighted more heavily in the average hourly demand savings of 0.91 kW. Fewer events are three or four hours long, leading to less weight in the overall savings estimate. In aggregate, summer 2019 impacts were largest during the first event hour and diminished in subsequent hours. This downward trend is typical of thermostat DR programs that use a setback strategy.

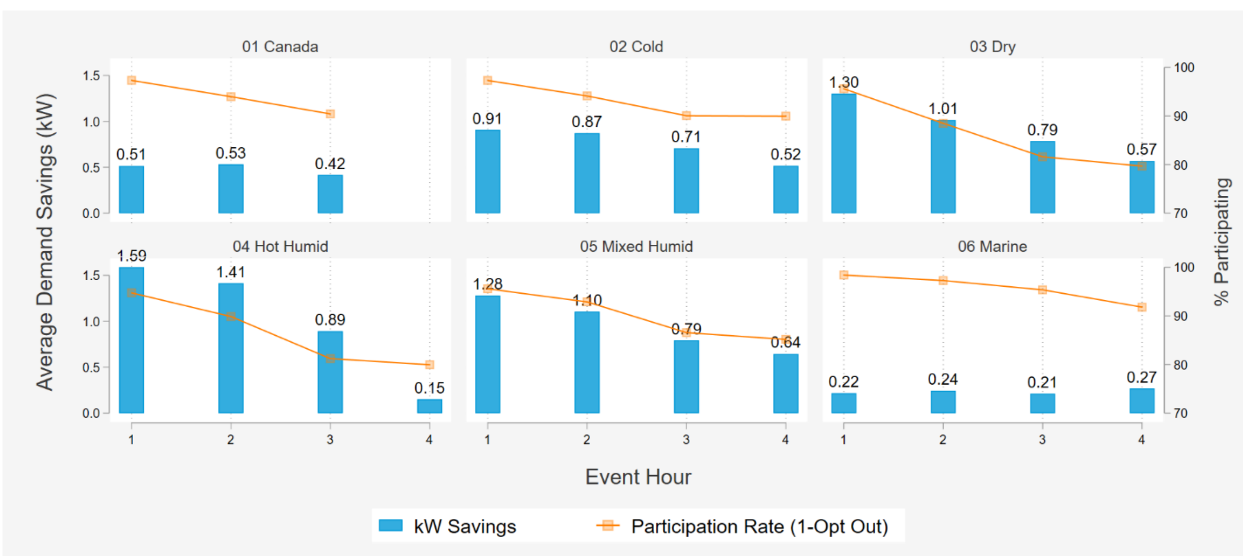

Figure 5: 2019 Demand savings and participation rates by event hour

There were 28 demand response events in the summer 2020 analysis. Demand response events ranged from three to four hours in length and were called either at 2pm or 4pm local time. Event days were chosen based on market research of existing DR programs and temperature triggers were used to select the days on which to dispatch DR. We estimate a weighted average DR savings of 1.12 kW per opt-in thermostat across all event hours. For the EE/DR regions in the RED population during summer 2020, this translates to approximately 29 MW of peak demand reduction capability if all regions were dispatched on the same day.

Figure 6 summarizes the LATE results of all demand response events called during summer 2019 and summer 2020. In general, the per-device load reductions increase with outdoor temperature. At any given temperature condition, load reductions are highest during the first hour of a demand response event and decrease in each subsequent hour.

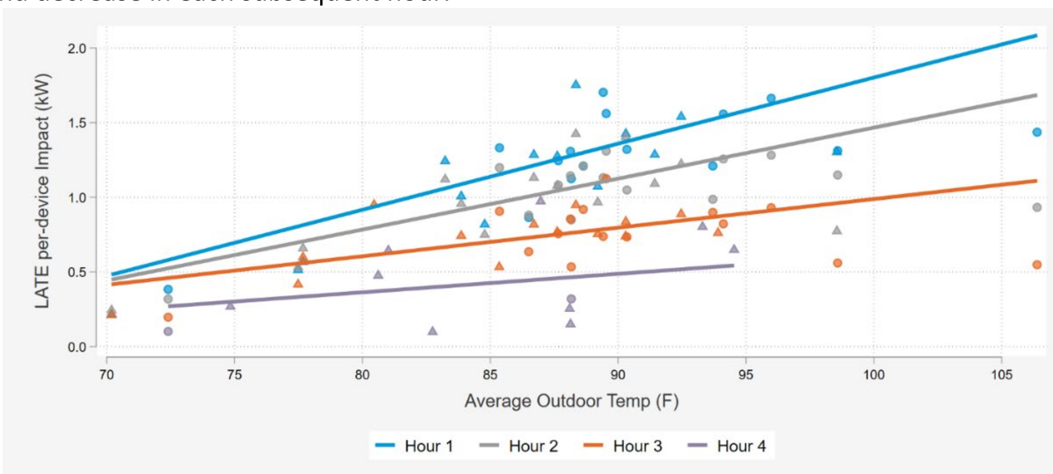

Figure 6: Per-device load reduction versus outdoor temperature by event hour

Time of Use Analysis

While the eco+ DR feature adjusts thermostat settings in response to specific event calls, the Time-of-Use (TOU) optimization features react to the price signals provided by a participant's utility rate. The TOU algorithm shifts cooling load to periods where electricity rates are lower by activating a pre-cooling function in anticipation of the higher priced hours. During periods when electricity rates are higher the eco+ algorithm engages a setback mode which reduces compressor runtime. Figure 7 illustrates a three tier rate. The algorithm engages in pre-cooling in the hour prior to both rate increases, but the setback mode is concentrated in the peak rate period.

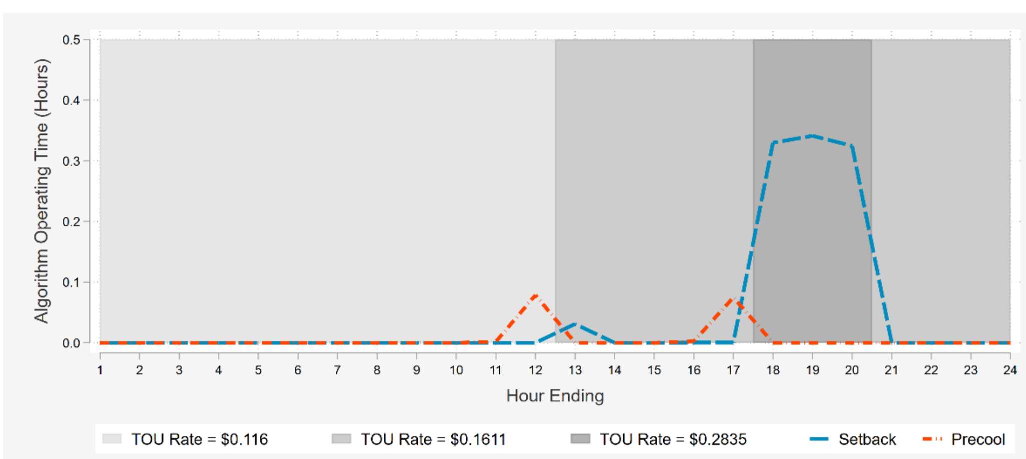

Figure 7: TOU feature operation example

As a result of the thermostat operating in pre-cooling and setback mode the compressor runtime for the HVAC unit is reduced during that period as illustrated in Figure 8. From a utility perspective shifting the cooling load shaves the peak demand. For customers reducing runtime during peak rate periods will result in bill savings. The associated bill savings are primarily from moving cooling load to periods when electricity use is less expensive, but TOU participants also see bill savings from reducing overall consumption.

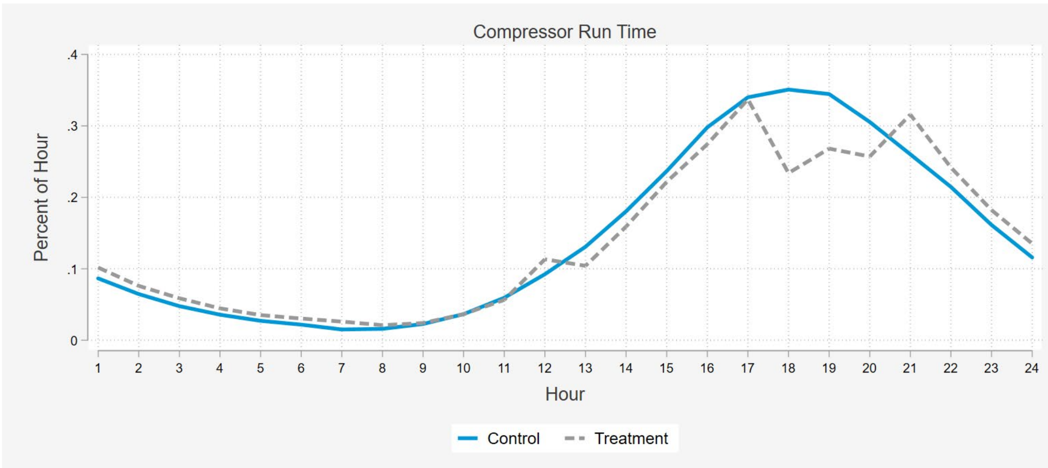

Figure 8: Comparison of compressor runtime between TOU participants and control group

TOU acceptance rates were too low to use the primary RED analysis technique, so we selected a matched control group (from the full randomized control group) based on pre-treatment compressor runtime. A positive side effect of this approach is that it enabled us to include participants from outside the TOU cells who were receiving TOU treatment. As discussed previously, there was no inherent difference in the eco+ offering across cells. Because many of the participants in the DR/EE cells were on TOU rates they could be added to the TOU analysis treatment pool which increased TOU treatment counts. For example, the majority of the Hydro One TOU participants ultimately analyzed came from the 01 Canada region cell as opposed to 07 Canada TOU.

For summer 2019, we analyzed four separate rates from five climate zones. The regression modeling approach for the TOU feature mirrored the DR analysis except for the RT_MidDay term, which was omitted from the TOU regression. Results by rate are displayed in Table 5. In all results, the impacts are relative to an ecobee thermostat without eco+.

Table 5: TOU high level results – summer 2019

Rate	Climate Region	Peak Duration (hours)	Price Ratio (Peak: Off-Peak)	Average kW Savings During Peak Period	On-Peak Percent Savings	Overall Percent Energy Savings	Percent Savings on Cooling Energy	Daily Bill Savings
Hydro One Res TOU	Canada	6	2	0.18	36%	3.4%	8%	$0.09
FPL RTR-1	Hot Humid	9	5.8	0.22	13%	5.0%	10%	$0.39
SMUD Res TOD	Hot Dry	3	2.4	0.25	23%	3.5%	8%	$0.19
PG&E EV-A	Mixed Dry	6	3.7	0.18	28%	8.8%	19%	$0.50
PG&E EV-A	Marine	6	3.7	0.10	20%	4.0%	11%	$0.23

In general, bill savings associated with the eco+ TOU feature are larger when participants have higher cooling usage and more expensive peak electricity prices. Even though the percentage of on-peak savings is highest in Canada, the magnitude of savings is low compared to the other regions due to the cheaper energy prices and limited air conditioning usage. Another factor affecting the average on-peak percent savings is the duration of the peak. Shorter peak hours yielded larger average demand impacts (kW) but less overall energy savings (kWh). The largest energy expenditure savings were found on the PG&E rate in part because this rate had TOU pricing on weekends. The PG&E rate was also substantially higher than other rates. In fact, the PG&E off-peak rate is higher than the Hydro One on-peak rate if the CAD to USD exchange rate is considered. For the summer 2020 season, we retained one of the rates from the previous summer, SMUD Residential TOD, and selected three new rates for analysis. These new rates provided additional variation in region and rate structure. The 2020 TOU analysis used the same methodology as the previous summer. Table 6 presents the results of the 2020 TOU analysis by rate.

Table 6: TOU high level results – summer 2020

Rate	Climate Region	Peak Duration (hours)	Price Ratio (Peak: Off-Peak)	Average kW Savings During Peak Period	On-Peak Percent Savings	Overall Percent Energy Savings	Percent Savings on Cooling Energy	Daily Bill Savings
Duke Energy RT	Mixed Humid	6	1.2	0.25	20%	8%	9%	$0.11
PacifiCorp EV-TOU	Cold	5	3.3	0.43	33%	15%	23%	$0.51
SMUD Res TOD	Dry	3	2.4	0.28	21%	3%	7%	$0.18
Tucson Electric Power Demand TOU	Dry	4	1.7	0.46	25%	6%	9%	$0.17

Generally, the 2020 results followed the same patterns as 2019. All rates experienced the greatest impacts in the first hour of the peak period. Tucson Electric Power had the highest on-peak average demand savings across all analyzed rates due to Arizona's extreme summer weather and large cooling loads. Two of the 2020 rates included demand charges, where customers incur an additional charge based on their peak demand within a certain time period. Since the eco+ analysis uses HVAC runtime data rather than whole home data, we cannot identify the peak demand hour and calculate the effect of the TOU algorithm on this billing determinant. For these rates, the percent bill impacts only apply to the energy portion of the bill, and do not take into account the demand charge portion of the bill.

Figure 9 shows the effect of users' slider level on the TOU algorithm. The TOU algorithm implements more aggressive pre-cooling and setbacks during on-peak hours when users select a higher slider level.

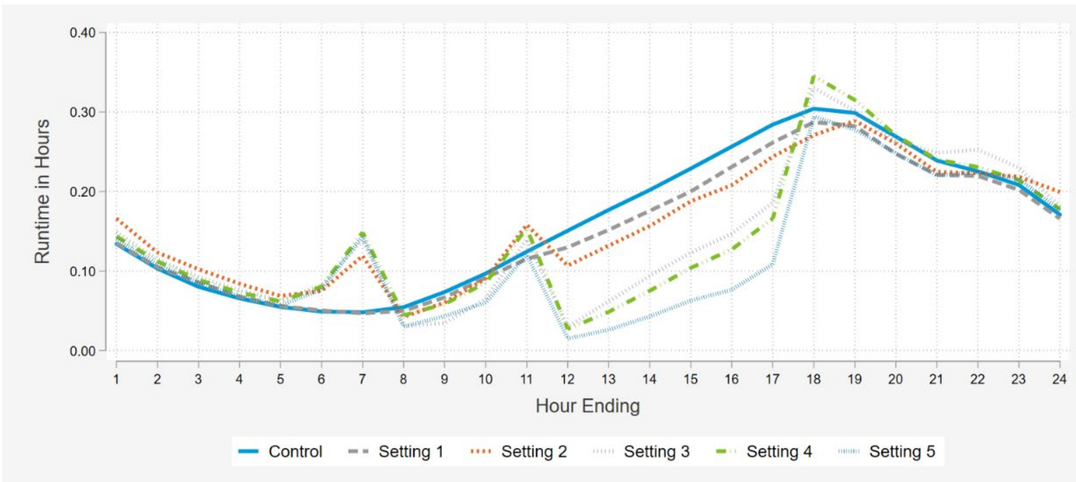

Figure 9: Hydro One average weekday runtime by slider level

Energy Efficiency Analysis

Figure 10 shows the percent difference between ITT and controls in average daily cooling runtime, by date, for the Mixed Humid Eastern climate zone. The blue curve shows the pre-period, the gray curve shows the post-period from Summer 2019, and the orange curve shows the Summer 2020 impacts for this region. The black bar indicates a time gap between Summer 2019 and Summer 2020.

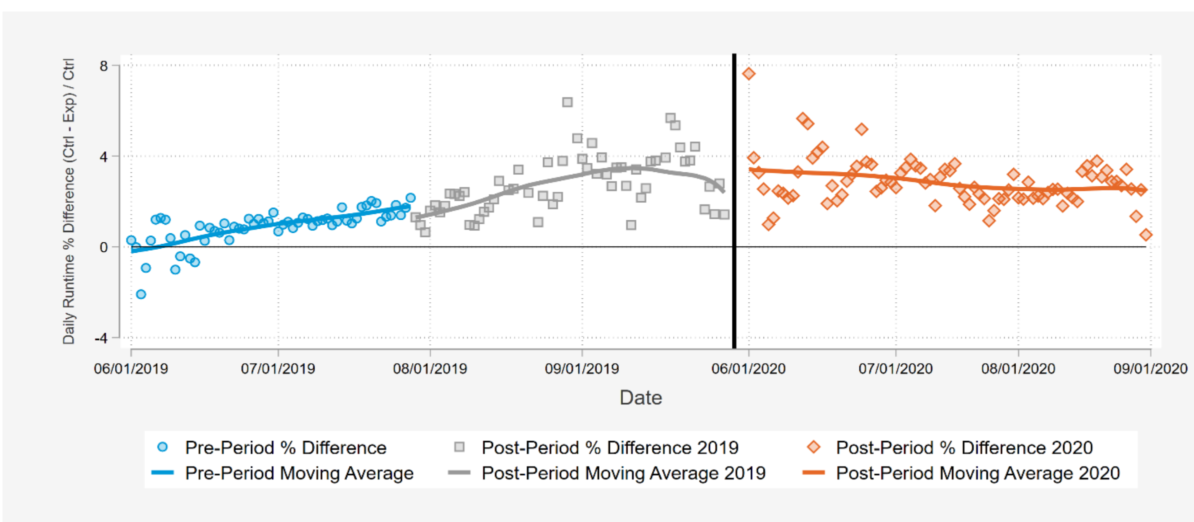

Figure 10: Daily percent difference in cooling runtime – Mixed Humid Eastern zone

Over the course of July and August 2019, devices were slowly onboarding to the EE platform. The impacts gradually increased as more customers accepted the eco+ software and as more features were implemented. By the evaluation period of summer 2020, all features were enabled and devices in the treatment group were more reliably being controlled by the eco+ platform. Results presented for 2020 show the impact of the full EE bundle and provide the results of the more established offering – albeit during a pandemic. All eco+ DR event days were removed from the EE analysis. Any devices with TOU enabled in the EE/DR cells were allowed to remain in the EE analysis since they were also offered EE features.

We tested a variety of regression techniques to model the energy efficiency impacts of the eco+ platform. As expected with a large RED, the results were robust to model specification. All models were implemented with thermostat-level fixed effects and cluster-robust standard errors. The fixed effects approach controls for the time-invariant characteristics of each device. Time-invariant attributes like which cell the device was randomized into are omitted from the regression because there is no variation. This approach is generally well-equipped to net out pre-treatment differences between groups and produces consistent estimates of uncertainty because standard error calculations are based on panel size rather than the number of observations.

Before selecting the final model, the team tested a series of model specifications. The parameter of interest (treatpost coefficient) and its standard error remain relatively stable across model specifications. The models with additional explanatory variables show a modest improvement in standard errors compared to the simple difference-in-differences (DID) model. Figure 11 shows the impact coefficients for each of the seven model specifications across each of the six study regions. We selected model #7, which includes variables for CDH60, CDH build-up, day-of-week, and relative humidity because all terms were highly statistically significant across all six EE/DR regions and have a sound underlying relationship with cooling loads.

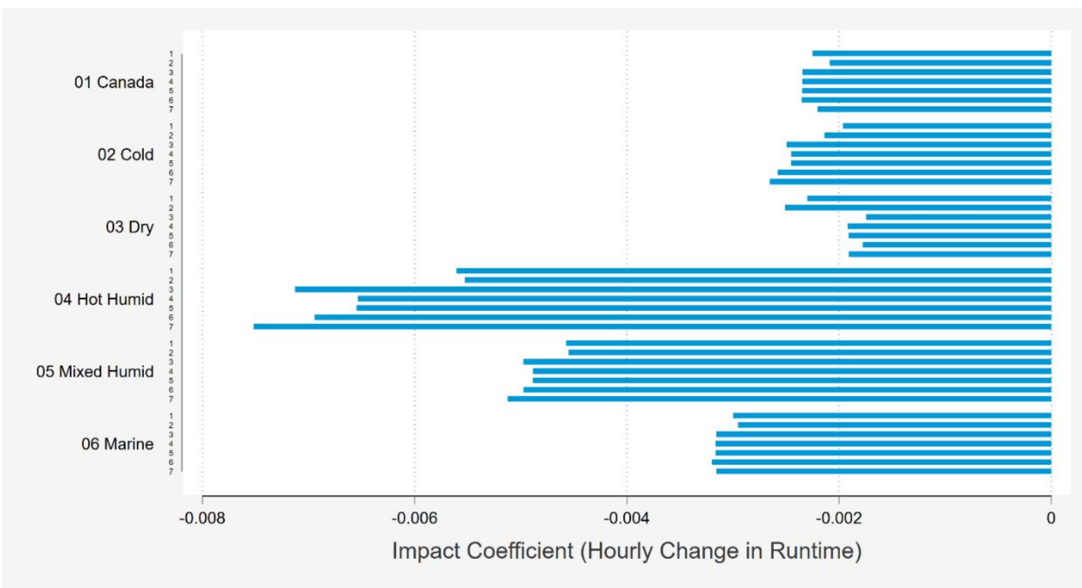

Figure 11: Regression model impact coefficient comparison, by region – summer 2019

The primary EE modeling was done using hourly thermostat runtime data. All savings are relative to an ecobee thermostat without eco+. We also tested daily models where the total daily runtime for each thermostat was modeled as a function of weather and day-of-week. The daily modeling exercise returned very similar impact estimates to the hourly modeling. Table 7 shows the LATE energy efficiency results, by region and month, along with the margin of error at the 95% confidence level for summer 2019. Table 8 shows the results for summer 2020.

Table 7: Summer 2019 LATE energy and demand savings with margin of error at 95% confidence level

Region	August Per-Device kWh	September Per-Device kWh	Total kWh	Peak kW Savings (Weekdays 2-6pm)
01 Canada	19.0 ± 10.5	5.0 ± 11.8	23.9 ± 15.8	0.08
02 Cold/Very Cold	22.2 ± 7.3	16.8 ± 6.5	38.9 ± 9.8	0.07
03 Hot-Dry/Mixed-Dry	17.5 ± 16.2	10.9 ± 17.3	28.5 ± 23.7	0.02
04 Hot Humid	56.3 ± 13.7	59.5 ± 11.6	115.9 ± 18.0	0.08
05 Mixed Humid	33.9 ± 7.1	33.3 ± 6.7	67.2 ± 9.8	0.11
06 Marine	26.6 ± 14.6	15.0 ± 10.4	41.6 ± 17.9	0.06

Table 8: Summer 2020 LATE energy and demand savings with margin of error at 95% confidence level

Region	June Per-Device kWh	July Per-Device kWh	August Per-Device kWh	Total kWh	Peak kW Savings (Weekdays 2-6pm)
01 Canada	13.2 ± 10.4	29.9 ± 15.1	16.1 ± 12.7	59.2 ± 22.3	0.08
02 Cold/Very Cold	16.1 ± 6.3	28.8 ± 8.7	19.8 ± 7.7	64.8 ± 13.2	0.11
03 Hot-Dry/Mixed-Dry	21.7 ± 9.8	36.3 ± 11.7	39.0 ± 12.9	96.9 ± 20.0	0.10
04 Hot Humid	41.4 ± 14.2	38.4 ± 15.8	31.2 ± 16.2	111.0 ± 26.7	0.06
05 Mixed Humid	21.5 ± 7.6	31.3 ± 9.3	21.4 ± 8.5	74.2 ± 14.7	0.11
06 Marine	16.0 ± 7.5	16.9 ± 8.6	22.4 ± 11.2	55.3 ± 16.0	0.08

Not surprisingly, regions that experience the hottest summers and have the most air conditioning load showed the largest per-device energy savings. However, when we consider the impacts on a percent basis in Table 9 we see that hotter regions save less than milder regions.

Table 9: Summer 2020 LATE percent savings by month and region

Region	June Percent Savings	July Percent Savings	August Percent Savings
01 Canada	4.8%	6.0%	4.5%
02 Cold	3.8%	4.2%	3.6%
03 Dry	4.8%	6.1%	5.7%
04 Hot Humid	4.6%	3.5%	2.8%
05 Mixed Humid	3.9%	3.8%	3.2%
06 Marine	12.1%	10.4%	8.7%

Efficiency professionals should always be careful when comparing percent impacts across climate zones or periods with different weather conditions. In extreme conditions, the vast majority of cooling energy is still required. In milder conditions the “trimming” effect appears much larger because air conditioning usage is less extreme. An algorithm like eco+ trims HVAC usage around the edges. Table 10 demonstrates this idea. If we think of AC runtime as an approximately linear function of the temperature differential between outdoor temperature and setpoint (Delta T), raising the setpoint by one degree effectively reduces the differential by one degree and creates the corresponding reduction in runtime. The percent change is larger in more mild conditions (80-degree outdoor temperature) than extreme conditions (100-degree outdoor temperature).

Table 10: Percent impacts in mild and extreme conditions

Outdoor Temp (F)	Original Setpoint (F)	Original Delta T	New Setpoint (F)	New Delta T	Percent Change
100	70	30	71	29	3.33%
80	70	10	71	9	10%

The eco+ EE algorithm follows this conventional wisdom. As expected, the Marine region has the mildest summer weather and the highest percent impacts.

Conclusion

This research demonstrates that optimization algorithms can make a smart thermostat even smarter and personalize savings features to respond to price signals and grid conditions. Smart thermostats are already a recognized energy efficiency measure in most jurisdictions across North America. The energy efficiency results suggest program administrators should consider increased savings assumptions and potentially incentive amounts for smart thermostats with advanced optimization features compared to an "off the shelf" smart thermostat. The demand response and time-of-use results demonstrate that smart thermostats can be an important tool for load flexibility as utilities seek to manage a power system with increasing levels of intermittent renewable generation.