REPORT

Evaluation of Southern California Edison's HER Persistence Pilot

Prepared for:

Southern California Edison Co.

SCE EM&V Advisor and Project Manager: Galib Rustamov, Ph.D.

September 2, 2020

Principal authors:

Eric Bell, Ph.D., Principal

Aimee Savage, Consultant II

John Ensley, Data Scientist

Robert Gottlieb, Project Analyst

CALMAC Study ID SCE0447

Contents

EVALUATION OF SOUTHERN CALIFORNIA EDISON'S HER PERSISTENCE PILOT

Savings from paper Home Energy Reports (HERs) persist for at least one year after the discontinuation of reports

Discontinued customer electric savings are equal to 98% of continued customer savings during the 12 months of the Persistence Pilot, and 82% for an additional eight-month period. However the difference in energy savings between the two groups is not statistically significant in either time period.

DAILY kWh SAVINGS

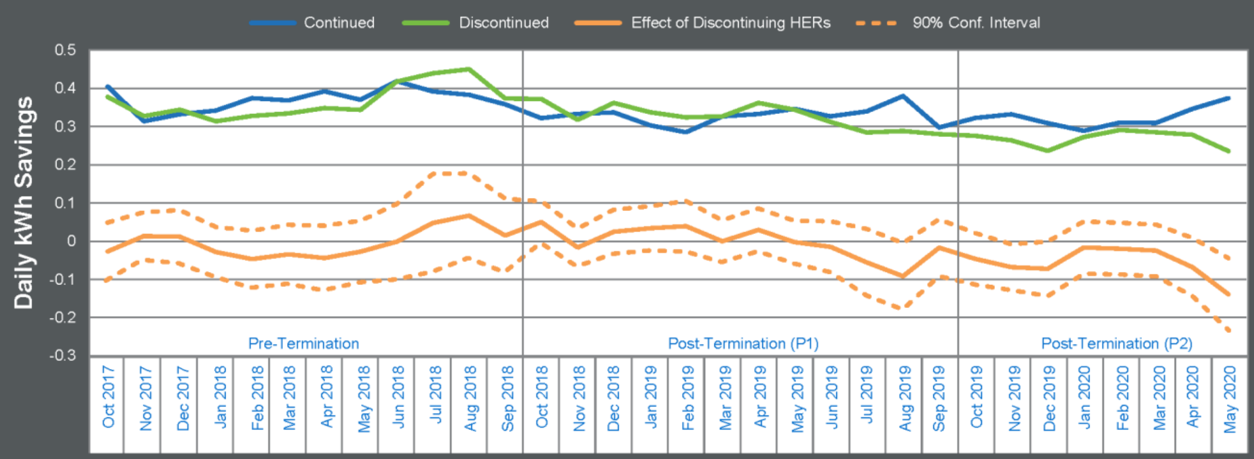

WAVE & SAMPLE SIZE

Wave 3 Group	eHER Status	Continues receiving print HERs	No longer receiving print HERs
Paper-Only	Has never received eHERs	27,000	27,000
Paper + Email	Continues to receive eHERs	39,000	39,000
Total households		66,000	66,000

CONSISTENCY WITH OTHER STUDIES

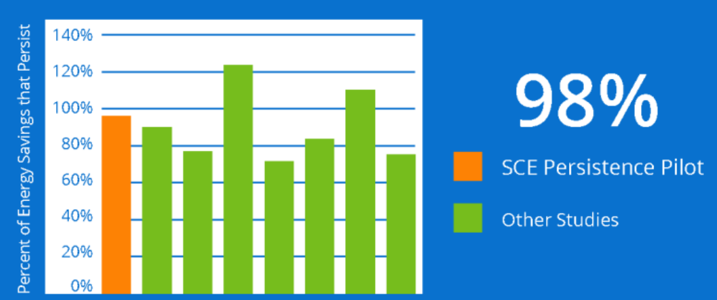

PAPER-ONLY PERSISTENCE

PAPER + EMAIL PERSISTENCE

Nexant

CUSTOMER REACTIONS TO HERs

Customers with the highest and lowest pre-termination savings move towards the overall average savings of around 0.3 kWh/day in the post-termination period. **This indicates that savings are indeed persistent**, but the variance among the savings may be decreasing over time.

HIGH & LOW SAVERS

0.45 kWh/day

0.1 kWh/day

Savings are strongly correlated with pre-treatment consumption:

Low-consumption customers are predicted to have the least amount of savings. In fact, households with less than around 20 kWh daily pretreatment average consumption were predicted in increase usage. Savings increase as per-treatment consumption increases from 20 kWh up until about 45 kWh, at which point it levels off. Higher levels of pre-treatment consumption after 45 kWh per day are not correlated with higher levels of savings.

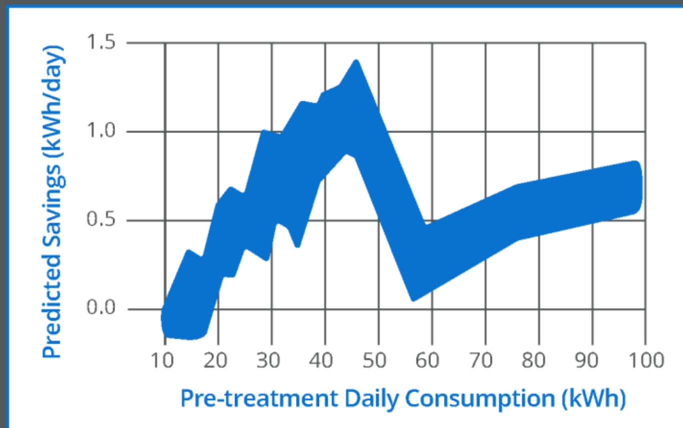

ADDITIONAL FINDINGS

- When studied in isolation, daily load shape and seasonal usage patterns did not have a significant effect on the level of persistence; nor did household age or income.
- However, the largest savings can be found in households with greater than average incomes and those in which the head of household is between 55 and 65 years old; Younger, lower-income households tend to produce the least savings.

1 Executive Summary

This document summarizes the evaluation of the Southern California Edison Co. (SCE) Home Energy Report (HER) Persistence Pilot. The primary objective of SCE's pilot is to understand what happens to customer behavior and energy savings if customers stop receiving HERs. The pilot was launched in October 2018 within SCE's "Wave 3" HER cohort, which began treatment in September 2015. Customers in this wave are segmented into two treatment cells: those who receive paper HERs only, and those who receive email HERs in addition to paper HERs. Customers in each treatment cell were randomly assigned to two groups under a robust experiment designed to test the persistence in savings: 50% of customers continue to receive paper HERs, and 50% no longer receive paper HERs. Customers in the Paper + Email cell continue to receive electronic HERs. Customers in the discontinued group resumed treatment in June 2020, twenty months after the Persistence Pilot began. The additional time before resuming treatment presented an opportunity to examine savings persistence beyond the originally planned 12-month period. Accordingly, the persistence of energy savings for the additional eight months between October 2019 and May 2020 is also included in this report. This report presents findings from the first year and eight months of the pilot, October 2018 through May 2020.

1.1 Key Findings

Key findings from the Persistence Pilot evaluation include:

- Savings from paper HERs persist for at least one year and eight months after the discontinuation of reports. Discontinued customer electric savings are equal to 98% of continued customer savings during the first year of the persistence study, and 82% for an additional eight-month period. Indeed, the difference in energy savings between the two groups is not statistically significant.
- Customers in the Paper-Only treatment group showed 90% persistence across the first 12-months of paper report discontinuation. However, the difference in energy savings between the continued and discontinued customers is not statistically significant.
- The Paper + Email group, who continued to receive electronic HERs but did not receive paper-based HERs, exhibited a small, but not statistically significant increase in energy savings during the first 12-months of the study.
- The findings from SCE's Persistence Pilot are consistent with other studies across the country including the initial study conducted by DNV-GL on SCE's Opower-1 group in 2017.[1] In this study, savings were estimated to persist at the 99% level after the first year, and 75% after the second year.
- Energy savings attributable to HERs are strongly correlated with pretreatment consumption: higher usage households tend to produce more savings.

[1] DNV-GL (2017), Review and Validation of 2015 Southern California Edison Home Energy Reports Program Impacts (Final Report), CALMAC Study ID: CPU0156.01

- When studied in isolation, daily load shape and seasonal usage patterns did not have a significant effect on the level of persistence; nor did household age or income. However, the largest savings can be found in households with greater than average incomes and those in which the head of household is between 55 and 65 years old. Younger, lower-income households tend to produce the least savings. It is important to note, however, that age and income are highly correlated with pretreatment energy consumption.
- The distribution of energy savings suggests there are two types of responders: high savers (0.5 kWh per day savings on average) and low savers (0.1 kWh per day). However, customers with the highest and lowest pre-termination savings move towards the overall average savings of around 0.3 kWh per day in the post-termination period. This indicates that savings are indeed persistent, but the variance among the savings may be decreasing over time.

1.2 Conclusions and Recommendations

Given the high level of persistence found in this evaluation, there are several opportunities to increase the cost-effectiveness of SCE's HER program, some of which may require further study.

- First, it is likely that HER savings would persist in SCE's other HER treatment cohorts. SCE may consider discontinuing treatment in these waves, perhaps temporarily, to maintain savings and reduce costs.
- Reducing the frequency of reports would likely have as similar effect. To meet the effective useful life (EUL) of 1 year, one or two HERs could be mailed instead of four—claiming the same savings but cutting the mailing and transaction cost by half or more.
- Alternatively, a staggered program design could be considered, where mailers are provided for one year, and not in another, with savings being evaluated and claimed for off-years.
- Finally, SCE may consider a shift to lower-cost channels by transitioning to email-only HERs after treating with a combination of paper and electronic reports.
- The causal forest analysis also identified that customers with low levels of pretreatment usage (less than 20 kWh per day) tend to not provide energy savings. Forty-one percent of Wave 3 customers used less than 20 kWh per day, however these customers only provide approximately ten percent of the aggregate MWh energy savings. Customers with low average daily usage could be considered for exclusion from future treatment waves— reducing cost, and resulting in higher per-customer savings levels for the average customer receiving treatment.

These findings have implications to current program implementation, and in the future under a 3rd party program framework. Under the current program implementation structure with utilities administering the programs, there are opportunities to improve the cost effectiveness through implementing strategies to pause treatment for customers after receiving several years of treatment. The entire treatment cost for these customers could be saved for the year, while energy savings are still realized. Alternatively, the savings could be reinvested to allow for treatment of additional customers, increasing the reach of the program. Under the 3rd party model this information provides value, as it allows for the 3rd parties to develop innovative solutions that build in pauses in treatment, and potentially more targeted treatment cells in order to optimize program cost effectiveness. These findings will also allow decision makers to better understand the potential implications for creative program designs that may be proposed by potential 3rd party implementers during the procurement process.

2 Introduction

This document summarizes the evaluation of the Southern California Edison Co. (SCE) Home Energy Report (HER) Persistence Pilot. The primary objective of SCE's pilot is to understand what happens to customer behavior and savings if customers stop receiving HERs. Key research questions include:

- How long does the treatment effect persist after the cessation of HERs?
- What is the energy savings decay rate per year after the cessation of the treatment with paper HERs?
- Which types of customers are the most persistent savers?
- How are SCE's persistence results similar or dissimilar from previous findings?

This evaluation seeks to identify the persistence effect within the pilot. The pilot was launched within SCE's "Wave 3" HER cohort. Treatment customers in this cohort are divided into two populations, those with email addresses in SCE's database and those without. Both groups have been receiving paper HERs since September 2015 and the email group receives eHERs in addition to paper reports. A randomly selected group of 66,000 treatment households experienced a twenty-month pause in paper reports from October 2018 through May 2020.[2] This represents half of the active treatment population at that time, and includes 27,000 non-email customers and 39,000 email customers (who continue to receive eHERs in the absence of paper reports). Table 2-1 summarizes the design of the Persistence Pilot.

Table 2-1: Persistence Pilot Design

Wave 3 Group	eHER Status	Continues receiving print HERs	No longer receives print HERs
Paper-Only	Has never received eHERs	27,000	27,000
Paper + Email	Continues receiving eHERs	39,000	39,000
Total households		**66,000**	**66,000**

This evaluation includes persistence estimates for the period from October 2018 through September 2019 (P1) and October 2019 through May 2020 (P2). This time period covers one year and eight months after the cessation of paper HERs for the discontinued groups. The continued groups and the Wave 3 control group were used to establish baseline energy savings. Next, persistence was estimated for the eHER and Paper-Only groups both separately

[2] Discontinued customers resumed treatment with paper HERs in June 2020.

and combined. Additionally, persistence in energy savings was estimated within specific customer segments and clusters.

The persistence of energy savings from HERs was estimated using a series of regression models. First, the energy savings from the continued group from October 2018 through May 2020 was estimated to establish baseline energy savings.[3] This is the denominator when estimating the percent of savings that persist. Second, the difference in energy consumption between the continued and discontinued groups was estimated for the same time period. This represents the change in energy savings attributed to the discontinuation of the paper HER treatment. This process was conducted for the Paper-Only group, the Paper + Email group, and the two groups combined. A similar process was used to estimate persistence among specific customer segments and clusters. Energy savings and persistence were estimated using a lagged dependent variable (LDV) model in which pretreatment energy consumption is an explanatory variable.

The remainder of this report is organized as follows:

- Section 3 describes the methodology used to estimate the persistence of savings;
- Section 4 presents high-level persistence estimates for the Paper-Only and Paper + Email groups;
- Section 5 presents the findings of the segmentation and clustering analysis;
- Section 6 presents the causal forest results; and
- Section 7 provides a comparison to other persistence studies.

[3] Savings and persistence levels were estimated separately for two time periods. P1 includes the first 12 months of the pilot, October 2018 through September 2019. P2 includes the following eight month period from October 2019 through May 2020.

3 Methodology

This section summarizes the methodological approach used to estimate the persistence of energy savings after the cessation of treatment with HERs. The discussion is organized into three sections summarizing the approach for estimating the persistence of savings, the segmentation and clustering of customers, and the causal forest analysis.

3.1 Estimation of Persistence

This section summarizes the methodology that was used to address estimate the persistence of energy savings for the Persistence Pilot population. Persistence was estimated separately for each population (Paper-Only and Paper + Email) and for various customer segments outlined in the following subsection. For the purposes of this analysis, customers in the pilot were defined in three ways:

- **Control** customers are those who have never received HERs and are statistically equivalent to those in the following two groups.
- **Continued** customers are those who have received HERs since the launch of Wave 3 (September 2015) and will continue to receive reports through at least May 2020.
- **Discontinued** customers are those who received HERs from the launch of Wave 3 through September 2018. They did not receive reports between October 2018 and May 2020.

Table 3-1 presents summary statistics for the four customer groups identified in Table 2-1. The number of accounts reflect the actual number of customers who were assigned into each of the treatment groups, and who also had complete datasets at the start of the pilot. The average daily usage during the pre-termination period when all customer groups included in the table were being treated is approximately 23.0 kWh for the Paper + Email group, and 22.2 kWh for the Paper-Only group. The relatively large standard error indicates wide variation in usage patterns across customers. The minimum and maximum usage is also presented. Negative minimum values are due to customer generation flowing back onto the electric grid from net energy metered (NEM) customers— typically due to excess solar generation.

Table 3-1: Pre-Termination Usage by Discontinued, Continued, and Home Energy Report Group

Treatment Group	Report Group	# of Accounts	Daily Usage (kWh)				
			Average	Median	Standard Error	Minimum	Maximum
Discontinued	Paper-Only	27,032	22.2	20.3	10.2	-21.9	228.8
Discontinued	Paper + Email	38,577	23.0	21.4	10.9	-25.8	191.8
Continued	Paper-Only	27,144	22.2	20.2	10.2	-20.0	252.9
Continued	Paper + Email	38,513	23.0	21.5	11.1	-91.7	311.2

Table 3-2, Table 3-3, and Table 3-4 provide similar summary statistics broken out by age, income level, and home type. Average daily energy usage tends to increase with age, particularly after 70 years old. Average daily usage tends to be higher at the lower and higher income brackets, and lower in the middle income brackets. Under home type the single family home is the most prevalent, accounting for nearly 93% of the population.

Table 3-2: Pre-Termination Usage by Age

Age Bin	# of Accounts	Daily Usage (kWh)				
		Average	Median	Standard Error	Minimum	Maximum
Under 20	175	21.1	22.0	8.7	-3.5	54.6
20-29	4,938	20.5	21.0	10.2	-13.7	154.2
30-39	16,782	21.1	21.3	10.1	-21.5	147.7
40-49	27,882	20.8	22.0	11.0	-25.8	311.2
50-59	37,137	21.2	21.3	11.1	-31.9	257.3
60-69	24,261	21.5	20.3	10.7	-91.7	228.8
70-79	12,818	22.6	19.8	10.0	-21.2	252.9
80-89	5,660	22.9	19.0	9.0	-16.7	125.4
90-99	1,613	25.4	18.9	9.1	-7.0	93.1

Table 3-3: Pre-Termination Usage by Income Level

Income Level	# of Accounts	Daily Usage (kWh)				
		Average	Median	Standard Error	Minimum	Maximum
Less than $15,000	3,006	25.9	19.8	9.8	-13.7	186.1
$15,000 - $19,999	3,075	25.4	19.0	9.4	-14.4	154.2
$20,000 - $29,999	6,947	28.3	19.6	9.0	-12.5	138.5
$30,000 - $39,999	7,978	20.2	19.5	9.1	-13.0	205.3
$40,000 - $49,999	10,319	19.4	19.7	9.3	-21.2	188.3
$50,000 - $74,999	31,672	22.6	20.0	9.7	-31.9	252.9
$75,000 - $99,999	24,602	22.9	21.1	10.4	-27.3	311.2
$100,000 - $124,999	13,420	25.9	21.4	10.5	-48.0	202.0
Greater than $124,999	30,247	25.4	23.5	12.6	-91.7	154.7

Table 3-4: Pre-Termination Usage by Home Type

Home Type	# of Accounts	Daily Usage (kWh)				
		Average	Median	Standard Error	Minimum	Maximum
2-4 Unit Duplex/Triplex/Quad	46	25.9	21.7	12.3	12.5	57.0
Apartment	5	25.4	25.5	9.1	15.4	38.7
Condo	21	28.3	22.5	17.0	7.5	67.3
Miscellaneous	583	20.2	18.9	7.7	-6.0	52.5
Mobile Home	4	19.4	20.3	5.9	11.4	25.6
Single Family Dwelling Unit	121,729	22.6	21.0	10.7	-91.7	311.2
Unknown	8,878	22.9	21.2	10.7	-11.9	252.9

The persistence of energy savings from HERs was estimated using a series of regression models. Energy savings were estimated for the continued group for two time periods. The first time period, P1, includes the first twelve months of the pilot (October 2018 through September 2019) and the second time period, P2, includes the remaining eight months before treatment resumed for customers in the discontinued group (October 2019 through May 2020). Energy savings were estimated using a lagged dependent variable model in which monthly energy consumption for continued and control customers will be estimated using consumption data from the pretreatment and treatment periods. The outcome of this model was used to establish baseline energy savings separately for the Paper-Only and Paper + Email populations (along with other customer segments) to which the discontinued groups can be compared. The regression specification is presented here with definitions for each term shown in Table 3-5.

$$kWh_{it} = a + b_t + c_t \cdot treatment_i + d_t \cdot pretreatment\_kwh_{it} + \varepsilon_{it}$$

Table 3-5: Lagged Dependent Variable Model Definitions

Variable	Definition
kWh_{it}	Customer *i*'s usage in month *t*.
a	The estimated constant for energy consumption (average for all customers in all periods).
b_t	The estimated coefficient for the month indicator variable.
c_t	The estimated coefficient for the month indicator variable for treatment customers. This is the treatment effect for a particular month *t*.
$treatment_i$	The treatment indicator variable for customer *i*. Equal to 1 for treatment customers and 0 otherwise.
d_t	The estimated coefficient for pretreatment consumption on a particular month *t*.
$pretreatment\_kwh_{it}$	Pretreatment usage for customer *i* for month *t*. Pretreatment consumption for a particular month in the post treatment period refers to the same calendar month in the pretreatment period.
ε_{it}	The error term.

The second model estimated the difference in energy consumption between the continued and discontinued groups separately for the Paper-Only and Paper + Email populations (along with other customer segments) for P1 and P2. The outcome of this model was the incremental difference in energy savings for each segment. Using a separate model made it possible to determine if the differences in energy savings between the continued and discontinued groups were statistically significant.

The following model specification was used to estimate the difference in energy consumption between the discontinued and continued groups is nearly identical to the specification above, with some small differences:

$$kWh_{it} = a + b_t + c_t \cdot discontinued_i + d_t \cdot pretermination\_kwh_{it} + \varepsilon_{it}$$

Table 3-6: Lagged Dependent Variable Model Definitions

Variable	Definition
kWh_{it}	Customer *i*'s usage in month *t*.
a	The estimated constant for energy consumption (average for all customers in all periods).
b_t	The estimated coefficient for the month indicator variable.
c_t	The estimated coefficient for the month indicator variable for discontinued customers. This is the incremental treatment effect for a particular month *t*.
$discontinued_i$	The discontinued indicator variable for customer *i*. Equal to 1 for discontinued customers and 0 otherwise.
d_t	The estimated coefficient for pre-termination consumption on a particular month *t*.
$pretermination\_kwh_{it}$	Pre-termination usage for customer *i* for month *t*. Pre-termination consumption for a particular month in the post treatment period refers to the same calendar month in the pre-termination period.
ε_{it}	The error term.

3.2 Segmentation and Clustering Analysis

In addition to estimating the persistence of energy savings for the entirety of the persistence pilot population, the persistence of different customer segments was also explored. Customer segments were created based on observable characteristics from the pretreatment period. Particularly, the average daily load shape and average monthly load shapes were used to isolate customers who have peak usage during different hours of the day and at different times of the year, respectively.

To cluster by daily load shape, AMI data was leveraged to estimate each customer's average hourly usage on non-holiday weekdays over the course of the pretreatment year. This yielded an average daily load profile for each customer, which was then normalized by dividing by the customer's total daily load. These customer load shapes, in terms of percentage of total load, were then used as the input into a k-medians clustering algorithm. K-medians clustering identified different usage patterns within the data, shown in Figure 3-1, and grouped customers into one of the three usage groups based on which load shape they align most closely with. Descriptions for each of these groups are presented in Table 3-7. A total of three different groups was selected because this number provided an optimal balance between distinct customer groups and the sample size per group being large enough to allow for meaningful estimates.

Figure 3-1: Daily Load Shape Clusters

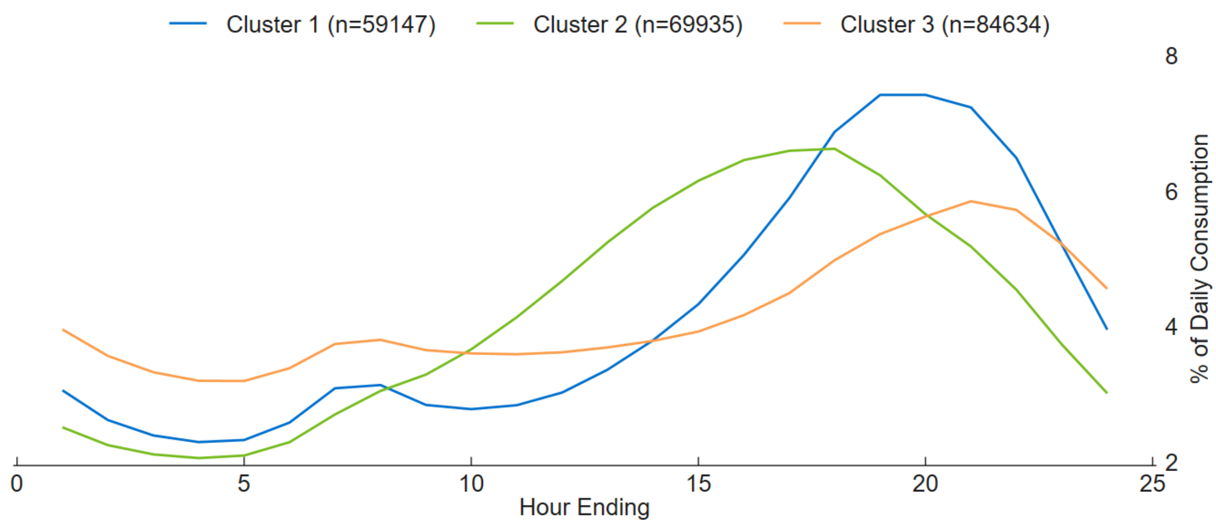

Table 3-7: Daily Load Shape Clusters Identified by K-Medians

Cluster	Cluster Description	Proportion of Customers
1	Typical users: peak usage occurs between 6:00pm and 7:00pm	27.7%
2	Early users: peak usage occurs between 5:00pm and 6:00pm	32.7%
3	Late users: peak usage occurs between 8:00pm and 9:00pm	39.6%

A similar process was followed to cluster customers by seasonal consumption, except that pretreatment data was aggregated at the monthly level for each customer rather than at the hourly level. This monthly data was then normalized for each customer using their total annual consumption. These seasonal customer load shapes were then fed into a k-medians clustering algorithm to produce three distinct usage groups, shown in Figure 3-2. Descriptions of these groups are presented in Table 3-8. A total of three different groups was selected for the seasonal load shape groups because the addition of more groups did not add any variety of seasonal shapes, but with three groups, a distinction could still be made between typical, low, and high summer usage groups.

Figure 3-2: Seasonal Load Shape Clusters

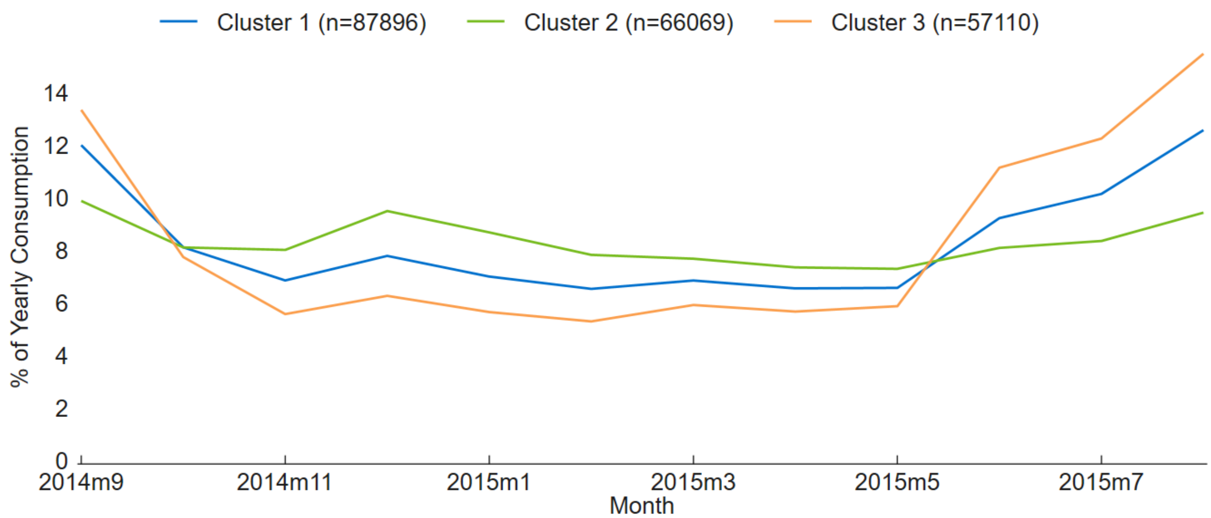

Table 3-8: Seasonal Load Shape Clusters Identified by K-Medians

Group	Group Description	Proportion of Customers
1	Typical users: load slightly higher in summer months	41.6%
2	Constant users: small difference between summer and winter month load	31.3%
3	Summer users: significantly higher usage in summer months	27.1%

Finally, customers were separated into segments based on their age and income. For the household age segmentation, customers were divided into those with an average household age over 50 years old and those under 50 years old. This cutoff was selected to maximize sample sizes. Similarly, households were divided into two income segments: those earning over $75,000 per year and those earning less than $75,000 per year. Energy savings persistence was estimated within each age group and income level using the same methodology described in Section 3.1.

Table 3-9: Age and Income Segmentations

Segmentation Type	Group Description	Proportion of Customers
Age	Average household age under 50 years	43.5%
	Average household age 50 years or older	56.5%
Income	Household income under $75,000	47.6%
	Household income $75,000 or more	52.4%

3.3 Causal Forest Algorithm

There is an emerging area of study in the statistical literature concerning the application of machine learning algorithms to causal inference problems such as the estimation of heterogeneous treatment effects. In particular, Athey and Imbens (2016) introduced the “causal tree” estimator, which applies the concept of regression trees to the estimation of conditional average treatment effects (CATEs). Just as random forests are ensemble models made up of large quantities of regression trees, Wager and Athey (2018) then introduce the “causal forest”, an ensemble of causal trees.

In this section, we apply the causal forest method to the Persistence Pilot data.[4] While direct estimation of the treatment effect (i.e. savings due to receiving a HER) is not feasible, the causal forest provides a model with which we can obtain *predictions* of the individual treatment effects. This large set of predictions grants a much richer understanding of the variation in effect sizes than a single overall average treatment effect (ATE) does. It also allows us to investigate the expected savings over any subset of the population.

Causal forests are a collection of causal trees. To grow a causal tree, first a random subsample of the full set of households is selected. That subsample is split according to a certain value of one of the explanatory variables. For instance, a split might be defined at 100 kWh of average monthly pretreatment consumption. All households with 100 kWh or less would be placed on one side of the split, and the rest on the other. The algorithm searches over all variables and values to find the split that maximizes the difference between the ATEs of each subgroup, subject to certain penalties to ensure the quantities of treatment and control households in each subgroup do not become too imbalanced. Each subgroup is then split again, and the process continues recursively to form a “tree” of splits. If splitting a certain subgroup in any way would not result in a better overall fit, that group is not split and forms a “leaf” at the end of the tree.

Individual trees are prone to overfitting on the subsample of the data on which they were trained. To alleviate this problem, thousands of trees are grown, each on a different subsample of households. To ensure these trees are sufficiently different from one another, each is limited to choosing splits based on a random subset of variables in the overall dataset. This collection of trees is then used to generate predictions using the following method. For every household, we find the set of trees that did not use that household for their initial growth. Then the household traverses each tree, following the splits as appropriate based on its characteristics, until we end up at a leaf node with a corresponding ATE. Finally, this large number of predictions is aggregated into a single estimate of the household’s treatment effect using a weighted average of other household’s predictions, weighted by the frequency with which other households fell into the same leaf node as the household of interest. The result is that households that occur in the same leaf, which will generally be quite similar to one another, get more weight in determining the final prediction.

There are several parameters that affect the estimates produced by the causal forest algorithm. They are particularly sensitive to the *minimum node size*, which defines the smallest number of households that are allowed to make up a leaf. In our analysis, we select this threshold by

[4] The additional eight-month analysis period which includes the months from October 2019 through May 2020 (P2) is not included in the causal forest analysis.

searching the parameter space and choosing the value that minimizes the 5-fold cross-validated error.

To assess the persistence of treatment effects, we grew a total of four causal forests, each using a different response variable:

1. Pretreatment vs combined pre- and post-termination consumption

$$y = \overline{kWh}_{preterm\,\&\,postterm} - \overline{kWh}_{pretrt}$$

2. Pretreatment vs pre-termination consumption

$$y = \overline{kWh}_{preterm} - \overline{kWh}_{pretrt}$$

3. Pretreatment vs post-termination consumption

$$y = \overline{kWh}_{postterm} - \overline{kWh}_{pretrt}$$

4. Pre-termination vs post-termination consumption

$$y = \overline{kWh}_{postterm} - \overline{kWh}_{preterm}$$

Note that in models 1 through 3, the "treatment" refers to whether a household initially received a home energy report. The treatment customers are compared with customers who never received HERs (the control group). In model 4, however, households who have never received a HER (the control group) are dropped from the dataset, and the "treatment" refers to whether HER delivery to a household was discontinued. In other words, the customers whose paper HERs were discontinued are the treatment group, and the customers who continued to receive paper HERs act as the control group.

4 Persistence of Energy Savings

This report section summarizes the energy savings impacts for the different test cells in the persistence pilot (continued and discontinued, Paper-Only, and Paper + Email). Energy savings were estimated for the post-termination periods for each group. The pre-termination period includes the twelve months prior to cessation of paper reports. The post-termination period is divided into two time periods, P1 and P2. P1 represents the first year of the pilot (October 2018 through September 2019) and P2 includes an additional eight-month period from October 2019 through May 2020. Discontinued customers resumed treatment with paper reports in June 2020, marking the end of the pilot.

Figure 4-1 shows the daily kWh savings estimates for each month in the pre-termination and post-termination periods for all customers in the study. A negative effect of discontinuing HERs represents a reduction in savings. There were small differences in energy savings between the continued and discontinued groups during the pre-termination period. Given that the discontinuation of reports was randomly assigned and that the differences are not statistically significant, it is likely that these variations in savings are due to random chance.

After May 2019, the discontinued group shows smaller savings compared to the continued group. However, the difference is not statistically significant in any month in P1, and is only statistically significant in three months of P2 (November 2019, December 2019, and May 2020). This indicates that the savings from paper HERs persisted for more than one year.

Figure 4-1: Daily kWh Savings and Savings Impact by Month – All Customers

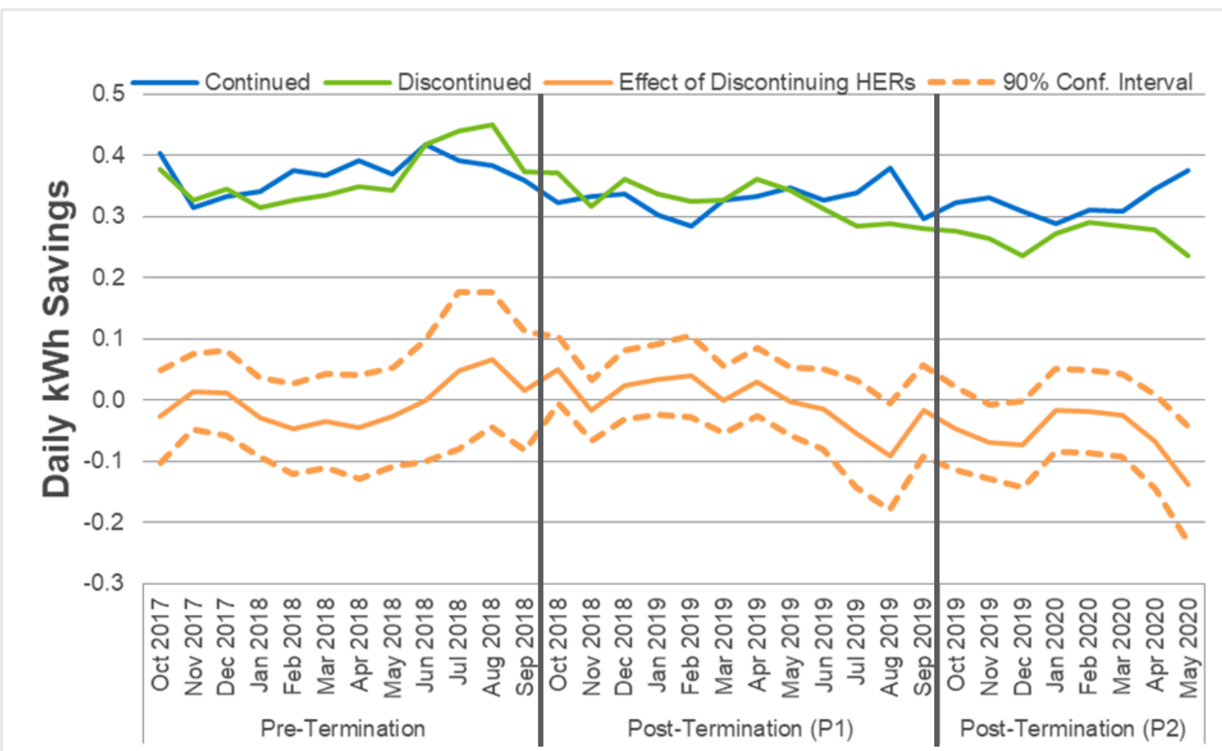

Figure 4-2 presents post-termination annual kWh energy savings for the full persistence pilot population. Average energy savings per customer are presented separately for P1 and P2. Energy savings are lower in P2 because the estimate only includes eight months, rather than one full year. The blue bars represent the continued group, and the green bars represent the discontinued group.

Baseline energy savings for the combined continued group during the first year of the persistence study (P1) were equal to 119.7 kWh, or about 0.3 kWh per day on average. Savings for the discontinued group were estimated to be 117.9 kWh (also about 0.3 kWh per day). The difference between the blue and green bars, and indications regarding the statistical significance of the difference, are presented in Figure 4-3.

Figure 4-2: Post-Termination Cumulative Energy Savings, per Customer

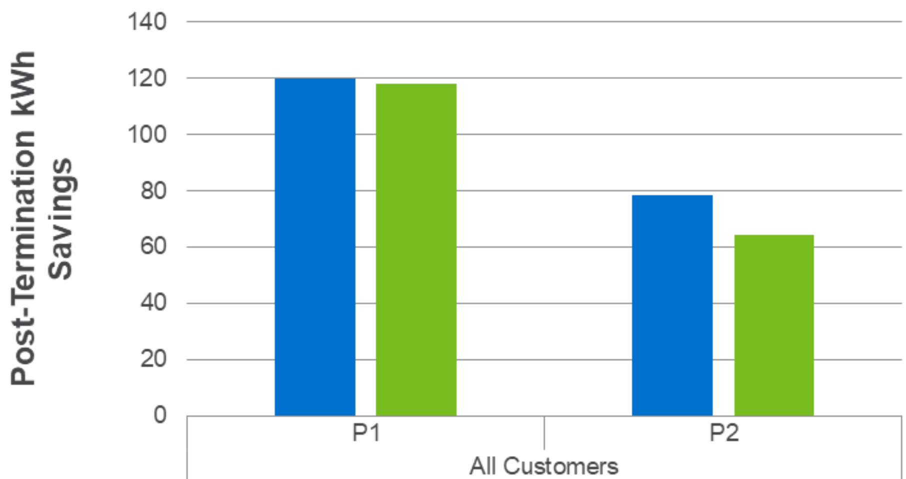

	P1	P2
Continued	119.7 kWh	78.6 kWh
Discontinued	117.9 kWh	64.4 kWh

Figure 4-3 shows the level of persistence for the first twelve months of the persistence pilot (P1) and the additional eight-month period (P2). The lines bisecting the top of the orange bar in the figure show the 90% confidence band for the persistence estimate. If the confidence band includes 100%, it means the estimated difference in energy savings between the continued and discontinued groups is not statistically different from 0 at the 90% level of confidence.

Discontinued customers exhibited savings equal to 98% of the continued group savings in P1 (October 2018 through September 2019) and 82% in P2 (October 2019 through May 2020). While the level of persistence is trending downward from P1 to P2, the difference in energy savings between the continued and discontinued groups was not statistically significant in either time period. In other words, the savings attributable to paper HERs persist for at least one year and eight months after discontinuation when viewed from this high-level time perspective.[5]

Figure 4-3: Persistence of Energy Savings

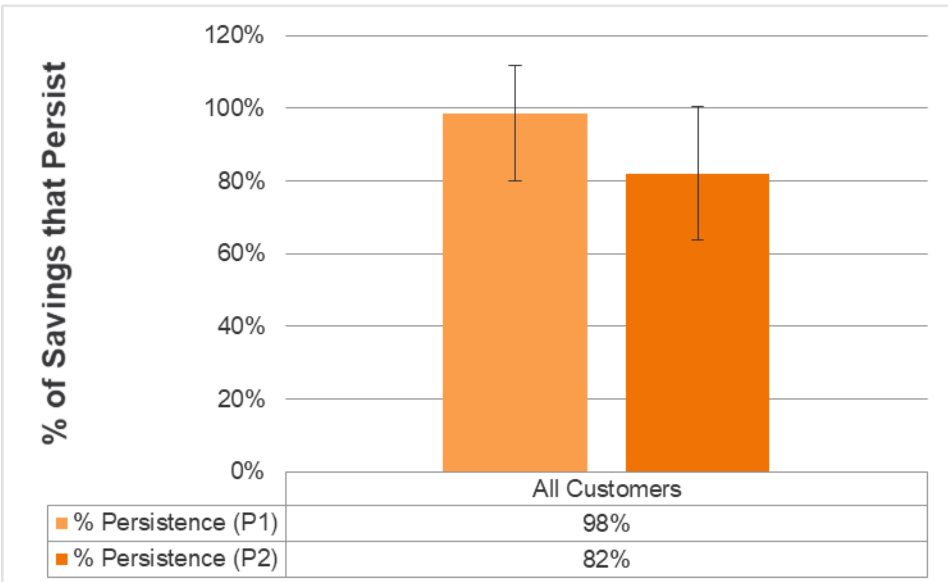

[5] As noted above, some individual months showed statistically significant differences.

Figure 4-4 presents the post-termination energy savings estimates for the Paper-Only and Paper + Email populations separately. Customers who received electronic HERs in addition to paper HERs had greater energy savings in both time periods, which is likely due to a combination of differences in treatment and differences in the underlying populations.

Figure 4-4: Post-Termination Energy Savings by Treatment Type, per Customer

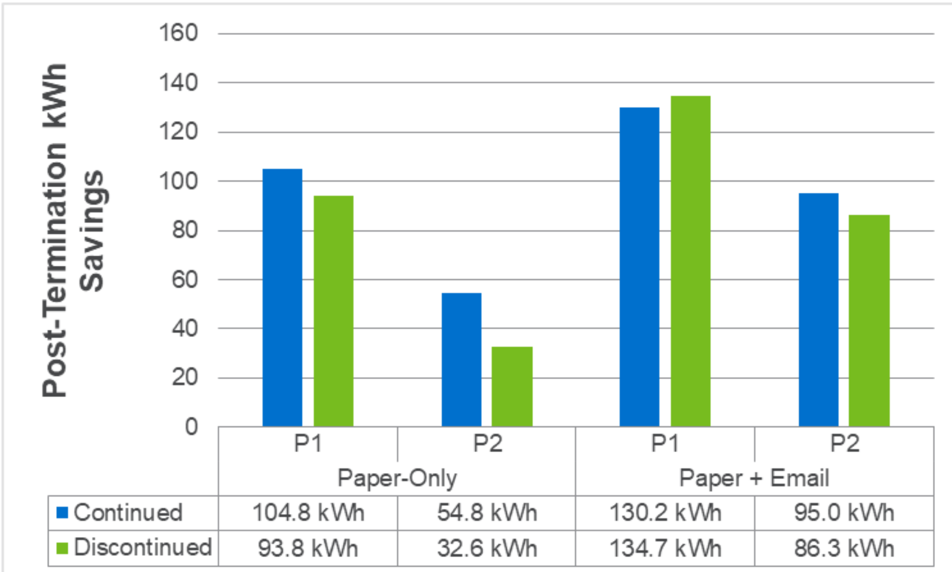

	Paper-Only P1	Paper-Only P2	Paper + Email P1	Paper + Email P2
Continued	104.8 kWh	54.8 kWh	130.2 kWh	95.0 kWh
Discontinued	93.8 kWh	32.6 kWh	134.7 kWh	86.3 kWh

Figure 4-5 shows the level of persistence in the first year of the post-termination period (P1) and an additional eight-month period (P2). Customers who only received paper HERs showed a small decline in energy savings in P1 (10%), and a larger decline in P2 (40%). The decline in energy savings was only statistically significant in P2 (October 2019 through May 2020). The Paper + Email group, who continued to receive electronic HERs, exhibited a small, but not statistically significant increase in energy savings during P1. The decline in energy savings in P2 was small and not statistically significant for the Paper + Email population.

Figure 4-5: Persistence of Energy Savings by Treatment Type

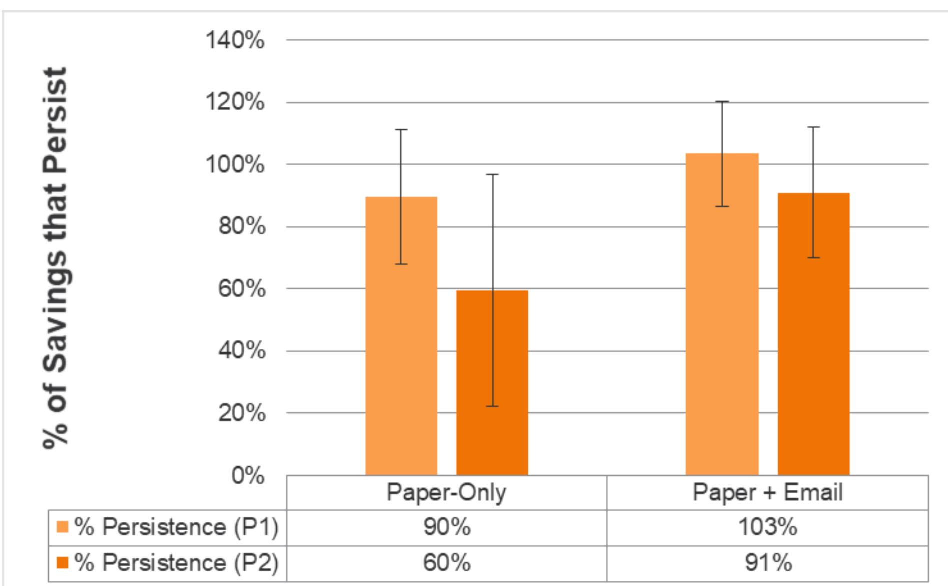

	Paper-Only	Paper + Email
% Persistence (P1)	90%	103%
% Persistence (P2)	60%	91%

A key research question asks how the persistence effect varies among those who receive email HERs and those who do not. Because the group of customers who provide email addresses to SCE do not have the same baseline energy savings as those who do not as shown in Table 3-1, the two groups are not directly comparable. While the difference in pre-termination energy usage is between the Paper-Only and the Paper + Email groups is relatively small, the difference is statistically significant. There were also differences in savings levels between the two groups as well. For example, continued customers in the Paper + Email group had greater savings in the post-termination period. Continued Paper + Email customers saved 130.2 kWh in the first twelve months of the post-termination period (P1), while continued Paper-Only customers saved 104.8 kWh.

Although the two treatment populations had differences in baseline energy savings, they had similar levels of persistence in P1. The difference in the levels of persistence between the Paper-Only and Paper + Email groups (90% and 103%, respectively) is small and not statistically significant. While being cognizant of the dissimilarities in underlying populations, this small difference in energy savings persistence indicates that layering electronic HERs with paper ones does not lead to more persistent savings.

5 Segmentation Results

Figure 5-1 presents post-termination energy savings for each of the three daily load shape clusters, and Figure 5-2 presents the level of persistence within each cluster. Paper + Email early and late peaking customers had similar baseline energy savings in P1 (124.4 kWh per year). Typical usage customers had slightly smaller annual baseline savings during the same period (109.5 kWh).

Figure 5-1: Post-Termination Energy Savings by Load Shape Cluster, Per Customer

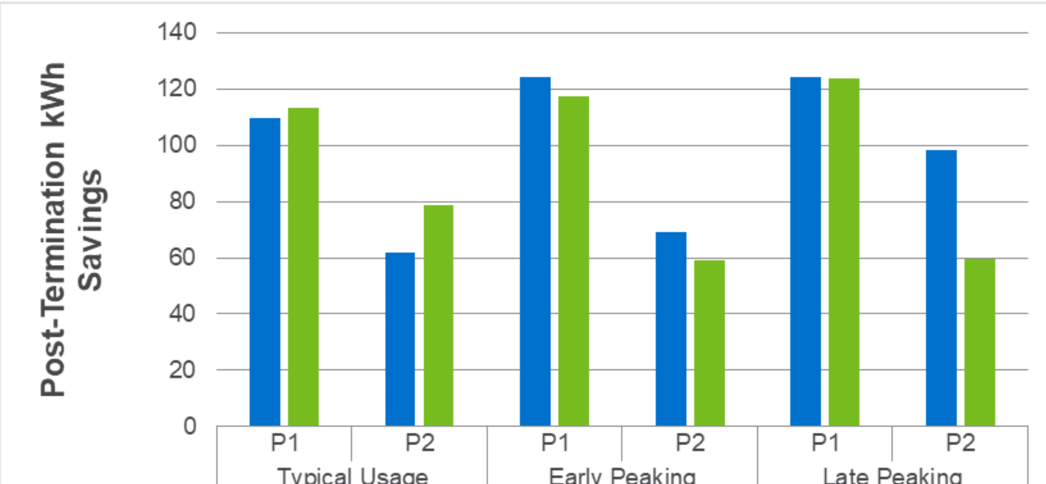

	Typical Usage P1	Typical Usage P2	Early Peaking P1	Early Peaking P2	Late Peaking P1	Late Peaking P2
Continued	109.5 kWh	62.0 kWh	124.4 kWh	69.3 kWh	124.4 kWh	98.3 kWh
Discontinued	113.4 kWh	78.6 kWh	117.2 kWh	59.0 kWh	123.6 kWh	59.4 kWh

The level of persistence is not statistically significantly different from 100% for most of the three load shape clusters for either time period. The exception is customers in the late-peaking cluster, who showed a statistically significant decline in energy savings in P2 (a persistence level of 60%). In P1, the differences in persistence levels between clusters are small and not statistically significant. In P2, the differences between the clusters are more notable; the difference between the typical usage and late-peaking customers is statistically significant, indicating that late-peaking customers have less persistent savings than typical energy users.

Figure 5-2: Persistence of Energy Savings by Load Shape Cluster

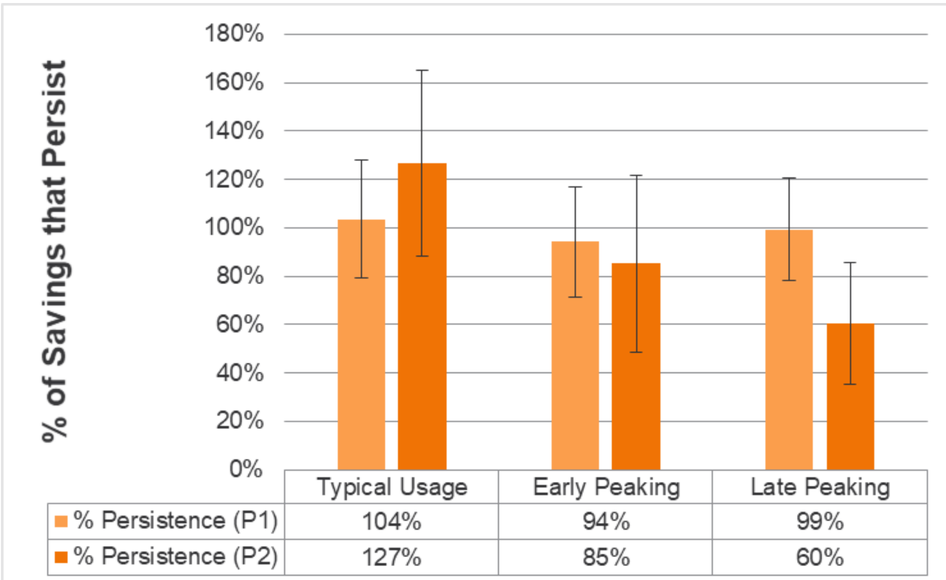

Figure 5-3 presents the post-termination energy savings for each seasonal usage cluster for P1 and P2. The typical usage cluster had the smallest baseline energy savings during the first year of the study (79.7 kWh per year) and the constant usage segment had the greatest baseline energy savings (145.9 kWh per year).

Figure 5-3: Post-Termination Energy Savings by Seasonal Usage Cluster, per Customer

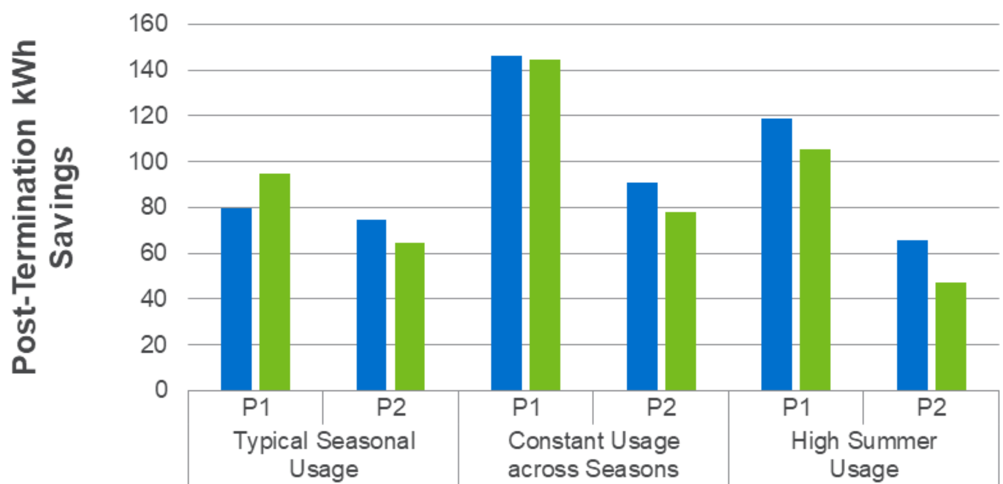

	Typical Seasonal Usage P1	Typical Seasonal Usage P2	Constant Usage across Seasons P1	Constant Usage across Seasons P2	High Summer Usage P1	High Summer Usage P2
Continued	79.7 kWh	74.8 kWh	145.9 kWh	91.0 kWh	118.7 kWh	65.4 kWh
Discontinued	94.6 kWh	64.7 kWh	144.3 kWh	77.8 kWh	105.2 kWh	47.3 kWh

Figure 5-4 presents the level of energy savings persistence (relative to continued customers) in each seasonal usage cluster. No segment had statistically significant reductions in savings in either post-termination period, indicating that the savings persist in each seasonal usage category. In fact, customers in the typical seasonal usage cluster increased their annual savings during P1, but not by a statistically significant amount. Additionally, the differences in the levels of persistence across clusters is not statistically significant. In other words, one seasonal usage cluster does not show more persistent savings than another.

Figure 5-4: Persistence of Energy Savings by Seasonal Usage Cluster

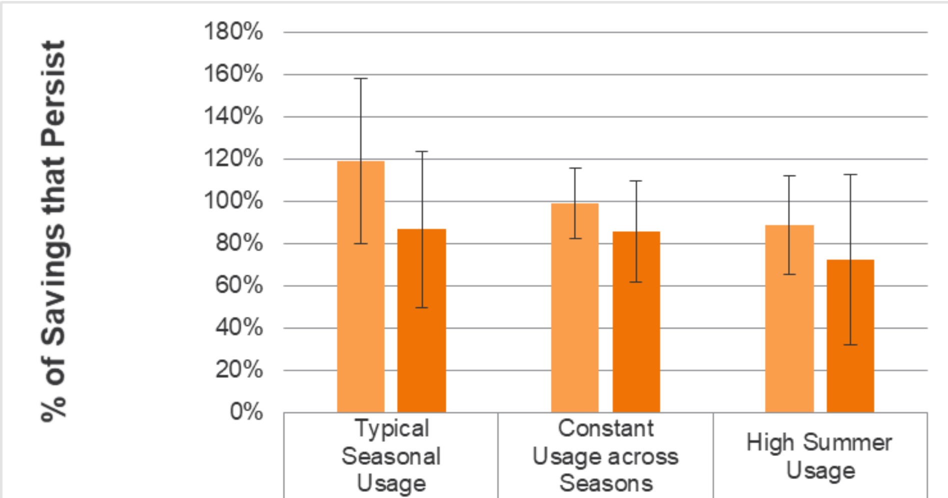

	Typical Seasonal Usage	Constant Usage across Seasons	High Summer Usage
% Persistence (P1)	119%	99%	89%
% Persistence (P2)	86%	85%	72%

Figure 5-5 shows post-termination energy savings for two age groups: customers under 50 years old and customers over 50 years old. On average, it appears that the older customer segment saves more energy than their younger counterparts (about 163.3 kWh per year in P1 versus 50.3 kWh per year, respectively).

Figure 5-5: Post-Termination Energy Savings by Age, per Customer

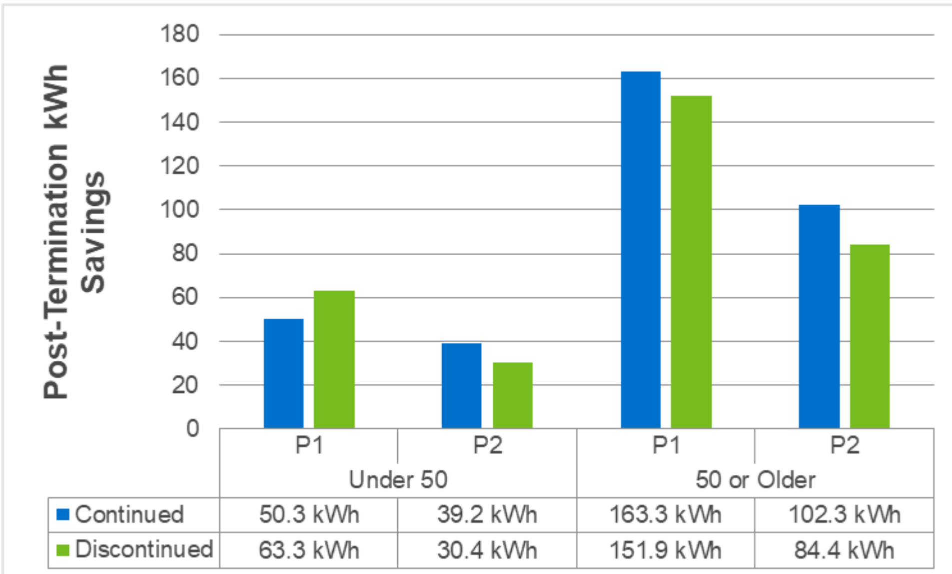

While the group of customers with an average household age of under 50 years old appears to show an increase in savings in P1 of the discontinuation period, it is important to note that the confidence band on the estimate is very wide and includes 100%. The group of customers with an average household age of over 50 years old showed a statistically significant decline in energy savings in P2, with a persistence level of 82%.

Figure 5-6: Persistence of Energy Savings by Age

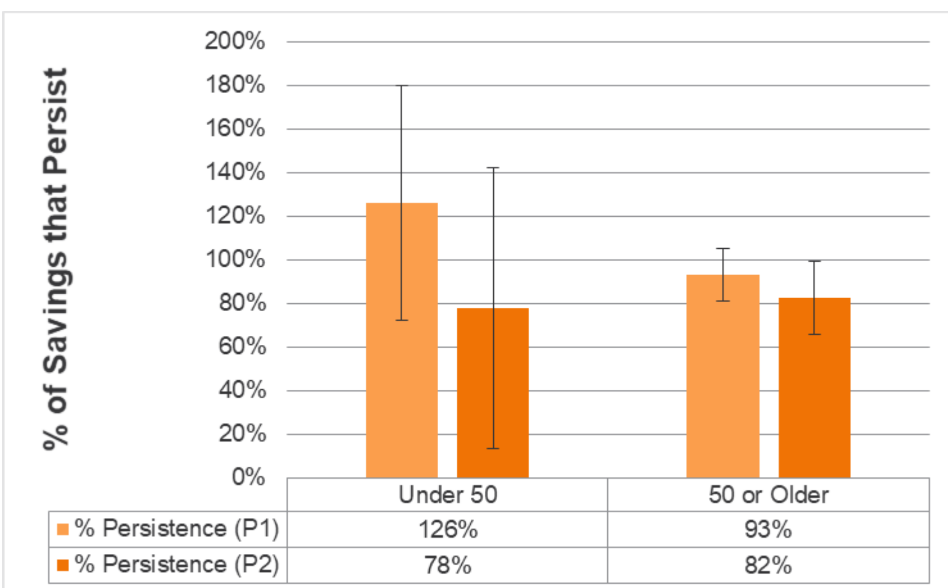

Finally, Figure 5-7 presents the post-termination energy savings for two income levels during P1 and P2. The group with incomes less than $75,000 per year had smaller energy savings versus households who earned more than $75,000 per year (91.3 kWh per year compared to 145.5 kWh per year in P1).

Figure 5-7: Post-Termination Energy Savings by Income, per Customer

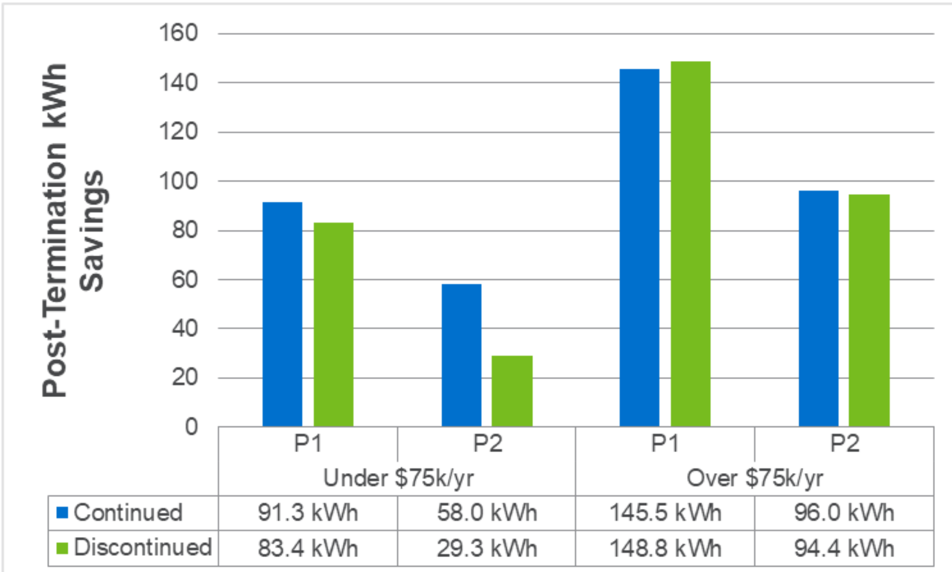

	P1	P2	P1	P2
	Under $75k/yr		Over $75k/yr	
Continued	91.3 kWh	58.0 kWh	145.5 kWh	96.0 kWh
Discontinued	83.4 kWh	29.3 kWh	148.8 kWh	94.4 kWh

Figure 5-8 shows the percent of energy savings that persist in the first year of the pilot (P1) and the additional eight-month period (P2) for each household income segment. Customers with household incomes under $75k per year had high levels of persistence in P1 (91%) but showed a statistically significant decline in P2 (51%). Customers with higher incomes had very small changes in energy savings that were not statistically significant in either post-termination period (102% and 98%).

Figure 5-8: Persistence of Energy Savings by Income

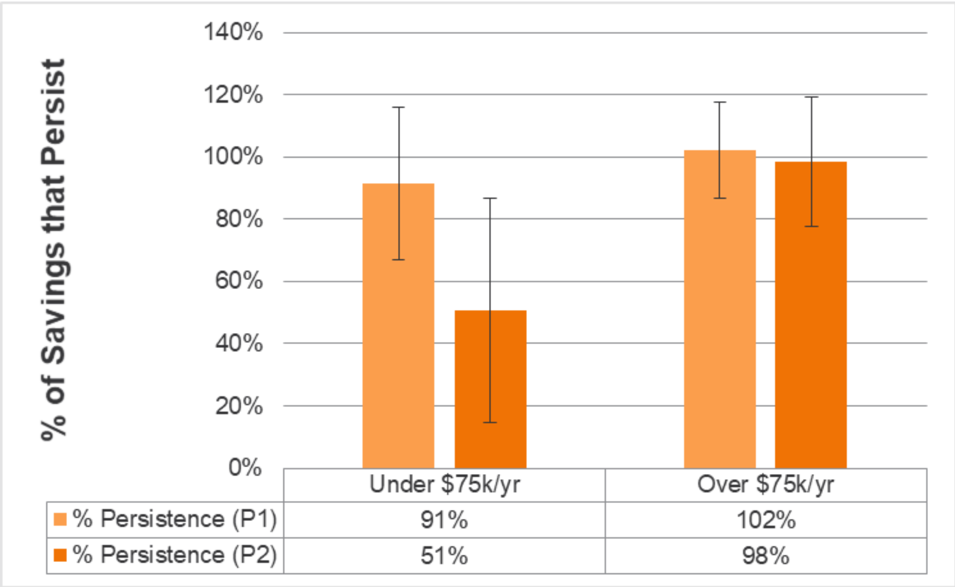

The findings in this section indicate that differences in income, age, and certain load shape patterns did not influence the persistence of energy savings during the first year of the pilot. However, there were notable differences in overall savings attributable to these customer characteristics. In order to gain additional insights into the patterns of persistence and overall savings, an additional approach that is able to test many more customer characteristics systematically through a machine learning algorithm is included in the following section.

6 Causal Forest Results

The causal forest analysis offers valuable insights into behavior that aren't available through the segmentation and regression-based approach in the prior sections. One of the primary benefits of the causal forest analysis is that model itself is able to identify the most meaningful way to separate customers into groups for comparison. This is done quickly, and iteratively, and allows for identification of patterns that may not otherwise be immediately obvious to the researcher.

Post-treatment savings based on the causal forest analysis during the pre-termination period are estimated to be 0.33 kWh per day (+/- 0.06 kWh per day).[6] This is consistent with the difference-in-differences estimation presented in Section 4. Figure 6-1 shows the distribution of all household-level predicted savings. This reveals much more about the differing responses to treatment than the single value provided by the overall average treatment effect (ATE). There are two peaks in this distribution, indicating a substantial group of households that are estimated to have seen savings of around 0.45 kWh and another group that saw savings of only around 0.10 kWh. A small number of households saw predicted savings of 1 kWh or more, corresponding to 365 kWh over a full year. Conversely, 9% of households were actually predicted to *increase* consumption following treatment.

Figure 6-1: Distribution of Predicted Daily Energy Savings

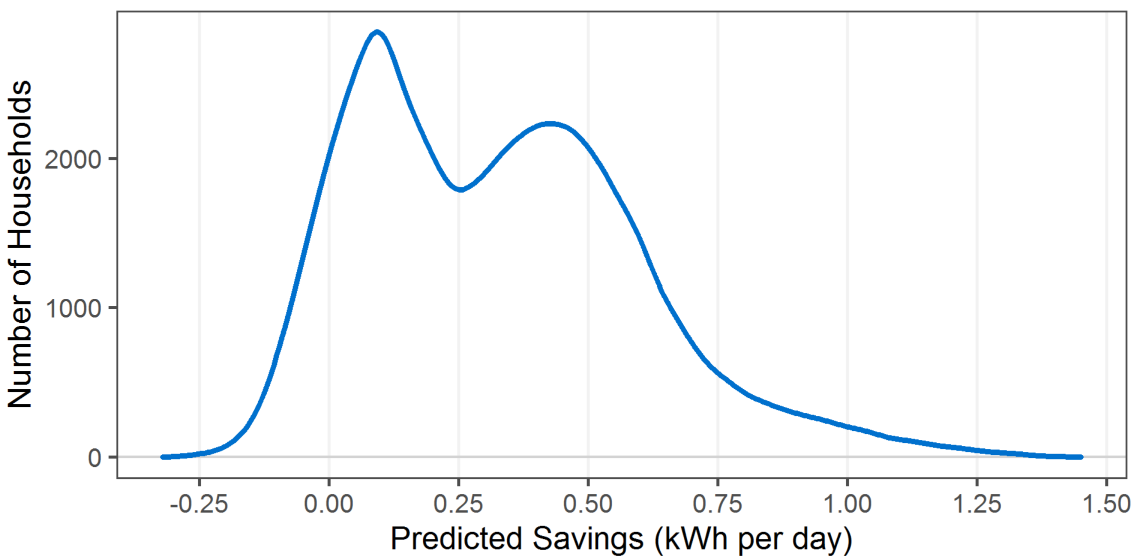

[6] The additional eight-month analysis period which includes the months from October 2019 through May 2020 (P2) is not included in the causal forest analysis.

Figure 6-2 shows the distribution of predicted savings by treatment type. The orange distribution includes customers in the Paper + Email population and the blue distribution includes customers in the Paper-Only population. The bimodal distribution shown in Figure 6-1 is driven by a similar distribution in the Paper + Email group. The Paper-Only population, on the other hand, does not have this characteristic. Pre-termination savings based on the causal forest analysis are estimated to be 0.35 kWh per day (+/- 0.08 kWh per day) for the Paper + Email groups and 0.28 kWh per day (+/- 0.08 kWh per day) for the Paper-Only group.

While the Paper-Only group has smaller predicted savings on average, the distribution is wider than the Paper + Email group and includes more customers with high predicted savings (over 1.5 kWh per day) and low predicted savings (less than -0.25 kWh per day).

Figure 6-2: Distribution of Predicted Daily Energy Savings by Treatment Type

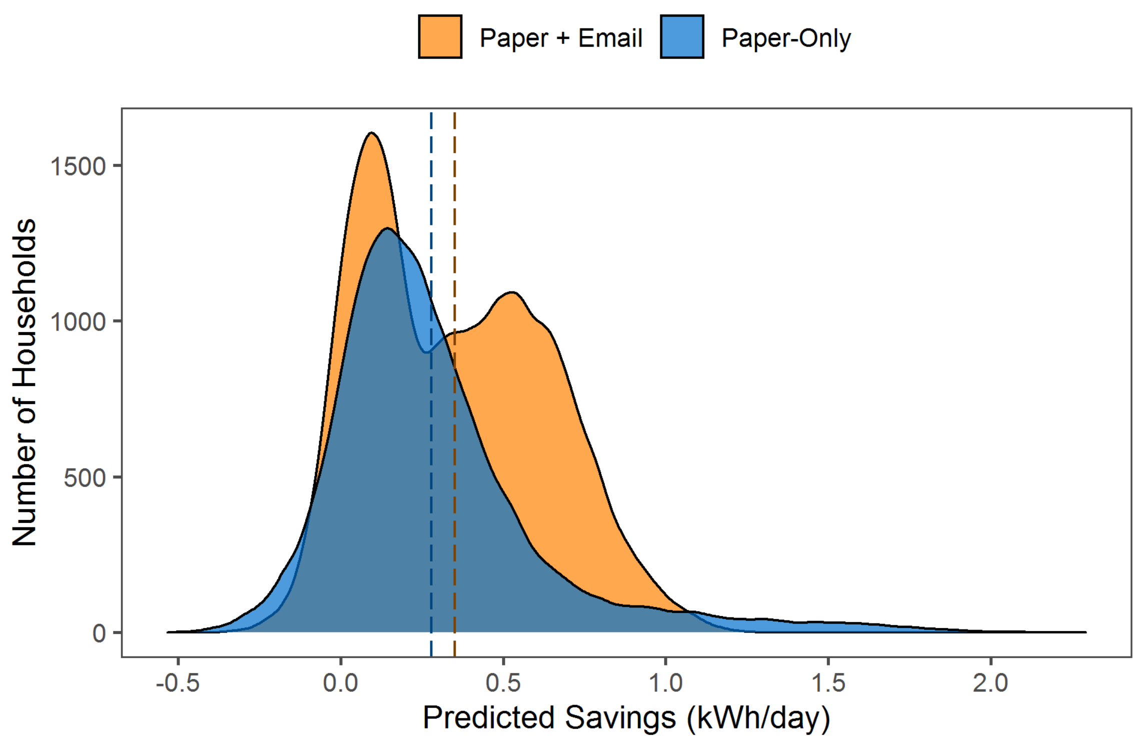

The causal forest algorithm finds splits in the data that do the best job of partitioning households into high and low savers. As more trees are grown, the variables that are frequently used for the first few splits in each tree can be interpreted as those that are most “important”, in that they consistently yield the largest separation among savings estimates. Figure 6-3 summarizes the relative importance of a subset of model inputs. The baseline pretreatment consumption, “pre_kwh”, is the most important variable by this metric, while age and income bracket are a distant second and third, respectively. Note that for readability, some variables that were less important overall have been omitted from this figure.

Figure 6-3: Importance of Select Variables regarding Model Accuracy

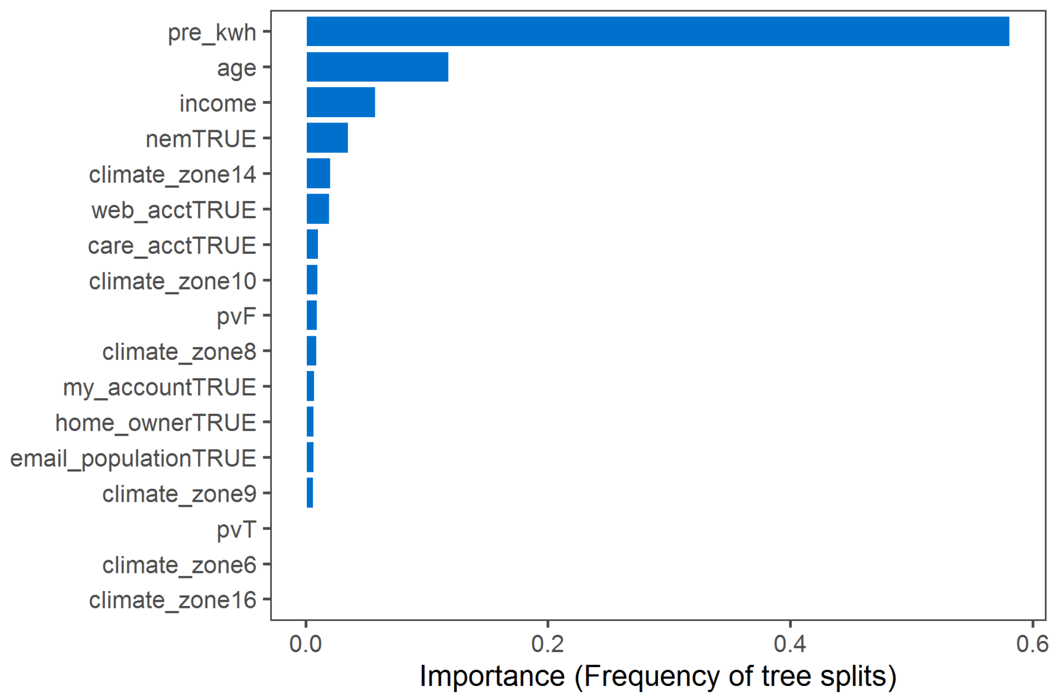

Figure 6-4 shows the relative importance of a subset of model inputs for the Paper + Email and Paper-Only populations. Baseline pretreatment consumption is the most important variable for both groups, followed by age. Positive NEM status is the third most important variable for customers in the Paper + Email population, but income is more important for the Paper-Only group. However, there is a notable amount of overlap in the most important variables between the two populations.

Figure 6-4: Importance of Select Variables regarding Model Accuracy by Treatment Type

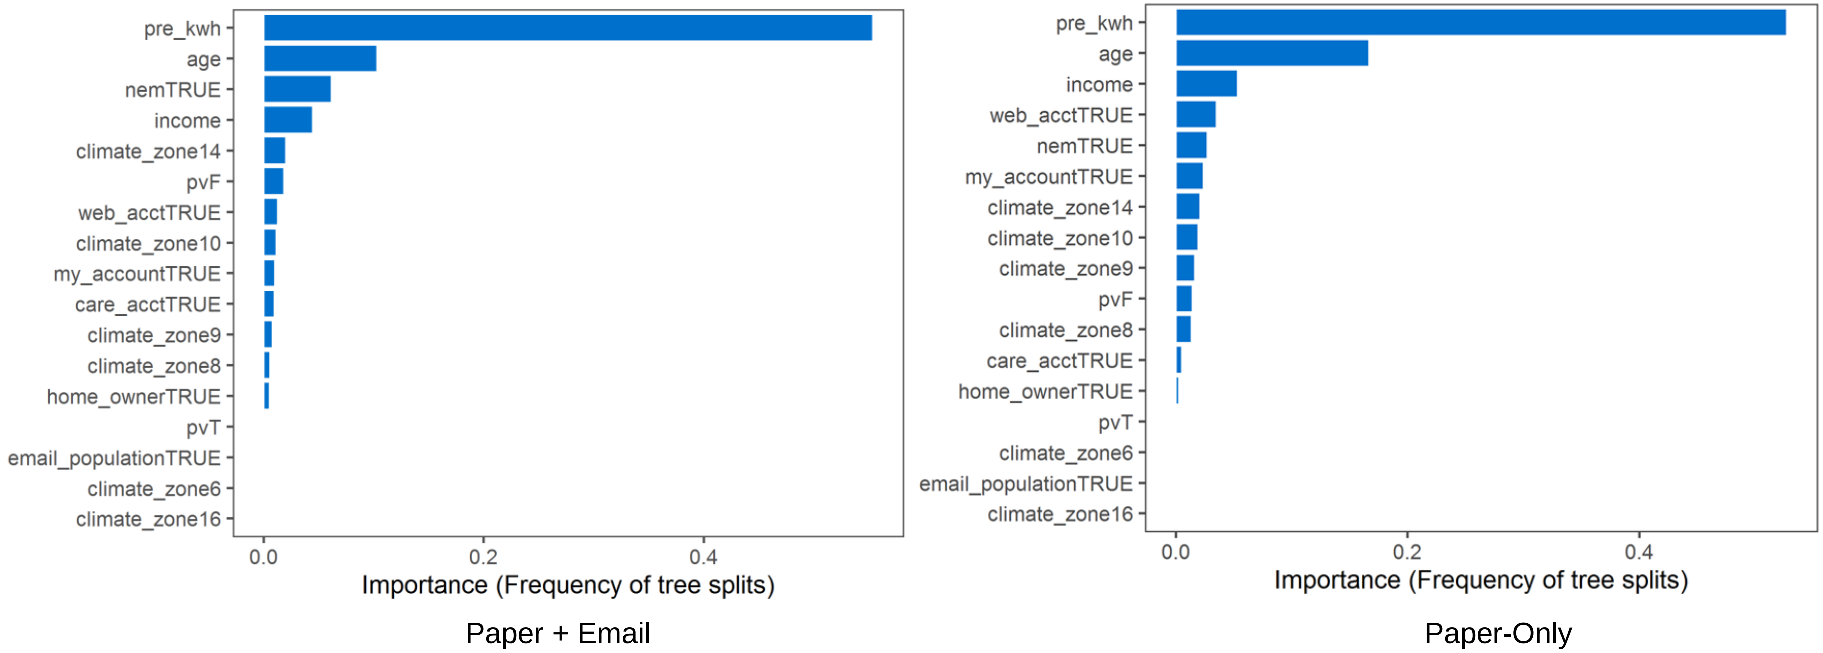

The advantage of the model-based estimation approach is that the relationship between savings and any number of other characteristics can be examined using household-level predicted savings. The figures that follow focus on this relationship for the three most important variables identified above: pretreatment consumption, head of household age, and household income.

Figure 6-5 shows the predicted treatment effect for each household against that household's pretreatment consumption. The number of customers represented at each point is indicated by the color scale on the right, with dark blue points including zero to 200 customers, and yellow points including more than 600. Low-consumption customers circled in green are predicted to have the least amount of savings. In fact, nearly all households which were predicted to increase usage (circled in red) fell below a daily pretreatment average consumption of around 20 kWh. In the orange circle, savings increase as per-treatment consumption increases up until about 45 kWh, at which point it levels off. Higher levels of pretreatment consumption after 45 kWh per day are not correlated with higher levels of savings.

Figure 6-5: Predicted Energy Savings vs. Pretreatment kWh, per Customer

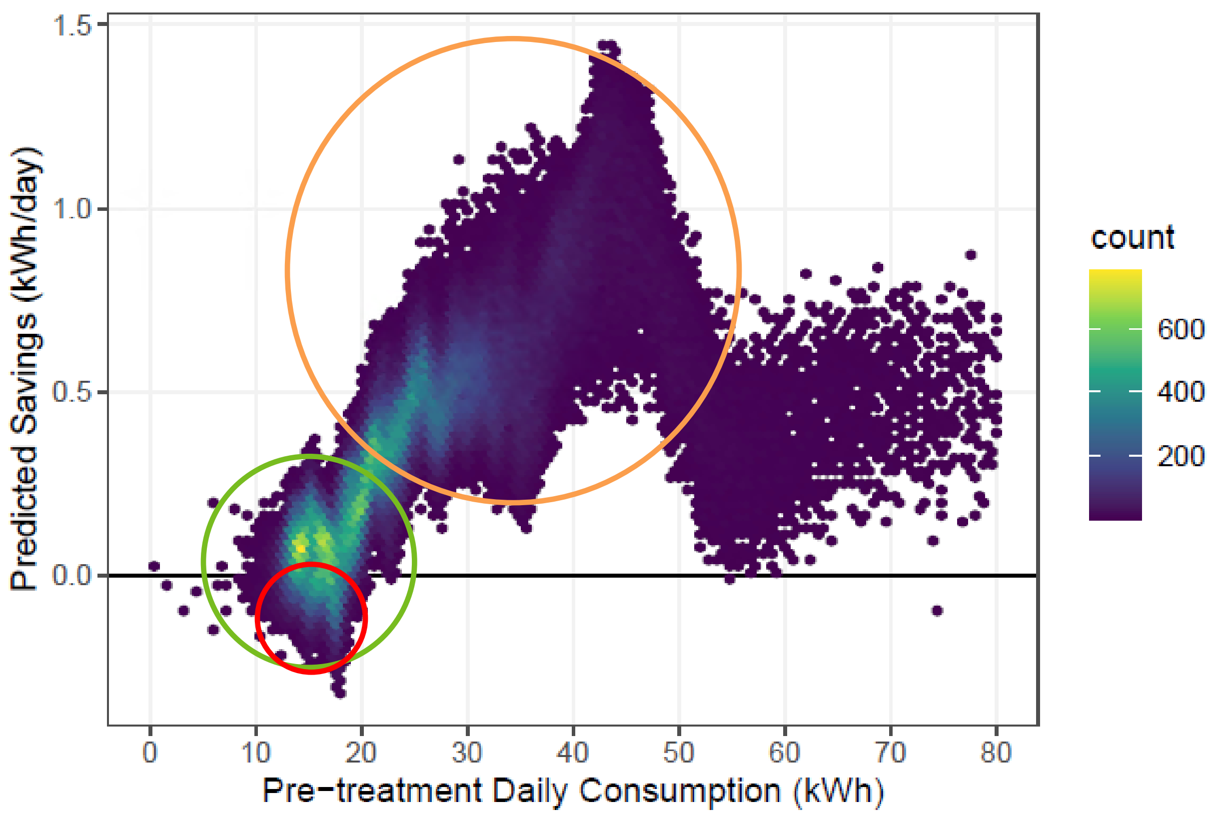

Figure 6-6 shows the predicted treatment effect versus pretreatment consumption for each population (Paper + Email and Paper-Only). There are several key differences between the two groups, the first being a wider distribution in predicted savings among Paper-Only customers versus Paper + Email. Second, there is a stronger positive relationship between pretreatment consumption and predicted savings in the Paper-Only group, especially at higher values of pretreatment consumption (over 35 kWh per day). Finally, nearly all Paper-Only customers with pretreatment consumption greater than 45 kWh per day are predicted to have positive savings. Paper + Email customers, on the other hand, have a group of high pretreatment energy users with negative savings effects.

Figure 6-6: Predicted Energy Savings vs. Pretreatment kWh by Treatment Type, per Customer

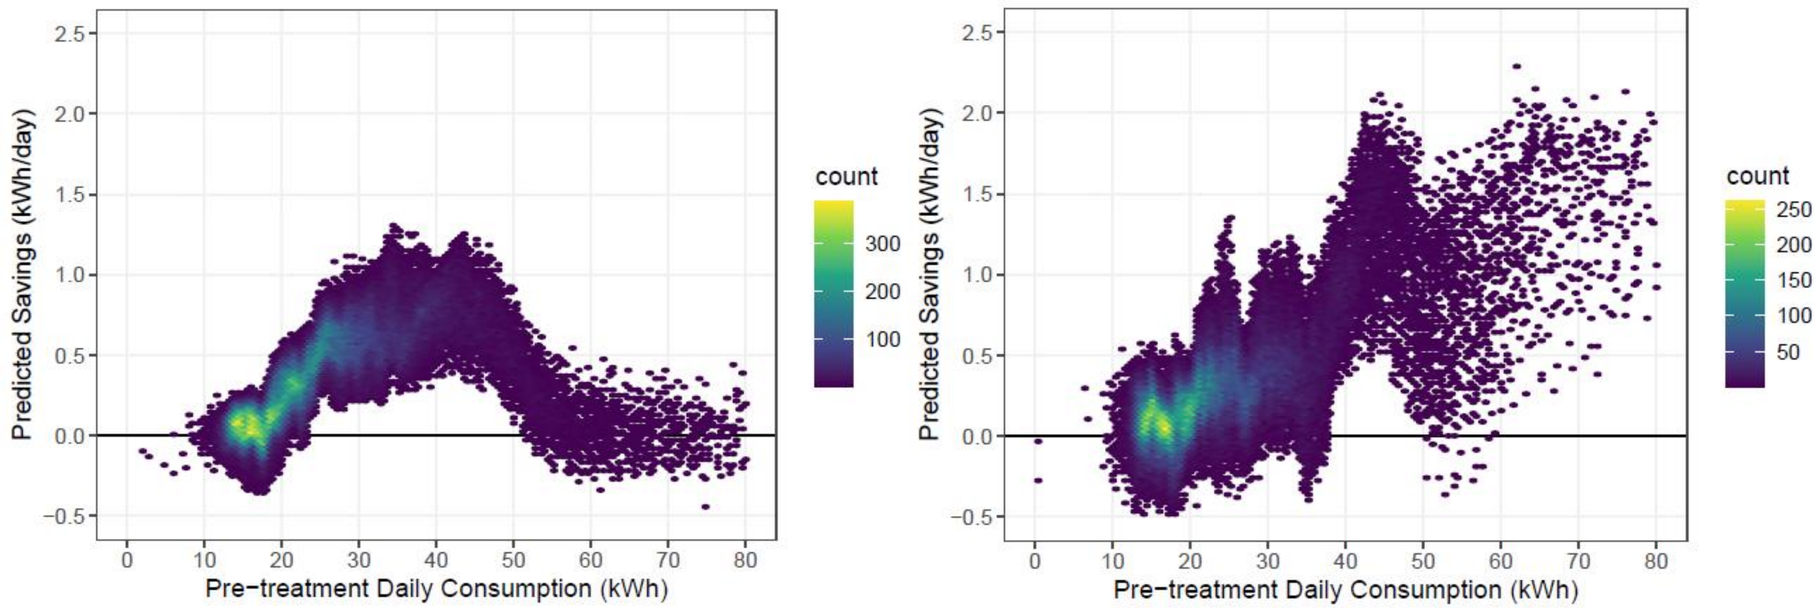

Paper + Email Paper-Only

Based on the findings above, it would appear that there may be an opportunity to be more selective regarding which customers receive treatment in future waves. Notably, the customers using less than 20 kWh daily do not provide high levels of savings on a per-customer basis. In fact, a significant portion of these customers are increasing usage rather than savings—customers (circled in red in Figure 6-5). That said, the population of customers using less than 20 kWh per day is nearly 40% of the overall Wave 3 population, as shown in Figure 6-7. Therefore, it is important to further examine this population to understand their contribution to the aggregate annual energy savings for the treatment wave.

Figure 6-7: Distribution of Pretreatment Daily Consumption (Paper + Email and Paper-Only Populations Combined)

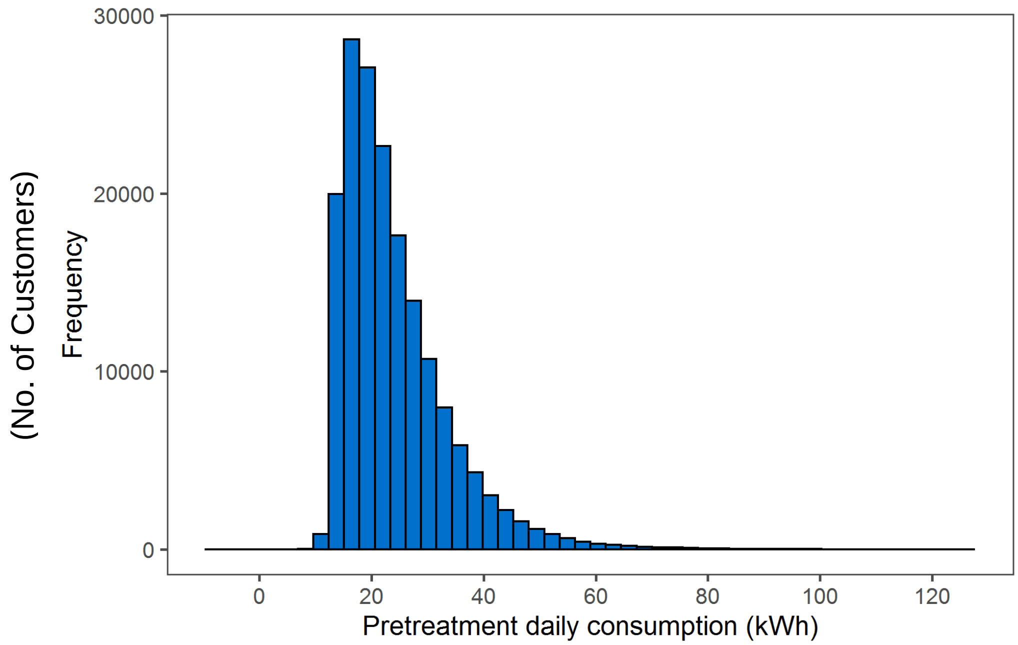

Table 6-1 presents contribution to the annual MWh savings from the Wave 3 customers split out by pretreatment usage level. Customers who typically use less than 20 kWh per day account for 41.3% of the population. However, they only account for 10.1% of the savings. Conversely, those customers using more than 20 kWh per day represent 58.7% of the population, and nearly 90% of the savings. From a cost effectiveness perspective, it is worth considering exclusion of customers using less than 20 kWh per day from future treatment waves in order to decrease treatment costs and increase the average savings per customer.

Table 6-1: Contribution to Annual MWh Savings by Pretreatment Usage Level

Pretreatment Average Daily Usage	Number of Customers	% of Customers	Annual Aggregate Savings (MWh)	% of Aggregate Annual Savings
Less than 20 kWh	54,167	41.3%	1,573	10.1%
Greater than 20 kWh	76,975	58.7%	14,022	89.9%
Total	**131,142**	**100.0%**	**15,595**	**100.0%**

Figure 6-8 presents average predicted daily energy savings by age and income. Each cell in the grid is shaded according to the average predicted savings value for the corresponding age and income categories. Yellow indicates a treatment effect closer to zero (less savings), and blue indicates effects with larger magnitude (greater savings). This enables inspection of how savings vary across multiple characteristics at once. It is clear that savings are coming primarily from incomes greater than $50,000 and ages 55 to 65. On the other hand, young and low-income households generate the least amount of savings according to the model. However, age and income are likely greatly correlated with pretreatment consumption.

Figure 6-8: Average Predicted Energy Savings by Age and Income

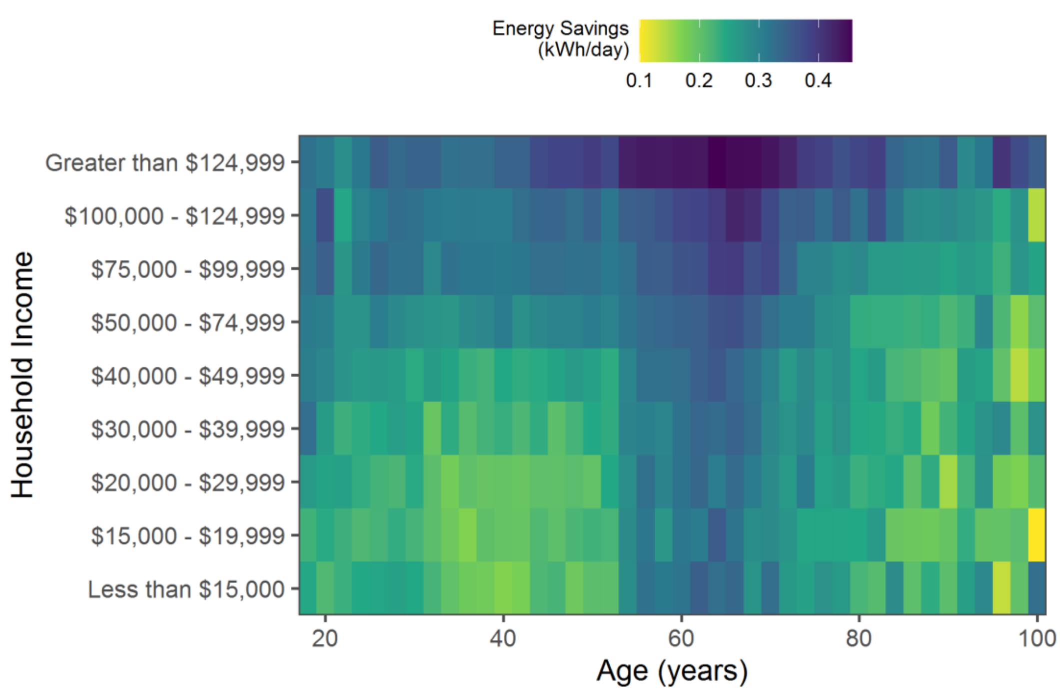

Figure 6-9 shows average predicted daily energy savings by age and income for the Email + Paper and Paper-Only populations separately. There is a striking difference between the two groups. Generally speaking, the lower-income Paper-Only customers who generate savings over 0.1 kWh per day fall between ages 40 and 80. In the Paper + Email population, customers of all ages generate savings at lower incomes. In this group, the relationship between savings, age, and income is not as apparent.

Figure 6-9: Average Predicted Energy Savings by Age, Income, and Treatment Type

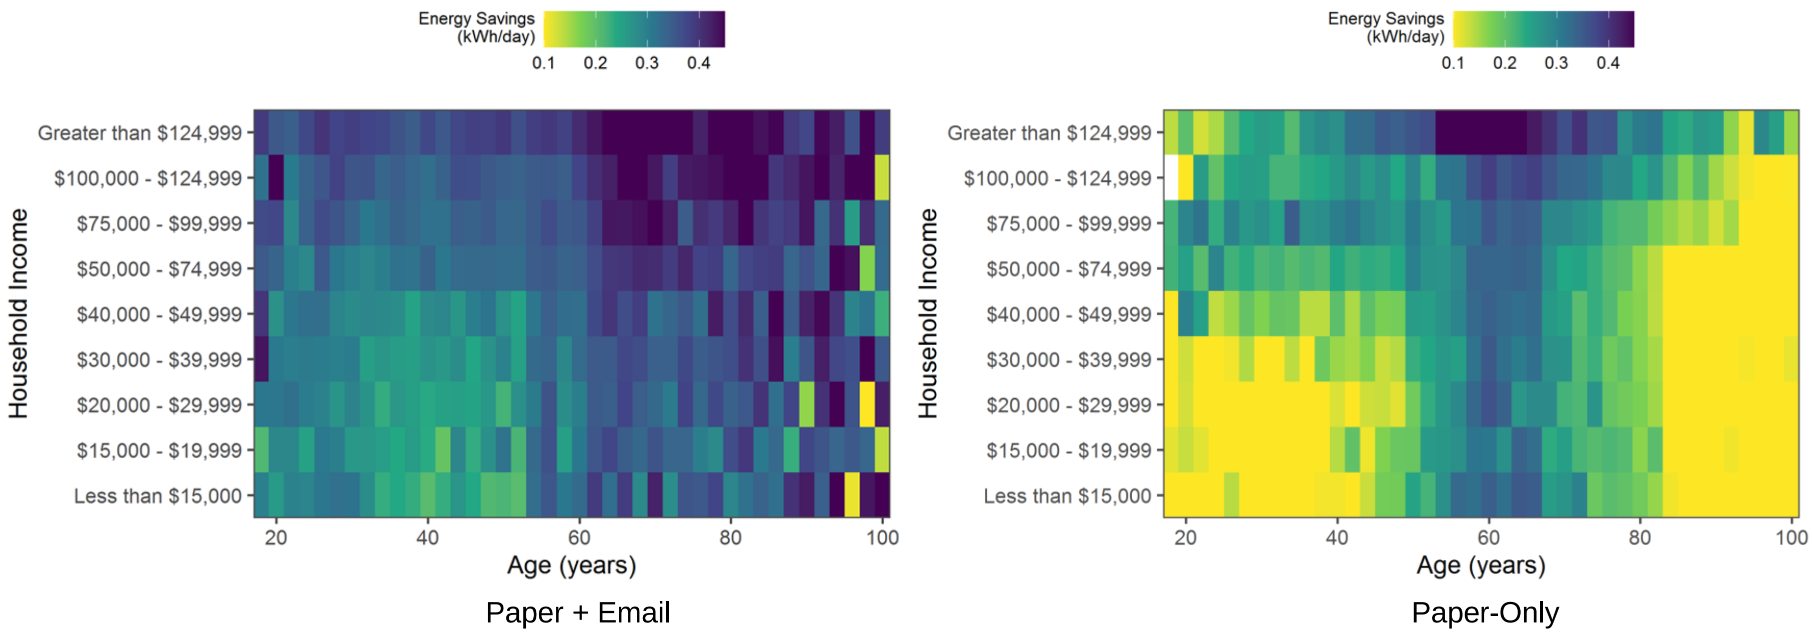

As Figure 6-1 shows, there are two peaks in the distribution of predicted savings. The distinction between these two peaks occurs at about 0.25 kWh per day. This suggests a natural split of the households into two categories: high savers who are predicted to save more than 0.25 kWh per day, and low savers who are predicted to save less. To examine potential differentiating factors, averages for all household characteristics were calculated for high and low savers separately and then compared. Figure 6-10 illustrates the differences for select characteristics. The differences are shown in terms of standard deviations to eliminate differences in magnitudes across attributes. In the figure, positive values indicate that the mean value among low savers was greater than that of high savers, and negative values indicate that high savers had the larger mean for that particular attribute. This analysis was performed for the Paper + Email and Paper-Only groups separately as well, but there were not any economically meaningful differences.

By far the largest difference occurred in the baseline usage. Households in the high saver category had significantly higher baseline usage on average than those in the low saver category. This aligns with the strong correlation depicted in Figure 6-5. Beyond this, high saver households are also more likely to receive emails, are slightly more likely to own their home, and tend to have higher incomes. Low saver households, on the other hand, are more frequently CARE customers.

Figure 6-10: Mean Comparisons between Low and High Savers

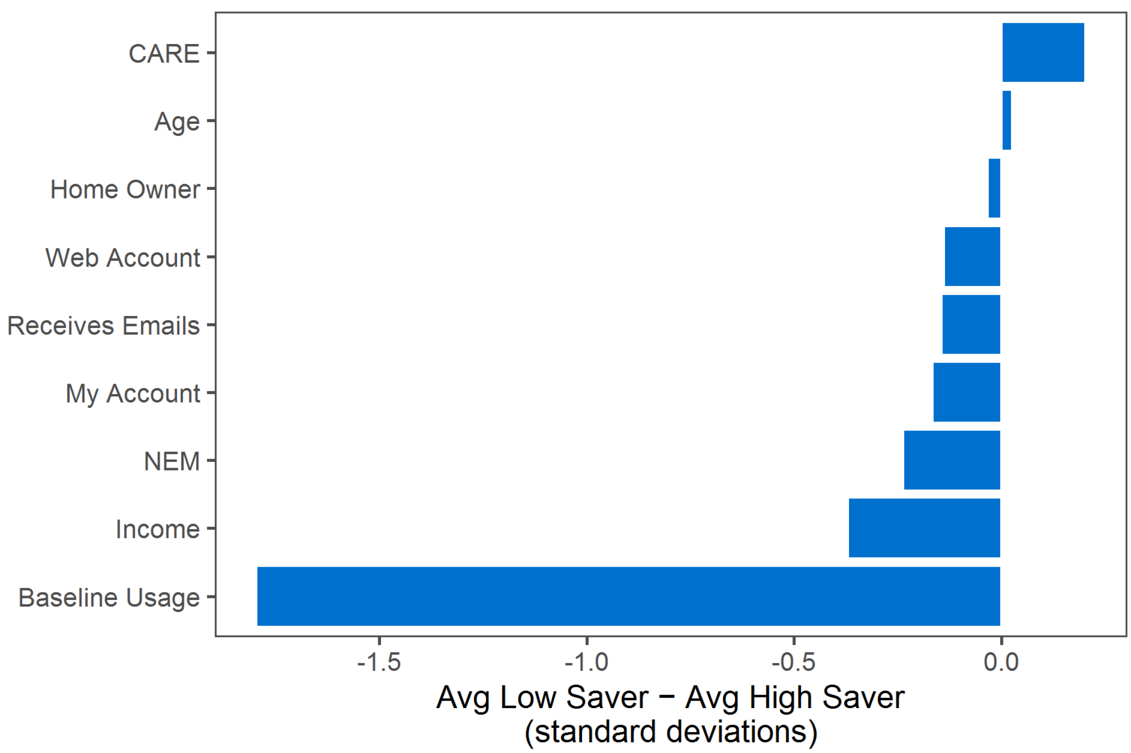

Figure 6-11 shows the predicted energy savings from the causal forest analysis attributable to the HERs by zip code. The blue colored zip codes indicate savings, with the darker blue indicating a higher level of savings. From this map, it appears that only two zip codes do not show savings, noted in yellow. This color yellow is very close to zero, meaning the HER treatment was largely not effective in these areas.

Figure 6-11: Average Daily Savings (kWh) per Customer by Zip Code

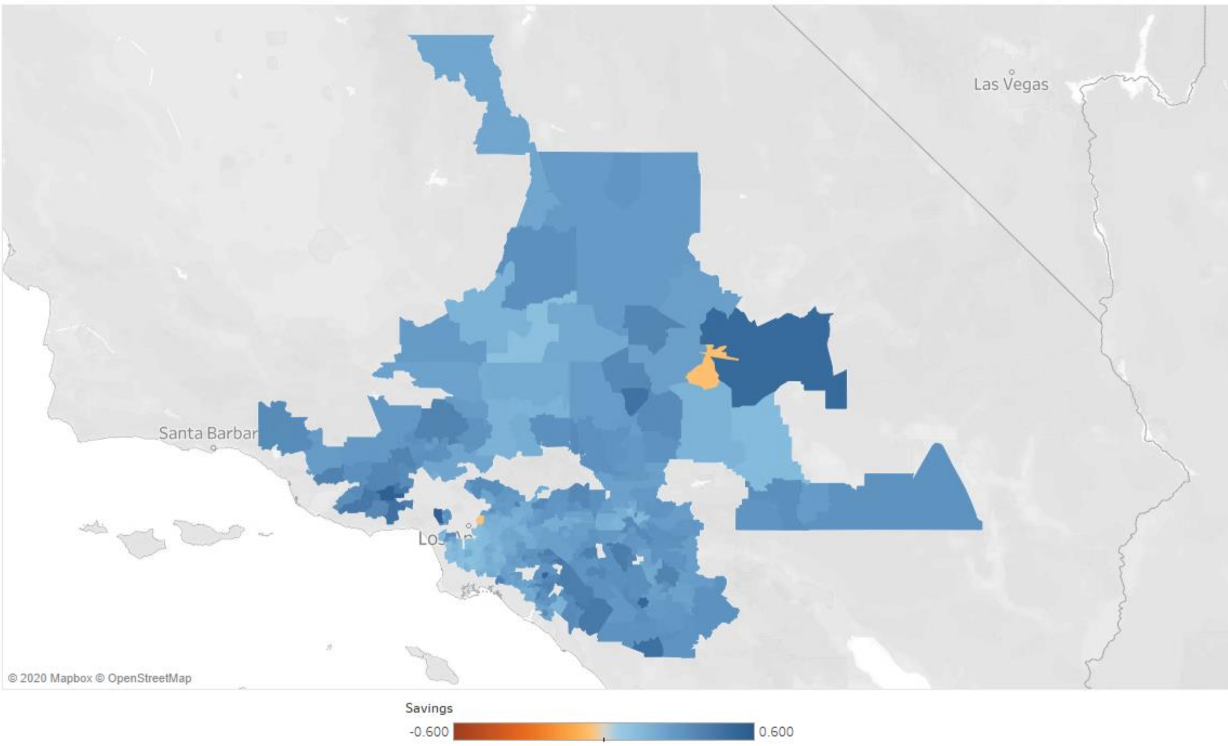

Predicted energy savings among discontinued customers for the pre-termination and post-termination periods separately are illustrated in Figure 6-12. In the pre-termination period, the bimodal nature of the savings is once again present, indicating that the general pattern of a “high-saving” group and a “low-saving” group still holds when looking at the discontinued group separately. In the post-termination period, however, the distinction between the two is largely lost. The range of pre-termination (year 1) effects is wider than that of post-termination (year 2) effects. Intuitively, one would expect this to be the case; the savings caused by a single intervention, the HER delivery, would decrease in magnitude as time passes. Notice also that the adverse effect of increased consumption does not carry over from year 1 to year 2.

Figure 6-12: Distribution of Predicted Daily Energy Savings in Pre- and Post-Termination Periods

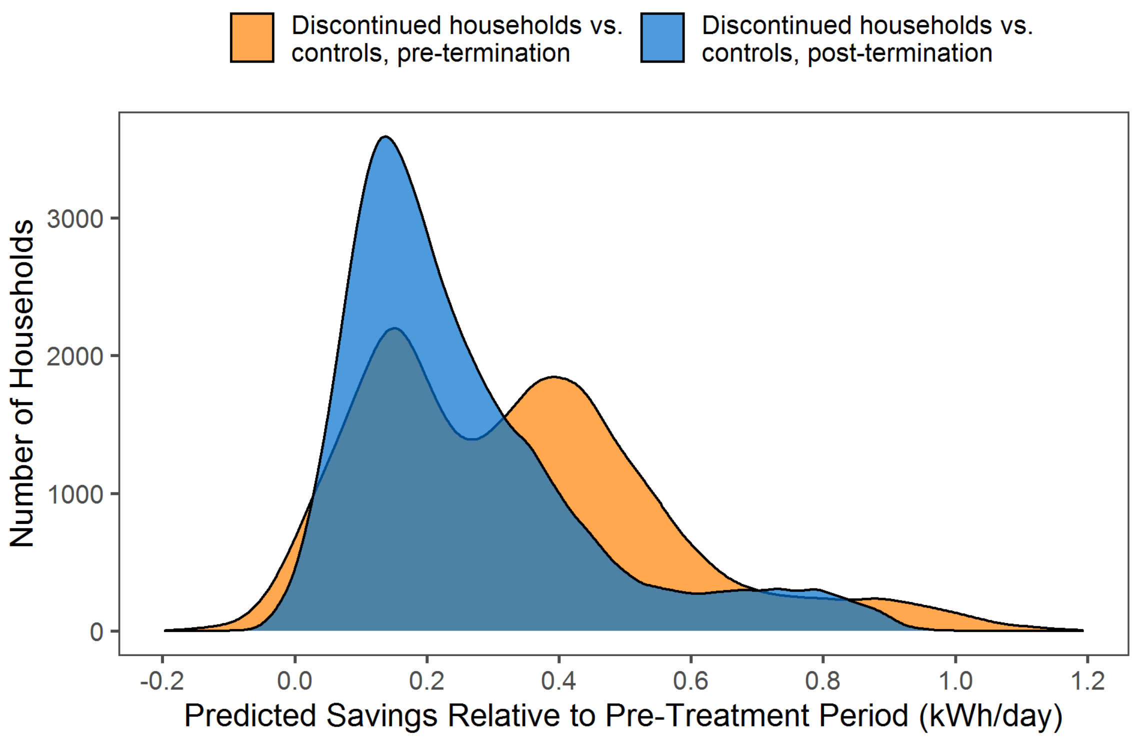

When examining the daily energy savings from the discontinued group during the pre-termination and post termination period the overall savings in the pre-termination period is 0.33 kWh. In the post-termination period, it is 0.26 kWh. According to this modeling approach, which examines year-over-year changes in savings, 79% of savings persisted during the first year of the persistence pilot. However, it should be noted this value does not take year-over-year differences for the continued customers into account. Accordingly, these observed differences could be attributable to weather or other factors.

To determine whether some of the savings reduction can be explained by the discontinuation, or other factors, the fourth model discussed previously is used. This final estimation involves dropping households that never received an HER (the control group) and estimating the discontinuation effect, which effectively controls for year-over-year differences with the continued group savings. This model found an overall effect on energy savings of 0.002 kWh, which is not statistically different from zero. This finding is consistent with the results presented in Section 4, which found no statistically significant difference in savings between the continued and discontinued groups during the post-termination period. The distribution of individual effects in Figure 6-13 shows no significant heterogeneity in persistence at a customer level. This leads to the conclusion that the HER discontinuation did not significantly impact savings in general.

Figure 6-13: Changes in Savings after HER Discontinuation

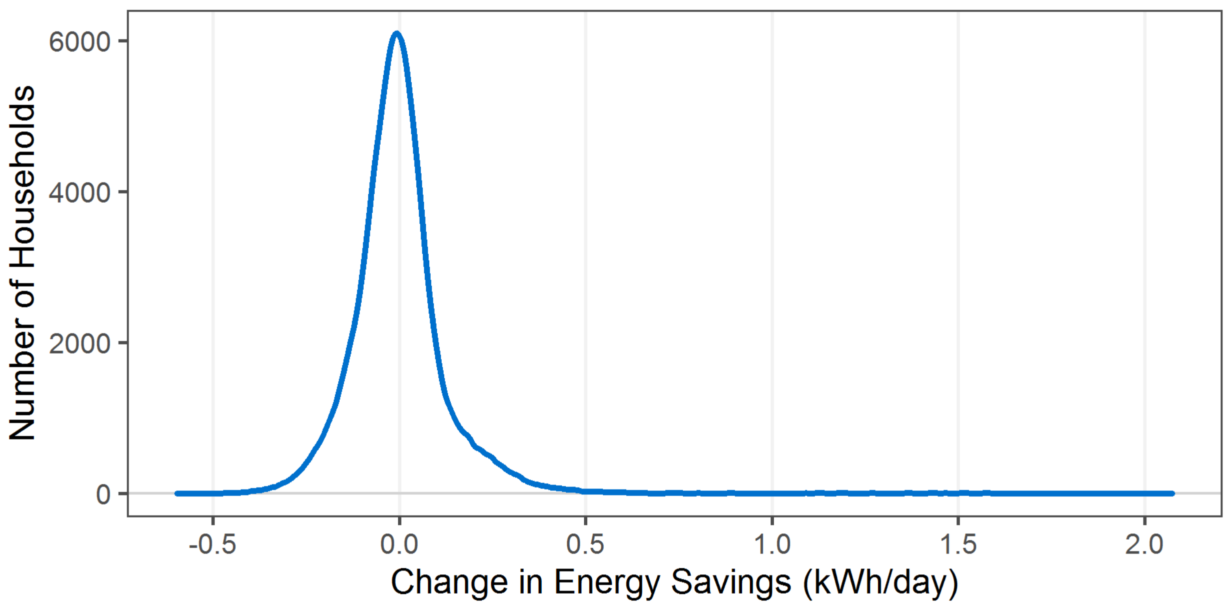

Figure 6-14 presents the persistence as a percentage of the pre-termination average daily savings by zip code. The levels of persistence are color-coded by quintile, with the green indicating higher levels of persistence and the yellow indicating lower levels of persistence. However, it should be noted that when persistence is examined at the zip code level, there is very little difference between the top and bottom quintiles. For example, the top quintile is persistence of greater than 100.3%, and the bottom quintile is persistence of less than 99.8%. There does not appear to be any pattern in persistence geographically, and with the quartile range being so narrow it does not appear that location is a meaningful driver of persistence.

Figure 6-14: Persistence of Average Daily Savings by Zip Code

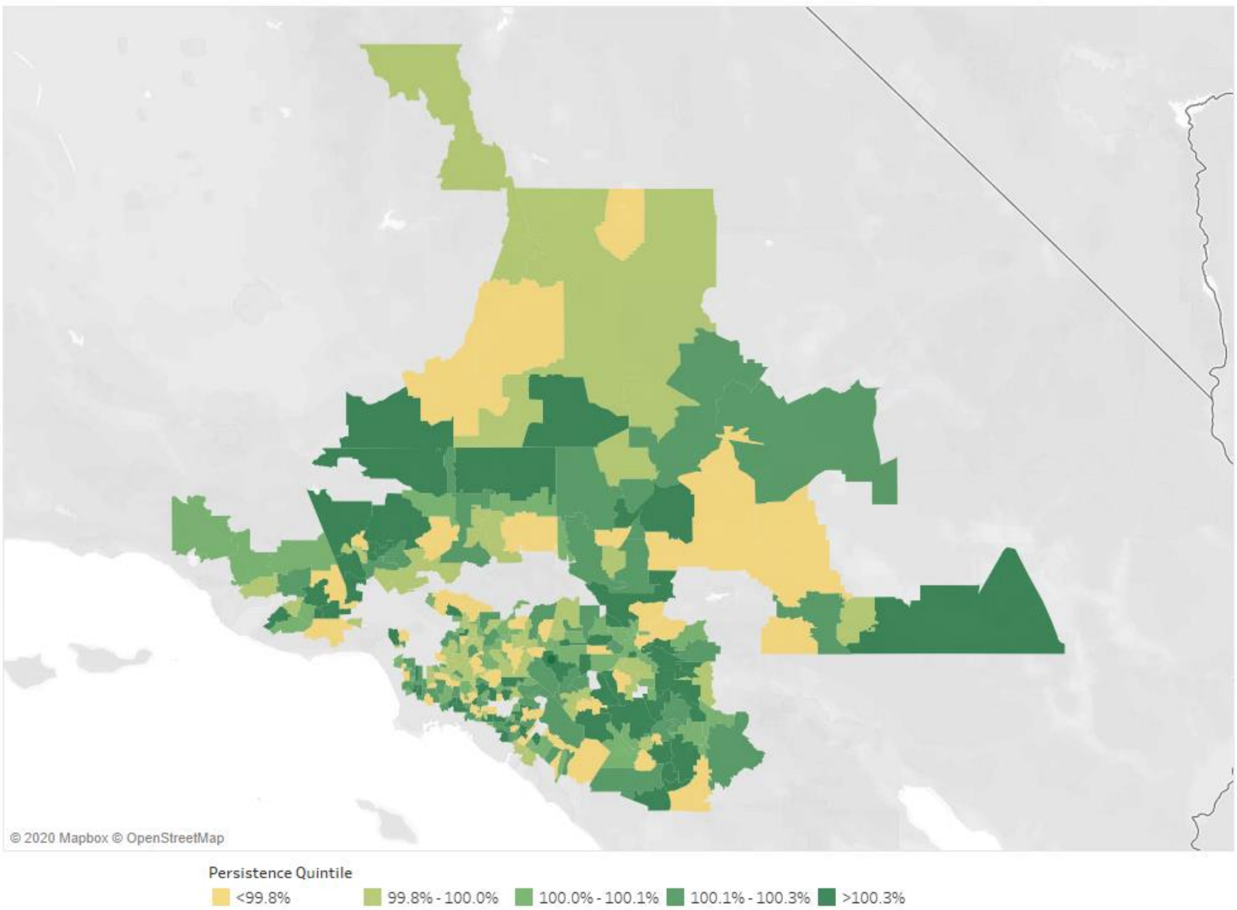

7 Comparison to Other Studies

The prospect of persistent energy savings is of interest to the utilities who include HERs in their energy efficiency portfolio. If energy savings persist at a high level, there is potential for significant cost-savings while maintaining valuable energy savings among residential customers. Several utilities across the county have examined the persistence of energy savings attributable to HERs. A selection of those studies and their estimated level of persistence for one year after report discontinuation is presented in Table 7-1.[7] The level of persistence found in SCE's Persistence Pilot, included in the last three rows of the table, falls within the range of results from previous studies.

Table 7-1: Previous HER Persistence Studies[8]

Utility	Experiment Name	Treatment Frequency	Approx. Years of Treatment prior to Discontinuation	Persistence[9] (1 Year)
ComEd	Wave 1	Bi-Monthly	4.5	96%
	Wave 3	Bi-Monthly	2.5	98%
	Wave 5	Bi-Monthly	1.5	78%
Eversource	Monthly Group	Monthly	1	125%
	Quarterly Group	Quarterly	1	70%
	Persistence Group	Monthly Abbreviated	< 1	74%
PG&E	Gamma Standard	Bi-Monthly	2.5	82%
	Gamma Reduced Freq.	Quarterly	2.5	120%
Puget Sound Energy	Legacy	Monthly & Quarterly	2	79%
SCE	**Wave 3 (All)**	**Bi-Monthly**	**3**	**98%**
	Wave 3 (Paper-Only)	**Bi-Monthly**	**3**	**90%**
	Wave 3 (Paper + Email)	**Bi-Monthly**	**3**	**103%**

[7] Sources are cited in Appendix A

[8] Some persistence studies are designed as rigorous experiments, while others measure year over year changes in energy savings without controlling for exogenous factors like weather or the economic climate. Only studies with an RCT design are included in the table.

[9] Relative to the continued group.

Figure 7-1 provides a visual representation of the SCE Persistence Pilot in green to the other studies in blue. While there is a range of persistence across the other studies, the SCE Persistence Pilot findings appear to be well within the observed range. Importantly, all of the studies are showing a notable level of persistence, with the lowest at approximately 70%. This provides strong evidence that savings persistence is real, and that there may be opportunities to leverage this finding to improve the cost effectiveness of the program through more strategic treatment strategies over time that may include pausing treatment for customers. By pausing treatment on customers, the funds that would have otherwise been spent to treat those customers can either be saved, reducing costs, or repurposed to treat other customers who would not have otherwise been treated due to budget limitations.

Figure 7-1: Comparison between SCE HER Persistence Pilot and Other Studies

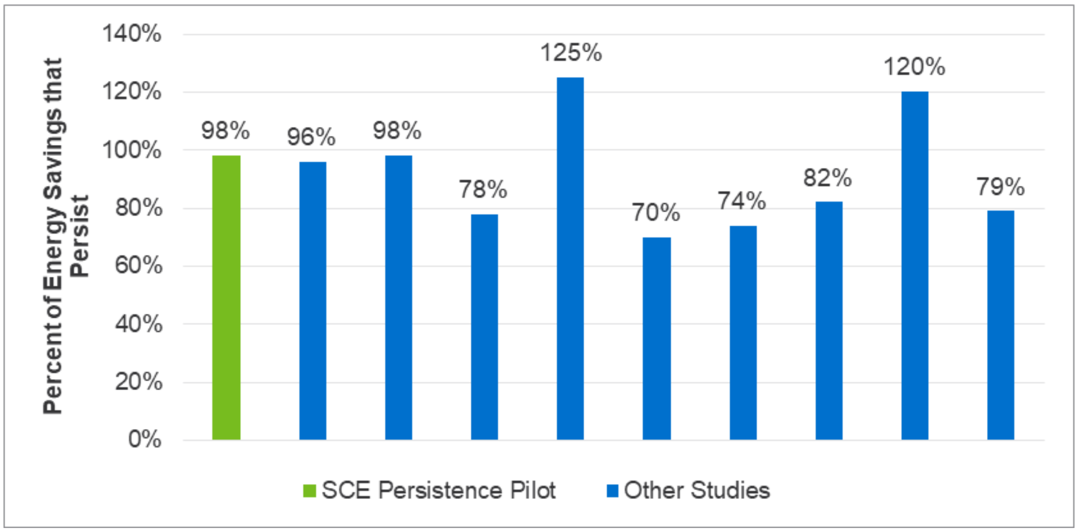

Appendix A   Studies Referenced in Section 7

Studies included in the table in Section 7 are cited below. Each of these persistence studies were designed as RCTs in which HER treatment customers were randomly assigned to continued or discontinued groups. The evaluators' chosen methodologies are in line with generally accepted evaluation practices and they are reliable comparisons to the SCE study.

Navigant. "ComEd Home Energy Report Program Decay Rate and Persistence Study – Year Two." July 20, 2016.

NMR Group. "Eversource Behavior Program Persistence Evaluation." April 9, 2017.

Nexant. "PG&E HER 2018 Energy and Demand Savings Early EM&V." March 10, 2020.

DNV KEMA. "Puget Sound Energy's Home Energy Reports Program Three Year Impact, Behavioral and Process Evaluation." April 20, 2012.

Additional studies that were not included the comparison in Section 7 are listed below. References to these studies are included in this appendix because they may appear in a search when researching HER persistence. However, these studies do not meet the level of rigor that warrants inclusion in this report; they were not designed as RCTs. Pre-post analyses that do not include a baseline (continued) group, like those in the studies below, cannot control for exogenous factors such as weather or economic changes and cannot be reliably compared to the SCE Persistence Pilot. However, they may provide valuable insights, nonetheless. These studies include the following:

Opinion Dynamics. "Massachusetts Cross Cutting Evaluation Home Energy Report Savings Decay Analysis." September 9, 2014.

Nexant. "Residential Behavioral Program Persistence Study." December 15, 2015.

DNV GL "Review and Validation of 2015 Southern California Edison Home Energy Reports Program Impacts (Final Report)." May 5, 2017.

Integral Analytics. "Impact & Persistence Evaluation Report Sacramento Municipal Utility District Home Energy Report Program: Program Years 2008-2011." November 2012.

Appendix B Regression Model Outputs

Table B-1 summarizes the regression model outputs for the baseline energy savings models. This model compares average daily kWh among continued treatment customers (***treatment*** = 1) to control customers (***treatment***=0) during the post-termination period. Two columns are included for each explanatory variable (***treatment***, ***pretreatment daily kWh***, and a constant): the coefficient on that variable and the standard error. Statistically significant coefficients are highlighted with asterisks. Each row in the table represents a separate model estimated for each customer segment included in this evaluation report. The regression specification is below:

$$kWh_{it} = a + b_t + c_t \cdot treatment_i + d_t \cdot pretreatment\_kwh_{it} + \varepsilon_{it}$$

Several boxes are included in the table to assist the reader in interpreting the regression outputs. First, the orange box highlights the coefficient on the ***treatment*** variable. This corresponds to the term c_t in the equation above and can be interpreted as the HER treatment effect on average daily kWh consumption. The value -0.328 indicates that treatment customers used 0.328 kWh less energy than control customers, per day, during the post-termination period. This is equivalent to 119.7 kWh in energy savings across the twelve-month post-termination period (as shown in Figure 4-2). There are three asterisks included in this cell, signaling that the estimate is statistically significant at the 99% confidence level. The value in the green box (0.0514 kWh) is the standard error on this point estimate. Finally, the values in the navy and purple boxes, 0.692 and 4.558, correspond to the coefficient d_t and the constant b_t in the regression specification above, respectively. Both of these estimates are also statistically significant at the 99% confidence level. The number of observations included in the estimate and the R^2 value are included in the last two columns.

Table B-1: Regression Model Output for Continued Treatment vs. Control Customers (Post-Termination Period 1)

Segment	Treatment		Pretreatment Daily kWh		Constant		Obs.	R-squared
	Coeff.	SE	Coeff.	SE	Coeff.	SE		
All Customers	-0.328***	(0.0514)	0.692***	(0.00321)	4.558***	(0.0741)	1,230,015	0.421
Paper-Only	-0.287***	(0.0689)	0.717***	(0.00471)	4.116***	(0.108)	508,116	0.483
Paper + Email	-0.357***	(0.0728)	0.676***	(0.00428)	4.814***	(0.101)	721,899	0.387
Typical Usage	-0.300***	(0.0895)	0.719***	(0.00482)	3.871***	(0.108)	345,838	0.432
Early Peaking	-0.341***	(0.0955)	0.698***	(0.00459)	4.106***	(0.117)	391,732	0.465
Late Peaking	-0.341***	(0.0822)	0.671***	(0.00665)	5.292***	(0.147)	492,433	0.366
Typical Seasonal Usage	-0.218**	(0.1000)	0.695***	(0.00453)	4.212***	(0.107)	314,463	0.475
Constant Usage Across Seasons	-0.400***	(0.0801)	0.698***	(0.00485)	4.399***	(0.118)	518,540	0.434
High Summer Usage	-0.325***	(0.0902)	0.679***	(0.00807)	5.072***	(0.175)	396,819	0.343
Under 50	-0.138	(0.0864)	0.696***	(0.00512)	4.846***	(0.120)	466,336	0.406
50 or Older	-0.447***	(0.0637)	0.689***	(0.00411)	4.380***	(0.0941)	763,679	0.432
Under $75k/yr	-0.250***	(0.0698)	0.681***	(0.00484)	4.744***	(0.104)	586,746	0.378
Over $75k/yr	-0.399***	(0.0749)	0.698***	(0.00427)	4.429***	(0.105)	643,269	0.443

Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

Table B-2 summarizes the regression model outputs for the incremental savings models used to estimate the persistence. The value in the orange box can be interpreted as the difference in daily kWh consumption between the continued and discontinued groups (or c_t in the regression specification below). A positive value indicates that discontinued customers used more energy than continued customers during the post-termination period. In the example highlighted below discontinued customers used 0.00511 kWh more per day than continued customers. Interpreted another way, discontinued customers saved 0.00511 kWh less than continued customers. This estimate was not statistically significant (there are no asterisks in the cell), indicating there is not a statistically significant difference between the continued and discontinued groups and the persistence level is very high.

$$kWh_{it} = a + b_t + c_t \cdot discontinued_i + d_t \cdot pretermination\_kwh_{it} + \varepsilon_{it}$$

Table B-2: Regression Model Output for Discontinued Customers vs. Continued Customers (Post-termination Period 1)

Segment	Discontinued		Pre-termination Daily kWh		Constant		Obs.	R-squared
	Coeff.	SE	Coeff.	SE	Coeff.	SE		
All Customers	0.00511	(0.0264)	0.756***	(0.00283)	3.849***	(0.0613)	1,530,549	0.685
Paper-Only	0.0299	(0.0377)	0.761***	(0.00332)	3.724***	(0.0710)	632,779	0.695
Paper + Email	-0.0124	(0.0364)	0.753***	(0.00402)	3.924***	(0.0881)	897,770	0.679
Typical Usage	-0.0108	(0.0445)	0.752***	(0.00524)	3.621***	(0.107)	430,708	0.703
Early Peaking	0.0199	(0.0475)	0.767***	(0.00275)	3.617***	(0.0640)	490,786	0.720
Late Peaking	0.00221	(0.0442)	0.746***	(0.00629)	4.249***	(0.133)	607,732	0.634
Typical Seasonal Usage	-0.0409	(0.0521)	0.740***	(0.00364)	3.648***	(0.0822)	389,989	0.713
Constant Usage Across Seasons	0.00446	(0.0404)	0.755***	(0.00309)	3.883***	(0.0700)	645,588	0.698
High Summer Usage	0.0368	(0.0460)	0.782***	(0.00905)	3.717***	(0.184)	493,117	0.638
Under 50	-0.0357	(0.0450)	0.754***	(0.00387)	4.047***	(0.0870)	578,419	0.676
50 or Older	0.0311	(0.0325)	0.757***	(0.00394)	3.735***	(0.0835)	952,130	0.690
Under $75k/yr	0.0216	(0.0371)	0.744***	(0.00448)	3.905***	(0.0930)	732,731	0.654
Over $75k/yr	-0.00894	(0.0375)	0.762***	(0.00365)	3.873***	(0.0821)	797,818	0.703

Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

Table B-3: Regression Model Output for Continued Treatment vs. Control Customers (Post-Termination Period 2)

Segment	Treatment		Pretreatment Daily kWh		Constant		Obs.	R-squared
	Coeff.	SE	Coeff.	SE	Coeff.	SE		
All Customers	-0.323***	-0.1	0.648***	0.0	5.574***	-0.1	782341	0.282
Paper-Only	-0.225***	-0.1	0.690***	0.0	4.816***	-0.1	323198	0.356
Paper + Email	-0.391***	-0.1	0.621***	0.0	6.093***	-0.1	459143	0.241
Typical Usage	-0.255***	-0.1	0.641***	0.0	5.668***	-0.2	220552	0.232
Early Peaking	-0.285***	-0.1	0.669***	0.0	5.005***	-0.2	247726	0.321
Late Peaking	-0.404***	-0.1	0.634***	0.0	6.011***	-0.2	314055	0.269
Typical Seasonal Usage	-0.308***	-0.1	0.615***	0.0	6.204***	-0.2	198274	0.188
Constant Usage Across Seasons	-0.375***	-0.1	0.659***	0.0	5.397***	-0.1	330028	0.293
High Summer Usage	-0.269***	-0.1	0.654***	0.0	5.355***	-0.2	253896	0.306
Under 50	-0.161*	-0.1	0.626***	0.0	6.468***	-0.2	294989	0.242
50 or Older	-0.421***	-0.1	0.662***	0.0	5.013***	-0.1	487352	0.309
Under $75k/yr	-0.239***	-0.1	0.600***	0.0	6.212***	-0.2	371675	0.212
Over $75k/yr	-0.395***	-0.1	0.668***	0.0	5.364***	-0.1	410666	0.314

Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

Table B-4: Regression Model Output for Discontinued Customers vs. Continued Customers (Post-termination Period 2)

Segment	Discontinued		Pre-termination Daily kWh		Constant		Obs.	R-squared
	Coeff.	SE	Coeff.	SE	Coeff.	SE		
All Customers	0.0582	0.0	0.776***	0.0	4.296***	-0.1	974133	0.506
Paper-Only	0.0912*	-0.1	0.805***	0.0	3.676***	-0.1	403275	0.539
Paper + Email	0.0356	0.0	0.759***	0.0	4.668***	-0.2	570858	0.488
Typical Usage	-0.0683	-0.1	0.774***	0.0	4.207***	-0.2	275070	0.491
Early Peaking	0.0423	-0.1	0.806***	0.0	3.756***	-0.1	310354	0.555
Late Peaking	0.160***	-0.1	0.748***	0.0	4.845***	-0.3	387876	0.469
Typical Seasonal Usage	0.0416	-0.1	0.723***	0.0	4.773***	-0.2	245729	0.428
Constant Usage Across Seasons	0.0545	-0.1	0.773***	0.0	4.371***	-0.1	411396	0.521
High Summer Usage	0.0744	-0.1	0.804***	0.0	4.012***	-0.3	315806	0.525
Under 50	0.036	-0.1	0.765***	0.0	4.828***	-0.1	366250	0.463
50 or Older	0.0739*	0.0	0.782***	0.0	3.983***	-0.2	607883	0.535
Under $75k/yr	0.118**	-0.1	0.750***	0.0	4.474***	-0.2	464518	0.443
Over $75k/yr	0.00652	0.0	0.788***	0.0	4.287***	-0.2	509615	0.544

Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

Headquarters
45 Stevenson Street, Suite 700
San Francisco CA 94105-3651
Tel: (415) 369-1000
Fax: (415) 369-9700
www.nexant.com