Understanding the Role of Weather in Air Conditioning Use Behavior and Demand Response Program Participation (BEH 2)

Revised Draft Final Report

Prepared for:

The Electric and Gas Program Administrators of Massachusetts
Part of the Behavioral Evaluation Program Area

Submitted by:
Navigant Consulting, Inc.
1375 Walnut Street
Suite 100
Boulder, CO 80302

303.728.2500
navigant.com

Reference No.: 183406
October 2, 2018

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

EXECUTIVE SUMMARY

Background and Objectives

The Massachusetts Program Administrators (PAs) engaged Navigant Consulting, Inc. (Navigant) and Illume Advising, LLC (the evaluation team) to study the effects of weather on air conditioner (AC) use behavior and demand response (DR) program participation. Understanding how weather affects customer AC usage will allow the PAs to improve the design, implementation, and customer messaging around these programs. Similarly, understanding the role of weather on a participant's likelihood to opt out of DR events will allow the PAs to balance opt-out rates in calling DR events.

Key study objectives include the following:

1. Develop models of central AC use that show how weather affects the timing of first and last use of AC during a given season and how weather affects AC use throughout the season.
2. Understand the role of weather on DR event participation and, specifically, how the propensity to opt out or override preset thermostat controls varies with weather and program parameters.

This study was undertaken in a staged approach and combined several unique datasets, including advanced metering infrastructure (AMI) data and thermostat telemetry data from National Grid's Smart Energy Solutions (SES) pilot in Worcester, as well as thermostat telemetry data from National Grid's Residential Wi-Fi Thermostat DR demonstration project.

- **Phase I** included a literature review, data management, and preliminary data exploration activities. The literature review yielded important findings from related studies and suggested areas for investigation, particularly with respect to the different ways people use AC. Through these activities the evaluation team came to a solid understanding of the datasets it had to work with in this study, their applicability to answering the research questions posed, initial insights into AC user types, and the types of regressions which might be useful in Phase II.
- **Phase II** built upon the research and analysis performed in Phase I to refine and finalize the methodology to meet the study's objectives. First, the evaluation team was able to define quantifiable AC use behaviors and metrics describing them, using both the AMI and thermostat data, and providing actionable and valuable insights for the PAs. The team then conducted regression analysis to estimate the effects of weather on AC use—both ways in which weather affects the first time and last time residential customers use their AC during the cooling season, and more general AC use behavior throughout the season. Finally, the team combined qualitative interviews with exploratory and regression analysis of thermostat telemetry data to derive meaningful relationships between DR program participant opt-out behavior and weather factors and program parameters.

Phase I Research

Phase I consisted of an in-depth review of existing literature, initial data management, and preliminary exploratory data analysis. The results of this phase were intended to inform the analysis to be performed in Phase II.

Literature Review

The evaluation team reviewed over 20 published sources of information on thermostat and AC use in response to weather and weather effects on DR program participation. These sources included both print and online documents spanning academia, program evaluation reports, conference proceedings, journal articles, and publicly available manufacturer information.[1]

Key Findings

The existing literature suggests AC use is strongly related to temperature and how people interact with their thermostats. A growing body of research categorizes user types, such as “constant on,” “set it and forget it,” “on-off switch,” and “fiddlers.” These findings encouraged us to dig deeper in Phase II to explore metrics quantifying AC use behavior and their relation to user types.

There was little previous research on the relationship between weather (especially extreme weather), incentives, and DR program opt-out rates. The literature suggests a wide range of opt-out rates and while opt-out rates appear to be correlated with weather, few studies make a causal inference and fewer still relate behavior to extreme weather. The wide range of opt-out rates and lack of consensus values in existing literature suggested the team should account for a wide range of potential factors to adequately explain opt-out behavior.

Data Management and Exploration

The evaluation team collected, compiled, and cleaned data from several different sources for this study including AMI and thermostat telemetry data from National Grid’s SES pilot, and thermostat telemetry data from National Grid’s Residential Wi-Fi Thermostat DR demonstration project (Thermostat DR project). Data from the SES pilot was explored in addressing the study’s first objective on AC use behavior. The team’s intent in Phase I was to determine whether the AMI and thermostat telemetry datasets should be used singly or jointly to study the effects of weather on AC use. Data from both the SES pilot and the Thermostat DR project were explored in addressing the study’s second objective on DR program opt-out rates. The team’s intent was to determine which dataset(s) were most appropriate to achieve study research goals.

Key Findings

The SES AMI data contained complete 2014 records for over 15,000 premises, of which 54 had complete cooling season thermostat telemetry data available. Due to the granularity of the thermostat data, the evaluation team determined this data could be used to gain insight into AC use behavior and user types, and to understand the relationship between weather and when customers first turn on and turn off their ACs during the cooling season. Preliminary modeling suggested AMI data alone could be used successfully to extract weather-dependent energy use from total energy use and identify appropriate balance points. As a result, the evaluation team primarily relied on AMI data to analyze the relationship between weather and AC use behavior throughout the cooling season.

The SES pilot data did not contain enough opt outs to be useful in the DR event participation analysis. As a result, the team worked solely with the Thermostat DR program data. This data showed substantial

[1] Navigant attempted to conduct interviews with several thermostat vendors during Phase II, all of whom declined to participate or comment on how their technologies’ control strategies and optimization algorithms interact with weather.

variation in the timing of events and full versus partial opt outs, suggesting Phase II analysis would yield meaningful results.

The data exploration resulted in several other key findings of interest, including:

- In general, the amount of time thermostats are in cooling mode is intuitive, increasing as temperatures increase throughout the summer and then decreasing in the early fall. Few devices are cooling more than 10% of the time before June and after August.
- Some users (4%) never switch into cooling mode and an additional 7% are in cooling mode five or fewer days (120 hours) throughout the summer. A typical device cools for just over 10% of the time during the hottest months of the year, but less than 10% most of the time, and almost no customers fall into the "always on" category, with their AC system operating in cooling mode almost all of the time.
- Customers typically use their AC for the first time in May or June with a wide distribution. However, the distribution for last use was highly concentrated over a couple of weeks, typically in September or October.

These findings helped the team determine its Phase II approaches to modeling first and last AC use, and quantifying different AC use behaviors within the AC user type analysis.

Phase II Research

Phase II of this study built upon Phase I to refine and finalize the methodology to meet the study's objectives and answer these key research questions:

1. How do people use their thermostats, to what extent can customers be identified by different AC user types, and what can the evaluation team learn about AC use by examining specific thermostat use behaviors?
2. How does the weather affect the timing of people's first and last use of their AC systems during the cooling season?
3. How does the weather affect people's AC use throughout the cooling season?
4. How do the weather and program design parameters affect people's participation behavior (i.e., propensity to opt out) in DR programs?

Thermostat Use Behaviors and AC User Types

> How do people use their thermostats; to what extent can customers be identified by different AC user types; and what can the team learn about AC use by examining specific thermostat use behaviors?

Navigant first studied thermostat use behaviors and whether the AC user types identified in the literature review could be identified through the team's research.

The evaluation team analyzed thermostat telemetry data from 54 thermostats with complete cooling season data to explore thermostat use behavior and AC user types. The relevant data points analyzed included mode, state, setpoints, and use of hold settings. Using exploratory data analysis the team, for example, examined the distribution of cooling setpoints and setpoint changes to develop an understanding of thermostat use behavior.

Key Findings and Considerations

Takeaway	Key Findings and Considerations
Message to Behaviors, not User Types	**Key Finding:** There is enormous diversity in how people use their thermostats and AC systems. Sophisticated thermostats allow a broad range of behaviors, and the team's results show that a static "user type" definition may no longer be useful as most customers change their behavior significantly throughout the summer cooling season.
	Consideration: Based on these findings, the team suggests considering focused messaging and targeted behavioral interventions on observed user behaviors rather than user types. Observed behaviors are quantifiable, non-static, and less subjective and arbitrary than user types.
Use Behavioral Norming Based on Actual Massachusetts AC User Data	**Key Finding:** Most customers in the team's sample avoid extreme AC usage behaviors, running their systems a limited amount of time even during the hottest months. Most run their AC less than 10% of the time during the summer months, and almost no customers would be considered "always on" users, choosing setpoints so low their AC is running nearly all of the time. However, a small segment of customers have their AC system running over 25% of the time during the hottest weeks of the summer.
	Consideration: A key behavioral lever is social norming. By letting people know what is "typical" or "normal" in their area, they tend to want to adjust their behavior to conform with the norm. As a result, the PAs can use the information provided in this study on how Massachusetts customers actually use their AC to educate customers about what is normal, thereby causing higher users to dial back consumption. This finding also may be taken into consideration by the PAs to help set expectations for planning around the potential range of savings for thermostat-related program offerings, given relatively low AC use by most customers.
Teach Users who Over-use the "Hold" Setting to Set More Rational Setpoints	**Key Finding:** Many customers use the hold setting excessively, signaling they either have not set comfortable and practical setpoint schedules that fit their lifestyle and the weather, or their thermostat fails to account for humidity in regulating indoor temperature to maintain comfort. Many customers have hold engaged 100% of the time during a given day or week, and the average customer engages hold 50% of the time throughout the summer.
	Consideration: Through targeted messaging, the PAs can help educate people about saving money on their energy bills by setting more reasonable setpoint schedules that keep them comfortable during the heat of the day, rather than relying on lengthy hold periods to keep the air running. Additionally, they may consider revisiting training of auditors to ensure customers are being educated about how to program their thermostats efficiently. Some thermostat models revert to regular schedules when the next scheduled setpoint change is reached rather than remaining in hold indefinitely—a consideration would be to consider prioritizing incentivization of these thermostat models. Additionally, the PAs may consider incentivizing only thermostats that take humidity into account in managing a home's indoor temperature to minimize inefficiency, as much of the unwanted hold behavior may be related to the thermostat's inability to effectively incorporate humidity into its algorithms.

Takeaway	Key Findings and Considerations
Message Differently to those with High vs. Low Average Setpoints	**Key Finding:** One of the most interesting findings on temperature settings was that by mid-summer, there is a bimodal distribution of average temperature settings. In July, as the temperature heats up, a large number of users schedule their average cooling setpoint around 83°F, while another large portion of users set it closer to 73°F.
	Consideration: The PAs can message differentially to these sets of customers to maximize the effect of behavioral interventions. For instance, it does not make sense to message to someone who reliably has their setpoint at 83°F throughout the summer to be more energy-conscious and further set back their thermostat—they are already doing it. However, by identifying the set of users with setpoints in the low 70s to upper 60s throughout the summer, the PAs could nudge these customers to set back their thermostat one or two degrees, thereby saving energy. Additionally, the PAs may consider revisiting training of auditors and installers to ensure they are educating customers about efficient setpoints.
Many Thermostats are Not Smart Enough, so Incentivize Truly Smart Models	**Key Finding:** Current smart thermostats are not smart enough. They allow customers who think the device is a valve to set thermostats at temperatures below what the household has chosen as an optimal temperature. A smart thermostat should have the capability of learning this optimal setting and preventing the user from adjusting the set point below it when indoor temperature is within the household's preferred range. This would prevent inefficient "valving" use. With such a configuration, customers could still reduce the set point if they were uncomfortable below the optimal setpoint, but only by increments of two degrees or less, to prevent erratic behavior and energy efficiency losses. Many smart thermostats also do not measure humidity, and as a result, a static set point may not always provide desired indoor comfort. Smart thermostats should either include a humidity sensing component or develop an algorithm using local NOAA weather conditions to adjust setpoint to reach a temperature *and* humidity-based comfort setpoint.
	Consideration: The PAs may consider incentivizing certain thermostats and not others based on the degree of smartness and key energy-efficiency compatible abilities. For instance, they might incentivize thermostats that only allow customers to reduce their thermostat set point only a couple of degrees lower than their preferred indoor temperature, to avoid wild "valving" temperature swings and associated inefficiency. They may also choose to incentivize smart thermostats that take advantage of indoor humidity level information or local weather station information on humidity to adjust set points to create combined temperature and humidity-based set point adjustments, rather than relying on temperature alone.

Weather and First/Last AC Use

How does the weather affect the timing of people's first and last use of their AC systems during the cooling season?

The evaluation team investigated what motivates users to turn on their AC for the first time in the cooling season and what motivates them to stop using it for the rest of the year. Navigant used graphical data exploration and regression analysis of thermostat telemetry data for 54 devices, paying particularly close attention to what role weather plays in these decisions. For visual analysis the evaluation team relied mainly on correlation plots between first/last AC use and weather variables. The team then performed a logistic regression analysis to understand the magnitude of causal effects and to discern more complete relationships between variables of interest.

Key Findings and Considerations

Takeaway	Key Findings and Considerations
Nudge People to Turn on their AC Later in the Season	**Key Finding:** People are responding to the temperature when deciding to turn on the air for the first time during the season, and not humidity. While humidity plays a role in ongoing AC use, the decision to begin using AC during the summer is primarily driven by temperature. The team observed a large degree of variation in when people first used their AC, with most spanning May to June. There are two primary scenarios leading to first AC use: 1) heat drives the customer to switch their thermostat to cooling mode, and sets a low enough setpoint to turn on the air; 2) a customer switches their thermostat to cooling mode more in response to the calendar than weather early in the season, establishes preferred setpoints, and the AC finally turns on at some point during the summer when the indoor temperature finally exceeds the setpoint.
	Consideration: PAs can use this knowledge to affect customer behavior and encourage energy efficiency. Especially moving forward with an increasing portion of customers having Wi-Fi and smart thermostats installed, the PAs may work with vendors to identify customers who turn their AC on early in the season and encourage them to set conservative cooling setpoints so the AC kicks in later in the season. Customers that respond in the moment to the heat, actively turning on their AC, may respond to messaging suggesting they leave their home, go to the pool, go out for ice cream or a trip to the mall, rather than turning on their AC the first time they feel uncomfortable. Some of these customers may respond to the behavioral competition and encouragement through incentives like coupons for ice cream retail discounts for delaying first time AC use. Through these and similar interventions the PAs can encourage overall energy efficiency savings and potentially avoid spikes in demand on hot days where a large portion of customers would have turned their AC on for the first time all at once. Note that delaying setting the thermostat to cooling may have the unintentional effect of the home continuing to be heated longer on cold nights, if the thermostat lacks an “off” mode, or it is not used. The PAs should carefully consider this tradeoff in using this later in the season nudge approach to cooling.
Get People out of their Homes on the Weekend to Postpone First AC Use	**Key Finding:** People are more likely to turn on their AC for the first time on a weekend day rather than a weekday.
	Consideration: Through behavioral messaging and incentives the PAs may help people postpone turning on their AC for the first time in the season—especially for those who actively choose to turn on their AC for the first time each season in response to discomfort. By offering encouragement to get out of the house on the weekends, people may turn on their AC later in the season.

Takeaway	Key Findings and Considerations
Encourage People to Shut their AC Off Earlier in the Season	**Key Finding:** Last AC use in the season was concentrated during a couple weeks in September. A smaller portion of customers continued to use their AC through October and November, even into December.
	Consideration: Two strategies may help PAs encourage customers to shut off their AC earlier in the season. Those customers identified as setting their system to cooling mode early in the season and then letting the system turn the AC on when it finally gets hot enough inside might be encouraged to turn their AC system off earlier in the season simply by messaging that it is time to consider switching your system to heating mode. This simple message may be enough to encourage some of these users to switch modes, ensuring their AC does not automatically turn on again for the remainder of the cooling season. The small portion of customers that continue to run their AC through the fall, and even into December, may respond to social norming. Messaging letting them know the way they use their AC is not typical may encourage them to act differently. The PAs should note the possibility that this strategy may have the unintended consequence of causing customers' heating to come on earlier in the season if an off mode is not utilized.

Weather and AC Use throughout the Cooling Season

How does weather affect people's AC use throughout the cooling season?

The evaluation team used both exploratory data analysis and regression analysis based on AMI data from approximately 9,000 homes that have AC to explore AC use in response to weather throughout the cooling season. Due to the granularity of the household-level AMI data the team was able to create a detailed fixed-effects model to estimate the effects of weather and other key variables on AC use throughout the summer. The results provide insights into how temperature, humidity, and additional variables such as day of week affect AC use.

Key Findings and Considerations

Takeaway	Key Findings and Considerations
Strategies to Reduce AC Use Throughout the Summer Should Focus on Both Heat and Humidity	**Key Finding:** In contrast to the team's findings on first AC use, research results suggest that people respond strongly to both temperature and humidity, alone and in combination, when deciding how to use their AC throughout the season. With respect to temperature alone, on average, for every 3°F temperature increase above 68°F, customers use an additional 1 kWh of AC. Doubling the humidity alone increases daily AC use by over 1 kWh. Moreover, temperature and humidity combined increase people's AC use even more.

Takeaway	Key Findings and Considerations
	Consideration: Targeted messaging with information about high efficiency dehumidifiers and tips to decrease home humidity could help reduce AC use throughout the summer. Providing customers with messaging around strategies for being out of the home during the hottest, most humid times of the day could also reduce AC use. In the days leading up to predicted hot, humid weather, PAs can message strategically to customers who use their AC in different ways, to maximize energy efficiency. As an example, customers who keep their setpoint very high throughout the summer might be pushed out of their comfort zone by the added humidity and more likely to override settings to engage their AC. As more and more customers have devices reporting telemetry data, the team proposes the PAs consider working with thermostat vendors to identify these customers and use targeted messaging to help them save energy. PAs may also continue expanding AMI coverage as AMI data will also help identify users who use AC in different ways. In line with earlier suggestions, the PAs may also decide to incentivize thermostats that effectively account for and manage humidity. The current usage pattern many customers follow of waiting until a critical comfort threshold is exceeded before turning on the AC exacerbates the problem of spiking, where many units come online at the same time, putting strain on the system. By incentivizing thermostats that take humidity into account, the PAs can encourage people to use their thermostats for humidity control well before peak humidity is reached, mitigating this issue.
Focus Less on Fatigue and More on In-the-Moment Interventions	**Key Finding:** Although periods of prolonged high temperatures increases AC use, the majority of use is determined by day-of temperatures and humidity. In other words, people are less sensitive to heat and humidity fatigue and are reacting more to the temperature and humidity conditions on a given day when making AC use choices.
	Consideration: Providing customers with coping mechanisms for heat and humidity on a day-to-day basis is more important than creating strategies to help them deal with prolonged heat spells. Providing tips about how to stay comfortable in the heat, how to reduce humidity in their homes, how to time activities to be out of the house during the worst heat and humidity may be the most effective. Providing coupons for out of the home activities on hot days can help customers cope and use less energy. Messaging encouraging them to dress for the weather, eat cooler meals, and other coping strategies can all help relieve the urge to turn up the AC. The PAs may also consider the benefit of having a DR program in place to further manage loads and usage spikes on hot and humid days.
Interventions Related to Heat in the Spring and Fall are Unlikely to Reduce AC Use	**Key Finding:** Electricity use was unaffected by temperature increases during the shoulder months of May and October, signaling that AC use is also relatively temperature insensitive during this time. This observation might be related to behavioral considerations around the timing of first and last AC system turn-on, or insensitivity to the relatively mild temperatures.
	Consideration: The PAs can focus targeted messaging and behavioral modifications around hot weather and AC use in the June through September period. Even though temperatures reach highs in the 70s in May and October, temperatures are generally low enough that customers are not using their AC much during these months, and when they are, it is at a baseload level and not responsive to heat.

Weather and DR Program Participation

How do weather and program design parameters affect people's participation behavior (i.e., propensity to opt out) in DR programs?

The evaluation team's research explored how weather and program design factors influence DR program participant behavior, particularly their propensity to opt out of events. Influential weather factors include temperature and humidity, while relevant program design factors include program implementation details such as thermostat type, event duration, notification methods, consecutive events, day of the week on which events occur, and frequency of events. The evaluation team paired in-depth interviews with Thermostat DR program participants and implementers with data analysis, including regression analysis, to understand these relationships from both qualitative and quantitative perspectives.

Key Findings and Considerations

Takeaway	Key Findings and Considerations
Shift Events Earlier in the Week and Earlier in the Day	**Key Finding:** Participants are more likely to opt out at the end of the week compared to mid-week. The likelihood of opting out is 3.6% higher on Thursdays and 3.8% higher on Fridays compared to Wednesdays. The team also found that participants are more likely to opt out of an event later in the day. The likelihood of opting out is 2.4% higher for each hour later in the day the event starts.
	Consideration: The evaluation team recognizes that DR events must correspond with capacity needs, and that the PAs may have little control over when events are called. However, these opt-out rate drivers should be taken into account when predicting impacts of DR events. All other things equal, plan for more events to take place on Mondays, Tuesdays, and Wednesdays, and fewer events on Thursdays and Fridays. Similarly, plan for events to take place earlier in the day. Both strategies are likely to decrease opt outs.
Call for More Frequent but Shorter Events	**Key Finding:** Longer events resulted in higher opt-out rates. For each additional hour of event duration, the likelihood of opting out increased by 3.3%.
	Consideration: Avoid pushing the limits on participant's comfort for diminishing returns. Customers mentioned in the surveys that back-to-back events might make them tired of being uncomfortable in their house. If shorter events end before customers become uncomfortable or before they even arrive home, then it may be possible to eliminate the problem with more frequent events. An alternate consideration for the PAs is to time events such that the first event hour corresponds with the system peak, alleviating maximal load, rather than having the third or fourth hour of some events coincide with system peak.

Takeaway	Key Findings and Considerations
Arrange for Better Management of Participants' Comfort During Events	**Key Finding:** Regression analysis revealed that hotter and more humid days resulted in higher opt-out rates. Interviews with participants and implementers echoed this finding. Our results do not generalize to how opt-outs may relate to very extreme temperatures, as the weather was mild during our analysis period. Non-linearities in the relationship between temperature and drop-outs would likely be magnified if the study were updated to include data from summers with more extreme temperature variation.
	Consideration: Whenever possible, consider ways to manage participant's comfort, such as encouraging efficient pre-cooling. As DR events are often necessary on hot and humid days, strategies like adaptive algorithms can help diminish the effects of heat and humidity on customer comfort. A snooze button that offers temporary relief from the event without having to opt out of the entire event could be appealing to customers with an interest in saving money and energy but become too uncomfortable to participate consistently for the duration of the event.[2]
Offer More Motivating Incentive Packages for Event Participation	**Key Finding:** Opt-out rates for thermostats enrolled in the Connected Solutions (CS) program (Honeywell and ecobee thermostats) are significantly lower relative to the Rush Hour Rewards (RHR) program (Nest thermostats). One key difference is that CS participants face significant incentives to participate. The CS incentive structure requires participation in at least 75% of the DR events to receive the $25 annual reward. Another key difference to consider is equipment usability, and how easy the thermostat interface makes opting out. An interface making it easier to opt out of Nest events could have contributed to higher opt outs in the RHR program.
	Consideration: Create more compelling incentive structures for all DR programs to discourage opt outs and increase full participation. PAs may also consider focusing on technologies that make it more difficult to opt out, while understanding the tradeoff that this may lead to greater customer frustration or dissatisfaction.
Integrate Gamification Strategies and Adaptive Algorithms to Decrease Opt-Outs	**Key Finding:** CS program participants expressed awareness of the participation requirement, but they still prioritized comfort over what they considered a small financial incentive. To make the most of the financial incentive's impact, the incentive should be sized appropriately to motivate participation. Moreover, the team found that customers were not opting out in reaction to certain triggers but rather opting out at any point during the event when they became uncomfortable.

[2] National Grid is working through EnergyHub in 2018 to implement the "Firm Load Dispatch" adaptive algorithm for its DR events. The team suggests the PAs analyze the effectiveness of this approach and modify or continue this type of approach based on effectiveness. The approach is described here: http://www.energyhub.com/firm-load-dispatch-paper

Takeaway	Key Findings and Considerations
	Consideration: The integration of gamification strategies could be used to increase customer motivation to participate for the entire event or for longer than they would have otherwise. As an example, the PAs may consider incentives that scale up with the extremity of the weather. For instance, customers may receive increasing incentives as the heat and humidity increase. Alternately, the PAs may encourage a "last man standing" competitive mentality in customers by showing the number of customers still engaged in the DR event on a dashboard accessible by customers. This fosters a sense of competition, and customers may be discouraged from opting out because they see others are still enduring the heat and discomfort. Customers may be receptive to features that make them feel more involved in the program's success. Some strategies may include developing personalized goals to engage participants throughout the event, a status bar that informs the participant of the remaining time in the event or push notifications that offer tips to increase comfort during the event. Additionally, the PAs may consider or continue pursuing adaptive algorithms which help manage temperature adjustments to minimize opt outs while maintaining savings.[3]

Conclusions and Future Research

This study presents new findings on the effects of heat and humidity on AC use and people's participation behavior in DR programs. Moreover, it is one of only a handful to pair both quantitative, data-based analysis with qualitative survey methods to provide more context and a deeper understanding of results.

Key Study Conclusions and Contributions

- In line with previous research, the study quantifies the strong effect of heat on first AC use, use throughout the season, and people's likelihood of opting out of DR events.
- The team's results demonstrate that humidity is not a key factor in determining first AC use, but that it does play a strong role in determining AC use throughout a season. Moreover, humidity also plays an important role in causing people to opt out of DR events.
- The evaluation team finds that categorization of thermostat users into AC user types is less valuable than focusing on quantifiable observed behaviors people exhibit when using their thermostats, and focusing targeted messaging or other behavioral interventions on users with specific thermostat use traits rather than user type categories is likely to be more effective.
- By combining thermostat and AMI data, this study also estimates a model to translate electricity usage into AC runtime, which may prove a useful starting point in future studies related to AC use where thermostat data is not available.
- This study also shows that program configuration considerations such as duration and timing of events, thermostat type and vendor, and proximity of events may have as large an effect on DR event opt-out behavior as weather variables do. This suggests the PAs carefully consider program characteristics in addition to weather variables when planning DR events to optimize savings.

[3] See reference above to National Grid's Firm Load Dispatch adaptive algorithm trial in 2018.

Areas for Future Research

- This study illustrates the wide variability in ways people use their thermostats and AC systems in response to increasingly complex technological capabilities of these systems. People can and do use their thermostats in nearly as many unique ways as the team had thermostats to study. To successfully begin categorizing user types in a meaningful way, future studies would need sample sizes of thousands rather than dozens with complete season thermostat data and could use machine learning methods to begin categorizing users by type.
- This study reveals that people's first and last use of AC during the cooling season are not governed by the same processes and, therefore, cannot be modeled identically. People's last AC use is highly concentrated over a short period of time, whereas first AC use instances occur more gradually as the summer progresses. The evaluation team's limited last-use findings suggest a lengthier dataset might provide more valuable insights into this behavior, offering the ability to observe last AC use instances over a period of multiple years rather than concentrated in a single month of one year.
- Another area for future research is to explore the effects of weather on an even more granular sub-daily level. While the modeling complications introduced by hourly or sub-hourly models outweighed their benefit in the current study, future studies might use this one as a starting place to study weather's effect on AC use throughout the season using AMI data, focusing particularly on sub-daily relationships.
- Finally, the largest constraint faced by this study was the limited sample size of AMI dataset customers overlapping with thermostat data. Given a much larger sample size, future research might be able to draw even more granular findings on which to base recommendations.

1. INTRODUCTION

1.1 Background and Objectives Summary

The Massachusetts Program Administrators (PAs) engaged Navigant Consulting, Inc. (Navigant) and Illume Advising, LLC (the evaluation team) to study the effects of weather on air conditioner (AC) use behavior and demand response (DR) program participation. Understanding how weather affects customer AC usage will allow the PAs to improve the design, implementation, and customer messaging around these programs. Similarly, understanding the role of weather on a participant's likelihood to opt out of DR events will allow the PAs to balance opt-out rates in calling DR events.

Key study objectives include the following:

1. Develop models of central AC use that show how weather affects the timing of first and last use of AC during a given season and how weather affects AC use throughout the season.
2. Understand the role of weather on DR event participation and, specifically, how the propensity to opt out or override preset thermostat controls varies with weather and program parameters.

This study was undertaken in a staged approach and combined several datasets, including advanced metering infrastructure (AMI) data and thermostat telemetry data from National Grid's Smart Energy Solutions (SES) pilot in Worcester, as well as thermostat telemetry data from National Grid's Residential Wi-Fi Thermostat DR demonstration project.

- **Phase I** included a literature review, data management, and preliminary data exploration activities. The literature review yielded important findings from related studies and suggested areas for investigation, particularly with respect to different types of AC users. Through these activities the evaluation team came to a solid understanding of the datasets it had to work with in this study, their applicability to answering the research questions posed, initial insights into AC user types, and the types of regressions which might be useful in Phase II.
- **Phase II** built upon the research and analysis performed in Phase I to refine and finalize the methodology to meet the study's objectives. First, the evaluation team was able to define quantifiable AC use behaviors and metrics describing them, leveraging both the AMI and thermostat data, and providing actionable and valuable insights for the PAs. The team then conducted a regression analysis to estimate the effects of weather on AC use—both ways in which weather affects the first time and last time residential customers use their AC during the cooling season, and more general AC use behavior throughout the season. Finally, the team combined qualitative interviews with exploratory and regression analysis of thermostat telemetry data to derive meaningful relationships between DR program participant opt-out behavior and weather factors and program parameters.

2. PHASE I RESEARCH

Phase I consisted of an in-depth review of existing literature, initial data management, and preliminary exploratory data analysis. The results of this phase informed the analysis in Phase II.

2.1 Literature Review

The team conducted an in-depth literature review early in Phase I of this study, which provided insights into both the relationship between weather and AC use and weather's effects on DR program participation. These results helped frame and refine the team's approach to Phase II research.

2.1.1 Background

The literature review ensured the team made full use of all secondary material currently available and did not spend time answering questions that were already well-researched. It also helped shape Phase II research. The review explored existing research in the following subject areas:

Figure 1. Key Literature Review Topics

AC Use Behavior

- When in the year customers first begin to use their AC and the relationship with weather
- When in the year customers stop using their AC and the relationship with weather
- How weather affects people's AC use throughout the season
- What different AC user-types have been identified in the literature and how might their AC use differ

DR Participation

- How weather influences DR participant's propensity to override programming and opt-out of events
- What existing literature can tell us about the importance of heat versus humidity in affecting DR participation

2.1.2 Methodology

The team's in-depth literature review included over 20 published sources of information on how weather affects AC use and DR program participation. These sources included both print and online documents spanning academia, program evaluation reports, conference proceedings, journal articles, and publicly available manufacturer information.[4]

2.1.3 Findings

In-depth literature review findings inform the approach to researching weather-related AC use and DR program participation in Phase II. Findings on the relationship between weather and AC use help the team anticipate challenges in modeling and questions that have not been well addressed by the existing literature. For instance, the effects of weather variables other than temperature such as humidity are not as well understood. On the DR participation side, secondary research suggests the team's findings on the relationship between weather and opt-out behavior may help the PAs increase savings by optimizing the balance between savings and weather-related opt-out rates.

2.1.3.1 AC Use and Weather

Existing literature suggests that people's AC use is strongly related to temperature. For example, a 2017 Massachusetts baseline study by Navigant finds that most residents report beginning to use their AC when the indoor temperature reaches approximately 75°F and when the outdoor temperature reaches 80°F,[5] as shown in Figure 2. This suggests that even during summers with relatively moderate weather, the team may find significant responses in temperature related AC use. Stewart (2014) finds that thermostat setpoints may vary significantly between summer and winter, with higher cooling season and lower heating season setpoints.[6] This reinforces the idea that people's AC use behaviors change based on the temperature.

Figure 2. Temperature and AC Use, Massachusetts Survey Responses

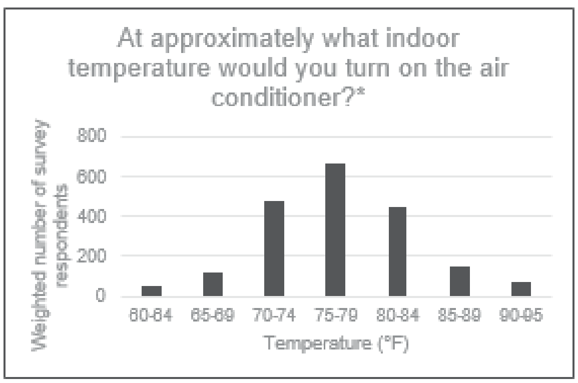

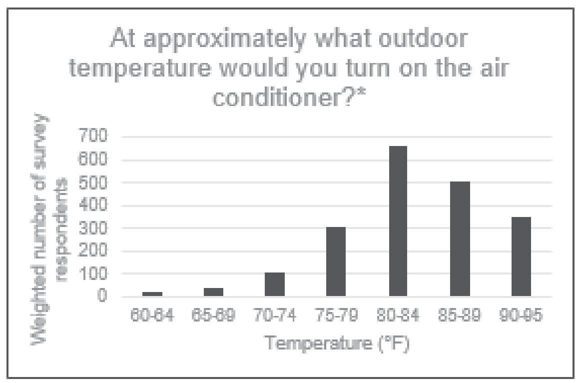

Source: Navigant. (2017) Massachusetts Baseline Study. Prepared by Navigant Consulting, Inc. for the Massachusetts Program Administrators and Energy Efficiency Advisory Council

[4] Navigant attempted to conduct interviews with several thermostat vendors during Phase II, all of whom declined to participate or comment on how their technologies' control strategies and optimization algorithms interact with weather.

[5] Navigant. (2017) *Massachusetts Baseline Study*. Prepared by Navigant Consulting, Inc. for the Massachusetts Program Administrators and Energy Efficiency Advisory Council.

[6] Stewart, J. (2014) *Energy Savings from Honeywell Connected Thermostats*. 2014 Behavior Energy and Climate Change (BECC) Conference, Conference Proceedings.

Humidity is also shown to influence AC use behavior. A study of California homes by Lawrence Berkeley National Laboratory (Berkeley Lab) finds that thermostat setpoints vary based on humidity, with lower setpoints in humid environments to help reduce humidity. The lab found that while the typical setpoint in humid environments is 74°F, in drier climates that setpoint is closer to 80°F.[7] Similarly, Navigant (2017) finds that Massachusetts residents report starting to use their AC when humidity is above 60%, though a large portion of survey respondents report waiting until humidity exceeds 80%.[8] These findings suggest considerable variation in observed humidity may be needed to identify its effect on AC use through regression modeling, and that people's decision to turn on their AC for the first time in a given cooling season may be driven more by temperature than humidity given a relatively lower temperature tolerance.

Figure 3. Humidity and AC Use, Massachusetts Survey Respondents

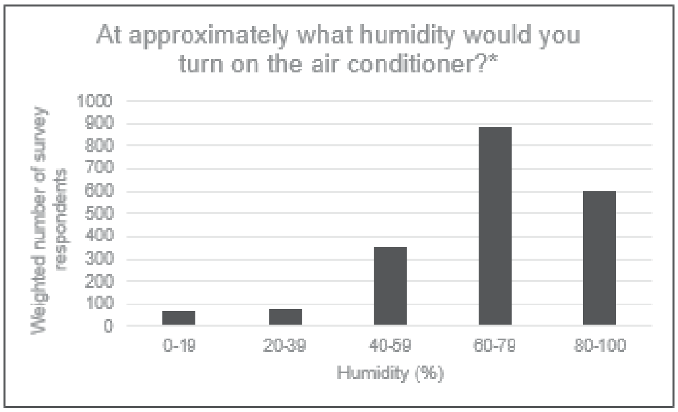

Source: Navigant. (2017) Massachusetts Baseline Study.

People use their thermostats to alter AC use in relation to weather. Reed (1991) finds that both AC runtimes and thermostat setpoints vary with temperature, and that while many households manipulate their thermostat during very hot days, others leave the setpoint constant.[9] Another study finds that many people adjust their thermostats to avoid entering an environment that is uncomfortably hot or cold when returning home. Some study participants increased their gas usage after installing a Wi-Fi thermostat as a result of this preemptive behavior.[10]

A growing body of research suggests the existence of different AC user types. Cadmus (2013) finds that switching from a programmable thermostat to a Wi-Fi-enabled one increases the likelihood of setting a programmed schedule from 59% to 69%.[11] Customers who set and adhere to a preprogrammed thermostat schedule are typically referred to as set it and forget it users in the literature. A 2016 study by Bonneville Power Authority found that participants who switched from a manual to a Wi-Fi-enabled

[7] Walker, I. S., & Meier, A. K. (2008). *Residential Thermostats: Comfort Controls in California Homes*. Lawrence Berkeley National Laboratory Berkeley, California March 2008 LBNL-938E.

[8] Navigant. (2017) *Massachusetts Baseline Study*. Prepared by Navigant Consulting, Inc. for the Massachusetts Program Administrators and Energy Efficiency Advisory Council.

[9] Reed, J. H. (1991). *Physical and Human Behavioral Determinants of Central Air-Conditioner Duty Cycles*. IEPEC, 208–217.

[10] Cadmus. (2013). *Liberty Utilities Wi-Fi Programmable Thermostat Pilot Program Evaluation.* Prepared for Liberty Utilities; Peffer, T., Pritoni, M., Meier, A., Aragon, C., & Perry, D. (2011). *How people use thermostats in homes: A review.* Building and Environment, 46(12), 2529–2541. https://doi.org/10.1016/j.buildenv.2011.06.002

[11] Cadmus. (2013). *Liberty Utilities Wi-Fi Programmable Thermostat Pilot Program Evaluation*. Prepared for Liberty Utilities.

thermostat were also more likely to fiddle with their setpoints to achieve their desired temperature.[12] These users are often referred to in the literature as fiddlers.

A different set of AC users, often referred to as on-off switch users, use their programmable thermostat as an on-off switch, forcing their AC to run manually rather than relying on automatic settings. Several studies confirm that between 25% and 50% of US households use their thermostat simply as an on-off switch, regardless of whether their thermostat is manual or programmable, while some use their thermostat as a valve to modify AC runtimes for comfort.[13] Research suggests that many households with a programmable thermostat avoid using the programming feature. As an example, Meier (2010) finds that 89% of participants in a utility programmable thermostat pilot study self-report that they rarely or never used their programmable thermostat to set or follow a program. The same authors cite a study by a thermostat manufacturer that finds out of over 35,000 programmable thermostats monitored, only 47% were typically in program mode, while 53% were in hold mode.[14]

Based on these and the previous findings, the evaluation team has documented the four most common AC user types described in the literature. These secondary literature-based AC user types are detailed in Figure 4 and inform the team's investigation thermostat use behaviors and AC user types in Phase II.

Figure 4. AC User Types

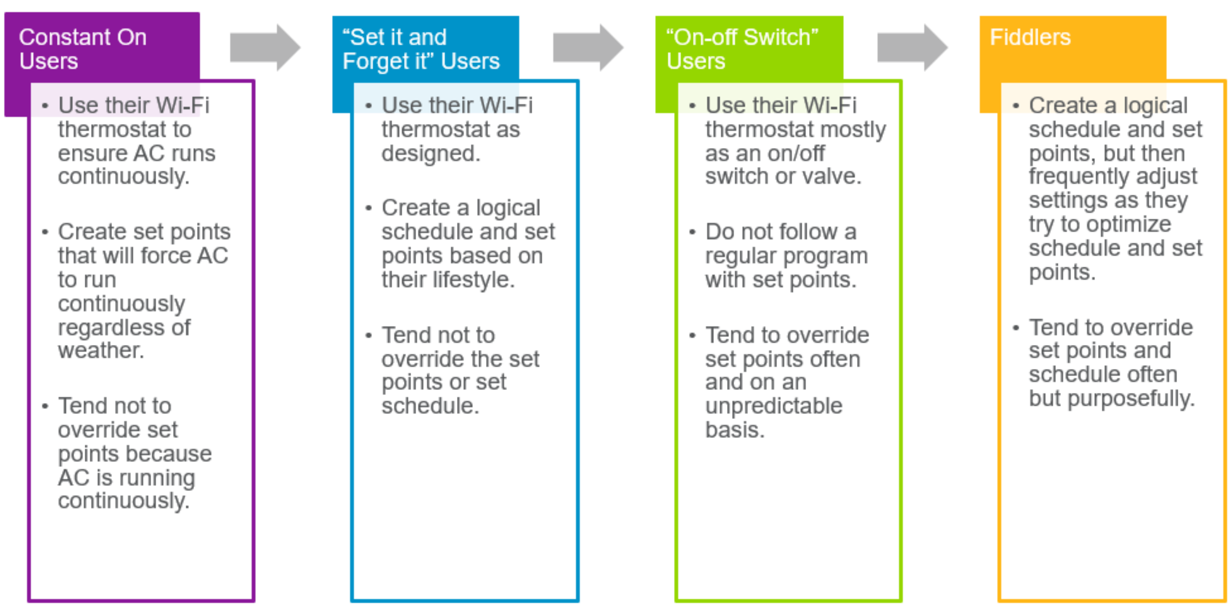

Source: Navigant

---

[12] BPA. (2016). *Nest Learning Thermostat Pilot Program Savings Assessment.*

[13] Kempton, Willett. (1986) *Two theories of home heat control*. Cognitive Science 10(1): 75-90; Meier, A., Aragon, C., Hurwitz, B., Mujumdar, D., Peffer, T., Perry, D., & Pritoni, M. (2010). *How People Actually Use Thermostats*. Controls and Information Technology, 2, 193–206; Rathouse, K., & Young, B. (2004). RPDH15: *Use of Domestic Heating Controls*. Watford: Building Research Esablishment (UK); Reed, J. H. (1991). *Physical and Human Behavioral Determinants of Central Air-Conditioner Duty Cycles*. IEPEC, 208–217; Walker, I. S., & Meier, A. K. (2008). *Residential Thermostats: Comfort Controls in California Homes*. Lawrence Berkeley National Laboratory, California March 2008 LBNL-938E.

[14] Meier, A., Aragon, C., Hurwitz, B., Mujumdar, D., Peffer, T., Perry, D., & Pritoni, M. (2010).

2.1.3.2 DR Program Participation and Weather

The literature review finds few studies draw causal relationships between weather variables and DR program opt-out rates. While existing literature points to a strong correlation between temperature and opt-out rates, these findings are often context-specific and not generalizable. Limited research exists on the relationship between humidity, the effect of varying incentives and penalties, or extreme weather on DR program opt-outs.

A summary of existing studies reporting opt-out behavior for different programs suggests a wide range for opt-out rates. While these opt-out rates appear correlated with weather, many studies do not explicitly examine the relationship between temperature or humidity and propensity to opt out of events. Other program setup parameters vary widely and may drive as much of the variation in opt-out rates as weather, given most studies provide correlations rather than causal estimates. Figure 5 presents an overview of opt-out rates, weather conditions, and other program parameters from recent DR program studies, in which observed opt-out rates vary between 2% and 50%.

Figure 5. DR Program Opt-Out Rates and Temperature Conditions

Authors and Studies	Average Temp during Event	Number of Events	Total n	Avg % Override
National Grid, 2017 (Preliminary Findings)	Average outdoor temperature has been greater during Rush Hour Rewards events leading to larger setbacks (3 °F) compared to Connected Solutions events (2 °F)	32 during summer of 2016 (not all events included all participants)	1,782 residential participants (1414 “Rush Hour Rewards”; 368 “Connected Solutions”)	15%
Navigant, 2015 & 2016	75 °F in 2015; 76 °F in 2016 (setbacks 3-4 °F in 2015; 2-3 °F in 2016)	40 (20 each year--2015 & 2016)	10,301	2%
Navigant, 2015	85 to 90 °F (setbacks between 0-3 °F based on customer setting choice--customers have the choice of "Max Comfort" setting where setback is 0 °F)	5 critical peak events; day-ahead High/Medium/Low pricing signals throughout summer	1,047	20-30%
KEMA, 2006	75 °F (setback 4 °F)	12/summer season	3,936	20%
Egan-Annechino, 2005	Above 90 °F (control accomplished by 50% off-cycle during event)	5-10/year for residential (2002 & 2003)	1,600 residential participations	27% for residential participants (via survey)
		Three simulated days for commercial (2004; no control days called by NYISO)	2,259 commercial participants	13.5% for commercial participants
Wang, Swisher, & Stewart, 2005	96 °F (setbacks 4 °F)	12/summer season	202	20%
Agnew, Goldberg, & Rubin, 2004	68-82 °F (setbacks between 3-5 °F)	12/summer season	100	Less than 20%-47%

Source: Navigant

A research study by San Diego Gas & Electric (SDG&E) finds that between 75°F and 85°F, a 1°F increase in temperature increases the override rate by 3.6 percentage points. This study finds the opt-out rate increases significantly with hotter weather, thereby decreasing savings overall, concluding that stronger opt-out penalties could decrease opt outs but would pose a challenge to recruiting.[15]

Navigant's 2017 National Grid DR study found that opt-out rates vary differentially with temperature based on manufacturer. While two manufacturers showed no clear relationship between temperature and opt outs, thermostat participants with one of the thermostat brands did show higher opt-out rates with increased temperature.[16] These findings inform the way the team structured its regression analysis in Phase II and suggest that in a program with mixed vendor types, accounting for vendor type and program parameters along with weather is critical to accurately estimating weather impacts on opt-out behavior.

Figure 6. Temperature and DR Program Opt-Outs by Manufacturer[17]

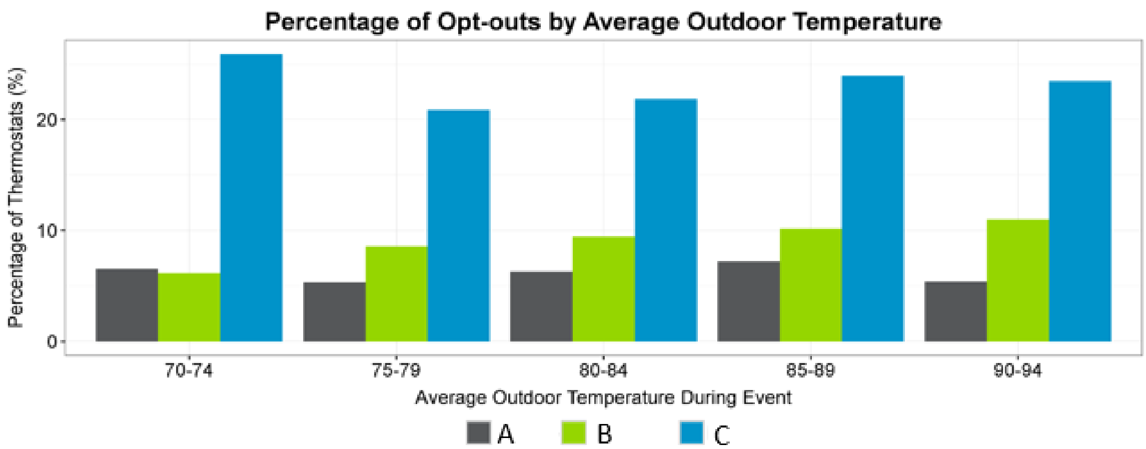

Source: Navigant. (2017). National Grid 2016 Residential Wi-Fi Thermostat DR Evaluation Final Report.

Both the SDG&E (2004) and Navigant (2017) DR studies find that savings varies with weather-dependent opt-out rates; optimal savings trades off between higher savings at higher temperatures (due to more avoided AC use) versus higher opt outs at higher temperatures (leading to lower participation and lower savings). Navigant (2017) finds that how savings vary with temperature differs by manufacturer, in part due to differences in the opt-out rate by manufacturer, as demonstrated in Figure 7. Because program parameters vary substantially by manufacturer, these differences may be more in response to variations in program parameters (e.g., event duration) than inherent manufacturer-specific effects. As a result, the team's Phase II regression analysis is careful to control for manufacture type and the timing of changes in program design parameters when estimating weather impacts on opt-out rates.

[15] Agnew, K., Goldberg, M., & Rubin, R. (2004). *You're Getting Warmer: Impacts of New Approaches to Residential Demand Reduction*. ACEEE Summer Study on Energy Efficiency in Buildings. Retrieved from http://www.eceee.org/library/conference_proceedings/ACEEE_buildings/2004/Panel_2/p2_1

[16] Navigant. (2017). *National Grid 2016 Residential Wi-Fi Thermostat DR Evaluation Final Report*

[17] This chart has been anonymized. This the original version in the National Grid 2016 report includes manufacturer names, we felt it was misleading to include them here, as in this context the increased propensity to opt-o

Figure 7. Savings Variation with Temperature, Differences by Manufacturer

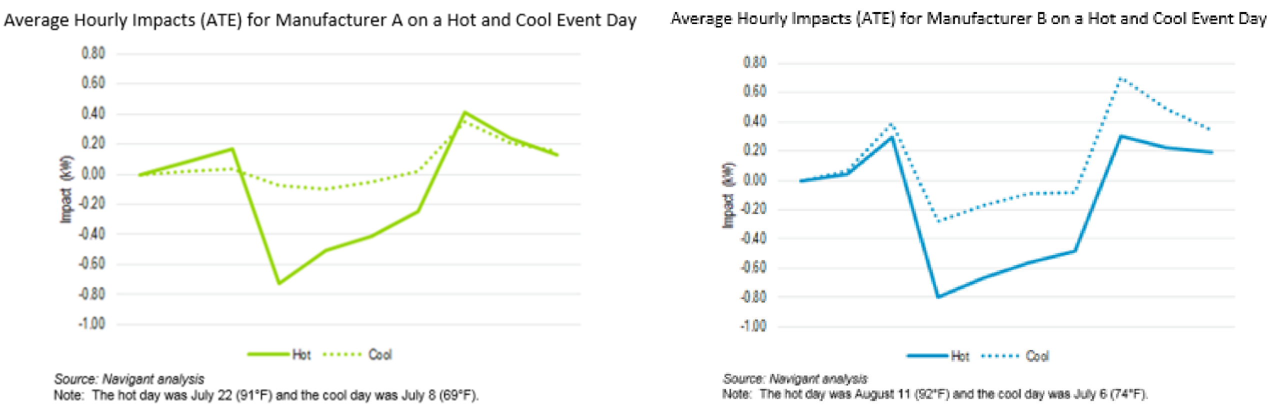

Source: Navigant. (2017). National Grid 2016 Residential Wi-Fi Thermostat DR Evaluation Final Report.

2.2 Data Management and Exploration

The evaluation team collected, compiled, and cleaned data from several different sources for this study. This included AMI and thermostat telemetry data from National Grid's SES pilot and thermostat telemetry data from National Grid's Residential Wi-Fi Thermostat DR demonstration project (Thermostat DR project). Data from the SES pilot was explored in addressing the study's first objective on AC use behavior. The team's intent in Phase I was to determine whether the AMI and thermostat telemetry datasets should be used singly or jointly to study the effects of weather on AC use. Data from both the SES pilot and the Thermostat DR project were explored in addressing the study's second objective on DR program opt-out rates. The team's intent was to determine which dataset(s) were most appropriate to achieve study research goals.

2.2.1 Background

For use in the AC portion of this study, the evaluation team obtained AMI data for approximately 10,000 Worcester, Massachusetts customers in 2014—post-installation of AMI metering but before the onset of the SES pilot program treatment in 2015. Similarly, the team possessed thermostat telemetry data from a subset of over 100 of those AMI customers, where a thermostat was installed prior to the onset of the SES pilot program treatment. For this set of customers, the team has both AMI and thermostat telemetry data throughout 2014 before SES treatment began. Both datasets could be used singly or jointly to study the effects of weather on AC use.

To study weather effects on DR participation, the evaluation team began with 2016 thermostat telemetry data from National Grid's Thermostat DR project, and data from its SES pilot spanning 2014 through 2016.[18] Since both could potentially be used to study weather effects on DR participant opt-out behavior, the team's intent in Phase I was to determine which dataset was most appropriate to achieve its research goals.

[18] For Phase II the team augmented this study with 2017 DR program data as well, but that data was not part of these Phase I data exploration activities.

The evaluation team undertook Phase I data preparation and exploration tasks in answering the research questions summarized in Figure 8.

Figure 8. Phase I Data Exploration Summary

Task Objective

- Clean and prepare data for exploratory analysis and output descriptive statistics
- Determine suitability of available (1) AMI and thermostat data to study AC use behavior, and (2) thermostat data to study DR participation.

Research Questions

- How suitable are the 2014 thermostat and AMI data for studying AC use behavior?
- What are the advantages and disadvantages of using thermostat versus AMI data in the AC use study?
- How suitable is the 2014-2016 thermostat data for studying DR program opt-out behavior?

Source: Navigant

2.2.2 AC Use Exploration and Findings

The evaluation team created descriptive statistics and graphics to better understand the completeness of each dataset. Review of the AMI data revealed complete data for approximately 14,000 premises from mid-January 2014 through January 2015 (see Figure 9).

Figure 9. AMI Data Completeness by Month

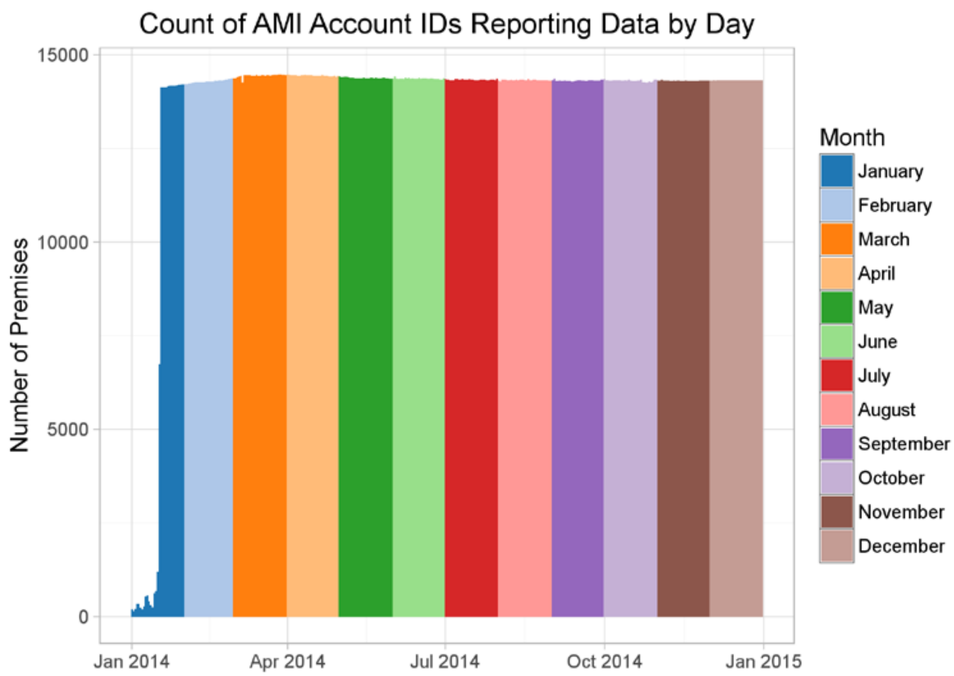

Source: Navigant

Similar exploration of the thermostat data (shown in Figure 10) revealed that, beginning in May 2014, the team had over 100 devices with complete records through the end of October 2014. The analysis period focuses on the hotter months—May through October—where thermostat use is likely to provide the best information on first, last, and overall AC use. Because thermostats were just being installed during this period, there were a larger number of devices reporting complete data later in the season relative to earlier months. These results suggested that the team would have a relatively small number of thermostats to use to explore the relationship of weather and first AC use, but a much larger sample to study weather's effect on AC use throughout the season.

Figure 10. Thermostat Data Completeness

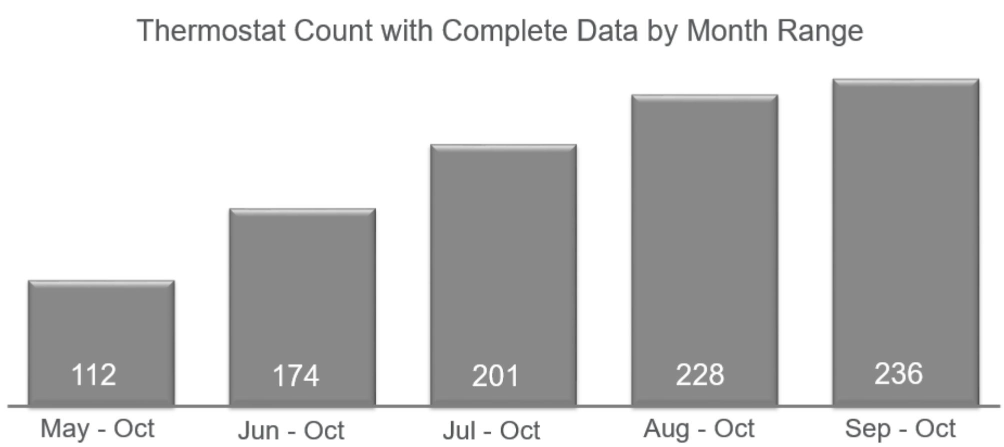

Source: Navigant

After narrowing the data to include sufficiently complete records, the evaluation team cleaned both datasets, removing extreme outliers, faulty data, duplicate records, and other anomalies. Thermostat telemetry data preparation included converting the data from state-change records to time-series entries formatted at the 15-minute interval level. After converting the data, the team summarized the time spent in cooling state by thermostat during each 15-minute interval. Investigating the overlap between thermostat and AMI data revealed 250 premises in total with both AMI and thermostat data. These premises were associated with 381 total thermostats. Of the 151 premises with only one thermostat (thus mappable to the AMI data), 54 had both sufficiently complete thermostat and AMI data available for the full summer season (May 1-October 1, 2014).

After summarizing both datasets, the team matched the AMI data to weather data for Worcester, Massachusetts[19] to create a basic model extracting weather-dependent energy use from total energy use. While Phase II would further refine this model, this demonstrated the team's ability to locate the approximate balance point upon which to base heating and cooling degree days (HDDs and CDDs) and to separate AC use from other energy use in the AMI data.

The evaluation team also inspected the distribution of thermostat records throughout the analysis period, assessing the correlation between the system mode (e.g., auto) and the system state (e.g., cool) and summarizing the amount of time spent in cooling state for all thermostats. Figure 11 provides a summary view of the distribution of cooling state records across all thermostat records. In general, the cooling state records are intuitive, increasing as the heat increases throughout the summer and then decreasing in the early fall. July and September are identified in this graph as the months where thermostats in the sample reported the most cooling records, though notable peaks and dips between July and September are likely correlated with cool or hot weather events.[20] Figure 12 conveys similar information from a different perspective. Instead of focusing on cooling records, this figure shows the percentage of time spent in the cooling state by week for each thermostat throughout the cooling season.

[19] The SES pilot was located only in Worcester, Massachusetts, hence the team's thermostat telemetry data is all from that region.

[20] One explanation for fewer cooling records in August could be unseasonably cool weather. Weather records for Massachusetts show that while the average temperature decreased from 80°F to 76°F then 70°F from July to August to September, during the same period the maximum temperature actually dipped during August; The maximum temperature for those three consecutive months, respectively, were 88°F, 84°F and 86°F, indicating a dip in August. Weather information obtained through US Climate Data website at https://www.usclimatedata.com/climate/worcester/massachusetts/united-states/usma0502/2014/7. Moreover, more people tend to take vacations in August, so may be away from their homes producing less cooling records during this month.

Figure 11. Thermostat Records from 2014 where System State is Cool

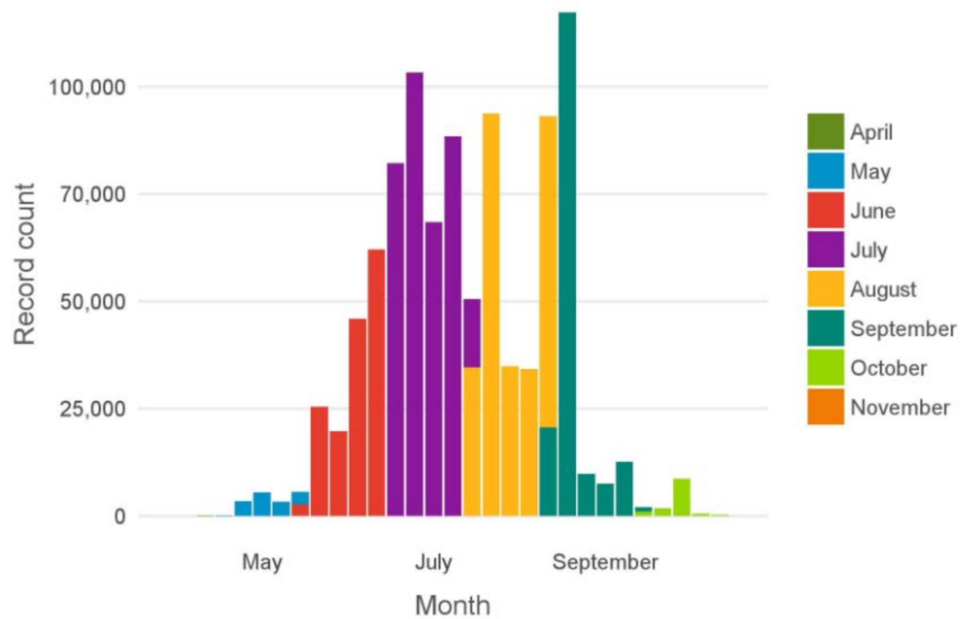

Source: Navigant

Figure 12. Weekly AC Runtime by Thermostat - All Thermostats[21]

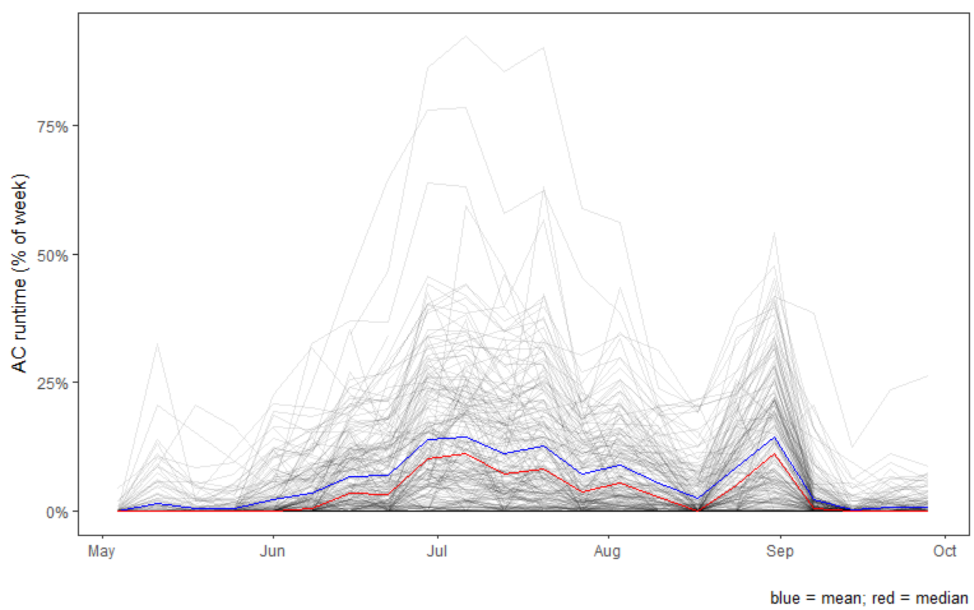

Source: Navigant

[21] In this figure "all thermostats" refers to all thermostats with any data between May 1 and October 31, 2014. No filters were applied to remove premises with multiple thermostats or thermostats with incomplete data during the cooling season. Figure 43. reproduces this figure using only thermostats used for the first AC use analysis.

Figure 13 shows the percentage of time spent in cooling mode by month across all thermostats. [22] It reveals that 4% of users never switch into cooling or auto[23] mode, so their AC is essentially disabled, while an additional 7% are in cooling mode 5 or fewer days (120 hours) throughout the summer. As expected, fewer users spend 0% of the time in cooling mode between June to September relative to the other months of the year.

Based on the setpoints customers have entered and the ambient temperature, even if these systems are in cooling mode, the AC system will only turn on part of the time—when the temperature exceeds the setpoint threshold. The right side of Figure 13 shows the percentage of time with AC running (time in cooling state). While the percentage of time AC is running increases during the hottest months and decreases again in the early fall, overall it is quite low.

Few devices are cooling more than 10% of the time before June and after August. Even during the hottest months, almost no customers fall into the “always on” category, with their AC system running in cooling state almost all the time. While the team observes systems cooling up to 40% of the time during mid-summer, these are extreme outliers. A large proportion of systems never enter the cooling state at all throughout the season, and as shown in Figure 14, the typical device cools for just over 10% of the time during the hottest months of the year, but less than 10% most of the time.

Figure 13. Percentage of Time in Cooling Mode or State by Month, All Thermostats

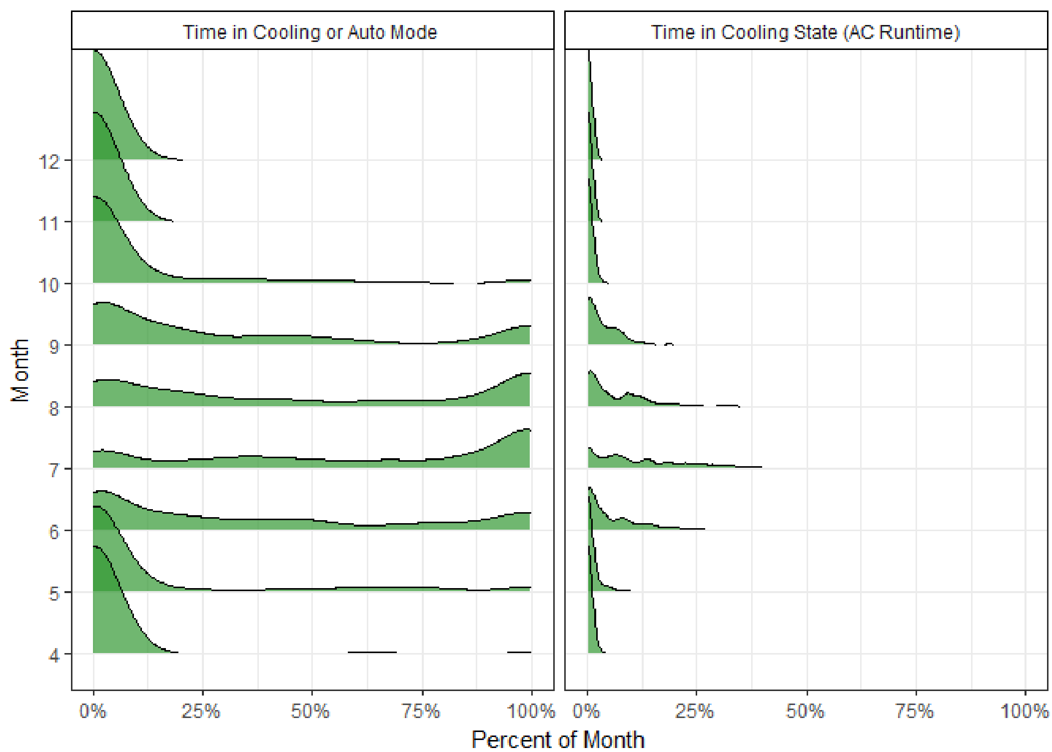

[22] This figure refers to “all thermostats” as the set of thermostats with any data during the period of May 1, 2014 to October 31, 2014.

[23] Auto mode is a mode in which the thermostat will engage either heating or cooling as necessary. This mode is most commonly used, or most useful, in regions where the temperature varies dramatically such that many days require switches between heating and cooling. However, in Arizona, auto mode and cooling mode are more or less identical during the summer, as cooling will usually be the only state engaged when auto mode is chosen.

Source: Navigant

The evaluation team merged the thermostat and AMI datasets, which involved joining cooling coefficient predictions from the AMI data to thermostat runtime calculations. The team then plotted the preliminary measure of AMI-based AC energy use versus the thermostat-based measure of cooling time to analyze the correlation between the two. Plotting the measures together, as shown in Figure 14, suggested that using only a basic modeling approach, the AMI-based and thermostat-based measures of AC use were weakly correlated. As a result, the team suspects more advanced modeling of AMI-derived AC use is needed in Phase II.

Figure 14. Preliminary Correlation of AMI and Thermostat-Derived AC Use Measures

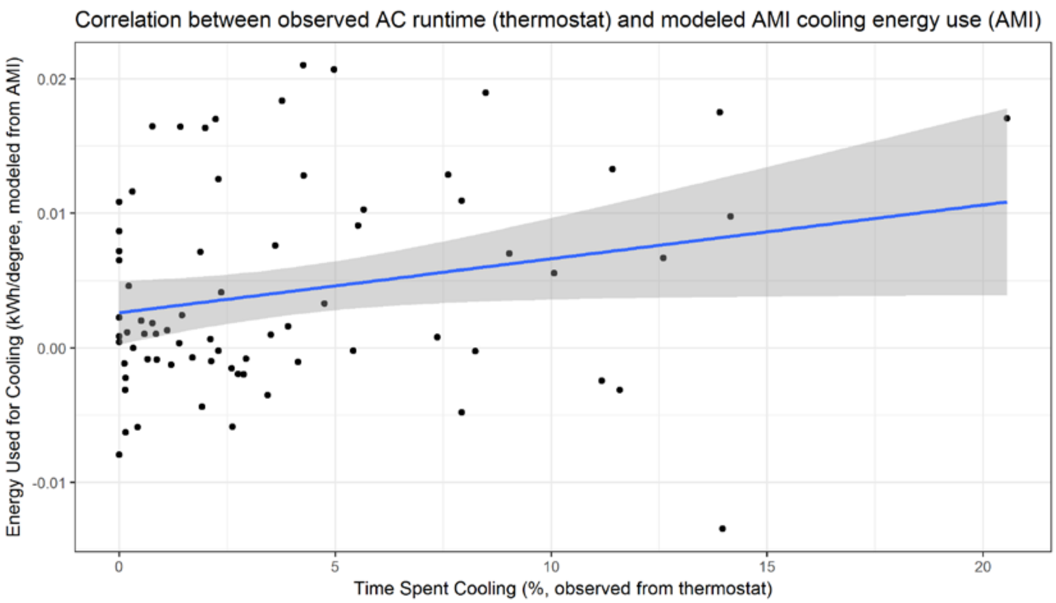

Source: Navigant

The evaluation team plotted average afternoon energy use (3 p.m. to 7 p.m.) for a cool month (April) versus a hot month (July) to get a sense of how different energy use profiles might be under substantially different weather conditions. The team found that overall energy use varied only slightly between these months, as evidenced by the large amount of overlap between the blue and red regions in Figure 15. This suggests that using the large sample sizes afforded by the AMI data may be essential to identify the small overall signal of weather effects on AC usage with statistical precision.

Figure 15. Small Difference in Total Energy Use between Hot and Cool Months

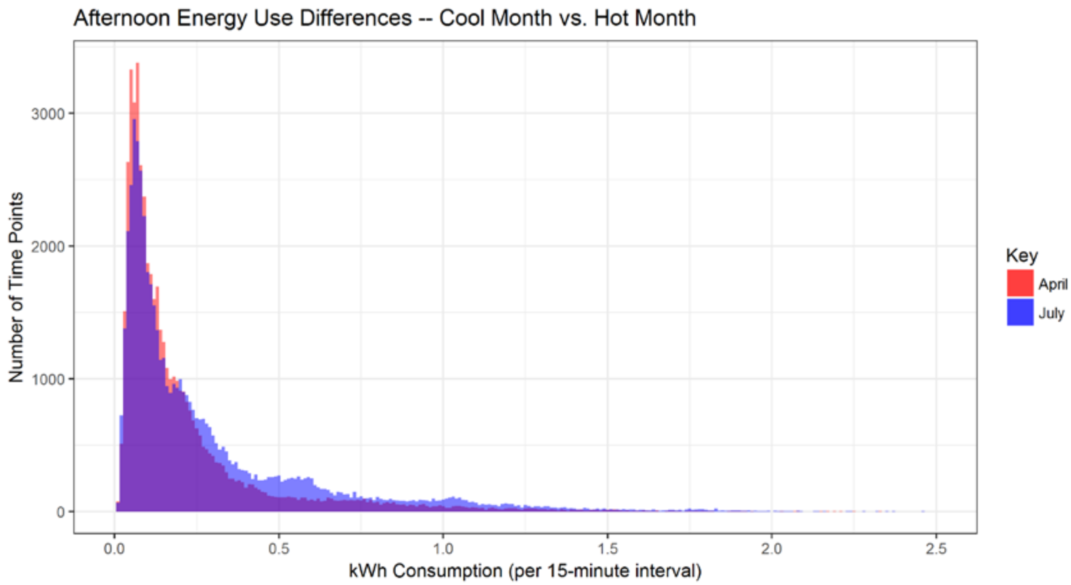

Source: Navigant

The evaluation team began its initial exploration of weather's role in a household's decision to turn on their AC for the first/last time in a season by plotting the occurrence of a customer's first/last cooling records from the thermostat data by month throughout the summer. As shown in Figure 16, first AC use typically occurred in May or June but with considerable variation in timing. Last use typically occurred in September or October, but peaked during certain weeks. This may suggest a more uniform, calendar-based decision point for turning off AC systems, or a discrete decision point based on the abrupt onset of cool weather. It could also suggest that while the AC system remained in cooling mode, the onset of cool weather made it such that the AC never turned on again based on customer setpoints.

Figure 16. Distribution of First and Last Cooling Occurrences

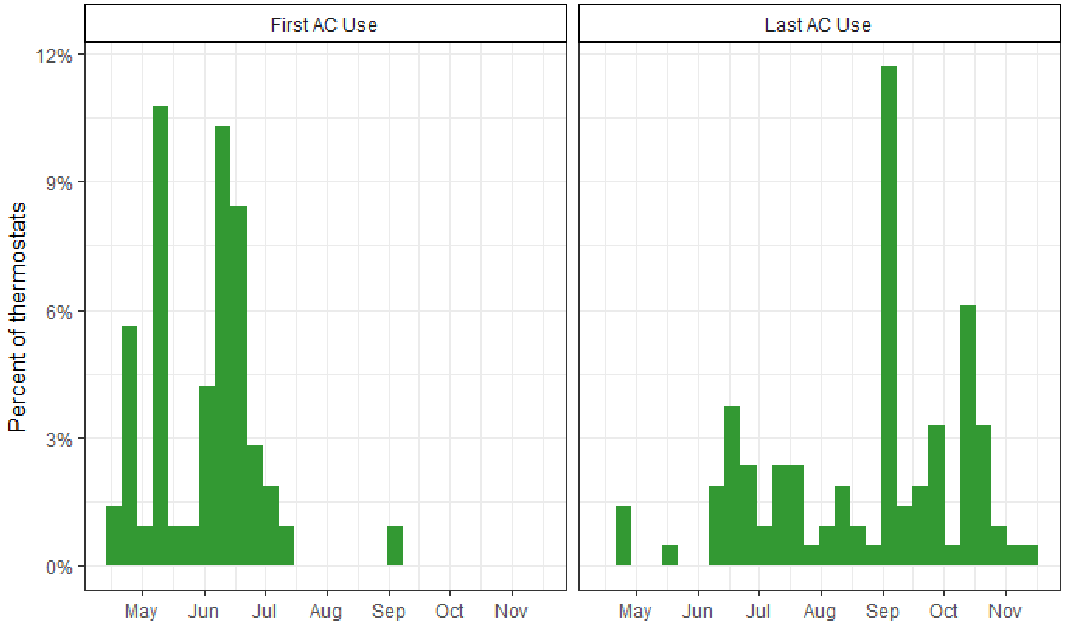

Source: Navigant

The team explored first and last AC use further by creating a scatterplot (see Figure 17) that shows the distribution of first AC use throughout the season relative to last AC use. This preliminary analysis reveals large variation in the timing of first and last use. While user types could be assigned with somewhat arbitrary boundaries, the behavior is really a continuum, without clear cutoffs or distinct groupings.

Figure 17. Initial Exploration of AC Use Groupings

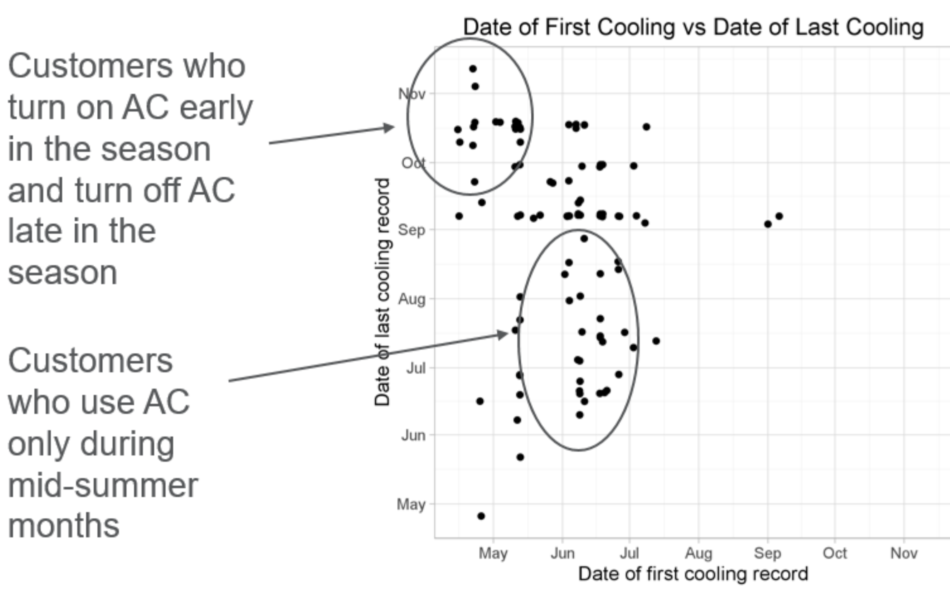

Source: Navigant

The evaluation team's preliminary findings from Phase I data exploration and analysis suggested that the team use the thermostat data to explore the effects of weather on the timing of first/last AC use during the season and the AMI data to investigate weather effects on AC use throughout the season. However, these results also suggested that the team may find additional benefits in marrying the two datasets to jointly explore these two topics in greater depth in Phase II.

2.2.3 DR Program Participation Exploration and Findings

While the evaluation team initially considered both the 2016 SES data and the 2016 National Grid Thermostat DR project data to study weather's effect on DR program opt-out behavior, the SES data did not contain a sufficient proportion of opt outs to be useful in this analysis. As a result, the team proceeded to work solely with the 2016 Thermostat DR project data.

Using the full set of telemetry data from the 2016 Thermostat DR data as its starting place, the team cleaned and prepared the data, removing commercial accounts, accounts with devices unconnected to AC, devices located outside National Grid's territory, and devices linked to multiple homes. The team also removed devices missing participation for more than three DR events during the season and devices with significant overall data quality/reporting issues and removed otherwise anomalous observations. The resulting dataset retained 1,888 thermostats with data spanning July through September 2016.

Despite a robust and large dataset, weather data revealed that the 2016 DR season had relatively few hot days. This limited the team's ability to draw statistically significant, meaningful findings relating weather to participant opt-out behavior. As a result, the team chose to postpone the DR research until after the 2017 DR season was complete, when it would have access to Thermostat DR project data for both the 2016 and 2017 event seasons.

The team began analyzing the 2016 data during Phase I to draw some initial conclusions around the prevalence of opt outs and their distribution by event and throughout the season. Tabulating the number of participating devices by event throughout the DR season, the team saw significant variation exists across the season between events in terms of event timing, the timing of subsequent events, and the number of participating devices (see Figure 18). This significant variation across time would enable the team to use regression analysis to model participation behavior.

As shown in Figure 19, opt outs are relatively frequent in the data, making the data suitable to study participation behavior, and the team expected that an additional year of data would further improve this.

The evaluation team expected that additional weather variation provided by adding 2017 to its study period, combined with the larger sample size and increased variation in opt-out behavior introduced by the additional data, would enable the team to draw statistically significant conclusions about weather variables' effect on DR event opt-out behavior. The high quality, large sample size final dataset the team would have to work with, spanning the 2016 and 2017 DR seasons, suggested it might also consider degree of participation. In other words, the team might be able to provide insights related to some customers opting out toward the end of events, while other customers opt out at the start.

Figure 18. 2016 Season Thermostat DR Project Participation

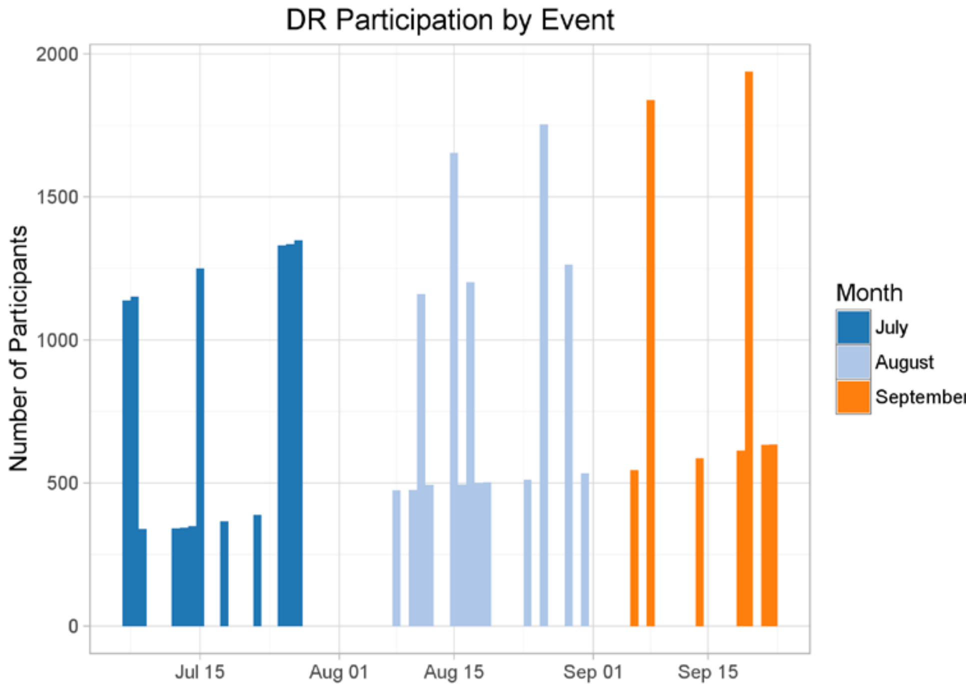

Source: Navigant

Figure 19. 2016 Season Thermostat DR Project Opt-Outs

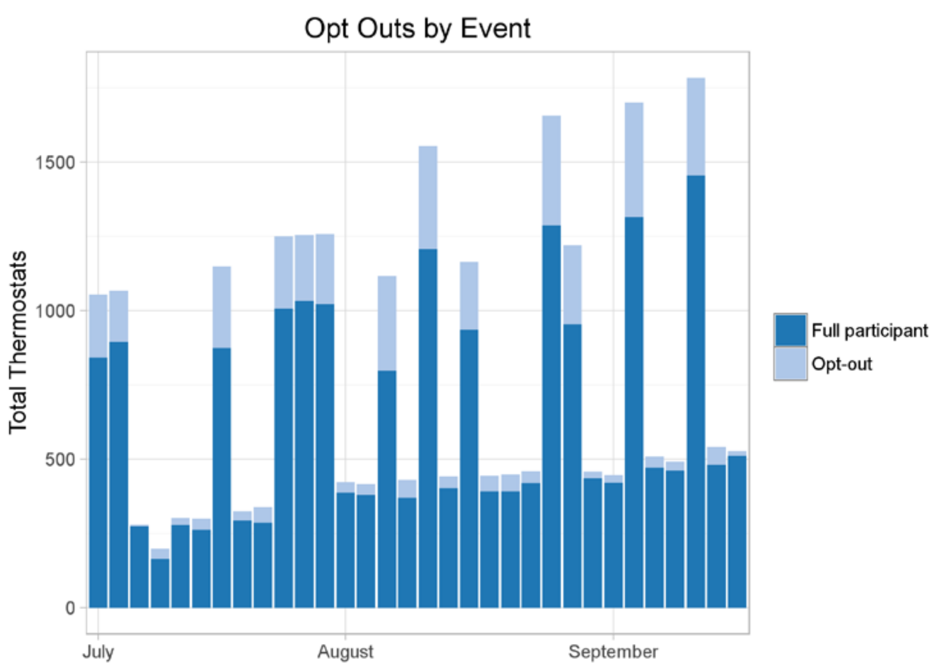

Source: Navigant

3. PHASE II RESEARCH

Phase II of this study uses the preliminary findings from the evaluation team's exploratory data analysis and literature review findings from Phase I to answer its key research questions:

- How do people use their thermostats, to what extent can customers be identified by different AC user types, and what can the team learn about AC use by examining specific thermostat use behaviors?
- How does the weather affect the timing of people's first and last use of their AC systems during the cooling season?
- How does the weather affect people's AC use throughout the cooling season?
- How do the weather and program design parameters affect people's participation behavior (i.e., propensity to opt out) in DR programs?

Accordingly, this section is organized as follows:

- Section 3.1: Thermostat Use Behaviors and AC User Types
- Section 3.2: Weather and First AC Use
- Section 3.3: Weather and AC Use throughout Cooling Season
- Section 3.4: Weather and DR Program Participation

3.1 Thermostat Use Behaviors and AC User Types

Smart thermostats afford users a high degree of control over the operation of their AC systems. As a result, the usage profile for each thermostat is effectively unique—the combination of cooling setpoint values, scheduled setpoint changes, weekend versus weekday operation, away and vacation holds, and manual overrides can and do vary significantly between homes. While the secondary literature suggests that users can be categorized as AC user types, the team's key finding in this portion of the research is that placing users into user type categories is less useful than simply looking at specific user behaviors, understanding specific ways customers are using their thermostats, and using that information to inform messaging or other interventions. Due to the increasingly complex functionality of thermostats, users display an awe-inspiring array of behaviors, suggesting that focusing in on the specific actions customers take with their thermostats and planning messaging and engagement around these is more actionable than classifying and targeting user types.

The distinction between cooling mode versus cooling state is important to keep in mind while reading the following sections. Cooling mode is when cooling is enabled (so the AC can turn on when the indoor temperature exceeds the cooling setpoint); cooling state is when the AC is actually running.

3.1.1 Methodology

To understand what behavioral characteristics could be gained from the thermostat data, it is first important to understand what data thermostats collect and transmit and which of these are shared with evaluators. While the data shared by thermostat vendors varies considerably, the following are relevant data points that were made available from each device to the evaluation team for this study:

- Device ID
- Date-time of each record
- Indoor temperature
- Cooling setpoint
- System mode (off, heat, cool, or auto)
- System state (off, heat, or cool)

Using the data points made available by the thermostat vendor, the evaluation team derived the following variables, which were used to better understand the specific actions the team observed customers taking in relation to their thermostats.

- Mean, minimum, and maximum cooling setpoint
- Percentage of time in cooling mode and cooling state
- Mean and median number of weekly cooling setpoint changes
- Mean and median number of weekly system mode changes
- Percent of time with Hold setting engaged
- Percent of time spent in irregular use-pattern

3.1.2 Findings

A key finding from the data is that understanding concrete actions people take with respect to their thermostats may be more informative than classifying customers as user types. Due to the functionality of modern thermostats, users have so many options that the range of observed behaviors is immense. Moreover, fewer and fewer customers display extreme behavior such as “always on” behavior (running the AC most or all of the time). In fact, most customers display a mix of different user type actions throughout the summer—rather than falling neatly in a category, they exist at different points on a spectrum of user types at different times.

3.1.2.1 Variation in Thermostat Use Behaviors

Figure 20 demonstrates the large variety of behaviors employed by customers in their use of thermostats with respect to cooling mode. Each faint black line traces the behavior of an individual thermostat throughout the season. The range of behaviors is almost limitless, and rather than clustering around distinct behavior types, users display a spectrum of behaviors.

Figure 20. Range of Observed Cooling Mode Behaviors

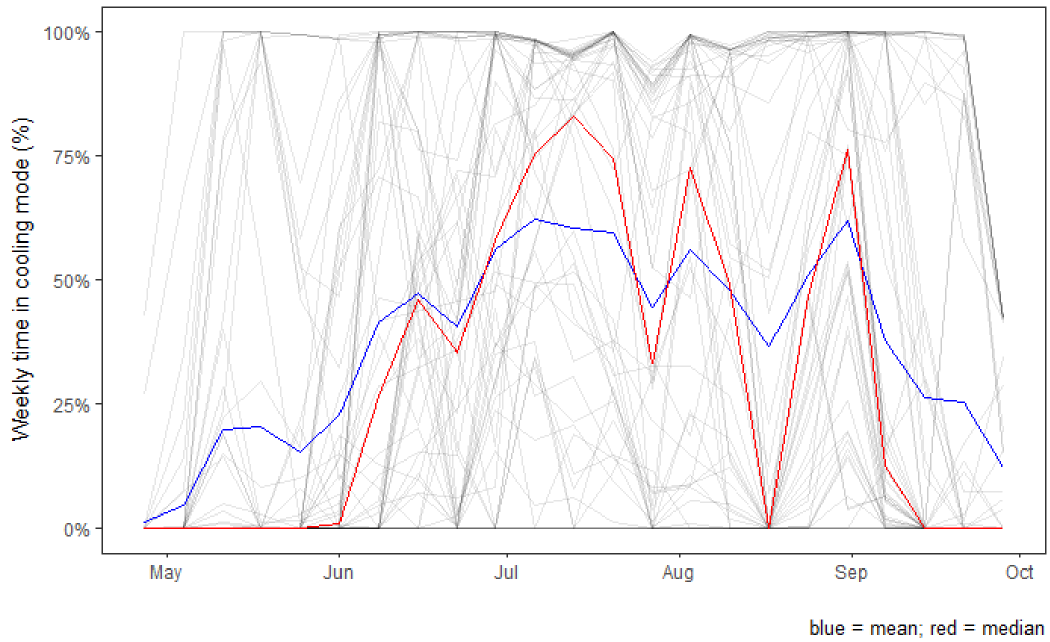

Source: Navigant

Figure 21 through Figure 26 illustrate the thermostat and AC usage patterns of six randomly selected thermostat users in the team's sample during July. The figures highlight the huge amount of variation in user types and behavior patterns. The variation in behavior is nearly infinite—almost no two devices behave the same throughout the month, and the behavior of most devices changes considerably week to week. The team observed nearly as many behavior patterns as there are thermostats in its sample.

Figure 21. Thermostat 3 July Behavior

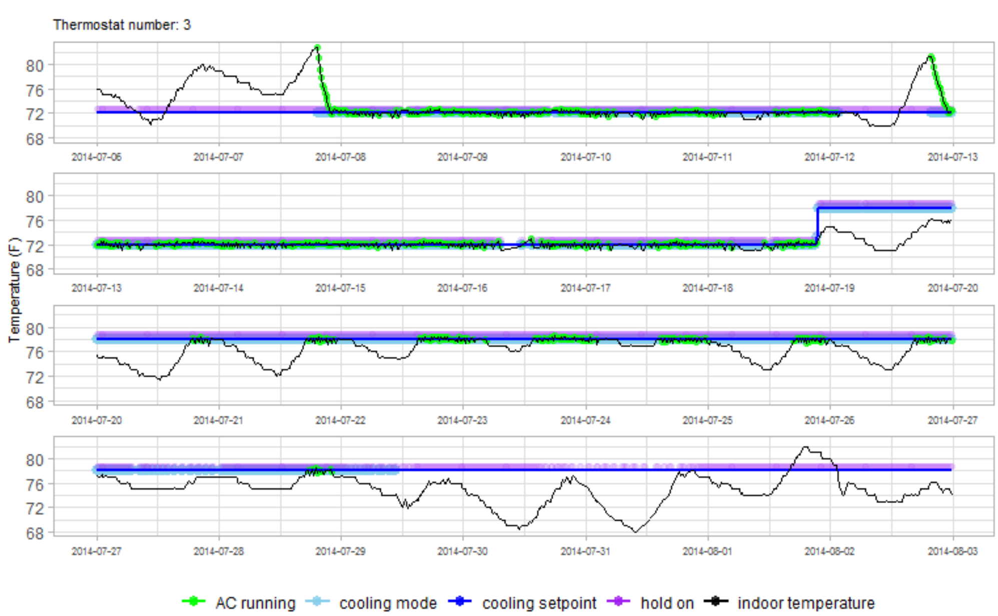

Source: Navigant

Figure 22. Thermostat 4 July Behavior

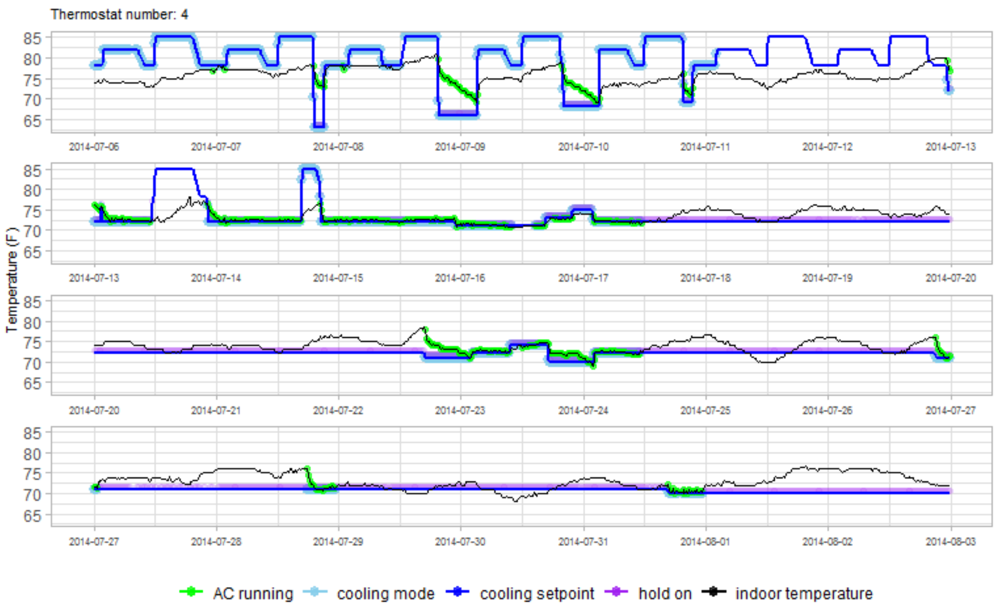

Source: Navigant

Figure 23. Thermostat 5 July Behavior

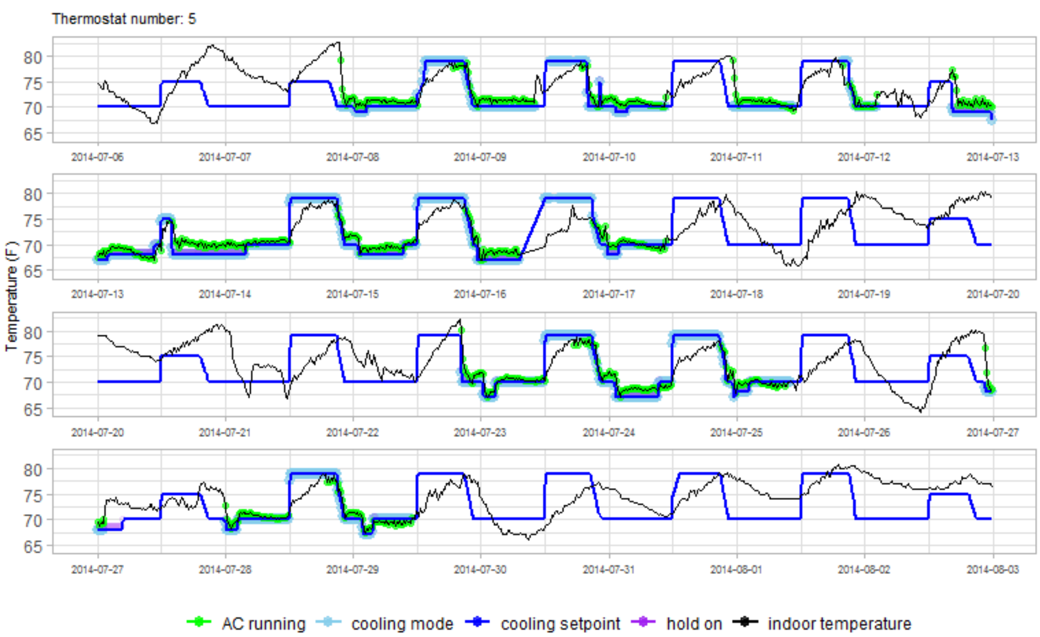

Source: Navigant

Figure 24. Thermostat 6 July Behavior

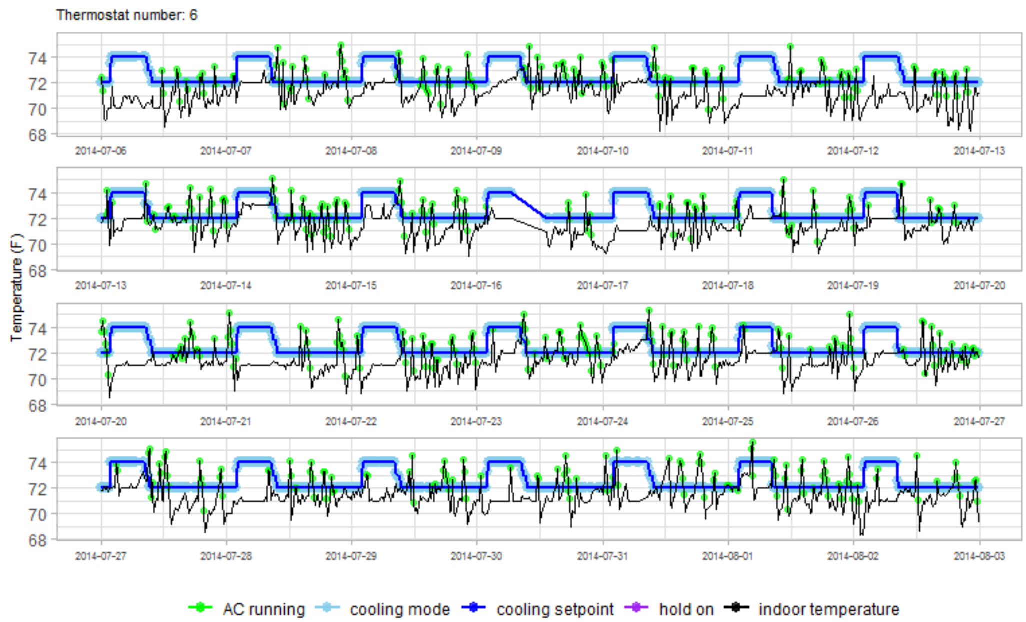

Source: Navigant

Figure 25. Thermostat 7 July Behavior

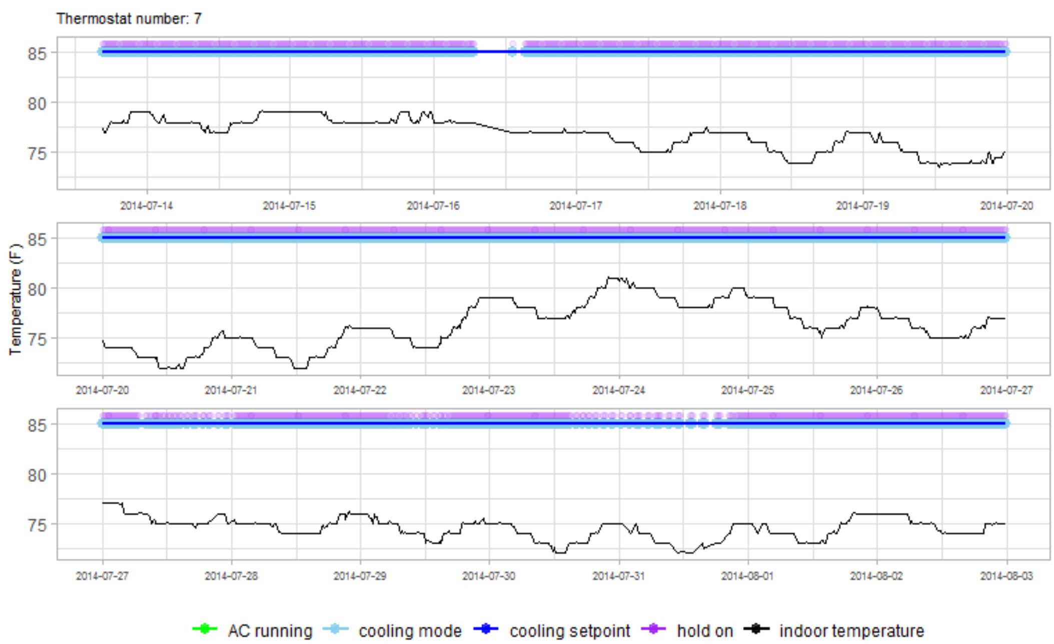

Source: Navigant

Figure 26. Thermostat 8 July Behavior

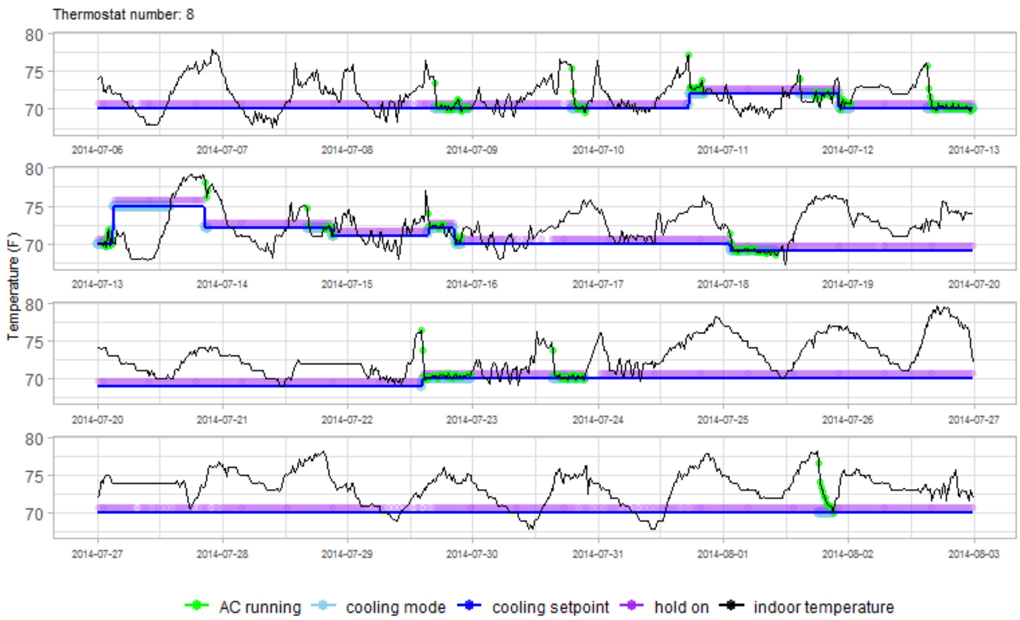

Source: Navigant

Due to the large amount of variation in how people use their thermostats, the team concluded that 54 thermostats with complete cooling season data is not a large enough sample size to categorize thermostats into meaningful user types. With a sample size of thousands rather than dozens, future research might employ machine learning techniques such as random forest models to assign AC user types. However, given the findings, this might not be warranted. Focusing on observed, quantifiable thermostat use behaviors may yield more actionable results than attempting to classify users.

The thermostats analyzed in this study were all newly installed, so the team observed how people use their thermostats while they are learning their preferences. If conducted, a thorough study of AC user types should observe users who have had their thermostats installed for varied lengths of time to understand how behaviors may change. The findings in this study are best understood as a detailed look at how people use their thermostats during the first cooling season after installation, understanding this behavior may change over time.

The following sections focus on observed behaviors: cooling setpoints, mode changes, proportion of time spent in different modes or states, and hold setting behaviors.

3.1.2.2 Cooling Setpoint Behaviors

Cooling setpoints are programmed into the thermostat, so that when the thermostat is in cooling or auto mode (i.e., AC is enabled), as soon as the indoor temperature exceeds the setpoint, the AC will turn on and the thermostat will enter a cooling state (i.e., the AC runs). Customers may choose their own setpoints or leave the factory settings in place on their thermostat. The thermostats included in this study allowed the user to program up to four distinct setpoints per day.

The evaluation team generated and analyzed weekly summary plots of the cooling setpoints for all thermostats. In Figure 27, each thin black line traces weekly setpoint metrics for a single device over time. These metrics are the weekly high, weekly average, and weekly low setpoint for each device.[24]

Figure 27. Weekly Maximum, Mean, and Minimum Cooling Setpoint Values by Device

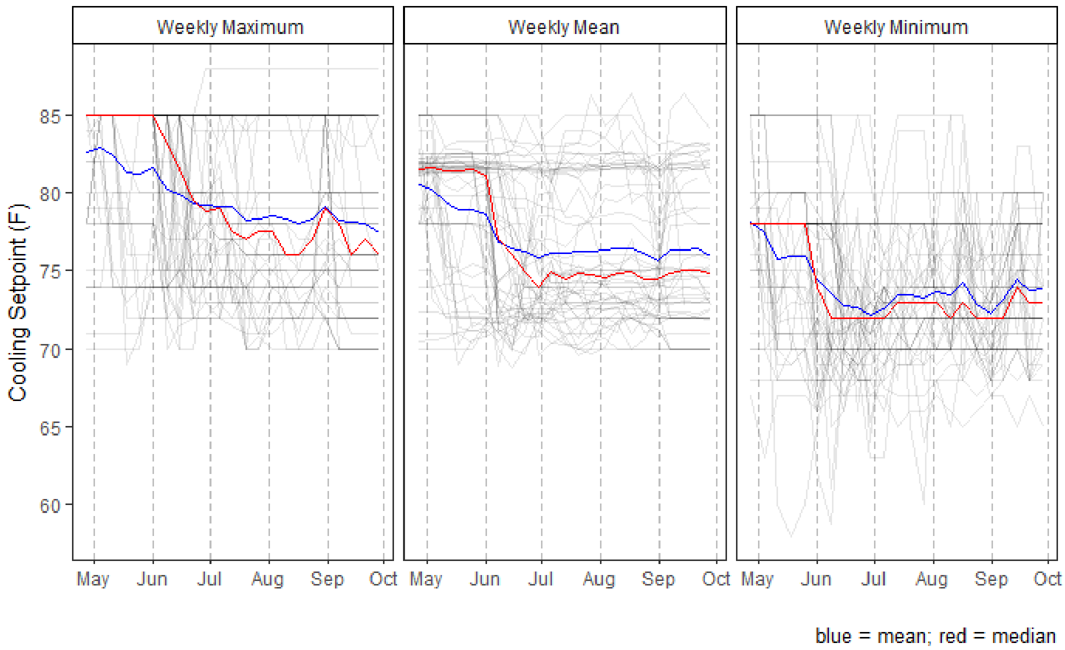

Source: Navigant

In each of the three subplots of Figure 27, the median weekly setpoint (in red) across all devices is constant until mid-June. Because these thermostats were newly installed, this delay suggests that most users left their thermostat in their default factory setting for cooling setpoint schedules until warmer temperatures motivated them to adjust them. The mean weekly setpoint (in blue) across all devices in each plot gradually declines over the first few months of data, indicating that even from the first week some users were adjusting away from the factory default settings.

The team observed a split mid-summer between those with high and low average weekly setpoints. Beginning in July, there is a clustering of mean setpoints around 83°F and another, less distinct clustering between 70°F and 75°F, indicating a large group of users kept their average weekly setpoint relatively high throughout the summer, while another large group chose much lower setpoints.

Despite most devices eventually lowering their maximum weekly cooling setpoint there were some devices for which the maximum weekly setpoints remained at 85°F throughout the summer. These users spent at least part each week throughout the summer with 85°F as their maximum setpoint. These could be people who purposefully choose 85°F as their maximum setpoint at least some of the time during the week, but could also reflect people who never altered the default settings.

Minimum and maximum setpoints in mid-summer tended to be clustered cleanly around discrete temperatures such as 80°F, 82°F, or 84°F. The evaluation team observed the strongest maximum

[24] In order to be recorded as a weekly maximum or minimum setpoint, the team required the setpoint duration to be at least 15 minutes.

setpoint clusters at 74°F, 80°F, and 85°F, and the strongest minimum setpoints clusters at 70°F, 72°F, and 78°F. The most common weekly maximum and minimum setpoints across all thermostats were 85°F and 78°F, respectively.

The team's findings of large differences in setpoint behavior by different groups of users suggest the PAs may be able to target messaging to customers with different setpoint behavior to optimally impact behavior.

3.1.2.3 Time in Cool or Auto Mode

Thermostat modes are cool, heat, auto, or off, and in cool or auto mode the AC will start running whenever the indoor temperature exceeds the cooling setpoint. In other words, cooling mode is when the AC is enabled. When the AC system is running, this is referred to as being in cooling state.

The thin black lines in Figure 28 show the weekly percentage of time spent in cool or auto mode for each device throughout the summer. During the warmest part of the summer, fewer people had their devices set in off mode, more were in cool or auto mode. Near the end of summer with the onset of cooler fall weather, more of the devices returned to off mode once again. This trend is illustrated by the upward and then downward trending mean and median (blue and red) lines.

Figure 28. Weekly Time in Cool or Auto Mode by Thermostat

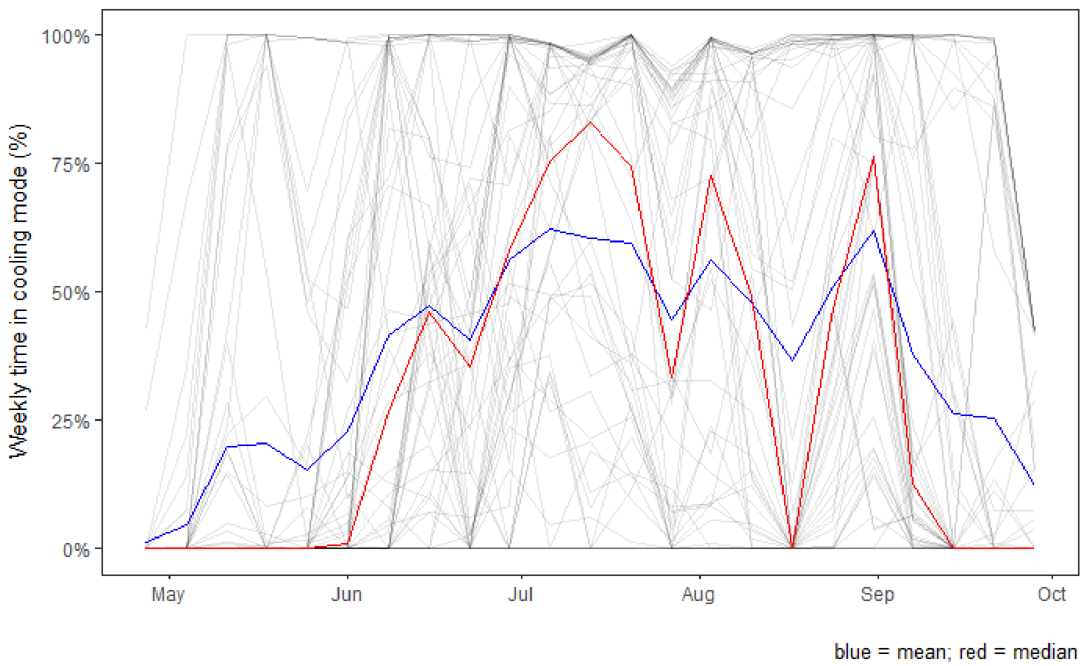

Source: Navigant

Figure 29 shows monthly histograms of the proportion of thermostats spending a given percentage of their time in cool or auto mode. This plot shows that a large proportion of customers spend 100% of their time in cool or auto mode (i.e., with AC enabled) during the hottest summer months. By contrast, throughout the summer, a considerable portion of customers have AC enabled (i.e., system in cool/auto mode) 25% or less of the time.

The bimodal distribution of customers between those that have AC enabled (i.e., in cool or auto mode) most of the summer, versus those that usually have it disabled, suggests that PAs may target these customers with different messaging to affect behavior efficiently.

Figure 29. Distribution of Time in Cooling Mode During the Cooling Season

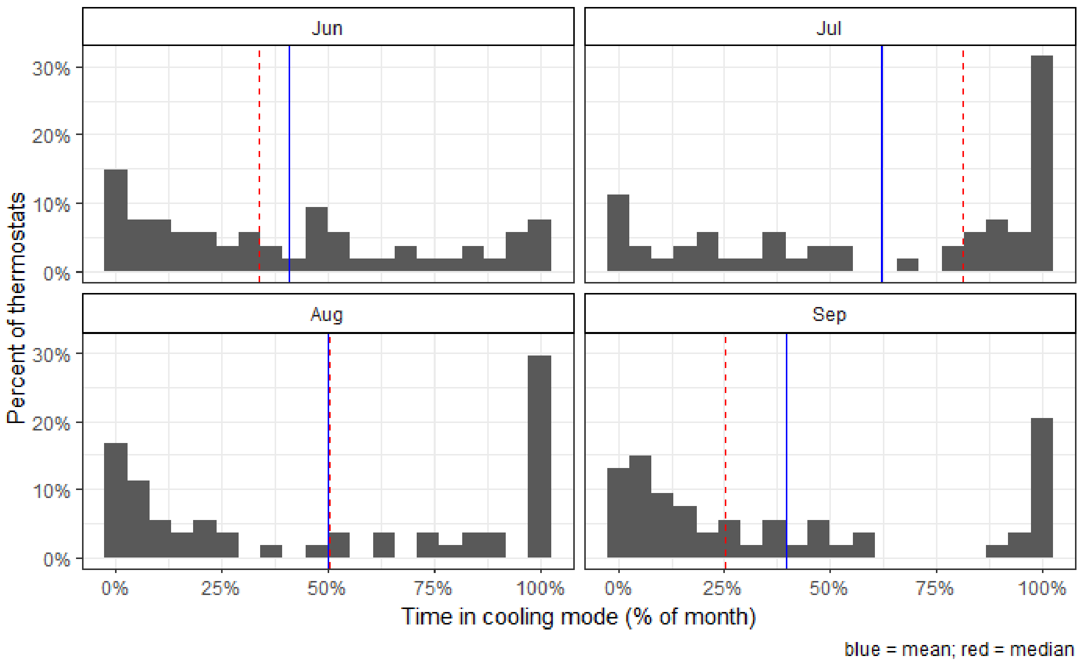

Source: Navigant

3.1.2.4 Time in Cooling State

Cooling state means the AC is running. This can occur when the thermostat is set in auto or cool mode (i.e., AC is enabled) and the indoor temperature exceeds the cooling setpoint, or when the customer manually adjusts their system to force the AC on.

Each of the faint black lines in Figure 30 traces time spent cooling by each device per week throughout the summer. As expected, devices spend more time with the AC running during the hottest summer months. The mean and median (blue and red) lines, during their peak in July, show the average user has their AC running between 10% and 15% of the time during the hottest part of the year, which is a relatively little. By contrast, a small proportion of AC units are running more than 40% of the time during the hottest weeks. In this sample of data, the team did not see any devices running for 100% of the time for any given week. In other words, the “always on” user type is not generally a relevant user type categorization in Massachusetts. In fact, this figure shows moderation is more the norm in this sample of Massachusetts users than extremes.

Figure 30. Weekly AC Runtime by Home

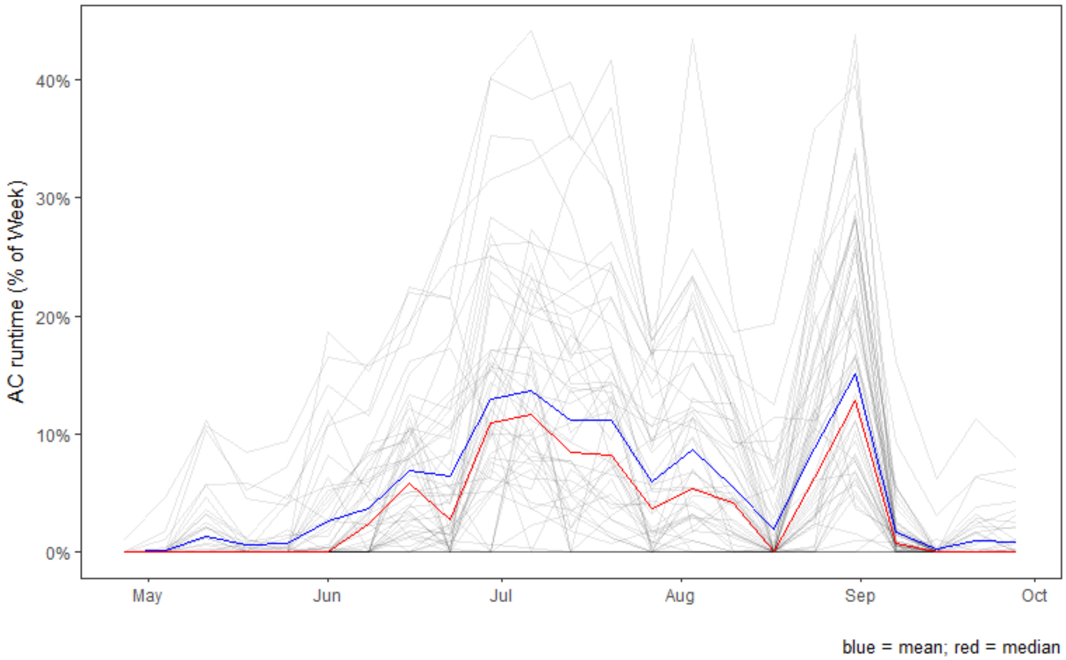

Source: Navigant

Figure 31 shows the proportion of thermostats with AC running (i.e., in cooling state) for a given percentage of time by month. All monthly plots are left-skewed, with most customers running their AC less than 10% of the time throughout the entire summer. Homes running their AC more than 25% of the time are extreme outliers. The average customer has their AC running about 5% of the time in June, August, and September, but closer to 10% during July, the hottest month.

The team's findings on time spent cooling suggest that most customers are already using their AC moderately. The PAs might consider targeted messaging or revisit the training of thermostat installers to educate customers who have their AC running during summer months 25% of the time or more on more efficient thermostat use and settings.

Figure 31. Distribution of AC Runtime during the Cooling Season

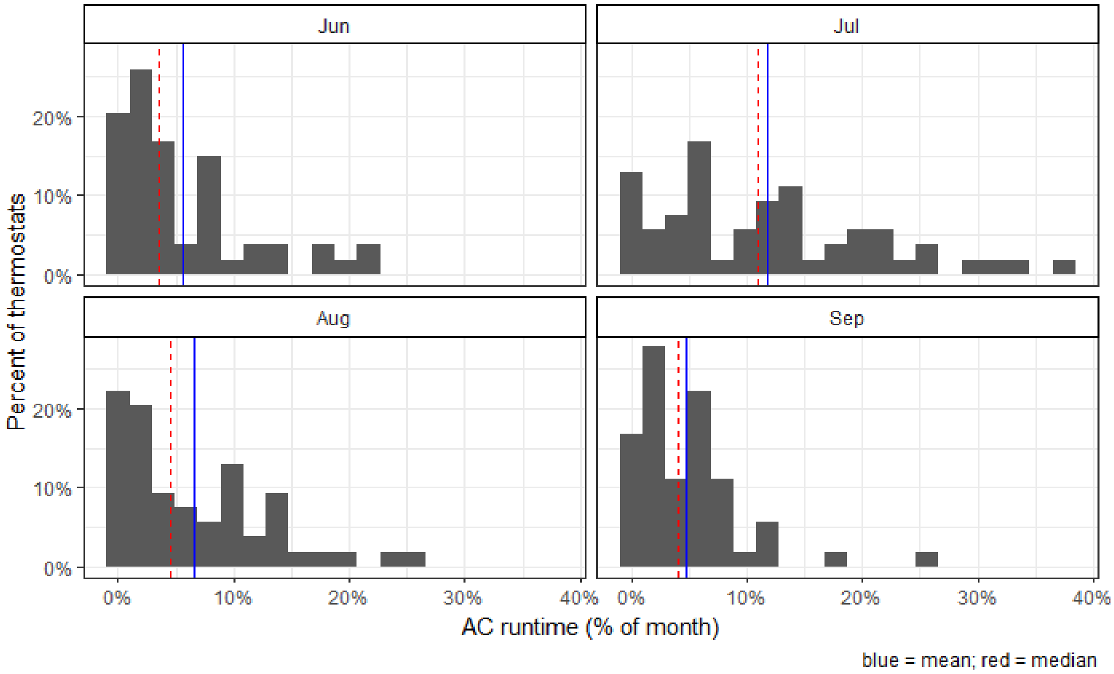

Source: Navigant

3.1.2.5 Mode Change Behaviors

The evaluation team did not find evidence that many people use the system mode to control their AC unit instead of using setpoint schedules. Assuming these people were using their AC systems at least 5 days per week and were only switching the thermostat in and out of cool mode once per day for cooling, the team would expect to see 10 mode changes per week. However, the data shows this behavior is rare. As illustrated in Figure 32, the average customer changes their mode less than 5 times per week, even during the hottest months. Both the mean and median (blue and red) number of changes decrease during the hotter weeks. This suggests there may be a threshold temperature at which people stop manually changing the mode of their thermostat and instead leave it in one mode, relying on the setpoint schedule to dictate when the AC turns on. They may also engage the hold setting when they reach certain temperature thresholds. The use of hold settings is explored further in Section 3.1.2.7.

Figure 32. Count of Weekly Thermostat Mode Changes by Customer

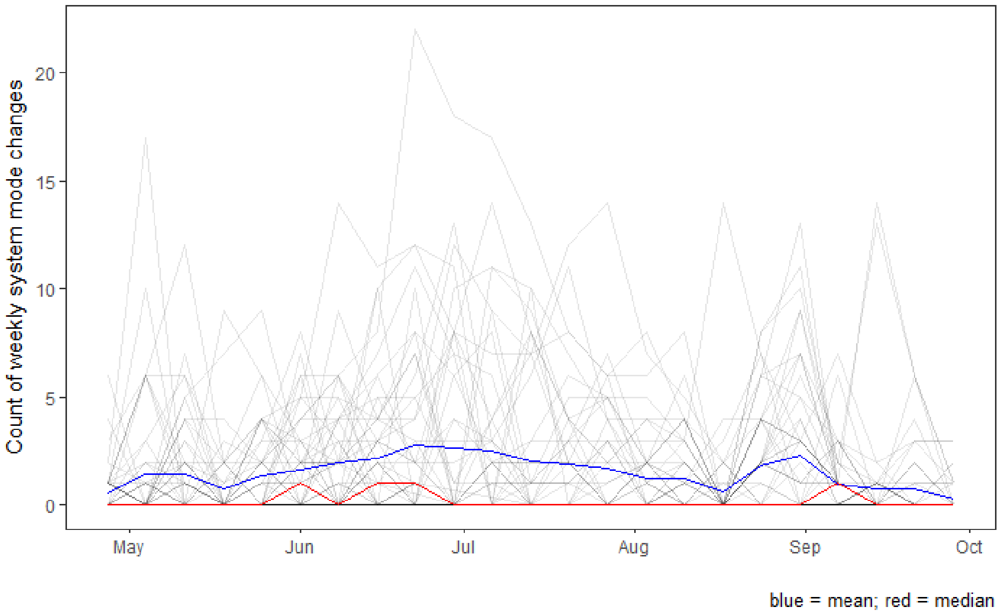

Source: Navigant

3.1.2.6 Setpoint Change Behaviors

Setpoints are the temperatures chosen and programmed into the thermostat which dictate when the AC or heat will run. During the cooling season, whenever the AC system is enabled (i.e., mode is auto or cool), as soon as the indoor temperature exceeds the chosen cooling setpoint, the AC will begin running to cool the home.

The black lines in Figure 33 trace the number of times setpoints are changed weekly for each thermostat throughout the summer. Users who leave their device at a constant setpoint throughout a week would have zero setpoint changes in these plots. Users who generally follow a programmed schedule with four setpoint changes each day of the week (e.g., wake up, away at work, home in the evening, and overnight) would have approximately 28 setpoint changes per week (shown by the orange dashed line).

The plot shows most users have fewer than 28 weekly setpoint changes throughout the summer. Thermostats remaining in a set schedule with no manual interventions would appear as long horizontal lines extending across the graph. Instead, almost every thermostat trace shows variation, suggesting most users are making some modifications to their setpoints throughout the summer. Dark line clusters around specific values early in the summer likely reflect factory default settings the users did not yet adjust. The number of average weekly setpoint changes decreases throughout the summer, perhaps indicating that customers who were fiddling with their setpoint temperatures (manually overriding the schedules) early in the season eventually settled on a schedule that worked well for them, requiring less fiddling, or acclimated to the summer heat. This may also reflect more AC systems being shut off for the season (e.g., mode set to off), hence less changes occurring overall.

Figure 33. Count of Weekly Cooling Setpoint Changes by Home

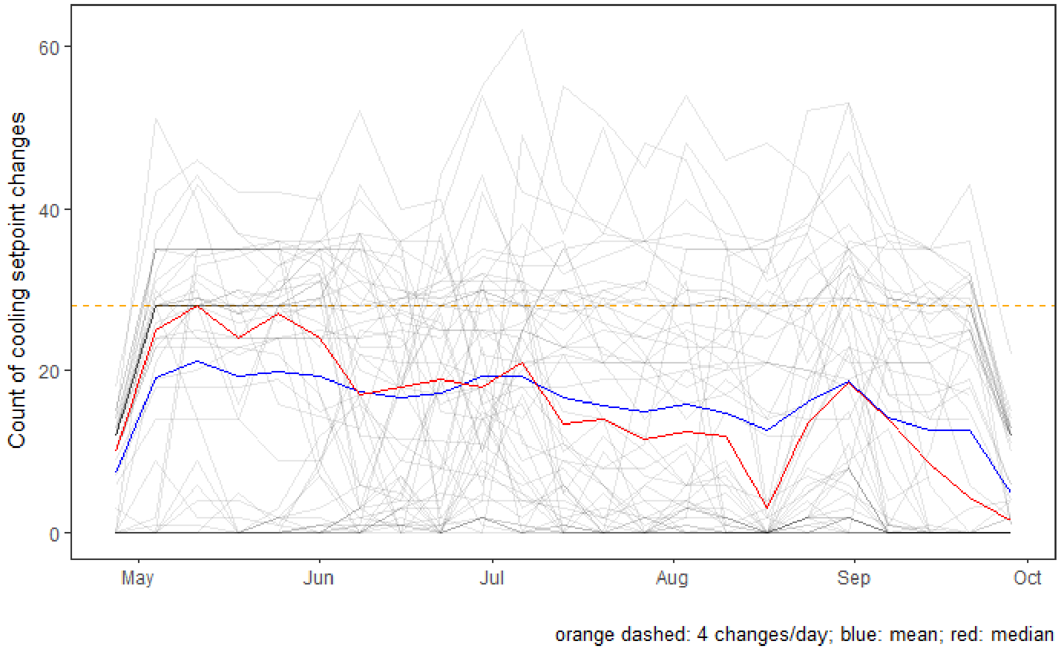

Source: Navigant

Figure 34 summarizes the median weekly number of setpoint changes across the summer.[25] Over 15% of the study population typically keeps their setpoint at a constant temperature without any manual or scheduled setpoint changes. By contrast, approximately 13% of the population follows a four setpoint per day schedule. While many users fall outside these two categories, zero and four setpoint changes per day are the most prevalent categories across all users.

[25] The team used median rather than mean in this case because the mean is highly sensitive to outliers; as a result, devices with even short periods of many setpoint changes at some point in the summer would not appear as typically having zero setpoint changes using the mean. The team feels the median presents a more balanced picture of behavior than the mean in this case, particularly when examining the behavior of those with typically few changes.

Figure 34. Distribution of the Median Count of Cooling Setpoint Changes per Week

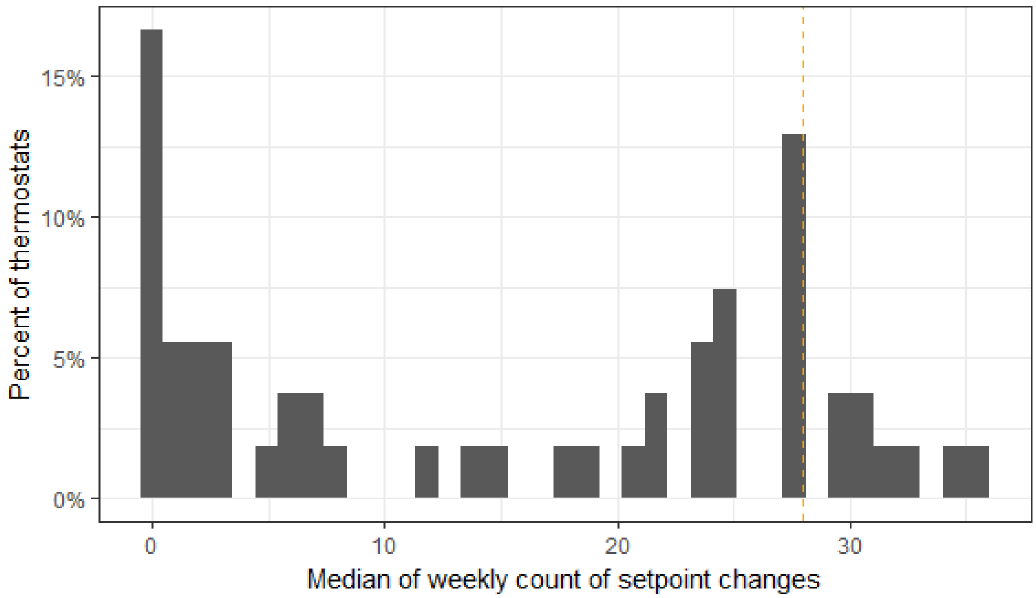

Source: Navigant

These findings suggest the PAs may achieve greater impact in messaging by targeting messaging differently to customers displaying constant setpoint behavior, versus those demonstrating a regular schedule with four distinct settings per day, as a large proportion of customers fall into these categories. Messaging suggesting customers choose higher setpoints during the day while they are typically away, and during the night while they are asleep will not resonate with customers that maintain a single setpoint throughout the day. Moreover, the PAs may consider updating training for thermostat installers to ensure they are educating customers who might set a single, low setpoint about the benefits of setting more efficient multiple setpoints.

3.1.2.7 Hold Setting Behaviors

Most programmable thermostats have a “hold” function that enables the user to override preset schedule setpoints and changes in response to their comfort considerations in the moment or to save energy. For example, a customer might have a programmed schedule set to increase the setpoint to 80°F at 7 p.m. when they typically open their windows to let in the cooler evening air. But perhaps this user has company one evening and wants to keep the temperature steady without opening windows—they may engage the hold setting to keep the temperature at 75° F until their company leaves, then release the hold. A less efficient use of hold would be a customer that does not have a schedule programmed, chooses 65°F, and then presses hold to keep the air running and the indoor temperature low.

As shown in Figure 35, many people have hold engaged either zero or 100% of the time during a given day. While the mean is stable and right around 50% of the time during a day in hold, the median is highly variable, and this variability appears most in the hottest summer months and late in the fall shoulder season. This suggests hold use might be correlated with specific weather events. It also suggests customers who over-use the hold function may not have set practical setpoint schedules, and so find themselves constantly overriding them by using hold.

Figure 35. Percent of Time with Hold Setting Engaged

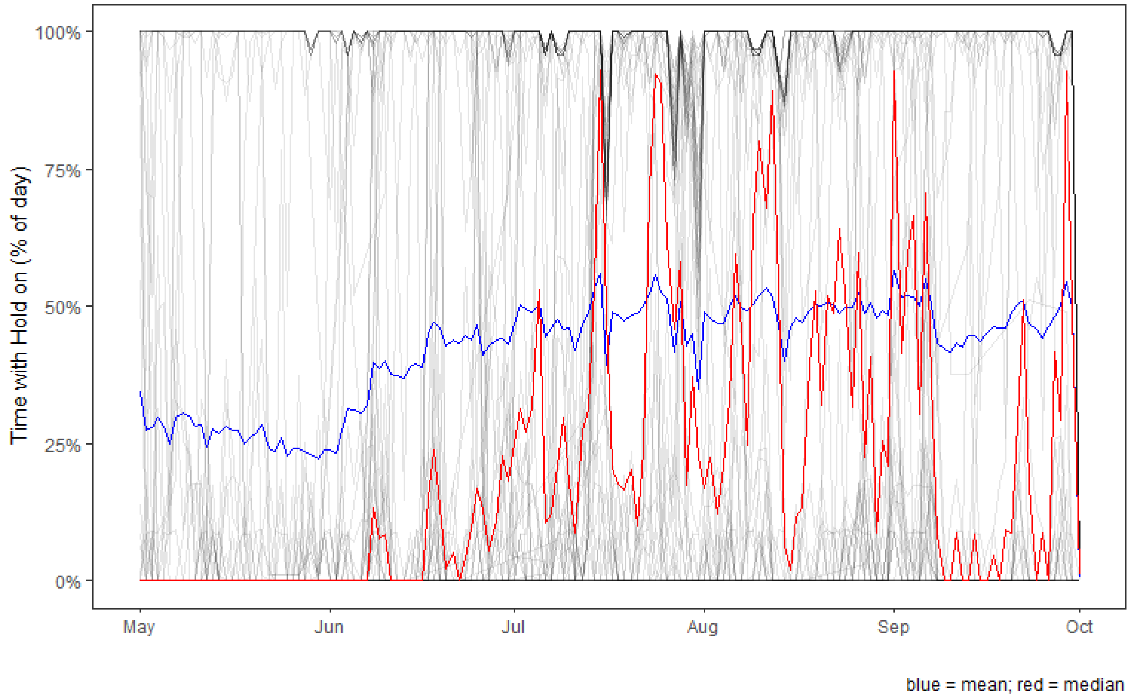

As shown in Figure 36, cooling setpoints are generally lower when hold is on versus off, but the difference is not large. This finding would be consistent with people mostly using hold for lower setpoints to override preprogrammed schedule changes that would have increased the temperature.

Figure 36. Comparison of Cooling Setpoints when Hold is On vs. Off

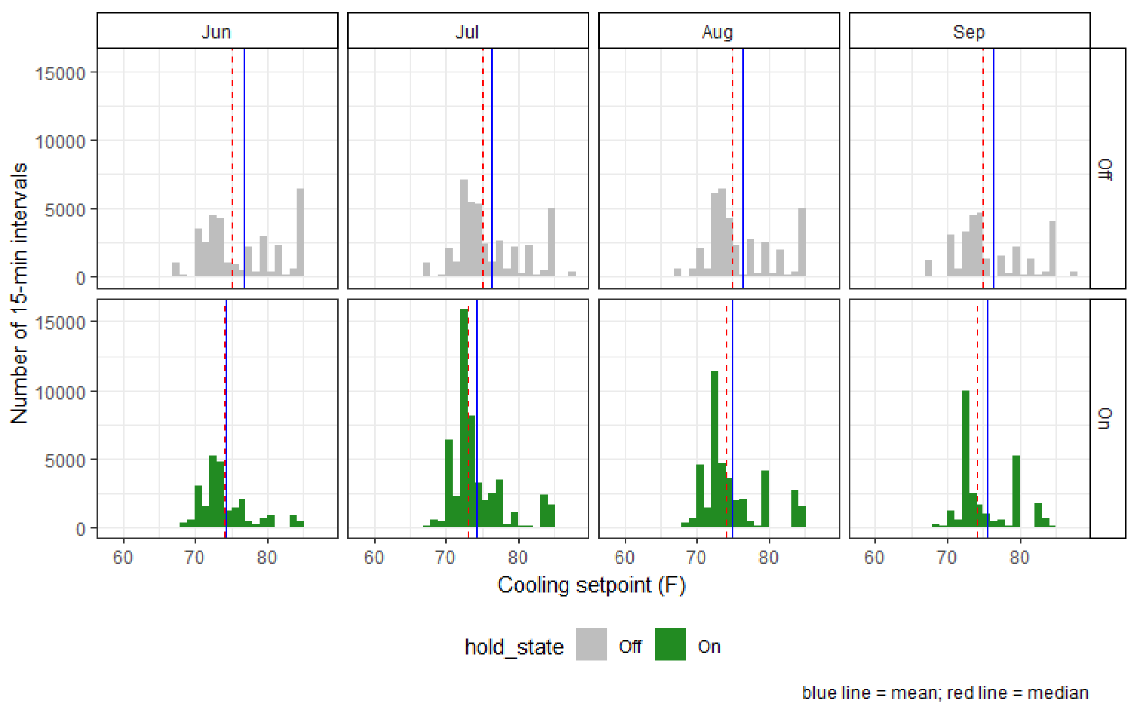

As shown in Figure 37, it appears that many customers use the hold setting to keep the AC running—mostly during the hotter periods of the summer—while many also use hold to keep the AC off, a behavior which is distributed more evenly throughout the summer.

Figure 37. Summary of Hold Setting and Runtime Throughout Season

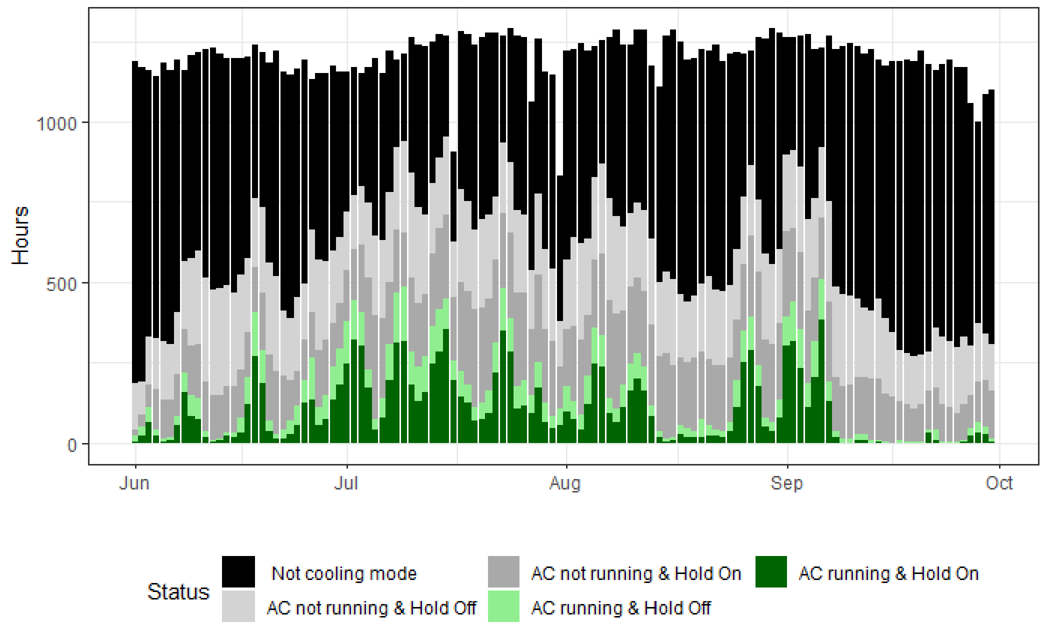

Some thermostat models revert from hold to the next scheduled setpoint by design, enabling more efficient use and avoiding prolonged or indefinite holds. The PAs may consider prioritizing incentivization of these models given this evidence that a large portion of users may over-use the hold feature.

3.1.3 Thermostat Use Behaviors and AC User Type Considerations

The following paragraphs draw conclusions based on the evaluation team's research on how people use their AC systems and thermostats. The team formulated its discussion around these findings as a series of actionable considerations.

Message to user behaviors, not user types.

A key finding in this study was the enormous diversity in how people use their thermostats and AC systems. Increasingly sophisticated thermostats allow a broad range of behaviors, and the team's results show that static "user type" definitions may not be as useful as previously thought, because most customers change their behavior significantly throughout the summer cooling season. A sensible approach is to message to and target interventions toward specific, quantifiable thermostat behaviors.

The team suggests the PAs may consider focusing messaging and targeted behavioral interventions based on observed user behaviors common in their user population. As an example, messaging in bill inserts may ask:

- “Do you find yourself using “Hold” a lot to override your schedule? You can save energy by setting a schedule that’s comfortable for you and letting the thermostat do the work for you!”
- “Do you often get cold in your house and have to adjust the thermostat? Pay attention to when, and you can reschedule your thermostat with warmer setpoints to save energy!”

The PAs may also consider revisiting training for auditors who install thermostats through PA programs to ensure they are educating customers about how to choose efficient setpoints, and how to program schedules and use features efficiently.

Use behavioral norming based on actual Massachusetts AC user data.

One of the evaluation team’s key findings is that most Massachusetts customers avoid extremes, running their AC a limited amount of time, even during the hottest months. In fact, most Massachusetts customers run their AC less than 10% of the time during the summer months, and almost no customers would be considered “always on” users, choosing setpoints so low their AC is running nearly all of the time. However, a small proportion of customers have their AC system running nearly 40% of the time during hot spells—meaning these customers may be using much more energy to cool their homes than their peers.

Research shows a key behavioral lever is social norming. By letting people know what is “normal” in their population, they tend to want to adjust their behavior to conform with the “norm.” As a result, the PAs can use the information provided in this study on how Massachusetts customers actually use their AC to educate customers about what is normal, thereby causing higher users to dial back consumption. Many customers may be surprised to learn that most of their peers run the AC less than 10% of the time, even during the hottest summer months—especially if they are customers that have their AC running 40% of the time during hot spells. The surprise factor of learning their behavior is not aligned with the “norm” may change the way they use their AC.

Teach users who over-use the “hold” setting to set more rational setpoints.

Many customers use the hold setting excessively, signaling they have not set comfortable and practical setpoint schedules that fit their lifestyle and the weather. In fact, the average customer has hold engaged 50% of the time during a given day.

The team observed the highest variability in hold use during the hottest weeks of the summer, suggesting many customers use hold to keep the air running during the heat of the day. Hold was also used more when the current setpoint was low, consistent with customers pressing hold to keep the temperature low and the air running. Customers who over-use the hold function may not have set practical setpoint schedules, and so find themselves constantly overriding them by using hold. Through targeted messaging, the PAs can help educate people about saving money on their energy bills by setting more reasonable setpoint schedules that keep them comfortable during the heat of the day, rather than relying on lengthy “hold” periods to keep the air running.

Some of people’s inefficient use of hold may be exacerbated by the thermostat’s inability to effectively account for humidity in keeping people’s homes at a comfortable temperature. The PAs may consider incentivizing only thermostats that take humidity into account in managing a home's indoor temperature to minimize inefficiency, as much of the unwanted hold behavior may be related to the thermostat's inability to effectively incorporate humidity into its algorithms.

Message differently to those with high vs low average setpoints.

By mid-summer a bimodal distribution of average temperature settings appears. Beginning in July, as the temperature heats up, a large number of users schedule their average cooling setpoint around 83°F, while another large portion of users set it closer to 73°F.

The PAs can message differentially to these groups of customers to maximize the impact of behavioral interventions. For instance, it does not make sense to message to someone who reliably has their setpoint at 83°F throughout the summer to be more energy-conscious and further set back their thermostat—they are already doing it. However, by identifying the set of users with setpoints in the low 70s to upper 60s throughout the summer, the PAs could nudge these customers to set back their thermostat a few degrees, thereby saving energy. As noted above, social norming can be a strong behavioral lever for influencing the way customers use their AC—by letting especially low average setpoint users know their behavior is not typical, the PAs can nudge them to adopt more efficient settings.

3.2 Impact of Weather on First and Last AC Use

The evaluation team investigated what motivates users to switch their thermostats into cooling mode for the first time during the season and then to switch it off again for the remainder of the year. The team paid particularly close attention to the role weather plays in these decisions. Having a better understanding of what causes people to turn their AC on and off for the season could help the PAs encourage energy efficiency savings, and potentially avoid costly demand spikes as large numbers of customers first turn on their air on the same days during the summer.

While the team investigated both first and last AC use, it did not obtain strong results correlating weather and last use; thus, most of the discussion that follows focuses on the first-use question.

3.2.1 Data Exploration

The evaluation team plotted the distributions of likely predictors of first AC use, and correlations between those predictors and occurrences of first use, to better understand these relationships. Thermostats in the dataset were installed throughout the summer, so not all records were complete during the first summer months. Visual data exploration helped refine the choice of explanatory variables in the evaluation team's regression models and determined appropriate cutoffs for the datasets to ensure complete early or late season observations to use for either the first or last-use models.

Based on this analysis, the team limited first and last AC use modeling to 54 thermostats that were installed before May 1, 2014, and with at least 80% data completeness during the cooling season.[26]

[26] For more information, completeness of thermostat data is discussed in Section 2.2.2.

3.2.1.1 Findings—First Use

Figure 38 summarizes maximum daily cooling degree hours (CDH), maximum daily temperature, maximum daily humidity, and the daily count of first AC use instances throughout the cooling season. Dots indicate days when first AC use events occurred, and the color of each dot represents the number of first uses occurring that day—the lighter the green, the more AC turn-ons that day. Accordingly, dots located atop weather spikes indicate that both the weather variable and AC turn-ons spiked on the same day, and the lighter green the dot, the greater the number of turn-ons.

The bottommost subplot shows that most first AC uses occurred in May and June, with fewer occurring later in the season because most homes that use AC would have already turned them on. There were 3 days in May and June where five or more AC devices turned on for the first time in the season, and all corresponded with high temperatures shown in the maximum temperature subplot. For example, over the course of 4 days in May 13 ACs ran for the first time. During those same 4 days, the daily high temperature increased from below 60°F to above 80°F. This was a major temperature swing and many people in this sample of the data responded by turning on their ACs for the first time.

Contrary to the team’s initial expectations, the days with the largest number of first-use occurrences were also some of the days with lowest maximum humidity. Similarly, the days with highest humidity had few first use counts. It should be noted that this suggests humidity does not play a strong role in determining the timing of when people first turn on their AC for the season—it does not suggest that humidity and ongoing AC use are not correlated.

This apparent negative correlation between humidity and AC turn-ons may be due to high humidity occurring on rainy or overcast days, during which few people would be motivated to turn on their ACs for the first time. It may also suggest that temperature sends a stronger signal to people to turn their AC on for the first time in the season than does humidity. High humidity spikes are distributed fairly uniformly throughout the cooling season from start to finish. This means many of the high humidity spikes occurred on days when the maximum temperature was low (in the 60s). Some of the high temperature spikes occurred on days with low humidity and a high number of first AC turn-ons. This further suggests heat is the primary signal motivating people to turn on their AC systems for the first time each season.

Figure 38. Daily Relationship between Weather and First AC Use

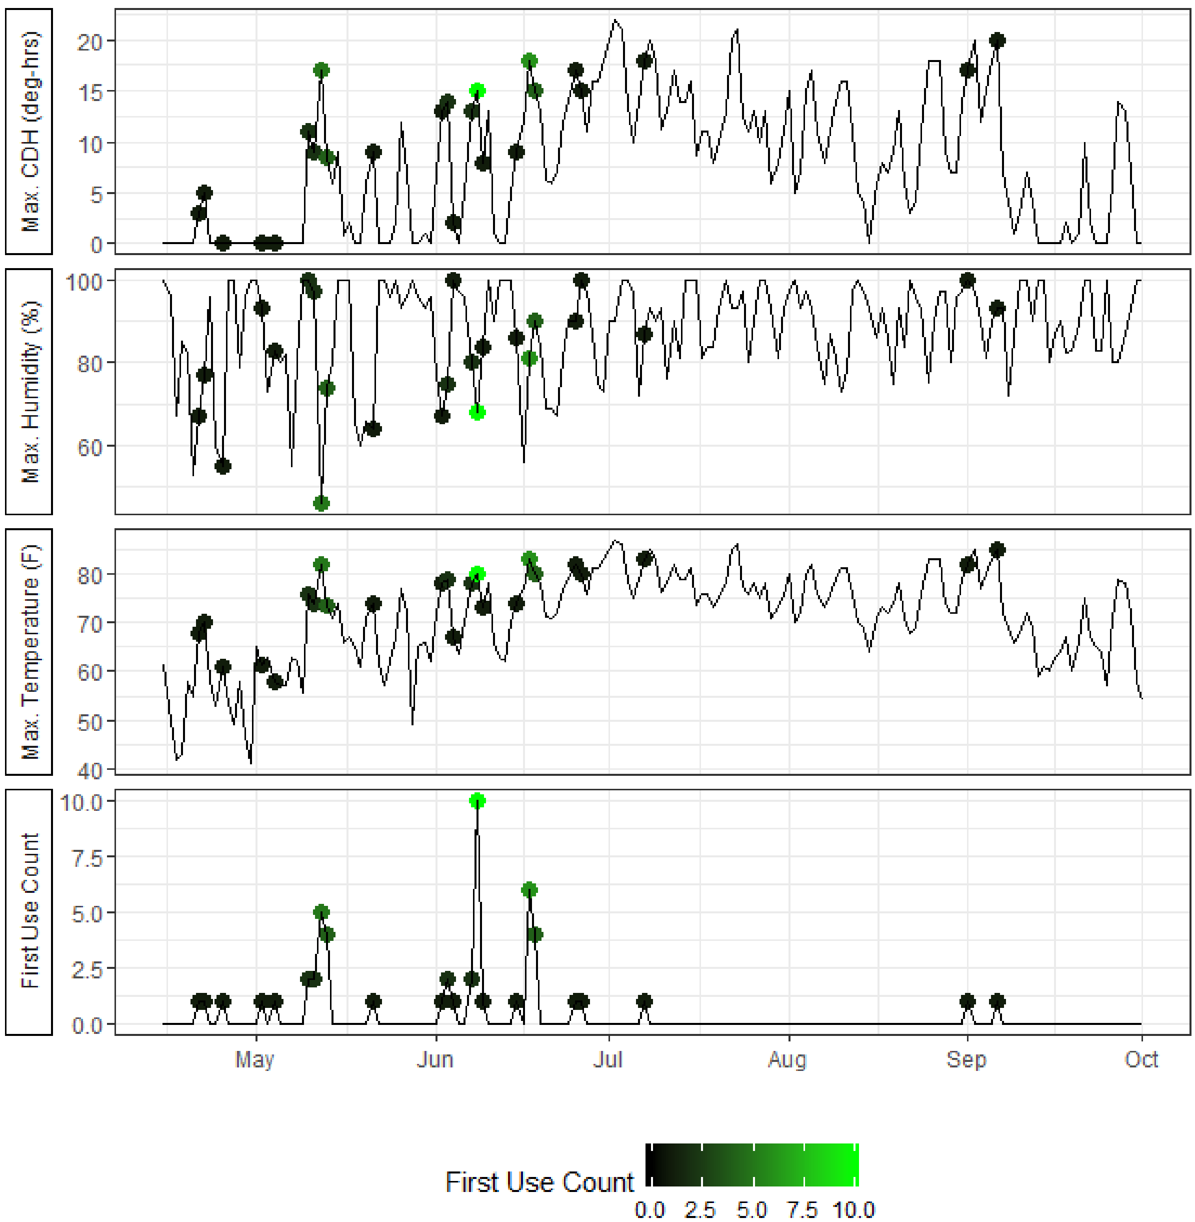

Source: Navigant

3.2.1.2 Findings—Last Use

An obvious difference between the first-use plots in Figure 38 and the last-use plots in Figure 39 is that last use is much less spread throughout the season. Last use is tightly clustered, mainly in September, and the majority appear to occur in response to a single cold snap. While not conclusive, this suggests that people may be looking to turn off their AC (and save money) for the season and react strongly to a single temperature drop if they believe it signals the onset of cooler weather for the year. As an example, people demonstrate the same type of behavior in the spring when looking for the right temperature signal to suggest it is time to turn on their sprinkler systems—they look for evidence of the last hard freeze. People may be using the same type of method in deciding when to turn their AC off for the season.

These findings come with a significant caveat, however, as they are based only on correlation. Analyzing last AC use in relation to weather simply through data visualization is difficult in that these plots do not indicate whether users actually turned their ACs off (i.e., switched the thermostat from cool mode to off mode) or if the sub-80°F outdoor temperatures were just not warm enough to raise the indoor temperatures above the cooling setpoints, causing the AC to not turn on. In other words, it is impossible to tell from visual inspection of these plots whether the last AC use was a conscious choice or if it was the interplay of temperature and setpoints. While the following sections explore these relationships between weather and last use more deeply using regression analysis, these results suggest modeling last AC use as a function of weather may be more difficult than first use and require a lengthier period of observation.

Figure 39. Daily Relationship between Weather and Last AC Use

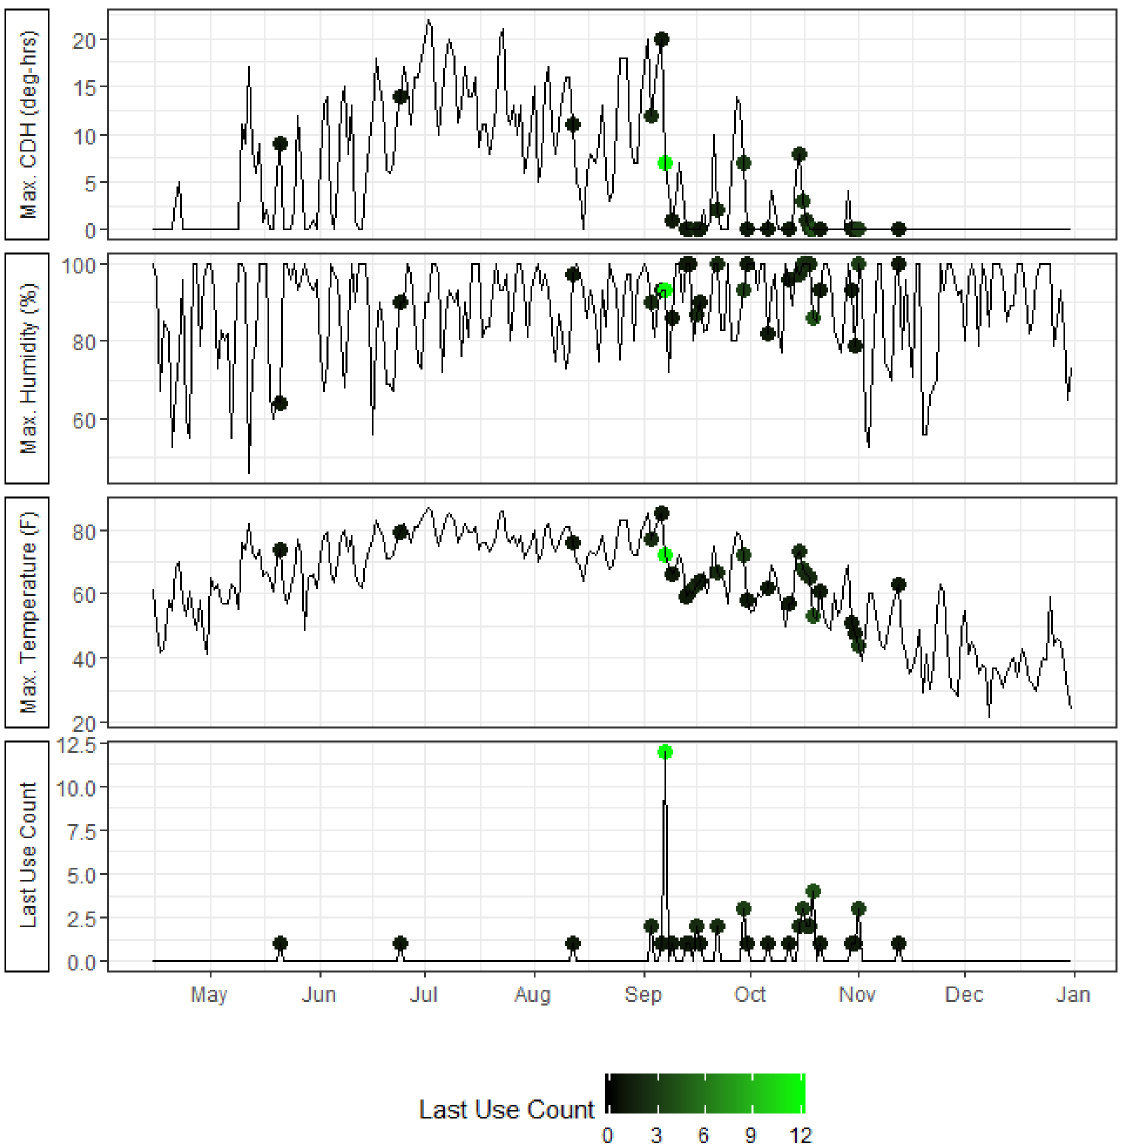

Source: Navigant

3.2.2 Regression Analysis

Building on the learnings gained through data visualization, the evaluation team conducted regression analysis to gain a more precise understanding of the causal effect of weather variables on when people turn on their AC systems for the first time during a cooling season. Moving beyond visual inspection of correlations, regression analysis allows the team to quantify effects and uncertainty, draw causal inferences, and disentangle the interactive effects between predictive variables. As discussed above, last-use modeling did not produce meaningful, significant results, and so these regression results are presented in Appendix Section A.4.

3.2.2.1 Methodology

For both the first and last AC use models, the dependent variable is a binary variable indicating whether first use (or last use) has occurred. Given the binary outcome variable, the evaluation team chose a logistic regression as the basis for the model. The team compared several different model formulations for estimating the impacts of weather on first AC use before arriving at the preferred model. To evaluate model performance, the evaluation team compared regression-based predicted values for the cumulative AC turn-on rate with actual cumulative AC turn-on values throughout the season for several model variations.[27] The team used a method known as cross-validation analysis to evaluate model performance and demonstrate that the preferred model performed well. Full methodological details including model specification and cross-validation results are provided in Appendix Section A.2.

3.2.2.2 Findings

Due to the nature of the logit model regression formulation, it is difficult to interpret the meaning of its regression coefficients directly. Instead, the evaluation team translated these to easily interpretable average marginal effects (AME), shown in Table 1.

[27] Logistic regressions are a form of classifier model. A typical approach to measuring performance of these classifier-type models is to specify a probability threshold (e.g., 50%) above which one considers the model's prediction to be positive (in this case, first use occurring) and then construct a confusion matrix (rates of true-positive, false-positive, true-negative, and false-negative) or the area under the curve (AUC) of the receiver operating characteristic (ROC) curve. Both of those approaches look at the model's ability to correctly classify specific observations as either a first use or not first use day. However, for the team's analysis, *the accuracy of the model at the individual observation level is less important than the cumulative accuracy across all observations*. Stated another way, the cumulative probability of all ACs having experienced first use is more useful than estimates for each individual home on each individual day. For this reason, the team chose to compare cumulative predicted turn-on rates with actual rates to evaluate performance.

The AMEs of the key variable are all statistically significant, and their magnitude and direction make intuitive sense. The team's findings suggest that the probability of someone turning on their AC unit for the first time in the season increases with the temperature above 65°F, captured through the MaxCDH term.[28] Additionally, as shown by the Weekend term, the model suggests people are more likely to turn on their AC for the first time if it is a weekend day rather than a weekday. As evidenced by the MaxTemp variable's small negative effect on the probability of turn on, temperatures below 65°F do not motivate people to turn their AC on for the first time in a season.

Aligned with the exploratory analysis findings in Section 3.2.1.2, humidity does not appear to motivate people to turn on their AC for the first time during the cooling season. The negative effect of humidity on the probability of first AC turn-on reflects that by the time hot temperatures cause people to turn their AC on for the first time, they have already experienced significant high humidity events. In other words, heat, not humidity, is the main motivation for people in the sample to turn their AC on for the first time during the cooling season.[29]

Table 1. Average Marginal Effects of the First-Use Model

Term	AME	Std. Error	Z Score	p.value
MaxCDH	0.00503	0.00083	6.08	1.24E-09
MaxHumidity	-0.00053	0.00016	-3.35	8.18E-04
MaxTemp	-0.00098	0.00024	-4.06	4.83E-05
Weekend[30]	0.01959	0.00748	2.62	8.82E-03

Source: Navigant

The interpretation of each AME is discussed in more detail below:

- **MaxCDH:** For every 1°F-hour increase in the day's maximum CDH value, the probability of an AC turning on for the first time increases by approximately 1 percentage point.
- **MaxHumidity:** The coefficient on maximum daily humidity suggests that for every 1% increase in the day's maximum relative humidity level, the probability that a given household will turn on their AC for the first time in the season decreases by 0.05 percentage points.[31]

[28] Recall, as discussed earlier, the base for the CDH calculation is 65°F, so maximum daily CDH captures the number of hours above 65°F on a given day.

[29] As discussed in detail in Section 3.3, even though humidity is not a strong driver of first AC use, it does play an important role in ongoing AC use throughout the season.

[30] As specified in the model formulation (Section 3.2.2.1), the model contains the Weekend term interacted with MaxCDH. While there exists a coefficient for this interaction of the terms, AMEs are (by definition) only calculated for the constituent terms. Thus, the AME for the Weekend term here truly represents the average difference in probability between weekend and weekdays.

[31] This negative AME of humidity may be the result of rainier, more humid cool days early in the season being correlated with lower probability of first time AC use. However, it may also simply signal that people are not responding to humidity when making the decision to turn on their AC for the first time. Because humidity peaks far earlier in the season than temperature, and there are not many turn-ons correlated with high humidity days (see Figure 49), it more likely indicates high humidity does not motivate people to turn on their AC for the first time in the season. As shown in Section 3.3, the relationship between humidity and general AC use throughout the season (as opposed to predicting first turn-on) is much stronger and positive.

- **MaxTemp:** For every 1°F increase in the day's maximum temperature, the probability of an AC turning on for the first time decreases[32] by approximately 0.1%.
- **Weekend:** Weekend status has a positive effect on the likelihood that an AC will be turned on for the first time. The AME indicates that, all other conditions being equal, customers are 2 percentage points more likely to turn on their ACs for the first time on the weekend than they are on weekdays.

3.2.3 Considerations Around Weather and First AC Use

The following section draws conclusions based exploratory research and regression modeling of how weather and other factors affect when people turn on their AC for the first time and the last time in a season. The evaluation team formulated its discussion around these findings as a series of actionable considerations.

Nudge people to turn on their AC later in the season.

The team's research suggests customers are responding to temperature when deciding to turn on the air for the first time during the season, and not humidity. While humidity plays a role in ongoing AC use, the decision to begin using AC during a given summer season is primarily driven by temperature. The team also observe a large degree of variation in when people first use their AC, with most first use spanning May to June.

The team's findings suggest two likely scenarios leading to first AC use:

(1) Early in the cooling season the customer switches their thermostat to cool or auto mode simply as a calendar-driven response to spring; their AC remains off for days or weeks until the indoor temperature exceeds the cooling setpoint, signaling the AC to start running.

(2) During the summer heat, the temperature drives the customer to switch their thermostat to cool or auto mode and choose a cooling setpoint low enough that the AC begins running immediately.

PAs can use this knowledge to affect customer behavior and encourage energy efficiency. For those who turn their systems to cooling mode early in the season, the PAs may consider encouraging them to set conservative cooling setpoints. By nudging them to choose setpoints a couple degrees higher, their AC will kick in later in the summer, saving energy. The PAs can message differently to customers that respond to heat in the moment by actively turning on their AC. Those customers may respond to messages suggesting they leave their home, go to the pool, go out for ice cream or a trip to the mall, rather than turning on their AC the first time they feel uncomfortable in the season. Some of these customers may respond to the behavioral levers of competition and goal setting—encourage them to see if they can avoid turning on their AC until a given date and provide encouragement to leave their houses.

[32] While this result may at first appear counterintuitive, in combination with the main CDH variable, the maximum temperature variable is intended to pick up the effect of temperatures below the 65°F CDH scale; hence, its negative sign makes intuitive sense—customers are less likely to turn on their AC on a given day when the temperature is low. The positive effect of temperature is captured through the main temperature term, CDH_max, which shows a larger and positive effect.

Get people out of their homes on the weekend to postpone first AC use.

This study finds people are more likely to turn on their AC for the first time on a weekend day, relative to during the week.

By offering people alternatives, encouraging them to escape the heat by going to the beach, going to the mall or the theater, the PAs may help people postpone turning on their AC for the first time in the season—especially for the segment of customers who actively choose to turn on their AC for the first time each season in response to thermal discomfort. By offering encouragement to get out of the house on the weekends, PAs can nudge people to turn on their AC later in the season, saving energy.

Encourage people to shut their AC off earlier in the season.

The team's analysis revealed that unlike first use, last AC use in the season was not spread out, but concentrated during a couple weeks in September. A smaller portion of customers continued to use their AC through October and November, even into December.

Two strategies may help PAs encourage customers to shut off their AC earlier in the season. Those customers identified as setting their system to cooling mode early in the season then letting the system turn the AC on when it finally gets hot enough inside might be encouraged to turn their AC system off earlier in the season simply by messaging that it is time to consider switching your system to heating mode. This simple message may be enough to encourage some of these users to switch modes, ensuring their AC does not automatically turn on again for the remainder of the cooling season. The small portion of customers that continue to run their AC through the fall, and even into December, may respond to social norming. Messaging that lets customers know the way they use their AC is atypical may encourage them to act differently. By letting them know that the majority of users stop using their AC in September, for example, may cause some users to react to the slight social pressure by shutting their AC off earlier and not extending use into early winter.

In considering these strategies, the PAs should be aware of potential unintended consequences. For these strategies to work most effectively, customers should have and utilize an off setting on their thermostat. Otherwise, encouraging them to postpone switching to cooling may lengthen the duration of night-time heating, thereby increasing heating use.

3.3 Weather and AC Use throughout the Cooling Season

Customers use their AC systems to maintain comfort throughout the summer in the face of changing outdoor weather conditions. As the temperature rises outdoors, AC systems must work harder to keep homes comfortable, consuming more energy across a given day. In this section the team explores how weather affects people's AC use throughout the season.

3.3.1 Data Exploration

While the evaluation team expected that temperature and humidity would correlate strongly with AC use, the team conducted preliminary exploration of the weather and AMI data to refine its understanding of these relationships prior to regression modeling. Constructing relevant variables from the weather data such as cooling index (CI), total heat index (THI) and heat build-up (HBU), the team then created explanatory graphics including histograms, correlation plots, and load shapes to better understand relationships between weather variables and energy use. The more basic models and data exploration helped inform later, more complex regression analysis work. A more detailed explanation of the methodology employed in this section is provided in Appendix Section A.5.

3.3.1.1 Findings

AMI-based load shape analysis revealed that nearly half of customers do not have usage patterns strongly correlated with the heat index, and many customers use AC sparingly. A few customers have load shapes that indicate their cooling demand maximizes at 100% of their system capacity during hot weather, though not all customers reach 100% AC demand even on the hottest days.[33] Despite mild weather in 2014 and relatively few hot days, the granularity of the AMI data and the large sample sizes allowed the team to draw strong conclusions about people's reaction to weather.

Figure 40 shows the correlation between daily average temperature and electric usage for all customers in the sample, broken out by month. The green density plot in the background shows the distribution of usage within each temperature bin, while the plotted squares show the mean usage for the population at given temperatures. The percentage of customers in each kWh bin within each vertical green stripe adds to 100%. At moderate temperatures, the highest daily kWh usage bins (40 kWh/day-50 kWh/day) within the vertical lines are lighter, showing 1% or less of customers fall in each bin. By contrast, higher percentages of customers fall within each kWh bin in the range of 5 kWh per day to 15 kWh per day, shown by the darker green portion of the vertical bar. In May, the farthest left vertical green bar represents days where the average daily temperature was around 45°F; in that case, average daily electricity usage was just below 20 kWh, and the bulk of users were centered around 10 kWh/day, as indicated by the darker green.

As illustrated in Figure 40, the team found that electricity use rises as temperature increases, particularly above daily average temperatures of 60°F. Where the team observes a clear increase in electricity usage due to weather, as in July through September, the increase appears as a spreading of the electric use distribution at higher temperatures (the dark green dissipates as the distribution of usage across the population spreads out.). In other words, the higher the temperature, the greater the variation in how much energy households consume. In August, for example, the far-right vertical bars indicating temperatures above 70°F are mostly light green, showing the distribution of electricity usage is spread out relative to the cooler temperature bars to the left with dark green sections.

[33] One potential explanation for some of these observations could be that instead of central AC systems, some homes with multiple window AC units entered our analysis. While we can generally distinguish between the two based on load thresholds, this remains a possibility. In that case, it could be that the window units "max out" during hot spells rather than a central AC system.

Figure 40. AMI vs. Temperature, Broken Out by Month

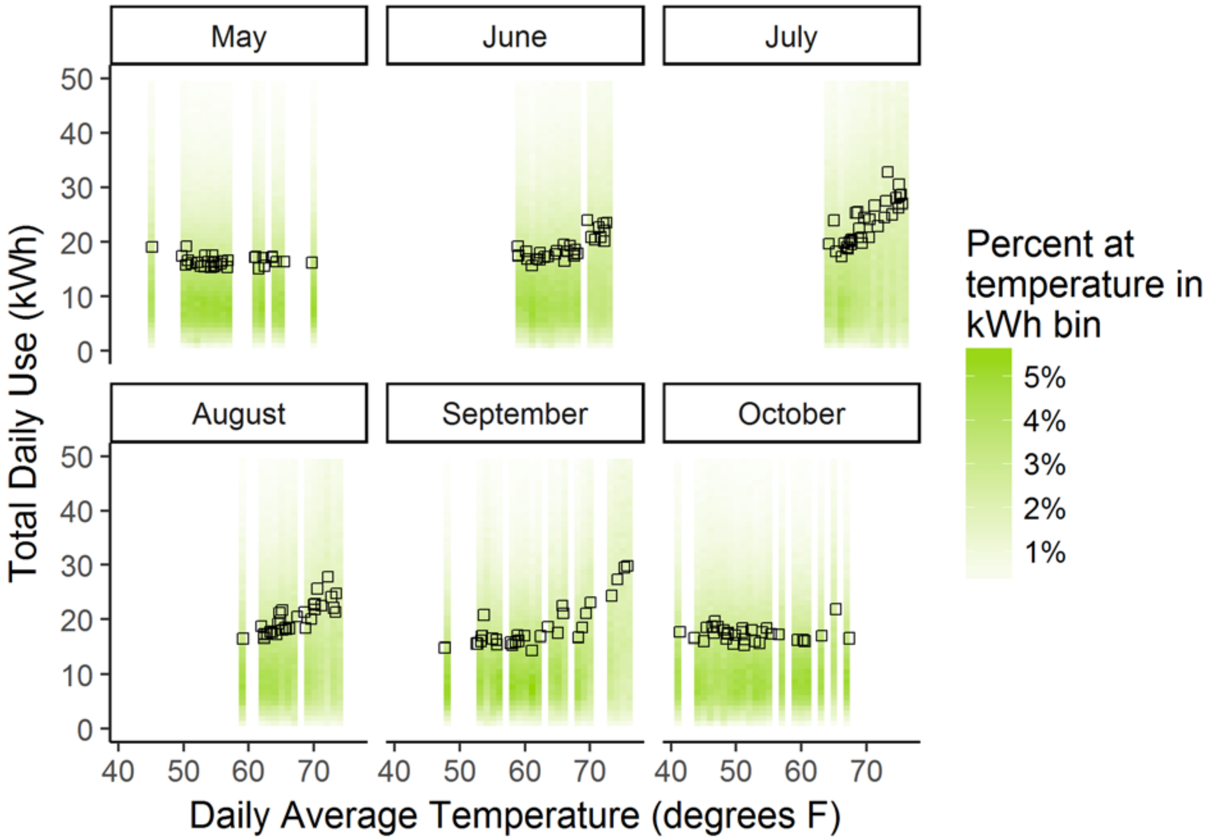

Source: Navigant

3.3.2 Regression Analysis

Building on the exploratory data analysis, the evaluation team conducted AMI-based regression analysis to quantify causal relationships between energy consumption, AC use, and weather variables throughout the summer. The complex cycling behavior of AC systems combined with large variation in the timing of cooling by customers with different AC use behavior made modeling the effects of weather on electricity use at the 15-minute level afforded by the AMI data problematic. It also introduced behavioral and system effects that could bias the results without improving the model's predictive power. For these reasons, the evaluation team used fixed-effects regression models aggregating data to the daily level. The team used only late spring, summer, and early fall data to avoid weather-dependent usage related to heating. This constraint reduced the sample size from approximately 16,000 to just over 15,000.[34] The team approached the regression analysis in two stages.

- Individual regressions to identify those with AC
    - The team ran two simple regressions on each participant with the 15-minute interval data to extract the effect of CI on electricity consumption. For robustness, the team used the results of both to determine which customers have AC systems.

[34] The team started with exactly 16,915 homes, 1,654 of which did not meet this criteria, leaving a final sample size of 15,261 homes.

- Pooled fixed-effects regression to quantify weather effects on AC use throughout the season on those with AC
    - The team subset the full dataset to only those customers evaluated in stage 1 to have AC to estimate a regression model across all AC users to quantify the relationship between their AC use and weather variables—this subset consisted of roughly 9,000 customers.
    - Cross-validation and robustness checks comparing predicted values with actual values and analyzing residuals verified the model performed well and was not overfit.

3.3.2.1 Findings

The findings from the model are intuitive, suggesting heat and humidity both play significant roles in AC use throughout the summer. Table 2 presents the regression outputs for key variables in this model, predicting energy use as a function of weather variables.

Table 2. Key AC Use Regression Output Results

Term	AME	Std. Error	t score	p.value
MaxCI	0.374	0.006	58.648	<1.00E-9
AvHBU	0.070	0.001	85.480	<1.00E-9
AvHum	0.030	0.001	53.385	<1.00E-9

Source: Navigant

Key findings and results are discussed below:

- **Customers use their AC more in response to increased heat.**
    - For every 3°F over 68°F, the model predicts roughly 1.1 kWh more usage per customer.
- **Humidity works in combination with heat to increase AC use.**
    - Above 77°F, high humidity progressively adds to the CI, adding an additional 6 degrees at 90°F and 90% humidity, translating to or 2.2 kWh additional AC usage.
- **Consistently high temperatures across a day (as opposed to a temperature spike) have more effect on total AC use than the maximum temperature reached on a day.**
    - A hot night and all-day leadup period of hot weather can add 8 kWh to a given day's usage value.
- **Beyond increasing the CI, humidity increases AC usage by itself.**
    - A difference of 50% humidity alone (e.g., the difference between 40% and 90% humidity), not in combination with heat, increases daily use by 1.5 kWh.
- **Day-of heat and humidity cause the majority of AC use, with the weather on preceding days increasing AC usage by a small amount.**
    - Lagged temperature and humidity terms contribute less than one-tenth of their day-of counterpart variables.
- **Day of week does not have a large effect on AC use relative to heat and humidity, but it does display interesting patterns, suggesting behavioral components are worthy of future research.**

  - While the effects of day of week on AC use are relatively small, at high temperatures Monday and Tuesday appear to have the highest usage, while Thursday has lower usage. At low temperatures, Saturday and Sunday have the highest usage.

The evaluation team created four different, realistic weather scenarios and inputted them into its regression model to analyze the resulting AC use behavior predictions. This should allow the PAs to draw conclusions about how weather may affect customer AC usage under different weather conditions and be helpful for planning purposes.

For each of the four weather scenarios, the team assumed a constant daily temperature and humidity profile for all days of the week. This assumption gives the team the ability to look at both the effect of varying the temperature and humidity and at the resulting predicted pattern of consumption across the week. Each weather scenario is described below:

- **Baseline Weather (black):** A CI and HBU of zero, with typical shoulder season humidity.
- **Warm Normal Weather (orange):** A CI below 10 and low HBU, with typical humidity.
- **Hot Dry Weather (red):** A CI between 15 and 20, moderate HBU, and very low humidity.
- **Hot Humid Weather (green):** Same temperature as Hot Dry Weather, but with higher humidity raising the CI to above 25 and increasing the HBU.

As shown in Figure 41, the model predicts higher AC usage on hot, humid days relative to hot, dry days. Warm days under normal humidity conditions raise AC use only marginally above the baseline. Temperature alone contributes 10 kWh, while high humidity adds an additional 8 kWh to daily electric consumption, as seen by comparing the difference between the black baseline line with the red and green lines, respectively. Not surprisingly, normal warm weather keeps AC use fairly constant throughout the week, while hot and humid days show higher AC usage toward the weekends and lower usage mid-week.

Figure 41. Four Example Predictions Using Preferred Model for Different Weather Scenarios

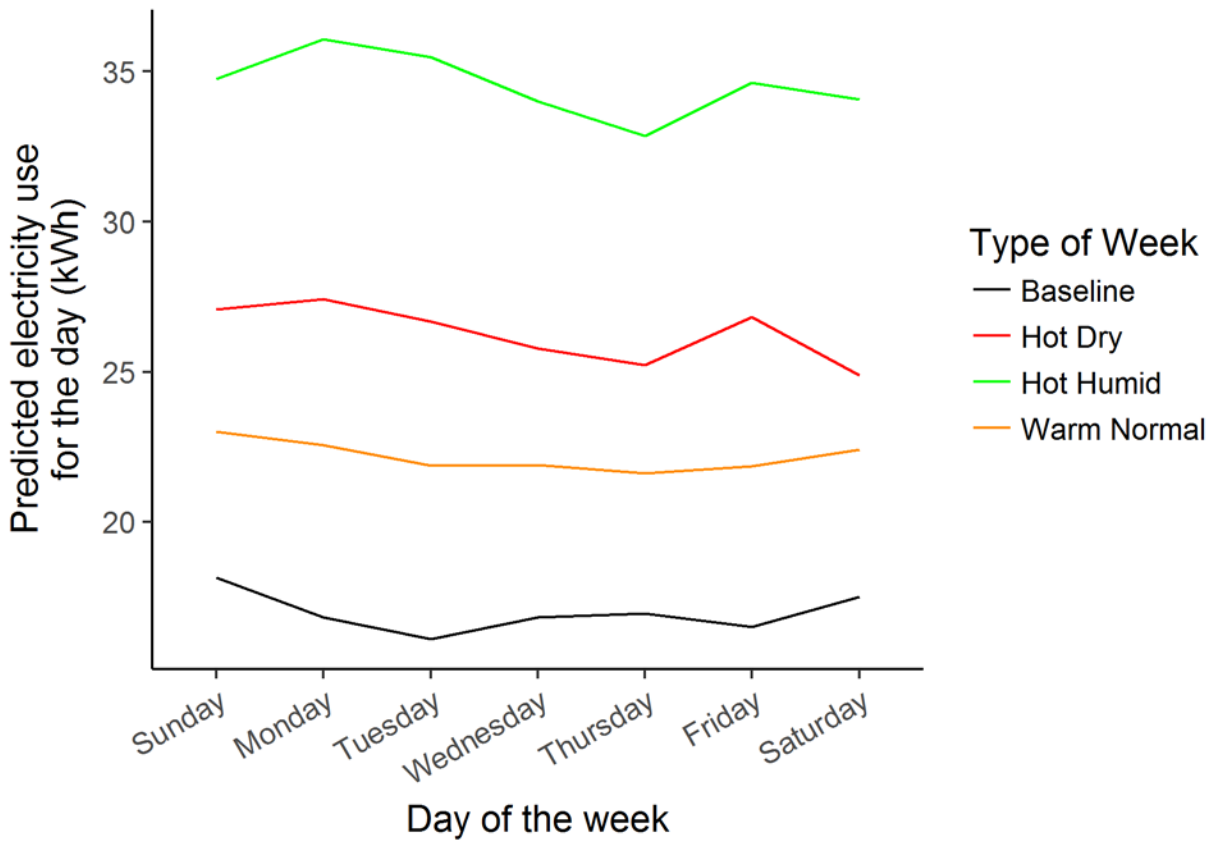

Source: Navigant

3.3.3 Considerations around Weather and AC Use throughout the Cooling Season

Strategies to reduce AC use throughout the summer should focus on both heat and humidity.

In contrast to the team's findings on first AC use, research results suggest that people respond strongly to both heat and humidity, alone and in combination, when deciding how to use their AC throughout the season. With respect to heat alone, on average, for every 3°F temperature increase above 68°F, customers use an additional 1 kWh of AC. In terms of the effect of humidity alone, doubling the humidity increases daily AC use by just over 1 kWh. Moreover, heat and humidity combined increase people's AC use even more.

Targeted messaging with tips for reducing humidity in the home, such as purchasing energy efficient dehumidifiers could help reduce AC use throughout the summer. Providing customers with messaging around strategies for being out of the home during the hottest, most humid times of the day could also reduce AC use. In the days leading up to predicted hot, humid weather, PAs can message strategically to customers who use their AC in different ways, as discussed in Section 3.1.3, to maximize energy efficiency. As an example, customers who keep their setpoint very high throughout the summer might be pushed out of their comfort zone by the added humidity and more likely to override settings to engage their AC. By providing these customers with tips on how to stay comfortable, incentives to install energy efficient dehumidifiers, options for leaving the home during the hottest, most humid part of the day, and other motivations to make other behavioral changes might cause these customers to consume less AC.

Focus less on fatigue and more on in-the-moment interventions.

The team's research findings show that while a period of prolonged heat does increase AC use the majority of use is determined by day-of temperatures and humidity. In other words, people are less sensitive to heat and humidity fatigue and reacting more to the temperature and humidity conditions on a given day when making AC use choices.

This finding suggests that providing customers with coping mechanisms for heat and humidity on a day-to-day basis is more important than creating strategies to help them deal with prolonged heat spells. Providing customers with tips about how to stay comfortable in the heat, how to reduce humidity in their homes, how to time activities to be out of the house during the worst heat and humidity during the day may be the most effective. Providing coupons for out of the home activities on hot days can help customers cope and use less energy. Messaging encouraging them to dress for the weather, eat cooler meals, and other coping strategies can all help relieve the urge to turn up the AC. The PAs may also consider the importance of having a DR program in place to further manage loads and usage spikes on hot and humid days.

Interventions related to heat in the spring and fall are unlikely to reduce AC use.

During the shoulder months of May and October, electricity use showed no increase with temperature, signaling little effect of heat on AC use. This might be related to behavioral considerations around the timing of first and last AC system turn on or insensitivity to the relatively mild temperatures. Moreover, electricity use and temperature are not well correlated at temperatures below 70°F.

Knowing this, PAs can focus on June, July, August and September for any targeted messaging and behavioral modifications around hot weather and AC use. Even though temperatures reach highs in the 70s during these months, most temperatures are low enough that customers are not using their AC much during these months. It may also be helpful for the PAs to understand that customer AC use appears to really show responsivity to temperature above 70°F, and not much below. In other words, targeted messaging and behavioral interventions related to AC use and temperature are not likely to be effective outside of June through September and on days cooler than 70°F.

3.4 Weather and DR Program Participation

Phase II DR program participation research explores how weather and program design factors influence DR program participant behavior, particularly their propensity to opt out of events. Influential weather factors include heat and humidity, while relevant program design factors include program implementation details such as thermostat type, event duration, notification methods, consecutive events, day of the week on which events occur, and frequency of events. The evaluation team paired in-depth interviews with DR program participants and implementers with data analysis including regression modeling to understand these relationships from both qualitative and quantitative perspectives.

3.4.1 DR In-Depth Interviews

The evaluation team conducted in-depth phone interviews with DR service providers (DRSPs) and DR program participants to understand how weather influences people's participation in DR programs, specifically their choices to opt out of or override settings. The team conducted in-depth interviews with three DRSPs and 20 program participants in National Grid's Thermostat DR project and SES pilot. These interview results were analyzed and synthesized to help explain the team's quantitative findings.

3.4.2 Program Overview

The following sections summarize key characteristics of the Thermostat DR project and SES pilot as context for the interview results that follow.

3.4.2.1 Residential Wi-Fi Thermostat DR Demonstration Project Overview

The Thermostat DR project is an opt-in demonstration project that includes two offerings: CS and RHR. Program enrollment can occur through a bring your own thermostat (BYOT)[35] approach or as part of the home energy services (HES)[36] program through which a smart thermostat is installed as the result of an audit process.

Table 3. Thermostat DR Program Offerings

Category	CS	RHR
State	Massachusetts and Rhode Island	Massachusetts
Type of Program	Opt-in	Opt-in
Types of Thermostats	ecobee, Honeywell	Nest
Total Program Duration	108 hours	52 hours

[35] BYOT programs allow customers the flexibility to purchase their own smart thermostat, which can include a variety of brands, and then they enroll in their utility's DR program.

[36] The HES program offers customers an opportunity to receive a Home Energy Assessment (HEA) conducted by a National Grid Energy Specialist. The Energy Specialist produces a custom home energy report outlining recommended energy efficiency improvements. The specialist will also install, at no cost, a number of energy-saving products, including a discounted Wi-Fi thermostat.

Category	CS	RHR
Event Duration	3 hours [37]	3-4 hours [38]
Advance Notification	Day of, >2 hours (customer notified)	Day of, >2 hours (customer notified)
DR Event Opt-Out Option (Before Event, During Event)	No, Yes (ecobee) Yes, Yes (Honeywell)	Yes, Yes
Intended DR Setpoint Range[39]	+/- 3°F [40]	+/- 3°F
Participant Incentive	BYOT: $25 for signup HES: Free thermostat and installation Both: $25 per year if complete 75% of events	$40 for signup No event requirement or incentive
Program Opt-Out Rates	ecobee: <15% Honeywell: few opt outs; 20% connectivity issues[41]	>20% of users opted out of events[42]

Source: Adapted from 2016 Residential Wi-Fi Thermostat DR Final Report

The *2016 Residential Wi-Fi Thermostat DR Final Report* summarizes program process and impact findings during program year 2016 and analyzes trends related to event opt outs. The report notes that most opt outs came from a small number of thermostats. Notably, the researchers did not find evidence of program fatigue as opt outs did not demonstrate a clear pattern of increasing over time. The researchers also found no correlation between opt-out rates and event frequency or day of week. The researchers did find a positive correlation between opt-out rates and temperature for Honeywell thermostats and opt-out rates and event duration for Honeywell and ecobee thermostats.[43] As these results are correlation-based and not causal, regression analysis in this should provide more definitive causal results.

---

[37] The event duration for CS changed from 2-4 hours in 2016 to 3 hours in 2017.

[38] The event duration for RHR changed from 4 hours in 2016 to 3-4 hours in 2017.

[39] The allowed temperature adjustment during a given event. The temperature setting will return to its prior temperature following an event.

[40] The intended DR setpoint range for CS changed from 2°F in 2016 to 3°F in 2017.

[41] According to the *2016 Residential Wi-Fi Thermostat DR Final Report,* reasons for connectivity issues varied: 1) There were 41 Honeywell thermostats (about 11%) without connectivity for all events. 2) Honeywell had server issues at the start of an event on July 12. 3) There were 19 ecobee thermostats (about 7% of participating thermostats without connectivity for all event days. 4) ecobee had notable connectivity issues on July 15 (reason not noted).

[42] The *2016 Residential Wi-Fi Thermostat DR Final Report* identifies the follow possible reasons for the higher opt-out rate for RHR: 1) Nest thermostats had a setback of 3° during events compared to 2° for the other thermostats. 2) RHR participants received the full incentive with no participation requirements. CS required participation in at least 75% of events to receive the incentive. 3) RHR events were hotter, on average, than CS events. 4) There may be differences in device usability, the types of users, or the method used to opt out.

[43] While these findings are informative, they are all correlation-based and not based on regression analysis; therefore, they cannot be interpreted as causal but merely directional. Moreover, they are based on a single year of observation.

3.4.2.2 Smart Energy Solutions Program Overview

SES is an opt-out dynamic pricing pilot[44] that offers customers technology packages, two of which include a Wi-Fi thermostat. Customers with a Wi-Fi thermostat are included in peak event days in which the program adjusts the thermostat temperature setpoints to help with load curtailment during peak periods. Customers are able choose between two dynamic pricing options: Smart Rewards Pricing and Conservation Day Rebate. The Smart Rewards Pricing tiered rate structure combines time of use (TOU) and critical peak pricing (CPP), where customers pay higher rates during peak periods and peak events and pay lower prices during off-peak periods. The Conservation Day Rebate uses a peak time rebate (PTR), where customers are provided a rebate for reducing their energy consumption during peak events.

Table 4. SES Pilot Offering, 2015-2016

Category	SES
State	Massachusetts
Type of Program	Opt-out dynamic pricing with programmable communicating thermostat option. Customers choose between a TOU CPP rate and a PTR rate along with one of four technology packages.
Total Program Duration	20 peak events in both 2015 and 2016
Event Duration	Peak events lasted 4-8 hours
Advance Notification	Day before event
Tstat Setback	2015: 3° or 4°; 2016: 2° or 3°
Customers with a Wi-Fi Thermostat	Approx. 280

Source: National Grid Smart Energy Solutions Pilot: Final Evaluation Report

The *2016 National Grid Smart Energy Solutions Pilot: Final Evaluation Report* summarizes impact and process considerations for the program during the 2015-2016 Pilot. The report notes that in 2015 customers were more likely to override the peak event thermostat settings as the summer progressed, which differed from 2016 where event participation was fairly consistent throughout the summer. The team found that 2015 had more events in September than 2016, possible reasons for this difference may be that the weather was hotter and more humid later in the summer and customers were reacting to those weather conditions; it is also possible that customers grew tired of the hotter, more humid weather as the summer progressed. The report found that the rate of overrides and opt outs was most strongly correlated with the length of the peak event and were not correlated with the size of the degree setback.[45]

3.4.3 Methodology

The evaluation team interviewed DRSPs and program participants to gain insights from both the perspective of the program implementer and customers enrolled in the program.

---

[44] Dynamic pricing programs offer alternative rate structures that aim to make the cost of electricity reflect the cost to produce it, which can vary based on the time of day due to changes in electric demand.

[45] While these findings are informative, they are all correlation-based and not based on regression analysis; therefore, they cannot be interpreted as causa, but merely directional. Moreover, they are based on a single year of observation.

The team conducted three in-depth interviews with DRSPs to understand program design and implementation, and to discuss what factors, from their perspective, are most likely to influence DR event participation.[46] The team recruited DRSP interview participants from willing industry contacts, and all resulting interviews were transcribed for analysis.

During the interviews, respondents were prompted to comment on the role of weather (both discrete weather events and trends over time) and event characteristics on participant opt-out behavior. Interviewers also discussed how DSRP smart controls change settings to understand whether these controls enhance or work counter to DR objectives.

In addition to DRSPs, the evaluation team conducted 20 in-depth interviews with Thermostat DR project and SES pilot participants. The team interviewed 12 who participated in the Thermostat DR project and eight who participated in the SES pilot. The main goal of these interviews was to understand factors that contributed to customer decisions to opt into the program initially and factors that contributed to the decision to opt out of an event or override preset controls, with a specific focus on the influence of weather and event characteristics.

The samples for each of the programs were created using self-reported opt-out and override behavior from post-event and end-of-season surveys. This information was used to identify a set number of participants to target based on their opt-out behavior to ensure the team spoke with both people who never opted out or overrode settings and people who did. To better understand participant motivations to opt out or override thermostat setpoints, the majority of interviews targeted customers that had opted out or overrode thermostat setpoints during a peak event, while fewer interviews targeted customers that did not. During recruitment, prospective interviewees were asked to verify their opt-out and override behavior, improving the team's confidence in the results. All resulting interviews were transcribed for analysis.

The interview guides were informed by findings from the *2016 National Grid Smart Energy Solutions Pilot: Final Evaluation Report* and the *2016 Residential Wi-Fi Thermostat DR Final Report*.

To reduce the burden on customers, the team did not recruit customers who participated in other focus groups or in-depth interviews related to the programs. Interview participants received a $50 gift card incentive for their time. Table 5 summarizes the final completed interviews by program and opt-out status.

Table 5. Participation Status Summary for Completed Interviews

DR Program	Had Never Opted Out	Had Ever Opted Out
SES	2	6
RHR	2	4
CS	2	4
Total	**6**	**14**

Source: Navigant

[46] Note that throughout this section "DR event participation" or "DR participation" refers to participation either in SES or the Thermostat DR program interchangeably. While the team may refer to National Grid's SES pilot and its Smart Thermostat DR program separately, when the team refers to DR event participation, it refers more generally to both or either.

3.4.4 Interview Findings

Interviews with DRSPs and customers suggest four broad categories of DR participants based on their behavior during events, which the team refers to as:

- Comfort Seekers
- Hold Outs
- Best-Effort Participants
- Reliable Participants

Before diving into the different participant types identified through the research, the team explains some of the key characteristics and behaviors noticed in a wide range of participants—both by participants themselves and the DRSPs that provide DR programming for them. These more general findings helped to the team define and explore the more detailed groupings that follow.

Customer are generally willing to participate in events to support their community or help their utility, but their decision whether to participate or opt out is ultimately motivated by their life circumstances and comfort considerations during an event. The team found that weather, especially high temperatures and high humidity, create conditions that motivate customers to opt out of events due to discomfort. DSRP interviewees also noted that the length of time that a customer has participated in the program affects when and how often they opt out of an event. Customers that have been active in the program for longer are more tolerant and are more likely to opt out toward the middle of the event. On the other hand, new participants are not as familiar with how the event will affect their home and are more likely to opt out toward the beginning of the event.

> *"Those customers who are experiencing an event for the first or second time, for them it becomes a challenge. We find that the high correlation in opt-outs versus average experience in being active in the program. I think if you look at a more stable customer base who's been with the program for a while, for them it will be more towards the middle when it is really hot. When there have been consecutive events is when they will opt-out." – DRSP*

In nearly all cases, participants were indifferent toward events if they were not home during them, regardless of weather conditions or event parameters. Most respondents indicated they had no issue letting an event run its course when they were away from home. When participants are home throughout an event or they come home while an event is still occurring, they are more likely to opt out, especially in more extreme weather. While some participants monitor their home temperature remotely to ensure the pets and household items are not negatively affected, they generally will not opt out if they are away from home.

> *"That's the nice thing about Mass Save that if it is a conservation day, they control it. Usually by the time I come home, the conservation is done. If you're home, you want to be comfortable. You don't want to be sweating in your house." – SES Participant*

> *"Give them coupons for ice creams to encourage people to leave the house and that they feel that it's worth it to participate and that they have a strategy to cope with their home getting really hot." – DRSP*

Program participants across all three programs indicated they can predict when peak events will occur based on the weather forecast, but they cannot predict their opt-out behavior based on the weather forecast. Program participants are willing to participate in an event regardless of the weather as long as they are comfortable in their home. Since program participants are willing, and in some cases eager, to help the utility by participating in peak events, the utilities may be able to lengthen participation before opt outs or prevent opt outs by equipping participants with tools and resources that help them mitigate and adapt to the effect of events on their comfort.

3.4.4.1 DR Event Participant Types

The following section describes the main categories of DR program participant identified through the in-depth interviews with DRSPs and program participants.

1. The Comfort Seekers

> *"I don't base everything on the events, you know what I mean? If I'm uncomfortable, certainly my comfort is going to come before one of the events." – CS Participant*

Comfort Seekers are program participants that heavily prioritize their comfort over their participation in an event. Comfort Seekers will not go out of their way to make changes to their day to accommodate an event; they will opt out as soon as they come home or feel uncomfortable. These customers are accustomed to manually managing their thermostat when they get "a little too hot" and will not wait for their scheduled settings to adjust the temperature automatically.

> *"We have people who just really love overrides. They're used to managing their thermostat. They're just used to managing their thermostat by overriding all the time. So, we have some people who are just super used to always touching their thermostat." – DRSP*

> *"It doesn't happen very often, but really on one of those hot, humid days that the conservation went down to ... was from 12pm to 7pm. You get home at 5:00 from working, the last thing you want to do is to come into a house that's steaming. Yeah. That's when I opted out." – SES Participant*

> *"I've turned it off, especially on days where we're not in the home for a few days at a time and if I know that there's no extreme temperatures, then I turn it off so then it doesn't really regulate it." – CS Participant*

The DRSPs reported that they do not gather systematic feedback from customers on why they chose to opt out of an event, but DSRPs suspect that a connection exists between households opting out and people coming home. This trend is consistent with comments from participants who indicated that they prioritized their comfort when they come home over continuing to participate in the event. Comfort Seekers may come home from work or an outdoor activity during an event, and their immediate instinct is to adjust their thermostat to get comfortable instead of waiting 15 or 20 minutes for their body temperature to regulate. Programs could offer participants an option to snooze an event when they need some relief and then automatically readjust their thermostat setpoint after a set period of time to curtail complete opt outs in these cases.

"It seems to trend with people that are coming home…there's been analysis that's been done to indicate that, okay, occupancy is triggered within the home, then we're also seeing a little bit of that drop out. So, it could be that people are coming home and then they notice it…There's always gonna be a portion of the population who are 'oh, okay, I'm home and I don't want this happening today' so they'll be people who kinda do it the first hour they start to notice"-- DRSP

2. The Hold Outs

A subset of the Comfort Seekers, as identified by the DRSPs, are the Hold Outs. These participants set an indefinite hold on their thermostat, which automatically opts them out of an event, even if it was not directly in response to the event. Because their normal thermostat use behavior is to have the system in hold mode a large portion of the time, by coincidence, when an event occurs, they are usually in hold mode and the de facto result is they do not participate in the event. Because this behavior is not in response to events yet precludes participation, this group of customers is not well suited for participation in a DR program and is not likely to adjust their behavior to benefit events.

"they would set their thermostat at one setpoint forever…the hold might have been set two, three, four days before, a week before. So, it wasn't necessarily a hold in response to an event, it was just they had already had it on there, and it just stayed. So, we have these groups, the people who just have a hold behavior that is not conducive to the DR event." - DRSP

3. The Best-Effort Participants

Best-Effort Participants are program participants that will try to push through an event by leaving their home, opening windows, or grilling instead of cooking inside; however, they may opt out if they exhaust their coping options. These participants do their best to avoid opting out of an event, but if the temperature exceeds their comfort threshold and they do not have the option of leaving the home, they will opt out of the event.

"Sometimes it was because I needed the house cooler and other times I had already been in it and it just got too hot. It was unbearable, so I opted out in the middle of it and got the air back on… most of the times I tried to go with it." – SES Participant

DRSPs noted that a participant's perceived ease of opting out of an event affects their likelihood of opting out. Participants that need to take additional steps such as make a phone call to opt out rather than by simply adjusting the temperature on their thermostat are more likely to remain in the event.

"I have opted out, especially in the situations where I'm home and I don't have plans to go anywhere. Other situations where I know a Peak Event's coming and I know maybe I can leave the house and not utilize the air conditioner that much or the pool, for example. Maybe go and take the kids out somewhere else and just be out of the house during Peak Events." – CS Participant

"Like I said, if I knew it was going to be on a day that I would be around the house, just before that, I would lower the temperature in the house. Then I'd shut the air conditioner off and keep all the doors and windows closed…I'm in and out. I'm retired, so I'm either going here or going there. I don't have to be home, but in the summertime, I'm probably out in the yard if I'm home instead of in the house." – SES Participant

In response to this observation, some DRSPs created speed bumps that customers must overcome before opting out of an event. These small hurdles create the perception that it is more difficult to opt out and lower opt-out rates. However, programs should be strategic when using speed bumps because they may increase participant frustration. As an alternative, programs may choose to integrate gamification or other positive reward strategies like personalized goal setting or interactive data presentations to motivate customers to fully participate in each event.

4. The Reliable Participants

The final group of DR participants identified are the Reliable Participants. These are program participants that will remain in an event regardless of personal comfort impacts or convenience considerations. Often these are customers with flexible schedules, who may be retired, and are easily able to adjust their plans to accommodate events. They may notice that their home is slightly warmer but will not make any adjustments once they realize that an event is occurring.

"Even if we give the customer the option to opt-out where they can go online or call our call center, and we can opt them right then and there for that event – in those cases it's very low. Where the opt-out option is made conveniently available on the thermostat or on their app, those opt-outs can be extremely high. Some of those can be done unwittingly because they may not know that they have opted themselves out of that event." – DRSP

"No every time the event happened, if it was getting too hot and I would like almost turn it on, it would basically be the end of the rush hour and then the AC would then just kick on right then." – RHR Participant

Because these DR participants are dedicated to participation and do not consider the program an inconvenience, the program provider does not need to strategize to reduce their opt-out rate. However, the program implementer may use these potential DR participants to their advantage by undertaking research to better understand and identify this group of potential participants to recruit to DR programs to maximize savings.

3.4.4.2 The Role of Weather and Event Timing

Participant opt-out decisions are made based on comfort in the moment.

This section explores in-depth interview findings relating weather and event timing to DR participant opt-out behavior.

The evaluation team's research findings suggest DR a program participant's decision to opt out of an event is typically based more on comfort and their circumstances during an event than directly on the weather. Participants reported that they could usually anticipate when a peak event would be called based on the weather forecast, but that they typically would not decide to opt out until their comfort threshold was exceeded during the event.

"It's occasionally in the summer months we do get high humidity to where one cannot be outside, so those are the days, definitely, I opt out because we have to have the air conditioning on. But I think those are the days National Grid has chosen to do the peak events." – SES Participant

When asked if their opt-out behavior changed by the end of the summer relative to the beginning or middle, participants did not indicate a strong difference in their behavior during different time periods. Participants often mentioned that July and August typically have more of the hot and humid days that might drive them to opt out of an event. However, they caveated this statement by admitting that those types of days are sporadic, can happen at any time during the summer, and that their behavior would be the same regardless of when the hot and humid day occurred.

"Well, if they're saying it's going to be in the 90s and humid today, almost always, it's a peak event day. Any day that's high temperatures with high humidity, it ends up being a peak event." – CS Participant

When asked if the weather leading up to an event influenced their opt-out behavior, participants commonly indicated it did not impact their event participation. They suggested their opt-out decision was more closely related to the weather on the event day. Responses suggested that the weather leading up to an event would only be a factor if multiple events were called in a row.

Participants connected their discomfort during events to high humidity.

"I think, well, probably heat and humidity. They come all together. I notice through if it's hot and not humid, it doesn't trigger a Peak Event. Or, if it does…they are shortened, maybe a 4-hour [even] as opposed to a 6- or 8-hour [event]. So, I think from looking at the pattern it definitely seems to happen more when it's hot and humid." - CS Participant

"In the house, uncomfortable would be 80 or higher. For the house to get to 80, that typically means that outside temperature would have to be almost 90 with high humidity. So, if it's very humid, then that normally is a deal-breaker for us." - SES Participant

"it was just on the days when it was particularly hot or humid. We get a lot of humidity here. If it would be particularly hot and humid, I would end the Rush Hour thing and cool myself off a little bit." - RHR Participant

Participants pointed to humidity as the weather condition that affected their comfort the most, especially when coupled with high temperatures. Overall, participants noted that days with high humidity and temperatures in the 90s were likely going to trigger peak events. Participants suggested they had methods to mitigate high temperatures, such as turning on ceiling fans or opening windows, but they did not have solutions for coping with humidity apart from adjusting their thermostat.

This is consistent with responses from the DRSP interviewees who noted that days that are both hot and humid are most likely to affect customer participation in events. Their responses confirmed the suggestion by participants that the key factor is humidity combined with heat rather than either in isolation.

> *"The humidity within the home, that's driving it for customers. Far less the actual temperature creep that's happening during demand response events, but where the humidity is increasing rapidly within the space…It can be a really hot day, but if the humidity is really low then maybe customers are more likely to continue participating, versus finding it uncomfortable if it's even a slightly milder day but not humid." – DRSP*

Back-to-back events can lead to program fatigue and frustration.

> *"The peak events happened like back-to-back-to-back. And they keep going later, and later into the evening where you have to monitor areas, and not have anything on until like eight o'clock at night."*
> *– SES Participant*

Most participants expressed that program fatigue from back-to-back events could push or had pushed them to opt out because they grew tired of being uncomfortable in their home and/or adjusting their lives to accommodate events. This insight is consistent with responses from the DRSP interviewees, which also indicated that fatigue from back-to-back events negatively affected participation.

> *"Yes, that happens. It's usually, you know New England weather, we might have a week long where it's hot and humid, 90 degrees, maybe for four days in a row, or something like that and they do seem to come back to back to back at that point. That has a lot to do with maybe opting out, because there's only so many things you can do outside the house."*
> *– CS Participant*

> *"For a second, you're kind of like ugh. And then you're looking like, 'Okay, I got it. It's rush hour. That's fine.' But like I said, if it happens a lot, I think it would become much more frustrating because you'd be like, 'I'm not comfortable. I'm not comfortable every day in my house.' So, at that point, you're going to be, like, 'No. I don't want do that.'"*
> *– RHR Participant*

This finding should be considered by utilities when making the decision whether to schedule back-to-back events. The program implementer or the utility faces a tradeoff between increased customer frustration and opt outs relative to the incremental demand savings achieved through the consecutive events.

> *"Yeah, actually humidity plays a bigger role sometimes. If it is really humid, it creates problems for customers, so I would say that if the temperature is high that's one option, but if it is a high temperature and high humidity, that creates a bigger problem for our customers." – DRSP*

> *"Hot and humid days are probably the worst; an ideal day would be it's hot…I would think that humidity causes quite a bit of problems for a lot of customers…we have not had any major problems with mass opt outs because of temperature going high. It's frequency of events that sometimes that gets us in trouble." – DRSP*

3.4.4.3 Commonalities Between Programs

The team's research revealed that some customer experiences and behaviors are specific to the program, how the program is delivered, the technology used, messaging, and other factors. However, some findings were revealed to be universal across program types. This section explores common findings between programs, while Section 3.4.4.4 examines differences.

Key participant motivations are common between program types.

The DRSPs interviewed noted that the two key drivers motivating customers to enroll in a DR program are: 1) to get a free or reduced-price smart thermostat, and 2) to save money or receive an enrollment incentive. These drivers were consistent with the responses provided by the participants interviewed.

During interviews, most SES participants indicated that receiving a smart thermostat was their primary motivation for enrolling in the program. They cited saving money and being able to track their energy usage as secondary motivations. RHR and CS participants responded similarly. Most indicated they used the utility rebate to obtain the smart thermostat technology and then enrolled in the DR program as an additional way to save money or to help the utility.

> *"The offer is compelling. They get a thermostat, which provides them with mobility. They can change their temperature from wherever they are." – DRSP*

> *"I would say for the most part it's people wanting the technology, and this was an easy way to get it...And then it's just a matter of really going to a website, and then going through the enrollment flow that way." – DRSP*

The team interviewed several participants in different programs who indicated similar confusion about program details (e.g., incentive structure, how much their thermostat was adjusted, what it meant to opt out). Continued engagement with the customer during the event season or improved education about the program during enrollment could increase event participation by increasing customer awareness, making them feel more connected to and invested in the program and decreasing the likelihood of opt out.

Customizable, remote temperature control is perceived as a key benefit.

Participants across all programs commented that the ability to remotely monitor and control the temperature of their home was the primary benefit of their thermostat. Program participants discussed setting up temperature schedules on their smart thermostats to reflect when different household members

> *"So, what I've done in the past for some of the peak events, where I know it could be a tad bit uncomfortable, is I would just set the temperature on the thermostat a little bit below normal, so that it averages out and still stay in a comfortable range. If it's multiple days in a row where one of those days is gonna be worse than the other, then you try to make sure that the thermostat doesn't go too far outside of the comfort zones." – SES*

> *"I'll also use the vacation setting where it kind of maintains a temperature throughout the 24-hour period that's not detrimental to your things that are in the house or cats or other things." – RHR Participant*

were home. Many participants reported scheduling their thermostat setpoint to either increase or decrease during the night to increase comfort or use less energy.

"I have different settings for the winter versus the summer…the thermostat will be set at a higher temperature during the day if no one is present and then adjust so that it's cool and comfortable by the time someone is getting home…In the summer, during the day, it's set at 80, 85 if no one is present and then if someone is present at home, it's reduced to 72, 73. Unless it's a peak event, then I would adjust accordingly." – SES Participant

Some participants mentioned that they would typically override the set schedule if they were coming home early so that the home would be cool when they arrived. Program participants also mentioned turning off their AC completely when they are on vacation or when the weather was relatively mild. All of these observations speak to the importance customers place on being able to customize temperature settings in their home through their thermostat.

Customers enroll to save money and want to see evidence of their savings.

Overall, program participants expressed skepticism about whether the DR programs they participated in led to savings. SES participants expressed interest in seeing what they would have paid had they not been enrolled in the program. Participants in RHR and CS were less concerned with whether their participation resulted in bill savings; rather, they expressed disbelief that any significant savings were being generated by the program as a whole. This belief that the program was not generating savings made it easier for participants to justify opting out of events when they were uncomfortable. Despite these doubts, customers expressed commitment to their respective programs, often expressing a sense of duty to help the utility or contribute to the greater good.

"I do like the peak/off peak rates advantage for the smart rewards but again I don't see ... I think I'm saving money but I don't see that transparency. Like, okay, this is what it would have cost you if you were not part of the program and this is what it's costing you today." – SES Participant

Participants that have a sense of duty or loyalty to their utility are motivated to participate in DR events. However, these same customers do not always understand why the program is necessary or how to cope with its negative impacts on comfort and convenience. Providing customers with resources and strategies to mitigate and adapt to the effects of DR events would help these participants balance their desire to be comfortable, while also supporting their interest in helping their utility. For instance, the utility or program provider could help customers develop DR event participation-conducive behaviors such as opening windows, using fans, leaving the house, and pre-cooling as alternatives to opting out.

3.4.4.4 Differences Between Programs

While the previous section showed considerable overlap in the program experiences of customers between different DR programs, the team's interview findings revealed several key differences between programs, which are explored in this section.

SES participants were often unclear how they became enrolled.

"I mean honestly, like I said, I don't even know how I got involved in the program, but somehow I did. I just, I kind of go along with it…Because they were so quiet about keeping that quiet that you could get out of it. We had no idea that we could get out of it." – SES Participant

SES participant interview responses suggested they were less familiar with the program's details, especially with respect to how they became enrolled in the program. By contrast, the RHR and CS participants typically had clear memories of enrolling in the program. In some cases, this was a point of frustration for SES participants because they knew neighbors that were not in the program but did not understand how they became involved or how to unenroll.

SES participants adjust their daily routines to accommodate events.

A major pain point communicated by SES participants was event duration, especially when events persisted later into the evening. Many SES participants commented that events that continued late into the evening made it difficult for them to balance reducing their energy use and the needs of their family. Many SES participants said they would prepare for peak events by cooking food ahead of time or not doing dishes. However, they experienced frustration when events lasted until 8 p.m. and had a difficult time cooling their home down to a comfortable temperature before going to bed.

"Well, if it's gonna be an uncomfortable day, where it's gonna be 90 with super high humidity, that will require a little bit of prep for the day or that just might require a day out of the house. So, you're just kind of looking at the weather forecast and living in New England for a little while, understanding what that means, and just planning accordingly." – SES Participant

"Well, we find it really hard to sleep in humidity and heat. You know, when ... if we go to bed and the house is still really hot and humid. That's really hard because if we only had since 9 P.M., 8 P.M. to cool it down, we probably can't get it cool enough for comfort to sleep. So, that's really our main level of like just comfort, is sleep." – SES Participant

"Yeah, I pre-cool the house in the summer so that prior to a peak event we reduce the uses of electricity for, in terms of dish washing and laundry and ironing and things like that. Occupants in the house are notified that ... So my family will be told so that they are aware that there's a peak event. Generally, lighting is not a factor since it's such low usage, it's more about what else are we doing. If needed, for cooking, we can tend to grill those days, things like that." – SES Participant

Most SES participants mentioned they would pre-cool their home, cook food ahead of time or cook outside on a grill, and avoided running dishwashers, clothes washers, or dryers until after the event. These pre- and during-event actions are likely specific to SES participants due to the peak pricing aspect of the program, which encourages participants to minimize overall energy consumption during events. By contrast, CS and RHR participants do not have that same financial motivation to reduce consumption. As a result, few RHR or CS participants reported taking any specific actions before or during an event.

3.4.4.5 Additional Key DRSPs Interview Findings

The following section outlines additional key findings from the three DRSP interviews, providing further insight into enrollment success factors and optimization strategies provided by DRSPs to help utilities maximize savings while minimizing customer discomfort or frustration.

Trust in the utility and DRSP plays a key role in enrollment.

The DRSPs provided key insights into success factors for customer enrollment and how customer enrollment is tied to utility goals and objectives. The DRSPs the team interviewed primarily discussed the opt-in program design within which customers have the option to opt out of events and how within this structure, event characteristics and incentives are customized to meet the utility's goals and needs. These DRSPs indicated that they typically work with utilities to promote the program by notifying customers that their utility offers a DR program at the time they install a smart thermostat in the home. In addition, many CS and RHR participants the team interviewed mentioned that after they installed their smart thermostat, they received a notification that their utility offered the DR program, which then motivated them to enroll. In both cases, trust of the utility and the DRSP appear to be a key consideration leading to increased enrollment.

> *"We usually see them get numbers (enrollment) from [mentioning DR programs during the thermostat install process]. I think because people both trust their utility and they trust [the DRSP]." – DRSP*

Optimization strategies offer a solution to meet the needs of both the customer and the utility.

The DRSPs discussed the importance of striking a balance between customer satisfaction and the utility's load shifting needs. DRSPs often offer optimization of temperature settings at the individual household level, referred to as adaptive algorithms, as a means of helping utilities attain this balance.

> *"Our system takes in outdoor temperature, humidity and wind speed. We'll provide 5-minute run time data for each thermostat serial number. Inside that it will tell you whether they participated in the demand response event in that 5-minute window, what the indoor temperature was and then the outdoor temperature and wind speed and all of that. So that even evaluators like yourself can start to correlate what's happening outside versus what's happening inside and the impact, especially where you have utilities that have vast service territory. You can get a really good sense of what's happening at different locations." – DRSP*

DRSP optimization strategies often use basic household information (e.g., year built, square footage) provided by the customer during the enrollment process and smart thermostat data collected over time to understand the home's thermodynamics and build home-specific models to inform these adaptive algorithms. In some cases, weather data is coupled with the thermal model of the home so that the DRSP can simulate the effects of adjusting the

> *"We call it adaptive algorithms, so I think we use those kinds of techniques, precooling, as well as adaptive algorithm that tailors to an individual's needs." – DRSP*

AC for different lengths of time throughout the day, while keeping the home's temperature track within the customer's comfort band.

3.4.5 DR Program Data Exploration

The quantitative data analysis portion of the DR participation evaluation work included both an exploratory data analysis and regression analysis. The exploratory data analysis complemented and informed the more focused regression analysis discussed in the next section. The initial data visualization step gave the team an understanding of the trends between each potential regression attribute and the outcome of interest, DR event opt outs. Additionally, because of the time invested in data exploration, the team was more confident that the available data provided enough variation in weather and event design parameters to identify statistically significant influences on opt-out rates, despite relatively mild weather.

3.4.5.1 Methodology

Using participation data from the CS (Honeywell and ecobee) and RHR (Nest) components of the Thermostat DR project, the evaluation team conducted initial data exploration to better understand the influences of weather and event parameters on a customer's propensity to opt out of DR events. The thermostat type most represented was Nest, followed by Honeywell and ecobee. Table 6 summarizes the counts of devices represented after removing records that were unenrolled or experiencing connectivity issues.

Table 6. Counts of Enrolled, Connected Devices

Thermostat Type	2016	2017
Nest	1,331	2,804
Honeywell	300	1,007
ecobee	52	201
Total	**1,683**	**4,012**

Source: Navigant

The participation data included the summers of 2016 and 2017 and was composed of variables describing the event (i.e., event duration, event start time, length of time between events) and customer thermostat use (i.e., cooling setpoint, cooling status, indoor temperature at quarter-hourly intervals). The thermostat data was used to define customer participation status throughout the event. Participation data was available for event days and the three days preceding each event day. The participation data was combined with quarter-hourly, ZIP-code level weather data, including temperature, relative humidity, HBU, and temperature-humidity index. Weather data was available for the entire summer, not just event days.

Visual inspection of the correlation of key variables such as opt-out rates and certain independent variables (i.e., weather, event attributes, historical AC use) through scatterplots and other graphics played a key role in the exploratory data analysis.

Data exploration revealed that a customer's decision to opt out is primarily driven by factors that lead to discomfort in the home. Increased heat and humidity during DR events are often correlated with higher opt-out rates. This finding is aligned with what the team learned from participants in the in-depth interviews. Participants cite the heat, and even more so the humidity, as reasons to override an event. Humidity was even considered a deal breaker by some participants.

In Figure 42, the team observes more frequent events overall in 2016 compared with 2017, and as a result, more opt outs. In 2017, the second event in a pair of back-to-back events often had higher opt-out rates, which could be evidence of fatigue due to consecutive events. Both summers experienced similar daily high temperature ranges with few very hot (temperature greater than 100°F) days. Still, events with elevated opt-out rates corresponded to event days with high temperature and humidity.

Figure 42. Opt-Out Rates by Event with Weather Overlay for 2016 and 2017 Summers[47]

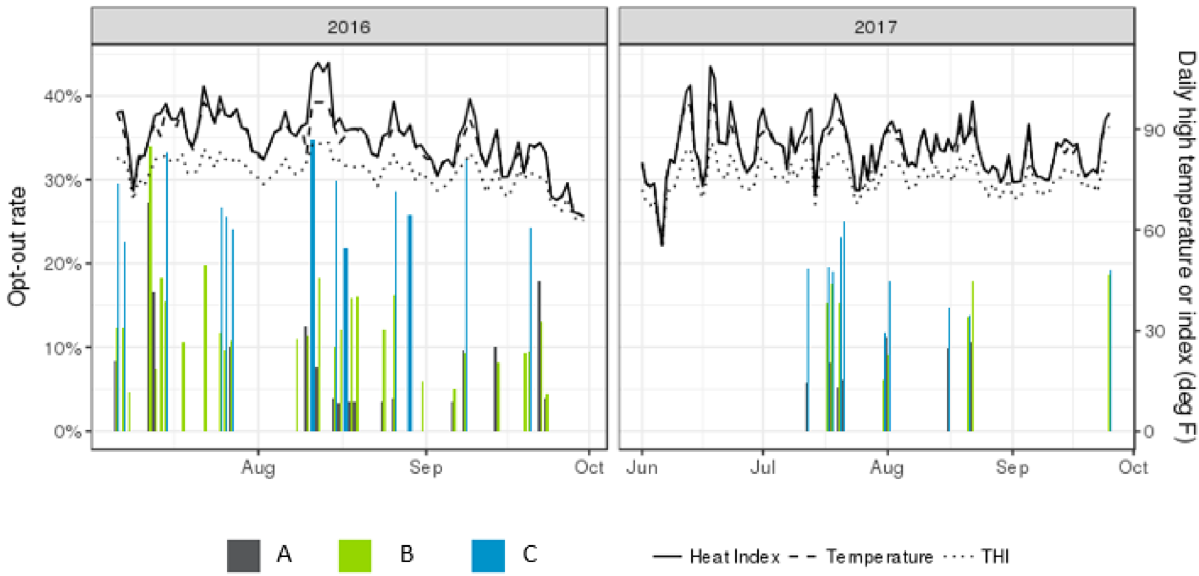

Source: Navigant analysis of CS and RHR data

To understand the opt-out habits of individual customers, the evaluation team created a histogram showing the proportion of events for which each device opted out. The in-depth interviews classified customers into four buckets based on their opt-out behaviors. Figure 43 shows the breakdown of these customer types by summer.

A large proportion (approximately 50%) of customers in each summer never opt out of any DR event. These could be people who were generally not at home during the event period or perhaps people who have a higher tolerance for heat and humidity. The in-depth interview findings classify these customers as the Reliable Participants.

Of those that opt out of some events but not all, the majority (82% in 2016, 81% in 2017) opt out less than half of the time. These are the Best-Effort Participants, who are interested in saving energy but will still opt out if they become too uncomfortable. A smaller proportion opt out of half of the events or more and would be considered the Comfort Seekers. Both groups of customers may decide whether to participate in each event based on their schedule, how their house heats up, and their personal comfort preferences, but the Comfort Seekers are more likely to opt out sooner and more often.

[47] The manufacturer names have been anonymized, as we feel program parameter details may be more predictive of these findings than inherent differences in technology between manufacturers.

A small group of customers (3%) always opt out when an event is called. The proportion of customers classified as Hold Outs who always opt out is stable year-over-year.

Figure 43. Devices that Sometimes, Always, or Never Opt Out of DR Events[48]

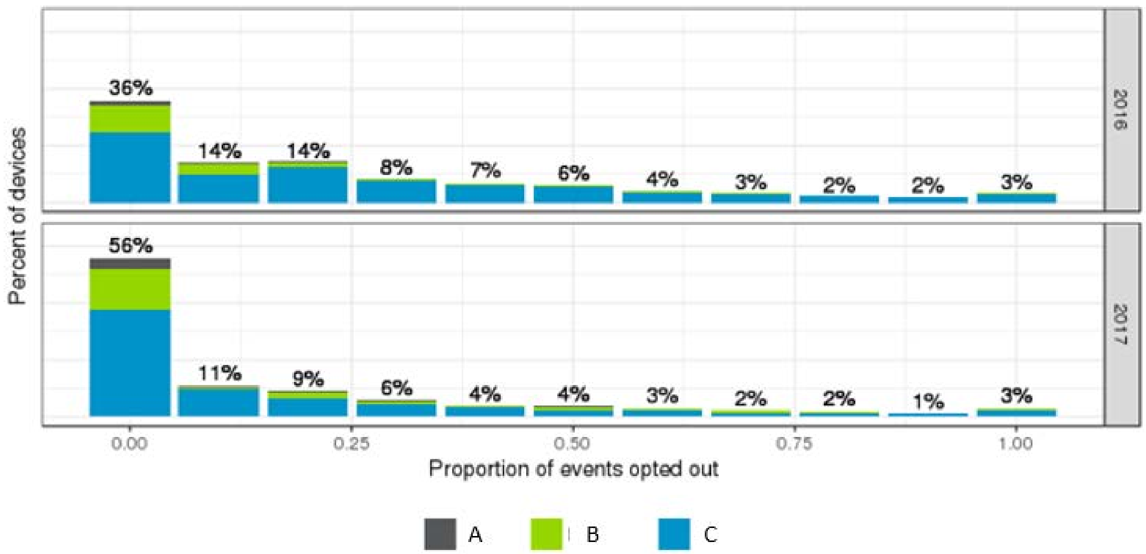

Source: Navigant analysis of CS and RHR data

Customers who opt out of events tend to do so uniformly throughout the duration of the event. Figure 44 shows the average cumulative opt-out rate throughout the events for 2-hour, 3-hour, and 4-hour events. The constant linear slope suggests that there are no specific triggers or thresholds that lead to many simultaneous opt outs or a sudden increase or decrease in opt outs. Rather, the point at which customers opt out during the event differs by household and may be reflective of heterogeneous comfort preferences.

The smooth cumulative opt-out profile could also reflect the distribution of times when people arrive home in the afternoon or evening. The DRSPs hypothesize that there is a relationship between people coming home and opting out of an event, but they have not collected information on the reasons for opting out to support this trend. Without access to occupancy data, it is not possible to know what proportion of people immediately opt out upon arriving home versus those that take more time to decide to opt out.

[48] As with the previous figure, we have anonymized manufacturers in this comparison, because we feel that program parameters, rather than inherent technology differences, may best explain differences shown in the diagram.

Figure 44. Average Cumulative Opt-Out Profile throughout Events[49]

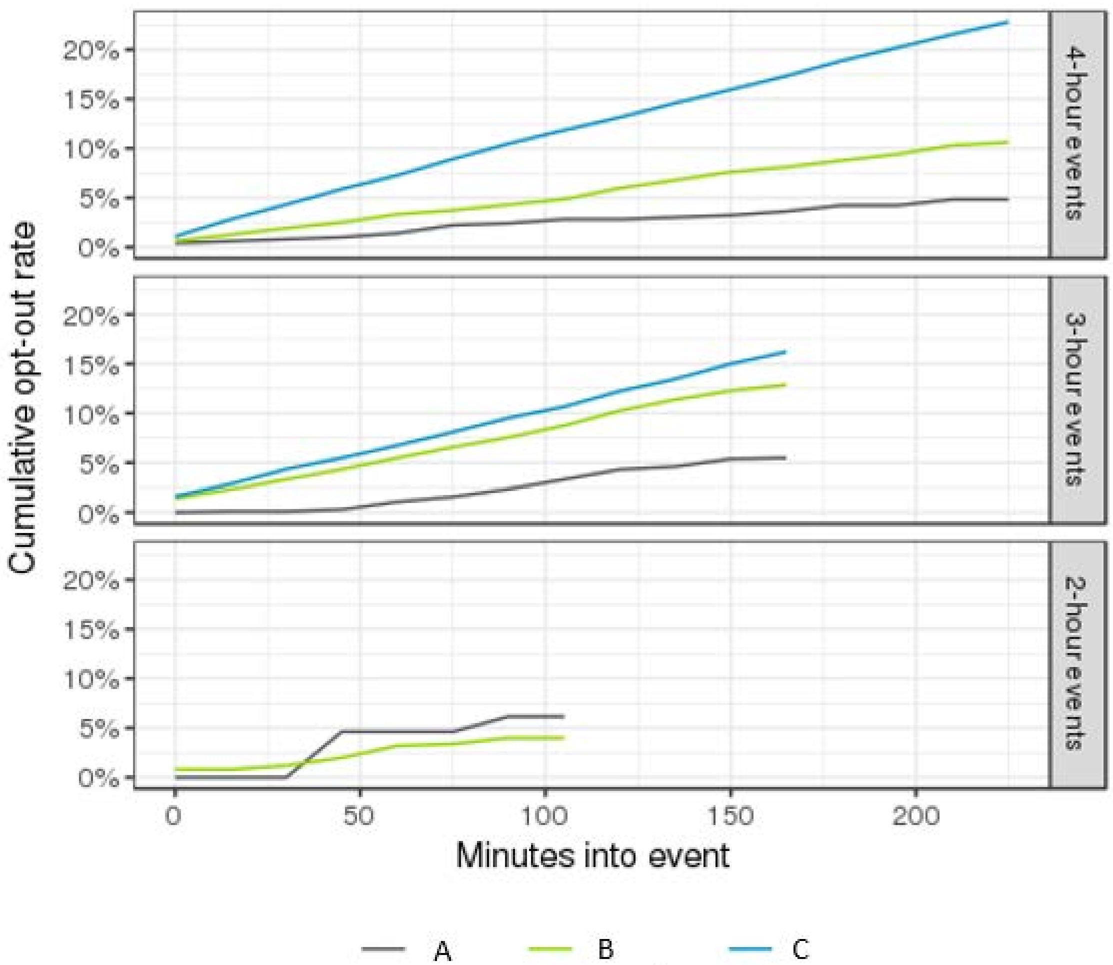

Source: Navigant analysis of CS and RHR data

[49] As with other graphics in this section, we have anonymized manufacturer, as we feel program parameters may have more effect on the findings in this case than inherent technological differences between products.

Figure 45. Number of Events and Opt-outs by Vendor and Event Duration, 2016 and 2017

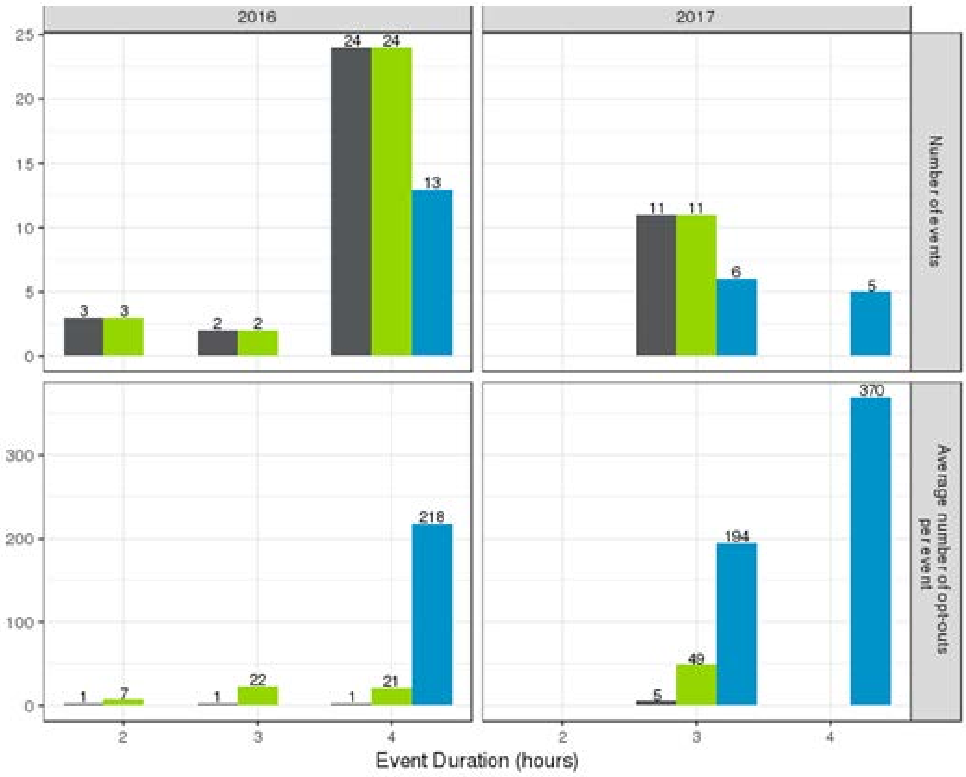

Source: Navigant analysis of CS and RHR data

Participants noted in interviews that they could predict, based on temperature and humidity, the days there would be an event. Although humidity was identified as the leading motivation to opt out, customers recognized that in their area, it is common for hot days to also be humid. Based on customer claims that becoming uncomfortable in their home led them to opt out, the team investigated how opt-out rates trend with heat and humidity-related variables.

Figure 46 plots the correlation between the opt-out rate and key daily high weather measures, including HBU, heat index, temperature, and temperature-humidity index. Each point represents one event and its size represents the number of participants in that event, excluding customers with connectivity issues or whose system was off or in heating mode. The plot also shows trend lines along with shaded 95% confidence intervals. For Nest and Honeywell thermostats, there is a positive correlation between higher heat and humidity-related weather indicators and higher opt-out rates. The team expects this to be the case as hotter, more humid weather leads to increased discomfort. In most cases, there is not enough data to determine the trend for ecobee.

Figure 46. Correlation between Daily Weather Variable Highs and Opt-Out Rates[50]

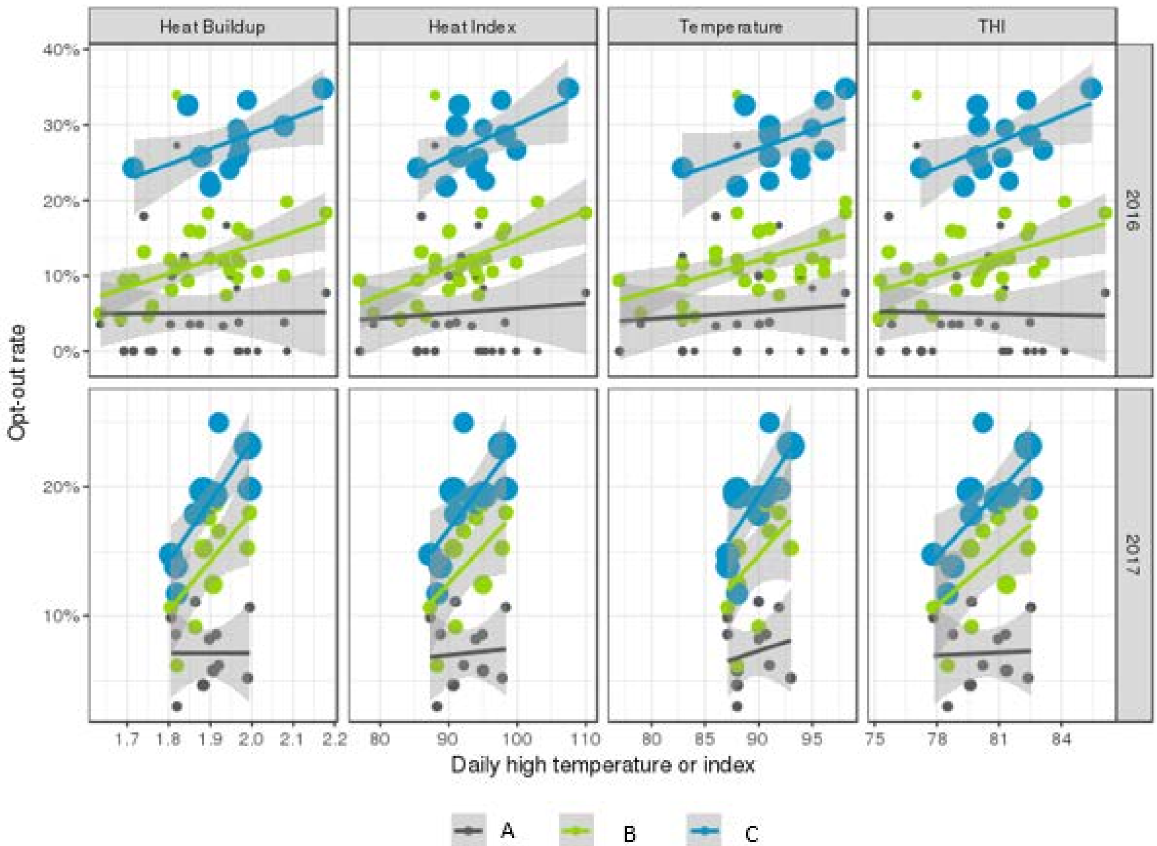

Source: Navigant analysis of CS and RHR data

3.4.6 Regression Analysis

Building on the learnings gained through exploratory data analysis and DR participant and implementer interviews, the evaluation team used regression analysis in this section to estimate causal relationships between weather and program variables and participants' propensity to opt out.

3.4.6.1 Methodology

The evaluation team used a linear probability model to estimate the likelihood that a DR program participant chooses to override preset controls. The model statistically relates the choice made prior to or during each event to various attributes (e.g., event attributes, weather attributes). The set of attributes considered was informed by the literature review, in-depth interviews, and exploratory data analysis. Navigant used participation data from the 2016 and 2017 summer DR programs, and used data from Thermostat DR demonstration project events spanning the summers of 2016 and 2017. The team's model includes variables characterizing the event, including event duration, event start time, length of

[50] As with other graphics in this section, we anonymized the manufacturers because we feel program parameters may explain as much of the differences as inherent differences in the technologies between manufacturer.

time between events, and customer participation status throughout the event at quarter-hourly intervals. Participation data was combined with quarter-hourly, ZIP-code level weather data, including temperature, relative humidity, HBU, and temperature-humidity index. The combined participation and weather data was aggregated to the event level, where one record represents one device on the day of an event. Table 7 shows the number of events called each summer.

Table 7. Number of Events Called

Thermostat Type	2016	2017
Nest	13	11
Honeywell, ecobee	29	11

Source: Navigant

This team is aware of structural changes made to the program over time; for example, changes to event frequency, cooling setpoint adjustments, event duration, and start time. It is also cognizant that many of these changes occurred between 2016 and 2017 and that program parameters also varied or changed differentially between Nest versus the vendor managing Honeywell and ecobee thermostat participation. The team's regression specification includes month-of-year fixed effects and interacts vendor with month of year to control for these structural program changes. The full model specification and additional methodological details are provided in Appendix Section 5.A.10.

3.4.6.2 Findings and Recommendations

The evaluation team's regression analysis revealed statistically significant relationships between event parameters, weather, and a participant's likelihood to opt out of DR events.

Table 8. Full Regression Outputs, DR Participation

Term	Estimate	Std. Error	Statistic	p-value
MOY201607	-4.35E-01	1.09E-01	-3.988	6.67e-05
MOY201608	-4.27E-01	1.11E-01	-3.856	0.000115
MOY201609	-4.59E-01	1.12E-01	-4.11	3.97e-05
MOY201707	-3.90E-01	1.06E-01	-3.681	0.000232
MOY201708	-3.84E-01	1.06E-01	-3.609	0.000308
MOY201709	-3.39E-01	1.09E-01	-3.122	0.001800
Monday	9.51E-03	6.19E-03	1.535	0.124844
Tuesday	7.28E-03	8.01E-03	0.908	0.363644
Thursday	3.55E-02	7.45E-03	4.765	1.90e-06
Friday	3.79E-02	7.79E-03	4.867	1.14e-06
Event Duration	3.32E-02	8.19E-03	4.046	5.21e-05
Event Start Hour	2.35E-02	6.76E-03	3.475	0.000512
Total Daily CDH65	2.65E-04	3.90E-05	6.787	1.16e-11
Event Max RH	3.82E-04	1.65E-04	2.31	0.020917
Lag Total Daily CDH65	-1.09E-04	3.57E-05	-3.044	0.002338
Lag Daily High RH	1.16E-04	1.46E-04	0.799	0.424364

Term	Estimate	Std. Error	Statistic	p-value
Consecutive Event	7.27E-03	9.87E-03	0.736	0.461488
MOY201607:Nest	1.53E-01	1.19E-02	12.811	<2e-16
MOY201608:Nest	1.42E-01	1.08E-02	13.185	<2e-16
MOY201609:Nest	1.57E-01	1.52E-02	10.341	<2e-16
MOY201707:Nest	3.93E-02	1.08E-02	3.643	0.000270
MOY201708:Nest	4.15E-02	1.19E-02	3.487	0.000490
MOY201709:Nest	2.69E-02	2.43E-02	1.105	0.269282
Consecutive Event:Nest	-2.41E-03	1.06E-02	-0.227	0.820336

Source: Navigant

Table 9 provides detailed information on the impacts of each independent variable based on the team's regression outputs.[51]

Table 9. Statistically Significant[52] Predictors of Likelihood to Opt Out of a DR Event

Independent variable	Impact on propensity to opt out
Day of week	Propensity to opt out *increases* by 3.6% on Thursdays relative to mid-week (Wednesdays).
	Propensity to opt out *increases* by 3.8% on Fridays relative to mid-week (Wednesdays).
Event duration	Every hour of additional event duration *increases* a participant's likelihood of opting out by 3.3%.
Event start hour	Starting an event an hour later in the day (say 3 p.m. rather than 2 p.m.) *increases* a participant's likelihood of opting out by 2.4%.
Cooling degree hours (CDH65)	Every degree-hour the temperature in a customer's ZIP code is above 65 degrees Fahrenheit on the day of an event *increases* that customer's likelihood of opting out by 0.03%. For example,[53] an event day that is 3.7°F hotter during all hours of the day increases the likelihood of opting out by 2.6%, relative to the 3.7°F cooler day.

[51] Full LPM regression output can be found in Appendix Section A.1.1 including all coefficients, standard errors, and p-values.

[52] Regression results are considered statistically significant in this study at the 95% confidence level (p value <= 0.05).

[53] An additional 3.67 degrees over the course of a day results in 3.67 * 24 = 88 more cooling degree-hours. The likelihood of opting out of an event on the hotter day increases by 88 * 0.0003 = 2.6%. Note that 88 cooling degree-hours is one standard deviation in total CDH65 on the day of an event.

Independent variable	Impact on propensity to opt out
	Every degree-hour the temperature in a customer's ZIP code is above 65°F on the day before an event *decreases* that customer's likelihood of opting out by 0.01%. For example,[54] a day before an event day that is 3.7°F hotter during all hours of the day decreases the likelihood of opting out by 0.9%, relative to the 3.7°F cooler preceding day.
Relative humidity (RH)	An increase in the maximum relative humidity during an event by 1% (absolute) *increases* the likelihood of opting out by 0.04%. For example,[55] the likelihood of opting out of an event with 17% (absolute) higher maximum relative humidity during the event increases by 0.7%, relative to the 17% less humid day.
Thermostat type[56]	In 2016, Nest users were 15% more likely to opt out relative to Honeywell and ecobee users.
	In 2017, Nest users were 4% **more** likely to opt out relative to Honeywell and ecobee users.

3.4.7 Weather and DR Program Participation Considerations

The following paragraphs draw conclusions based both on the regression analysis findings and the in-depth interview results. The evaluation team formulated its discussion as a series of actions to consider aimed at increasing customers' full participation in DR events.

Shift events earlier in the week and earlier in the day.

Consider the tradeoff between decreasing opt-outs by shifting events earlier in the week with higher savings potential on higher use days such as Fridays and weekends. The utility may consider planning for more events to take place on Mondays, Tuesdays, and Wednesdays, and fewer events to take place on Thursdays and Fridays to decrease opt-outs but will have to balance this with overall lower savings potential on these days due to lower usage. The team found that participants are more likely to opt out at the end of the week compared to mid-week. The likelihood of opting out is 3.6% higher on Thursdays and 3.8% higher on Fridays compared to Wednesdays. The trend at the beginning of the week could not be determined at a statistically significant level. While this approach would increase participation, it has to be

[54] An additional 3.74 degrees over the course of a day results in 3.74 * 24 = 90 more cooling degree-hours. The likelihood of opting out of an event the next day decreases by 90 * 0.0001 = 0.9%. Note that 90 cooling degree-hours is one standard deviation in total CDH65 on the day before an event.

[55] A 17% (absolute) increase in relative humidity increases the likelihood of opting out of an event by 17 * 0.0004 = 0.7%. Note that 17% (absolute) is one standard deviation in event maximum relative humidity. Also note that we believe program parameters specific to how the program was fielded for different manufacturers may explain more of the differences observed than inherent differences in the technologies.

[56] The team recognizes that program parameter choices in how programs for different vendors are administered may play a larger role in opt-out propensity than any underlying differences in technology between thermostats.

balanced with the utility need to lower load during peak use periods. To the extent that Mondays, Tuesdays and Wednesdays are lower use days overall, the utility would have to optimize the tradeoff between higher participation on Mondays through Wednesdays with higher costs on Fridays and weekends.

Similarly, plan for events to take place earlier in the day. The team found that participants are more likely to opt out of an event later in the day. The likelihood of opting out is 2.4% higher for each hour later in the day the event starts.

The finding from the regression analysis that higher opt-out rates occur at the end of the week and later in the day is aligned with the sentiment uncovered by the in-depth interviews. Many customers cited their schedules as very influential to whether they would opt out; namely, they would not opt out if they were not home. Intuitively, the team would expect more people to be coming home early, working from home, or taking the day off on a Friday. With a greater number of people at home at the end of the week and later in the day, it makes sense to expect higher opt-out rates.

Call more frequent but shorter events.

Avoid pushing the limits on participant's comfort for diminishing returns. The team found that longer events resulted in higher opt-out rates. For each additional hour of event duration, the likelihood of opting out increased by 3.3%. However, there is no statistically significant evidence that more frequent events cause higher opt-out rates due to fatigue.

Customers did mention that back-to-back events might make them tired of being uncomfortable in their house. If shorter events end before customers become uncomfortable or before they even arrive home, then it may be possible to eliminate the problem with more frequent events.

Arrange for better management of participants' comfort during the events.

Whenever possible, consider ways to manage participant's comfort, such as encouraging efficient pre-cooling. The team found that hotter and more humid days resulted in higher opt-out rates, and these are often the most important days to include as DR events.

It may not be reasonable to completely avoid calling events on hot and humid days, but strategies like adaptive algorithms and household thermal models can help diminish the effects of heat and humidity on customer comfort. A snooze button that offers temporary relief from the event without having to opt out of the entire event could be appealing to customers with an interest in saving money and energy but become too uncomfortable to participate consistently for the duration of the event.

Offer more motivating incentive packages for event participation.

The CS incentive structure requires participation in at least 75% of the DR events to receive the $25 annual reward. The team found significantly lower opt-out rates for thermostats enrolled in the CS program (Honeywell and ecobee thermostats) compared to the RHR program (Nest thermostats).

While the team cannot prove causation without a randomized control trial and it has not controlled for other program differences, the team hypothesizes that having incentives tied to event participation leads to a lower opt-out rate. CS program participants expressed awareness of the participation requirement when interviewed, though they still prioritized comfort over what they considered a small financial incentive. To make the most of the financial incentive's impact, the incentive should be sized appropriately to motivate participation.

Integrate gamification strategies and adaptive algorithms to decrease opt outs.

Gamification strategies might increase customer motivation to participate for the entire event or for longer than they would have otherwise. Gamification strategies are new in the area of DR programs and the industry has not yet thoroughly researched and evaluated these approaches. However, researchers and vendors are exploring how to use the motivating aspects of gamification with inspiration from devices like FitBit, miles per gallon feedback provided by some cars, the points and levels that are ubiquitous in online games, and social networks. Several utilities have piloted DR programs internationally with gamification elements with early indicators of positive results.[57] AlSkaif et al (2018) outline a theoretical framework for applying gamification strategies to residential DR. Key components of a gamification approach include:

- **Rewards:** Beyond bill savings, these rewards could include, for example, earning points for participation or sustained participation; achieving "levels" based on participation, and tiered rewards based on amount of participation.
- **Social connections:** Connections can motivate customers to sustain their participation. Competitions provide normative social influence; communication and collaboration among participants can motivate participants to work toward collective reductions.
- **Personalized goals:** Participants can set and track progress toward meeting goals for demand reductions and event participation.
- **Timely feedback:** Participants can track their progress toward their goal and track points/levels/rewards earned in real time.

The examples of gamification for DR that the team examined came from vendors providing apps that include all the rewards, goals, and communication elements and require near real-time meter data. However, some elements of gamification could possibly be incorporated through mobile/SMS notifications or through more timely notification via email of participants' accomplishments and rewards earned during events.

[57] For two examples of vendor-reported results, see: http://www.openenergi.com/news-posts/gengame-partnership-demonstrates-power-gamification-drive-domestic-dsr-success/ and https://www.utilitydive.com/news/gamification-for-the-grid-inside-bidgelys-australian-demand-response-pilo/419230/ More research is needed before the team can definitively declare these approaches successful.

The team sees promise in using gamification strategies to increase customer motivation to participate for the entire event or for longer than they would have otherwise. The team found that customers were not opting out in reaction to certain triggers but rather opting out at any point during the event when they became uncomfortable. The motivating features of gamification may nudge those customers to continue their involvement. This finding suggests that customers, in particular the Best-Effort Participants, may be receptive to features that make them feel more involved in the program's success or give them additional motivations to continue their participation. The gamification strategies may also attract some Hold Outs and Comfort Seekers to at least partial participation.

The utilities may use these behavioral insights to design incentive programs that encourage customers to refrain from opting out during the most extreme weather conditions, precisely when the utility most needs their participation in DR. The PAs may provide incentives that scale with the degree of discomfort (i.e. incentives increase with increasing heat and humidity) or may offer the chance of a large incentive to those that stay engaged for the duration of an event during the worst weather conditions. These tactics may help the utility keep customers from opting out when they value their participation the most.

Additional research will be needed to assess the effectiveness of gamification strategies and if those strategies need an integrated app experience or can be accomplished via other communication channels.

Additionally, the PAs may consider or continue pursuing adaptive algorithms that help manage temperature adjustments to minimize opt outs while maintaining savings. Adaptive algorithms can help diminish the effects of heat and humidity on customer comfort. A snooze button that offers temporary relief from the event, for example, without having to opt out of the entire event could be appealing to customers with an interest in saving money and energy but become too uncomfortable to participate consistently for the duration of the event.[58]

[58] National Grid is working through EnergyHub in 2018 to implement the "Firm Load Dispatch" adaptive algorithm for its DR events. The team suggests the PAs analyze the effectiveness of this approach and modify or continue this type of approach based on effectiveness. The approach is described here: http://www.energyhub.com/firm-load-dispatch-paper

4. CONCLUSIONS AND FUTURE RESEARCH

The evaluation team's study provides new findings on the effects of heat and humidity on AC use and people's participation behavior in DR programs. Moreover, it is one of only a handful to pair both quantitative, data-based analysis with qualitative survey methods to provide more context and a deeper understanding of results.

The most important finding from Navigant's research on AC user types and thermostat use behavior is that users display an enormous range of thermostat-related behaviors. Because thermostats have become sophisticated, allowing a wide range of customization, customers use their thermostats in nearly as many unique ways as the number of thermostats the team had available for study. Rather than conforming to an easily identifiable AC user type category, the team's research finds that identifying and quantifying key thermostat use behaviors is more actionable. The PAs can then base targeted messaging or other program interventions in the future on customers who display identified behaviors in order to maximize the potential for behavior change. Future research using a sample size of thousands of thermostats and machine learning might be able to categorize users in a meaningful way by AC user type, but this may not be as useful as identifying and targeting specific thermostat use behaviors.

In line with previous research, this study quantifies the strong effect of heat on first AC use, use throughout the season, and people's likelihood of opting out of DR events. It also demonstrates that, while humidity is not a key factor in determining first AC use, it does play a strong role in determining AC use throughout a season and in causing people to opt out of DR events.

This study also shows that program configuration considerations such as the duration and timing of events, thermostat type and vendor, and proximity of events may have as large an effect on DR event opt-out behavior as weather variables do. Thus, PAs may want to carefully consider program characteristics in addition to weather variables when planning DR events to optimize savings.

This study reveals that people's first and last use of AC during the cooling season are not governed by the same processes and therefore cannot be modeled identically. People's last AC use is highly concentrated over a short period of time, whereas first AC use instances occur more gradually as the summer progresses. The team's limited last-use findings suggest a lengthier dataset might provide more valuable insights into this behavior, offering the ability to observe last AC use instances over a period of multiple years rather than concentrated in a single month of one year.

Another area for future research is to explore the effects of weather on an even more granular sub-daily level. While the modeling complications introduced by hourly or sub-hourly models outweighed their benefit in the current study, future studies might use this as a starting place to study weather's effect on AC use throughout the season using AMI data, focusing particularly on sub-daily relationships.

Finally, the largest limitations to the evaluation team's study were the limited sample size of AMI dataset customers with thermostat data and the relatively mild temperatures in 2014. Given a much larger sample size and potentially wider temperature variation, future research might be able to draw even more granular findings on which to extend findings more reliably to very hot temperatures.

5. BIBLIOGRAPHY

Agnew, K., Goldberg, M., & Rubin, R. (2004). You're Getting Warmer: Impacts of New Approaches to Residential Demand Reduction. ACEEE Summer Study on Energy Efficiency in Buildings. Retrieved from http://www.eceee.org/library/conference_proceedings/ACEEE_buildings/2004/Panel_2/p2_1

AlSkaif. T., Lampropoulos, I., Van den Broek, M. Van Sark, W.. (2018). Gamification-Based Framework for Engagement of Residential Customers in Energy Applications. Energy Research & Social Science, 44, 187–195.

BPA. (2016). Nest Learning Thermostat Pilot Program Savings Assessment.

Cadmus. (2013). Liberty Utilities Wi-Fi Programmable Thermostat Pilot Program Evaluation. Prepared for Liberty Utilities.

Kempton, Willett. (1986) Two theories of home heat control. Cognitive Science 10(1): 75-90.

Meier, A., Aragon, C., Hurwitz, B., Mujumdar, D., Peffer, T., Perry, D., & Pritoni, M. (2010). How People Actually Use Thermostats. Controls and Information Technology, 2, 193–206.

Navigant. (2017). National Grid 2016 Residential Wi-Fi Thermostat DR Evaluation Final Report

Navigant. (2017) Massachusetts Baseline Study. Prepared by Navigant Consulting, Inc. for the Massachusetts Program Administrators and Energy Efficiency Advisory Council.

Peffer, T., Pritoni, M., Meier, A., Aragon, C., & Perry, D. (2011). How people use thermostats in homes: A review. Building and Environment, 46(12), 2529–2541. https://doi.org/10.1016/j.buildenv.2011.06.002

Rathouse, K., & Young, B. (2004). RPDH15: Use of Domestic Heating Controls. Watford: Building Research Esablishment (UK).

Reed, J. H. (1991). Physical and Human Behavioral Determinants of Central Air-Conditioner Duty Cycles. IEPEC, 208–217.

Stewart, J. (2014) Energy Savings from Honeywell Connected Thermostats. 2014 Behavior Energy and Climate Change (BECC) Conference, Conference Proceedings.

Walker, I. S., & Meier, A. K. (2008). Residential Thermostats: Comfort Controls in California Homes. Lawrence Berkeley National Laboratory Berkeley, California March 2008 LBNL-938E.

APPENDIX A. ADDITIONAL MATERIAL

A.1 Supplemental Thermostat Data Exploration

In this section the evaluation teams provide additional graphics exploring the thermostat data, which may provide additional context given the complexity of thermostat data.

Figure 47 highlights a key factor of the thermostat dataset, that the number of installed thermostats increased throughout the study period. During some periods (e.g., mid-July), the number of installed thermostats was increasing at a similar rate to the number of devices which had reported first AC use. This suggests that ACs may have been used for the first time the same day they were installed. To deal with the complexity this trend introduced, the team limited its analysis to devices installed before May 1 in the modeling of first AC use.

Figure 47. Counts of Thermostats Installed and Thermostats with First AC Use

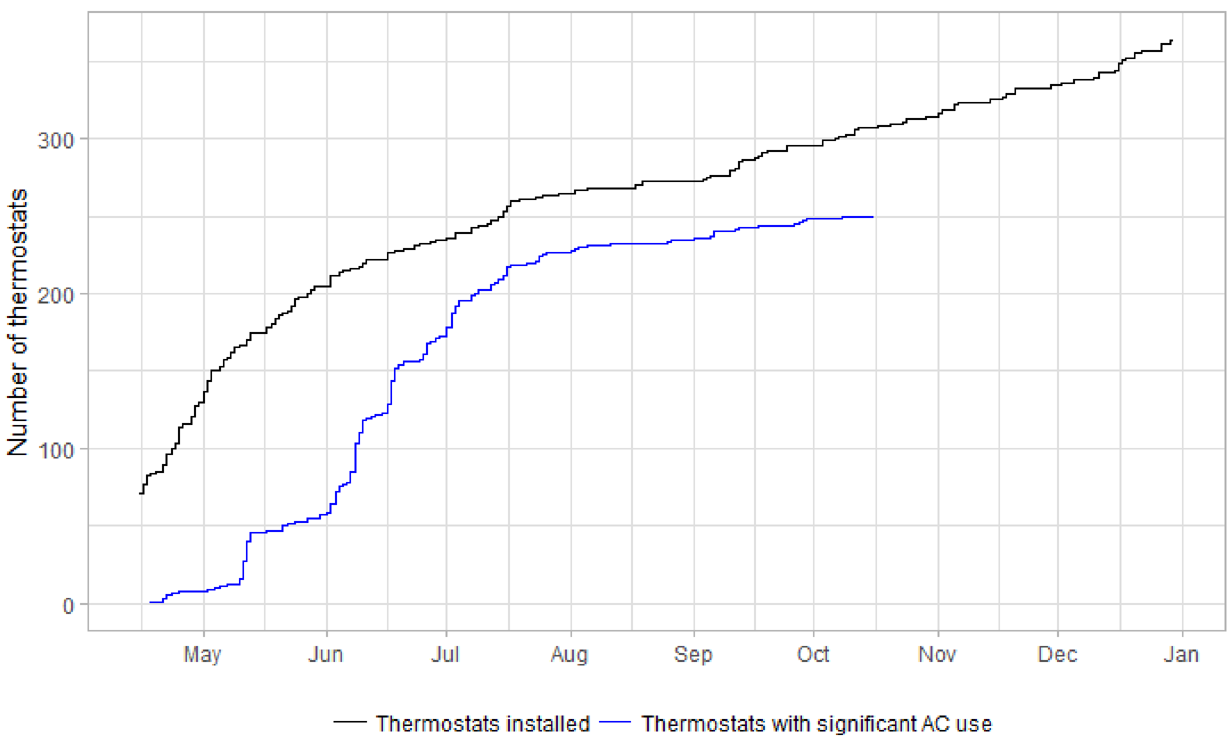

Source: Navigant

Figure 48 shows similar data to that provided in Figure 47, but rather than cumulative installed and cumulative first use thermostats, as seen in the upper plot, the count of thermostats reporting any data (black) and reporting any AC runtime (blue) by day. The decline in the black line starting in June indicates that some devices were either not communicating with the data collection service or had been uninstalled. However, since devices with a substantial amount of missing data were removed from the analysis during data cleaning, the team does not anticipate this to affect the results negatively. However, it does suggest the sample size would have been larger had this attrition not occurred. The lower plot simply illustrates that the spikes in the AC use correlate with spikes in temperature, offering reassurance that AC use is responsive to temperature.

Figure 48. Comparison of Thermostats Reporting AC Runtime against Temperature and THI

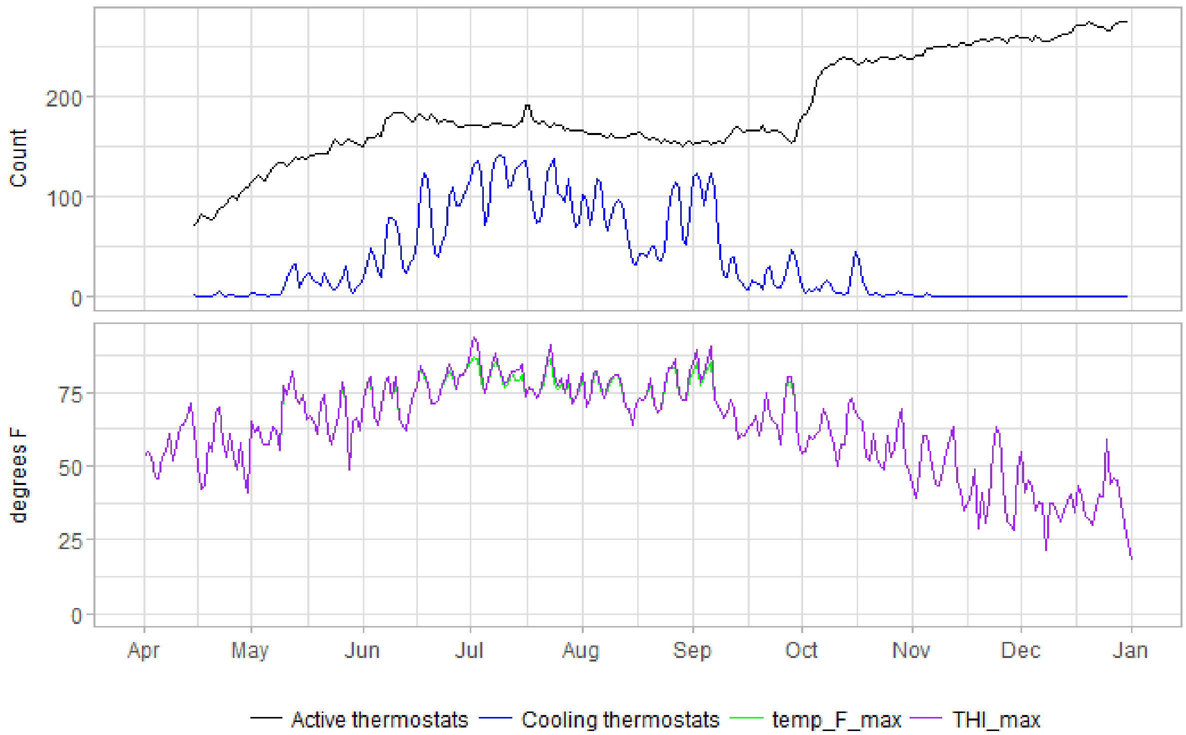

Source: Navigant

Figure 49 provides another perspective on the relationship between cooling mode and AC runtime (cooling state). Additional views of this data are shown in Figure 13, Figure 29, and Figure 31. Figure 49 enables a better understanding of how these variables relate. The subplot titles in this figure are the month from which the data was taken. The blue lines are contour plots indicating the relative density of the black plots. For example, in June, customers that spend nearly 100% of their time in cooling mode (grey dots stacked on the far right of the plot) generally spend around 13% of their time with the AC running (center of the blue circles), but that some spend as much as 100% of their time with the AC running. It is also interesting to look at the distribution of percent-runtime values for customers who spend less than 100% of their time in cooling mode. Some of these customers have higher percentage runtimes than the 100%-cooling-mode customers (e.g., in June, the grey dots near 75% cooling mode and near 50% AC runtime). This is likely due to lower cooling setpoints among these customers who are nonetheless switching their devices into off mode for some portion of the month.

Figure 49. Relationship Between Cooling Mode Duration and AC Runtime by Month

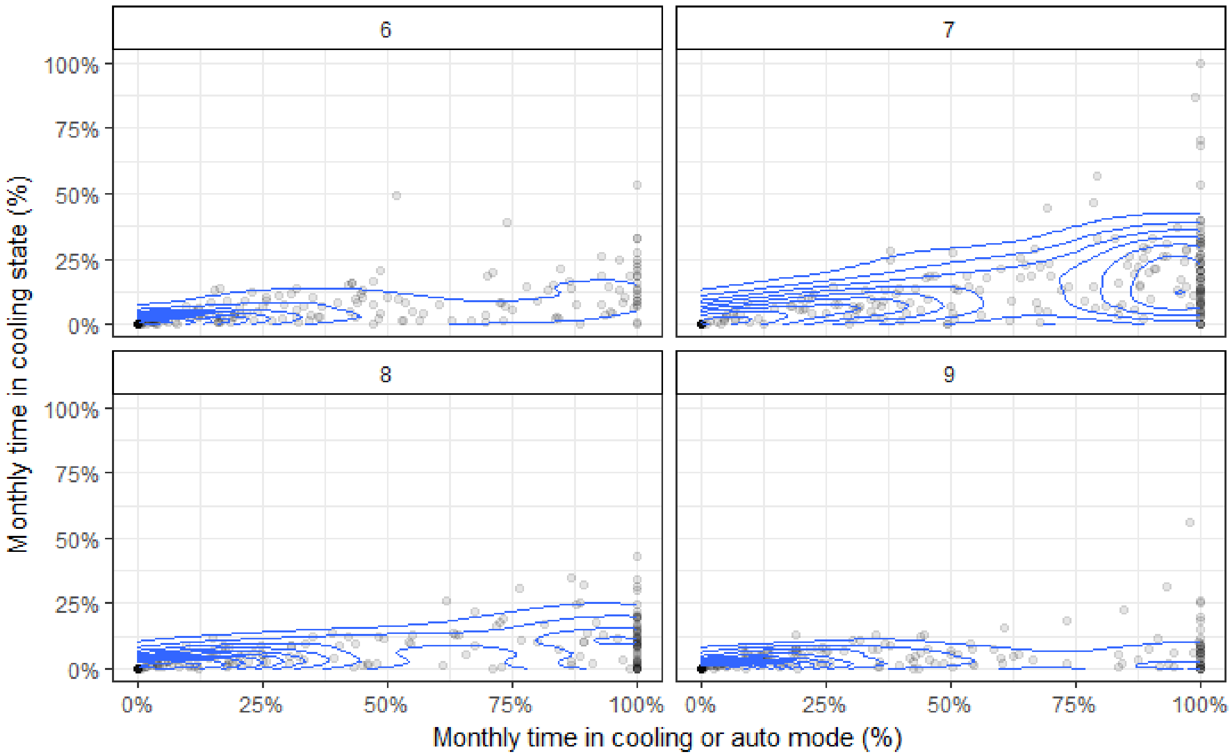

Source: Navigant

A.2 First AC Use Methodology

For both the first and last AC use models, the dependent variable is a binary variable indicating whether first use (or last use) has occurred. Given the binary outcome variable, the evaluation team has chosen to use a logistic regression as the basis for the model. The team's preferred model specification is given below:

$$Logit(Y_{it}) = \beta_1 MaxCDH_t + \beta_2 MaxTemp_{t-1} + \beta_3 MaxHumdity + \beta_4 (MaxCDH_t \times Weekend_t) + \varepsilon_{it}$$

Y_{it}	= a binary variable taking on a value of 1 if thermostat *i* experienced first AC use for the season on day *t*
$MaxCDH_t$	= the daily maximum of the hourly CDH (base 65) values for day *t*
$MaxTemp_t$	= the maximum temperature for day *t*
$MaxHumidity_t$	= the maximum humidity for day *t*
$Weekend_t$	= binary indicating whether day *t* was a weekend
ε_{it}	= an error term for thermostat *i* on day *t*

The day of first use only occurs once, so none of the data following the day of first use provides useful information. It would be inaccurate to classify days subsequent to first use for thermostat *i* as not first use days (i.e., false in the binary variable) since first use has already occurred. For this reason, the logistic regression is conducted on the data leading up to and including first use for all 54 thermostats in the sample.[59]

The team compared several different model formulations for estimating the impacts of weather on first AC use before arriving at the preferred model specified above, inspecting the coefficients for alignment with the team's intuition regarding sign and magnitude of the effects and quantitatively comparing performance of model variations for robustness.

To evaluate model performance, the evaluation team compared regression-based predicted values for the cumulative AC turn-on rate with actual cumulative AC turn-on values throughout the season for several model variations.[60] The team used a method known as cross-validation analysis to:

- Evaluate model performance by comparing the predicted cumulative AC turn-on rate to the actual
- Assure the preferred model specification is not over-fit
- Compare error between the model variations to assure the preferred model error is reasonably low

In cross-validation the evaluation team first trains the regression models on a subset of the data and then generates the predicted cumulative probability curve using the full dataset. The team repeats this exercise many times using different subsets for the training data. The error of each fit is calculated as the square root of the mean of the squared difference between the predicted curve and the actual curves—often called the root mean squared error (RMSE). Lower RMSE indicates a better fit.

Figure 50 illustrates the closeness of the predicted versus actual cumulative AC turn-on rates based on the cross-validation outcomes. Each of the faint grey lines is the cumulative probability curve from one iteration of the cross validation. The third subplot from the top on the left of Figure 49 is the team's preferred model specification. Visual inspection reveals it is the specification that best predicts the cumulative rate of first AC turn-ons relative to actual.

Calculating the RMSE between each of these lines and the true (blue) line provides the data behind the box plots in Figure 51. In this figure, the box plots give a sense of the distribution of errors associated with each model variation. Lower errors indicate better fit, and narrower boxes indicate less sensitivity to which subset of the data was used for training the model in the cross-validation analysis. High sensitivity to which subset of the data is used for training would be indicative of overfitting. The team's preferred specification is the box plot third from the top in Figure 51. The low RMSE and low sensitivity to training subset reassure the team that its preferred specification has relatively low RMSE and does not suffer from overfitting.

---

[59] The complete thermostat dataset includes more than 54 devices, but as described previously, the team narrowed its analysis of first use to devices with 80% or greater complete data between May 1 and October 1, which narrowed the sample to 54 devices.

[60] Logistic regressions are a form of classifier model. A typical approach to measuring performance of these classifier-type models is to specify a probability threshold (e.g., 50%) above which one considers the model's prediction to be positive (in this case, first use occurring) and then construct a confusion matrix (rates of true-positive, false-positive, true-negative, and false-negative) or the area under the curve (AUC) of the receiver operating characteristic (ROC) curve. Both of those approaches look at the model's ability to correctly classify specific observations as either a first use or not first use day. However, for the team's analysis, *the accuracy of the model at the individual observation level is less important than the cumulative accuracy across all observations*. Stated another way, the cumulative probability of all ACs having experienced first use is more useful than estimates for each individual home on each individual day. For this reason, the team chose to compare cumulative predicted turn-on rates with actual rates to evaluate performance.

Figure 50. Modeled and Observed First Use Cumulative Probability Curves

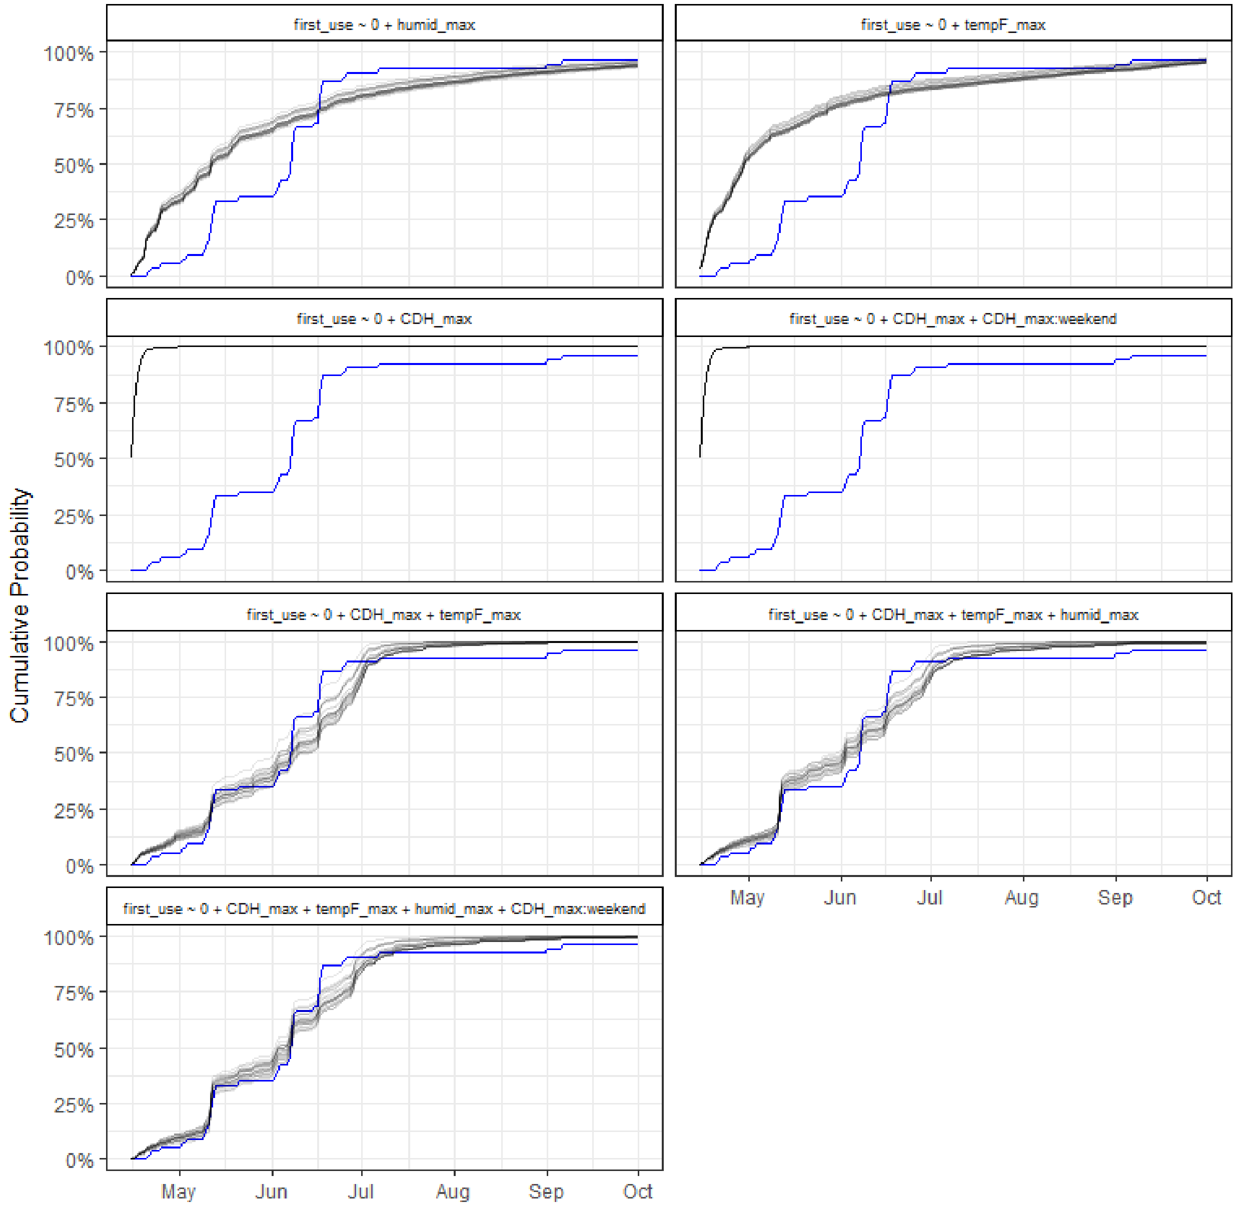

Source: Navigant

Figure 51. Comparison of First Use Model RMSE

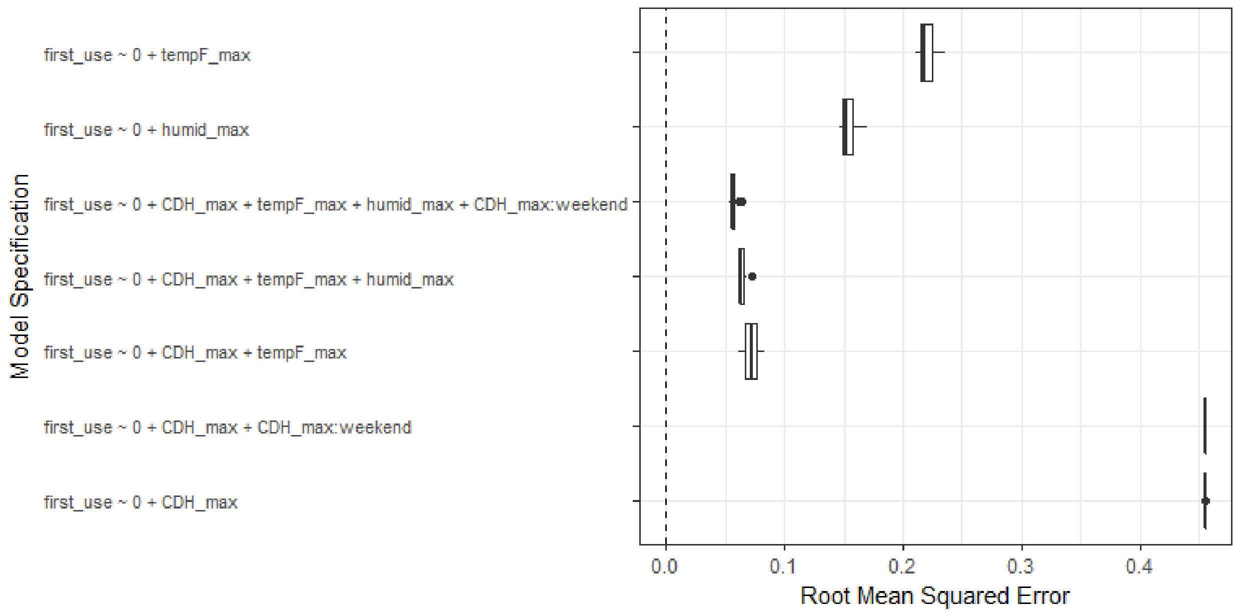

Source: Navigant

A.3 Supplemental First AC Use Findings

In this section, the team shares additional findings associated with the investigation into what determines the timing of when customers first use their ACs during a season.

In Figure 52 the team shows scatter plots for a variety of weather variables versus the cumulative number of ACs which had experienced first use on the day of that weather observation. Within each subplot, there are many faint, grey dots. Each dot represents the weather and count of first AC uses on one day. Clusters or overlapping dots become darker grey and then black as the number of overlaps increase. The titles of the subplots contain the names of their weather variables. Superimposed on the scatter plots is a regressed trend line.

The team used this plot (and other like it) to visually identify strong correlations between the weather variables and first AC use. Trend lines which slope up to the right suggest a positive correlation between that weather variable and the first use. Trend lines that curve or have sharp bends are suggestive of nonlinear relationships between the variables (e.g., “humid_avg”, second row, sixth column). These trends informed the team’s choices of model variations to compare using the previously discussed cross-validation techniques.

Figure 52. Exploratory Graphical Relationship between Weather Variables and First AC Use

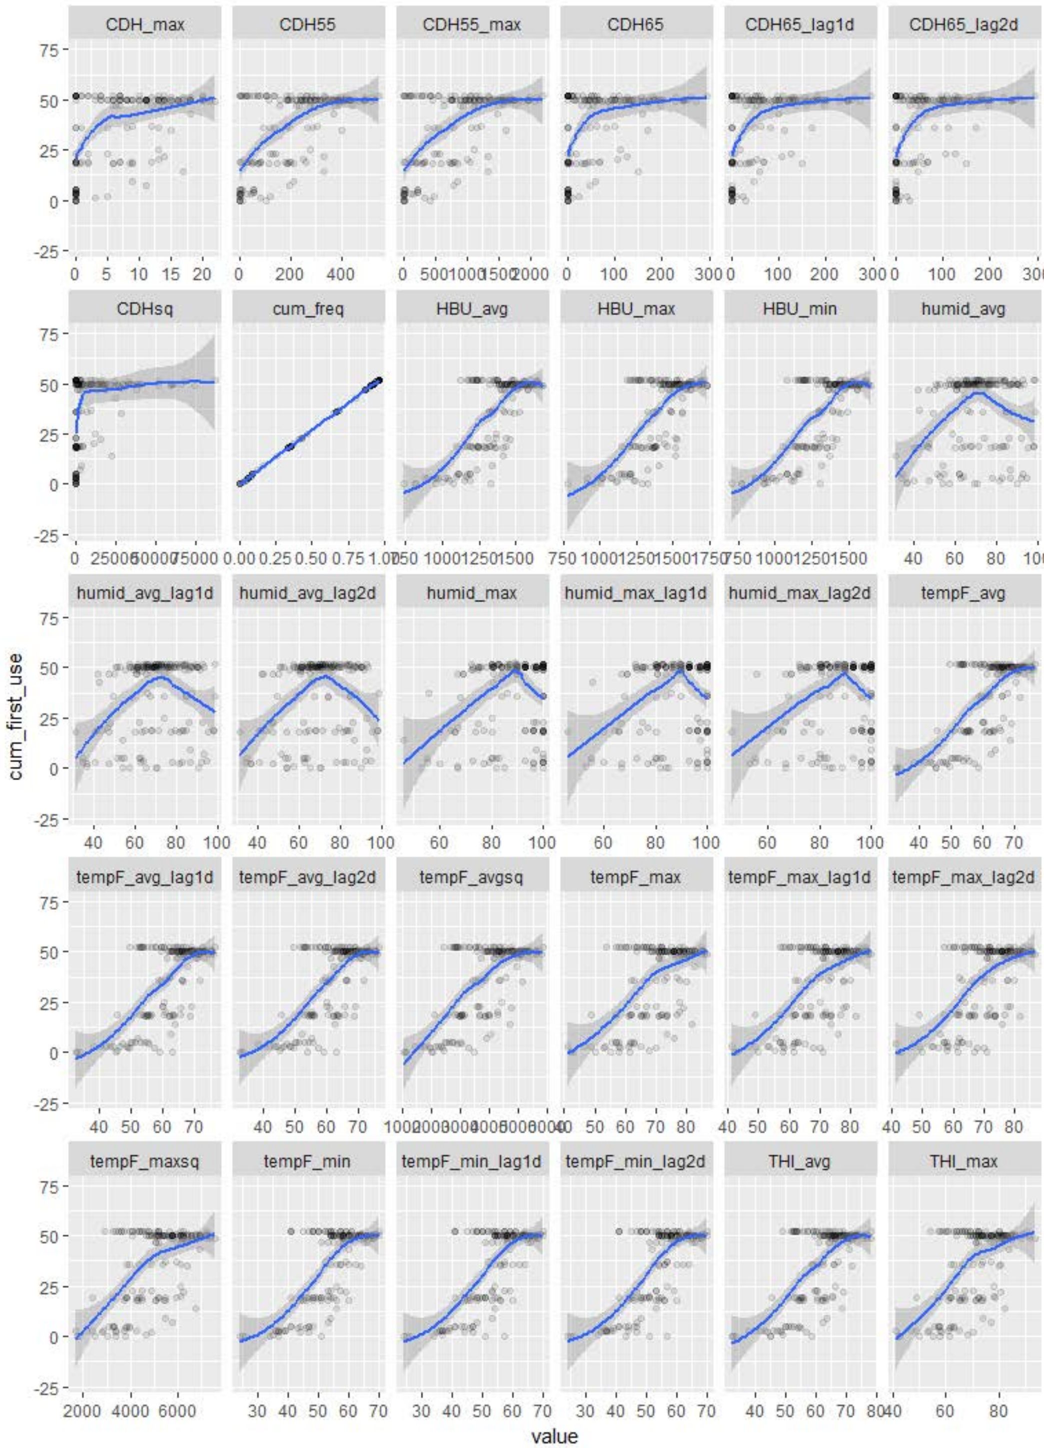

Source: Navigant

As discussed in Section 3.2.2.2, the AMEs of logit models are more useful for understanding the relationship between the predictor variables and the response variable. However, for completeness, the team has included the logistic coefficients of the preferred model specification (see Section 3.2.2.1 for the full specification) in Table 10. These coefficients would be useful for predicting the probability of that any ACs that have not yet turned on for the first time during a season would turn on under given (likely forecasted) weather conditions. Note, however, that to determine how many ACs would likely be running on a given day, the cumulative probability of first use occurring must be calculated, and thus weather assumptions must be made and probabilities calculated for all days in the season leading up to the day of interest.

Table 10. Logit Coefficients of the First Use Model

Term	Estimate	Std. Error	Statistic	p.value
MaxCDH	0.324	0.040	8.052	8.14E-16
MaxTemp	-0.054	0.013	-4.261	2.03E-05
MaxHumidity	-0.030	0.008	-3.522	4.28E-04
MaxCDH × Weekend	-0.071	0.023	-3.108	1.88E-03

Source: Navigant

A.4 Supplemental Last AC Use Findings

In this section, the evaluation teams share findings associated with its investigation into what determines when customers use their AC systems for the last time during a cooling season. As discussed in the main body of the report in Section 3.2.1.2, the results of the modeling of last AC use were not statistically meaningful and demonstrated that the same model could not successfully be applied both to first and last AC use as a function of weather. The team's findings on the relationship between last AC use and weather are directional and suggestive, not causal or predictive. However, the team has included these for informational value and as a starting place for future research in this area.

Figure 53 shows distributions of a variety of weather variables grouped into last use days and non-last use days. Within each subplot, one sees two violin-plots (i.e., horizontally mirrored histograms): one for the last use data (indicated with a one on the vertical axis) and one for the non-last use data (indicated with a zero). The titles of the subplots contain the names of their weather variables. Differences between the violins in a given subplot suggest there could be some correlation between that weather variable and the occurrence of the AC last use. For example, the last use distribution in the tempF_max subplot (fourth row, fourth column) is shifted lower than the non-last use distribution. This suggests that last use is more likely to occur on days with lower maximum temperature (conversely, few last uses occurred on high maximum temperature days).

Figure 53. Relationships between Weather Variables and Last AC Use

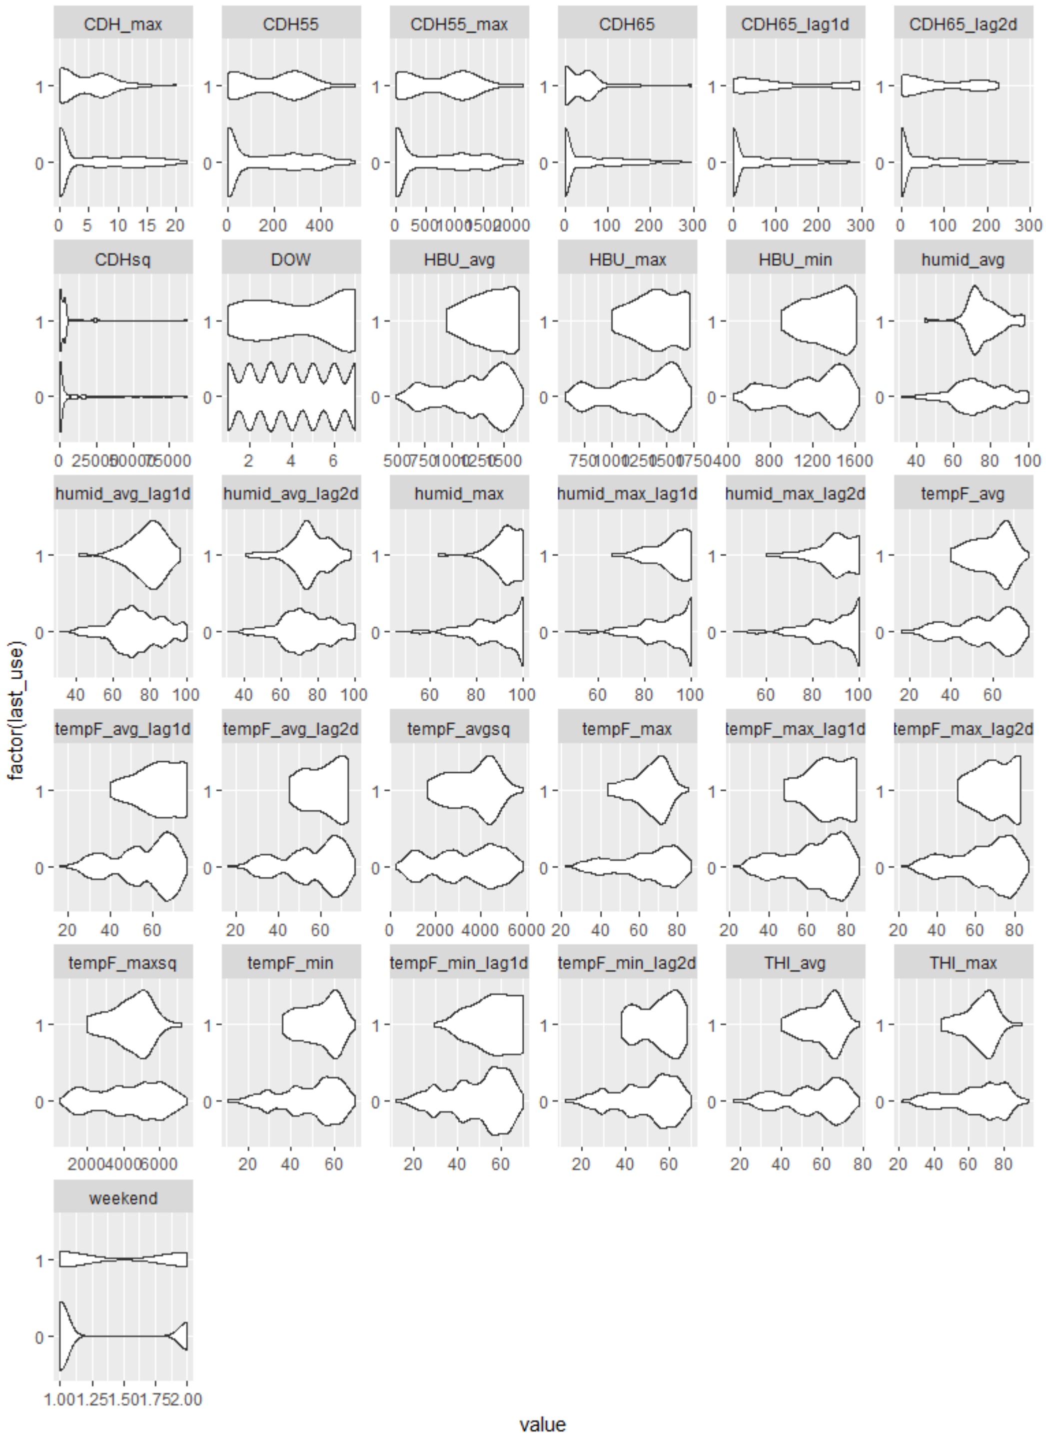

Source: Navigant

Although the fit and predictive power of the last use models was not nearly as good as the first use models, useful information can still be gleaned from the modeling outputs. It is likely that additional model tuning and feature engineering could yield a model with greater predictive power, but these formulations do give us insights into the correlation between readily available weather variables and last AC use. Below is the team's preferred formulation of the model:

$$Logit(Y_{it}) = \beta_1 AvgTemp_t + \beta_2 AvgTemp_{t-1} + \beta_3 AvgHBU + \beta_4 MaxCDH65_t + \beta_5 TotalCDH65_{t-1} + \varepsilon_{it}$$

Y_{it}	= a binary variable taking on a value of 1 if thermostat *i* experienced last AC use for the season on day *t*
$AvgTemp_t$	= the average temperature in Fahrenheit of day *t*
$AvgTemp_{t-1}$	= the average temperature in Fahrenheit of day *t-1*
$AvgHBU_t$	= the average heat buildup of day *t*
$MaxCDH65_t$	= the maximum of hourly CDH (base 65) values during day *t*
$TotalCDH65_{t-1}$	= the total daily CDH (base 65) day *t-1*
ε_{it}	= an error term for thermostat *i* during day *t*

Similar to the first use findings the team shared in Section 3.2.2.2, for last use the team compared regression-based predicted values for the cumulative AC turn-off rate with actual cumulative AC turn-off values throughout the season for several model variations. In Figure 54 the team presents the modeled and observed cumulative probability curves for variations of the last use model followed by, in Figure 55, the RMSEs from cross validation.

Figure 54. Modeled and Observed Last Use Cumulative Probability Curves

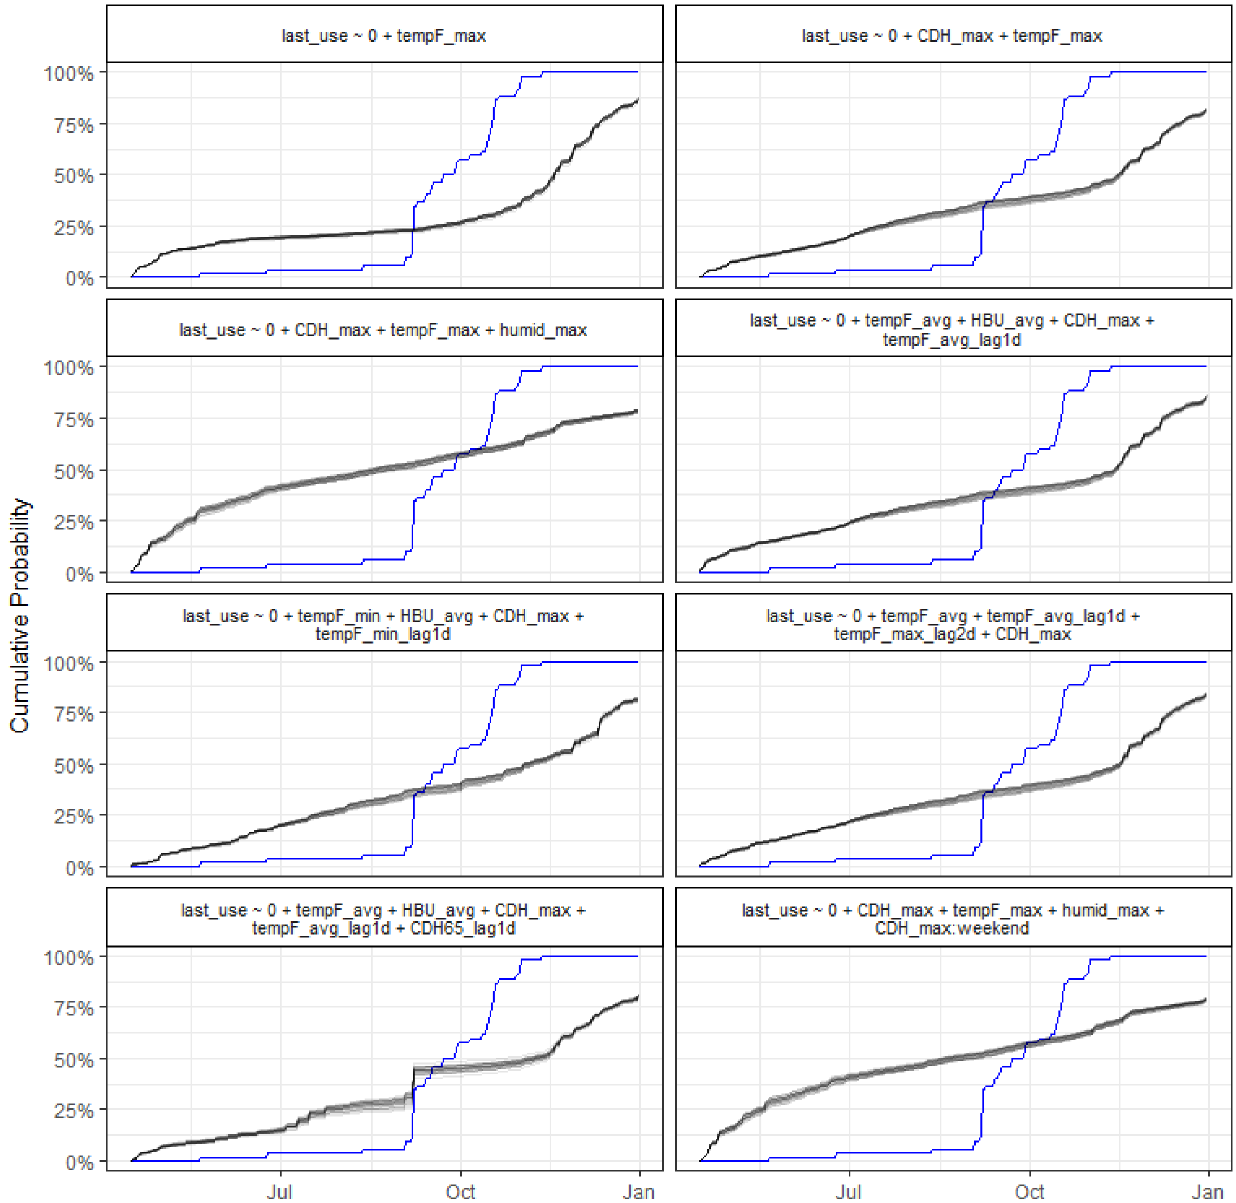

Source: Navigant

Figure 55. Comparison of Last Use Model RMSE

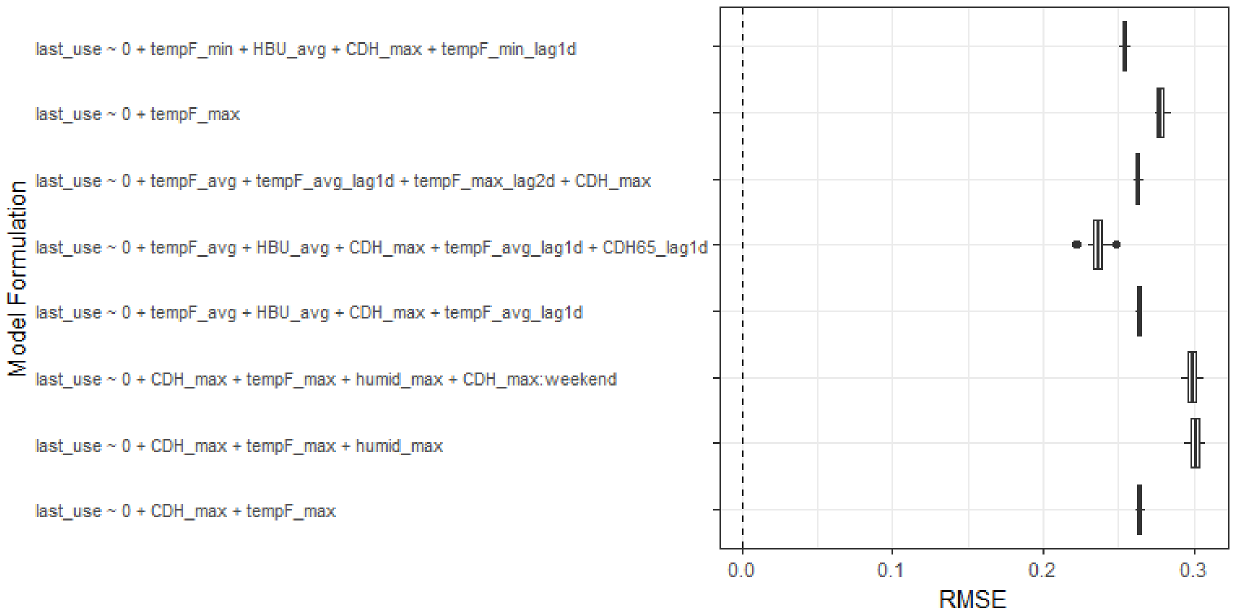

Source: Navigant

While none of the models performed especially well, the model that included 1-day lagged terms for both average temperature and maximum CDH did partially anticipate the abrupt increase in last uses that occurred in mid-September (also shown in Figure 39). From this the team can conclude that the previous day's weather did have some influence on the sudden increase in last AC use cases on the following day. This aligns with the team's findings from the exploratory data analysis (see Section 3.2.1.2).

The preferred specification is below, and the variables have the same definitions provided in Section 3.2.1.1:

The logit coefficients and AMEs of the team's preferred last use mode specification are shown in Table 11 and Table 12.

Table 11. Logit Coefficients for the Last AC Use Model

Term	Estimate	Std. Error	Statistic	p.value
AvgTemp	0.028	0.059	0.469	6.39E-01
AvgHBU	-0.011	0.006	-1.824	6.82E-02
MaxCDH	-0.136	0.046	-2.951	3.17E-03
AvgTemp*t-1*	0.089	0.080	1.114	2.65E-01
CDH*t-1*	0.029	0.002	13.268	3.54E-40

Source: Navigant

Table 12. AMEs for the Last AC Use Model

Factor	AME	Std.Error	z	p.value
MaxCDH	-0.00086	0.00031	-2.78	5.41E-03
CDH*t-1*	0.00018	0.00002	8.67	4.25E-18
AvgHBU	-0.00007	0.00004	-1.82	6.95E-02
AvgTemp	0.00018	0.00037	0.47	6.39E-01
AvgTemp*t-1*	0.00056	0.00050	1.11	2.66E-01

Source: Navigant

These AME values can be interpreted as follows:

- **MaxCDH:** For every 1°-hour increase in the daily maximum CDH value, the probability of an AC being used for the last time on that day decreases by 0.1 percentage points. Stated another way: warmer days are less likely to be the days of last use.
- **CDH*t-1*:** This is the 1-day lagged CDH. This AME indicates that for every one-degree-hour increase in the previous day's cumulative CDH, the likelihood of a AC being used for the last time increases by 0.01 percentage points. In essence: if the day before was warmer, the next day is (slightly) more likely to be the last AC use of the season.
- **AvgHBU:** The AME here is very small but indicates that warm days which are part of a series of warm days that are less likely to be the last AC use.
- **AvgTemp** and **AvgTemp*t-1*:** These are the daily average temperature and the previous day's average temperature. Both indicate that an increase in their values leads to an increase in the likelihood of a last use occurring on that day – 0.02 and 0.06 percentage points, respectively. Similar to the first use model's AMEs, these effects may seem counterintuitive, but it is likely that these are interacting with the CDH effects to only apply to temperatures less than 65. For temperatures above 65, the MaxCDH AME cancels out that of the AvgTemp and AvgTemp*t-1* terms (see the discussion in Section 3.2.2.2).

As mentioned, it is likely that additional model tuning and feature engineering could result in better understanding of what leads to last use and improve the predictive accuracy of the last use model. Some areas for future investigation include:

- Additional investigation into the joint interaction between customer mode selection (i.e., switching their thermostat to off) and weather-driven last use (i.e., thermostat mode is cool, but air temperatures are below setpoints).
- Additional lagged variables, possibly considering the net change between days (e.g., temperature swings; cold snaps) rather than the cumulative cooling values.
- Inclusion of HDH values—this is counter intuitive given that these are summer day, but small values of HDH may be highly predictive
- Utilizing advanced machine learning techniques and locally interpretable model-agnostic explanations (LIME) to detect complex relationships and non-linear dependencies.

A.5 Detailed Weather and AC Use Data Exploration Methodology

The team explored the relationship between both AMI and thermostat data with several temperature variables constructed from weather data available through the National Oceanic and Atmospheric Administration (NOAA).[61]

For 2014, all participants in the dataset were located in Worcester, Massachusetts, so the team limited its weather data to that most representative of Worcester.[62] The team checked participant ZIP codes to ensure that the group was not highly dispersed and verified that weather data from a single station would suffice. Windspeed, temperature, humidity, precipitation, and sky condition (e.g., cloudy, clear) were extracted at the 5-minute level. Weather station measurements are not always provided at regular intervals, so the data were averaged and interpolated to align with the 15-minute interval AMI data. The evaluation team aggregated this data to both the hourly and daily level for exploratory analysis and modeling.

The team attempted to acquire pollen data but was unsuccessful in finding a monitoring station anywhere in the Northeast with a historical record reaching back to 2014. The team had also hoped to obtain customer information beyond program tracking metrics, but this proved beyond the scope of this study. The possibility of using additional datasets should be left open for future work.

After aggregation, the evaluation team computed a variety of secondary weather characteristics, including Thermal Heat Index (THI), Cooling Index (CI), Heat Build-Up (HBU), and lagged weather variables.

Thermal Heat Index (THI): Thermal Heat Index (a.k.a. Sultriness, Heat Index, apparent temperature, or temperature-humidity index) is a temperature index accounting for the extent to which humidity aggravates physiological effects of high temperature; it is derived from several biometeorological equations and studies. In other words, it is the apparent hotness of the air to humans, in degrees, adjusted upward at high humidity and temperature. While not theoretically based in building heat transfer, the evaluation team has found in prior work that the corrections to temperature within THI align it with AC use better than temperature alone, particularly when customer behavior is taken into account. The equation for THI is:

$$\mathrm{THI} = \mathrm{c}_1 + \mathrm{c}_2\mathrm{T} + \mathrm{c}_3\mathrm{R} + \mathrm{c}_4\mathrm{TR} + \mathrm{c}_5\mathrm{T}^2 + \mathrm{c}_6\mathrm{R}^2 + \mathrm{c}_7\mathrm{T}^2\mathrm{R} + \mathrm{c}_8\mathrm{TR}^2 + \mathrm{c}_9\mathrm{T}^2\mathrm{R}^2$$

where *R* is the relative humidity in units of percentage and *T* is the ambient dry-bulb temperature.[63]

In , the black lines represent the daily averages for temperature and humidity, with the surrounding green band indicating the daily minimum and maximum values. On the temperature graph, the THI is also plotted behind the temperature, in red. On days where the THI maximum exceeded the temperature maximum by more than 3°, a dashed vertical line has been drawn to align with humidity. For most of the summer, THI does not exceed temperature because the corrections take effect above 77°F and temperatures in the summer of 2014 rarely exceeded this value.

[61] NOAA weather data is available through the following website: https://www.ncdc.noaa.gov/data-access

[62] Worcester, Massachusetts weather data was obtained from the NOAA website for weather station WBAN:94746

[63] The c_x constants are the results of a polynomial fit to a calculated Heat Index table and can be found in a Technical Document published by the National Weather Service, available at: http://www.wpc.ncep.noaa.gov/html/heatindex_equation.shtml

Figure 56. Daily Average Humidity and Temperature for 2014

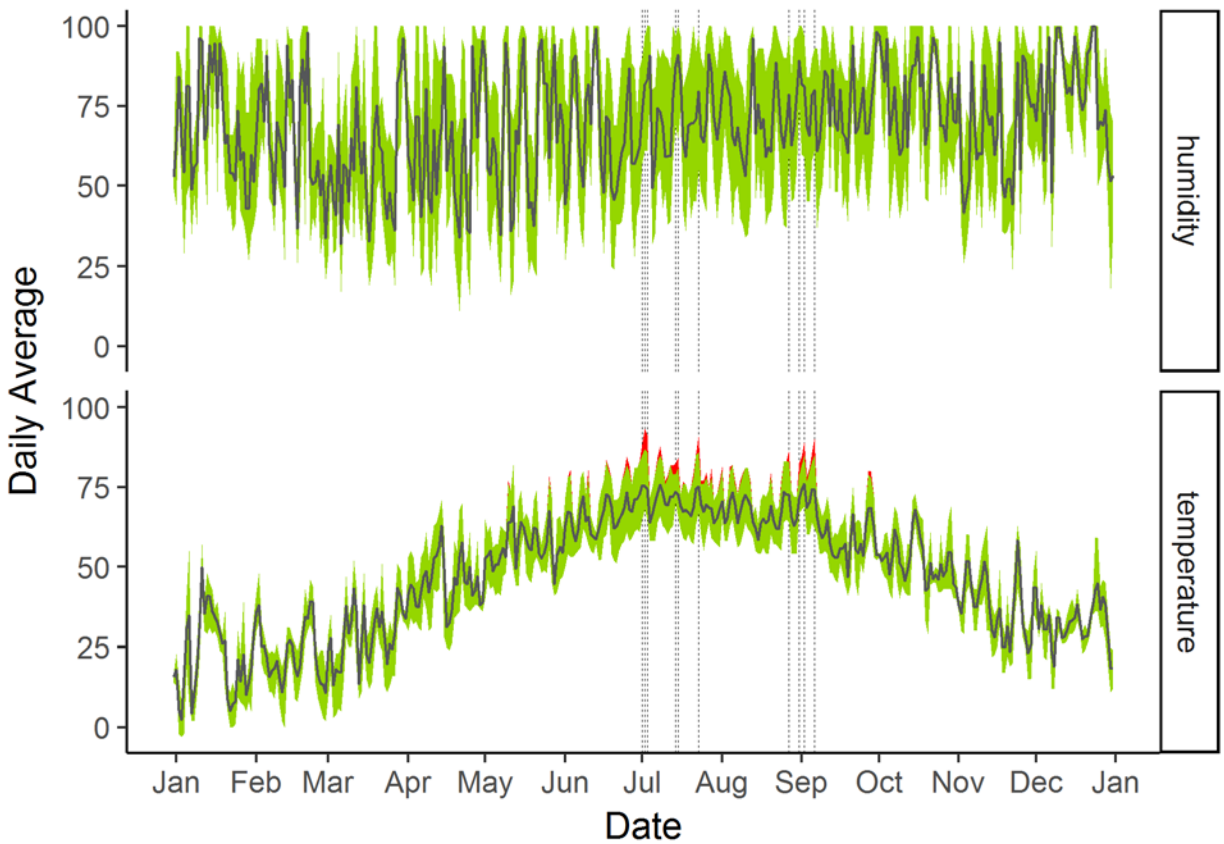

Source: Navigant

Cooling Index (CI): The evaluation team calculates cooling index based on THI, similar to the calculation of CDH. The team assumes a balance point of 68°F, with temperature counting up from zero above it and remaining zero below it. The representation of heat through a CI variable allows heat to correlate with AC use without the use of censored regression model—the uniform zero value below 68°F matches the behavior of the AC, which stops running around this temperature.

Heat Buildup (HBU): HBU captures the recent history of hot weather. It is the sum of the temperature history, with each step backward from the time in question weighted less than the step prior. The formula for HBU is:

$$\mathrm{HBU_0} = \sum_{t=1}^{\infty} (0.96)^{t} \cdot (\mathrm{CI}_{0-t})$$

where t is the number of hours prior to current time and CI_{0-t} corresponds to the CI t hours before the current time. This quantity is useful for capturing behavioral effects related to long periods of hot weather or short spikes in the temperature, as HBU will be high in the former case and low in the latter.

Lagged variables: A number of lagged variables are included in the models. In daily models, 1- and 2-day lags for temperature and humidity variables were tested. In the hourly models, lagged 3-hour averages were employed to account for heat accumulation in buildings over time.

As with the weather data, the evaluation team explored AMI data at the 15-minute, hourly, and daily level to investigate trends with respect to other variables. First, the team sampled individual customers and aggregated their hourly load shapes by month to search for relevant usage patterns. Then the team moved down to a more granular level, examining individual daily load shapes to explore the direct effect of weather on selected customers. Once the team had gained familiarity with individual customer usage patterns, it aggregated the customer data to search for population-wide features and relationships with weather.

Based on a visual inspection of load shapes, nearly half of customers do not appear to have usage patterns strongly correlated with temperature (THI), and many customers appear to use AC sparingly.

A few customers have load shapes that indicate their cooling demand maximizes at 100% of their system capacity during hot weather. provides a good example of this type of AC user. The two plots shown in the figure are constructed from a sample of daily AMI load shapes for a single customer. On the left, only days with maximum THI between 65°F and 70°F are sampled versus on the right, which only includes days where THI peaked above 80°F. Average hourly electricity use in kilowatt-hours (kWh) is represented by the black lines on either side. ACs operating below 100% capacity cycle over time to minimize starting and stopping their compressors. That behavior is evident in the load shape, which oscillates with increasing frequency as the temperature increases, before flattening out at 100% AC demand. In other words, the plot on the left is an example of a system that does not max out because the maximum temperature those days was not high versus the plot on the right of the same system that does max out on the days shown due to hot temperatures and the AC cycling behavior described.

Not all customers reach 100% AC demand on the hottest days. This may indicate that those customers have either appropriately sized or oversized systems. The team also noted during its data exploration that only a handful of days had truly hot weather in 2014. As a result, there are a limited number of days the team was able to observe people’s reaction to very hot weather. However, the granularity of the AMI data and large sample sizes still allow the team to draw strong conclusions about people’s reaction to weather.

Figure 57. A Sample of Daily AMI Load Shapes for a Single Customer

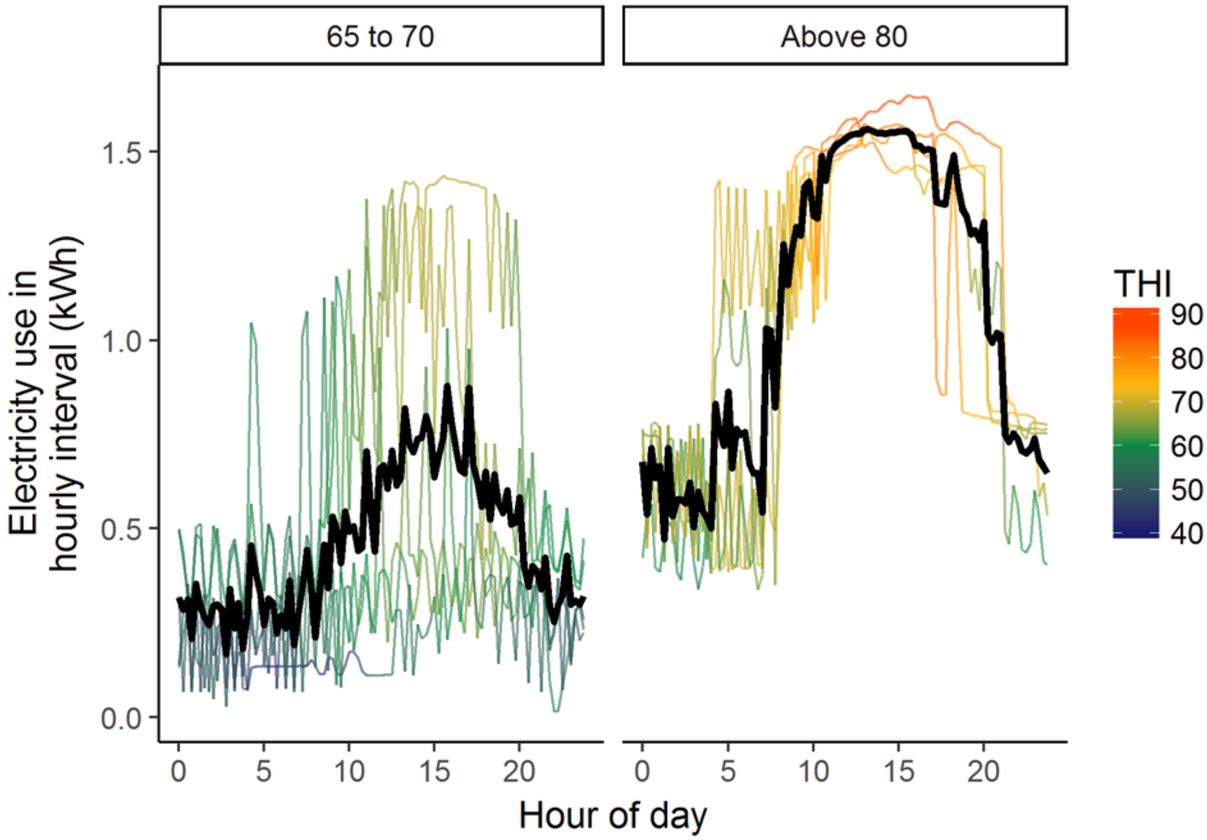

Source: Navigant

The complex cycling behavior of AC systems, combined with a large variation in the timing of cooling by customers with different AC use behavior, makes modeling the effects of weather on electricity use at the 15-minute level problematic. For example, some customers turn their AC system on as soon as they get home, while other customers follow a set schedule that may or may not correspond with occupancy. Modeling at that level of granularity runs the risk of introducing behavioral and system effects that could bias the results. More importantly, this level of granularity would not improve the team's ability to quantify the relationship between weather variables and AC use. For these reasons, the evaluation team used aggregated daily models.

A.6 Detailed Regression Methodology and Robustness Checks—AC Use throughout the Season

For the regression analysis, the team used only late spring, summer, and early fall data to avoid weather-dependent usage related to heating. This constraint reduced the sample size from approximately 16,000 to just over 15,000.[64] The team approached the regression analysis in two stages.

- Individual regressions to identify those with AC

[64] The team started with exactly 16,915 homes, 1,654 of which did not meet this criteria, leaving a final sample size of 15,261 homes.

- The team ran two simple regressions on each participant with the 15-minute interval data to extract the effect of CI on electricity consumption. For robustness, the team used the results of both to determine which customers have AC systems.

- Pooled regression to quantify how weather variables affect AC use
    - The team subset the full dataset to only those customers with AC to estimate a regression model across all AC users to quantify the relationship between their AC use and weather variables.

Individual Models – Identifying Customers with AC

The first regression form, which is referred to as the single-stage model for individual regressions, was:

$$kWh = \beta_0 + \beta_1 \overline{CI}_{t,t-3} + \sum \beta_{m,h} month \cdot hour$$

where $\overline{CI}_{t,t-3}$ is the average CI across the three prior hours, *month* is a categorical variable corresponding to the month of the year, and *hour* is a categorical variable corresponding to the hour of day. The interacted *hour* and *month* terms absorb weather-insensitive usage at each hour of the day for any given month, while the CI term quantifies weather-sensitive usage.

In the second approach for individual regressions, which is referred to as the two-stage model, the team first subtracted out baseline usage to arrive at adjusted kWh usage values for each household. To do this, the team modeled baseline usage by estimating the following model only during times when $\overline{CI}_{t,t-3}$ was between 60°F and 65°F:

$$kWh = \beta_0 + \sum \beta_{w,h} dow \cdot hour$$

where *dow is* a categorical variable corresponding to the day of the week and *hour* is a categorical variable corresponding to the hour of day. This model was used to predict baseline usage for time intervals across all summer months, and then that prediction was subtracted from the measured usage. The resulting adjusted usage should only contain long-term time effects and weather-dependent effects. The team then used this adjusted energy usage value as the dependent variable, which was modeled using the following:

$$kWh_{adj} = \beta_0 + \beta_1 \overline{CI}_{t,t-3} + \sum \beta_m month$$

where there are no new terms in the model.

Two different views of these results are provided in and . The first figure plots the effect of temperature (CI) from both the single-stage and two-stage individual regression models together. Visual inspection shows the results are well aligned, though not identical. shows the distribution of temperature-related demand implied by the regression results, allowing an assessment of the distribution of weather-related usage.

Figure 58. Comparison of the Predicted Effect of Temperature on Cooling Load from the Single-Stage and Two-Stage Models

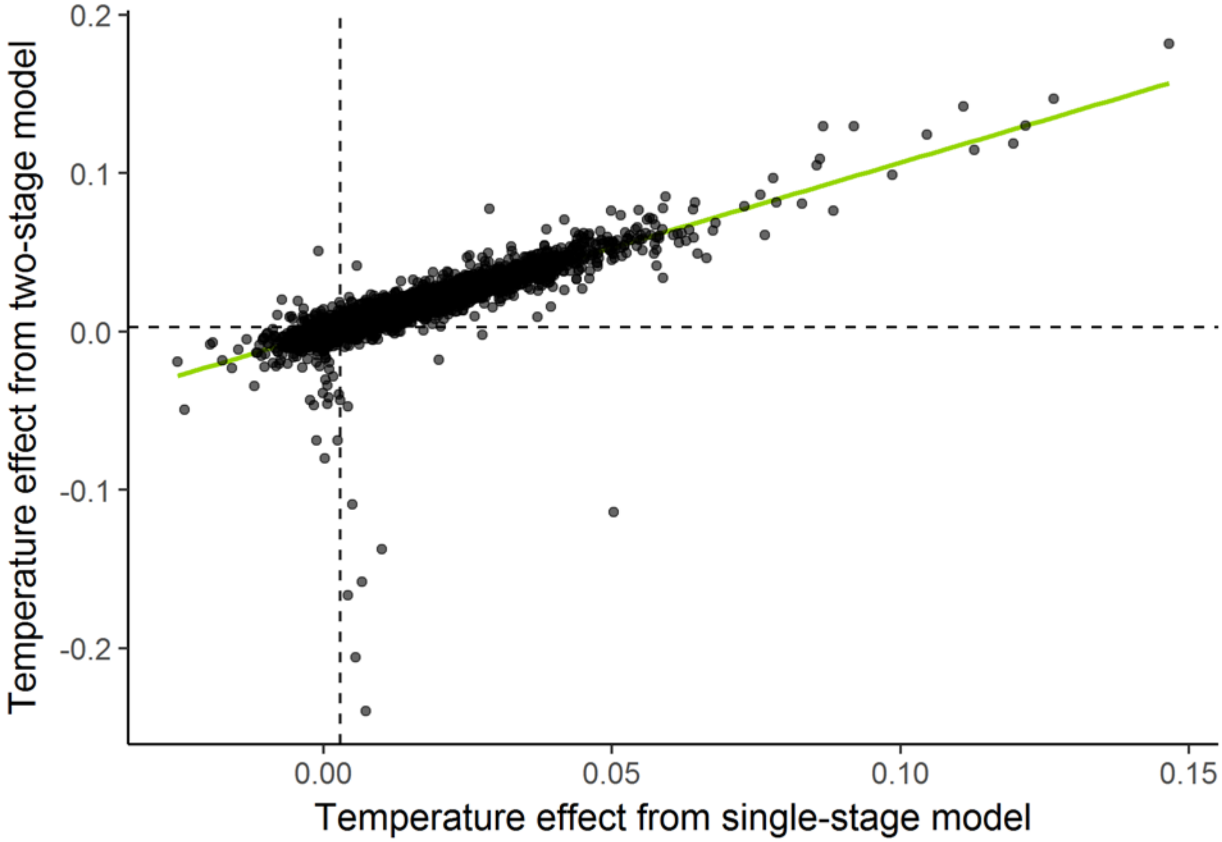

Source: Navigant

Figure 59. Number of Accounts with a Given Predicted Effect of Temperature on Electricity Consumption

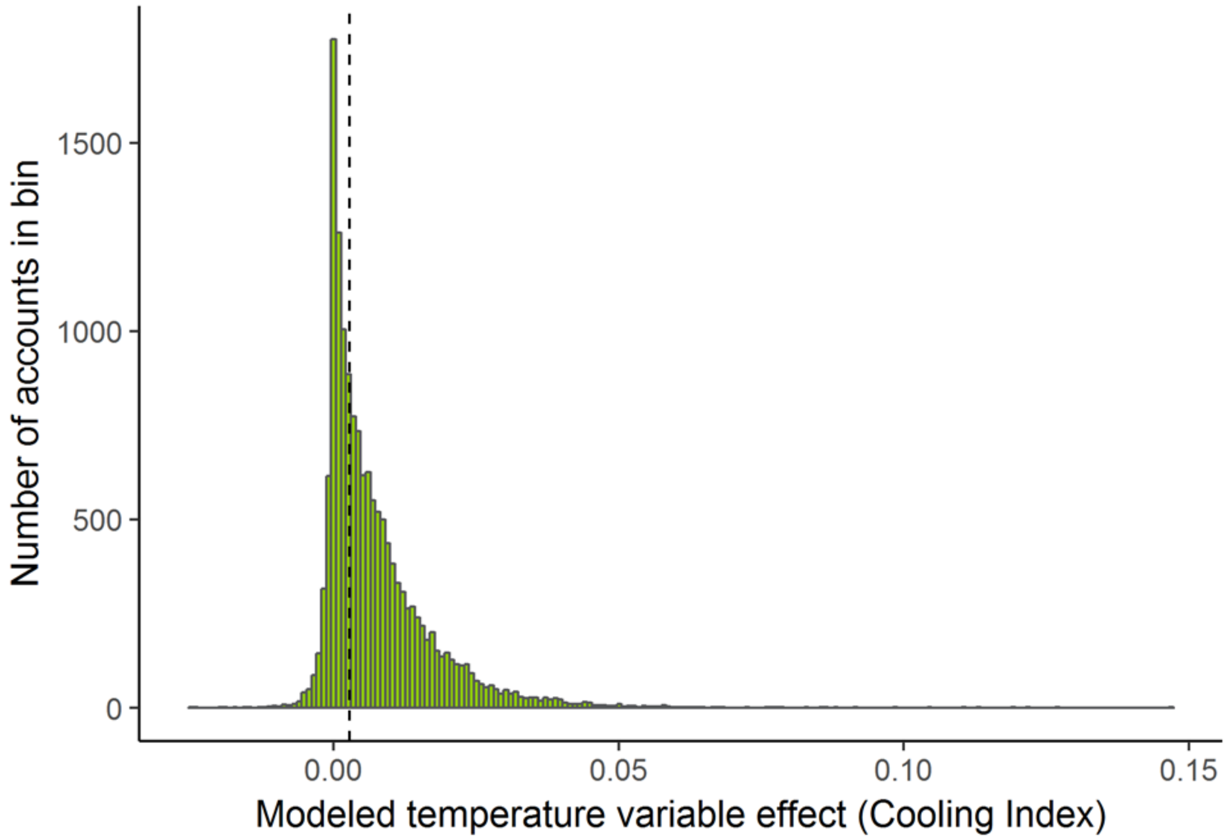

Source: Navigant

For a robust approach to defining customers in the AMI data as having AC, the evaluation team used a double-threshold approach, requiring customers' daily demand derived through *both* regression models to exceed 250 kW to be defined as having AC.[65] Customers satisfying this requirement are represented by the upper right quadrant of . Although the team observed no clean breakpoint in energy demand estimated by these regressions (see), the threshold was set at 250 kW (shown with a dotted line on the figure) because that value is half the power consumption of a typical small window AC.[66] This divided the sample of approximately 15,000 homes into roughly 9,000 with AC and 6,000 without.[67] The evaluation team only used those with AC in the analysis that follows.

The team next compared load shapes across the 12 months of 2014 for the two groups (see) to provide a further robustness check that the segmentation of customers with and without AC was successful. The team plotted the average load shape for each group across all months in 2014. While load shapes for both groups remained similar from April through October, the group defined as having AC developed a strong afternoon peak from June through September, as expected. In July, average demand at 4 p.m. for

---

[65] This value corresponds to an effect size for $\overline{CI}_{t,t-3}$ greater than or equal to 0.00284 for both regressions.

[66] Assuming 5,500 Btu/h and energy efficiency ratio (EER) equal to the federal minimum (11), converted with the equation $kW = {}^{BTU/hr}/_{EER}$

[67] The exact final sample sizes were 8,994 with AC and 6,267 without AC.

the group with AC was 1.1 kW more than the group without, while this difference was only 0.2 kW in April. This evidence strongly suggests the individual regressions successfully segmented the customers into AC users and non-users.

Electricity consumption for the group with AC was slightly higher across all months, which could be due to a variety of factors that could not be investigated with available data. These factors include home size, greater reliance on the HVAC system in all weather, the presence of a reversible heat pump, or a correlation between opting for AC and greater penetration of other electricity-consuming devices. Nonetheless, the load shape analysis clearly shows the double-threshold individual regression approach successfully segmented those with AC and those without.

Figure 60. Average Daily Electric Load Shapes for Customers Defined as Having vs. Not Having AC

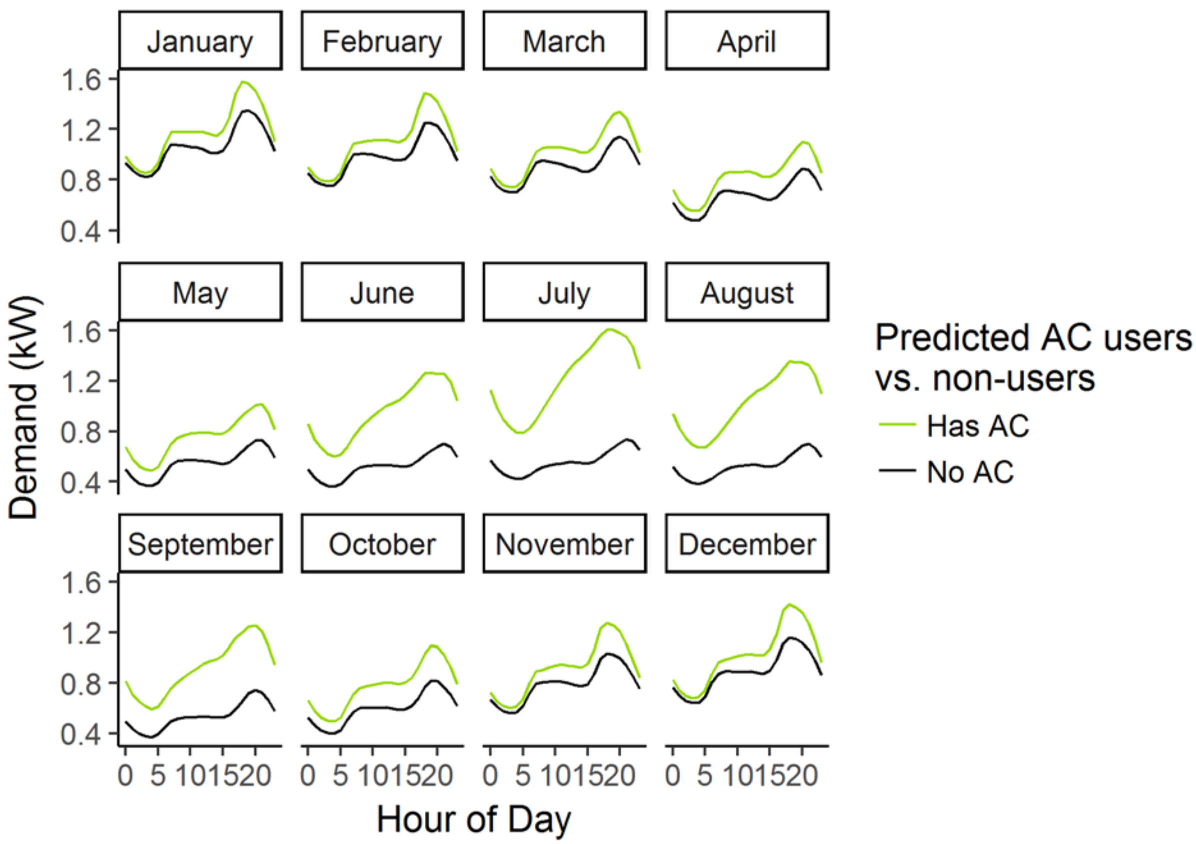

Source: Navigant

While it is always possible that a small proportion of households with no AC passed these multiple screening hurdles and were included in the regression analysis moving forward, the likelihood is small given the team's robust approach to screening. This guarantees the cleanest, highest quality results when quantifying the relationships between weather variables and AC use for this population.

Pooled Model – Estimating Weather Effects on AC Use

The evaluation team's overall goal in this section of the research was to quantify causal relationships between weather variables such as heat, HBU, and humidity with AC use. After sub-setting the AMI data to the set of customers defined as having AC systems, the team had a sample size of roughly 9,000 households to use in its main regression analysis.

Due to the large sample size and extremely granular AMI data, the evaluation team had the ability to develop a detailed fixed effects regression model to explain weather's impact on AC use. The team began with all the weather variables that made intuitive sense in the model to capture heat and humidity. However, each of these can enter the model in many forms (e.g., lagged variables, HBU, temperature, CI, etc.). To ensure the models perform well and are not overfitted, the team performed cross-validation as a robustness check. [68] Cross-validation involved training the model on different subsets of the data, then calculating the error from prediction on the remaining data over many iterations. The team then calculated the average error produced by each model variation to compare performance ensure the model performs well out of sample (i.e., is generalizable). The results of this cross-validation analysis for six of the model variations are shown in . Though the team's preferred model specification is that labeled "*Last Model,*" these results suggest all model variants performed similarly well (low error) and the narrowness of the boxes shows low variability across different data subsets. The blue dots shown depict model performance without holding any data back as a training set, and their closeness to the box plots adds to the robustness of the results.

[68] In total, the team tested 25 fixed effects model variants and the full box plot results of the cross-validation analysis for these 25 models are available in Appendix A.6 . The team subjected each prospective model to 200 iterations of 20% holdback cross validation, wherein 80% of the data was used to train the model and the RMSE of prediction on the remaining data is the diagnostic metric.

Figure 61. Box Plots of RMSE from Cross-Validation for Top X Models

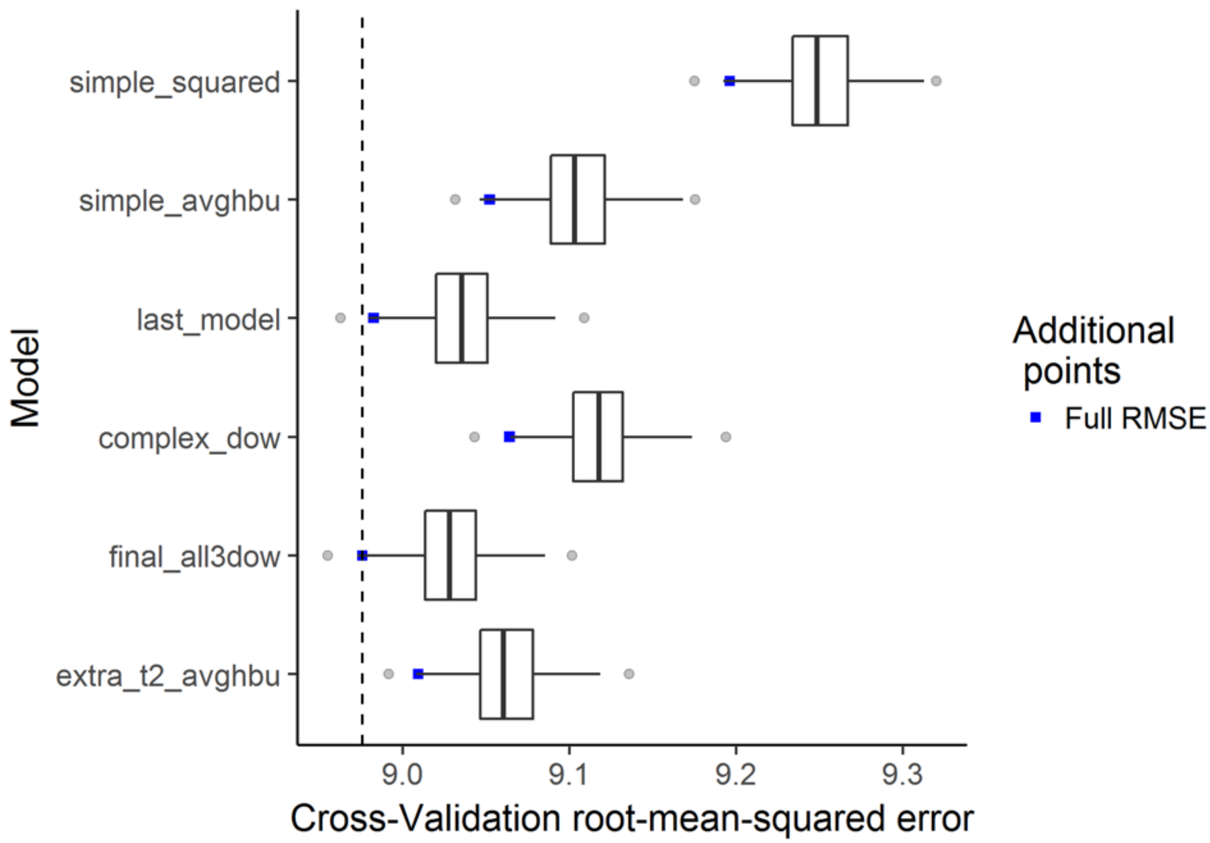

Source: Navigant

The team's preferred model is as follows:

$$kWh_{it} = \alpha_i + \beta_1 MaxCI_t + \beta_2 AvHBU_t + \beta_3 AvHum_t + \beta_4 MinTemp_{t-1} + \beta_5 MinTemp_{t-2} + \beta_6 MaxTemp_{t-1} + \beta_7 MaxTemp_{t-2} + \beta_8 AvHum_{t-1} + \beta_{9d} \sum_d DOW_t + \beta_{10d} \sum_d DOW_t \cdot MaxCI_t + \beta_{11d} \sum_d DOW_t \cdot AvHBU_t + \varepsilon_{it}$$

where:

kWh_{it}	= kWh electricity usage by individual *i* on day *t*
α_i	= Individual-level fixed effects
$MaxCI_t$	= The maximum cooling index reached on day *t*
$AvHBU_t$	= The average heat buildup across day *t*
$AvHum_t$	= The average relative humidity across day *t*
$MinTemp_{t-1}, MinTemp_{t-2}$	= The 1- and 2-day lagged minimum daily temperature
$MaxTemp_{t-1}, MaxTemp_{t-2}$	= The 1- and 2-day lagged maximum daily temperature
$AvHum_{t-1}$	= The 1-day lagged average daily humidity
DOW_t	= A categorical variable for the day of the week (e.g., Monday, Tuesday)
ε_{it}	= Individual- and time-specific error term

The evaluation team used the preferred specification to estimate weather effects on daily energy usage for the 9,000 customers in the sample defined as having AC, and the model performed well.[69]

For robustness, the team plotted the predicted values from the model against the actual values (shown in green and black in Figure 62), and these values aligned well. As shown in , the model correctly tracks energy usage peaks due to cooling demand in hot weather periods. It also remains in good agreement with baseline usage into the shoulder seasons. While the team observed some slight shape to the residuals, they are clustered around zero and the apparent structural variations are small across the different months, suggesting the model performs well throughout the summer.

Figure 62. Predicted and Actual Average Daily Use, with Residuals

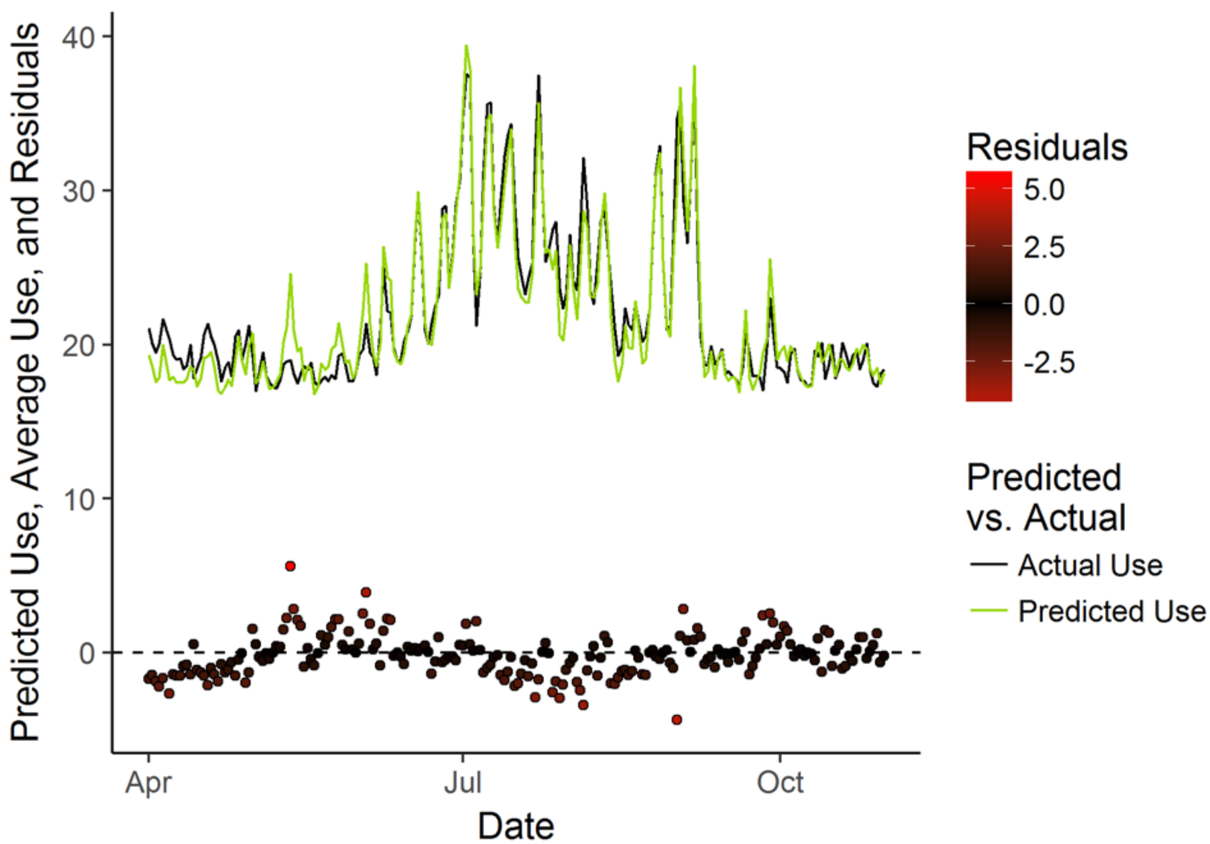

Source: Navigant

A.7 Robustness Checks on Modeling Weather Effects on AC Use throughout the Season

As a robustness check on the model estimating weather effects on AC use throughout the season, the team tested the model's ability to predict AC runtime from the thermostat data rather than kWh electric usage. While Appendix Section A.8 direction models AC runtime as a function of AMI data, this model variation is intended to leverage having customers that exist in both the AMI and thermostat datasets,

[69] The model successfully explained a large portion of variation in the data ($R^2 = 0.75$).

which can be used to assess how similarly models perform for these same customers when electricity use versus AC runtime are the dependent variables. There were 260[70] customers available for this analysis who were in both the cleaned thermostat and AMI datasets. For these 260 customers (a subset of the main 2014 regression analysis sample), the team started with its preferred regression model, but ran it first with AMI-based electricity use (kWh) as the dependent variable and then alternately with thermostat data-based AC runtime as the dependent variable.

This version of the model required a small modification to account for differences in the nature of the data. Because thermostat runtime data represents the amount of time the AC runs and cannot go below zero or above 100%, the model must account for the truncated nature of the data. The evaluation team employed a censored regression Tobit model for this purpose, which properly accounts for the dependent variable not exceeding those values. As shown in , the residuals from this Tobit model are centered around zero as expected, but the residuals from the Linear model are consistently above zero due to truncation of data points near zero runtime. Therefore, the team used the Tobit model for its AC runtime modeling.

Figure 63. Comparison of Residuals between Linear and Tobit Models

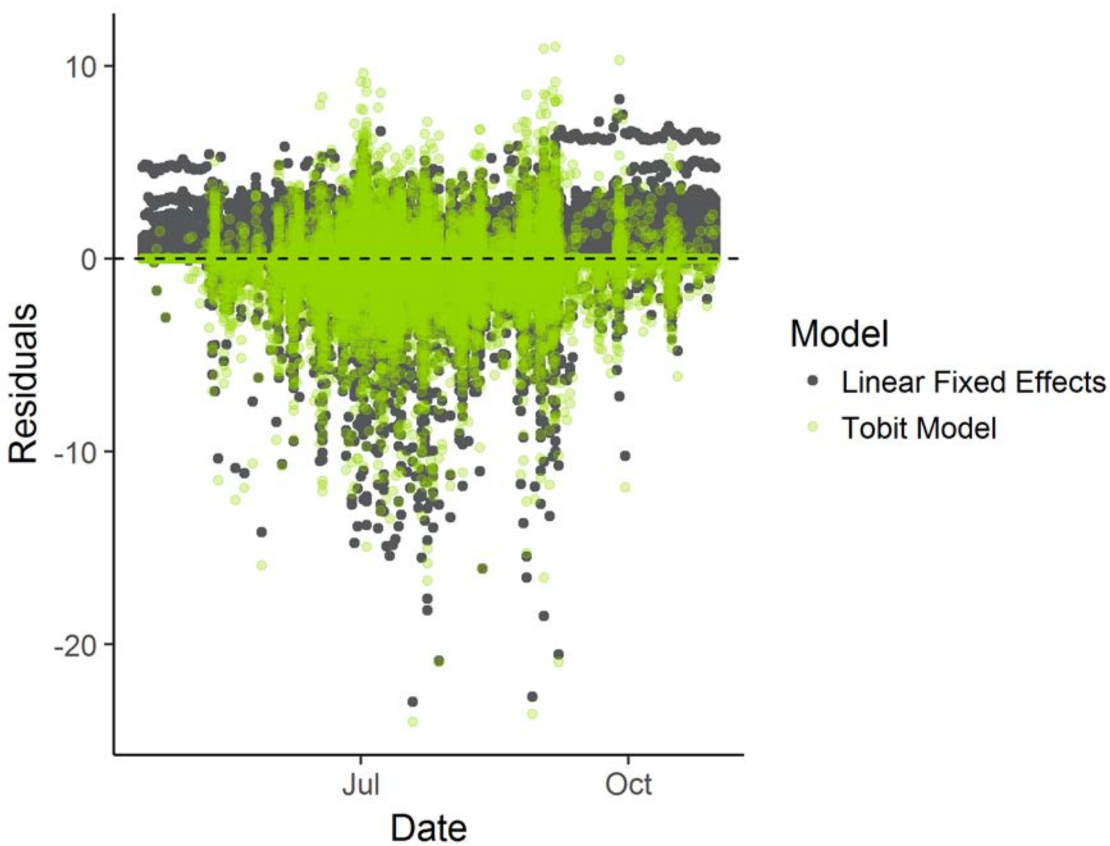

Source: Navigant

[70] This number is larger than the 54 homes used in the First AC Use model because that model was limited to customers that had installed their thermostat early in the season and had AMI data, while the number of installed thermostats, and hence the overlap between the thermostat and AMI datasets, increased throughout the summer.

Below are the regression output results for some of the key variables in the model predicting AC use as a function of weather, when estimated using all 260 customers with thermostat data for the summer first:

- With electric use from the AMI dataset as the dependent variable (Table 13)
- Then with AC runtime from the thermostat dataset as the dependent variable, on the same set of 260 customers (Table 14)

Table 13. Regression Results from the AMI-based Electric Use vs. Weather Model, Run on Only Customers with Thermostat and AMI Data

Term	Estimate	Std. Error	t-value	p-value
maxci	0.509	0.053	9.6	< 1e-9
avghbu	0.082	0.007	12.1	< 1e-9
avghum	0.035	0.005	6.7	2.53e-11
mintemp_l1	-0.004	0.020	-0.2	8.56e-1
maxtemp_l1	-0.079	0.016	-5.0	7.13e-7
avghum_l1	-0.004	0.006	-0.7	5.03e-1
mintemp_l2	0.002	0.016	0.1	8.95e-1
maxtemp_l2	-0.018	0.014	-1.4	1.75e-1

Source: Navigant

Table 14. Regression Results from the Thermostat-based AC Runtime vs. Weather Model, Run on Only Customers with Thermostat and AMI Data

Term	Estimate	Std. Error	t-value	p-value
maxci	0.383	0.021	18.1	< 1e-9
avghbu	0.005	0.003	1.6	1.09e-1
avghum	0.005	0.003	1.7	9.29e-2
mintemp_l1	0.116	0.012	9.3	< 1e-9
maxtemp_l1	0.115	0.011	10.7	< 1e-9
avghum_l1	0.000	0.003	0.1	8.87e-1
mintemp_l2	0.075	0.010	7.7	< 1e-9
maxtemp_l2	-0.017	0.009	-1.9	5.38e-2

Source: Navigant

Figure 64 and Figure 65 show predicted electric energy use (kWh) and predicted AC runtime, respectively, based on these two models.

Figure 64. Average and Participant-Level Predicted Electric Use for Customers with Thermostats, 2014

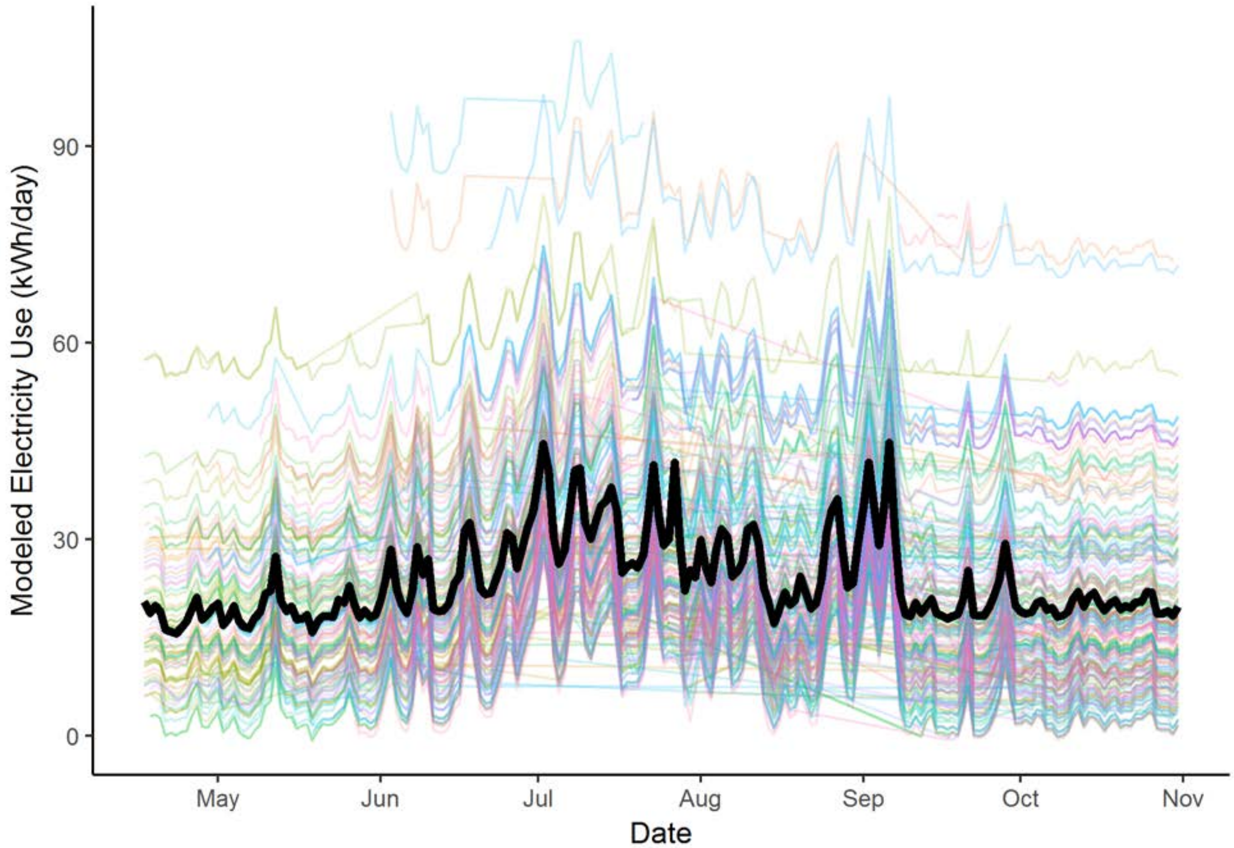

Source: Navigant

Figure 65. Average and Participant-Level Predicted AC Runtime for Customers with Thermostats, 2014

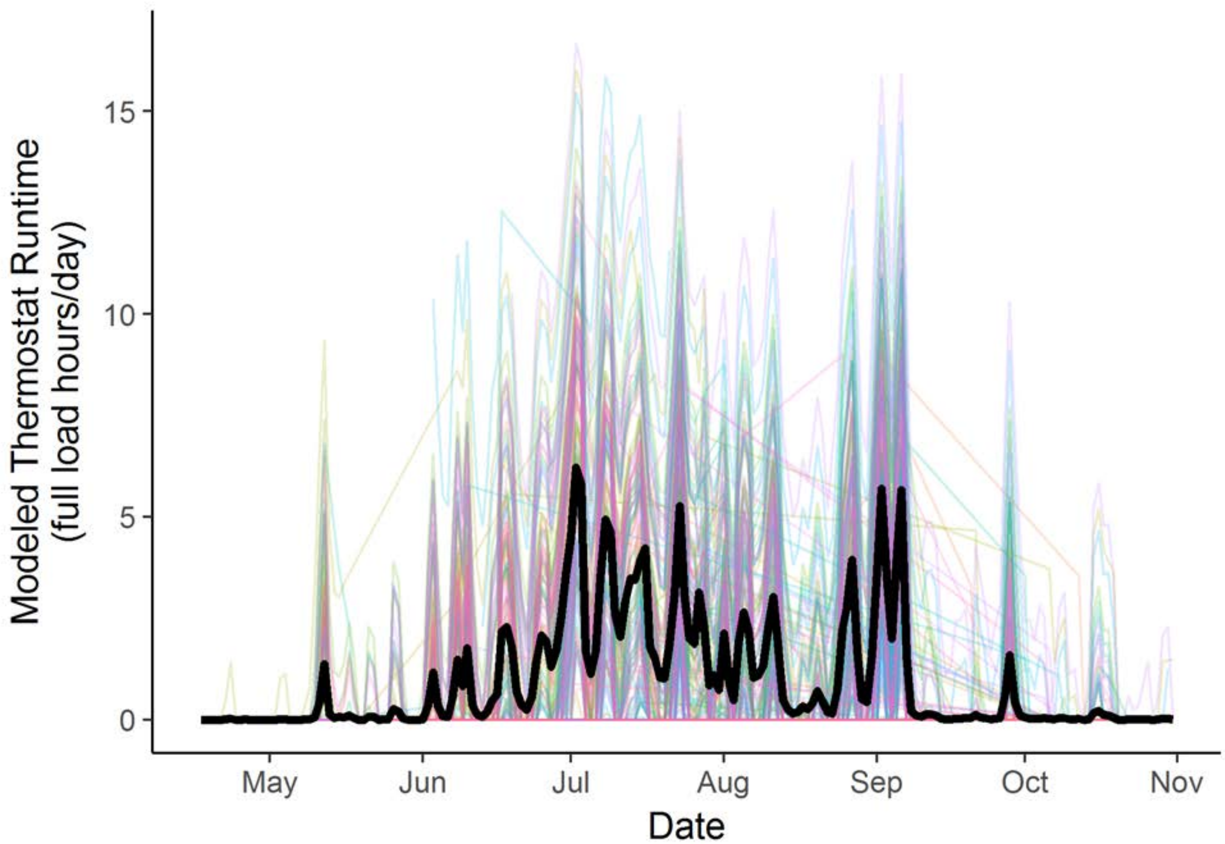

Source: Navigant

Although both models had fewer significant terms than the original model due to the much smaller sample size, they produced similar results.[71] The models show peaks both in electricity usage and AC runtime during key summer hot spells in July and September. Both models indicated that Monday and Tuesday had high weekday consumption, while Thursday had low weekday consumption. Sunday had higher consumption than Saturday in both models and higher consumption than Tuesday in the thermostat model. Despite overall agreement, some important differences were revealed through the models. The thermostat model did not predict any runtime for a warm normal week, suggesting that electric consumption at lower warm temperatures may be driven by non-AC weather-based use. Additionally, the terms for HBU and average humidity were not significant in the thermostat model though they were in the AMI-based electricity use model. Although it is difficult to determine the cause, this difference might indicate those terms are driving non-AC but still weather-based usage. Examples might include standalone fans or increased runtime from heat pumps that are not central AC, such as window AC units or refrigerators. As a result, for the Massachusetts PAs' purposes, the forecast maximum CI may, therefore, be the term to watch for planning DR.

[71] The team has omitted complete regression outputs for both models in the main text for brevity, but they are available in Table 21 and Table 22 in Appendix A.7.

As a final robustness check on the model's results, the team investigated predictions for a typical meteorological year (TMY) and for the years 2010-2017 based on data from the same Worcester weather station used in the main analysis. compares the predictions for these different weather years:

Figure 66. Predicted Electric Use Using 2014, 2010-2017, and TMY Weather Data

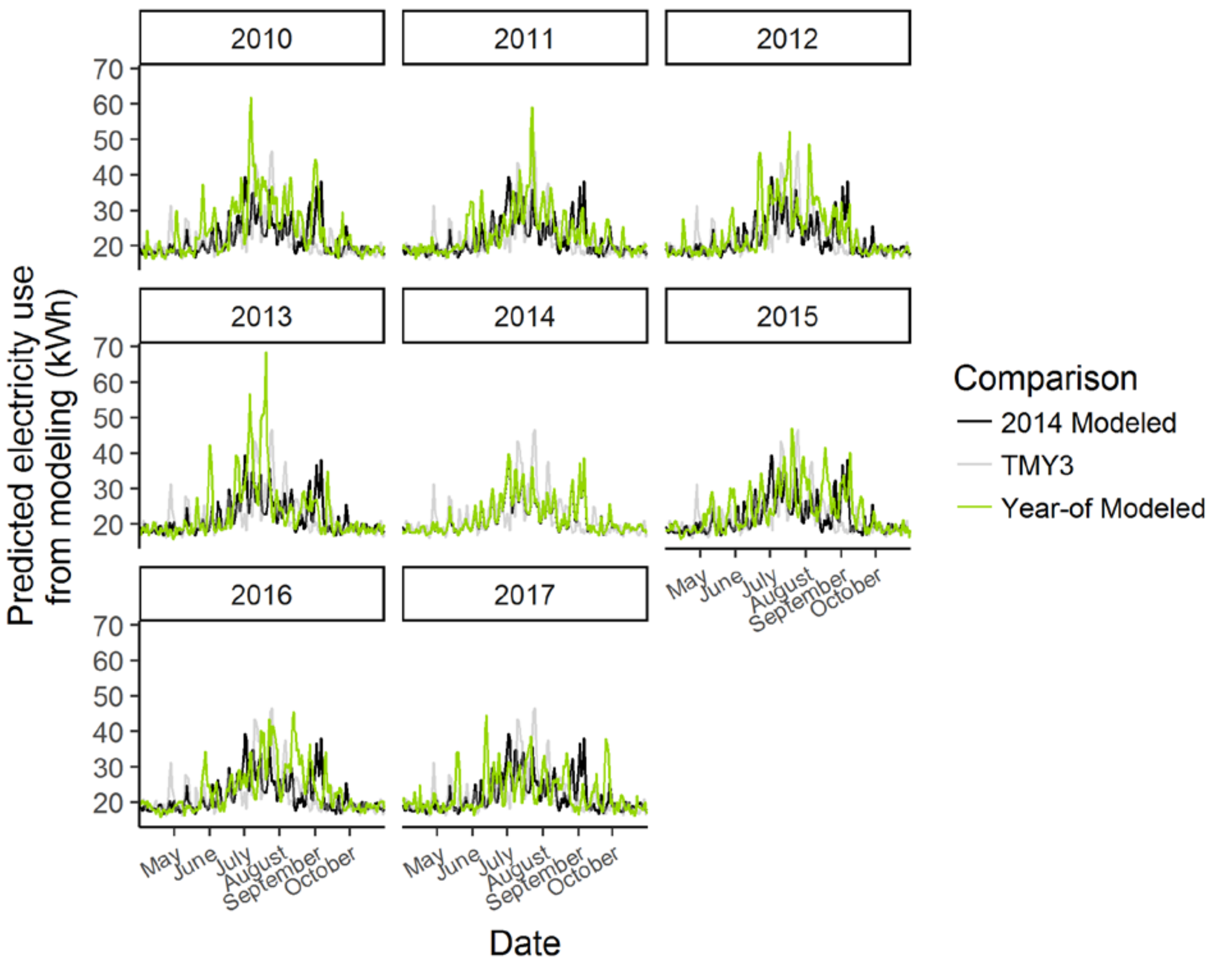

Source: Navigant

The analysis in this study is based on 2014 data, so it is important to understand the generalizability of the results to years with different weather. TMY and the 2014 data had similar weather conditions and, therefore, similar modeling results. The 2014 summer (the analysis year) was the mildest summer in the 8-year span from 2010 to 2017, which might suggest problems for using the evaluation team's model to predict AC use under much hotter weather conditions. However, reassuringly, most predicted values are reasonable for the 8-year span, with no extreme low or high predictions, and with extreme highs corresponding to extreme weather as expected. This lends credence to the idea that these model results are generalizable. There are a few exceptions. For example, on one day in 2013, the model predicts an average usage of 70 kWh, which is unlikely to be accurate as it is nearly twice the upper bound of the prediction for 2014. However, this appears to be the exception to the rule; generally, the model appears to make reasonable predictions under different actual weather-year scenarios.

A.8 Predicting AC Runtime Using Weather and AMI Data

In many evaluations, AC runtime needs to be converted to electricity use derived from billing/AMI data or vice versa. The evaluation team felt, given the wealth of information at its disposal, it might be helpful to the PAs for the team to provide a basic exploration of converting electricity use to AC runtime, accounting for weather. This would provide the PAs and future researchers with a benchmark method for making this often-needed conversion. To the extent that the customers in this study are a representative subset of the larger population of AC users, the coefficients from these models can be used to produce a model that can predict an unknown customer's AC runtime for given weather conditions and energy consumption. This could be useful, for example, in identifying high AC use customers without the need for smart thermostat data

Using all three datasets (AMI, thermostat, and weather), the team tested several model formulations that could aid in predicting AC use at the individual level given AMI-based electricity usage and weather forecasts or, alternatively, historical averages.

As with the other models, the evaluation team first looked at the overall relationship between some of the likely predictors and the length of AC runtime. This is shown in Figure 66. For several of the subplots, there is a strong clustering of values on the horizontal axis. This indicates there are many days during which the predictor variable was non-zero (e.g., daily kWh), but there were no hours of AC runtime for that specific customer. This may be due to the cool summer and the inclusion of customers which had no or very little AC use during the summer. Also evident is the possibly linear, or piece-wise linear, relationship in some of the variables. The average temperature subplot (second row, first column) shows AC runtime increasing linearly for temperatures above 60°F.

Figure 67. Relationship between AMI and Weather Data and AC Runtime

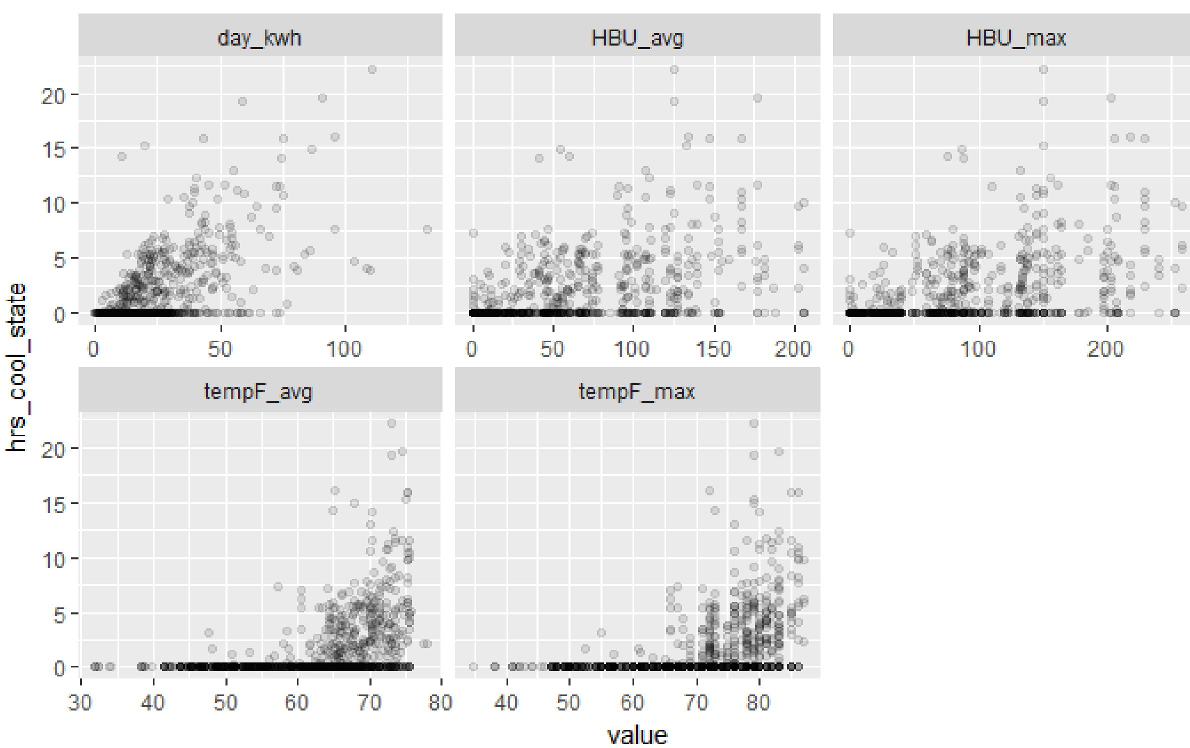

Source: Navigant

The evaluation team considered two categories of models to predict the AC runtime using the AMI and weather data, both hourly and daily models. For each category, the team used a censored Tobit fixed effect regression due to the discrete upper and lower bounds on AC runtime (zero to 24 hours). The team compared the RMSE results from several different model specifications within each category. However, given the small sample size, the hourly-level results were not statistically meaningful. As a result, a daily model was chosen.[72]

The team used customer-level data from both the AMI and thermostat datasets, aggregated to the daily level (e.g., daily total AC runtime and daily total kWh), in these models. As in the development of other regression models in this study, several model variations were tested using cross-validation to determine the most appropriate form for weather terms (e.g., CI, HBU, temperature, lagged variables) to enter the model. Of the model formulations compared within this category (Figure 68), the one based on total daily kWh, daily maximum HBU, and daily maximum temperature demonstrated the best predictive performance and is the team's preferred specification.[73] This preferred model specification is given below:

$$hours\_cool_{it} = \alpha_i + \beta_1 kWh_t + \beta_2 MaxHBU_t + \beta_3 MaxTemp_t + \varepsilon_{it}$$

where:

α_i	= Individual-level fixed effects
kWh_{it}	= kWh electricity usage by individual *i* on day *t*
$MaxHBU_t$	= The maximum heat buildup reached on day *t*
$MaxTemp_t$	= The maximum temperature reached on day *t*
ε_{it}	= Individual- and time-specific error term

The coefficients and fit statistics from the team's preferred model are shown in .[74]

Table 15. Regression Coefficients for Customer-Level Daily AC Runtime Model

Term	Estimate	Std. Error	Statistic	p-value
DailykWh	0.16	3.26E-03	48.87	0.00E+00
MaxHBU	0.03	7.88E-04	32.18	2.40E-222
MaxTemp	0.11	6.47E-03	16.30	2.19E-59

Source: Navigant

[72] Further results from the customer-level hourly model, including graphs showing box plots of the RMSE derived through cross-validation analysis are provided in Appendix A.7. However, these are not included in the main body of the report for brevity.

[73] By best predictive performance, as with the other regression model cross-validation results, the team implies the box plots are narrow, indicating the model is fairly generalizable when tested across various subsets of the sample and the overall error is sufficiently low. The mean RMSE for this model is slightly less than 1.25 hours (theoretical maximum: 24 hours). However, as the mean value of the daily AC runtime is 1.2 hours, the error as a percentage of the typical values is relatively large. This provides a caveat to drawing overly general conclusions from the results and again suggests the need for further research using a larger dataset.

[74] Because this model has been developed for the purpose of use in future research to better translate AMI-based energy use data to AC runtime in Massachusetts, the team does not discuss individual coefficient interpretations. However, to use this model and its outputs in future work, the team has made the full regression output, including intercepts and fixed effects coefficients, available in Appendix A.8. In this manner, future researchers can use this full set of coefficients along with the model specification and actual weather and energy use data to predict AC runtime.

Individual coefficient interpretations are not important in this case, although they all have intuitive magnitudes and directions. The purpose of this exercise is to provide the PAs and future researchers this model as a starting place to use AMI data combined with weather data to predict AC runtimes in research efforts where thermostat data is not available. The evaluation team feels these results are a starting benchmark for this translation, but they are preliminary, based on a relatively small sample size, and encourage future in-depth research to refine this translation.

Figure 68 shows cross-validation results for the main models tested for this exercise, with the team's preferred model shows as the second from the top.

Figure 68. Cross-Validation RMSE Results for the Customer-Level Daily AC Runtime Models

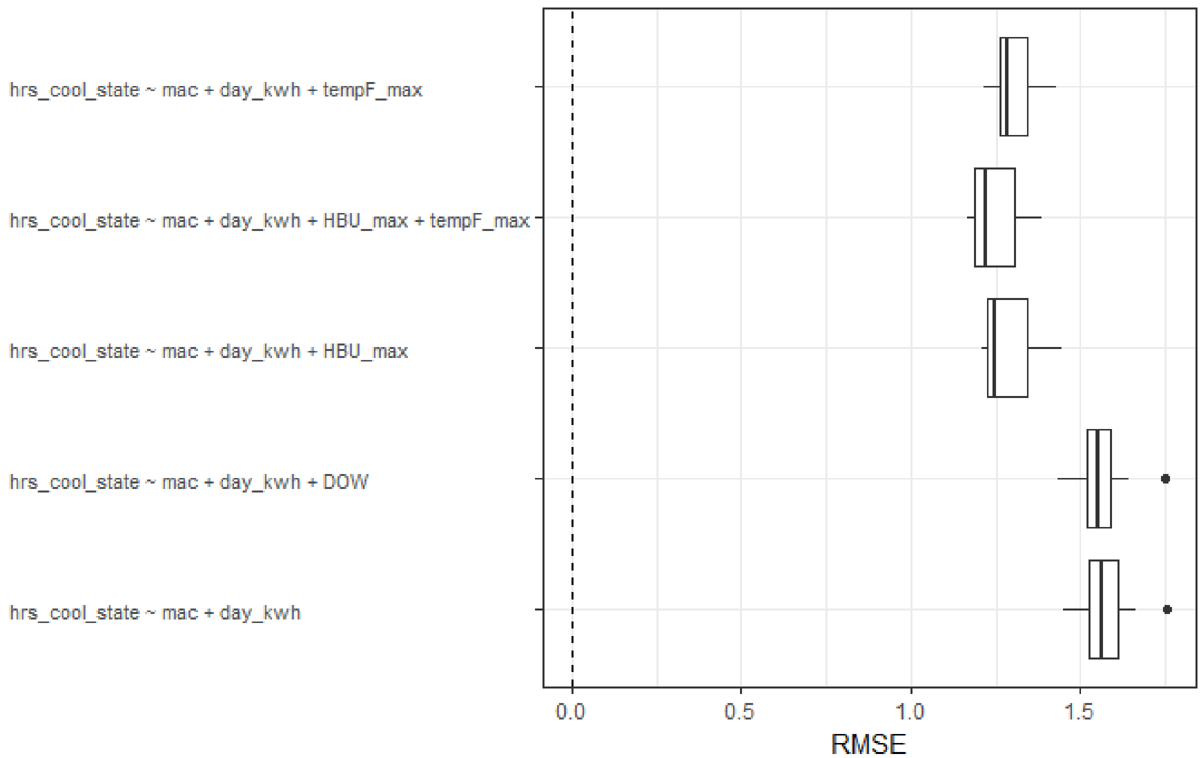

Source: Navigant

Table 16. Regression Results from the Full AMI vs. Weather Model

Term	Estimate	Std. Error	t-value	p-value
maxci	0.375	0.0064	58.6	< 1e-9
avghbu	0.070	0.0008	85.5	< 1e-9
avghum	0.030	0.0006	53.4	< 1e-9
mintemp_l1	0.008	0.0022	3.6	3.83e-4
maxtemp_l1	-0.076	0.0018	-43.0	< 1e-9
avghum_l1	0.003	0.0007	5.2	1.89e-7
mintemp_l2	0.023	0.0017	13.9	< 1e-9
maxtemp_l2	-0.033	0.0016	-21.0	< 1e-9

Source: Navigant

A.9 Additional In-Depth Interview Quotes

A.9.1 Interview Findings

Additional quotes from DRSPs and participants support the findings in Section 3.4.4.

“We come home, things are back to normal in the evening, and we may not even know that it happened in some instances. You still see the alert, but you never ... you're not home to have it affect you.” – CS Participant

“And also the fact that it's typically during the middle of a work day and I do often work from home a couple days out of the week. But if it's the rush hour and nobody's home and it's warm in the house and there's nobody there, then there's really not a huge impact on it.” – RHR Participant

A.9.2 DR Event Participant Types

Some additional quotes from DRSPs and participants support the findings in Section 3.4.4.1.

“You kind of get worn down by it. My husband, he's a contractor, and he works outside, so if he was outside all day and it was really, really hot and humid, I'd try and cool it off in here when he gets home, too. That might impact a Rush Hour, because usually the Rush Hour was, like I say, a little later afternoon. He'd get home and shower and try and cool off so I'd lower the temperature a little bit there for him, too.” – RHR Participant

“I can think of four or five times it happened and we were home for, I think, all of them except for that one little time that we came home and it was happening and that's why we cranked it up. And all the other ones, I remember looking at it and going, "Huh. It feels a little different." We looked at it and we realized we were in it. But it wasn't bad. It wasn't something that we had to opt out of.” – RHR Participant

“Just the situation where, if I'm out and I've already turned either the heat down or the A/C down, then I happen to have to stay longer at work, or stay longer wherever I am, I like to keep it turned down.” – CS Participant

A.10 Weather and DR Program Participation Methodology

The team chose to model opt-out behavior using an LPM[75] because the dependent variable is binary, taking a value of one when a participant opted out of the event and zero otherwise. LPM regression predicts a linear change in the probability of an outcome due to changes in the independent variables.

[75] In addition to the LPM, the team built a series of logistic regression models. The logistic regression models gave similar results, but the LPM provided the best fitting model and most easily interpreted coefficients. See Appendix Section A.1.2 for logistic regression output.

$$Y_{it} = MOY_t + \beta_{1t}\sum_t DOW_t + \beta_2 EventDuration_{it} + \beta_3 EventStartHour_{it} + \beta_4 TotalDailyCDH65_{it} + \beta_5 MaxRH_{it} + \beta_6 LagTotalDailyCDH65_{it} + \beta_7 LagMaxRH_{it} + \beta_8 ConsecutiveEvent_{it} + \beta_{9t}\sum_t (MOY_t \times Nest_i) + \beta_{10t}\sum_t (ConsecutiveEvent_{it} \times Nest_i) + \varepsilon_{it}$$

Y_{it}	= a binary variable taking on a value of 1 if customer *i* opted-out or overrode preset controls during event *t*
MOY_t	= a month-of-year fixed effect
DOW_t	= a set of binary variables indicating on which day of week event *t* took place[76]
$EventDuration_{it}$	= the duration of event *t* for customer *i*
$EventStartHour_{it}$	= the hour of day of the beginning of event *t* for customer *i*
$TotalDailyCDH65_{it}$	= the total daily CDH (base 65) in customer *i*'s ZIP code on the day of event *t*
$MaxRH_{it}$	= the maximum relative humidity in customer *i*'s ZIP code during event *t*
$LagTotalDailyCDH65_{it}$	= the total daily CDH (base 65) in customer *i*'s ZIP code on the day before of event *t*
$LagMaxRH_{it}$	= the maximum relative humidity in customer *i*'s ZIP code on the day before event *t*
$ConsecutiveEvent_{it}$	= a binary variable taking on a value of 1 if customer *i* was called for an event on the day prior to event *t*
$Nest_i$	= a binary variable taking on a value of 1 if customer *i* is a Nest thermostat user
ε_{it}	= an error term for customer *i* during event *t*

The estimated coefficients, β_x, represent the effects of specific weather variables or event attributes on the propensity of a participant to opt out of a DR event.

The model fit using the above specification was cross-validated against several other model specifications. The team conducted 50 iterations of each model, each time reserving 20% of the data as a test set. In comparing the models, the team calculated the distribution of the cross-validation errors using RMSE as the comparison metric.[77] Lower RMSE implies better model performance. shows box plots of the cross-validation RMSE for a selection of candidate models. All LPM specifications have similarly low error.

[76] The set of binary variables did not include Saturday or Sunday, as no events took place on weekends. The regression model eliminated the Wednesday indicator such that all other day of week effects are relative to Wednesday (i.e., the day of week with the lowest opt-out rate).

[77] Note that in earlier regression cross-validation work in this study the team performed 200 iterations of the cross-validation algorithm but found the results were virtually identical using 50 iterations. As a result, the team used 50 iterations in this and following cross-validation efforts.

Figure 69. Cross-Validation Errors for Logistic and LPM Candidate Models

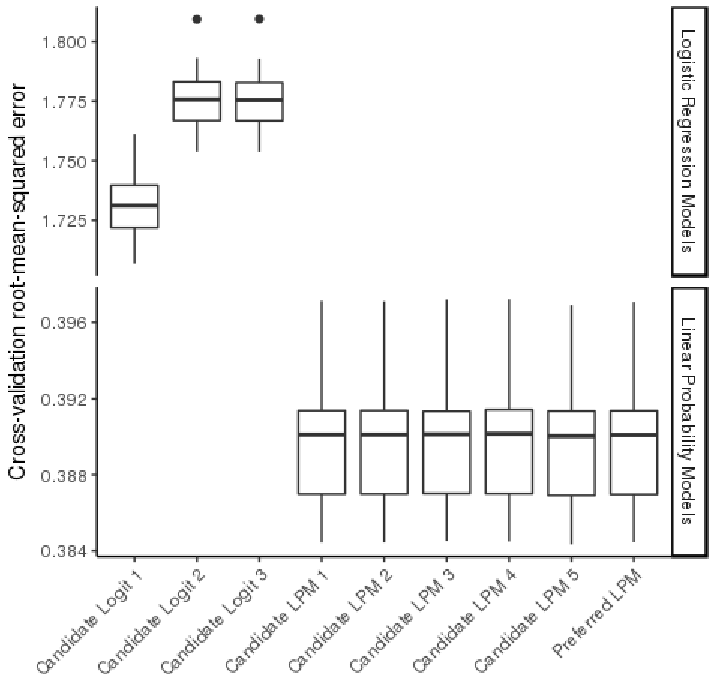

Source: Navigant

A.11 Weather and DR Program Participation Detailed Findings

While a truncated table with the most important coefficient results is included in Section, full results are presented in Figure 70 for all regression coefficients.

Figure 70. LPM Regression Full Coefficient Output

```
Coefficients:
                                  Estimate Std. Error t value Pr(>|t|)    
moy201607                       -4.353e-01  1.091e-01  -3.988 6.67e-05 ***
moy201608                       -4.269e-01  1.107e-01  -3.856 0.000115 ***
moy201609                       -4.587e-01  1.116e-01  -4.110 3.97e-05 ***
moy201707                       -3.904e-01  1.060e-01  -3.681 0.000232 ***
moy201708                       -3.835e-01  1.063e-01  -3.609 0.000308 ***
moy201709                       -3.387e-01  1.085e-01  -3.122 0.001800 ** 
monday_flag                      9.506e-03  6.194e-03   1.535 0.124844    
tuesday_flag                     7.277e-03  8.010e-03   0.908 0.363644    
thursday_flag                    3.549e-02  7.447e-03   4.765 1.90e-06 ***
friday_flag                      3.789e-02  7.786e-03   4.867 1.14e-06 ***
event_duration                   3.315e-02  8.191e-03   4.046 5.21e-05 ***
event_start_hour                 2.349e-02  6.761e-03   3.475 0.000512 ***
total_daily_CDH65                2.647e-04  3.900e-05   6.787 1.16e-11 ***
event_max_RH                     3.820e-04  1.654e-04   2.310 0.020917 *  
lag_total_daily_CDH65           -1.085e-04  3.566e-05  -3.044 0.002338 ** 
lag_daily_high_RH                1.164e-04  1.456e-04   0.799 0.424364    
consecutive_event_flag           7.270e-03  9.872e-03   0.736 0.461488    
moy201607:nest_flag              1.527e-01  1.192e-02  12.811  < 2e-16 ***
moy201608:nest_flag              1.418e-01  1.076e-02  13.185  < 2e-16 ***
moy201609:nest_flag              1.572e-01  1.520e-02  10.341  < 2e-16 ***
moy201707:nest_flag              3.930e-02  1.079e-02   3.643 0.000270 ***
moy201708:nest_flag              4.149e-02  1.190e-02   3.487 0.000490 ***
moy201709:nest_flag              2.688e-02  2.433e-02   1.105 0.269282    
consecutive_event_flag:nest_flag -2.408e-03  1.060e-02  -0.227 0.820336    
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1
```

Source: Navigant

A.12 Supplemental Weather and DR Program Participation Findings—Logistic Regression Analysis

In addition to the LPM specification, the team ran a logistic regression version of the model. A logistic regression model predicts change in log-odds of an outcome due to changes in the independent variables. In place of the lags in CDH65 and RH the logistic regression model included HBU to account for weather history leading up to the event. Following is the logistic regression specification:

$$
\begin{aligned}
Logit(Y_{it}) = MOY_t &+ \beta_1 Monday_t + \beta_2 Tuesday_t + \beta_3 Thursday_t + \beta_4 Friday_t + \beta_5 EventDuration_{it} \\
&+ \beta_6 EventStartHour_{it} + \beta_7 TotalDailyCDH65_{it} + \beta_8 MaxRH_{it} + \beta_9 DailyHighHBU_{it} \\
&+ \beta_{10} ConsecutiveEvent_{it} + \beta_{11t} \sum_t (MOY_t \times Nest_i) \\
&+ \beta_{12t} \sum_t (ConsecutiveEvent_{it} \times Nest_i) + \varepsilon_{it}
\end{aligned}
$$

Y_{it}	= a binary variable taking on a value of 1 if customer *i* opted-out or overrode preset controls during event *t*
MOY_t	= a month-of-year fixed effect
$Monday_t$	= a binary variable taking on a value of 1 if event *t* took place on a Monday

$Tuesday_t$ = a binary variable taking on a value of 1 if event *t* took place on a Tuesday

$Thursday_t$ = a binary variable taking on a value of 1 if event *t* took place on a Thursday

$Friday_t$ = a binary variable taking on a value of 1 if event *t* took place on a Friday

$EventDuration_{it}$ = the duration of event *t* for customer *i*

$EventStartHour_{it}$ = the hour of day of the beginning of event *t* for customer *i*

$TotalDailyCDH65_{it}$ = the total daily CDH (base 65) in customer *i*'s ZIP code on the day of event *t*

$MaxRH_{it}$ = the maximum relative humidity in customer *i*'s ZIP code during event *t*

$DailyHighHBU_{it}$ = the maximum HBU in customer *i*'s ZIP code on the day of event *t*

$ConsecutiveEvent_{it}$ = a binary variable taking on a value of 1 if customer *i* was called for an event on the day prior to event *t*

$Nest_i$ = a binary variable taking on a value of 1 if customer *i* is a Nest thermostat user

ε_{it} = an error term for customer *i* during event *t*

The estimated coefficient, β, represents the impact of specific weather variables or event attributes on the propensity to participate in a DR event.

The logistic regression model would be preferable to a linear model if the relationships between the independent variables and the opt-out rates were nonlinear. However, the team did not find this to be the case. Figure 71 plots the marginal effects for each attribute in the logistic regression model. In each plot, the x-axis is the value of the attribute, and the y-axis is the model's predicted value of the opt-out rate. Because each variable presents a nearly perfectly linear trend with opt-out rate, the team determined that the added complexity of the logistic specification was not providing any benefit to the fit of the model and was unnecessary. Still, the model produces relationships between event attributes, weather, and opt-out rates that make intuitive sense.

Figure 71. Marginal Effects of Logistic Regression Coefficients

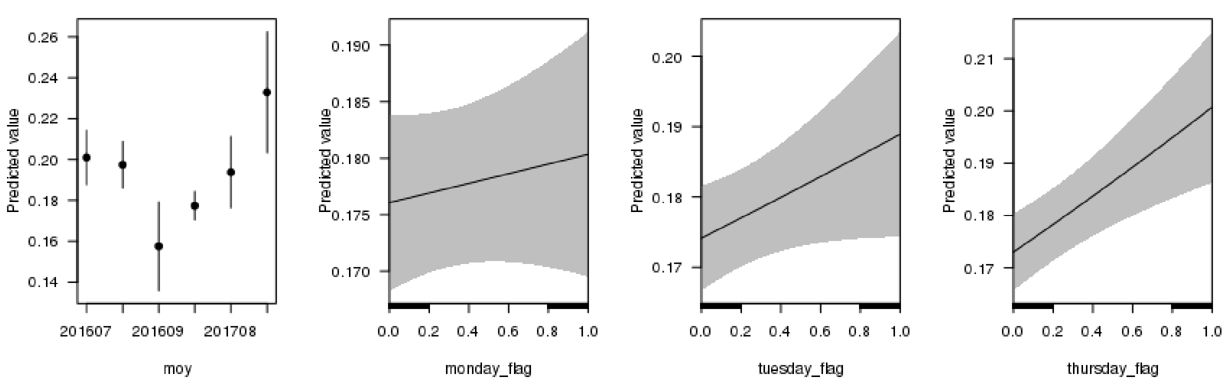

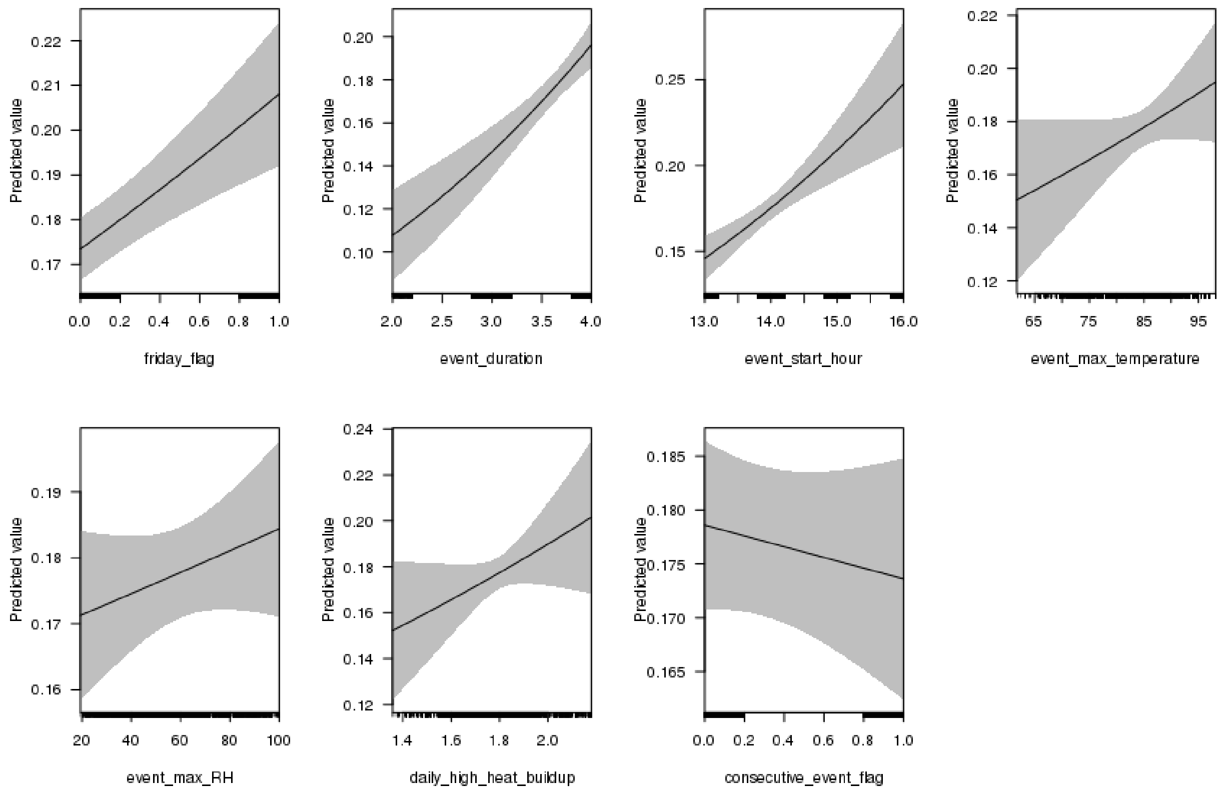

Source: Navigant

For completeness, the logistic regression coefficients are shown in Figure 72, presented as log-odds estimates. These coefficients match the LPM coefficients in direction, which provides a helpful robustness check for the team's preferred model.

Figure 72. Logistic Regression Coefficients

```
Coefficients:
                                Estimate Std. Error z value Pr(>|z|)    
moy201607                     -8.0503393  0.8418826  -9.562  < 2e-16 ***
moy201608                     -8.0467297  0.8463084  -9.508  < 2e-16 ***
moy201609                     -8.4391099  0.8584938  -9.830  < 2e-16 ***
moy201707                     -7.6425306  0.8123634  -9.408  < 2e-16 ***
moy201708                     -7.6285961  0.8153918  -9.356  < 2e-16 ***
moy201709                     -7.2760882  0.8302620  -8.764  < 2e-16 ***
monday_flag                    0.0291205  0.0408992   0.712 0.476462    
tuesday_flag                   0.0994717  0.0526202   1.890 0.058708 .  
thursday_flag                  0.1825407  0.0494342   3.693 0.000222 ***
friday_flag                    0.2246828  0.0501245   4.482 7.38e-06 ***
event_duration                 0.3513315  0.0670782   5.238 1.63e-07 ***
event_start_hour               0.2176893  0.0479711   4.538 5.68e-06 ***
event_max_temperature          0.0085773  0.0051904   1.653 0.098426 .  
event_max_RH                   0.0011146  0.0009611   1.160 0.246178    
daily_high_heat_buildup        0.4149982  0.2671712   1.553 0.120350    
consecutive_event_flag        -0.0021108  0.0734193  -0.029 0.977064    
moy201607:nest_flag            1.0086246  0.0873220  11.551  < 2e-16 ***
moy201608:nest_flag            0.9718706  0.0759486  12.796  < 2e-16 ***
moy201609:nest_flag            1.1386992  0.1127292  10.101  < 2e-16 ***
moy201707:nest_flag            0.2157714  0.0859989   2.509 0.012107 *  
moy201708:nest_flag            0.3491745  0.0905400   3.857 0.000115 ***
moy201709:nest_flag            0.1814859  0.1657202   1.095 0.273458    
consecutive_event_flag:nest_flag -0.0454375  0.0788026  -0.577 0.564210    
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Source: Navigant

A.13 Supplemental Weather and DR Program Participation Findings—Differences by Thermostat Type

Previous research[78] indicated that the relationship between weather and opt-out rate was stronger for Honeywell users compared to Nest and ecobee users. To investigate this hypothesis, the team built LPM models segmented by thermostat type for Honeywell and Nest cohorts. There was not enough data to produce a meaningful relationship between weather and opt-out behavior for ecobee users.

The variables included in the segmented models differed slightly from the model built on the full population. The month of year fixed effect was replaced with a year fixed effect because, when considering only one vendor, program changes are only made between the two program years. Given the smaller sample sizes, the relative humidity term became statistically insignificant. For this reason, the team removed the lagged relative humidity term. The interaction terms with the Nest thermostat indicator were replaced with an interaction term between year and event duration. Below is the LPM regression specification for the Nest and Honeywell segmented models.

$$\begin{aligned} Y_{it} = {} & Year_t + \beta_1 Monday_t + \beta_2 Tuesday_t + \beta_3 Thursday_t + \beta_4 Friday_t + \beta_5 EventDuration_{it} \\ & + \beta_6 EventStartHour_{it} + \beta_7 TotalDailyCDH65_{it} + \beta_8 MaxRH_{it} + \beta_9 LagTotalDailyCDH65_{it} \\ & + \beta_{10} ConsecutiveEvent_{it} + \beta_{11t} \sum_t (Year_t \times EventDuration_{it}) + \varepsilon_{it} \end{aligned}$$

[78] Navigant. (2017). National Grid 2016 Residential Wi-Fi Thermostat DR Evaluation Final Report.

Y_{it}	= a binary variable taking on a value of 1 if customer *i* opted-out or overrode preset controls during event *t*
$Year_t$	= a year fixed effect
$Monday_t$	= a binary variable taking on a value of 1 if event *t* took place on a Monday
$Tuesday_t$	= a binary variable taking on a value of 1 if event *t* took place on a Tuesday
$Thursday_t$	= a binary variable taking on a value of 1 if event *t* took place on a Thursday
$Friday_t$	= a binary variable taking on a value of 1 if event *t* took place on a Friday
$EventDuration_{it}$	= the duration of event *t* for customer *i*
$EventStartHour_{it}$	= the hour of day of the beginning of event *t* for customer *i*
$TotalDailyCDH65_{it}$	= the total daily CDH (base 65) in customer *i*'s ZIP code on the day of event *t*
$MaxRH_{it}$	= the maximum relative humidity in customer *i*'s ZIP code during event *t*
$LagTotalDailyCDH65_{it}$	= the total daily CDH (base 65) in customer *i*'s ZIP code on the day before of event *t*
$ConsecutiveEvent_{it}$	= a binary variable taking on a value of 1 if customer *i* was called for an event on the day prior to event *t*
ε_{it}	= an error term for customer *i* during event *t*

The coefficient of interest in determining the relative impacts of temperature is *TotalDailyCDH65*. Because the average opt-out rates for Nest and Honeywell are different, the coefficients from the segmented models are not directly comparable. This is because the same absolute increase in opt-out rate is a relatively larger increase for Honeywell customers, based on their lower average opt-out rate.

To compare the CDH-based coefficients and determine which users are more sensitive to weather, the team normalized the coefficients by dividing by the average opt-out rate. The normalized relative coefficients can be interpreted as the average percent increase in opt-out rate that results from a one degree-hour increase in total daily CDH. Table 17 shows both the original coefficients produced by the regression models and the normalized relative coefficients that are directly comparable between thermostat types.

Table 17. Relative Impacts of Weather on Nest and Honeywell Customers

Thermostat Type	Opt-Out Rate	CDH65 Coefficient		CDH65 Relative Coefficient	
		Estimate	95% CI	Estimate	95% CI
Nest	22.5%	0.00034	[0.00025, 0.00043]	0.00150	[0.00109, 0.00190]
Honeywell	12.6%	0.00020	[0.00008, 0.00032]	0.00156	[0.00062, 0.00251]

Source: Navigant

After normalizing for average opt-out rate of each thermostat type, the CDH relative coefficients are very similar between Nest and Honeywell. Even though the Honeywell relative coefficient is slightly larger than the Nest relative coefficient, the team cannot confirm the original finding that Honeywell users are more sensitive to weather because there is a large amount of overlap between the relative coefficient confidence intervals.

Copied below are the full sets of coefficients for the segmented model estimated on the Nest and Honeywell cohorts of thermostats. The coefficient named "*total_daily_CDH65*" displays the estimated coefficient shown in Table 17, and includes more information on the standard error, p-value and significance level.

Figure 73. Nest Segmented LPM Coefficients

```
Coefficients: (1 not defined because of singularities)
                          Estimate Std. Error t value Pr(>|t|)
year2016                -1.698e-01  1.123e-01  -1.512 0.130554
year2017                -2.316e-01  1.105e-01  -2.097 0.036039 *
monday_flag              1.827e-02  7.592e-03   2.407 0.016092 *
tuesday_flag             2.507e-03  9.929e-03   0.252 0.800662
thursday_flag            3.672e-02  9.587e-03   3.830 0.000128 ***
friday_flag              4.327e-02  1.052e-02   4.111 3.95e-05 ***
event_duration           2.555e-02  7.894e-03   3.236 0.001213 **
event_start_hour         1.545e-02  7.493e-03   2.062 0.039253 *
total_daily_CDH65        3.374e-04  4.650e-05   7.255 4.13e-13 ***
event_max_RH             7.145e-04  1.948e-04   3.668 0.000245 ***
lag_total_daily_CDH65   -1.320e-04  4.384e-05  -3.011 0.002605 **
consecutive_event_flag   9.858e-03  9.097e-03   1.084 0.278500
year2017:event_duration         NA         NA      NA       NA
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Source: Navigant

Figure 74. Honeywell Segmented LPM Coefficients

```
                          Estimate Std. Error t value Pr(>|t|)
year2016                -2.728e-01  1.391e-01  -1.961  0.04986 *
year2017                -2.296e-01  1.326e-01  -1.732  0.08324 .
monday_flag              1.109e-02  1.030e-02   1.077  0.28152
tuesday_flag             1.427e-02  1.220e-02   1.170  0.24213
thursday_flag            2.856e-02  1.313e-02   2.175  0.02964 *
friday_flag              3.318e-02  1.140e-02   2.911  0.00361 **
event_duration           1.427e-02  9.824e-03   1.453  0.14627
event_start_hour         1.846e-02  8.182e-03   2.256  0.02408 *
total_daily_CDH65        1.975e-04  6.086e-05   3.246  0.00117 **
event_max_RH             1.754e-04  2.142e-04   0.819  0.41308
lag_total_daily_CDH65   -3.061e-06  6.020e-05  -0.051  0.95945
consecutive_event_flag  -4.073e-04  9.769e-03  -0.042  0.96674
year2017:event_duration         NA         NA      NA       NA
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Source: Navigant

A.14 Supplemental Weather and DR Program Participation Findings—The Impact of Weather History

One interesting finding from the main regression model is the negative impact of the lagged *CDH65* variable. That is, when the day before an event is hotter, a participant's likelihood of opting out decreases. The team's hypothesis is that a hotter day before an event makes the event day feel cooler by comparison and decreases the opt-out rate. Conversely, a cooler day before an event makes the event day feel hotter by comparison and increases the opt-out rate.

To check the robustness of this finding, the team estimated a model which replaced the lagged *CDH65* variable with the lag ratio, defined as the ratio of total daily *CDH65* on the event day to *TotalDailyCDH65* on the day before the event. The same type of lag ratio variable was created and used to replace the relative humidity lag variable. Figure 75 explains how the model with the lag ratios ("Lag-ratio model") corroborates the findings presented from the main regression model ("Main model").

Figure 75. Weather History Impacts on Opt-Out Rates

Lag-ratio model

- ***Coefficient:*** Ratio of total daily CDH65 to its lag
- ***Direction:*** Positive
- ***Interpretation:*** When the day before an event is cooler than the event day (i.e., ratio > 1), the event day feels hotter by comparison, resulting in higher opt-out rates

Main model

- ***Coefficient:*** Total daily CDH65 lag
- ***Direction:*** Negative
- ***Interpretation:*** A cooler day before an event makes the event day feel hotter by comparison, resulting in higher opt-out rates

Source: Navigant

For completeness, Figure 76 shows the estimated coefficients for the model using lag ratios. Like the main model, the relative humidity history (RH lag ratio in this model, maximum RH lag in the main model) remains statistically insignificant. The trends the team observes for the remaining attributes do not show any substantial changes between the main model and the model with lag ratios.

Figure 76. LPM with Lag-Ratio Coefficients

```
Coefficients:
                                   Estimate Std. Error t value Pr(>|t|)
moy201607                        -4.740e-01  1.150e-01  -4.122 3.77e-05 ***
moy201608                        -4.643e-01  1.162e-01 | -3.996 6.44e-05 ***
moy201609                        -4.929e-01  1.171e-01  -4.209 2.57e-05 ***
moy201707                        -4.246e-01  1.115e-01  -3.809 0.000140 ***
moy201708                        -4.181e-01  1.120e-01  -3.734 0.000189 ***
moy201709                        -3.705e-01  1.140e-01  -3.250 0.001153 **
monday_flag                       1.013e-02  6.155e-03   1.646 0.099709 .
tuesday_flag                      8.225e-03  7.922e-03   1.038 0.299154
thursday_flag                     3.430e-02  7.374e-03   4.652 3.31e-06 ***
friday_flag                       3.954e-02  7.889e-03   5.012 5.41e-07 ***
event_duration                    3.513e-02  8.310e-03   4.227 2.37e-05 ***
event_start_hour                  2.333e-02  6.791e-03   3.435 0.000592 ***
total_daily_CDH65                 1.824e-04  3.249e-05   5.614 1.99e-08 ***
event_max_RH                      4.458e-04  1.510e-04   2.953 0.003146 **
CDH65_to_lag_ratio                2.105e-02  5.549e-03   3.793 0.000149 ***
RH_to_lag_ratio                   7.816e-03  2.481e-02   0.315 0.752754
consecutive_event_flag            8.046e-03  9.792e-03   0.822d 0.411236
moy201607:nest_flag               1.530e-01  1.191e-02  12.851  < 2e-16 ***
moy201608:nest_flag               1.429e-01  1.061e-02  13.468  < 2e-16 ***
moy201609:nest_flag               1.565e-01  1.519e-02  10.301  < 2e-16 ***
moy201707:nest_flag               3.786e-02  1.083e-02   3.497 0.000471 ***
moy201708:nest_flag               4.259e-02  1.189e-02   3.582 0.000341 ***
moy201709:nest_flag               2.701e-02  2.433e-02   1.110 0.266891
consecutive_event_flag:nest_flag -4.633e-03  1.059e-02  -0.437 0.661874
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Source: Navigant