PGE Commercial Smart Thermostat Pilot Evaluation

Final Report

Submitted by Evergreen Economics and Driftless Energy

November 18, 2022

Table of Contents

1 Executive Summary

Portland General Electric (PGE) offers direct install smart thermostats to small and medium business customers as part of a demand response pilot program. In September 2019, Energy Trust of Oregon (Energy Trust) initiated a pilot and began providing incentives for a subset of these smart thermostats to test their potential as a cost-effective energy efficiency measure. Smart thermostats are theorized to save energy in commercial buildings through temperature setbacks, improving fan mode scheduling, and adjusting settings during unoccupied hours.

Methodology

This report presents an evaluation of this pilot program to determine energy and peak demand savings of incentivized thermostats, the characteristics that relate to these savings estimates, and how these savings change over time. To accomplish these objectives, we received and analyzed hourly interval energy electricity usage data and monthly gas usage data from PGE. Our billing analysis provides estimates of the change in electricity consumption (kWh), peak demand (kW), and gas consumption (therms) attributed to the installation of the smart thermostats.

A total of 182 distinct sites (with 410 smart thermostats) had sufficient utility billing data and reliable baseline models (based on our model fit criteria) to be used to estimate energy savings and peak demand impacts. A total of 153 distinct sites (with 318 smart thermostats) had sufficient utility billing data and reliable baseline models to be used to estimate gas savings.

Energy and demand savings estimates presented in this report may be influenced by the impacts of the COVID-19 pandemic. The pandemic led to significant changes in energy usage patterns for many businesses as they adapted to new public health mandates. We created a customized non-routine adjustment for each site, which adjusts the model predictions based on the sites' observed change in energy usage around the start of the pandemic. While this may be sufficient to produce accurate estimates of energy savings, the pandemic is a still a significant source of uncertainty that should not be ignored.

The evaluation team conducted a separate analysis of smart thermostat data from the two thermostat manufacturers (Ecobee and Pelican). This engineering analysis verified that the smart thermostats are operating as expected in terms of scheduling fan modes and temperature setpoint changes. We plan to issue a revised memo with the findings from the engineering analysis, addressing the lingering questions about how the thermostats are used in practice.

Findings

We reviewed a sample of sites that installed the smart thermostats prior to the start of the pandemic (9 sites with gas data and 11 sites with electric data) and found statistically significant savings for both gas and electric energy usage. The electric energy savings during the months leading up to the start of the COVID-19 pandemic align with the *ex-ante* savings estimate in Energy Trust's measure approval document. This provides evidence in support of the current *ex-ante* energy savings.

Across the full sample of sites in the billing analysis (n=182 for electric, and n=153 for gas), smart thermostats led to statistically significant reductions in energy usage across all days of the week. Additional findings include:

- Average per-thermostat electric savings of 13.7 percent (3,847 kWh per year);[1]
- Average per-thermostat gas savings of 11.8 percent (165 therms per year);
- Peak demand savings of 0.93 kW per thermostat (26.7%) at 8 p.m. on the coldest days of the year;
- Especially high savings during off-peak hours in both summer (0.54 kWh per thermostat, 20.6%) and winter (0.55 kWh per thermostat, 22.9%); and
- Higher energy savings for certain building types (medium offices and schools), and for Pelican smart thermostats (versus Ecobee).

While energy savings were correlated with various characteristics, significant energy savings of more than 10 percent of baseline energy usage were detected at around half of the sites (54% of sites for electric, 48% of sites for gas). We are confident that the smart thermostats are saving at least as much electricity as expected in in the measure approval document (*ex-ante* savings), though gas savings were less consistent.

[1] These energy savings estimates are based on analysis of site-level utility energy meters, divided by the number of thermostats installed at the site as an estimate of change in energy usage per thermostat.

Memo

Energy Trust of Oregon

To: Board of Directors

From: Sarah Castor, Evaluation & Engineering Manager
Wendy Gibson, Sr. Program Manager – Commercial
Jackie Goss, Sr. Engineer

cc:

Date: February 23, 2023

Re: Staff Response to the PGE Commercial Smart Thermostat Pilot Evaluation

Energy Trust's joint effort with PGE to install and test commercial smart thermostats began in 2019 as what Energy Trust now calls a coordinated research project, which is a test of a new technology or delivery method before expanding to a standard offer. This pilot evaluation was intended to inform Energy Trust's energy savings estimates for commercial smart thermostats and provide an understanding of how savings vary by thermostat manufacturer, settings, building and HVAC system characteristics and business type. PGE also used the pilot to test the demand response capabilities of these devices.

The pilot produced many smart thermostat installations between 2019 and early 2021, a period that was heavily influenced by the COVID-19 pandemic and related business closures, both temporary and permanent. The evaluated energy savings per thermostat were higher than expected, and sizable at around 14% of electric use and 12% of gas use. The evaluation was not designed to include a control or comparison group of similar customers and buildings that did not receive a smart thermostat, which makes it difficult to be certain how much of the evaluated savings are due to the smart thermostat and not related to reduced business operations during the pandemic. Results were also not sufficient to understand how the thermostats performed differently for various building/business types and sizes, thermostat settings (both before and after installation) or by thermostat manufacturer, given that there were relatively few Pelican thermostats in the evaluated projects. We believe more study of commercial smart thermostats is needed to confidently estimate savings from commercial smart thermostats.

Energy Trust is planning to conduct another evaluation in 2024 to assess savings based on thermostat performance in 2022 and 2023, after the more significant impacts of the pandemic had passed and with a larger group of participants. We will also use different analysis method to estimate savings, likely with a matched comparison group of nonparticipants or future participants. The Existing Buildings program will continue to offer the measure pending results of the evaluation in 2024.

2 Introduction & Background

Portland General Electric (PGE) began offering direct install smart thermostats to small and medium business customers as part of a demand response pilot program. In September 2019, Energy Trust of Oregon (Energy Trust) initiated a pilot and began providing incentives for a subset of these smart thermostats to test their potential as a cost-effective energy efficiency measure. Smart thermostats are theorized to save energy in commercial buildings through temperature setbacks, improving fan mode scheduling, and adjusting settings during unoccupied hours.

Study Objectives/Evaluation Goals

This study was designed to produce rigorous estimates of energy and peak demand savings to assess the value of smart thermostats as energy efficiency measures for commercial customers on non-event days. PGE has commissioned a separate evaluation to assess the demand response functionality of these devices and event participation by individual customers.

Table 1 provides a link between the research questions posed by Energy Trust and the evaluation tasks. The billing analysis provides estimated energy and demand savings realized at the meter, including variability in savings across and within sites. The engineering analysis dives deeper into the device schedules and temperature setpoints to help explain why the savings were (or were not) achieved.

Table 1: Map of Research Questions to Evaluation Tasks

	Evaluation Task	
Research Question	Billing Analysis	Engineering Analysis
What are the overall energy and demand savings of commercial smart thermostats?	✓	
What are the distributions of energy and demand savings by major bins (e.g., weekday afternoons in the winter)?	✓	
What are the trends in energy and demand savings results over time?	✓	✓
What are the energy and demand savings impacts by thermostat manufacturer, thermostat settings, building characteristics (i.e., HVAC capacity, floor area, percent conditioned space), and business type?	✓	✓

A separate memo will be issued with the methods and findings of the engineering analysis for each thermostat manufacturer, Pelican and Ecobee.

3 Methods

This section describes the Evergreen team's methodologies for assessing the impact of the commercial smart thermostats installed by Portland General Electric (PGE) that were also qualified for the Energy Trust study.[2] The Evergreen team conducted impact analysis utilizing the project data sources (e.g., PGE pilot participant audit data, utility billing data). The program data were used to summarize participation across thermostats by manufacturer, building and customer characteristics, incentives provided, and savings claimed. The billing analysis provides estimates of the change in kWh, kW, and therms attributed to the installation of the smart thermostats. The engineering analysis verifies that the smart thermostats are operating as expected, in terms of scheduling fan modes and temperature setbacks.

3.1 Data Sources

Table 2 provides a summary of every data source we utilized for this evaluation, fields provided, sample coverage (e.g., number of premises and date range), and how the data were used.

After receiving each data source, we conducted data quality checks before cleaning and preparing the data for analysis (e.g., flagging outliers, identifying and addressing missing values). Evergreen conducted all preparation and analysis using open-source software languages (R and PostgreSQL) to ensure reproducible results. All our code was tracked with version control software in a GitHub repository to provide a complete link between the raw data and the completed analysis presented in this report. We will deliver a copy of the final analysis dataset at the end of the evaluation.

[2] The Evergreen team includes Evergreen Economics and Driftless Energy.

Table 2: Data Sources for the Evaluation

Data Source	Unique Fields	Coverage	Intended Use
PGE Pilot Participant Audit Data	Customer and premise, device serial number, HVAC system details, existing thermostat type, installation date	n=593 premises Aug 2018-Feb 2021	Link between data sources (e.g., AMI, application programming interface [API], project tracking data); identify units eligible for Energy Trust incentives, define the thermostat installation date
Advanced Metering Infrastructure (AMI) Hourly Electric and Monthly Gas Billing Data	Hourly electricity and monthly gas consumption	n=591 premises[3] Oct 2018-May 2022	Billing analysis; estimates of energy and demand savings
Demand Response Events	Identify date and time of each Schedule 25 demand response event	Jun 2020-May 2022	Flag usage during demand response events in the AMI and API data
Measure Approval Document (MAD)	Cost effectiveness calculator, requirements, and measure analysis	Version 235.1 Valid May 2019-Dec 2020	Define *ex-ante* savings (kWh) by HVAC type, business type, and heating zone
Disqualified Thermostats List	Premise ID, install date, reason for disqualification	n=95 premises as of 8/6/21	Identify all thermostats that do not qualify for Energy Trust incentives
Solar Table	Premise ID, rate code, and net energy meter (NEM) flag	n=8 premises	Identify sites with solar to exclude from the billing analysis
Energy Partner Customer List	Premise ID, heating system type, and floor area	n=100,975 premises	Participant characteristics, for use in segmentation analysis
Project Tracking (PT)	Premise ID, product code, *ex-ante* savings, useful life	n=325 premises Sept 2004-May 2021	Identify other measures installed during the study period
National Oceanic and Atmospheric Administration (NOAA) Weather	Hourly interval outdoor air temperature	n=31 stations Jan 2018-May 2022	Weather normalization (actual weather)
Typical Meteorological Year (TMY3) Weather	Typical weather conditions, based on historical outdoor air temperature	n=1,020 stations in the NW	Weather normalization (typical weather)

3.2 Exclusion Criteria

Thermostats Qualified for Energy Trust Incentives

Our analysis was limited to the subset of thermostats incentivized by PGE that also met Energy Trust's eligibility criteria for the energy efficiency portion of the pilot. The following rules were used to *disqualify* thermostats from the Energy Trust portion of the pilot program:

- New construction (i.e., must be an existing building);
- Commercial floor area above 200,000 square feet (excluding large facilities with centralized control systems);
- Lodging with 24/7 operation (for the lack of savings opportunities);
- Semi-conditioned spaces;
- Heating Zone 2;
- Not enrolled in the PGE Schedule 25 thermostat program;
- Not one of the two approved thermostats (Pelican TS-200 and Ecobee EMS-SI); and/or
- Does not satisfy all of the installation requirements (single ducted HVAC system, 20 or fewer heating or cooling zones with independent controls, smart thermostats control at least 50 percent of these zones, etc.)

Billing Analysis Sample Attrition

Evergreen implemented a series of exclusion criteria for participants with insufficient billing data or concurrent changes at the facility (e.g., other measures) that would prevent us from deriving a reliable estimate of energy savings attributable to the smart thermostat. Specifically, we **excluded participants** that:

- Had less than 12 months of pre- or 9 months of post-installation data (which must include a full heating and cooling season and part of the shoulder seasons);
- Had account turnover during the study period;
- Were outliers in annual energy consumption (top and bottom 1 percent of sites);[4]
- Had net-metered solar PV system present;
- Installed other measures funded by Energy Trust during the study period above a minimum savings threshold of half of the *ex-ante* savings estimated for the smart thermostat (e.g., 940

[3] A premise is the equivalent of an electric account. A single site may have multiple premises that are associated with one or more account holders. Each premise is associated with one or more smart thermostats.

[4] All of the sites that qualified as outliers in energy consumption failed multiple exclusion criteria.

kWh for a non-grocery commercial building with resistance heat and cooling that installed two smart thermostats[5]);

- Poor model fit (i.e., less reliable for meter-based savings); or
- Had two or more statistically significant non-routine events (NREs) (i.e., a major shift in energy consumption not associated with the program intervention) that spanned more than 10 percent of the baseline or post-period, which would negatively impact the reliability of savings estimates derived from billing analysis.

We started with 588 sites with smart thermostats Table 3 provides a summary of participant attrition during the first phase of the analysis. Significant attrition was caused by installations that did not qualify for Energy Trust incentives (n=94 sites and 446 thermostats); leaving 413 sites that qualified for incentives. Another 75 sites (with 213 thermostats) were removed because they installed other energy efficiency measures during the study period (a two-year span) that had similar or higher *ex-ante* savings, making it too difficult to isolate the impact of the smart thermostats. Lastly, there were 24 sites that installed the thermostat so close to the start of the COVID-19 pandemic that it would not be possible to ascertain whether a change in energy usage was caused by the thermostat installation, the pandemic, or a combination of the two.

Table 3: Participant Attrition from Billing Analysis by Exclusion Criteria, Phase 1

Exclusion Criteria		Sites Dropped	Remaining Sites
Site Received a Smart Thermostat from PGE		-	588
Do Not Qualify (DNQ) for Energy Trust Incentives	Listed as DNQ by Energy Trust	94	494
	Commercial floor area above 200,000 square feet	0	494
	Not one of the two approved thermostat models	0	494
	Lodging with 24/7 operations, semi-conditioned space, or located in Heating Zone 2	52	442
	Not enrolled in the PGE Schedule 25 T-stat program	29	413
Unsuitable for Billing Analysis (regardless of fuel)	Concurrent installation of other measures incentivized by Energy Trust	75	338
	Installed the thermostat near the start of the COVID-19 pandemic (non-routine adjustment not feasible)	24	314

[5] This example is based on the kWh savings listed in the *Measure Approval Document for Direct Install Commercial Smart Thermostat PGE Pilot* as of September 2020.

Table 4 continues the attrition process, looking now at the data for each fuel type separately. With the updated utility billing data we received, a much smaller number of thermostats were removed for having less than 12 months pre- or 9 months of post-installation data at the time of the evaluation (n=6 electric and n=118 gas). Most of the gas sites that were removed for insufficient gas data do not have any gas service (i.e., no gas data were expected).

Another 101 electric and 23 gas sites were dropped from the measure analysis because the site-level billing model did not pass our accuracy standards. This occurs more often in commercial and industrial buildings where energy usage is driven more by business operations than external factors that we can measure and control for in the regression (weather, hours of daylight). These sites are difficult to predict with meter-based savings and may be a better fit for building simulations, engineering analysis, or metering. See Section 3.3.1 Baseline Model Fit for more information on the model accuracy criteria.

We adjusted for the impact of six different types of statistically significant non-routine events. We excluded 24 electric and 20 gas sites from the analysis because they had multiple statistically significant NREs, introducing too much uncertainty into the savings estimates. See Section 3.2.2 Non-Routine Events for more information on how NREs were identified and adjusted.

Table 4: Participant Attrition from Billing Analysis by Exclusion Criteria, Phase 2

<table>
<tr><th colspan="2" rowspan="2">Exclusion Criteria</th><th colspan="2">Electric</th><th colspan="2">Gas</th></tr>
<tr><th>Sites Dropped</th><th>Remaining Sites</th><th>Sites Dropped</th><th>Remaining Sites</th></tr>
<tr><td colspan="2">Received a smart thermostat, qualified for Energy Trust incentives, and suitable for analysis (Phase 1)</td><td>-</td><td>314</td><td>-</td><td>314</td></tr>
<tr><td rowspan="4">Unsuitable for Billing Analysis (by Fuel)</td><td>Had less than 12 months of pre- or 9 months of post-installation electric/gas data*</td><td>6</td><td>308</td><td>118</td><td>196</td></tr>
<tr><td>Had net-metered solar PV system present[6]</td><td>1</td><td>307</td><td></td><td>196</td></tr>
<tr><td>Model did not pass accuracy standards (i.e., poor explanatory power)</td><td>101</td><td>206</td><td>23</td><td>173</td></tr>
<tr><td>Have multiple statistically significant NREs (2+ major changes in the post-period)</td><td>24</td><td>182</td><td>20</td><td>153</td></tr>
</table>

* Most of the gas sites that were removed for insufficient gas data do not have any gas service (i.e., no gas data were expected).

[6] Four additional sites had solar PV present, but the sites were dropped in the previous step for failing a different exclusion criterion.

In addition to site exclusions, we **excluded individual observations of days**:

- With demand response events (as well as the day prior, if customers were notified the day ahead of events) from the models that are used to estimate energy efficiency savings, to avoid double-counting; and
- At the start and end of daylight savings (two days per year).

3.2.1 Outliers

Evergreen also identified outliers in energy consumption or those with unusual energy consumption patterns during the study period. To start, we defined an outlier as any kWh reading that was more than three times the distance of the interquartile range (IQR) from the median interval measurement, based on the full time span of hourly interval AMI data for each site.[7] Next, we manually reviewed all the flagged outliers with time-series plots and then adjusted any flags that appeared to be too sensitive (false positives) or not sensitive enough (false negatives) by site.

Finally, we estimated the baseline models with and without these flagged outliers to assess the relative baseline model fit and determine if they had a statistically significant impact on the estimated energy savings. Based on this process, we found that inclusion (or exclusion) of outliers had **no significant impacts on results**. Outlier values have been retained throughout our analysis.

3.2.2 Non-Routine Events

We used model residual cumulative sum plots to identify NREs at each participant site. Model residuals are the difference between the actual energy use and the model prediction. The cumulative sum of residuals should be randomly distributed around zero during the baseline period and then negative sloping during the post-period, under the assumption that savings are accumulating. A change in the slope is evidence of an event that is not captured in the model (i.e., a change other than weather, number of daylight hours, etc.). This option is useful at identifying the presence of an NRE as well as the start and end dates and the magnitude of the event. We have categorized NREs on the basis of three criteria:

1. **Cause** – Whether the NRE started around the beginning of the COVID-19 pandemic;
2. **Duration** – Whether the NRE was temporary or lasted until the completion of the study; and
3. **Period** – Whether the NRE started in the baseline (prior to the thermostat install) or post-period.

---

[7] This definition of an outlier is based on CalTRACK rule 2.3.6. The IQR is a measurement of variability. The rank-ordered data are divided into four equal parts called quartiles. The IQR measures the distance between the first and third quartiles, corresponding to the 25th and 75th percentiles, containing the middle 50 percent of observations.

Figure 1 shows an example of model residuals for a single customer during the post-period (black line). We have added trendlines to emphasize the average slope across the full period (blue line), the slope pre-COVID (green line), and the slope during the COVID-19 pandemic (red line). Customer 032 exhibited a prominent shift in its savings slope, with a steepening of the slope around March 2020 at the start of the COVID-19 pandemic (i.e., more of a reduction in energy usage than we observed prior to the start of the pandemic). This steepening continued from the start of the pandemic through the end of the post-period. This plot confirmed that the COVID-19 pandemic had a significant impact on energy usage.

Figure 1: Example of COVID-19 Pandemic NRE Detection for Customer 032

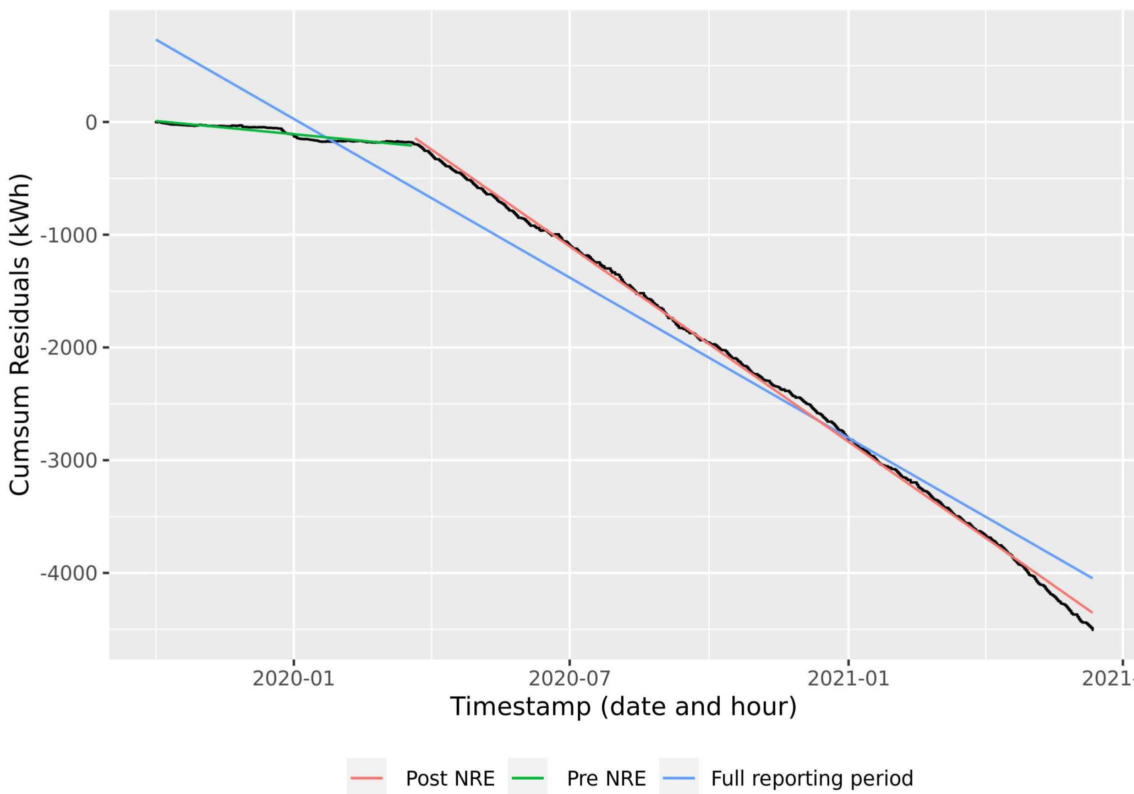

We identified some NREs that did not appear to be associated with the COVID-19 pandemic. Figure 2 shows an example of model residuals for a different customer (Customer 086). Customer 086 exhibited a prominent shift in its savings slope, with a sudden flattening of the slope around February 2021, well after the start of the COVID-19 pandemic (i.e., less of a reduction in energy usage than we observed prior to the NRE). This flattening continued from February 2021 through the end of the post-period. This plot confirmed that this non-COVID-19 NRE had a significant impact on energy usage.

Figure 2: Example of Non-COVID NRE Detection for Customer 086

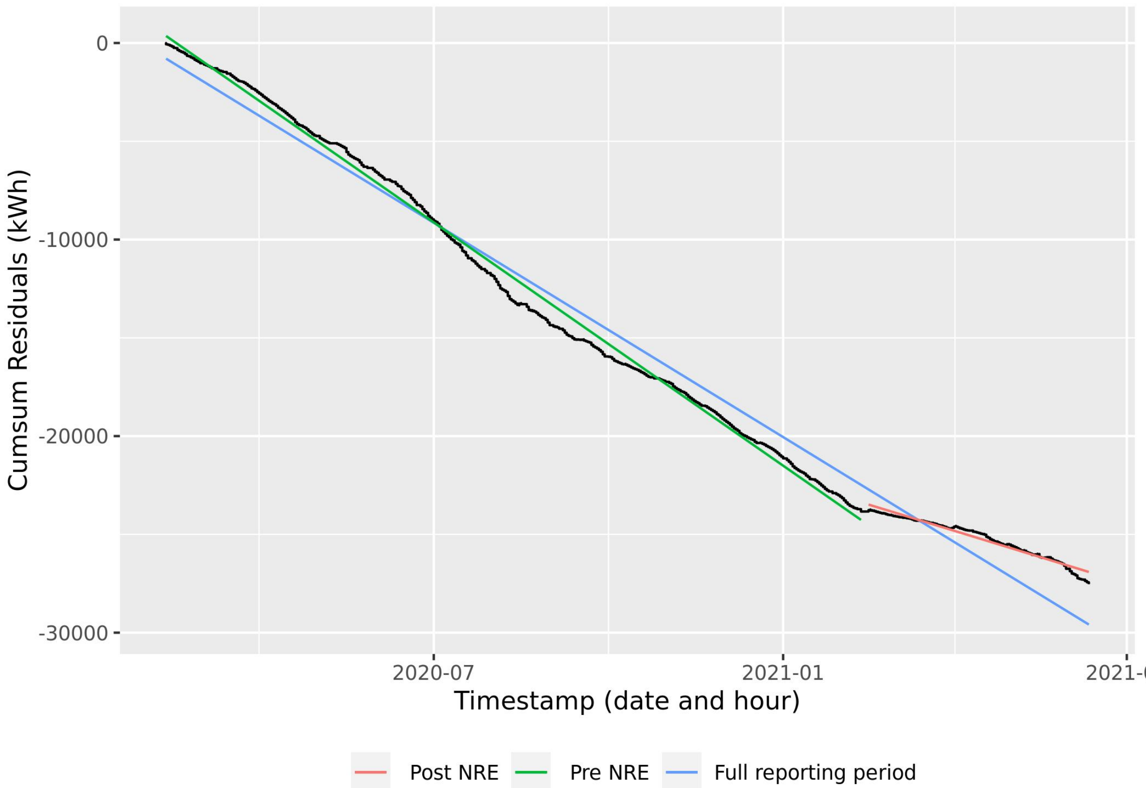

During this process, we identified

- 56 gas sites and 144 electric sites with a **single** significant NRE in the post-period for adjustment; and
- 20 gas sites and 24 electric sites with **multiple** NREs in the post-period, which were excluded from the analysis.[8]

For the subset of sites with a single NRE, we used a second regression model to estimate the impact of the NRE (whether or not it was related to the COVID-19 pandemic) on the energy usage of each site as shown in Equation 1. The purpose of this model is to estimate the incremental impact of NREs on the residuals in the post-period (i.e., the estimated change in energy usage after the program implementation). The interaction term captures any increase or decrease in the impact of the NREs over time. The statistical significance of the NRE coefficients tells us whether an NRE adjustment is necessary, and the value of the coefficients is our best estimate for the

[8] These sites were shown in the last line of Table 4: Participant Attrition from Billing Analysis by Exclusion Criteria, Phase 2.

impact of the NRE. This adjustment model was estimated separately for each participant and each hour of the day.

Equation 1: Non-Routine Adjustment Ordinary Least Squares (OLS) Regression Model

$$Residual_{i,t} = \beta_0 + \beta_{Ti}Timestamp_{i,t} + \beta_{Ci}NRE_t + \beta_{TCi}Timestamp * NRE_{i,t} + \varepsilon_{i,t}$$

Where:

$Residual_{i,t}$ = Difference between the model prediction and actual energy usage, for customer *i* during time interval *t* during the post-period

$Timestamp$ = Post-period time interval t

NRE = Dummy variable (0, 1) representing the period during the NRE

β_{Ci}= Average impact of the NRE on the energy usage of customer *i*

β_{TCi}= Average incremental impact of the NRE on the energy usage of customer *I* for each additional time interval *t*

β_{0i} , β_{1i} …= Coefficients estimated by the model for customer *i*

ε = Random error assumed to be normally distributed

We adjusted for the impact of one of six different types of NREs at each site where we found a statistically significant change in the model residuals for a single NRE. Table 5 shows each of the six types of NREs for which we adjusted, the number of sites affected, the percentage of the period that was affected, and the impact on estimated savings for kWh.

Table 6 shows the same information for gas savings. The second line shows that we identified and adjusted for temporary NREs, events that resolved on their own within the period, at 52 sites. These NREs spanned 34 percent of the period on average, or around 4 months for a site with a full 12-month baseline period. Before adjusting for this NRE, we started with a savings estimate of 9.3 percent, but this increased to 14.2 percent when we adjusted for the NRE. This happens when an NRE makes the baseline energy usage artificially high; once you correct for that temporary change in baseline usage during the event, you see a wider gap between baseline and post-period usage (greater savings). NREs can work in either direction, increasing or decreasing the estimated savings. Overall, our adjustments for COVID-19 NREs tended to decrease estimated savings while adjustments for non-COVID-19 NREs tended to increase savings (as shown with Customer 086 in Figure 2). While COVID-19 was an expected source of NREs (n=56 electric sites and n=13 gas sites), NREs of unknown cause (n=88 electric sites and n=43 gas sites) were the most common.

Table 5: Impact of Electric NRE Adjustments on Estimated Savings (kWh)

Cause	Duration	Period	Sites	% of Post During/After NRE	Original Savings as % of Baseline	Adjusted Savings as % of Baseline
N/A	N/A (no NRE)	N/A	38	N/A	12.5%	12.5%
Unknown	Temporary	Either	52	34%	9.3%	14.2%
	Lasting	Either	36	32%	7.2%	9.9%
COVID-19	Temporary	Baseline	34	38%	14.1%	23.8%
		Post	14	52%	36.3%	13.3%
	Lasting	Baseline	5	61%	9.4%	18.3%
		Post	3	74%	30.9%	-3.2%

Table 6: Impact of Gas NRE Adjustments on Estimated Savings (therms)

Cause	Duration	Period	Sites	% of Post During/After NRE	Original Savings as % of Baseline	Adjusted Savings as % of Baseline
N/A	N/A (no NRE)	N/A	97	N/A	8.7%	8.7%
Unknown	Temporary	Either	16	26%	-3.0%	8.7%
	Lasting	Either	27	11%	15.1%	29.5%
COVID-19	Temporary	Baseline	6	45%	8.3%	25.5%
		Post	3	41%	40.6%	19.1%
	Lasting	Baseline	3	55%	-20.8%	-91.9%
		Post	1	79%	54.6%	9.3%

3.3 Billing Analysis

We used the AMI Customer Segmentation (AMICS) modeling approach to estimate pilot impacts on both electric and gas energy usage. The AMICS approach estimates a separate usage profile for each service account (i.e., distinct customer and premise) or customer segment by season and day type while controlling for weather conditions and other differences across days (e.g., day-of-week, hours of daylight). We also used the AMICS modeling approach to estimate electric load shapes.

The AMICS approach uses segmentation to produce a portfolio of load shapes and then compares each day in the post-period against similar days in the baseline, as shown in Figure 3. A key benefit of the AMICS model is avoiding over-reliance on 'average day' conditions.

Figure 3: AMICS Modeling Approach

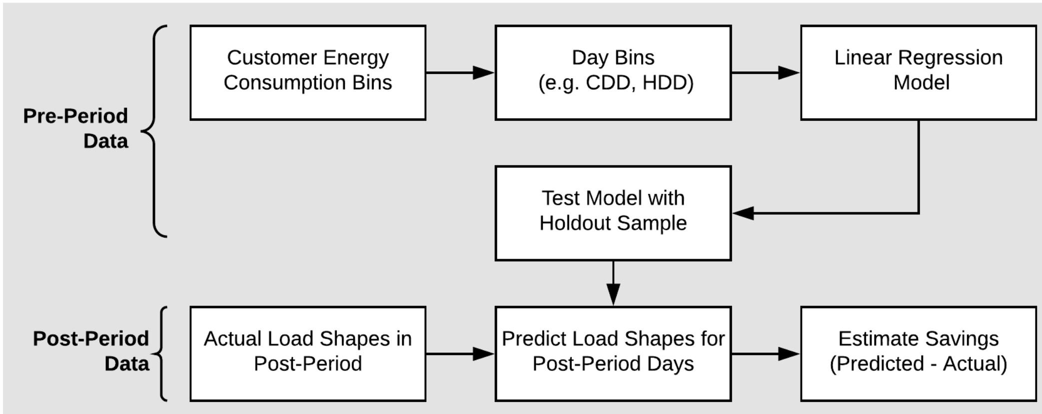

Each day of the study period is binned in terms of its weather, day type, and season.[9] Segmenting days by attributes such as their cooling degree-days (CDDs) explicitly incorporates these elements into our model, controlling for differences in energy usage across days. Each customer is assigned to a single bin (for a site-specific model), but because weather and day type change throughout the year, each customer has days that are assigned to many different bins.

Once the data are segmented, the AMICS model estimates an ordinary least squares (OLS) regression model for each customer and each day bin (weather and day type combination) that has a single dummy variable for each hour of the day as shown below:

Equation 2: Segmented OLS Regression Model

$$Usage_{i,t} = \beta_{H0,i}H00_{i,t} + \beta_{H1,i}H01_{i,t} + \beta_{H2,i}H02_{i,t} + \cdots + \beta_{H23,i}H23_{i,t}\varepsilon_{i,t}$$

Where:

$Usage_{i,t}$ = Energy consumption for a customer *i* during time interval *t*

$H00, H01, \ldots$ = Array of indicator variables (0,1) representing the hour of the day

$\beta_{H0,i}$, $\beta_{H1,i}$ …= Coefficients estimated by the model for customers *i*

ε = Random error assumed to be normally distributed

Most other methods provide one annualized savings number across all participants. The regression modeling approach employed by the AMICS model estimates a full unique set of slope coefficient

[9] The weather bins are created by calculating cooling degree-hours (CDH) for each hourly observation using a base temperature of 65 degrees Fahrenheit, and then taking the average of these hourly values to create a single cooling degree-day (CDD) value for each customer on each day (i.e., each "customer-day") in the study period. This process is repeated to assign these same days to heating degree-day (HDD) bins.

estimates for each customer segment for each day bin (weather and day type).[10] When applied to an entire program, the AMICS model provides separate load shapes (and thus separate savings estimates) for each customer segment, which makes it a useful tool for targeting. Binning the data and then estimating separate regression models for each bin enables the overall model to control for a greater amount of the variation across both customers and weather conditions. This is not a proprietary "black box" method, but rather a series of simple linear regressions that are estimated with open-source statistical software (R and PostgreSQL).[11]

These site-level models use the AMICS approach to explain the change in energy usage at each of the participant sites by time of day. We tested hundreds of model specifications that included wide-ranging controls for:

- Cooling degree-days (CDDs), with a base temperature of 65 degrees Fahrenheit;
- Heating degree-days (HDDs), with a base temperature of 65 degrees Fahrenheit;
- Hours of daylight, measured from dusk to dawn;
- Season (based on calendar season); and
- Day type (i.e., separating weekdays from weekends).

Additionally, we also tested a model variation for which pre-period data were restricted to one year prior to installation. We tested a variety of models, selecting from the controls listed above. For each site and fuel type, we selected the model that most accurately predicted pre-period energy usage as determined by the R-square, among models with coefficient of variation of the root mean square error [CV(RMSE)] less than 25 percent—while also ensuring at least 90 percent coverage of days in the reporting period.

The load shape estimates produced by the AMICS modeling approach provide hourly kWh energy usage; the segmentation enables us to calculate annualized, seasonal, and even winter weekday

---

[10] Application of the AMICS model to evaluate a large program often results in thousands of distinct sets of regression coefficients. However, this process is relatively fast and straightforward, as Evergreen has developed a custom R software package to automate each AMICS query and OLS regression from an integrated SQL database.

[11] The AMICS approach was extensively tested on residential HVAC programs in Phase I of the AMI Billing Regression Study. The Phase II study expanded this research to include a variety of commercial programs. The relative accuracy of AMICS estimates for both site- and segment-level models (as proposed in this analysis) were demonstrated in the 2019 NMEC Pre-Qualification Pilot Feasibility Study, conducted by Evergreen Economics under the supervision and guidance of the Emerging Products team at Southern California Edison (SCE).

We conducted a separate analysis of site level commercial HVAC savings for SCE in 2018 to demonstrate that the AMICS approach can be applied to individual commercial buildings and not just groups of program participants. This study included a side-by-side comparison and cross validation exercise that found no significant difference in prediction error between AMICS and the Temperature and Time of Week (TTOW) modeling approach developed by the Lawrence Berkeley National Laboratory (LBNL). Both AMICS and TTOW struggled with the same sites and the same days; switching from AMICS to TTOW is unlikely to have any impact on site attrition.

on-peak hourly load predictions and corresponding savings estimates. AMICS was also used to provide daily gas energy usage for annualized and seasonal gas savings. We present the savings estimates for individual customer segments (including thermostat, building, and customer characteristics) to identify scenarios with higher and lower than average energy savings.

3.3.1 Baseline Model Fit

We assessed the baseline model fit with a series of goodness-of-fit metrics:

- Normalized mean bias error (NMBE);
- Coefficient of variation of the root mean square error, CV(RMSE);
- Fractional savings uncertainty (FSU) (bias corrected); and
- Coefficient of determination, R-square.[12]

Figure 4 provides a visualization of these four error metrics across the 307 sites with thermostats qualified for Energy Trust incentives and sufficient billing data for the analysis (before NRE exclusions), while Figure 5 shows the 196 gas sites with sufficient data. Sites that fell outside the blue shaded area were considered too weak to provide reliable savings estimates. A high FSU means that it will be difficult for these models to reliably detect savings in the same magnitude as the *ex-ante savings*. However, if the savings estimate exceeds *ex-ante* (e.g., the customer saved more than expected), we would still see a statistically significant change in energy usage. We did not set an FSU threshold for this program because we were consistently seeing savings estimates in excess of *ex-ante*. Sites that fell in the blue shaded areas had extremely variable energy usage in the baseline that could not be sufficiently explained by outdoor air temperature and seasonality (n=101 for electric and n=23 for gas).[13]

[12] Low R-square values suggest that additional independent variables should be tested. If no additional variables are feasible (due to limitations in the data), then a low R-square value is acceptable.
[13] These sites were shown in the second to last line of Table 4.

Figure 4: Electric Baseline Model Fit

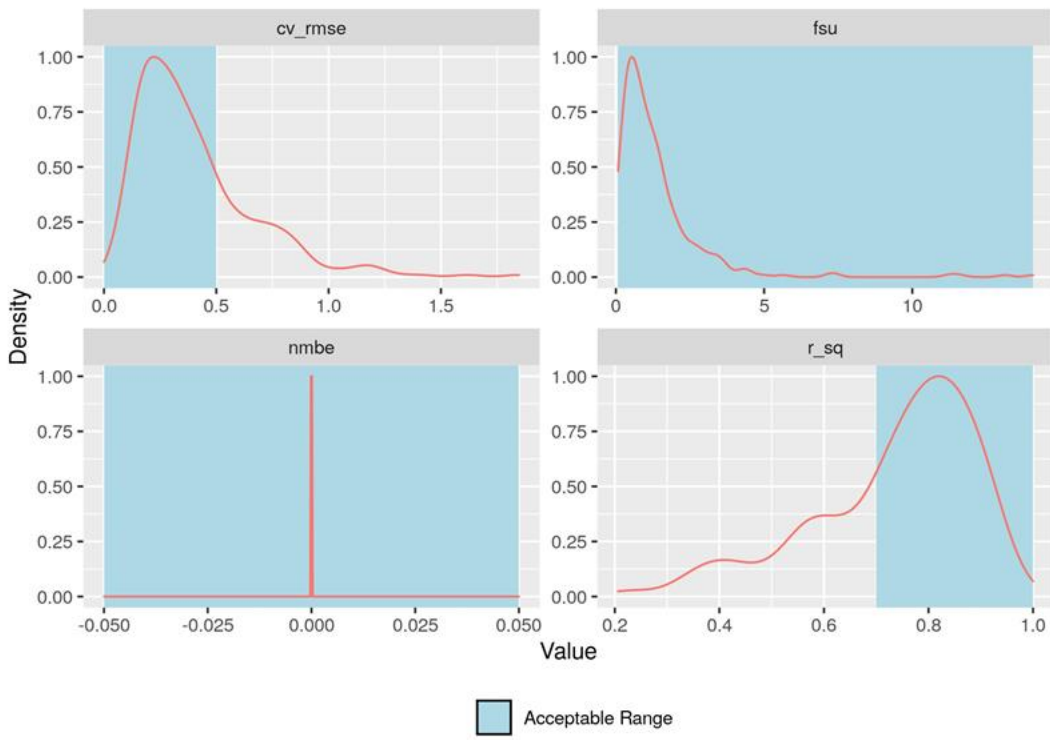

Figure 5: Gas Baseline Model Fit

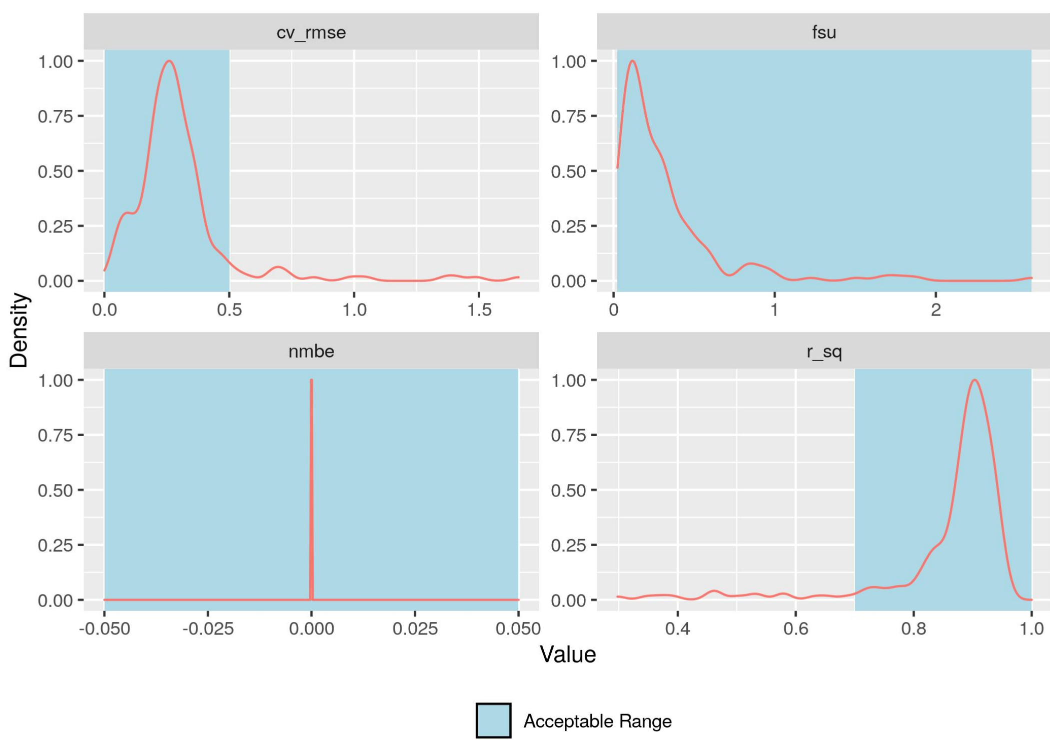

We reviewed the characteristics of sites that were dropped due to poor model fit, looking for any patterns that may have impacted the analysis findings. While thermostat type had no direct impact on attrition, our comparison of the brands could feasibly be affected if there is greater attrition in one than the other. Of the 101 electric sites that were dropped, 99 had Ecobee thermostats installed (98%). This is slightly higher than the proportion we have retained in the sample for electric analysis, which includes 173 sites with Ecobee smart thermostats (95%) and 9 with Pelican smart thermostats (5%). However, the thermostat brand is related to other factors, such as business type, where smaller businesses like small office (n=41 dropped sites), small retail (n=14 dropped sites), and non-hotel hospitality (n=10 dropped sites) are often installing the Ecobee.

While a somewhat disproportionate number of Pelican smart thermostats were dropped from the gas analysis (29% of Pelican sites versus 11% of Ecobee sites) again, this is correlated with other factors, namely HVAC heating capacity; seven of the 23 dropped sites had either very large systems (greater than 800,000 BTU) or no data on heating system size.

3.3.2 Computing Standard Errors

In the AMICS approach, we estimate individual regression models for thousands of customer-day segments, providing a kWh energy usage prediction for each hour.

Because the AMICS model is estimated using data from the baseline year (12 months prior to thermostat installation), we compute the relative variance for each hour of the day for each customer-day bin as the ratio of the variance to predicted hourly kWh usage. These relative variances are then applied to the post-period data to create confidence intervals for the model predictions of each hour of each customer-day in the post-period. With 24 hours per day and thousands of customer-day segments, we compute over 24,000 confidence intervals. For aggregated predictions, such as the annual and seasonal post-period load shapes, we use bootstrapping to estimate the relative variance for each hour, accounting for variation in the number of observations and relative kWh represented by each customer-day bin.

A bootstrap draws a series of random samples with replacement from the empirical distribution of values and then selects the 2.5 and 97.5 values ($\alpha/2$ and 1-$\alpha/2$) as the lower and upper values of the 95 percent confidence interval (1-α). In situations in which the empirical distribution of data is skewed and bounded (cannot fall below zero), bootstrap confidence intervals have been shown to be asymptotically more accurate than standard percentile-based methods, while retaining the desirable property of robustness. This approach adjusts for both bias and skewness in the distribution by estimating the density from the observed data, rather than assuming the data conform to a known parametric distribution. We are therefore confident that the confidence intervals we developed using the bootstrap method are at least as good (if not superior) in performance to standard percentiles. Specifically, we used the bootstrap method developed by

Bradley Efron.[14]

As with any confidence interval estimated from a small sample size, there is a potential for overstating confidence (i.e., estimating unrealistically tight error bounds) when the sample measurements are very similar by random chance. Larger samples improve the accuracy of both the load shape estimates and error bounds—the confidence intervals will not necessarily get tighter, but they will provide a more realistic estimate of the true variability across sites.

3.3.3 Energy Savings Estimation

Energy savings are estimated by comparing the model predictions for each customer and day bin to their actual energy usage on all similar days in that bin after the thermostat installations, as shown in Equation 3. We are comparing energy usage in the post-period to what the same customer used in the baseline period.

Equation 3: Energy Savings

$$Savings_{i,t} = \widehat{Predicted}_{i,t} - ActualPost_{i,t}$$

Where:

$\widehat{Predicted}_{i,t}$ = Energy consumption predicted for customer *i* for days in bin t

$ActualPost_{i,t}$ = Actual energy consumption for customer *i* on days in bin t

The predictions are based on the estimated regression coefficients. Depending on the analysis, we can aggregate the predictions based on normalized weather conditions from the typical meteorological year to produce our estimates for normalized electricity usage and savings (kWh or therms) by season, as shown in Equation 4. In this example, the normalized summer energy consumption of customer *i* is the sum of three days from day bin “112”, seven days from bin “214”, and so on. Each day bin has a set range of CDD, HDD, hours of daylight, and day type based on our model specification. We identify the distribution of days in the normalized weather year using the same specifications, and calculate the number expected from each day bin. The predicted energy consumption for each day bin comes from the hourly coefficients estimated by the model for customer *i* from days in the baseline period with the same CDD, HDD, hours of daylight, and day type. Every energy consumption estimate for the post-period is based on similar days in the baseline. We perform a similar aggregation of actual post-period energy usage.

[14] Efron, B. 1987. “Better bootstrap confidence intervals.” *Journal of the American Statistical Association*, 82, 171-185.

Equation 4: Example Calculations of Predicted Summer Energy Consumption and Savings

$$\widehat{Summer_i} = 3 * \left(\hat{\beta}_{H0,i,112} + \hat{\beta}_{H1,i,112} + \hat{\beta}_{H2,i,112} + \cdots + \hat{\beta}_{H23,i,112}\right) + 7 * \left(\hat{\beta}_{H0,i,214} + \hat{\beta}_{H1,i,214} + \hat{\beta}_{H2,i,214} + \cdots + \hat{\beta}_{H23,i,214}\right) + \cdots$$

$$Actual_i = 3 * kWh_{i,112} + 7 * kWh_{i,214} + \cdots$$

$$Savings_i = \widehat{Summer_i} - ActualSummer_i$$

Where:

$\widehat{Summer_i}$ = Predicted energy consumption for customer *i* during the summer of a normalized weather year (TMY3)

$\hat{\beta}_{H0,i}$, $\hat{\beta}_{H1,i}$ ...= Coefficients estimated by the model for customer *i*

$ActualSummer_i$ = Actual energy consumption of customer *I*, adjusted to reflect the summer of a normalized weather year (TMY3)

$kWh_{i,112}$ = Average actual energy consumption of customer *i* in the post-period for individual days in bin *112* (a set combination of CDD, HDD, hours of daylight, and day type)

3.3.4 Demand Savings Estimation

For this analysis, we defined utility peak hours based on PGE's time-of-use (TOU) rates, as shown in Figure 6.

Figure 6: PGE Peak Period Definitions

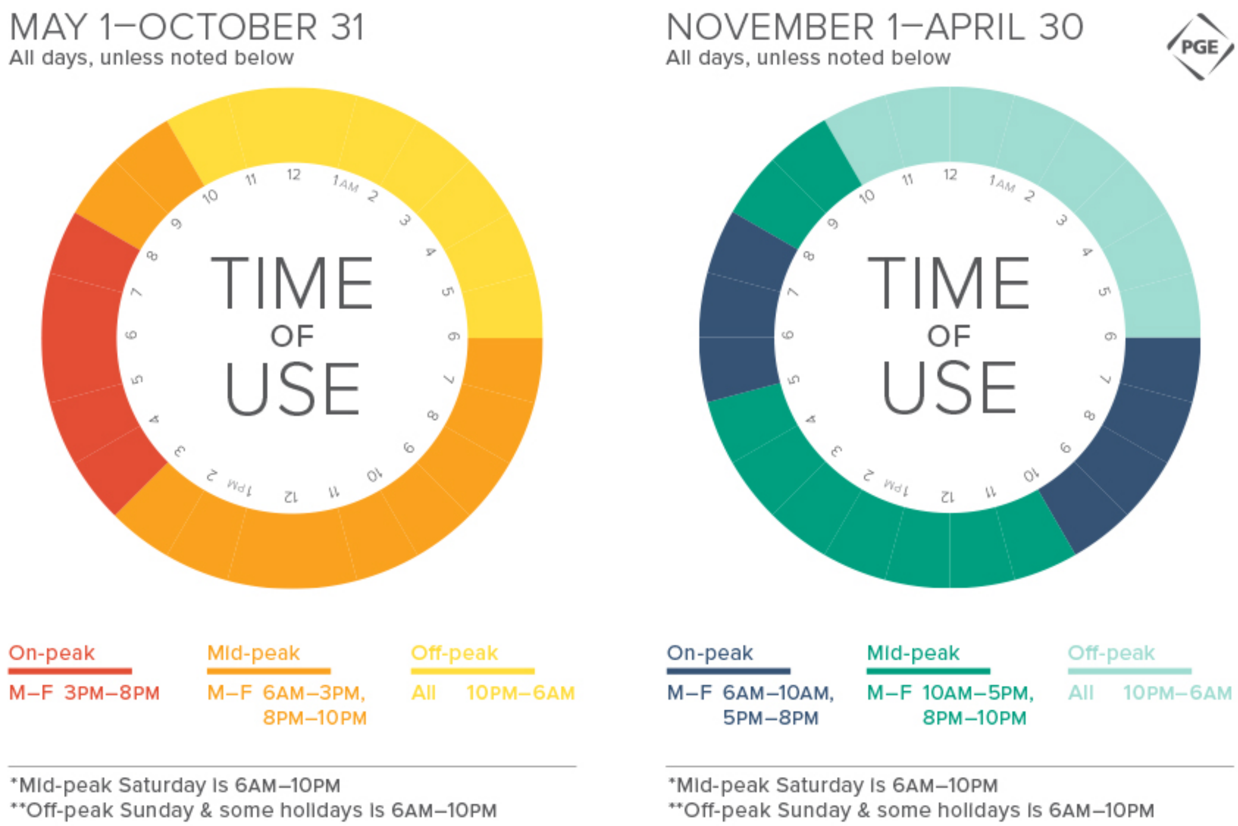

Image Source: Portland General Electric, *Time of Use Pricing: For Schedule 32 businesses and organizations*

For the peak demand (kW) estimate, we compared the load shapes estimated by the model for the hottest summer days (peak cooling) and coldest winter days (peak heating) to the actual energy usage of these buildings in the post-period, after the smart thermostats were installed. The model provides our best estimate for how each building would have used energy during utility peak hours on extreme weather days (e.g., day bin 911) if the smart thermostats had never been installed, as shown in Equation 5. We compare this prediction to the actual energy usage (kWh) during the utility peak hours on this extreme weather day as an approximation of the peak demand savings (kW).

Equation 5: Example Calculations of Predicted Summer Peak Energy Consumption and Savings

$$\widehat{Peak_i} = \left(\hat{\beta}_{H15,i,911} + \hat{\beta}_{H16,i,911} + \hat{\beta}_{H17,i,911} + \hat{\beta}_{H18,i,911} + \hat{\beta}_{H19,i,911} + \hat{\beta}_{H20,i,911}\right)/8$$

$$Actual_i = (kWh_{H15,i} + kWh_{H16,i} + kWh_{H17,i} + kWh_{H18,i} + kWh_{H19,i} + kWh_{H20,i})/8$$

$$Savings_i = \widehat{Peak_i} - ActualPeak_i$$

Where:

$\widehat{Peak_i}$ = Predicted energy consumption for customer *i* during the utility peak hours on the hottest days of the summer of a normalized weather year (TMY3)

$\hat{\beta}_{H15,i}$, $\hat{\beta}_{H16,i}$ …= Coefficients estimated by the model for customer *i* for hours 15, 16, …

$ActualSummer_i$ = Actual energy consumption of customer *i*, adjusted to reflect the summer of a normalized weather year (TMY3)

$kWh_{911,i}$ = Average actual energy consumption of customer *i* in the post-period for individual days in bin *911* (weekdays with the highest CDD, lowest HDD, and set hours of daylight)

3.4 Acknowledgement of Limitations

Our billing analysis was limited by several factors. First, our requirement of at least nine months of post-period data restricted the number of available sites. While this requirement was critical to ensure that we had sufficient observations of the peak heating and peak cooling periods at every site, this requirement still restricted the number of sites in our analysis.

A second limitation of this analysis stems from the various analysis-related impacts of the COVID-19 pandemic. One such impact was a delay in installation, which contributed to a limited number of sites having sufficient post-period data. Another impact was on the energy usage data of the sites with sufficient post-period data. At sites where the pandemic has led to non-routine energy usage, we have attempted to correct for this impact with NRE adjustments. While these adjustments should lead to more robust analysis, they also introduce uncertainty and may not fully account for the impacts of the NRE.

3.4.1 Impact of the COVID-19 Pandemic on Savings Estimates

This section gives additional detail on how the COVID-19 pandemic has impacted participant energy usage and the extent to which we were able to adjust for this impact in our estimates of energy savings attributable to the smart thermostats.

Among the many impacts of the COVID-19 pandemic has been a disruption of energy efficiency programs and their evaluation. This pilot was no exception, where:

1. Smart thermostat installations were put on hold in March 2020; and
2. Existing commercial energy usage patterns changed, ranging from immediate business closures with gradual reopening to more modest changes in energy usage resulting from new public health mandates (e.g., increased ventilation, reduced occupancy).

The pandemic is responsible for some of the original 588 sites being removed from this analysis because of its impact on the baseline model fit and/or its inconsistent impact on energy usage during the post-period (with more than one significant NRE). Despite these limitations, we identified 182 electric sites and 153 gas sites with qualified thermostats and sufficient data that could support billing analysis. Of these, 11 electric (6% of the electric sample) and 9 gas (4%) sites had extensive energy usage data prior to the start of the COVID-19 pandemic, creating an opportunity for us to compare the program impacts for these sites before the pandemic versus during the pandemic.

This analysis demonstrates that even before the start of the pandemic, these smart thermostat devices led to statistically significant reductions in energy usage similar to what was expected (per the *ex-ante* savings listed in the measure approval document). Table 7 (kWh) and Table 8 (therms) summarize the per-thermostat savings estimates on a daily level with 95 percent confidence intervals from October to March for the 11 electric sites and 9 gas sites with sufficient data. The 'Actual Energy Usage' column shows the average daily energy usage per thermostat across each time period, prior to and during the COVID-19 pandemic. The 'Adjusted Prediction' column represents the expected average daily energy usage after applying our adjustments for NREs. The values in this table are **weighted to represent only the observed weather of the period prior to the start of the COVID-19 pandemic** (10/09/2019 to 03/14/2020) for the 11 electric sites and 9 gas sites in order to create an apples-to-apples comparison. In other words, both of these estimates reflect the same sites, the same weather conditions, and the same seasonal energy consumption seen during winter and shoulder months from mid-October to mid-March. The only remaining difference should be the COVID-19 pandemic.

Our analysis of the timeframe prior to the start of the COVID-19 pandemic confirms that savings from qualified smart thermostats were both large and statistically significant. The devices are clearly saving energy. The *ex-ante* savings estimate for the 11 electric sites is 7.70 kWh per day, which is similar to our estimated daily energy savings of 7.61 kWh per day prior to the start of the COVID-19 pandemic. Similarly, gas savings were significant prior to the start of the pandemic, although considerably lower than the *ex-ante* savings value of 10.0 therms per day at gas sites with sufficient pre-pandemic data.[15] While this is not a perfect comparison, this heuristic demonstrates the **existence of non-zero energy savings that align with the *ex-ante* savings, prior to the introduction of additional uncertainty from the COVID-19 pandemic** and non-routine adjustments.

The estimates for savings during year 1 of the COVID-19 pandemic are higher, though not statistically significantly different. While the sample is limited to only 11 electric and 9 gas sites, this analysis suggests that the first year of the pandemic **may have inflated savings estimates** (from 7.55 kWh to 12.75 kWh; and from 0.42 therms to 0.80 therms) despite our efforts to adjust for the impact of the COVID-19 pandemic. Fortunately, this incremental **impact diminishes over**

[15] This daily savings estimate has not been adjusted to represent the same weather conditions observed prior to the COVID-19 pandemic (October 2019-March 2020). This is a simplified comparison using the best available estimate from the measure approval document, annualized *ex-ante* savings.

time as electric savings during year two of the COVID-19 pandemic dropped closer to the pre-COVID values (and gas savings became insignificant at these sites).

Table 7: Summary of Per-Thermostat Daily Electric Savings Impact for Selected Sites by Time Period (kWh)

Time Period	# of Sites	# of T-stats	Adjusted Prediction	Actual Energy Usage	Daily kWh Savings	% Savings Estimate
Prior to COVID-19 pandemic (Oct 2019 - Mar 2020)	11	22	85.66	78.11	7.55 ± 3.3	8.8% ± 3.9%
During COVID-19 pandemic Year 1 (Mar 2020 - Feb 2021)	11	22	85.66	72.90	12.75 ± 3.26	14.9% ± 3.8%
During COVID-19 pandemic Year 2 (Mar 2021- Feb 2022)	11	22	84.95	75.05	9.9 ± 3.72	11.7% ± 4.4%

Table 8: Summary of Per-Thermostat Daily Gas Savings Impact for Selected Sites by Time Period (therms)

Time Period	# of Sites	# of T-stats	Adjusted Prediction	Actual Energy Usage	Daily Therm Savings	% Savings Estimate
Prior to COVID-19 pandemic (Oct 2019 - Mar 2020)	9	22	5.14	4.71	0.42 ± 0.26	8.2% ± 5.1%
During COVID-19 pandemic Year 1 (Mar 2020 - Feb 2021)	9	22	5.14	4.36	0.78 ± 0.50	15.2% ± 9.6%
During COVID-19 pandemic Year 2 (Mar 2021- Feb 2022)	9	22	5.47	5.37	0.10 ± 0.38	1.8% ± 6.9%

To further highlight this point, we have selected two sites that clearly demonstrate two common scenarios observed in this study. The first example demonstrates that the COVID-19 pandemic did not cause the high energy savings estimates that we are seeing at many sites in our analysis, while the second demonstrates the complexity of isolating the impact of the smart thermostat installations at sites that were impacted by the COVID-19 pandemic.

Figure 6 shows the daily electric energy usage at Site A one year before and one year after the thermostat installations (dotted line). The solid vertical line marks the start of the COVID-19 pandemic. Site A is a non-grocery business with gas heat and cooling that installed five thermostats. As shown in the figure, the start of the COVID-19 pandemic is associated with little to no change in energy usage at this site, with similar levels of energy usage to the right and left of the solid line. This is contrasted with the energy usage observed after the thermostat installations (dotted line), where we suddenly see many days with energy usage below 100 kWh instead of around 200 kWh or above.

After normalizing for weather conditions with the regression model, we estimated savings of nearly 44,000 kWh per thermostat per year (38%) across the five installed thermostats, far above the *ex-ante* savings value of 1,275 kWh for this site. This site demonstrates that **savings in excess of *ex-ante* were observed at sites where there was little to no COVID-19 impact.** Hence, **our higher-than-expected savings estimates are <u>not</u> necessarily evidence that the COVID-19 adjustment is incomplete.**

Figure 6: Raw Daily Electric Energy Usage at Site A

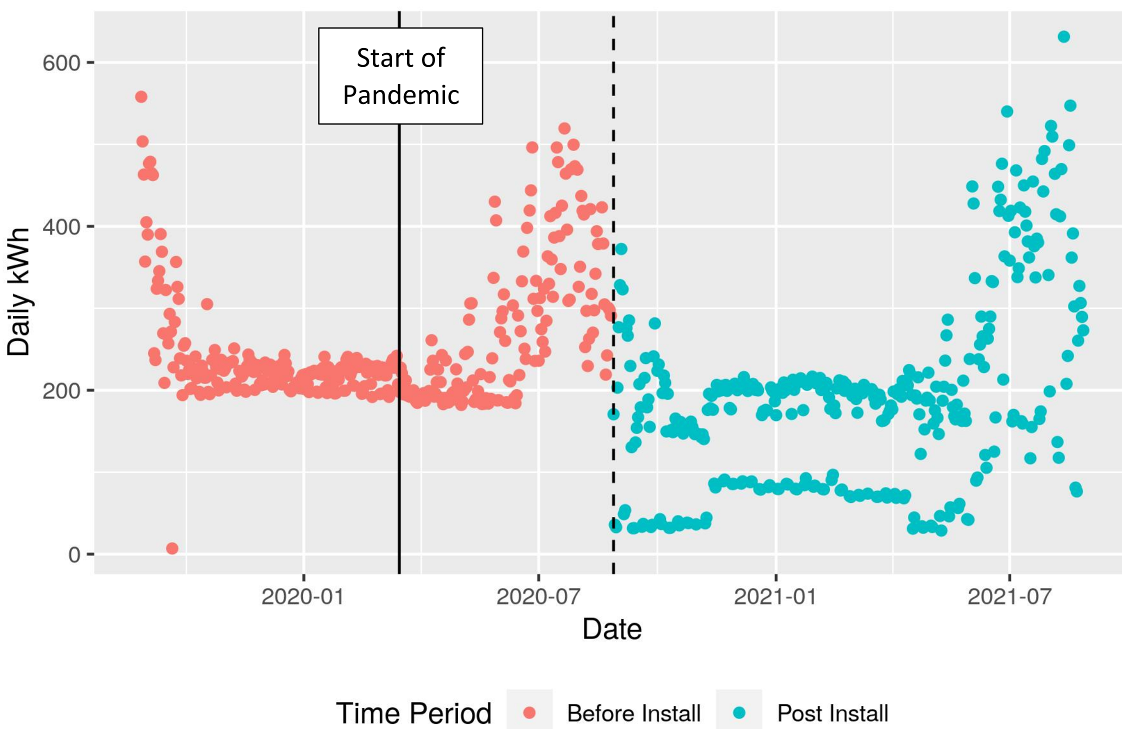

Figure 7 shows the raw daily energy usage at Site B one year before and one year after program participation. As with the previous figure, the dotted line shows the thermostat installations, and the solid vertical line indicates the start of the COVID-19 pandemic. Site B is a non-grocery business with gas heat and cooling with two thermostats installed. As opposed to Site A, electricity usage appears to decrease dramatically from around 75 kWh to around 25 kWh per day at the start of the pandemic (solid line). The weeks before and after the installation of the thermostats (dotted line) appear similar, at around 10 kWh per day. Shortly after, however, energy usage increases, returning to a more typical level for this site at around 50 kWh, though this is still much lower than the pre-pandemic baseline. One possible explanation for this pattern is that many businesses closed temporarily early in the COVID-19 pandemic when the first shelter-in-place orders were issued. The thermostats were installed during this period of low energy usage, which we presume is related to the COVID-19 pandemic. In August 2020, their energy usage ticks back up; this is a few months after the end of the shelter-in-place orders when the economy was gradually reopening. The big question at this site is why did energy usage not return to prior levels of around 75 kWh? Was this due to the smart thermostats, the continued impact of the COVID-19 pandemic, or both?

Our NRE detection method is capable of identifying the period from March 2020 to September 2020 as an NRE in need of adjustment at this site. Unfortunately, the installations occurred during this NRE. With such a short time before and after the installation *during* an event, we are unable to claim sufficient certainty in our estimates of the incremental impact of the thermostats on energy usage.

While we have corrected for this primary NRE, **there exists the possibility that the COVID-19 pandemic continues to impact energy usage after the end of this event** (March-September 2020) as the customer's energy usage increased substantially but did not fully return to levels observed prior to the start of the pandemic. This is reasonable for office buildings, as many have still not returned to pre-pandemic levels of occupancy, with more people choosing to work from home.

For Site B, our model estimated savings of nearly 5,000 kWh per thermostat per year (21%) across the two thermostats at this site—again, well above the *ex-ante* savings value of 510 kWh for this site. In this example, **the COVID-19 pandemic may be inflating or deflating our savings estimate.** This site clearly demonstrates the **complexity of isolating the impact of the smart thermostat installations at sites impacted by the COVID-19 pandemic**.

Figure 7: Raw Daily Energy Usage as Site B

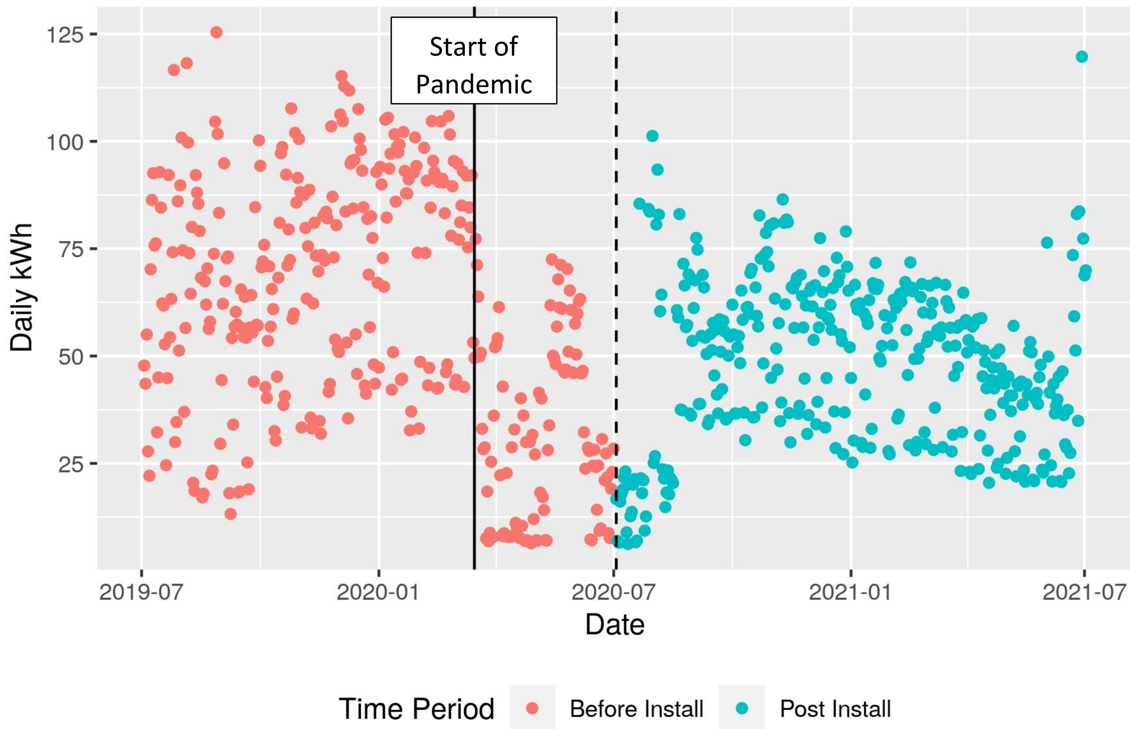

We attempted to adjust for the impact of the COVID-19 pandemic and subsequent NREs at all sites where it was required, but there is still a possibility that the savings estimates have been influenced by these events. This is primarily the result of variable impacts that the pandemic had on energy usage over time. For example, consider a scenario where a business closed during the initial shelter-in-place orders. It reopened two months later, increasing ventilation (per CDC guidelines), shortening operating hours, and increasing the time spent disinfecting surfaces between shifts. The temporary closure would have a large and statistically significant impact on energy usage (seen as a sudden increase in residuals that spans two months) that our NRE detection procedure would identify and then adjust for. There are three possibilities for the reopening phase:

1. Reopening is statistically significantly different from the baseline and seen as a continuation of the closure, where COVID has a sudden negative impact on energy usage (intercept) that gradually lessons over time (slope).
2. Reopening is statistically significantly different from the baseline and seen as an entirely separate NRE; this site would be dropped for having multiple NREs (and therefore too much uncertainty).
3. Reopening is not statistically significantly different from the baseline; no adjustment is made. Though there may have been changes at the site, they are not large enough to trigger a non-routine adjustment applied, despite its causal relation to the pandemic. In

this situation and for others like it in this analysis, **most of the impact of the COVID-19 pandemic has been corrected, but some lingering impacts may remain embedded in the savings estimate**. The challenge we are facing is that **there is no way to know for sure** what caused energy usage to change from the billing data alone. Without additional information, we are not able to tease out gradual changes caused by the COVID-19 pandemic from gradual changes in thermostat operation.

Our analysis of the savings from 11 electric and 9 gas sites with thermostats installed before and during the pandemic confirmed that savings appear higher during the COVID-19 pandemic, but the difference was not statistically significant.

Savings estimates presented in the following sections should be interpreted with caution due to the lingering uncertainty in the COVID-19 adjustment, with the caveat that savings may still be inflated but should slowly be approaching the long-term savings as businesses settle into a new normal.

4 Findings

This chapter provides findings from our analysis of participant billing data.

4.1 Overall Smart Thermostat Impacts

Our first section of results focuses on the broad savings impacts of the thermostat pilot based on the billing analysis.

4.1.1 Energy Savings

First, we present savings aggregated across all sites included in the billing analysis (n=182 electric sites) in the form of a load shape. Figure 8 shows the predicted post-period load shape (red) with the actual post-period load shape (blue) for the **average site on the average day** during the year. This prediction is based on the pre-period model and post-period weather data; it represents the expected load shape for these customers in the absence of the smart thermostat, after adjusting our predictions to account for non-routine events (NREs) (including the COVID-19 pandemic). The error around each hourly prediction is depicted as a 95 percent confidence interval in the shaded area around each estimate. Whenever the actual post-period load shape (blue line) falls outside the predicted post-period load region (red area), this indicates that a statistically significant change was observed during that hour. While later results are presented at the thermostat level, load shape analysis including Figure 8 are at the site level to maintain consistency with the underlying data. The Advanced Metering Infrastructure Customer Segmentation (AMICS) model found statistically significant reductions in the whole-building energy usage for pilot participants across all hours of the day. Overall, **participants reduced their pre-period energy usage by an average of 15.6 percent by installing one or more smart thermostats**.

Figure 8: Hourly Load Impacts of Smart Thermostats on the Average Site (a single site on the average day across a full year)

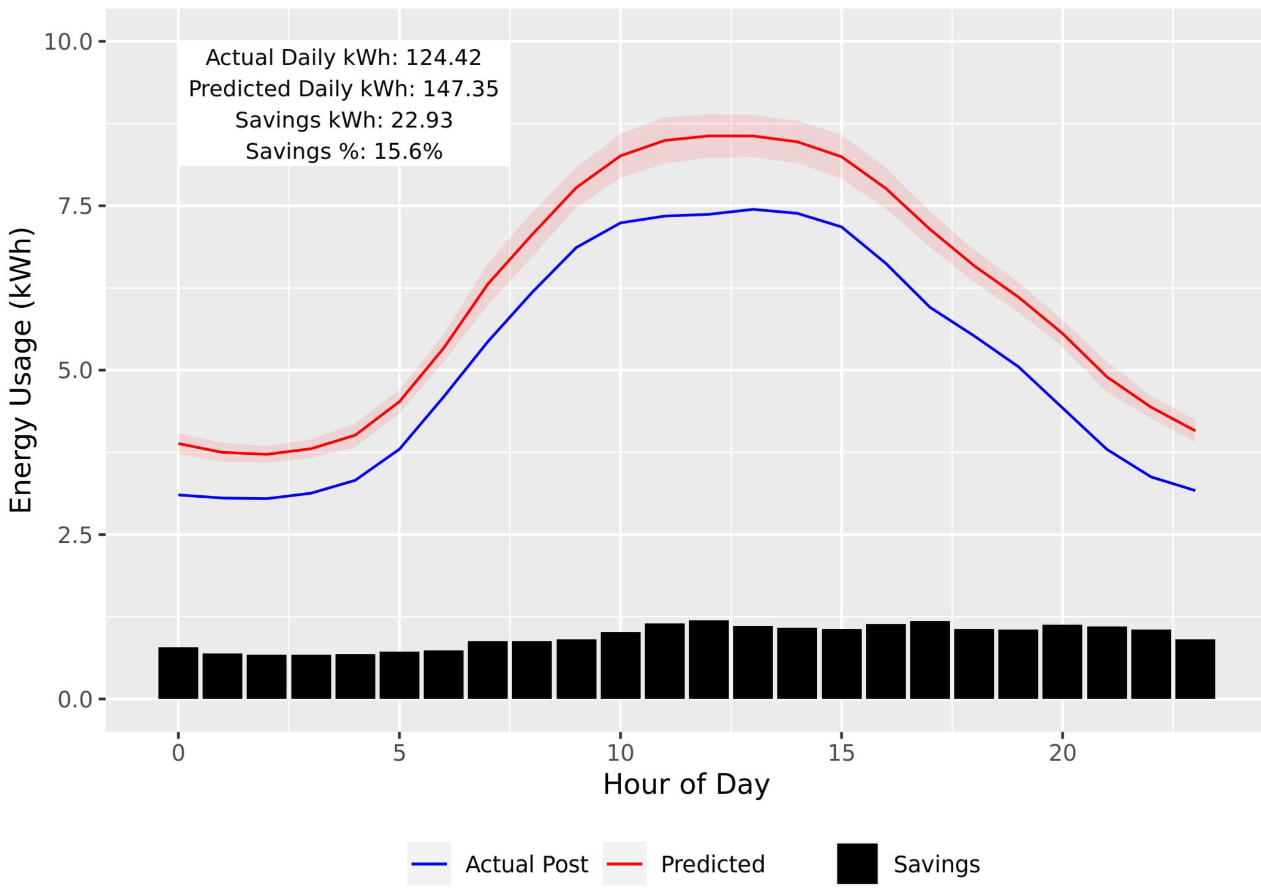

The site-level impacts are caused by one or more smart thermostats. Table 9 compares the *ex-ante* kWh savings from the measure approval document with our estimated, NRE-adjusted savings. NRE-adjusted savings account for most of the impacts of the COVID-19 pandemic, but some lingering impacts may remain embedded in these estimates. These results have been normalized to a typical weather year (TMY3) and represent savings per-thermostat and include a 95 percent confidence interval. A total of 410 thermostats were installed at the 182 sites included in this analysis. For all but nine of these sites, the measure type was consistent across all thermostats at the site, allowing us to directly compare estimated per-thermostat kWh savings with *ex-ante* values. At these nine sites, multiple types of measures were installed. Results for these sites have been omitted as they are not representative of any individual measure. In general, our annual kWh savings estimates greatly exceeded the *ex-ante* savings values. It is also worth noting that the majority of the sites included in our analysis of electric savings are gas heated.

Table 9: Annual Normalized *Ex-Ante* and *Ex-Post* Savings Estimates for kWh (TMY3 Adjusted)

Measure Type	# of Sites	# of T-stats	*Ex-Ante* Savings per T-stat (Annual kWh)	*Ex-Post* Savings per T-stat (Annual kWh)
Non-Grocery - Resistance Heat w/ Cooling (HZ1) - Smart Thermostat	9	10	940	1,239.0 ± 2,171.4
Non-Grocery - Heat Pump w/ Cooling (HZ1) - Smart Thermostat	19	33	580	4,856.6 ± 1,565.6
Non-Grocery - Gas Heat w/ Cooling (HZ1) - Smart Thermostat	144	315	255	3,429.3 ± 2,356.0
Grocery - Gas Heat w/ Cooling (HZ1) - Smart Thermostat	1	5	*	*
Multiple Measures	9	47	*	*

Note: While most of the impact of the COVID-19 pandemic has been accounted for with NRE adjustments, some lingering impacts may still inflate these savings estimates.

*We omitted customer groups with fewer than five customers, as the sample is likely too small to draw meaningful conclusions from. Results for multiple measures have been omitted as they are not representative of any individual measure.

Similarly, Table 10 compares *ex-ante* gas savings by measure with the AMICS-estimated, NRE-adjusted savings calculated in this analysis. These results have been normalized to a typical weather year (TMY3) and represent per-thermostat values with a 95 percent confidence interval. A total of 318 thermostats were installed at the 153 sites included in this analysis, almost all of which are gas heated. For all but five of these sites, the measure type was consistent across all thermostats at the site, allowing us to directly compare estimated per-thermostat savings with *ex-ante* values. While our annual gas savings estimates generally greatly exceeded the *ex-ante* savings values, in most cases they were not statistically significant.

Table 10: Annual Normalized *Ex-Ante* and *Ex-Post* Savings Estimates for Gas (TMY3 Adjusted)

Measure Type	# of Sites	# of T-stats	*Ex-Ante* Savings per T-stat (Annual Therms)	*Ex-Post* Savings per T-stat (Annual Therms)
Non-Grocery - Resistance Heat w/ Cooling (HZ1) - Smart Thermostat	1	1	*	*
Non-Grocery - Heat Pump w/ Cooling (HZ1) - Smart Thermostat	3	3	*	*
Non-Grocery - Gas Heat w/ Cooling (HZ1) - Smart Thermostat	144	285	31.0	154.8 ± 236.1
Multiple Measures	5	29		

* We omitted customer groups with fewer than five customers, as the sample is likely too small to draw meaningful conclusions from.

4.1.2 Demand Savings

We also analyzed savings on the basis of various day types including the hottest and coldest days of the year. Figure 9 shows the post-period predicted load shape (red) with the actual post-period load shape (blue) for analysis sites on the hottest days of the year at each site (with CDDs between 7 and 19 corresponding to average daily temperatures between 72°F and 84°F with daily highs up to 98°F). This prediction is based on the pre-period model and post-period weather data; it represents the expected load shape for these customers in absence of program pilot participation after adjusting to account for NREs. The error of each hourly prediction is depicted as a 95 percent confidence interval in the shaded area around each estimate. Whenever the actual post-period load shape (blue line) falls outside the predicted post-period load region (red shaded area), this indicates that a statistically significant change was observed during that hour. As the actual post-period energy usage (blue line) falls below the predicted (red line) during every hour, we can conclude that the AMICS model detected statistically significant reductions in the whole-building energy usage for pilot participants across all hours on the hottest days of the year. Overall, participants saved an average of 26.1 kWh or 14.3 percent on these extreme cooling days. In the peak hour of the day (5 p.m.), participants saved an average of 0.56 kW per thermostat or 8.5 percent, although this estimate is insignificant.

Figure 9: Load Shape During Peak Cooling, Hottest Days of the Year (an average site on the hottest day)

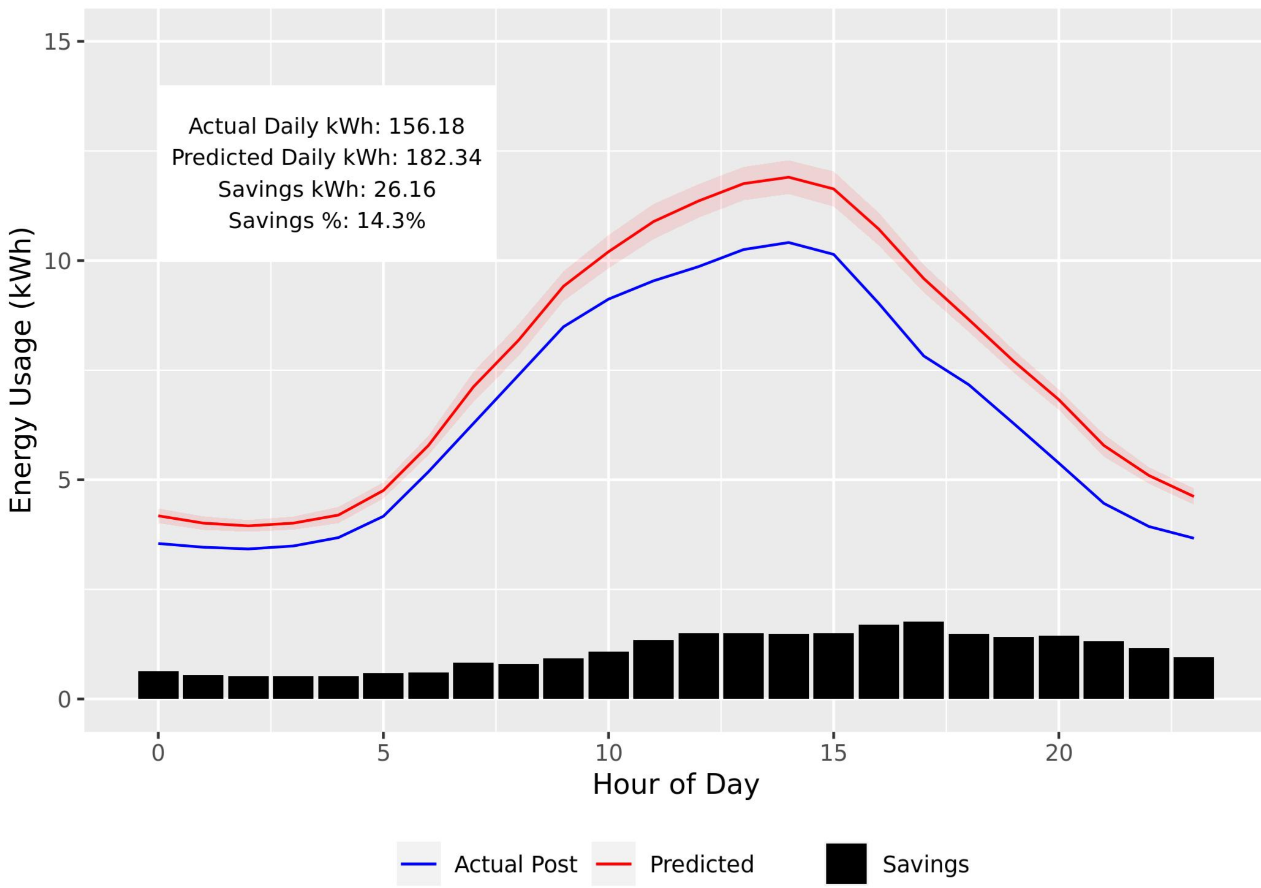

Figure 9 shows the actual post-period load shape (blue) against our predicted load shape (red) for participants on the coldest days of the year at each site (with HDDs between 27 and 34 corresponding to average daily temperatures between 38°F and 31°F, with daily lows as low as 24°F). Again, we find statistically significant reductions in the whole-building energy usage for pilot participants across all hours of the day on the coldest days of the year. Overall, participants saved 27.5 kWh or 22 percent on peak heating days. In the peak heating hour of 8 p.m., participants saved an average of 0.91 kW or 20.5 percent.

Figure 10: Load Shape During Peak Heating, Coldest Days of the Year

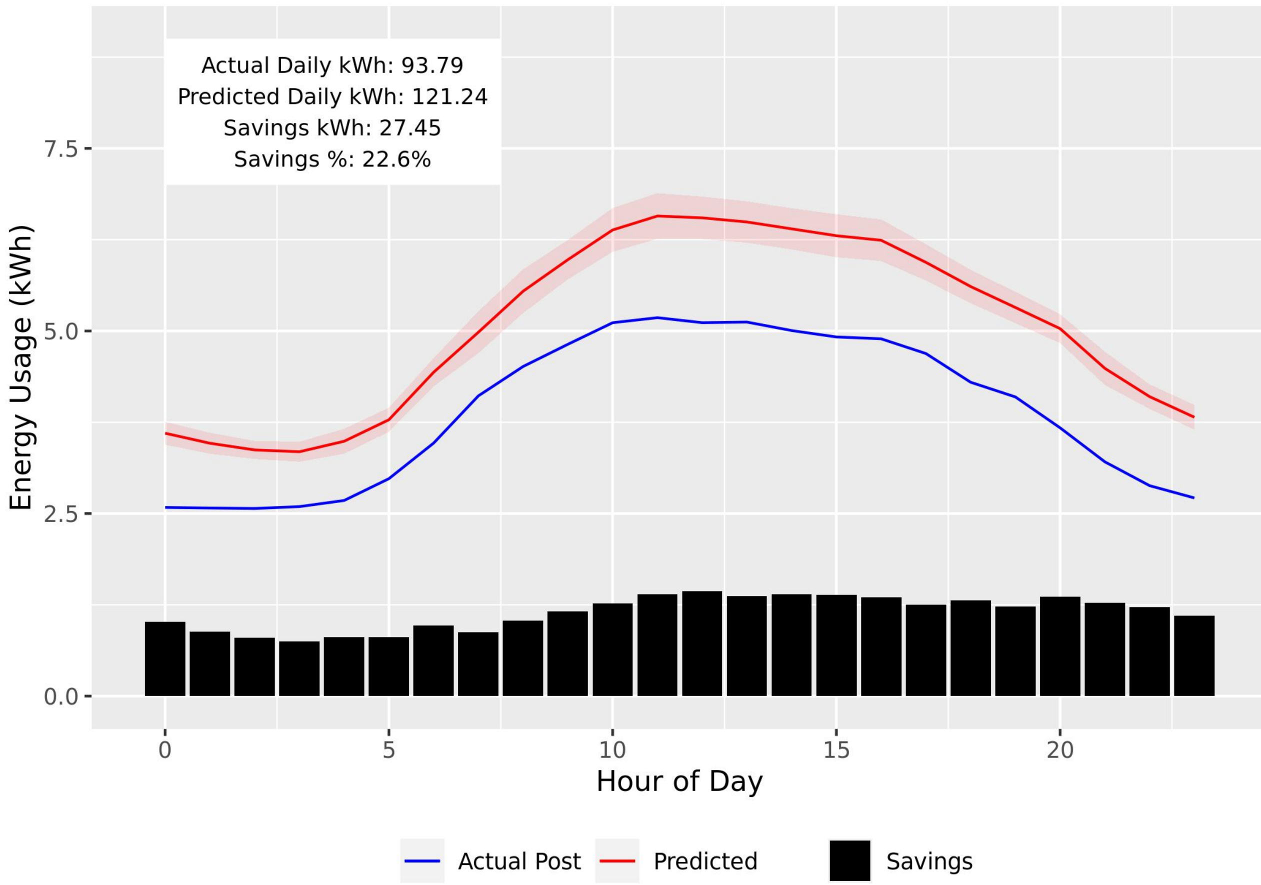

Table 11 summarizes the charts above and shows how savings varied on these day types for each thermostat installed. "Adjusted Prediction" includes the NRE adjustment (e.g., estimated increase in usage associated with the COVID-19 pandemic), while "Original Prediction" does not (i.e., our initial prediction for the post-period if the COVID-19 pandemic had never occurred). Peak heating represents the peak heating hours (6 a.m. to 10 a.m. and 5 p.m. to 8 p.m.) on the coldest days observed at each site. While this varied by site, this includes a range of HDDs between 27 and 34 (average daily temperatures between 38°F and 31°F, with daily lows as low as 24°F). Peak cooling represents the peak cooling hours (3 p.m. to 8 p.m.) on the hottest days at a site. Again, these values varied by site but represent a CDD ranging between 7 and 19 (average daily temperatures between 72°F and 84°F with daily highs up to 98°F). Days without these extreme temperatures are shown in the "All Other Days" row. Savings for all three day types were statistically significant for kW, with peak heating days having higher savings in terms of kW (0.87 and 0.62 kW) and as a percentage of baseline energy usage (22% and 12%). Savings estimates include a 95 percent confidence interval.

Table 11: Average Demand Savings per Thermostat on Peak Heating and Cooling Days (kW)

	Sites	T-stats	Original Prediction	Adjusted Prediction	Actual	kW Savings	% Savings
Peak Heating Day	182	410	3.80	3.90	3.04	0.87 ± 0.42	22.2% ± 10.7%
Peak Cooling Day			5.07	5.14	4.52	0.62 ± 0.50	12.1% ± 9.8%
All Other Days			3.34	3.44	2.99	0.45 ± 0.35	13.1% ± 10.1%

The demand savings results in Table 11 can be further broken out by the heating fuel used at each site. Table 12 shows how the demand savings varied between sites heated with gas and those heated with electricity. Sites where multiple measures were installed have been excluded. The highest electric savings in both kWh and percentage terms occurred at gas-heated sites during peak heating hours. While usage was considerably higher at electrically-heated sites during these same hours (4.80 kWh at electrically-heated sites, 3.76 kWh at gas-heated sites), demand savings were lower (0.52 kW for electrically-heated sites, 0.94 kW for gas-heated sites). Electric savings during peak cooling and all other (non-peak) days were similar between electrically-heated and gas-heated sites. Savings estimates include a 95 percent confidence interval.

Table 12: Average Demand Savings per Thermostat on Peak Heating and Cooling Days by Heating Fuel (kW)

		Sites	T-stats	Original Prediction	Adjusted Prediction	Actual	kW Savings	% Savings
Electric Heat	Peak Heating	28	43	4.74	4.80	4.28	0.52 ± 0.34	10.8% ± 7.1%
	Peak Cooling			4.93	4.61	3.99	0.62 ± 0.21	13.4% ± 4.5%
	All Other Days			3.78	3.68	3.18	0.50 ± 0.23	13.6% ± 6.1%
Gas Heat	Peak Heating	145	320	3.65	3.76	2.82	0.94 ± 0.42	25.0% ± 11.0%
	Peak Cooling			5.20	5.36	4.75	0.61 ± 0.48	11.3% ± 8.9%
	All Other Days			3.32	3.47	3.03	0.44 ± 0.36	12.7% ± 10.5%

4.2 Impacts by Season and Day Type

Next, we analyzed program impacts on the basis of season. Figure 11 shows the post-period predicted load shape (red) with the actual post-period load shape (blue) for analysis sites during each of the four seasons. As the actual load shape (blue line) falls below the predicted load shape (red line and shaded area), we see statistically significant reductions in the whole-building energy usage for pilot participants across all hours for each season. However, the magnitudes and hours of peak savings vary.

Figure 11: Average Hourly Savings by Season (for the average site)

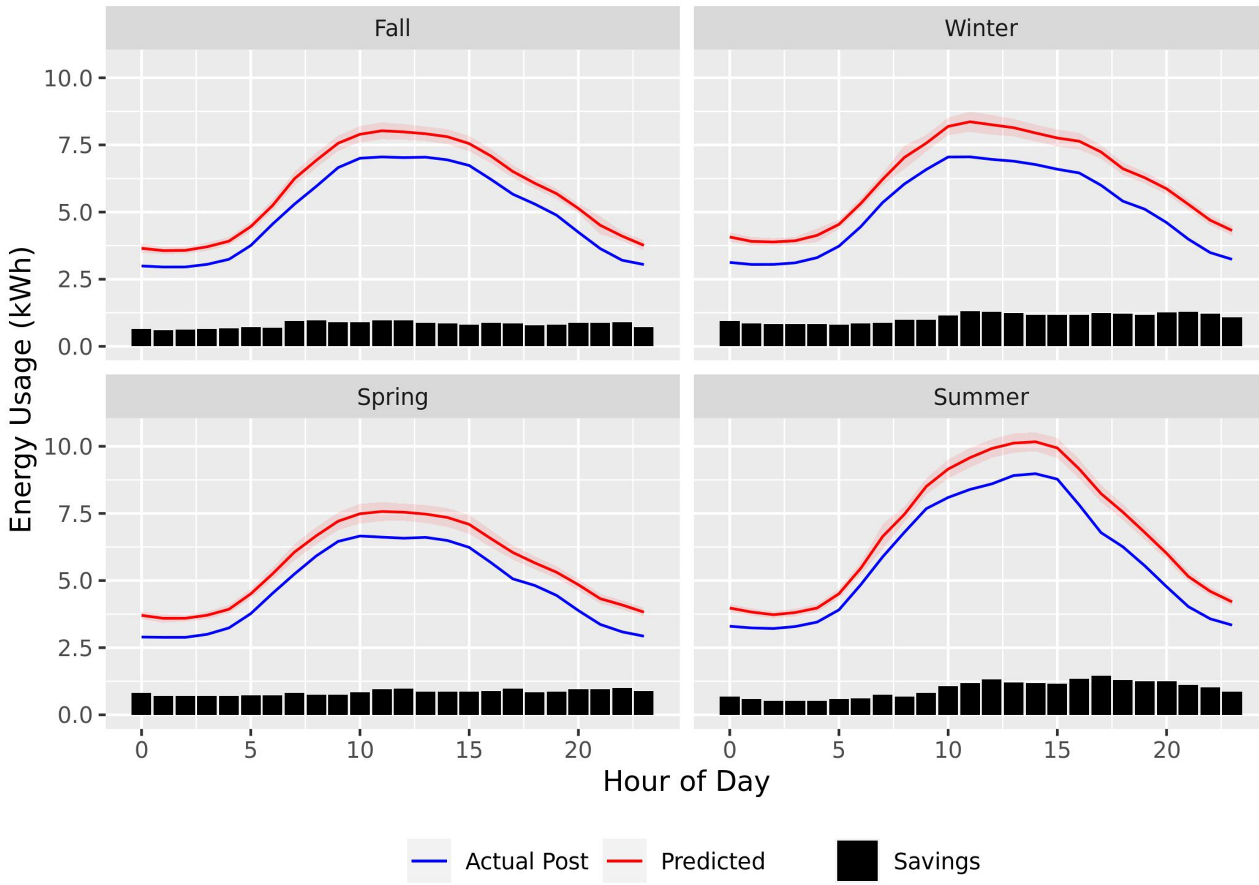

Table 13 summarizes Figure 11 and shows how savings varied by season. Results have been normalized to represent the full season within a typical weather year (TMY3) and include 95 percent confidence intervals. “Adjusted Prediction” accounts for NREs, while “Original Prediction” does not. Average savings were lowest in the summer both in terms of kWh and as a percentage of baseline energy usage, while winter had the highest average savings. Table 14 summarizes the similar information, but in terms of therms, for gas savings. Gas savings were largest during the winter and insignificant in the other seasons.

Table 13: Seasonal Electric Savings, kWh per Thermostat (TMY3 Adjusted)

Season	Sites	T-stats	Original Prediction	Adjusted Prediction	Actual	Seasonal kWh Savings (kWh per thermostat, TMY3 adjusted)	% Savings
Fall	182	410	7,438	7,434	6,501	934 ± 877	12.6% ± 11.8%
Winter	182	410	7,398	7,397	6,252	1,145 ± 696	15.5% ± 9.4%
Spring	182	410	6,326	6,994	5,999	995 ± 722	14.2% ± 10.3%
Summer	182	410	8,431	8,634	7,717	917 ± 860	10.6% ± 10.0%

Table 14: Seasonal Gas Savings, therms per Thermostat (TMY3 Adjusted)

Season	Sites	T-stats	Original Prediction	Adjusted Prediction	Actual	Seasonal Therms Savings (Therms per thermostat, TMY3 adjusted)	% Savings
Fall	152	318	336	313	295	17.4 ± 51.0	5.6% ± 16.3%
Winter	153	318	595	591	502	89.5 ± 78.6	15.1% ± 13.3%
Spring	153	318	316	335	300	35.2 ± 67.7	10.5% ± 20.2%
Summer	153	318	201	193	177	16.7 ± 50.4	8.6% ± 26.0%

Figure 12 shows the post-period predicted load shape (red) with the actual post-period load shape (blue) by day type and season, with hours shaded to reflect the time-of-use (TOU) rate that applies. The purpose of this analysis is to assess whether the smart thermostat savings occur during hours when the utility would prefer to shed load, with the highest priority being a reduction in on-peak hours (orange) and the lowest during off-peak hours (purple). There are statistically significant reductions in the whole-building energy usage for pilot participants across all hours of the day, day types, and seasons, although the magnitude of these savings varies. These smart thermostats do offer savings during weekday on-peak hours for both summer and winter, but the greatest savings occur off-peak on Sunday afternoons. Smart thermostats provide energy savings during peak hours through temperature setbacks in summer months and setups in winter months. The off-peak savings will also come from the thermostats reducing HVAC usage when the buildings are unoccupied, which occurs more often on weekends and overnight.

Figure 12: Electric Savings by TOU Peak Period (for the average site)

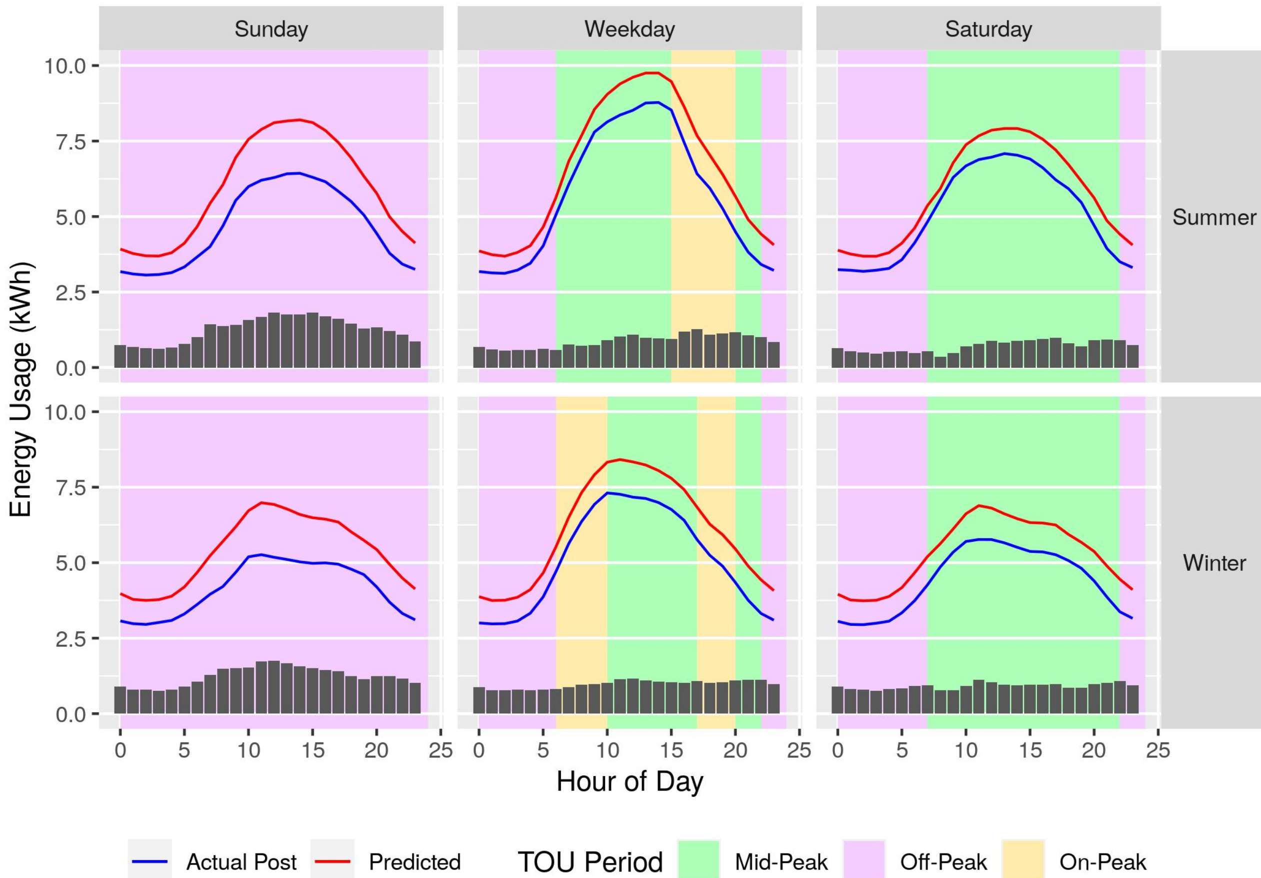

Table 15 summarizes Figure 11 and shows savings (with 95% confidence intervals) by TOU time period on an average hourly basis. "Adjusted Prediction" accounts for NREs, while "Original Prediction" does not. Savings across all of the TOU periods were statistically significant with the exception of summer mid-peak. The off-peak periods in both the winter and summer had the most variability, with wider confidence intervals. Off-peak savings appear higher than on-peak or mid-peak savings, though the differences are not statistically significant.[16]

[16] This analysis was not possible for gas because the usage data are in daily intervals.

Table 15: Hourly Thermostat Savings by TOU Period (per thermostat)

TOU Phase	TOU Period	Sites	T-stats	Original Prediction	Adjusted Prediction	Actual	Average kW Savings	% Savings
Summer	Mid-Peak	182	410	3.36	3.33	2.94	0.39 ± 0.39	11.7% ± 11.8%
	Off-Peak			2.64	2.64	2.09	0.54 ± 0.48	20.6% ± 18.3%
	On-Peak			3.61	3.48	2.98	0.50 ± 0.41	14.3% ± 11.7%
Winter	Mid-Peak	182	410	2.95	2.97	2.53	0.44 ± 0.31	14.8% ± 10.4%
	Off-Peak			2.39	2.40	1.85	0.55 ± 0.37	22.9% ± 15.6%
	On-Peak			2.90	2.94	2.51	0.43 ± 0.26	14.6% ± 8.9%

The results in Table 15 can be further broken out by the heating fuel used at each site. Table 16 shows how the winter TOU results varied between sites heated with gas and those heated with electricity. Sites where multiple measures were installed have been excluded. The highest savings in both kWh and percentage terms occurred at electrically-heated sites off-peak. However, off-peak savings are similar in percentage terms between gas- and electrically-heated sites, and on-peak savings are higher among gas-heated sites in both kWh and percentage terms.

Table 16: Hourly Thermostat Savings during Winter TOU Period by Heating Fuel (per thermostat)

Heating Fuel	TOU Period	Sites	T-stats	Original Prediction	Adjusted Prediction	Actual	Average kW Savings	% Savings
Electric	Mid-Peak	28	43	3.75	3.64	3.17	0.47 ± 0.32	12.8% ± 8.8%
	Off-Peak			3.24	3.18	2.47	0.71 ± 0.35	22.4% ± 10.9%
	On-Peak			3.83	3.74	3.45	0.29 ± 0.26	7.8% ± 6.8%
Gas	Mid-Peak	145	320	2.86	2.96	2.52	0.44 ± 0.32	14.8% ± 10.7%
	Off-Peak			2.25	2.33	1.81	0.52 ± 0.39	22.1% ± 16.6%
	On-Peak			2.81	2.90	2.46	0.45 ± 0.28	15.4% ± 9.7%

4.3 Impacts Across Customers

In addition to analyzing the average smart thermostat impacts, we also analyzed how impacts varied across building and thermostat characteristics.

As a basis for subsequent analysis, we first present savings estimates for the individual sites included in this analysis (n=182 electric sites and n=153 gas sites). Figure 13 (kWh) and Figure 14 (gas) show per-thermostat savings estimates and confidence intervals for each site in this analysis, ordered by the per-thermostat savings observed at the site. These figures highlight the variation in savings that we observed at a site level as well as the range of confidence intervals we see in this analysis, with wide confidence intervals around many of the highest savings estimates. The savings per thermostat were not corelated with the number of thermostats installed. Our subsequent analysis aggregates these findings by customer segment to identify characteristics that may be driving the variability in savings.

Figure 13: Annual Electric Savings per Thermostat (kWh)

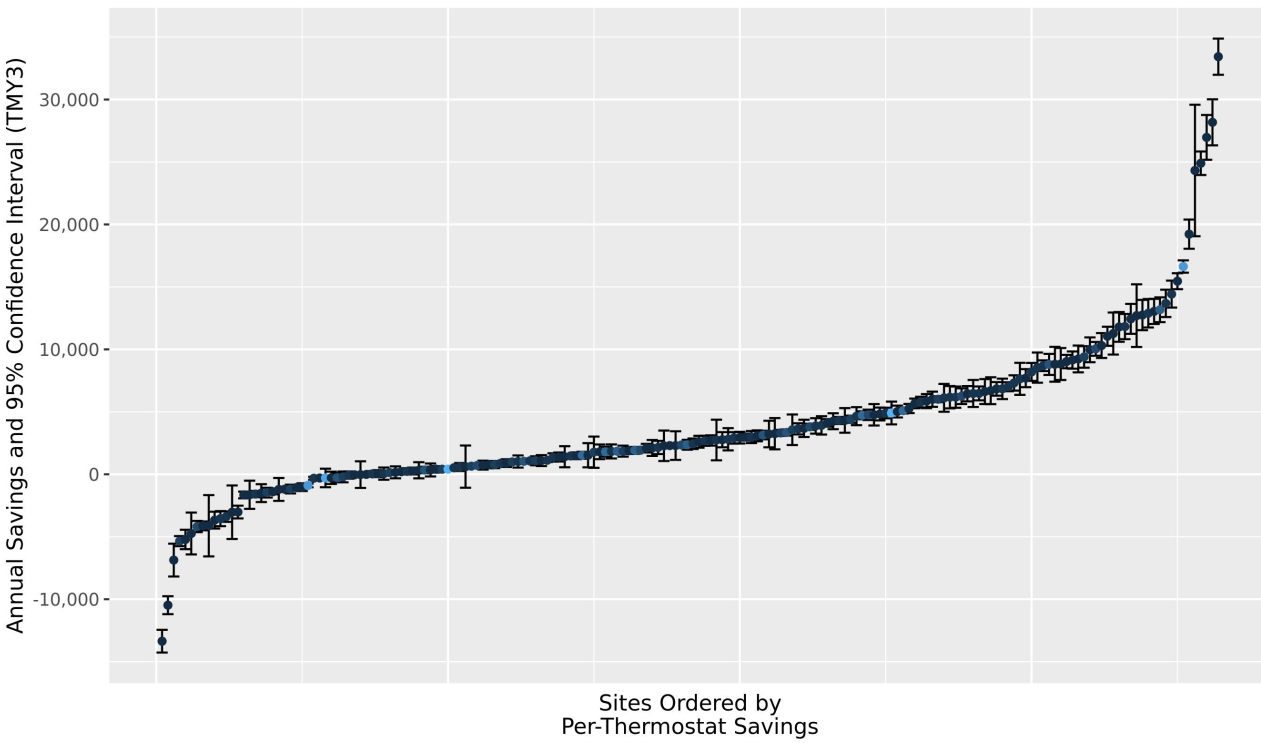

Figure 14: Annual Gas Savings per Thermostat (therms)

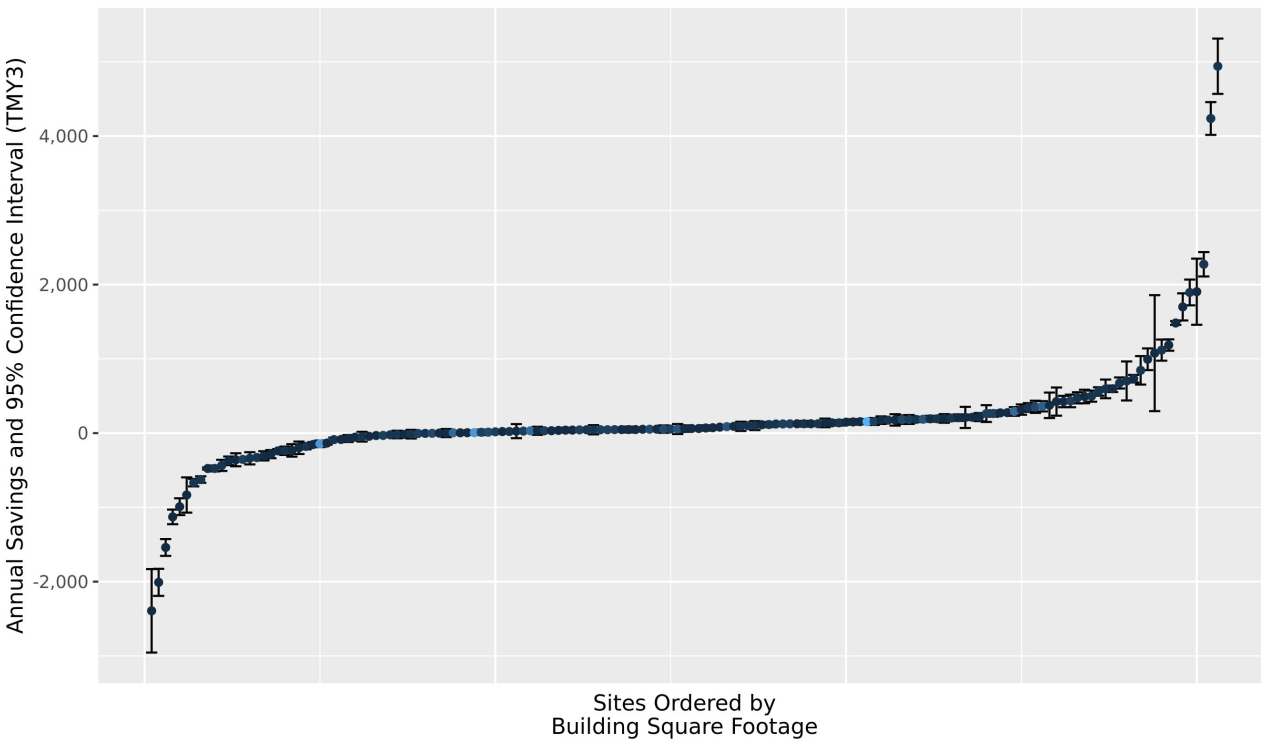

4.3.1 Thermostat Characteristics

We analyzed how estimated savings varied with the two types of qualified smart thermostats (Ecobee and Pelican) that were utilized by participants. However, there was considerable variation between the businesses that received Pelican smart thermostats and those that received Ecobee smart thermostats. First, in terms of building size, Pelican smart thermostats were generally installed at much larger sites (with an average total size of 16,533 square feet versus 3,089 square feet). Because of this, sites with Pelican thermostats tended to have many more thermostats installed at the site (8.6 thermostats per site versus 1.9 thermostats for Ecobee). On average, each Pelican thermostat served 2,024 square feet of the site whereas each Ecobee served 1,658 square feet. In terms of business type, Pelican thermostats were most common in medium offices and were very rarely installed in food service, small office, and small retail offices. In most instances, Pelicans were installed in larger, more complex buildings with central economizers.

Table 17 (kWh) and Table 18 (therms) summarize savings from the billing analysis for each thermostat type on an annual, per-thermostat basis. “Adjusted Prediction” accounts for NREs, while “Original Prediction” does not. For all sites included in this analysis, only one type of thermostat was installed, regardless of the total number of thermostats installed. Also, the majority of sites included in this analysis (173 of 182 for electric, 149 of 153 for gas) had Ecobee smart thermostats installed. Electric savings were significant for both types of thermostats. Baseline usage and kWh savings varied dramatically between the two thermostat types. For gas, savings for Pelican thermostats were large and significant, but were insignificant for Ecobee thermostats. Given the limited number of sites with Pelican thermostats (n=9 electric and n=4 gas), it is difficult to draw broader conclusions about Pelican thermostats from this sample.

Table 17: Annual Electric Thermostat Savings by Model (kWh)

Thermostat	Sites	T-stats	Original Prediction	Adjusted Prediction	Actual	Annual kWh Savings (per thermostat)	% Savings
Ecobee EMS-SI	173	333	30,204	31,248	27,201	4,048 ± 3,003	13.0% ± 9.6%
Pelican TS-200	9	77	16,681	14,636	11,653	2,983 ± 1,374	20.4% ± 9.4%

Table 18: Annual Gas Thermostat Savings by Model (therms)

Thermostat	Sites	T-stats	Original Prediction	Adjusted Prediction	Actual	Annual Therms Savings (per thermostat	% Savings
Ecobee EMS-SI	149	281	1,455	1,439	1,281	158 ± 235	11.0% ± 16.4%
Pelican TS-200	4*	37	989	1,095	879	216 ± 72	19.7% ± 6.5%

* Warning: As this sample has fewer than five sites, it is likely too small to draw meaningful conclusions from.

Based on the audit, we identified two additional explanations for the variation in savings: 1) if a schedule existed before installation and 2) if the new schedule was the same as the previous schedule. While there was no way to assess if the new schedules were the same as the previous schedules (information on existing schedules was not collected during the audit), we do know whether a schedule existed before installation for the sites in the billing analysis.

Table 19 (kwh) and Table 20 (therms) show how the energy savings estimates varied based on whether the customer reported having a schedule on their existing system before the smart thermostat was installed. This table excludes savings for sites where a schedule before the installation was listed as “not applicable” and for sites with multiple thermostats and inconsistent existing conditions (i.e., only a portion were scheduled). Savings were higher in terms of kWh for sites that did not have an HVAC schedule prior to the smart thermostat being installed. Results were insignificant for gas savings.

Table 19: Annual Electric Thermostat Savings by Existing Schedule (kWh)

Existing Schedule	Sites	T-stats	Original Prediction	Adjusted Prediction	Actual	Annual kWh Savings (per thermostat)	% Savings
Yes	99	217	25,706	27,505	23,934	3,571 ± 3,512	13.0% ± 12.8%
No	53	78	37,803	37,464	32,763	4,701 ± 2,284	12.5% ± 6.1%

Table 20: Annual Gas Thermostat Savings by Existing Schedule (therms)

Existing Schedule	Sites	T-stats	Original Prediction	Adjusted Prediction	Actual	Annual Therms Savings (per thermostat)	% Savings
Yes	79	148	1,216	1,366	1,185	180 ± 271	13.2% ± 19.8%
No	54	96	1,920	1,639	1,526	113 ± 226	6.9% ± 13.8%

4.3.2 Building Characteristics

Figure 14 (kWh) and Figure 15 (therms) show the annual site-level savings estimates and confidence intervals for the sites in this analysis by the building floor area in square feet. These point estimates have also been colored by the total number of thermostats installed at each site. While there is a general trend toward both greater savings and wider error bounds at larger sites, this pattern is not consistent, especially for gas. Trends in gas savings are similar to electric; however, fewer of the savings estimates by building size are statistically significant.

Figure 14: Annual Site-Level kWh Savings by Square Footage of Building

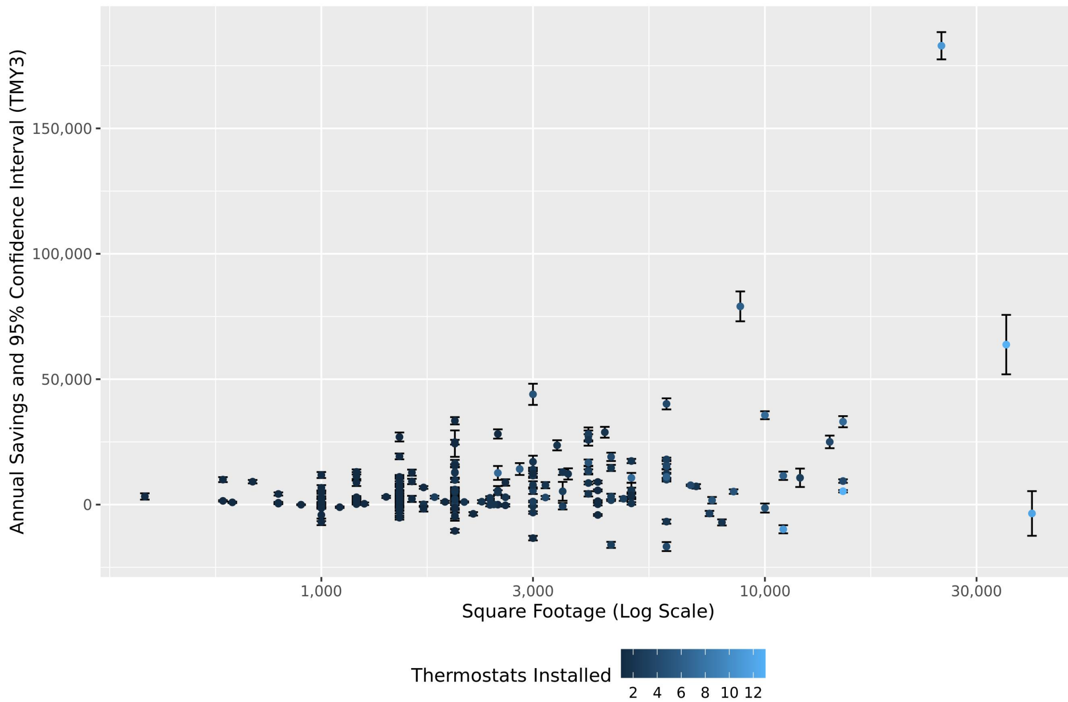

Figure 15: Annual Site-Level therms Savings by Square Footage of Building

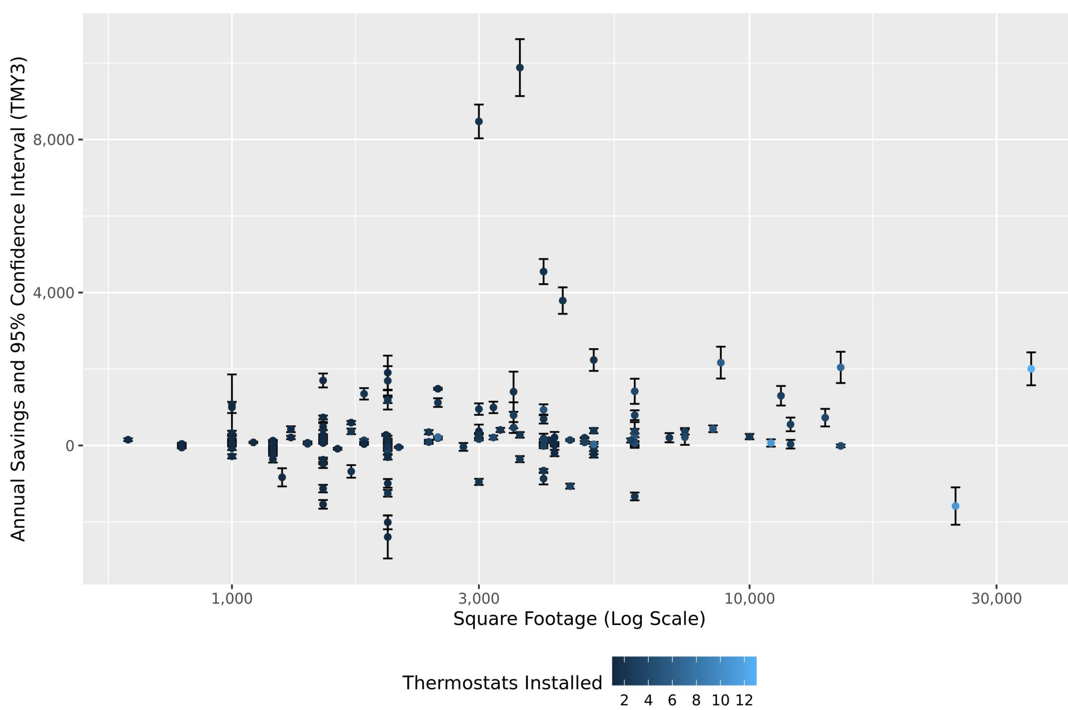

Next, we estimated thermostat-level kWh savings split by total HVAC cooling capacity (in tons) per site and heating capacity (in BTU) per site for the systems with smart thermostats installed. Table 21 provides our estimated energy savings by the total cooling capacity of the HVAC equipment with 95 percent confidence intervals.[17] There were no sites in the sample with between 70 and 149 tons of cooling capacity, and three sites did not provide any data on cooling capacity during the audit. As with building size, results for cooling tonnage tended to be mixed, with consistently significant savings, but without consistent trends across tonnage categories especially when considering the small sample sizes of certain tonnage categories.

Table 21: Annual Electric Thermostat Savings (kWh) by Total Cooling Capacity

Total Cooling Tonnage	Sites	T-stats	Original Prediction	Adjusted Prediction	Actual	Annual kWh Savings (per thermostat)	% Savings
0 - 10	138	196	32,834	33,891	29,752	4,139.0 ± 2,122	12.2% ± 6.3%
11 - 20	27	89	16,736	18,297	14,763	3,533.8 ± 1,484	19.3% ± 8.1%
21 - 30	9	52	22,000	22,729	18,589	4,140.4 ± 1,358	18.2% ± 6.0%
31 - 150	5	55	*	*	*	*	*
Unknown	3	18	*	*	*	*	*

* We omitted customer groups with fewer than five customers, as the sample is likely too small to draw meaningful conclusions from. The sites in the highest capacity group 31-150 tons were too varied to be useful; these sites had 35, 36, 41, 66, and 141 tons of cooling capacity.

Table 22 (kWh) and Table 23 (therms) provide a similar comparison of energy savings by total heating capacity in BTU for sites with gas heating. The 34 electric sites and 8 gas sites that do not have data on heating capacity are electrically heated. As with cooling capacity, kWh savings are significant for each heating capacity bin, but there is no broad trend across bins. For gas savings, on the other hand, only buildings with 200,001 to 400,000 total heating BTU were observed to have large and significant gas savings.

[17] We did not include comparable analysis for gas savings as we do not expect a relationship between cooling capacity and gas savings.

Table 22: Annual Savings per Thermostat by Total Heating Capacity (kWh)

Total Heating BTU	Sites	T-stats	Original Prediction	Adjusted Prediction	Actual	Annual kWh Savings (per thermostat)	% Savings
0 – 200,000	95	123	34,113	35,794	31,689	4,106 ± 2,210	11.5% ± 6.2%
200,001 – 400,000	38	101	21,926	22,760	19,080	3,680 ± 1,404	16.2% ± 6.2%
400,001 – 600,000	8	42	15,406	15,727	13,573	2,155 ± 948	13.7% ± 6.0%
600,001 – 800,000	7	67	14,287	13,973	10,079	3,894 ± 1,687	27.9% ± 12.1%
N/A, electric heat	34	77	31,778	30,920	26,429	4,491 ± 3,040	14.5% ± 9.8%

Table 23: Annual Saving per Thermostat by Total Heating Capacity (therms)

Total Heating BTU	Sites	T-stats	Original Prediction	Adjusted Prediction	Actual	Annual Therms Savings (per thermostat)	% Savings
0 – 200,000	105	141	1,397	1,454	1,339	114 ± 255.6	7.8% ± 17.6%
200,001 – 400,000	30	84	1,128	1,221	923	298 ± 206.8	24.4% ± 16.9%
400,001 – 600,000	7	35	598	598	559	39 ± 38.9	6.5% ± 6.5%
600,001 – 800,000	3	29	*	*	*	*	*
N/A, electric heat	8	29	4,293	2,955	2,573	382 ± 246.4	12.9% ± 8.3%

* We omitted customer groups with fewer than five customers, as the sample is likely too small to draw meaningful conclusions from.

We also analyzed how savings varied by business type, as shown in Table 24 (kWh) and Table 25 (therms). Electric savings rates varied considerably across business type, although savings for all business types were significant. Outside of “Other” businesses, kWh savings were highest on a percentage basis in schools and medium offices, and highest in kWh terms in food service businesses. Retail (both small and medium) and small offices tended to have lower savings. For gas, food service had especially high savings in terms of therms, while hospitality (including non-hotel) as well as medium offices had the highest percentage savings. A number of business types had insignificant gas savings, including small retail, which was the only business type to show negative gas savings.

Table 24: Annual Electric Savings by Business Type (kWh)

Business Type	Sites	T-stats	Original Prediction	Adjusted Prediction	Actual	Annual kWh Savings (per thermostat)	% Savings
Food Service	54	83	44,510	45,539	39,049	6,490 ± 2,942	14.3% ± 6.5%
Grocery	2	9	*	*	*	*	*
Hospitality Non-Hotel	7	12	25,685	28,090	23,781	4,309 ± 2,765	15.3% ± 9.8%
Logistics	2	5	*	*	*	*	*
Medium Office	20	104	12,507	12,735	10,531	2,204 ± 911	17.3% ± 7.2%
Medium Retail	5	12	16,215	15,987	14,466	1,520 ± 784	9.5% ± 4.9%
Other	10	45	18,669	19,993	13,488	6,505 ± 3,448	32.5% ± 17.2%
School	10	39	18,979	19,370	15,528	3,842 ± 1,494	19.8% ± 7.7%
Small Office	38	59	14,354	15,148	13,485	1,663 ± 1,228	11.0% ± 8.1%
Small Retail	34	42	39,724	40,658	37,391	3,266 ± 2,327	8.0% ± 5.7%

* We omitted customer groups with fewer than five customers, as the sample is likely too small to draw meaningful conclusions from.

Table 25: Annual Gas Savings by Business Type (therms)

Business Type	Sites	T-stats	Original Prediction	Adjusted Prediction	Actual	Annual Therms Savings	% Savings
Food Service	44	73	3,259	3,205	2,697	508.2 ± 421.75	15.9% ± 13.2%
Grocery	1	4	*	*	*	*	*
Hospitality	1	2	*	*	*	*	*
Hospitality Non-Hotel	6	12	472	668	358	310.0 ± 120.46	46.4% ± 18.0%
Logistics	2	5	*	*	*	*	*
Medium Office	12	57	553	553	453	99.7 ± 65.51	18.0% ± 11.8%
Medium Retail	2	5	*	*	*	*	*
Other	13	47	772	773	705	68.1 ± 134.70	8.8% ± 17.4%
School	9	20	979	985	923	62.1 ± 63.67	6.3% ± 6.5%
Small Office	46	72	472	465	438	27.5 ± 64.81	5.9% ± 13.9%
Small Retail	17	21	1,347	1,321	1,585	-263.7 ± 276.89	-20.0% ± 21.0%

* We omitted customer groups with fewer than five customers, as the sample is likely too small to draw meaningful conclusions from.

Figure 16 shows the load shape impacts for the most common business types in the sample. As with the previous load shape figures, our prediction is based on the pre-period model and post-period weather data for each season; it represents the expected load shape for these customers in absence of program pilot participation and has been adjusted to account for NREs. The error of each hourly prediction is depicted as a 95 percent confidence interval in the shaded area around each estimate. Whenever the actual post-period load shape (blue line) falls outside the predicted post-period load region (red area), this indicates that a statistically significant change was observed during that hour. The AMICS model finds statistically significant reductions in the whole-building energy usage for pilot participants across all hours of the day for the food service and medium office business types. The small retail business type had relatively low savings with a few insignificant hours, while the small office type had even lower savings, and many hours were insignificant.

Figure 11: Annual Load Shape Impacts for Select Business Types (average site on an average day)

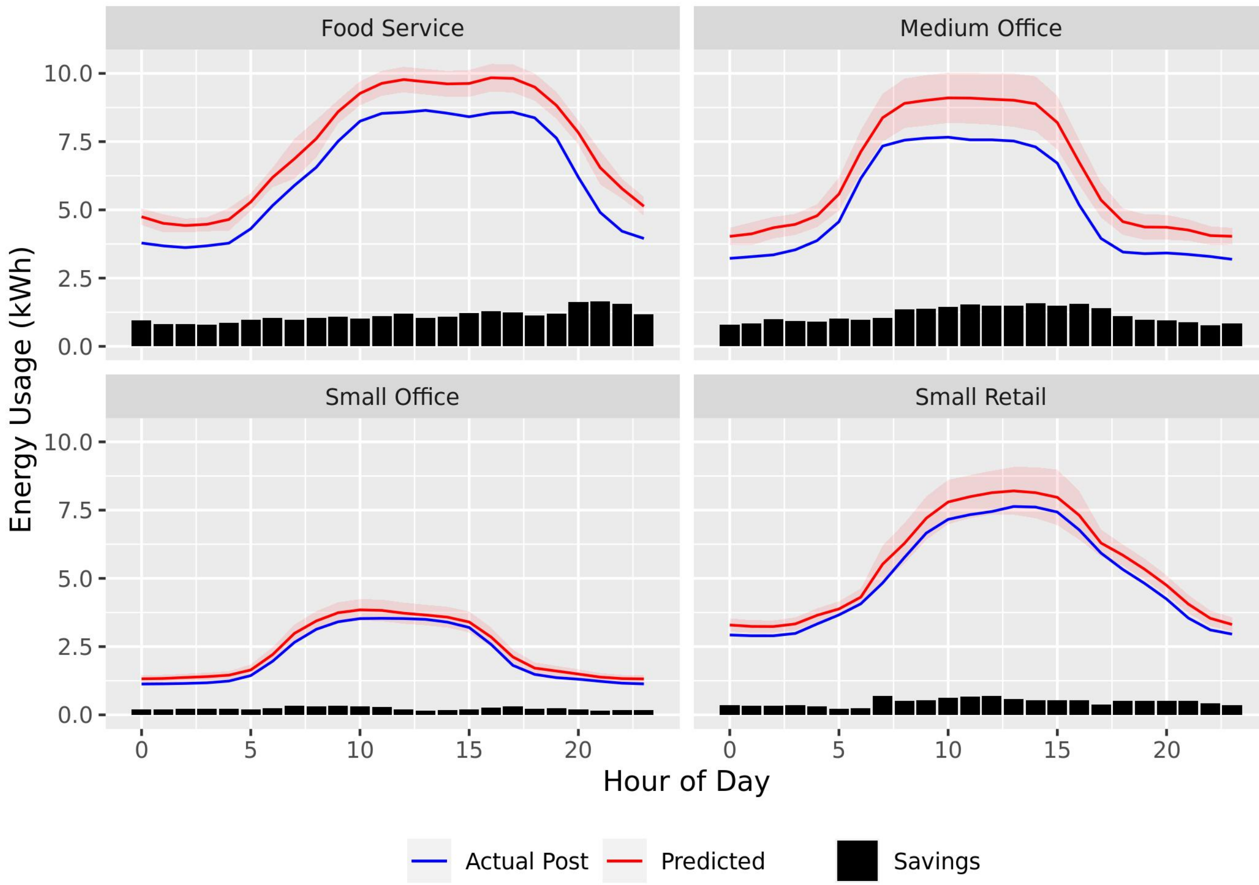

Overall, our analysis suggests that annualized per-thermostat savings, for both therms and kWh, do not correlate with many building type factors. The one possible exception to this is business type, where considerable variation in savings does exist. While the COVID-19 pandemic may obscure the exact magnitude of these results, our results consistently find that smart thermostats are saving significantly more energy than the *ex-ante* kWh savings estimate in the measure approval document.

4.4 Impacts Over Time

This section provides the estimated impact of the thermostat pilot over time in terms of the estimated energy savings from the billing analysis.

In addition to developing point estimates for thermostat impacts, we also used billing analysis results to assess how savings varied over time. Figure 17 (kWh) and Figure 18 (therms) show the average weekly energy savings per thermostat by primary heating fuel at the site. Results have been aggregated on a weekly basis to show general trends, minimizing the noise introduced by changes in operations throughout the week. The values represent savings during a typical weather year (TMY3) and include our adjustment for the COVID-19 pandemic. Sites with multiple measures and measures with few sites represented (i.e., gas savings for electrically-heated sites) have been removed from this analysis. Electric savings at sites heated with gas appear to be stable with minimal seasonal variation. This suggests that most of the electric thermostat savings in these cases are coming from ventilation or fan motors, rather than from cooling. Electrically-heated sites had much more variability in electric savings throughout the year, with the lowest savings in the fall (Sept-Nov) followed by a rapid increase in savings during two winter months (Dec and Jan). For gas (Figure 18), savings increased dramatically during part of fall and all winter months (Nov-Feb) with a low level of non-zero savings throughout the rest of the year. This is not surprising, as gas savings are expected to come from gas heating, reducing waste through scheduling.

Figure 17: Thermostat Electric Weekly Savings (kWh) by Primary Heating Fuel

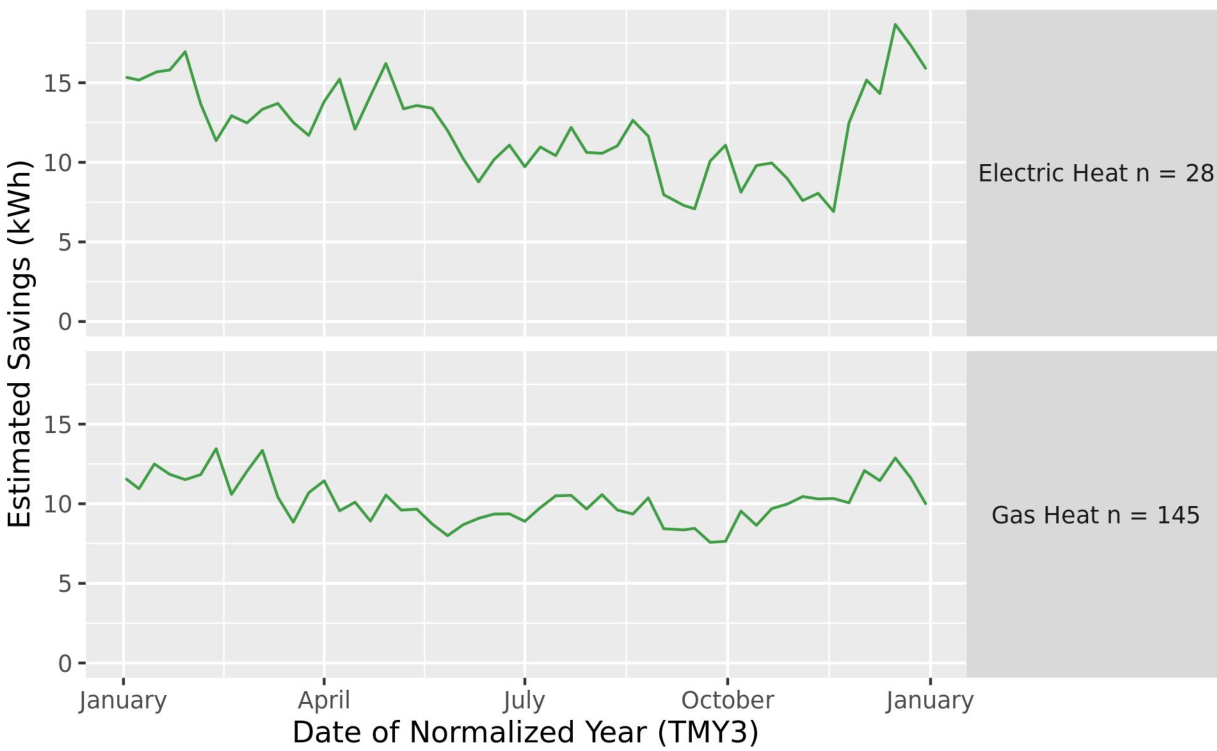

Figure 18: Thermostat Gas Weekly Savings (therms) for Gas Heat

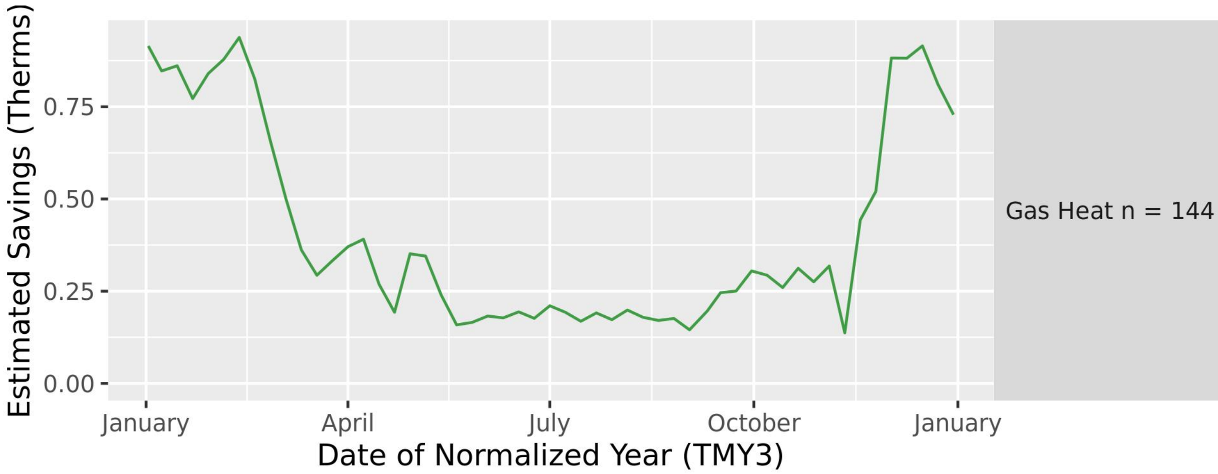

5 Conclusions

Smart thermostats are theorized to save energy in commercial buildings through temperature setbacks, improving fan mode scheduling, and adjusting settings during unoccupied hours.

Our energy and demand savings estimates from the billing analysis confirmed that the smart thermostats installed by Portland General Electric (PGE) that also qualified for Energy Trust of Oregon (Energy Trust) incentives led to statistically significant reductions in kWh across all hours of the day, seasons, and day types and statistically significant therm reductions under specific conditions (with insignificant or negative savings under others).

We recommend that Energy Trust continue to incentivize both Ecobee and Pelican smart thermostats for commercial buildings and then conduct a second impact evaluation to update the energy savings in the Measure Approval Document (MAD). Our results consistently find that these devices are saving significantly more energy than the *ex-ante* electric (kWh) savings estimate in the MAD (see Figure 6), with positive though statistically insignificant gas savings.

The COVID-19 pandemic is challenging to control for. We have made every effort to control for the impacts of the pandemic with customized site-level non-routine adjustments. However, due to the variable impact of the COVID-19 pandemic on energy usage in businesses over time (e.g., shelter in place orders, gradual reopening, and later Centers for Disease Control and Prevention (CDC) recommendations for increased ventilation), it is likely that the pandemic was not fully captured by our adjustments and will have some impact on our savings estimates (see Figure 7). The COVID-19 pandemic creates uncertainty around the magnitude of savings. However, **we are confident that electric savings meet or exceed *ex-ante* values**, as this was clearly observed prior to the start of the COVID-19 pandemic at many participating sites as well as at sites whose energy usage was not noticeably impacted by the COVID-19 pandemic. Please refer to the end of this chapter for our recommendations for a second impact evaluation that would address the remaining sources of uncertainty in the savings estimates for smart thermostats.

Below, we provide a summary of the key findings from the impact evaluation as they relate to each of the research questions:

1. *What are the overall energy and demand savings of commercial smart thermostats?*

   Acknowledgement of Limitations: Our analysis of 11 sites that installed the smart thermostats prior to the start of the COVID-19 pandemic confirmed that these smart thermostat devices led to statistically significant reductions in energy usage, which aligned with the *ex-ante* savings listed in the MAD before the start of the pandemic. We found higher electric and gas savings in the first year of the pandemic despite controlling for differences in weather and seasonality, though the difference was not statistically significant. This increase in savings appears to be tapering off, with the second year of the pandemic dropping closer to the pre-COVID savings. Savings estimates presented in this report should be interpreted with caution. Most of the impact of the COVID-19 pandemic has been corrected, but some lingering impacts may remain embedded in the savings estimate. Despite these limitations, we are very confident that these thermostats are achieving at least the *ex-ante* savings values, especially in terms of kWh.

A total of 410 thermostats were installed at the 182 sites included in the electric billing analysis, and 318 thermostats were installed at 153 sites for gas. In most cases, our weather-normalized electric savings estimates per thermostat greatly exceeded the *ex-ante* savings values listed in the MAD, and gas savings were insignificant.

- Most of the thermostats in our sample were installed in non-grocery buildings with gas heat and cooling (n=144 sites and 315 thermostats). They saved an average of 3,429.3 ± 2,356 kWh per year, far exceeding the *ex-ante* savings of 255 kWh. The gas savings were insignificant at 154.8 ± 236.1 therms per year.
- Savings for a grocery building with gas heat and cooling (n=1 site and 5 thermostats) were insignificant for both fuels and may not be representative.
- Savings for non-grocery buildings with electric heat and cooling varied by heat type, with heat pumps saving 4,856.6 ± 1,565.6 kWh per year and resistance heat saving 1,239.0 ± 2,171.4 kWh per year.

The smart thermostats reduced electric demand by 0.52 kW (11%) in sites with electric heat and 0.94 kW (25%) in sites with gas heat during utility peak hours on the hottest days of the year (daily average temperatures between 72°F and 84°F with a high of 98°F). On the coldest days of the year (average daily temperatures between 38°F and 31°F with a low of 24°F), the thermostats reduced demand by 0.62 kW and 0.61 kW (13% and 11%) for sites with electric and gas heat, respectively.

These findings suggest that smart thermostats are a useful measure for both energy and demand savings for commercial sites.

2. *What are the distributions of energy and demand savings by major bins (e.g., weekday afternoons in the winter)?*

We found statistically significant kWh savings in the whole-building energy usage for pilot participants across all hours of the day for each season, day type, and time-of-use period. However, the magnitudes and hours of peak savings varied. Gas savings were statistically significant only in winter.

- The highest electric and gas savings were observed in the winter.

- Smart thermostats do offer electric savings during weekday on-peak hours for both summer and winter, but the greatest savings occur off-peak on Sunday afternoons in winter.
- For sites with gas heat, gas savings were much higher during the winter heating months (Nov-Feb) with a low level of non-zero savings throughout the rest of the year. Electric savings were stable throughout the year, which suggests that the electric savings are from improvements to ventilation rather than cooling load.
- For sites with electric heat, energy savings were the lowest in the fall (Sept-Nov) and then rapidly increase in early winter (Dec-Jan).

Our findings are consistent with the theory that smart thermostats can provide energy savings during peak hours through temperature setbacks in summer months and setups in winter months. Off-peak savings will also come from the thermostats reducing HVAC usage when the buildings are unoccupied, which occurs more often on weekends and overnight.

3. *What are the trends in energy and demand savings over time?*

Electric savings were relatively flat throughout the year, except for an increase in savings during the winter for the subset of sites with electric heating. As expected, gas savings were much higher in the winter months corresponding with gas heat.

4. *What are the energy and demand savings impacts by thermostat manufacturer, thermostat settings, building characteristics (HVAC capacity, floor area, percent conditioned space), and business type?*

Thermostat manufacturer: The majority of smart thermostats included in this analysis were Ecobee EMS-SI (77% of thermostats in the electric analysis and 88% in the gas analysis) with the remainder being Pelican TS-200s. Electric savings were statistically significant for both thermostats although the count of Pelican thermostats was low. Baseline usage and kWh savings varied dramatically between the two thermostat groups, with average savings of 13 percent for Ecobee thermostats and 20 percent for Pelican thermostats, corresponding to annual per-thermostat savings of 4,048 kWh and 2,983 kWh, respectively. The gas savings for Ecobee thermostats were not statistically significant (11% or 158 therms), but Pelican thermostats saved 19 percent or 216 therms per thermostat per year.

It is difficult to draw robust conclusions about whether the observed differences in energy savings are driven more by the physical device, thermostat functionality (with only the Pelican thermostats capable of controlling economizers), square-footage controlled by the thermostat (with Pelicans controlling larger spaces on average), business type, program design, or random chance.

Thermostat settings: While there was no way to assess how the new schedules differed from the previous schedules (as this information was not collected during the audit), participants did report whether a schedule existed before installation. The electric savings estimates were very similar, and gas savings were not statistically significant. The future application programming interface (API) analysis will provide more insight into the similarities and differences in thermostat settings.

Building Characteristics

- **HVAC capacity:** On a per-thermostat basis, there was no consistent relationship between electric savings by total cooling capacity or total heating capacity (BTU) in total fuel savings or as a percentage of baseline energy usage.
- **Floor area:** We did not observe a consistent relationship between floor area and per-thermostat savings. We observed small and/or insignificant gas savings across nearly the full range of building sizes.
- **Percent conditioned space:** We did not explore savings by the percentage of space conditioned because the applicable MAD (version 235.1) lists semi-conditioned spaces as disqualified from receiving Energy Trust incentives.

Business type: Business type was a relatively strong determinant of savings. Electric savings were highest on a percentage basis in schools (n=10 sites) and medium offices (n=20), and highest in kWh terms in food service businesses (n=54). Small retail (n=34 sites), medium retail (n=5 sites), and small offices (n=38) tended to have lower electric savings. In terms of gas savings, food service (n=44 sites) had especially high savings in terms of therms, while non-hotel hospitality (n=6) and medium offices (n=12) had the highest percentage savings.

5.1 Recommendations for Future Research

We recommend that Energy Trust conduct a second impact evaluation in the future to refine the savings estimates by business type and heating type, avoiding much of the uncertainty caused by the COVID-19 pandemic and sample attrition. A future evaluation of these commercial smart thermostats should do the following to reduce sample attrition and limit sources of uncertainty:

- **Allow more time to elapse:** For the most accurate estimates, future evaluations could focus on thermostats **installed at least 12 months after** the end of the shelter-in-place orders (when businesses have settled into a new normal) so that the COVID-19 pandemic will have limited impact on the baseline or post-period of the analysis. Best practice is to require **12 months of post-installation** to observe a full year at every site. The results from this timeframe would provide more concrete estimates of the energy savings attributable to these smart thermostats that will be applicable to a post-pandemic economy. Where possible, additional data going back to before the start of the pandemic could help determine how much participants changed their energy usage at the start of the pandemic, which would make it possible to detect if/when this change stops. Identifying sites where pandemic impacts have normalized will aid considerably in attributing energy usage changes to the program.

- **Develop a matched comparison group:** To further aid in the ability of the evaluation to attribute savings to the program, we recommend matching participants to future program participants to form a comparison group. This may enable future evaluations to control for the pandemic-related energy usage changes that may affect sites either pre or post installation. This should be especially effective in aggregate, while site-level savings estimates may end up being less reliable given that the random usage patterns of two sites will affect savings. Future participants are preferable to non-participants as they have already consented to share data, and they are more likely to match on unobserved characteristics such as the desire for a smart thermostat and a propensity to adopt Pelican versus Ecobee smart thermostats.
- **Reduce attrition with daily models and aggregation:** We may see less attrition with daily models in the billing analysis, as commercial sites are more difficult to predict hour-by-hour than day-by-day. All of the annual and seasonal impacts by customer segment could have been derived from a daily model. An hourly model will still be required for the demand savings and energy savings by time-of-use. To further minimize attrition, we recommend focusing on aggregate impacts for each measure. Site-level models are more efficient for creating breakouts by customer segment, but this comes at the cost of attrition as some customers' energy usage cannot be adequately explained by variations in weather, season, time of day, and day of week.
- **Estimate savings by measure type:** The primary application for these measures appears to be non-grocery buildings with gas heat (n=144 sites). We had a single grocery site with gas heat, which had statistically insignificant negative savings. Our sample size was also too small to investigate trends in savings for sites with electric heat (n=9 resistance and n=19 heat pump) by customer segment. These measures are worth additional investigation
- **Identify required breakout groups:** To ensure that results from the evaluation are useful for updating the MAD and meet all other stakeholder needs, we recommend compiling a list of requested breakout groups (e.g., savings by building size and fuel type) early in the evaluation planning process. The evaluation staff then can monitor the list of incoming participants, waiting for sufficient sample in each of the breakout groups before initiating the start of the evaluation.

EVERGREEN ECONOMICS

SARAH MONOHON
SENIOR CONSULTANT
Office: 971.888.7478
Cell: 253.273.4890

1500 SW 1st Ave., Suite 1000
Portland, OR 97201
monohon@evergreenecon.com
www.evergreenecon.com

MEMORANDUM

Date: October 18, 2022

To: Sarah Castor, Energy Trust of Oregon

From: Sarah Monohon, Evergreen Economics and John Flotterud, Driftless Energy

Re: PGE Commercial Smart Thermostat Pilot Evaluation: Pelican and Ecobee API Analysis

---

This memo provides an addendum to the Portland General Electric (PGE) Commercial Smart Thermostat Pilot Evaluation draft report, adding insights derived from analysis of the application programming interface (API) data from Pelican and Ecobee smart thermostat. This analysis spans the first quarter of 2020 through the first half of 2021.

The first draft of the evaluation report issued in August 2021 included analysis of Pelican thermostat API data from January 2020 through July 2020 as well as data from October 2020 (August and September 2020 data were not included in the dataset).[1] The original Pelican data included 51 thermostats located in 14 facilities with 6 medium offices, 4 schools, 1 small retail, and 3 others.

The updated datasets for Pelican and Ecobee allowed us to compare the first quarter of 2020 to the first quarter of 2021 to identify any differences in the operation of facilities between the period before the COVID-19 pandemic and during the pandemic.

Pelican provided an updated API dataset that included data from October 2020 through December 2021. This additional thermostat data rounded out the 2020 heating season and added the 2021 cooling and heating seasons. There is, however, still a gap in provided Pelican data for August and September of 2020. Our analysis now incorporates data from Pelican thermostats through December 2021 except for two thermostats with data only through March 2021, one with data through April 2021, and two with data through June 2021. These additional data provide more insight into the change of setback/setup temperatures over time and per heating/cooling season.

[1] Evergreen Economics and Driftless Energy. 2021. *PGE Commercial Smart Thermostat Evaluation: Draft Report*.

Similarly, Ecobee provided an updated API dataset that included data from January 2021 through December 30, 2021; however, data were only analyzed through March 31, 2021 as this was sufficient to address the primary objective of characterizing the impact of the COVID-19 pandemic on thermostat settings. This updated analysis spans from January 1, 2020 in the original Ecobee data through March 2021 in the updated Ecobee data. The original Ecobee data included 347 thermostats in 193 facilities (19 medium office, 17 schools, 27 small retail, 1 logistics, 43 food service, 57 small office, 1 grocery, 5 medium retail, 7 hospitality-not hotel, 1 hospitality, 15 other). This analysis retains 317 of the 347 thermostats in the new dataset.

The distribution of the respective setback and setups for each smart thermostat for 2020 and 2021 is shown in Figure 1. We identified minimal change in these distributions when including additional data.

Figure 1: Distribution of Setback/Setup Ranges by Thermostat Type in 2020 and 2021

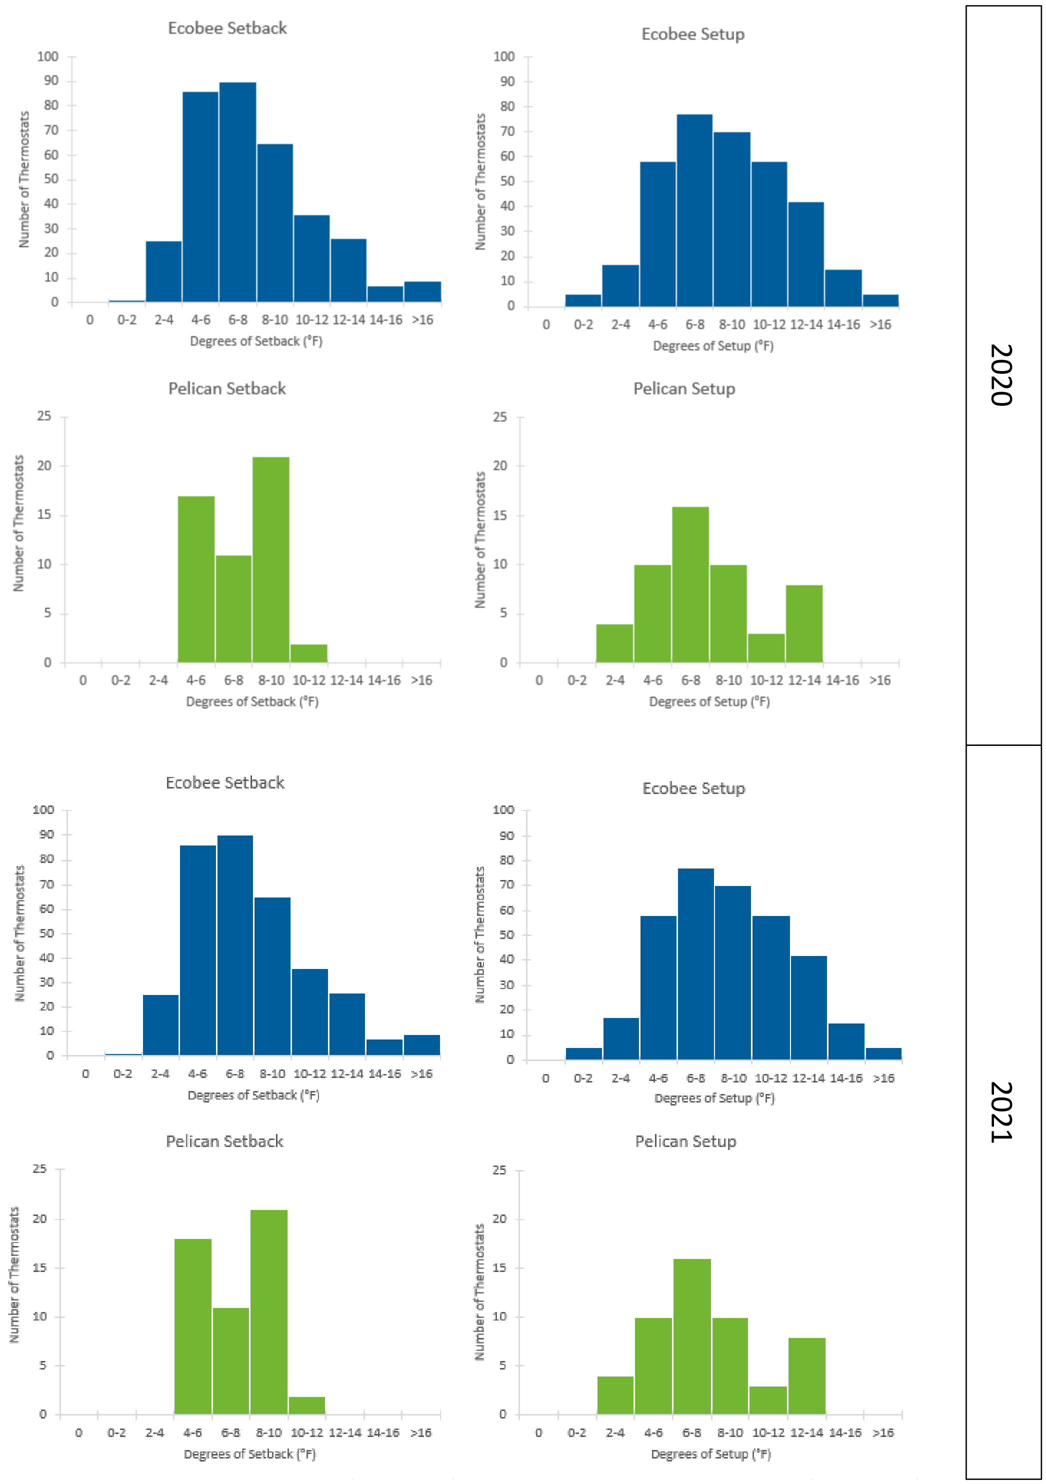

Methods

The methods and qualifying criteria used in the analysis of the additional 2021 data is the same as used in the Portland General Electric (PGE) Commercial Smart Thermostat Pilot Evaluation draft report. The list of qualifying sites was retained, and data sets were not re-evaluated to determine if the additional data would increase the sample of qualifying thermostats. The methods used to develop the qualifying list in the original report are outlined below.

The original API data provided by Energy Trust of Oregon included 120 Pelican and 658 Ecobee smart thermostats. We removed thermostats from the sample that were installed at sites that failed one or more of the following Energy Trust incentive eligibility criteria:

- The site was an existing building (i.e., not new construction);
- The site was a commercial site that had a floor area less of than 200,000 square feet;
- The site installed one of the two approved smart thermostat models (Pelican TS-200 and Ecobee EMS-SI);
- The site was not lodging with 24/7 operations (due to the lack of savings opportunities from setbacks), a semi-conditioned space, or in Heating Zone 2;
- The site was enrolled in the PGE Schedule 25 thermostat program; and/or
- The site was listed as Does Not Qualify (DNQ) by Energy Trust (e.g., manufacturing, non-qualifying gas rate code, mixed use).

We required a minimum of 60 days of API data from each thermostat. The final dataset from Pelican included 8 facilities and 51 smart thermostats, while Ecobee included 193 facilities and 347 smart thermostats that met Energy Trust program eligibility criteria with sufficient data for the API analysis. The Pelican thermostats were largely installed before 2020 with a few added in the first quarter of 2020. The Ecobee thermostats were largely installed up to August of 2020.

The structure and contents of the API data presented some additional challenges:

- The Pelican thermostats, while capable of controlling economizers, do not record any of the **economizer operating parameters** within the API data, preventing the team from verifying both the Pelican's performance with this feature and its compliance to the Measure Approval Document (MAD) requirements.[2]

- Both thermostats include ***primary* heating** time within their auxiliary heat metrics, but we were unable to investigate how often the thermostats implemented ***auxiliary* heat**.
- The Ecobee does not provide an indicator on how the **setpoints were changed**; that is, whether it was a scheduled change or a manual override. The Pelican does include an indicator for whether a setpoint change was initiated on the physical device or online but does not distinguish between setpoint overrides and holds.

The Pelican data document each instance when a setting is changed (including temperature). This method often results in an observation every few minutes, but there are many instances with large gaps in time. It is not possible to discern which of these large gaps (sometimes more than a month) are data losses as opposed to long time spans with no changes in thermostat settings. To limit the influence of these long and possibly erroneous time intervals, we limited the Pelican analysis to **intervals of five minutes or less**.

API Data

The provided API data included numerous variables that were consistent across the Pelican and Ecobee data sets, across both collected time periods. Unfortunately, there were several variables that were missing from one or more sources. The provided API data points are shown below in Table 1 for each of the two data extracts for both Ecobee and Pelican thermostats.

[2] During the audit, the following HVAC system parameters were recorded: HVAC Type (RTU/Package, Split System, Standalone furnace), Fuel type (Gas & AC, Electric Resistance, Heat Pump) nominal heating and cooling size. There is no variable indicating if economizer was present or enabled. Economizer control was not a measure requirement; however, this additional feature has the potential to increase (or decrease) savings for this thermostat compared to others without this feature.

Table 1: API Data Provided by Source

Variable	Data	Ecobee 2020	Ecobee 2021	Pelican 2020	Pelican 2021
HVAC Mode/Run Status	Off, Auto, Cool, Heat, Aux Heat	X	X	X	X
Run Status	Off, Cool, Heat			X	X
Zone Calendar	Auto, Hold, Fan on, Custom	X	X		
System	Auto, Off, Cool, Heat			X	X
Zone Climate/Status	Occupied, Unoccupied		X	X	X
Temperature	Indoor space temperature	X	X	X	X
Heating Temperature	Heating setpoint temperature	X	X	X	X
Cooling Temperature	Cooling setpoint temperature	X	X	X	X
Outdoor Temperature	Outside air temperature	X			
Humidity	Humidity		X		
Fan	Auto, On			X	X
Fan	Seconds	X	X		
Aux Status	Off, On			X	X
Aux Heat 1	Seconds	X	X		
Aux Heat 2	Seconds	X	X		
Aux Heat 3	Seconds	X	X		
Comp Heat 1	Seconds	X	X		
Comp Heat 2	Seconds	X	X		
Setback	Off, On			X	X
Setby	Schedule, Remote, Station			X	X

The analysis approach utilized variables that were consistent across both thermostat types. This impacted our methodology for determining occupied and unoccupied setpoints. The Ecobee data were missing the Occupied/Unoccupied variable in the first data set and did not include the setback variable provided by Pelican. The occupied and unoccupied temperatures were determined by looking for the average heating and cooling temperature each month and categorizing the values based on being greater than or less than the average temperature. For example, if the average heating temperature is 65, the setpoint value of 68 is categorized as occupied while the setpoint of 60 is categorized as unoccupied.

The Pelican data did not have any indicator for “Holds” nor any data related to economizer operation.

Temperature Setpoints

The Pelican and Ecobee smart thermostats are theorized to save energy by improving the efficiency of the HVAC operation through programming multiple temperature set points and schedules. These thermostats use four separate temperature setpoints:

1. Occupied heating
2. Occupied cooling
3. Unoccupied heating
4. Unoccupied cooling

Using separate set points for heating and cooling during occupied hours reduces the likelihood that the HVAC will turn on, as there can be a wide range of acceptable temperatures. For instance, instead of setting the thermostat to 70° Fahrenheit (F), one could have the system heat to 65°F and cool to 75°F. If the existing thermostat had a single setpoint instead of separate heating and cooling setpoints, then we would expect to see savings from the installation of the smart thermostat. More than half (63%) of the existing thermostats had schedules at the time of the audit, but there is no additional information on the schedule or set points.

The two unoccupied setpoints are designed to allow the indoor temperature to be more variable when occupant comfort is not a concern. For example, one might choose a reduced heating setpoint of 55°F (a setback) and an increased cooling setpoint of 85°F (a setup) to prevent the building from becoming extremely hot or cold, while cutting back on energy waste. Individual businesses will have varied preferences for all four temperature setpoints. Fortunately, as long as the thermostats are being set back (heating) and set up (cooling) during unoccupied hours (and the building does have unoccupied hours), we would expect to see some savings. The scheduling functionality is feasible with a programmable thermostat. The smart thermostats add features like remote access, more custom scheduling options, varying duration of holds, and integration with utility demand response programs. One of the expected savings mechanisms for smart thermostats are from limited duration overrides, to avoid permanent setpoint changes in response to temporary discomfort.

Pelican Smart Thermostats

The Pelican smart thermostats in our sample had average setbacks of 7.3 ± 1.7°F and setups of 8.0 ± 3.6°F from January through July of 2020 and in October 2020. Figure 2 shows that the decrease in Pelican setbacks that we had observed from March to July 2020 was temporary, and that the October 2020 setpoint we observed was more representative of where the setback would hold for the winter. The average setback of the Pelican thermostats through the end of 2021 was 7.3°F, and the average setup was 8.1°F.

Figure 2 shows that the setback/setup levels appeared to change in a cyclical nature, with increased setup over the summers of 2020 and 2021, which was reversed by October each year. Likewise, it appears that the setback increased in the winter and the setup decreased slightly in the winter.

Figure 2: Pelican Smart Thermostat Setback and Setup

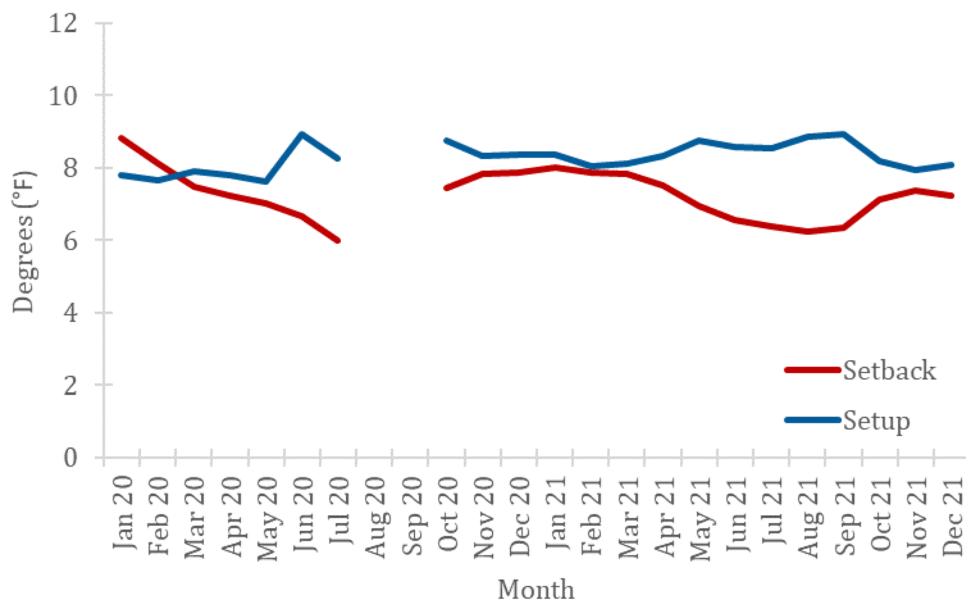

Setback and setup are calculated as the difference between multiple setpoint temperatures: the difference between the occupied and unoccupied setpoint for heating (setback) as well as the occupied and unoccupied setpoint for cooling (setup). These temperatures are shown in Figure 3 and Figure 4. The goal for this analysis was to understand if the seasonal changes in setpoints were due to changes in the occupied or unoccupied setpoints, or both.

The occupied heating setpoint shown in Figure 3 varied the most over time, increasing and decreasing with the seasons with an overall decrease in the setpoint. The unoccupied setpoint has been very stable, with an overall slight downward trend over the two-year analysis period. The occupied setpoint varied with the heating season, with the setpoint increasing during the winter and decreasing in the summer. This suggests that the occupants increased the temperature settings during occupied hours to improve comfort as the ambient outdoor temperatures got colder. They did not adjust the unoccupied setpoint, leading to an overall increase in the average setback. This combination led to the increased setback observed during the winter months, as the setback is based on a comparison between the occupied and the unoccupied setpoints. In other words, the increased setback in winter months was not due to reducing the temperature at night but was due to increasing the occupied setpoint temperature during the day.

Figure 3: Pelican Heating Setpoints

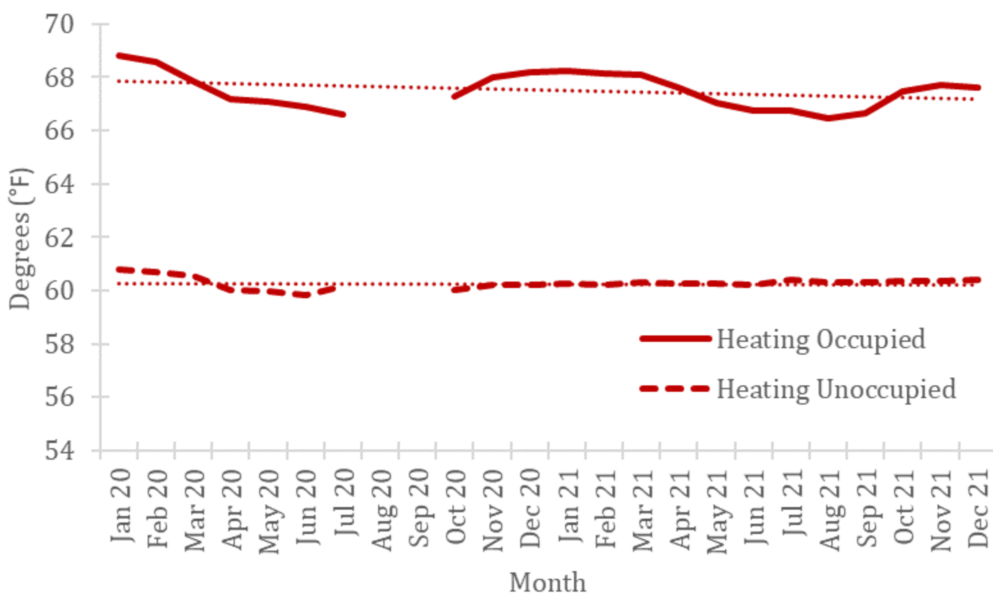

The cooling occupied and unoccupied setpoints shown in Figure 4 both decreased over time, which suggests an increase in cooling and a corresponding increase in energy usage (i.e., a reduction in energy savings) over time. The cooling setpoints also appear to have varied over time, with the greatest change occurring during the occupied period. The occupied setpoint decreased slightly faster than the unoccupied setpoint, resulting in a small increase in the average setup over time. The average setup increased from 8.0°F in 2020 to 8.1°F in 2021.

Similar to the heating setback trends, the cooling setups increased during the summer months, and this change was due to the lowering of the occupied setpoint during the day and not due to an increased setting back of the unoccupied cooling setpoint.

Figure 4: Pelican Smart Thermostat Cooling Setpoints

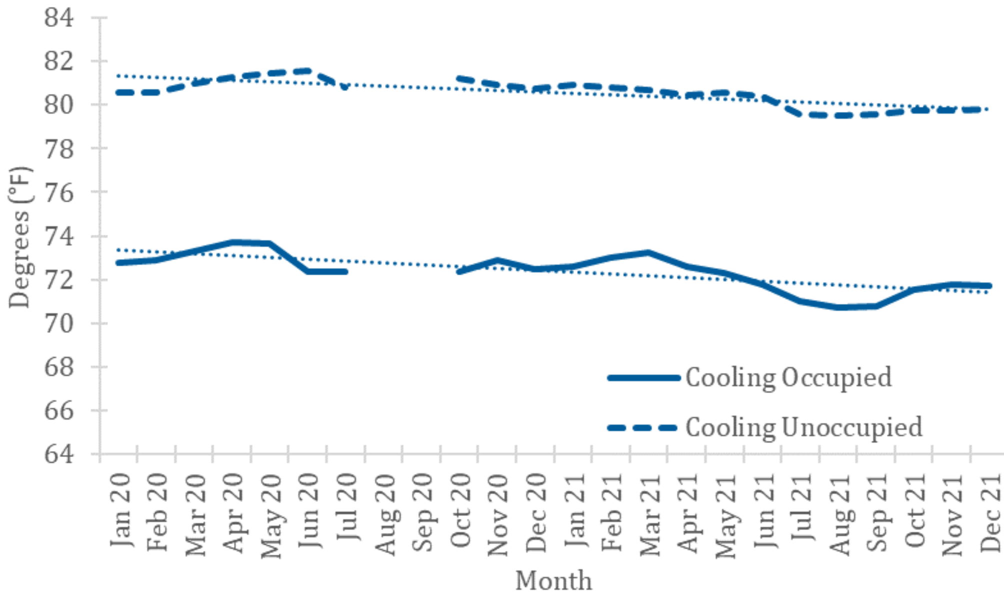

Ecobee Smart Thermostats

The Ecobee thermostats in our sample had average setbacks of 8.0 ± 3.7°F and setups of 8.7 ± 3.9°F from January through December of 2020. Figure 5 shows that after the initial drop in Ecobee setback observed in January and February of 2020, the setpoint stabilized to only a small decrease in setback over time. The average setback of the Ecobee thermostats through March 31, 2021 was 7.7°, and the average setup was 8.8°.

Figure 5: Ecobee Smart Thermostat Setback and Setup

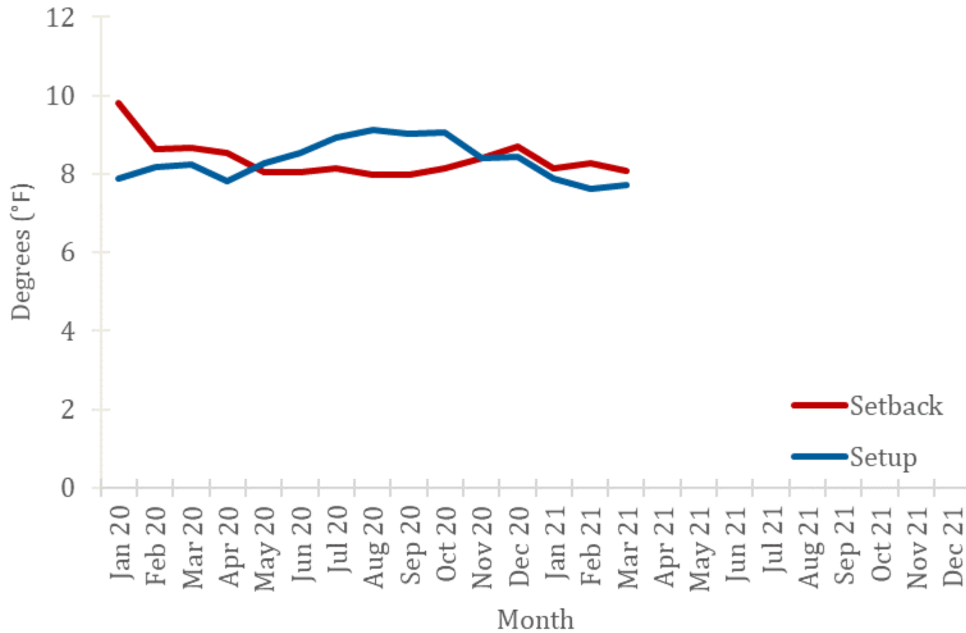

As mentioned earlier, setback and setup are calculated as the difference between multiple setpoint temperatures: the difference between the occupied and unoccupied setpoint for heating (setback) as well as the occupied and unoccupied setpoint for cooling (setup). These temperatures are shown in Figure 6 and Figure 7. The goal for this analysis was to understand if the seasonal changes in setpoints were due to changes in the occupied or unoccupied setpoints, or both.

The occupied heating setpoint shown in Figure 6 has been very stable over the entire dataset. The unoccupied setpoint has been trending upwards over the entire dataset and unlike in the Pelican data where the changes were corrected after each season, the Ecobee unoccupied setpoint has been consistently increasing, resulting in reduced savings. This has decreased the setback (the difference between the occupied setpoint [solid line] and unoccupied setpoint [dotted line]) over time from 8.0°F down to 7.7°F in 2021.

Figure 6: Ecobee Smart Thermostat Heating Setpoints

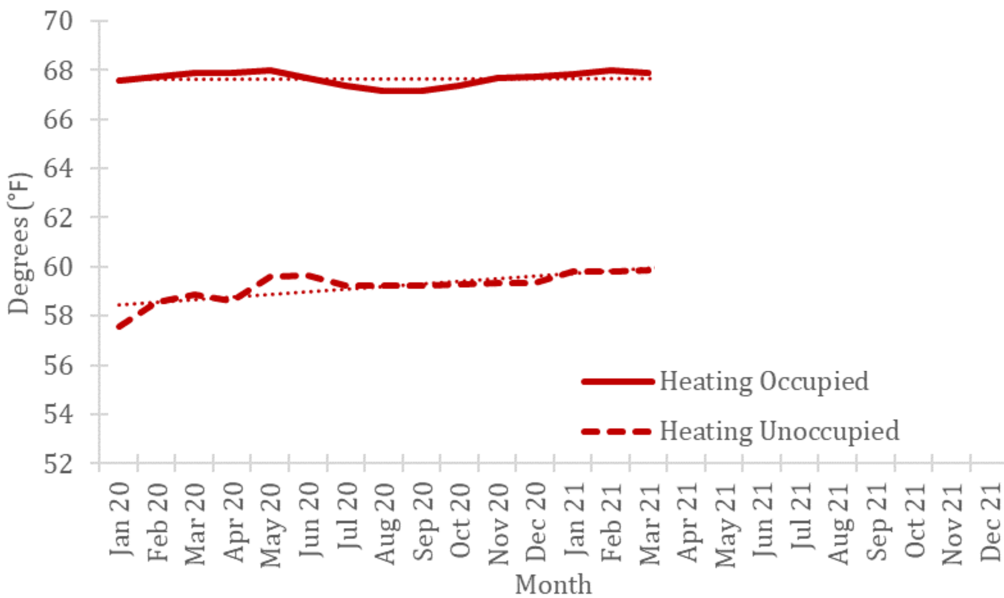

The cooling occupied and unoccupied setpoints shown in Figure 7 both decreased over time, which suggests an increase in cooling and a corresponding increase in energy usage (i.e., a reduction in energy savings) over time. The cooling setpoints also appear to have varied over time, with the greatest change during the occupied period, similar to what we observed in the Pelicans. The occupied setpoint decreased slightly faster than the unoccupied setpoint, resulting in a small increase in the average setup over time. The average setup increased from 8.7°F in 2020 to 8.8°F in 2021.

The Ecobee results for the setback are different from the results observed for the Pelican data and the setup data for Ecobee. The Ecobee occupied heating setpoint was found to be very stable with only small variations in setpoints between the heating and cooling seasons, while the Pelican and Ecobee setups observed the most variation in the occupied setpoint between the heating and cooling season. The unoccupied setpoints were found to be very consistent with nearly no changes for the Pelican heating setpoint and the cooling setpoints for both being very consistent with only small decreases in the setpoint over time. The Ecobee unoccupied heating setpoint was found to be consistently increasing over time.

Figure 7: Ecobee Smart Thermostat Cooling Setpoints

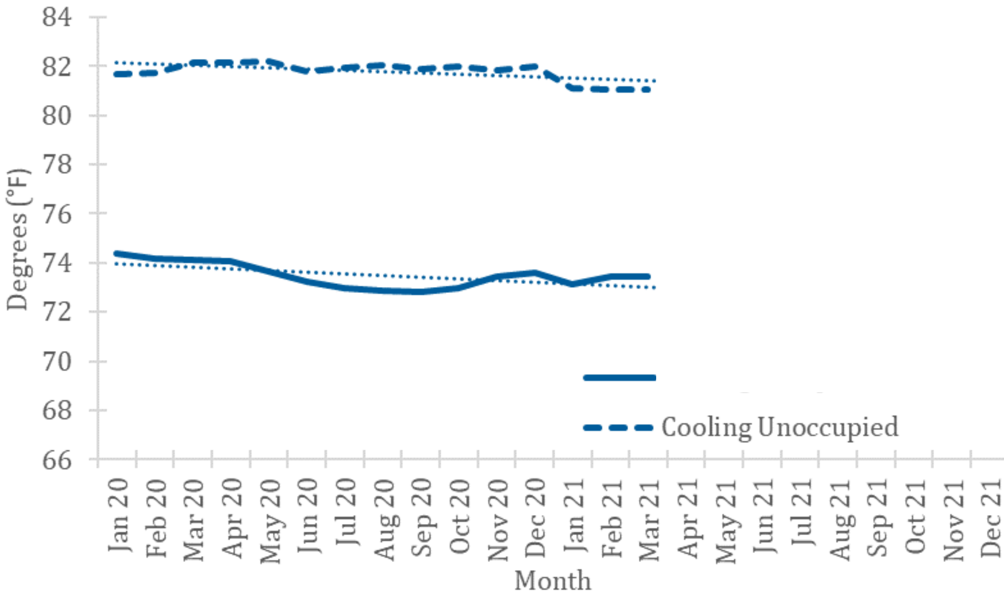

Unoccupied Hours

The overall trend shown in Figure 8 indicates that the proportion of occupied and unoccupied hours of Pelican and Ecobee smart thermostats were relatively consistent over the study period, except for the second quarter of 2020 for Pelican. However, there were some significant limitations in our analysis of unoccupied hours. An Ecobee smart thermostat logs data every five minutes, so any gaps in data from communication or device failures are easily identified as missing data. The Pelican logs data only when changes are made, which leads to sporadic time intervals between observations. We have no way to know whether a long time interval exists because there were no changes in the HVAC system operation or if the lack of data is due to communication or device failure. We imposed a filter on the Pelican data, limiting the time intervals between observations to a maximum of five minutes per data point (i.e., removing observations that are further apart), to be consistent with the Ecobee data for this comparison. If we remove the cap, the percentage of unoccupied hours increases for Pelican and align more closely with Ecobee, though this would introduce more uncertainty from large gaps in data. This increase in unoccupied hours makes sense, as there would be fewer changes during unoccupied hours, leading to more time intervals that exceed the five-minute threshold. This means that unoccupied hours are less likely to be retained in the filtered dataset, reducing the overall average percentage of hours flagged as unoccupied in the Pelican data. Figure 8 should be interpreted as a lower bound for the true percentage of unoccupied hours for Pelicans, with the true value falling somewhere between the values shown for Pelican and Ecobee smart thermostats.

Figure 8: Percent of Unoccupied Hours

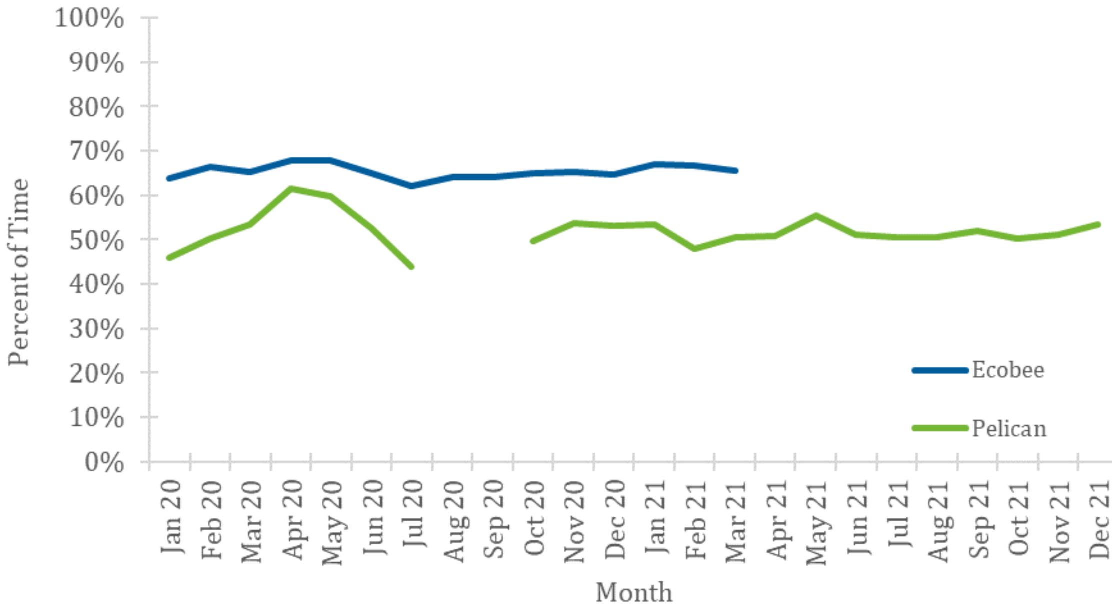

The percent unoccupied hours by building type is shown in Figure 9. There does not appear to be any drastic difference between Ecobee and Pelican thermostats on how they perform with regards to unoccupied hours based on building type.

Figure 9: Average Percent Unoccupied Hours by Building Type

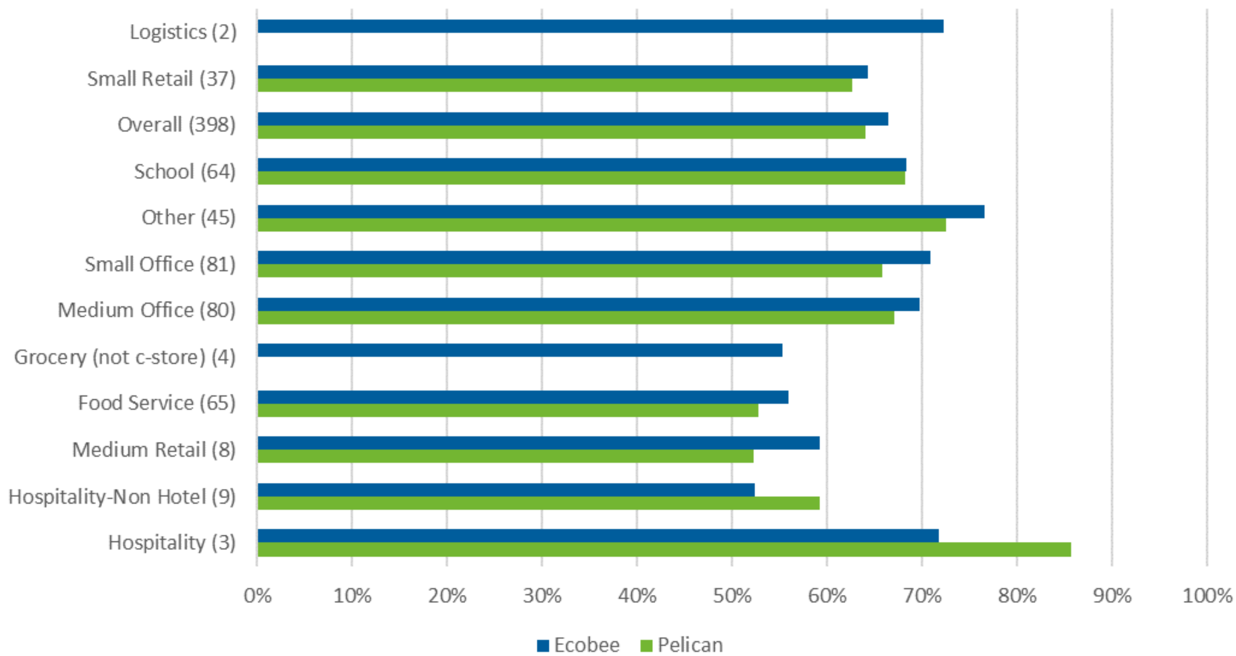

COVID-19 Pandemic

The evaluation of the commercial smart thermostats was complicated by the start of the COVID-19 pandemic in March 2020. The first three months of Pelican and Ecobee API data were logged prior to the first COVID-19 shelter-in-place orders, before there was a large societal response to the pandemic. Once shelter-in-place orders went into effect, the occupancy and use of many commercial facilities were altered, and it is unclear how these changes may have affected the savings attributable to commercial smart thermostats.

Figure 10 through Figure 13 compare the first three months of 2020 (before the COVID-19 pandemic) to the first three months of 2021 (during the COVID-19 pandemic) to look for any meaningful differences between the setpoints. The solid lines show the average setup and setback, while the dotted lines show a standard deviation. We did not observe any substantial or statistically significant differences in overall setback or setup points between the first quarters of 2020 and 2021. The downward slope of the 2020 setback temperatures is steeper than for 2021, though the difference is subtle, not statistically significant, and this reduction only appeared in the first month or two for both Pelican and Ecobee smart thermostats before stabilizing for the remainder of the period. It is hypothesized that some of this initial drop could be due to adjustments being made shortly after the thermostats were installed.

Figure 10: Comparison Between Pelican Setback in Q1 of 2020 and 2021

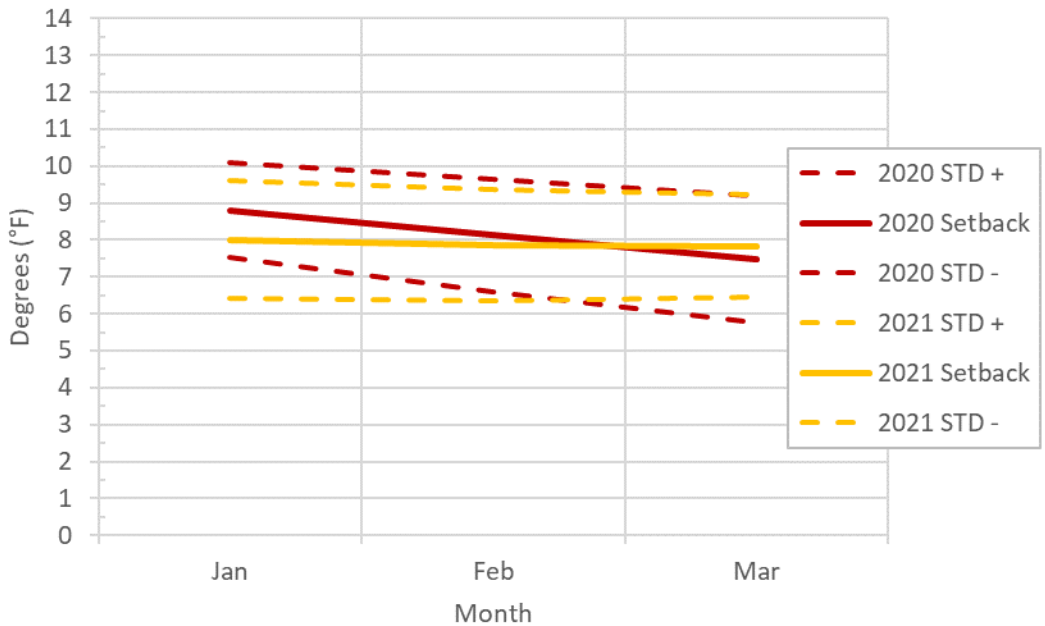

Figure 11: Comparison Between Pelican Setup in Q1 of 2020 and 2021

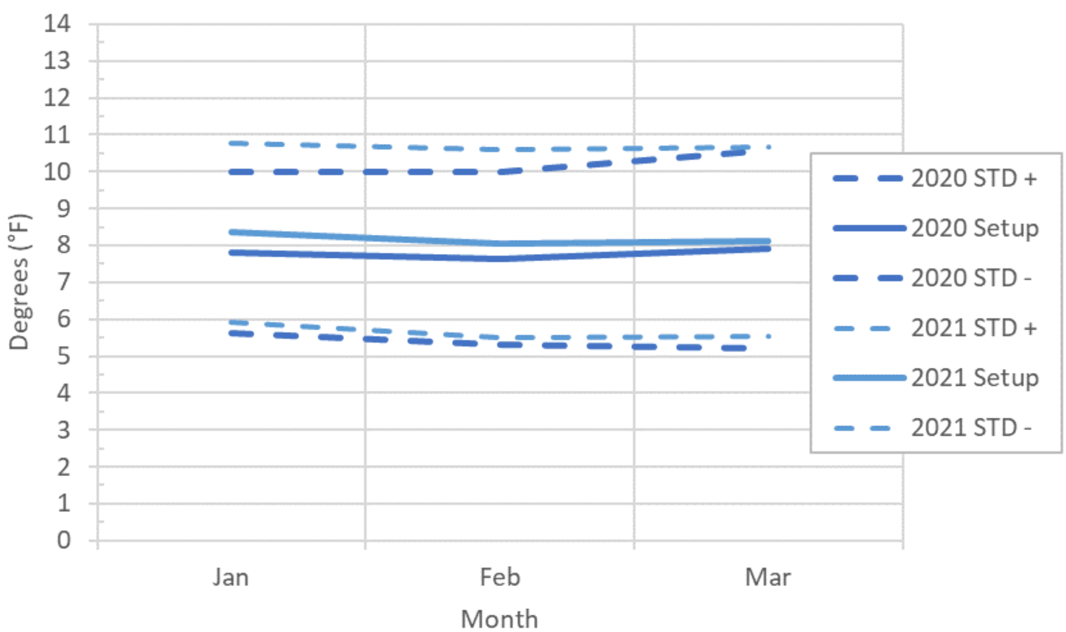

Figure 12: Comparison Between Ecobee Setback in Q1 of 2020 and 2021

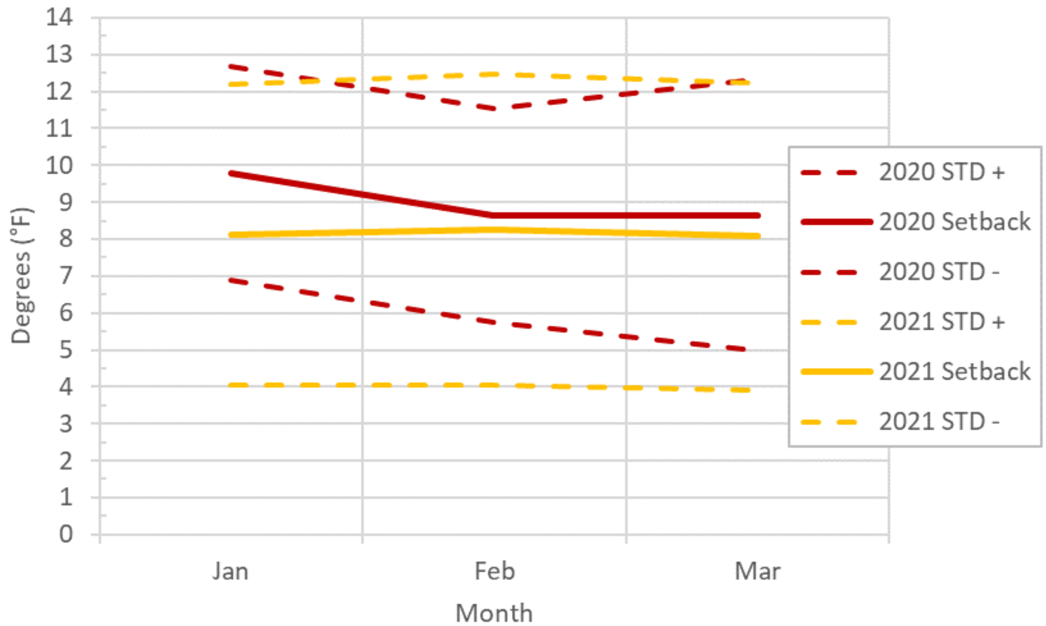

Figure 13: Comparison Between Ecobee Setup in Q1 of 2020 and 2021

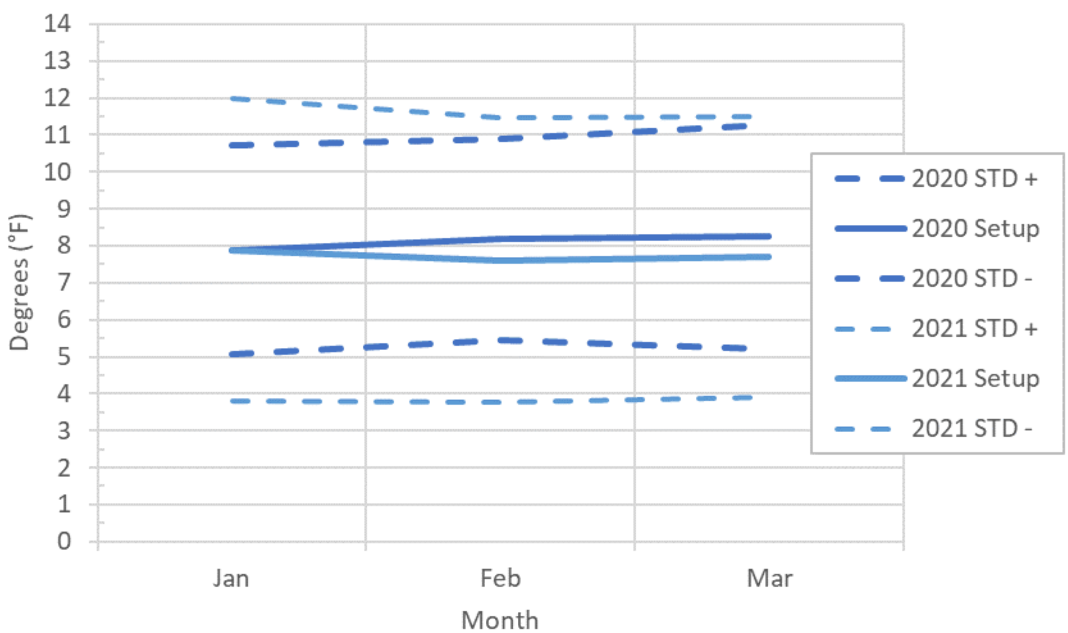

Figure 14 makes the same comparison of the first quarters of 2020 and 2021 for the percentage of unoccupied time. It was expected there was a possibility of unoccupied time increasing once the pandemic took effect as businesses reduced hours of operation or office hours. This, however, was not observed in the sample of thermostats, and the percentage of unoccupied time has remained over the entire dataset.

Figure 14: Comparison Between Unoccupied Time in Q1 of 2020 and 2021

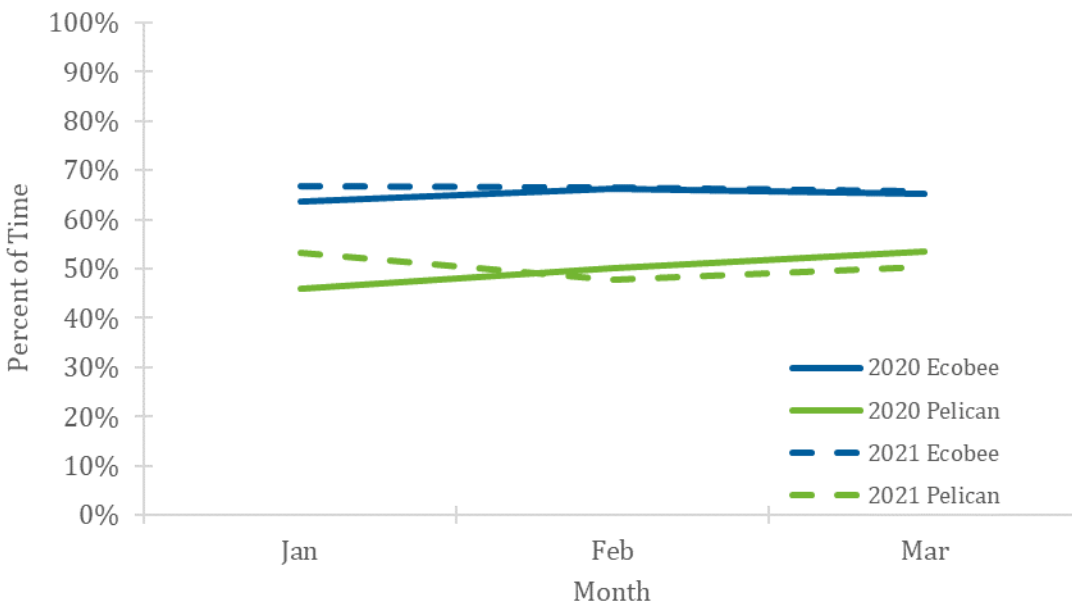

The results of this limited sample of thermostat data suggests that the COVID-19 pandemic had limited effect on the scheduling of the thermostat setpoints. This does not necessarily translate into energy usage, as the setpoints are only one component that influences the operation of HVAC equipment. While the usage and occupancy of buildings likely changed during the pandemic, it does not appear that the setpoints were significantly impacted.

Conclusions

The updated API data from Pelican and Ecobee rounded out the 2020 heating season for Pelican and added the 2021 cooling and heating seasons for both types of thermostats. These additional data made it possible to distinguish some of the seasonal/cyclical changes of setback and setup temperatures from longer term changes in setpoints over time; these trends will impact the magnitude and persistence of energy savings attributable to the commercial smart thermostats.

Below, we provide a summary of the key findings from the Pelican and Ecobee API analysis as it relates to each of the research questions from the broader evaluation:

What are the distributions of energy and demand savings by major bins (e.g., weekday afternoons in the winter)?

> The Pelican and Ecobee smart thermostats in the API analysis logged around 50 to 65 percent of hours as unoccupied, when the setback and setup temperatures are utilized. Furthermore, the percent of unoccupied hours remained consistent over the duration of the study.

What are the trends in energy and demand savings over time?

> The occupied setpoints of the Pelican smart thermostats were adjusted more frequently than the unoccupied setpoints, corresponding to the heating and cooling seasons with occupants adjusting the indoor temperatures to be more comfortable as outside temperatures became more extreme. One of the expected savings mechanisms was from limited duration overrides, but it appears that users were adjusting the setpoint instead of overriding. This would lead to an overall effect of increasing energy usage (i.e., reducing savings) of the HVAC systems during the coldest days of the winter and hottest days of the summer. The unoccupied setpoints were found to be adjusted to much less of an extent. This results in increased setback/setup increasing the energy savings for the temperature adjustment compared to the changes made during the occupied period. The changes to the occupied setpoints during the heating and cooling seasons generally reverted to their original setpoints before the end of the season.
>
> The occupied setpoints of the Ecobee smart thermostats were found to be adjusted less frequently than the Pelican setpoints, but the Ecobee unoccupied heating setpoint steadily increased over the study period. The overall reduction in setback was mitigated slightly by a small increase in occupied heating setpoint; however, increasing either one of these results in increased heating energy usage. The Ecobee setup setpoints were found to be consistent with the Pelican setup setpoints.

Overall, the thermostats performed well and were able to retain their occupied and unoccupied schedules. It is suspected that the ability to select a temporary “hold” or outright prevent permanent “holds” is a key advantage of smart thermostats over traditional programable thermostats. The smart thermostats did, however, see decays in both their setbacks and setups over time. This is due to changes in the occupied as well as the unoccupied setpoints. The performance of smart thermostats is still dependent on the inputs given to them and are susceptible to meddling of these setpoints, highlighting the importance of annual or seasonal review of the setpoints.

While API data is able to tell us what happened to the thermostat settings, it does not provide the context into *why* the change was made, which could be important for savings attribution. This data is also unable to provide context on how the users' operating behaviors may have changed with the new thermostats; whether we are observing pre-existing behaviors or a new behavior that is specific to the smart thermostats. During the application phase, it would be beneficial to collect additional data on the thermostats that were replaced and ask users about their operation and interactions with the existing thermostat as a baseline. Future research could interview users to ask direct questions about their behaviors surrounding adjustment of the thermostat settings and their decision-making processes to provide a deeper understanding of the impact this new technology has had on their behaviors. This exploration would be centered around whether the mechanism for energy savings aligns with our expectations.