Energy Efficiency Program Targeting:
Using AMI data analysis to improve at-the-meter savings for small and medium businesses

Final: Research Report

August 10, 2018

CALMAC Study ID PGE0421.02

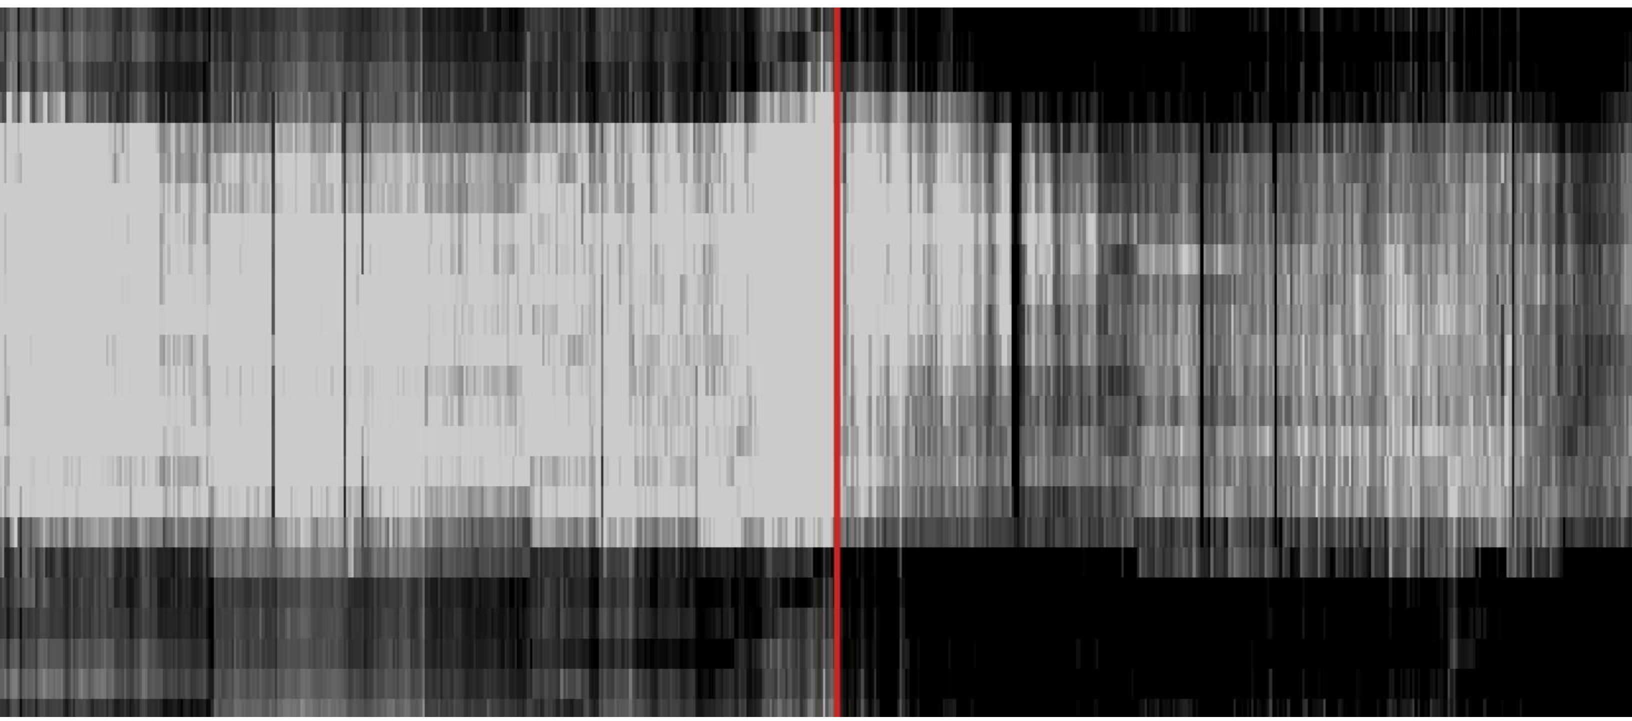

Convergence Data Analytics

Prepared by:
Sam Borgeson, Partner
Brian Gerke, Consultant
Convergence Data Analytics, LLC
For Pacific Gas and Electric Corp.
August 10, 2018

Contents

Cover image: heat map of pre and post intervention consumption for an anonymous DI customer, with hours of day along the y-axis and days of the year along the x-axis.

1 Introduction

This research paper investigates the potential for predictive targeting methods to improve "at-the-meter savings" outcomes for small and medium business (SMB) energy efficiency (EE) programs.

At-the-meter savings are computed based on changes in customers' metered consumption[1] before and after program-related interventions. The goal of targeting in this context is increased savings at lower cost through the identification of customers most likely to realize deeper savings, on average, than past participants and targeting those customers with encouragement to participate in future iterations of the program.[2]

This research is motivated by several questions: What magnitude and depth of savings "gains" over current programs are achievable through targeting? What metrics and methods best support the comparisons necessary to evaluate the performance of different strategies? How are these results determined by the methods of estimating at-the-meter savings? How much variation in optimal targeting strategies can be expected across different types of customers, programs, and measure types?

For this work, two longstanding PG&E programs were used to evaluate data-driven targeting schemes. The first is the Regional Direct Install Program (DI), which focuses primarily on lighting and refrigeration upgrades, and the second is the Commercial HVAC Quality Maintenance Program (HVAC), which services air conditioning equipment based on industry standard maintenance protocols. For both programs, we have identified customer characteristics and usage patterns, all available prior to the program start, that predict at-the-meter savings.

We observe that the best predictive targeting strategies depend on the nature of the interventions being made and the end-uses they impact. High total and peak usage are predictive of lighting savings, high total and baseload usage are predictive of refrigeration savings, and high temperature sensitivity and estimated disaggregated AC usage are predictive of air conditioning savings.

When applied as "targeting filters" that select sub-groups of customers based on threshold values, usage metrics related to baseload and total consumption can roughly **double average DI program savings of the remaining customers when targeting 1 out of every 2 customers and roughly triple average DI savings when targeting 1 out of 4 customers**. For the HVAC program, the results are even more dramatic. The program is fairly light touch and the average savings at the meter for the program were nearly indistinguishable from the background noise of other changes in consumption over time. Filters based on usage characteristics that estimate AC loads were able to **elevate average HVAC savings from 1 kWh/day to 13 kWh/day when targeting 1 out of every 2 customers and to 28 kWh/day when targeting 1 out of every 4 customers.** Those performance gains elevate savings well beyond the background noise of other changes.

Although this research applies targeting methods to two specific PG&E energy efficiency (EE) programs, the methods and results developed here are generalizable to a wide variety of EE programs and customer types. The insights gained through this work can be readily employed to guide future

[1] This research is based on whole premise revenue grade utility meters installed at customer premises.
[2] Energy efficiency programs are generally open to all utility customers. Targeting does not prevent customer participation but is intended to encourage those customers most likely to achieve the greatest savings.

interventions toward optimized savings results, both for participating customers and for the EE programs.

2 Background

In candid moments, many EE advocates and implementers can readily list interventions that have great potential to save, but nevertheless do not fit within existing program policies and evaluation rules. With the advent of new evaluation and implementation pathways, improved program targeting stands to help unlock some of that untapped potential. Whether it is through lowering acquisition costs by focusing recruitment on viable participants or improving per-customer outcomes by focusing program recruitment on those with the greatest expectation of savings, data-driven targeting and personalization have a lot of teach and offer EE administrators, policymakers, and implementers.

The passage of Assembly Bill 802[3] in California established existing conditions baselines[4] for many energy efficiency (EE) programs with the goal of allowing programs to incentivize and claim savings for the replacement of inefficient equipment with more efficient equipment. AB802 further specified that the savings should be estimated "…*taking into consideration the overall reduction in normalized metered energy consumption as a measure of energy savings."* The employment of existing conditions baselines directly aligns the savings attributable to a program with the change in a customer's metered energy usage. This has catalyzed interest in Normalized Metered Energy Consumption (NMEC) based programs.[5]

As a part of this interest, PG&E is running a pay for performance (P4P) pilot where implementers are paid proportionally to their NMEC savings on projects. NMEC-based evaluation will also be an option for implementers under California's expanded third party solicitation for EE programs. Under such rules, projects that deliver significant metered savings are valuable while projects whose impacts do not manifest at the meter are not. This is in contrast to the deemed savings model in which no direct reward exists for maximizing metered savings. The possibility (and perhaps business imperative) of optimizing metered savings in the P4P paradigm has major practical implications for program design. In particular, it motivates identification and targeting of customers with high savings potential.

3 Key concepts for at-the-meter programs

The central hypothesis evaluated by this work is that customer attributes can predict savings-at the meter, and furthermore, certain patterns in *electric consumption* are correlated with efficiency potential and savings in a manner that complements and extends the performance gains from

---

[3]AB 802 instructs the California Public Utilities Commission to authorize EE programs with savings measurement based on *"all estimated energy savings and energy usage reductions, taking into consideration the overall reduction in normalized metered energy consumption as a measure of energy savings"* https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=201520160AB802

[4]Before AB 802, the standard California baseline for EE measures was set to equal to current code requirements, or industry standard practice (ISP), in cases where no code exists. Before AB 802 EE programs could only incentivize the portion of savings above code. With existing conditions baselines, eligible programs can incentivize the full savings.

[5] At the time of this writing, the CPUC is taking comments on its proposed ruling on how NMEC-based evaluation will be handled under the expanded third-party solicitation process. The proposed ruling requires it to be subjected to the same level of review and oversight as custom projects.

targeting based on customer characteristics. Such potential has been widely recognized and discussed in the efficiency community but has rarely been directly tested and quantified.

3.1 At-the-meter vs. alternatives

To target customers for programs with savings determined through NMEC changes, it is first useful to review evaluation approaches currently in use and how they differ.

Programs that rely on **deemed savings** reward implementers for efficient customer recruitment and deployment of generally prescriptive measures at scale.

Programs that use **simulated outcomes** to compute savings are often utilized for large commercial, industrial, and agricultural customers with unique usage requirements and patterns. These types of custom programs rely on substantial engineering analysis that may make sense to facilitate large individual savings claims but are not practical for programs with high volume.

New approaches that have not (yet) been included in tables of deemed savings or simulation software, or readily incorporated into engineering analysis, such as behavioral measures, cannot be rewarded under traditional program rules.

Deemed and custom programs often have further philosophical complexities, including counterfactual projections and the determination of code or industry standard practice baselines, consideration of "early retirement" and corresponding dual-baseline periods, and a host of adjustment factors designed to tune savings such as realization rates, and expected in-service rate adjustments. In these cases, assigned savings are several steps removed from any change in metered energy use.

In contrast, programs and evaluations that determine at-the-meter savings will reward direct changes in observed energy usage. Program administrators and implementers should note that plugging in a deemed or custom program directly into an NMEC platform will not perform optimally. Instead, programs utilizing at-the-meter evaluation will realize deeper savings with new program designs and customer recruitment strategies tailored to at-the-meter performance.

3.2 At-the-meter program targeting

When program outcomes are assessed at the meter, evaluators typically observe a wide range of realized savings across participants. It is well understood that in many cases the magnitude and depth of EE project savings is influenced by customer characteristics like their business type, climate, efficiency of existing equipment, operating hours, levels of occupancy, etc. Performance minded program planners and administrators have long been aware that these characteristics can be used to better plan and target programs. However, most programs are evaluated with rules that assign fixed deemed savings for each participating customer without any possibility of reward for planners to be recognized for deeper achieved savings. As a result, the majority of program administrators and implementers are not well-versed in employing targeting schemes.

3.3 Program impacts and the challenge of attribution

Measured savings ideally accrue when the meter spins more slowly than it used to *due to program interventions*. Within billing analysis impact evaluations, this has typically been assessed by comparing *pre-program consumption* to *post-program consumption* after normalizing for known sources of load variability, like weather. With some exceptions, the pre-period is usually taken as the year just before the start of the program and the post-period as the year immediately following. Unfortunately, there are a number of reasons that both individual and collective customer load might naturally vary between those time periods. For example, companies might hire new workers, purchase new equipment, or redouble production to meet rising market demand; and the opposite might occur as well.

Along with normalization for quantifiable variables such as weather, accurate attribution of metered savings can benefit from a control (or comparison) group. Ideally the control group is designed to account for changes in metered consumption that would have occurred over time among program participants if they hadn't experienced the EE intervention. For this work, we generated matched control groups for each program studied with characteristics comparable to the respective participant groups. We ran the control groups through all the same methods used to measure program savings for the program participants. See the Methods section and Appendix B for a much more detailed explanation of the control group selection and comparison steps.

3.4 At-the-meter natural variability and program impacts

Figure 1a illustrates the pre/post "savings" for the DI program control sample. Here savings are defined as difference in average weather-normalized kWh/day between the pre- and post-periods. **By definition, the control customers are not program participants, so any variability or trends observed in their outcomes can only be explained as non-program effects, or what we call "natural variability."** While the average is close to zero, there are many customers with big swings in consumption (both positive and negative tails in the figure). The range of outcomes in the control group represents the variability against which the attribution of program savings must occur.

To build intuition for how program interventions overlay on top of observed natural variability, we plot two savings scenarios. Figure 1b illustrates hypothetical deemed savings impacts. Under deemed savings, we assume a fixed amount of savings per customer.[6] This is the equivalent of shifting the whole distribution (otherwise unchanged) toward positive savings on the right. In contrast, Figure 1c illustrates hypothetical "proportional savings" impacts. For proportional savings, we assume that the savings are proportional to some property of consumption, e.g. proportional to the pre-intervention average daily consumption. This has the effect of moving some customers further than others, introducing a positive bias and skew to the distribution (Figure 1c).

[6] We are aware that deemed savings can actually be differentiated by site and implementation details, but the basic concept remains that the point of deemed savings is to provide a fixed estimate for savings.

Figure 1: Potential savings outcomes relative to natural variability found in the control population. a: NMEC "savings" computed for the DI control group illustrating the magnitude of natural variability in the population. b: Savings distribution resulting from hypothetical treatment of non-participants with measures that produce fixed results assumed by deemed savings methods. c: Savings distribution resulting from hypothetical treatment of non-participants with measures that save proportionally to daily energy consumption.

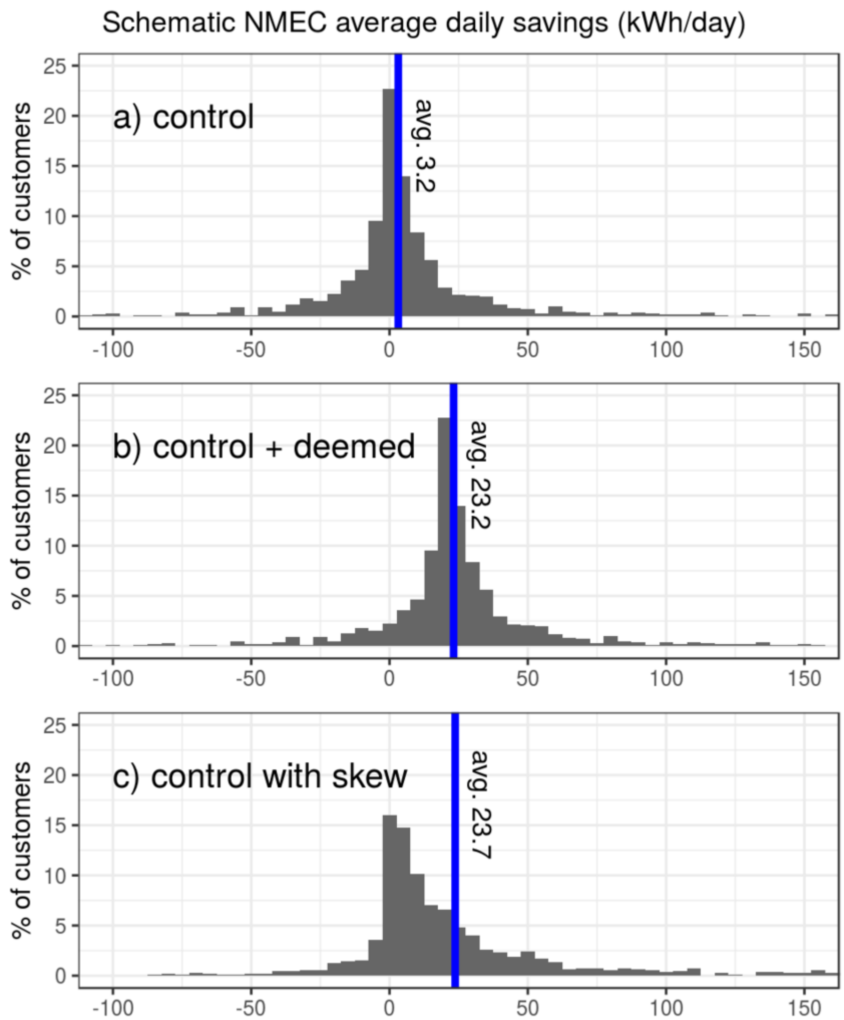

In practice, distributions of at-the-meter savings tend to look more like c than b because the same intervention strategy can easily have different impacts at different sites. The challenge for evaluators in assessment of metered-savings is to separate program-induced savings from natural variability.

4 Program characteristics

4.1 Regional Direct Install

Direct Install (DI) programs provide or assist with the direct installation of upgraded equipment in a participating customer's facility. The vast majority of these measures are lighting and refrigeration, with limited appliance, HVAC, and electronics measures. Each program offers a combination of deemed and custom measures. In contrast to rebate programs that provide financial incentives to participating customers *after* purchase and/or installation, DI programs generally apply their direct install incentives to the cost of installation (including labor and materials) *before* the customer incurs out-of-pocket costs. This is typically done through installation contractor incentives.

PG&E administers several Regional DI Programs, each of which is implemented by either an independent Third-Party Vendor or a Lead Local Partner (Implementers). The DI programs operate in distinct regions within PG&E's service territory and generally focus on serving hard-to-reach and underserved small and medium business (SMB) customers, as well as municipal facilities, special districts, and non-profits. These programs are designed to facilitate greater participation among underserved customers, who are less likely to participate in energy efficiency through other program channels. Program services may include direct customer outreach and marketing, energy assessments, installation assistance, quality control, financial incentives and rebate processing.

4.1.1 Customer Eligibility

To be eligible for Program participation, customers must meet the following criteria:[7]

- Must receive electricity and/or gas from PG&E
- Must pay Public Purpose Programs (PPP) charge on their PG&E utility bill
- Must be a non-residential PG&E customer
- Maximum billing demand (kW) at the customer Premise shall not have exceeded **200 kW** at any time within the past 12 months

It's worth noting that although Regional DI focuses on SMB customers, Large customers defined by usage (annual electric usage greater than 500,000 kWh or gas usage greater than 250,000 therms) may participate in the program if peak demand is under 200 kW. In addition, San Francisco County serves larger hospitality customers with peak demand up to 500 kW.

4.2 Commercial HVAC Quality Maintenance

The Commercial HVAC Quality Maintenance (HVAC) Program, now referred to as the Commercial HVAC Optimization Program, is based on the assumption that there are energy and demand savings achievable through the regular application of quality maintenance (QM) procedures applied to existing nonresidential HVAC equipment.

The Program offers a comprehensive, properly sequenced maintenance program based on the ASHRAE / ACCA Standard 180 (an industry standard for quality maintenance) that helps improve energy efficiency and reliability, equipment life, thermal comfort and indoor air quality. The program is driven by Service Agreements between customers and contractors and incorporates training, marketing and incentives to help contractors understand and communicate the value of HVAC quality maintenance and energy efficiency.

[7] There are exceptions to these rules, with the following groups NOT eligible for the SMP programs studied: lodging customers with greater than 100 kW maximum demand as well as customers that are part of an ownership structure with 10 or more locations within PG&E service territory; tribal casinos; grocery stores with multiplex refrigeration systems or have more than 10 stores total located within PG&E's service area; colleges and universities; wastewater treatment plants; medical facilities; dairy facilities; wineries, and common areas of mobile home parks.

Customers receive incentive payments at the time the unit is brought to ACCA 180 Standard baseline by the Contractor and over the course of the Program's service agreement. The Program also pays a portion of the repair and upgrade costs directly to the enrolled HVAC contractor. Additional incentives are available for optional retrofit add-ons, such as advanced controls and sensors, variable speed drives and high efficiency motors.

4.2.1 Customer Eligibility

The Program is designed for commercial rooftop units powered by electricity from PG&E. Units that may not be eligible for incentives include: units less than three tons, new units under warranty, units that participated in PG&E's AirCare Plus™ program within the past five years, units that participated in the Commercial Quality Maintenance program in the past eight years, and units that are in a serious state of disrepair.

5 Data

5.1 Meter, account, intervention, and weather data

A specific utility account is tied to a location through a service agreement. For our work, we define a "premise" as a unique account/location combination, which can span multiple sequential service agreements. This work started with assembling data from all premises that received interventions associated with either the DI or HVAC program for the 2014 and 2015 program years, with a max demand during the year prior to the program year of <200 kW. Although these two programs generally focus on SMB customers, there are larger customers that also participate. Distinct from the qualifying characteristics for SMB program eligibility, **customer size in this report is defined by annual electric and gas usage** as described below:

- **S**: electric usage less than 40,000 kWh and gas usage less than 10,000 therms
- **M**: electric usage 40,000-500,000 kWh and gas usage 10,000-250,000 therms
- **L**: electric usage greater than or equal to 500,000 kWh and gas usage greater than or equal to 250,000 therms

We pulled four types of data for this list of participating SMB customers:

Account data: targeting-relevant account and site characteristics (some of which, like rate-schedule, reflected values as of their program participation) including zip code, rate schedule, NAICS code, a pre-computed customer size code (corresponding to the above size definitions, where 'N' is assigned to customers with not enough data to compute an annual value), and a net metering indicator for sites with rooftop PV or other on-site generation.

Intervention data: All known program interventions (i.e. not just DI or HVAC) for the period spanning 2013-2016, with PG&E tracked characteristics including program type, install date, technology family (broad categories like lighting, HVAC, and refrigeration), technology (specific technologies within each family, like LEDs and Linear fluorescents), and various categorical assignments for each intervention related to program focus, geography, and method of delivery.

Meter data: 24 hourly electricity consumption readings of smart meter data per day for the period spanning 2013-2016 (with some starting later and/or ending earlier as customers move in and out of their SMB locations, change rates, or otherwise alter their Service Agreements).

Weather data: Using an open source Python tool[8], we downloaded hourly Quality Controlled Local Climatological Data (QCLCD) spanning 2013 through 2016 for every zip code for which we had customer data. The relevant data for the purposes of this project was hourly outside temperature to be used in weather normalization.

We also pulled a random sample of 20,000 non-participating SMB customers who did not participate in either the DI or HVAC program between 2014 and 2015. For this sample, we pulled account data, meter data, and weather data, as described for the participants. We also pulled any program intervention activity (by definition not in the DI or HVAC programs) that happened to be associated with these customers for the period 2013 through 2016. This sample of non-participants served as the pool from which our matched control samples were drawn for each program.

5.2 Premises and intervention counts

A series of data quality checks were applied to the raw customer data, ultimately producing at the final set of quality checked (QC'd) premises used in the analysis. Table 1 below provides the counts of the remaining premises passing each sequential data quality check, with the final numbers used for the savings and targeting analysis in bold. The "mixed" program category simply refers to premises that participated in both the DI and HVAC programs. They are absent from the first row because that row contains a tally of total participant count for each program prior to establishing which customers participated in both. The number of "mixed" customers is small and the independent program effects are difficult to tease apart from one another, so the "mixed" participation customers were excluded from the analysis.

We required 120 days of meter data prior to the program intervention(s) (pre) and 120 days of meter data after (post). To avoid changes too large to be caused by the program interventions between pre- and post-program periods, we dropped the minority of outlying customers whose energy consumption rose above 200% or dropped to less than 50% of pre-period consumption.

The control sample was also subjected to the same data quality requirements, and the number of available control customers (the same pool for both programs) is also documented in the table.

Table 1: Program participating premises passing each cumulative step in our data validation requirements

criteria	DI	HVAC	mixed
all participants	13428	3578	N/A
and occupied without missing data	10442	2004	118
and no PV	10363	1970	117
and 120 days of pre and post	7767	1304	66
and remove factor of 2 changes	**7497**	**1193**	**60**
controls passing the same criteria	6168	6168	6168

[8] https://github.com/sborgeson/local-weather

5.2.1 Premises with interventions completed per year

Table 2 provides the count of cleaned and validated premises completing program interventions in each year for both the DI and HVAC program.

Table 2: Number of data-validated customers completing interventions for each program and each program year

year completed	premise count (DI)	premise count (HVAC)
2013	1026	224
2014	2514	310
2015	2598	275
2016	1359	384

5.2.2 Program interventions by technology family

Table 3 and Table 4 below provide counts of interventions for cleaned and validated premises by the technology family. The DI program is dominated by lighting and refrigeration installations. These categories are significant enough to merit independent analysis. In the second table, we see that the HVAC program interventions focus mostly on unitary air conditioning and heat pump units, typically in the context of quality and maintenance measures.

Table 3: Count of premise-level DI program interventions grouped by technology family

technology family	intervention count
LIGHTING	5331
LIGHTING and REFRIGERATION	1005
REFRIGERATION	772
APPLIANCES	91
APPLIANCES and LIGHTING	85
ELECTRONICS AND IT and LIGHTING	50
NA and REFRIGERATION	45
APPLIANCES and REFRIGERATION	33
HVAC	27
APPLIANCES and LIGHTING and REFRIGERATION	23
ELECTRONICS AND IT and LIGHTING and REFRIGERATION	9
ELECTRONICS AND IT	7

Table 4: Count of premise-level HVAC program interventions grouped by technology (within the HVAC technology family)

technology	intervention count
UNITARY AC/HP	721
HVAC CONTROL and QUALITY MAINTENANCE and UNITARY AC/HP	127
HVAC CONTROL and QUALITY MAINTENANCE	96
QUALITY MAINTENANCE	90
CHILLER	45
AIR DISTRIBUTION and MOTORS PUMPS AND FANS and QUALITY MAINTENANCE and UNITARY AC/HP	29
AIR DISTRIBUTION and MOTORS PUMPS AND FANS and QUALITY MAINTENANCE	20
AIR DISTRIBUTION and QUALITY MAINTENANCE and UNITARY AC/HP	20
HVAC CONTROL and MOTORS PUMPS AND FANS and QUALITY MAINTENANCE and UNITARY AC/HP	15
AIR DISTRIBUTION and QUALITY MAINTENANCE	8
QUALITY MAINTENANCE and UNITARY AC/HP	8

6 Methods

The data sets we developed for this project were designed to test the impacts that different targeting strategies (informed only by data available during the pre-program-implementation period) would have had on the eventual outcomes. To quantify savings for a targeted sub-group of customers, we must first make customer-level savings estimates from the difference in pre/post Normalized Metered Energy Consumption (NMEC). We then compute the average savings of the targeted group after dropping, or filtering out, everyone else. We define targeted subgroups using known customer attributes, like business type (NAICS code), rate type, location, etc. and/or metrics derived from customer meter data, which we call "features." We define the program wide average impact as "all participants' savings" or "unfiltered savings" and the average savings values for sub-groups as "filtered savings." We define the percentage of customers *eliminated by the targeting filter(s)* as the "filter percentage." Finally, we define the improvement in daily kWh savings per-customer brought about by a filter as the "gain", i.e. with daily kWh units, and the percentage improvement above and beyond the unfiltered program performance as the "gain %."

Armed with the ability to compute average savings for the full population of participants and targeted subsets of participants, we conducted two different but related analyses. First, we segmented the participants by known characteristics, like NAICS business category codes and counties, and quantified the extent to which those categories influence savings. We also computed savings results according to the technology types of the measures deployed to quantify expected savings by technology. Second, we computed consumption features, like temperature sensitivity, baseload, and various load shape characteristics, for every customer and calculated the savings performance of every individual and every pair of features when used to filter customers by feature values.

In aggregate, these filtered savings results allow us to conduct an exhaustive search for the characteristics and features that produced the best savings results, and to examine the performance of different feature combinations. Our final step was to select a handful of important and high performing characteristics and features for further discussion.

This section provides an overview of the data preparation and analysis methods we undertook to perform this study. Appendix B walks through specific aspects of our methods in even more detail.

6.1 Defining and calculating at the meter savings

To assess the savings generated at the meter by the DI and HVAC programs, for each participating customer we compute the difference between the daily average energy consumption after the interventions (the post-period) and the weather-normalized baseline derived from data from before the interventions (the pre-period). We then optionally compare the resulting average or distribution of program savings to the average or savings distribution from a control group that has been sampled to resemble the participant group during the pre-intervention period, as described in the Comparison to controls section below. This computational effort was done in R, an open source statistical computing environment, using scripts that heavily leverage the core capabilities of the VISDOM package developed and maintained by Convergence Data Analytics staff.[9]

1. Cleaning data

To ensure that our analysis ran against occupied premises with operating meters, we required that all customers have an average consumption greater than 180W and that no more than 15% of meter readings are missing or zero. We also removed customers with net-metering rates (meaning that they have on-site PV), since they likely have local generation that could require different evaluation techniques than those employed here.

For both samples, we drop customers with less than 120 days of pre or post data and also exclude sites whose energy consumption changed by more than a factor of two between the pre and post periods, i.e. the post-period consumption is no greater than twice and no less than half the pre-period consumption.

The requirement of at least 120 days of pre- and post-data strikes a balance between the NMEC model's need for a significant amount of weather variability to estimate space cooling loads, and the desire for our statistical analysis to include as many customers as feasible. See Appendix B, "Step 1: Requiring sufficient time-series data for evaluation" for more details on our reasoning on this tradeoff.

The factor-of-two cutoff criterion was designed to address the reality that there are very large outliers in the savings distributions *whose apparent savings could not possibly have been delivered by the program*. Large changes (up or down) in consumption are simply not plausible based exclusively on the interventions implemented in the programs studied here. To avoid biasing our results, we wanted criteria that could be applied to outliers in both the negative and positive direction. We also wanted to avoid defining outliers based on absolute savings, since this would risk eliminating real savings values from very large energy consumers. And, once again, we wanted to preserve as many customers as

[9] It's open source. See for yourself: https://github.com/ConvergenceDA/visdom

possible for the rest of the analysis. See Appendix B: "Step 2: Eliminating observed changes that are impossible through efficiency" for more details on our approach to outlier trimming.

After cleaning the data, the steps taken to calculate at the meter savings were as follows.

2. Model Cooling Energy Usage

To isolate cooling energy from total customer load, a prerequisite for weather normalization and estimating weather sensitive loads, we run a weather normalization regression model that explains total daily kWh (KWH) as a function of daily cooling degree hours (CDH) and an indicator for weekend (WKND) or weekday.[10] A day's CDH is the sum of the degrees the outside temperature (Tout) is above 65°F (or 0 if cooler than 65°F) across all hours, h, in each day, d.

$$CDH_d = \sum_{h=1}^{24} max\big(0, (Tout_{h,d} - 65)\big)$$

$$KWH_d = c + \alpha \cdot CDH_d + \beta \cdot WKND_d + \varepsilon$$

The regression coefficient *c* is the expected daily energy consumption for weekdays with zero CDH. The coefficient α quantifies the cooling sensitivity of each customer and can be used to predict daily cooling energy given a computed CDH for day d. This weather normalization model is run for each customer twice, separately using data from pre- and post-intervention periods. We fit the "post-period daily CDH model" using daily meter (KWH_d) and weather data (CDH) from the post period and the "pre-period daily CDH model" using daily meter (KWH_d) and weather data (CDH) from the pre-period.

3. Compute Bulk Daily Energy Savings, aka 'NMEC total energy savings'

The pre-period model fit is used to provide a forecast for each day in the post period (i.e. using post-period days of the week and weather data as inputs). This forecast provides the expected daily energy consumption in the post period absent any EE program interventions or other changes at the site that impact energy consumption and serves a baseline for evaluating program savings. The observed post-period daily consumption (or alternately daily consumption estimates from the post-period model) is then subtracted from the baseline for each day in the post period to provide NMEC savings estimates, in units of daily kWh.

4. Compute cooling savings, aka 'disaggregated HVAC energy savings'

The α coefficients from both the pre- and post-period models can be used to predict disaggregated daily *cooling* energy. The predictions using the pre-intervention coefficient are used as the baseline (aka counterfactual) for how much cooling energy would have been required on post-period days if the efficiency intervention had not occurred. Savings estimates are made by computing the difference between baseline cooling and post-period cooling model estimates.

5. Compute daily average savings

Finally, the daily estimates for both total energy and cooling energy savings are averaged across all post-period days to obtain the expected daily total and daily cooling energy impacts per customer. With estimates available for every customer, average saving can be computed for any sub-group of

[10] This model is consistent with the venerable PriSM piecewise regression methodology, which has updated manifestations in IPMVP Option C, VISDOM, and CalTRACK/OpenEE Meter.

customers. The sub-groups can be based on customer attributes, or criteria derived from meter data consumption features.

6. Repeat for a comparable control sample

The same energy-savings calculations are repeated using a sample of customers who did not participate in the programs being studied. To ensure that the population of control-group participants resembles the program participant group, the control samples were matched 1-to-1 to the DI and HVAC participants using average pre-period daily total consumption and pre-computed site energy categories (S, M, and L, and N, as defined above). Appendix B provides a more detailed explanation of this process.

For the control samples (one for DI and one for HVAC), synthetic intervention dates were generated based on the distribution of real ones in the program participant data to divide control data into pre- and post-periods. As non-participants, they cannot achieve program savings. Any apparent savings are due to other programs and "natural variability" caused by non-program-related activities. The mean and standard error of the control savings can be used as a yardstick against which program savings can be measured while controlling for consumption trends in the general population of customers.

7. Compare to controls

The savings estimates for program participants are based on the assumption that all non-program changes in energy use average out across a sufficiently large set of customers. However, there are long term trends in technology adoption, energy use, and economic conditions that can cause shifts in the entire population of SMB customers. For example, even without program incentives, LED lighting is becoming the default choice for many businesses. The widespread adoption of LEDs is systematically lowering lighting energy use for customers outside of programs as well as inside. Under these conditions, a simple pre/post-NMEC gross savings analysis will include non-program changes and free ridership. To account for trends in the general population of SMB customers, we subtract control group "savings" from our computed program savings. The resulting difference can be more strongly attributed to the programs than simple pre/post savings estimates alone. **However, depending on program design, controls are not always included in NMEC calculations,[11] so our controls were used to verify the existence of savings, but not to compute the at-the-meter savings values reported throughout this work.** For more commentary on the potential role of controls in at-the-meter calculations, see "Synthetic controls for NMEC savings" in the discussion section.

8. Compute savings statistics

With individual savings estimates for every customer, we proceeded with slicing, dicing, and filtering customer groups. The resulting distributions, means, and standard deviations of savings values demonstrate significant variability across customers. These can stand on their own or be compared to the same metrics computed for the control sample.

6.2 Defining and calculating customer attributes and features

We define "customer attributes" as any information known about a customer through their account data. Examples of customer attributes include NAICS codes, rate plans, counties, and utility-assigned "size" of consumption.

---

[11] We note that our work provides some significant evidence that developing synthetic controls would improve the attribution of NMEC savings to program interventions, i.e. help untangle them from natural variability.

We defined "consumption features," or "features" for short, as any information about a customer derived from their pre-program meter data. For this project, we used the open source meter data analysis framework VISDOM to compute roughly 100 consumption features per customer. Most are part of VISDOM's set of "basic" and "weather" features, with several custom features implemented to support NMEC and pre/post modeling. Consumption features can be as simple as average consumption or the variance of consumption over a period of time, or the results of sophisticated estimation and classification algorithms. The coefficients of the NMEC weather normalization models are themselves features, and so are the disaggregated load estimates that they support. In fact, the average daily savings calculations themselves were treated as features to be derived one customer at a time using their pre- and post-intervention data.

We considered a total of 100 individual consumption features in this study and looked at their ability to predict either total daily energy savings (DI) or total daily disaggregated cooling energy savings (HVAC), the features considered are summarized in Appendix A. To facilitate further analysis, we subdivided these into four project-relevant categories.

- **Consumption** – metrics related to the total amount of energy consumed by each customer. Mean kWh/day, kWh in August, and the maximum hourly consumption are all examples of consumption features.
- **Variability** – metrics related to the variability of energy consumption over time. The difference between daily minimum and maximum consumption, the variance of each customer's meter readings, and the ratio of consumption overnight compared to mid-afternoon are all examples of metrics of variability.
- **Thermal** – metrics relating outside temperature to consumption. The correlation between meter readings and outside temperature, the ratio of average summer month consumption to winter month consumption, α - the cooling sensitivity regression model coefficient, and disaggregated total AC consumption are all examples of thermal metrics.
- **Baseload** – metrics that capture the magnitude of always-on loads. The mean of daily minimum consumption, the 3rd percentile of all meter readings, and the average consumption at 2am, are all examples of features that will typically reflect baseload consumption.

6.3 Filtering using customer features

Figure 2 below provides a density plot[12] of values of the *kw_mean* feature (which is the average of all available meter readings) for all DI program participants. The vertical lines are at the 10th, 25th, 50th, 75th, and 90th percentile values of *kw_mean*.

We "filter out" all customers whose *kw_mean* value is below each successive percentile line to eliminate those with lower consumption. In the terminology of this work, each percentile cut point is known as a "filter level" or "filter %." For example, filtering out all customers whose average consumption is below the 75th percentile would be said to be the 75% filter level using *kw_mean*. The resulting sub-group of customers would all have *kw_mean* values in the upper quarter of values and have just a quarter the number of all participants.

[12] A normalized and smoothed histogram-like distribution whose area is 1

Figure 2: The distribution of the kw_mean feature, which is the average consumption across all meter readings, for all DI participants, with the 10th, 25th, 50th, 75th, and 90th percentiles marked using vertical dotted lines. We identify sub-groups of customers whose kw_mean values are greater than each of these lines and say that the filter % for each group is the percentile that all its members exceed.

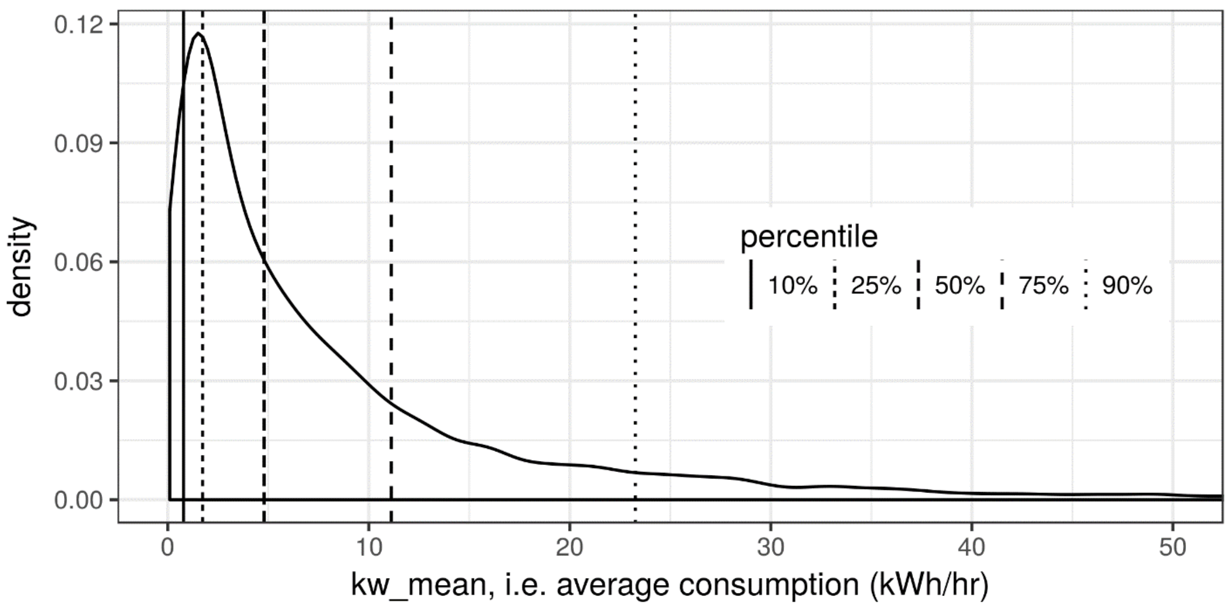

We then compute the mean pre/post NMEC daily savings (in kWh/day) for those remaining customers. As a reminder, we define program-wide average impact as "all participants' savings" or "unfiltered savings" and average savings values for the filtered sub-groups as "filtered savings." Finally, we define the *increase* in savings brought about by a filter as the "gain" of the filter and the percentage improvement above and beyond the unfiltered program performance as the "gain %." The filtered savings are also normalized by the pre-period consumption magnitude to quantify the depth of savings.

This framework of filtering customers based on feature values and computing the resulting gains can be done for any feature or combination of features. For this study, we have tested the performance of every single feature, and of all pairs of features,[13] as filters whose aim is to concentrate the savings of the resulting sub-groups compared to all program participants for both DI and HVAC programs. The results section provides highlights of our findings, including that *kw_mean* is a relatively high-performing filter feature for the DI program, but so are many other metrics that capture different aspects of consumption and baseload demand.

7 Results

7.1 Comparison to controls

Figure 3 shows the distribution of pre/post savings for both the DI (left/blue) and HVAC (right/red) programs, as compared to the "savings" of the non-participant control customers (black) matched to

[13] The pairs did not significantly out-perform single filters and are not discussed in detail in this report.

each. The plot for the DI program displays the change in *total daily consumption*, while the plot for the HVAC program displays the *pre/post change in daily cooling consumption*, as disaggregated by the weather normalization regression model described in the methods section.

One can see that both control group distributions peak near zero savings and display relatively symmetrical positive and negative tails. The peak of the DI distribution is clearly in positive territory, with a diminished negative tail and enhanced positive tail, indicating a positive savings impact (on average) for this program. The evident positive skew of the DI distribution (the blue line is below the black for negative savings and above the black for positive savings), compared to the control group, demonstrates positive savings impact for this program with a reduced proportion of negative savers and an increased proportion of savers. The HVAC program distribution exhibits a modest positive skew compared to its control group as well.

Figure 3: Savings distributions for the DI (left/blue) and HVAC (right/red) programs, compared to control samples selected to resemble the pre-period data of each participant group (black).

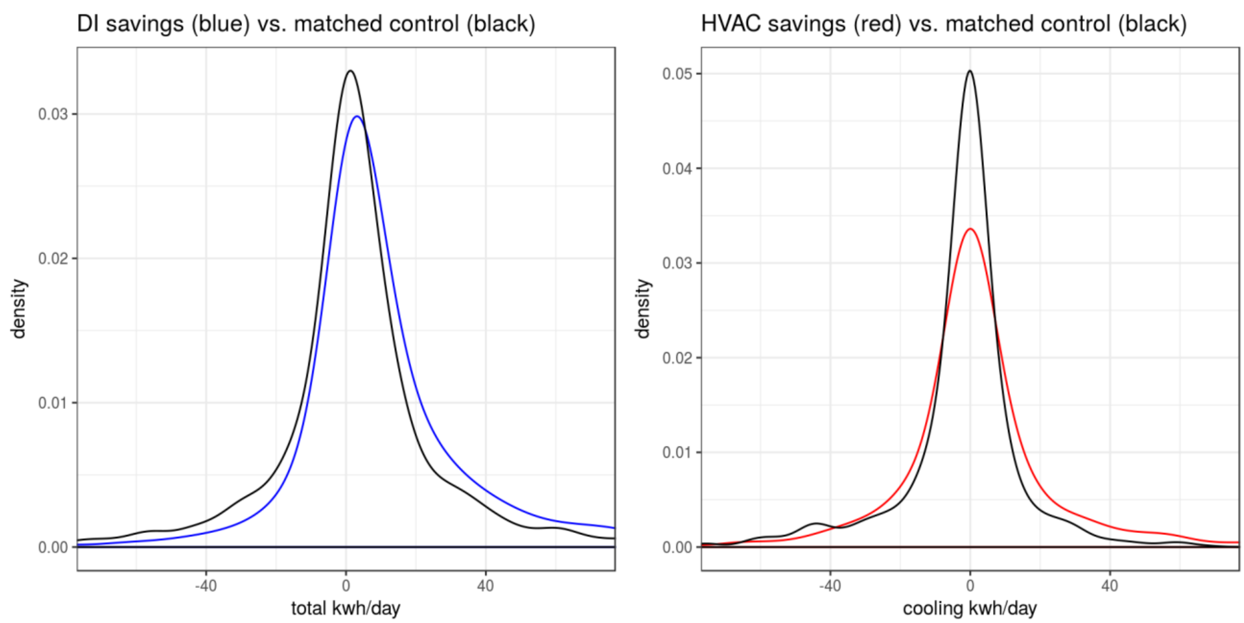

There is a broad distribution of "savings" for both program participants and matched control group. That finding is consistent with other meter-based savings research as it is a direct result of the natural variability in energy consumption of customers (both program participants and non-participants) over time.

To better quantify the overall impact of each program, we compute the average savings value for each participant group and each matched control group. We then compute the difference between the program and control average savings for each program (aka the net savings), as well as the uncertainty (standard error) in this difference. Table 5 shows the results. The uncertainty is significantly smaller than the net savings for both programs, so we conclude that they have both achieved positive and significant net program savings.

Table 5: Average reduction in energy consumption (kWh/day) for the DI and HVAC programs, following outlier rejection, and the net program savings, after subtracting the control-group savings.

Program	Load	# sites	Gross program reduction	Avg control reduction	Net savings	uncertainty
DI	total	7497	14.39	3.16	11.23	0.01
HVAC	cooling	1193	1.03	-3.47	4.51	0.05

Notably, the DI control group shows a reduction in *total energy* consumption (i.e. a trend toward energy reduction), which reduces the net savings estimate. The controls for the HVAC program see a modest increase in *cooling energy* use (negative savings). On their own, HVAC participants realize a small reduction in cooling energy consumption, made more significant by comparison relative to their control-group counterparts. The control group comparisons confirm that there are savings attributable to both programs, but also underscore the potential for significant control group corrections on gross pre/post NMEC savings estimates.

The control matching methods used to establish program level savings rely on the controls being good proxies, on average, for participants. This assumption is defensible for larger samples where the idiosyncratic consumption of a few outliers averages out with others. For smaller samples, on the other hand, matched controls can add more noise to the result than the plain pre/post estimates. The sections that follow focus on the drivers of *relative savings* across smaller sub-groups of program participants. For simplicity, clarity, and to avoid uncertainties caused by imperfect matching, their savings results are based on un-controlled pre/post changes in consumption.

7.2 Customer category/sub-category results

This section summarizes findings from computing average program savings for various customer sub-groups defined by attributes and characteristics expected to influence energy consumption and program outcomes. As seen in Appendix C, hundreds of sub-groups can be specified using one or two types of characteristics,[14] where the types of characteristics include:

- **climate zone** – The California climate zone, as defined by the California Energy Commission.
- **customer size** – Customer size categories S, M, L, characterizing the total amount of energy consumed at each premise annually, with N for customers whose account was opened too recently to have a value computed.
- **NAICS sector** – Standardized business type categories assigned to each premise.
- **rate** – The utility rate each customer is enrolled under.
- **tech family** – The technology family for the EE intervention(s) performed, including HVAC, lighting, refrigeration, electronics and IT, appliances, boilers and steam systems, industrial systems, motors, building envelope, pumps and fans, and food service technology.
- **program year** – The year each EE intervention was completed.

[14] This report excludes groups with fewer than 20 members. This helps to protect the anonymity of customers, but such small groups are also unreliable indicators for larger trends in the data – they are too likely to be dominated by a small number of outliers.

Figure 4 below presents a comprehensive look at the sub-group average savings (x-axis) and the sub-group average savings as a percentage of pre-consumption loads (y-axis), for the DI (red) and HVAC (blue) programs, across all sub-groups examined. Circle sizes correspond to the premise count of the sub-groups. The data for all the visualized groups are detailed line by line in the tables found in Appendix C. Based on this overview of customer sub-category results, we conclude:

- The savings associated with the DI program are more prominent than savings associated with the HVAC program.
- For both programs, the largest savings (further right) come from very small groups – some, no doubt, are dominated by a small number of outliers.
- DI results based on groups of 200 or more customers approach 40-50 kWh/day savings.
- The magnitude of savings as a percentage of pre-intervention consumption roughly correlates with absolute savings, but there are prominent exceptions.
- The highest savings as a "% of pre-intervention consumption" can be seen to be associated with below average absolute savings. This suggests that the smaller customers are more likely to have their consumption dominated by the end-uses addressed by the programs.

Figure 4: Sub-group average savings (x-axis) vs. average savings as a percentage of pre-intervention consumption (y-axis) for all customer sub-groups based on characteristic types, sized by the number of premises in each group, for both DI (red) and HVAC (blue) participants.

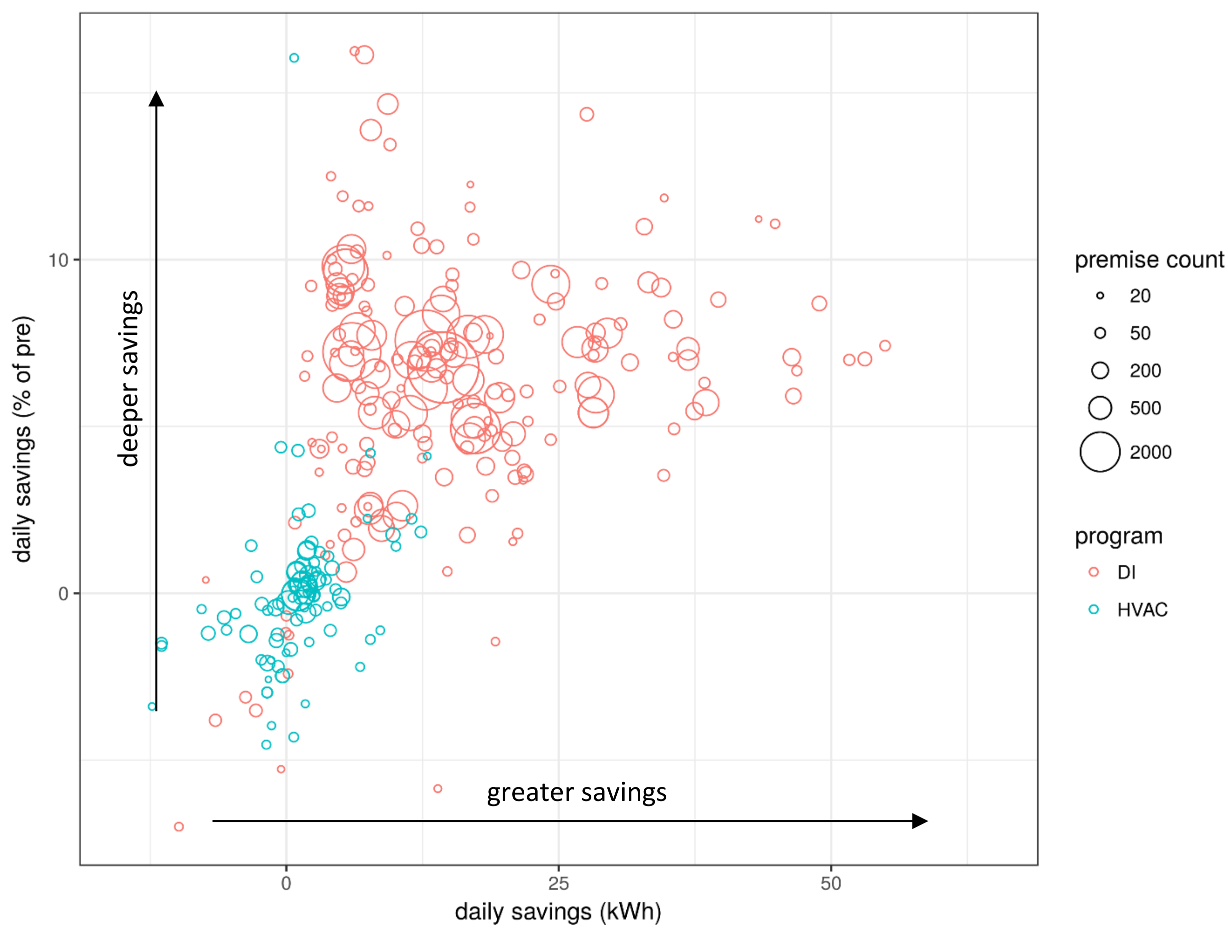

7.2.1 DI program customer sub-categories

Table 6 provides highlights of the key findings and best performing sub-groups within the DI program. See the "Tabulation of DI savings by customer characteristics" section of Appendix C for all the supporting details. **The summary table supports the following observations:**

- Fitting with long established targeting rules of thumb, DI savings track customer size (magnitude of consumption), categories well.
- Smaller customers tend to save a greater percentage of their pre-intervention consumption than larger customers, indicating deeper savings are being achieved at such sites, most likely because they are more likely to have consumption that is dominated by a program-accessible end use.
- Customers on the "Medium" rates (E19 and A10) significantly out-perform the average DI program participant, which is consistent with the correlation between customer size and savings.
- Customers on TOU rates tend to out-perform their peers. Customers on E19 save an average of 168% more than the average DI program participant! Possibly more striking, the group on the TOU rate amongst "Small" customers (A6) out-saves the general population by 19%, while the standard "Small" rate (A1) saves 58% *less* than the general population.
- Customers who received lighting measures alone out-save customers who received only refrigeration measures, but each of these sub-groups perform worse than the full DI program average. When lighting and refrigeration are done together, the savings are a notable 105% greater than the DI program average. Beyond the additive impacts of addressing both end-uses, this might have to do with the types of customers, like grocery stores and restaurants, that are eligible and elect to complete both interventions.
    - This general pattern holds across size categories, with average savings from "lighting and refrigeration" interventions for large (L) customers being a remarkable 269% greater than the DI program average.
- Lighting interventions tend to save a little over 7% of pre-intervention total consumption, but refrigeration saves just over 2%.
- Refrigeration savings as a % of pre-intervention total consumption tend to decrease with size (refrigeration loads are usually a smaller % of the total load for large customers). The pattern is a little less clear for lighting. At nearly 10%, lighting savings as a % of pre-intervention total consumption are greatest for size S customers, but lowest at 5% for size M.
- Within the lighting category, LEDs (24% higher than average) substantially out-save linear fluorescents (49% less than average), and CFLs (85% less than average). Projects involving both LEDs and linear fluorescents (55% higher than average) are observed to have the best savings performance.
- Savings from walk-in coolers are modest, but savings from walk-in coolers and controls together are nearly 140% greater than all refrigeration projects. Controls appear to be key drivers of walk-in cooler savings (or this result is dominated by a few outlying projects that addressed both).
- Savings from sports, entertainment, and recreation venues, non-department stores, and more technical manufacturing are all around 70% greater than DI projects in general.
- Savings go up and down a bit from year to year, but in expectation, every year will return savings at about the average for all DI. However, customers whose participation spans more than one year save 130% more than typical DI participants. This is likely partially due to the

cumulative impact of multiple interventions and partially to the self-selection effect of highly motivated customers participating multiple times in programs.

Table 6: Key DI program outcomes by customer characteristics. Recall that % gain is savings above and beyond the program average, so all positive values represent improvement over average outcomes.

category	Sub-category group	premise count	daily savings as % of pre-intervention usage (%)	daily savings (kWh)	% gain[15]
size	L	1025	5.41	28.17	96
	M	3403	4.94	17.36	21
	S	2595	9.65	5.44	-62
rate	E19 Medium general demand TOU	720	5.72	38.51	168
	A10 Medium general demand	1631	5.95	28.41	97
	A6 Small general service TOU	275	7.82	17.12	19
	A1 Small general service	4766	7.24	6.01	-58
DI tech. family	LIGHTING and REFRIGERATION	1005	7.80	29.46	105
	LIGHTING	5331	7.57	12.77	-11
	REFRIGERATION	772	2.32	10.10	-30
	LIGHTING and REFRIGERATION & size L	112	7.02	53.10	269
lighting	LED and LINEAR FLUORESCENT	645	9.7	19.8	38
	LED	2065	9.0	15.8	10
	LINEAR FLUORESCENT	995	3.3	6.5	-55
	COMPACT FLUORESCENT	40	-0.6	1.9	-87
refrigeration	REFRIGERATION CONTROL and WALK-IN COOLER	42	9.0	24.0	67
	WALK-IN COOLER	687	1.9	9.6	-33
program year	more than one	386	9.32	33.23	131
NAICS code	Arts, Entertainment, and Recreation[16]	218	8.75	24.74	72
	RETAIL TRADE - 1[17]	1786	9.26	24.27	69
	MANUFACTURING - 3[18]	62	4.60	24.26	69

[15] % savings greater than the average savings across all premises

[16] Arts, Entertainment, and Recreation is concert halls, sports venues, museums, etc. It includes all NAICS codes starting with 71: https://www.bls.gov/iag/tgs/iag71.htm

[17] Retail Trade - 1 is basically non-department stores. It is composed of all the NAICS codes starting with 44 at this location: https://www.bls.gov/iag/tgs/iag44-45.htm. All code starting with 45, basically department stores, are Retail Trade – 2.

[18] Manufacturing – 3 is more technical manufacturing, including the production of metal products, machinery, and electronics. It is all NAICS starting with 33 here: https://www.bls.gov/iag/tgs/iag31-33.htm

7.2.2 HVAC program customer sub-categories

Table 7 (below) provides highlights of the key findings and best performing sub-groups within the HVAC program. See the "Tabulation of HVAC savings by customer characteristics" section of Appendix C for all the supporting details. **The summary table supports the following observations:**

- HVAC program savings are loosely correlated with customer size, but the effect is not nearly as strong as it was for the DI program.
- One might expect HVAC saving to strictly correlate with hotter climate zones. However here we see that the hottest climate zones, cz12, and cz13 are not the strongest performers. The northern Central Valley, cz11, and northern coast including the Bay Area, cz03, perform best. Note that a quality maintenance program addresses other aspects of air and water distribution in addition to the AC units systems themselves.
- As with DI, the "Medium" rate class, especially the TOU version, out-performs the general population of HVAC program participants.
- Chiller projects dramatically out-perform other types, with unitary AC projects associated with above average savings. Notably, the quality maintenance interventions are associated with below average savings – the actual average is negative, but this is likely just a symptom of the variability being so much higher than the savings so that negative outliers can dominate average outcomes.
- As with DI, program year 2015 had noticeably better results than others, with program year 2013 returning noticeably worse results than others. Unlike DI, premises with interventions spanning more than one year did not perform better than their peers. At least some of the year over year variability in outcomes could be due to imperfect weather normalization.
- Accommodation and Food Service has the largest average savings by far, followed by Public Administration and Retail Trade – 2. None of these has a very large premise count, however.

Table 7: Key HVAC outcomes by customer characteristics.

category	Sub-group	premises	daily savings (% of pre)	daily savings (kWh)	% AC gain
size	L	694	0.27	1.52	48
	M	327	0.66	0.96	-7
	S	119	-2.47	-0.36	-135
climate zone	cz11	98	1.51	2.30	122
	cz03	282	0.55	2.06	99
	cz12	345	-0.59	1.78	72
	cz04	202	0.41	1.33	28
	cz13	163	-2.09	-1.75	-269
	cz02	65	1.43	-3.23	-413
rate	E19 Medium general demand TOU	249	0.36	2.76	167
	A10 Medium general demand	354	0.16	1.94	88
technology	CHILLER	45	0.7	7.2	593
	UNITARY AC/HP	721	0.3	1.9	83
	QUALITY MAINTENANCE	90	-0.4	-2.3	-323

category	Sub-group	premises	daily savings (% of pre)	daily savings (kWh)	% AC gain
program year	2015	235	-0.11	5.03	387
	more than one	189	-0.43	-0.98	-194
	2013	224	-1.22	-3.48	-437
NAICS code	Accommodation and Food Service	49	2.23	11.49	1011
	Public Administration	34	4.20	7.72	647
	RETAIL TRADE - 2	39	-1.39	7.70	645

7.3 Filter results

As described in the methods section under "Filtering using customer features," consumption feature values (computed using only pre-intervention data) were used to construct "feature filters" that define program participant sub-groups whose average savings magnitude or depth (relative to total consumption) can be compared to the average performance across all participants. A typical feature filter uses a threshold value to select a subset of customers, for example all customers whose mean daily consumption is greater than 100 kWh. The thresholds are selected so they eliminate, or filter out, 10%, 25%, 50%, 75%, and 90% of all customers when applied, keeping 9 in 10, 3 in 4, 1 in 2, 1 in 4, and 1 in 10 of the original customers, respectively. When a consumption feature correlates with program savings, the average savings of the filtered sub-groups are larger than the average for all participants and the feature filter can be said to yield savings gains.

7.3.1 Understanding how specific consumption feature filters perform

There were 100 consumption features computed for every premise and evaluated as feature filters for this project. The complete list of those features can be found in Appendix A: Definitions. Subsequent figures utilize a subset of features to illustrate the savings gains from consumption feature filtering. Table 8 lists and defines those features. Each performs well in predicting either savings magnitude or savings depth (or both) for either the DI or HVAC program (or both).

Table 8: The set of features used to illustrate filter performance in this section. These include some of the top performers for both DI and HVAC programs for enhancing savings and savings as a percentage of pre-intervention consumption.

feature	application	definition	units
Aug_range	DI	Aug. mean of daily range of demand (daily maximum – minimum)	kW
kw.mean	DI, HVAC	mean demand (all data)	kW
kw.tout.cor	HVAC	correlation between electric consumption and outside temperature	
kw.var.summer	DI, HVAC	electric demand variance (summer)	kW
sum2win	HVAC	ratio of total consumption during summer months to winter months	
mx2mn	DI	ratio of daily maximum consumption to daily minimum consumption	
pre_CDH	DI, HVAC	pre-period modeled temperature sensitivity	kWh/day/F
pre_CDH_pct	HVAC	pre-period temperature sensitivity as a % of daily kWh	%/F
pre_daily.cooling.kwh	HVAC	pre-period modeled daily cooling load	kW

feature	application	definition	units
discretionary	DI, HVAC	Non-baseload consumption (consumption above the daily minimum)	kW
discretionary_pct	DI, HVAC	discretionary consumption as a percentage of total consumption	%

7.3.2 Savings magnitude vs. depth

Consumption at the meter is determined by **three main factors**:

1. The **magnitude** of service demand
2. The **timing** of service demand
3. The **efficiency of the equipment** used to meet the service demand

Efficiency interventions can impact any or all of these factors. In some cases, the service being provided (like lighting) is greater than anyone requires. This results in waste that can be corrected by reducing the level service provided. For example, fixtures could be removed from an over-lit room. In other cases, the level of service is appropriate, but it is being provided at time when it is not needed. For example, lighting controls might switch or dim lights based on time of day or occupancy. Finally, whatever the service demand or its timing, there may be new equipment that can meet the demand using less energy. For example, LEDs can replace linear fluorescent lights.

For implementers, there are two main **strategies to increase project savings**:

1. **Focus on larger magnitude consumers.** For example, focus on factor 3 and perform the same LED for linear fluorescent lighting swap for customers with more square footage and lighting fixtures.
2. **Perform deeper improvements for customers with especially wasteful or inefficient systems.** For example, focus on factors 1 and 2 to identify and swap LEDs for incandescent lights or add controls to always-on fixtures.

If strategy 1 for increasing savings is dominant, outcomes will tend to have improved magnitude of savings, but fixed depth of savings as consumption feature filters go deeper. If strategy 2 is dominant, outcomes will tend to have improved depth of savings as consumption feature filters go deeper. Because project costs, engineering, and equipment profiles are different for each of the two strategies, a given implementer might choose to tune targeting efforts to better support one or the other strategy. For this reason, our filter results are presented in terms of both savings magnitude and depth.

7.3.3 Filter performance outcomes

7.3.3.1 DI program

The savings associated with the selected DI feature filters used to filter from 10% through 90% of customers are plotted in Figure 5. For both panels, the filter depth is on the x-axis. The left panel places mean daily savings of the filtered sub-group of customers (magnitude of savings) on the y-axis and the right panel places savings as a percentage of pre-intervention consumption (depth of savings) on the y-axix. The secondary y-axis of the left panel provides the gain %, or the percentage by which filtered groups exceed average program savings, e.g. a % gain of 100 indicates a doubling of savings. The dotted horizontal line is the DI program average savings. The dashed lines are the hypothetically perfect filter outcome where we use our *knowledge of actual program outcomes* to eliminate customers from lowest to highest savings (or percentage savings) – the equivalent of a consumption feature with perfect prediction of outcomes.

In the left panel it can be seen that three consumption feature filters, average consumption (*kw.mean*), daily non-baseload energy consumption (*discretionary*), and the average daily difference between min and max consumption in August (*Aug_range*) double the magnitude of savings (a gain % of 100 or greater) between 50 and 60% filter depth. *kw.mean* has an advantage that grows more prominent at even higher filter depths. It triples unfiltered savings just below 80% filter depth.

In the right panel, it can be seen that the features that amplified the magnitude of savings did not amplify the depth of savings. Instead the ratio of maximum to minimum consumption for each day (mx2mn) and the percentage of daily consumption represented by non-baseload (*discretionary*) are the top performers. These are both metrics of variability in consumption, suggesting that variability metrics are better at isolating deeper savings opportunities.

A metric of cooling loads that plays an important role in the HVAC results (*pre_CDH*) is notable for improving the magnitude of savings to some degree at the expense of depth of savings. HVAC loads are not addressed by the DI program, so the most likely explanation for this pattern is that HVAC loads tend to correlate with DI-relevant loads, but because DI doesn't address HVAC loads, customers with HVAC loads that make up a large fraction of their total consumption achieve smaller savings depths.

Figure 5: Filter performance as mean daily savings magnitude vs. depth for a representative sample of top DI feature filters. Dotted horizontal line is the unfiltered population average savings. Dashed line is the perfect filter outcome that truncates the savings distribution from least to most to achieve the filter % desired.

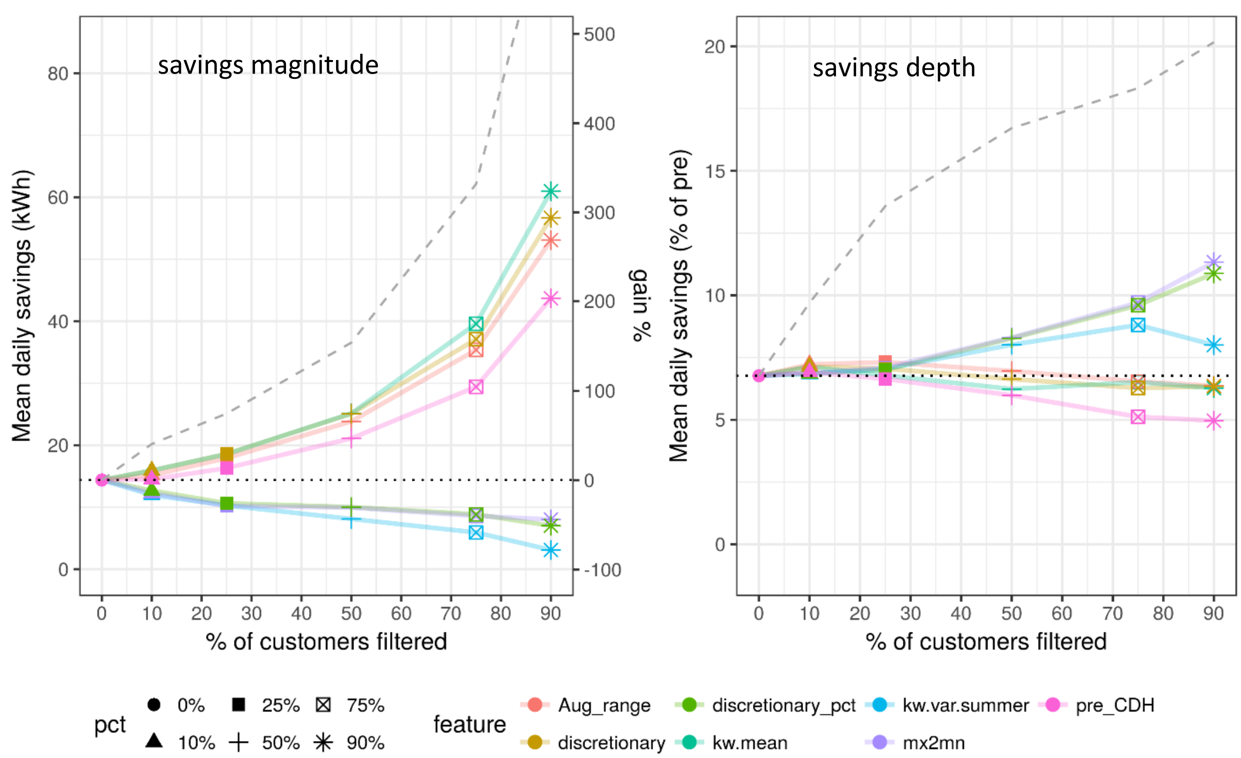

The DI program is dominated by two interventions for two separate end uses: lighting and refrigeration. The next sections present filter performance plots for each in isolation.

7.3.3.2 DI refrigeration outcomes

Refrigeration is an end use that reliably fits into the 24x7 baseload of a building. Some business types, like restaurants and grocery and liquor stores, will tend to have significant refrigeration loads while others, like office spaces, will not. Figure 6 presents the results of consumption feature filters applied to the subset of DI participants who's only intervention technology family was refrigeration. Among those participants, all of the filters that scale with consumption magnitude amplify savings. Even cooling loads and non-baseload scale well enough with baseload that filters based on them amplify refrigeration savings. The observation that refrigeration saving scale in lockstep with total consumption is bolstered by the lack of any stand-out filter in the area savings depth. Refrigeration products have long been shaped by codes and standards. At this point, it is probable that highly inefficient equipment is rare, so DI refrigeration projects deliver incremental efficiency gains whose magnitudes of savings are primarily a function of the underlying service demand for refrigeration, which in turn drives total consumption.

Figure 6: Filter performance as mean daily savings magnitude vs. depth for a representative sample of DI refrigeration feature filters. Dotted horizontal line is the unfiltered population average savings. Dashed line is the perfect filter outcome that truncates the savings distribution from least to most to achieve the filter % desired.

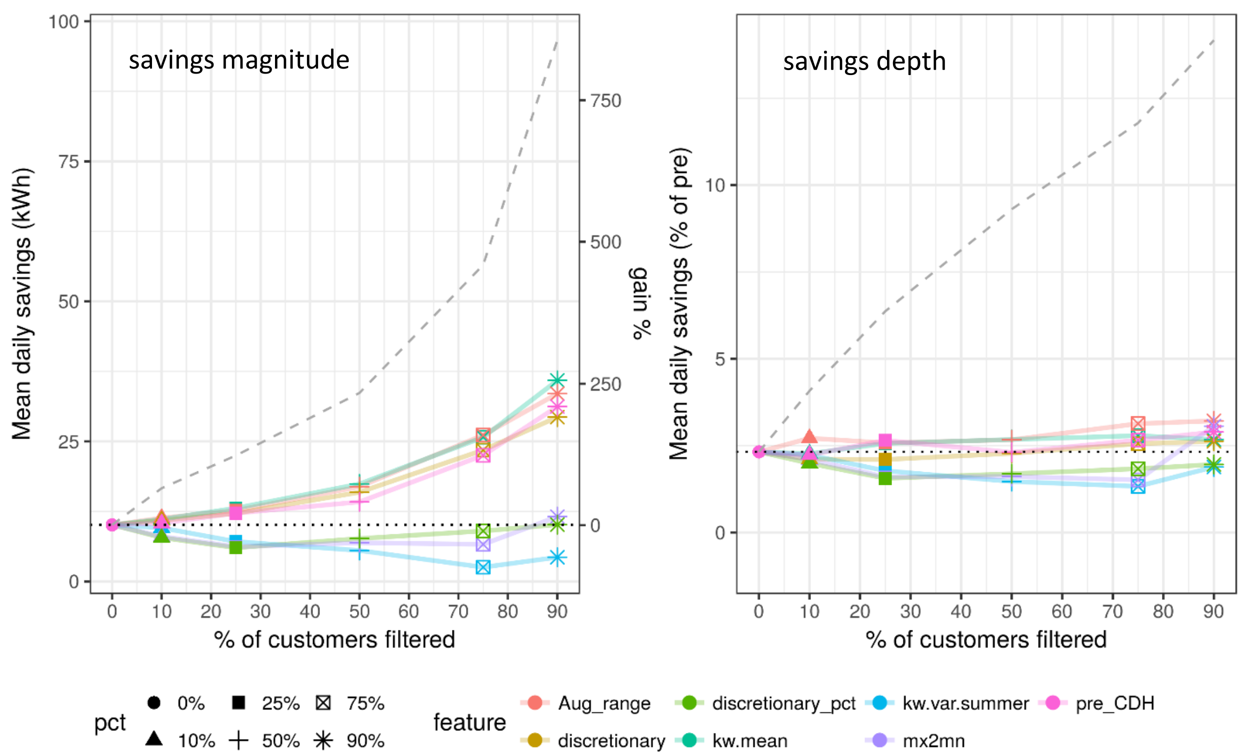

7.3.3.3 DI lighting outcomes

As illustrated in the discussion of savings magnitude vs. depth above, lighting efficiency can include changes to service demand (i.e. de-lamping), tighter controls, and/or device efficiency gains. In other words, some project approaches should be expected to achieve simple efficiency improvements that scale with total consumption, while others will go deeper with improvements to controls as well as luminous efficacy. Figure 7 presents the results of consumption feature filters applied to the subset of DI participants who's only intervention technology family was lighting.

The left panel shows significant gains in savings from consumption feature filters that quantify the magnitude of consumption. While the basic mean or total of overall consumption breaks ahead at the 90% filter depth (1 in 10 customers preserved), it is in a virtual dead heat with the *Aug_range* and *discretionary* features. These features both quantify the magnitude of the loads that occupants control, so it makes sense that they would correlate with lighting loads.

The right panel shows deeper savings can be achieved through concentrating on customers with elevated discretionary loads as a percentage of total loads (*discretionary_pct*) or with elevated maximum daily demand as a multiple of daily minimum demand (mx2mn). Both of those features relate to the fraction of total loads that are occupant-controlled, so this depth of savings finding is consistent with the understanding that most lighting loads (even automated ones) are ultimately occupant driven.

Figure 7: Filter performance as mean daily savings magnitude vs. depth for a representative sample of top DI lighting feature filters. Dotted horizontal line is the unfiltered population average savings. Dashed line is the perfect filter outcome that truncates the savings distribution from least to most to achieve the filter % desired.

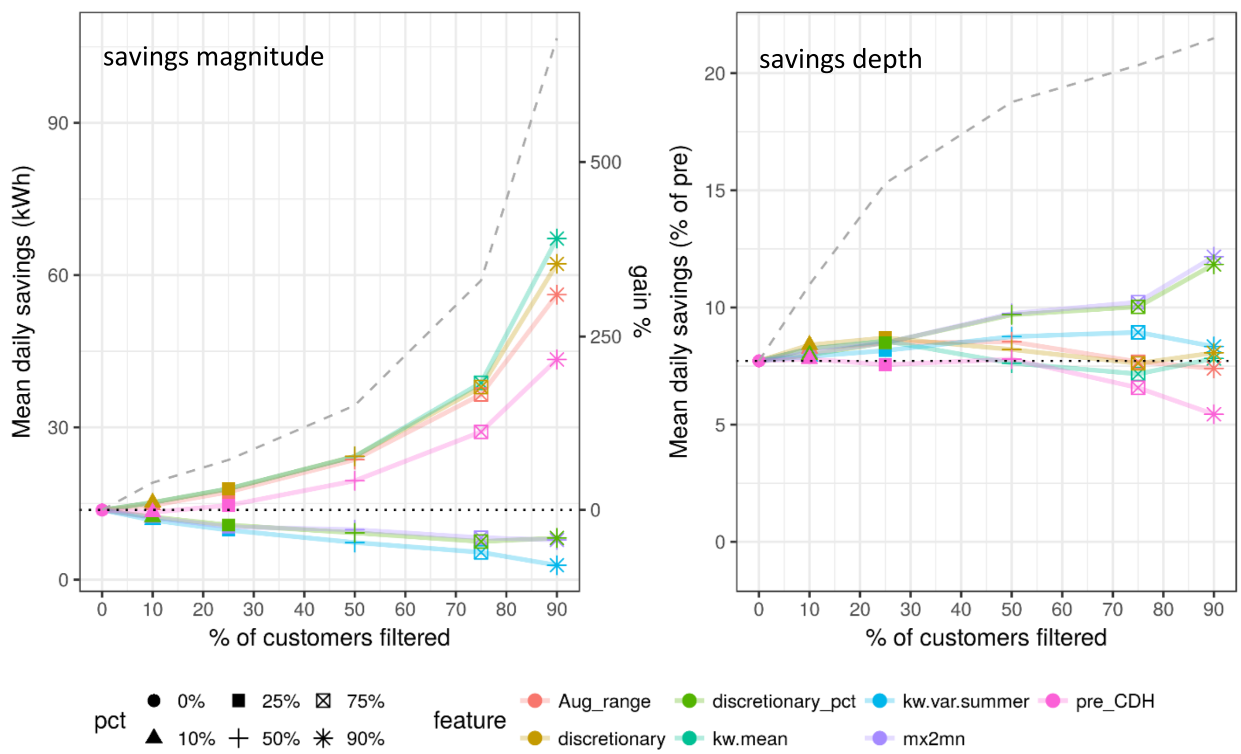

7.3.3.4 DI filters within customer category sub-groups

Program planners and implementers would be well advised to use every tool at their disposal in designing programs to maximize savings. In that spirit, this section presents the results of performing feature filtering within prominent customer and program categories.

7.3.3.4.1 CUSTOMER SIZE

The simplest framing of DI program results is that size, meaning total consumption, is a good predictor of program savings. If proxies for size are to be the focus of DI program targeting, one might wonder how well features can drive savings within the pre-computed customer size categories. Figure 8 illustrates the results of running feature filtering within the major consumption "size" categories S, M, and L, each in a separate panel. The dashed horizontal lines show the average savings for each size category.

Even with the data pre-grouped by size categories, the best performing of the example filters are all related to the magnitude of consumption, with average/total consumption turning in the best performance. Instead of capturing all of the targeting potential related to total consumption, the **size categories appear to be complimented by the more precise, individual consumption metrics**.

The figure also suggests that targeting cannot squeeze enough program performance out of small customers to match the magnitude of the large ones, but projects for smaller customers tend to have deeper savings. Filter depths of 50% or greater applied to size M customers allow their savings to surpass the average across all size L customers and there are 3 times as many M customers in the sample. For depth of savings, it is also noteworthy that each size category has a different top performing consumption feature filter. This suggests that different project strategies are being employed for different customer sizes.

Figure 8: Filter performance within customer size categories. Dotted horizontal lines are the average savings for each size category and dashed lines are the theoretically perfect filter performance for each size category.

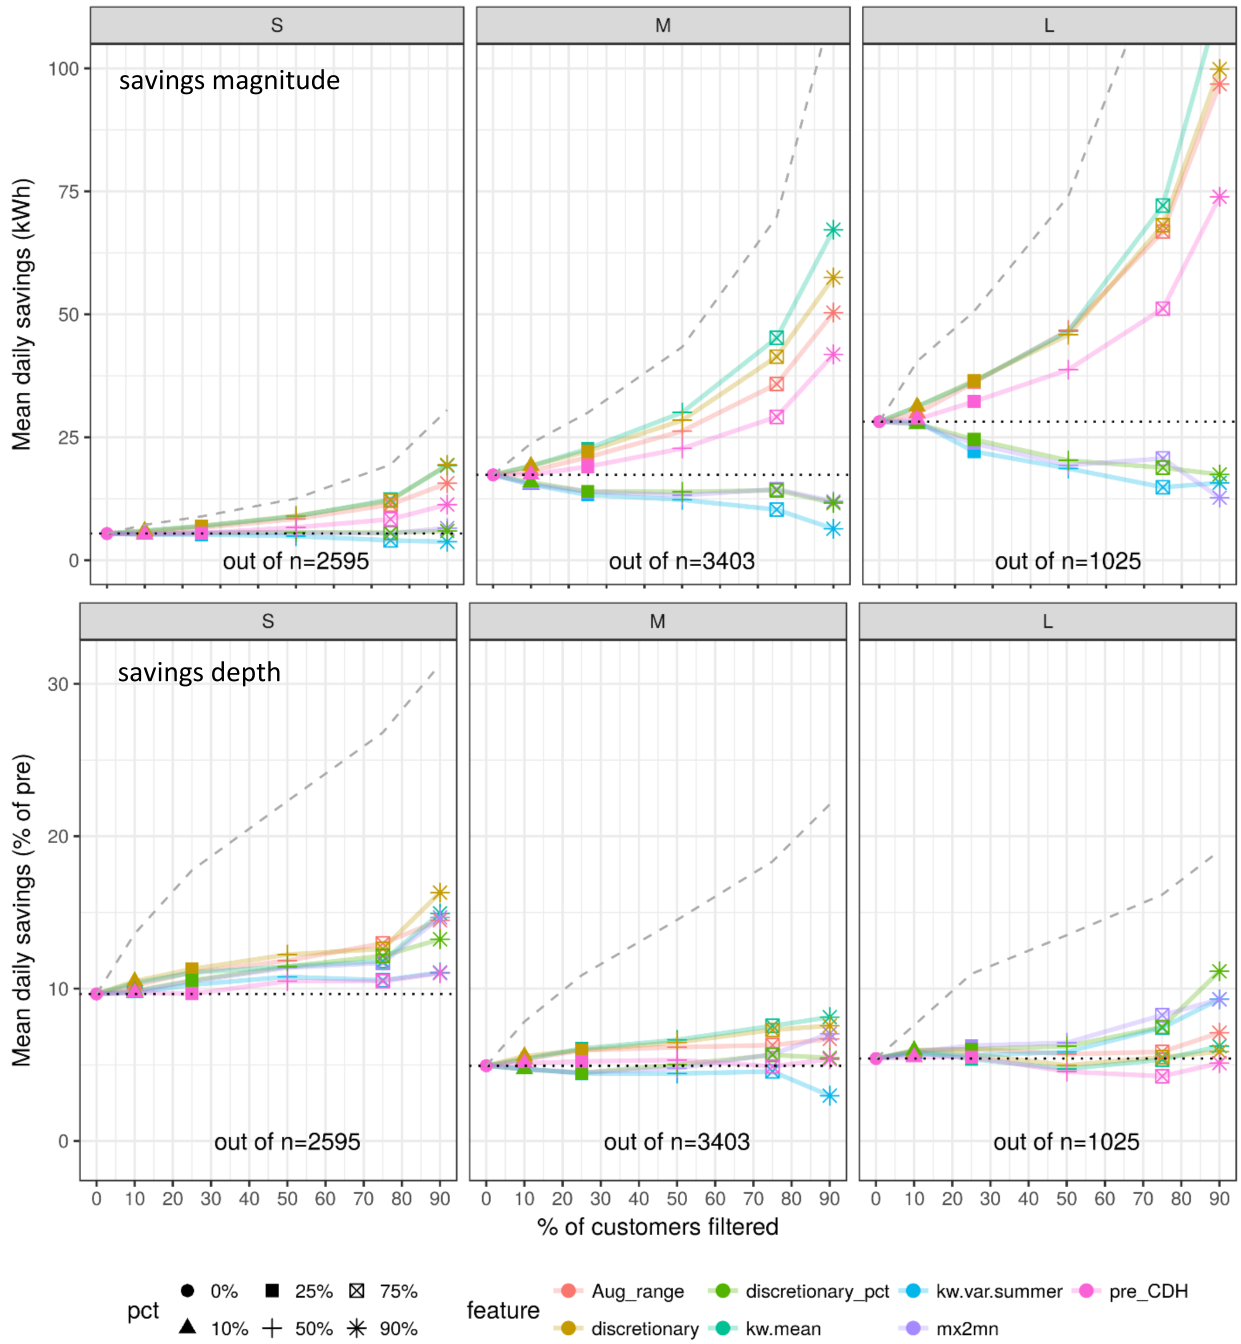

7.3.3.4.2 RATE TYPE

Another category with notable DI program performance gains is the utility rates of participating customers. Figure 9 presents the results of feature filtering within the rate categories. As with the size categories, the unfiltered averages are smaller for the small customers (rates A1 and A6) and larger for the medium customers (rates A10 and E19). The time of use (TOU) rate customers (A6 and E19) easily out-perform their standard rate counterparts (A1 and A10, respectively). It also appears that the smaller

A6 customers can match the performance of medium-sized A10 customers at filter depths of about 50%, but the sample of A6 customers is too small to have high confidence in those specific numbers. Due to the strong showing from TOU customers, it would be useful to disentangle the self-selection effect (enthusiastic customers who opt-in to TOU rates are possibly more likely to opt into EE programs and do well in them; TOU enrollment may also correlate with higher average consumption) from the pricing effect of the rates. As customers are steered into TOU rates by default, this distinction will determine the durability of the TOU correlated savings documented here.

Figure 9: Filter performance within rate type categories. Dotted horizontal lines are the average savings for each rate category and dashed lines are the theoretically perfect filter performance for each size category.

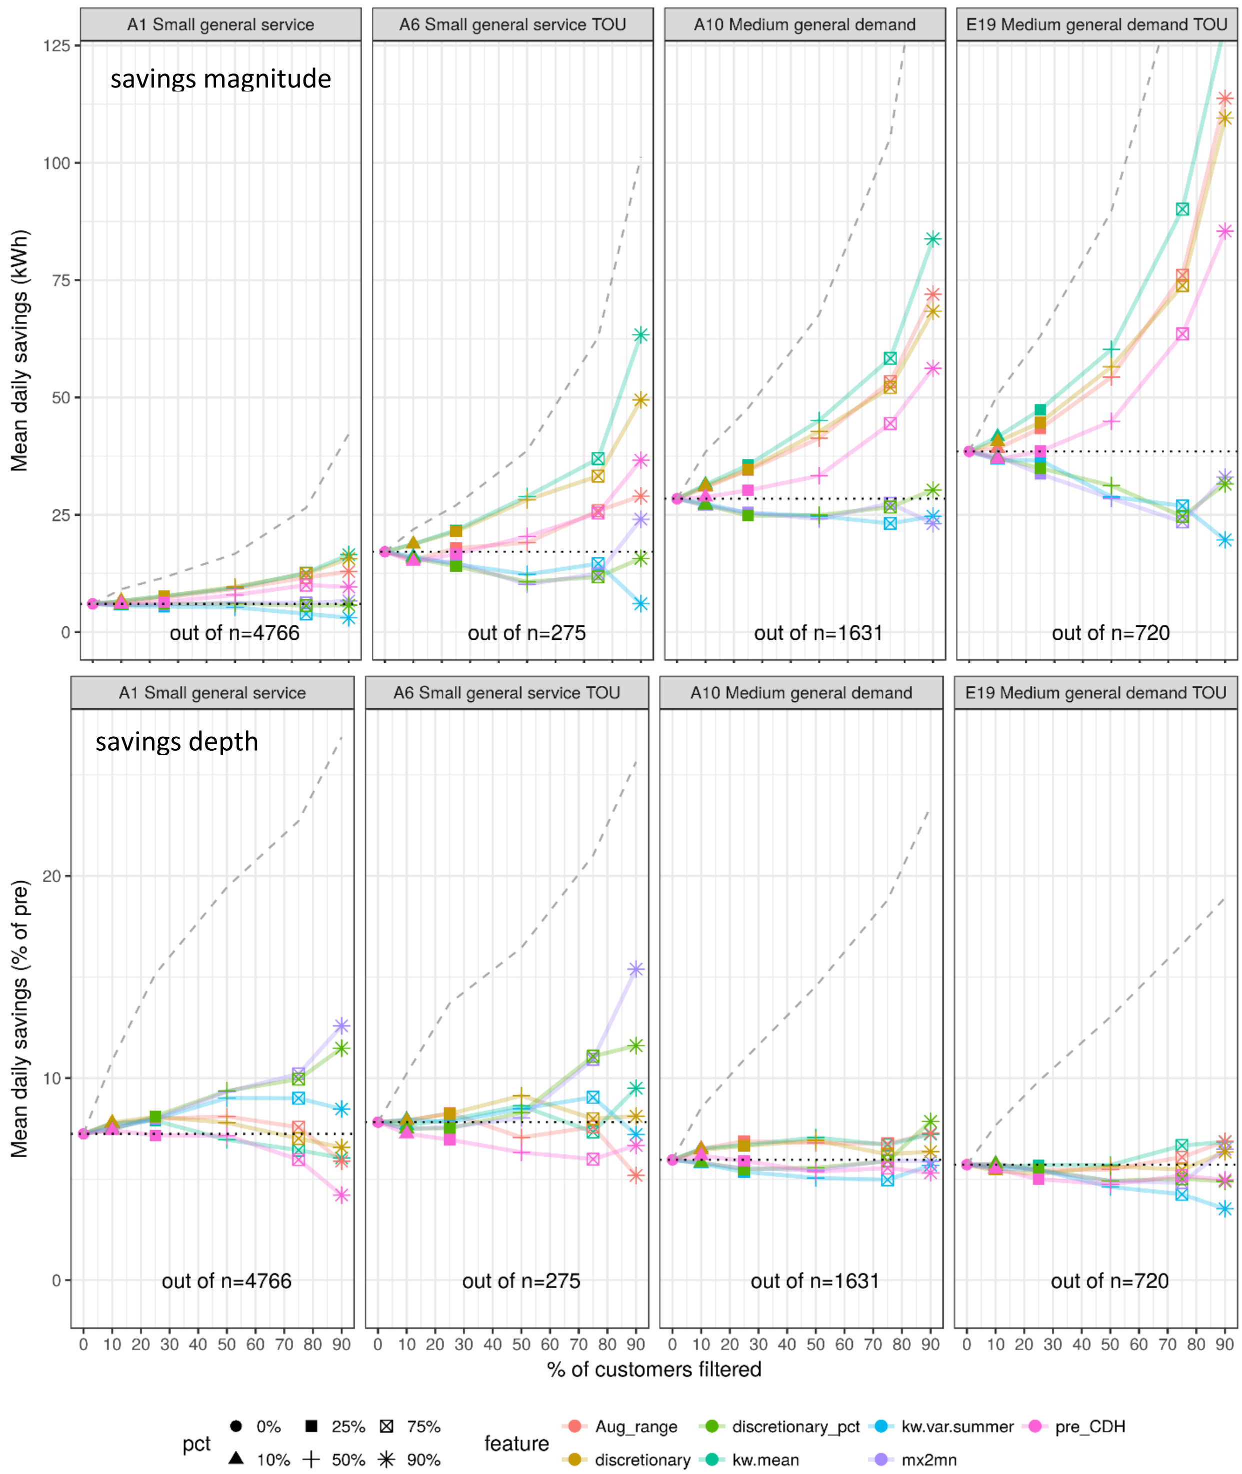

7.3.3.5 HVAC program

The example features for HVAC include more features that correlate with AC usage, like the degree to which consumption increases in the summer (*sum2win*), at higher temperatures (*kw.tout.cor*), and weather normalization outputs for temperature sensitivity (*pre_CDH* and its normalized form *pre_CDH_pct*) and AC consumption (*pre_daily.cooling.kwh*).

The example HVAC program features are plotted in Figure 10 with filter % from the 10th through the 90th percentile on the x-axis. The left panel places the magnitude of mean daily savings on the y-axis and the savings gain % on the secondary y-axis. The right panel places the savings as a percentage of pre-intervention total consumption on the y-axis. The dotted horizontal lines are the HVAC program average savings. The dashed line is the hypothetically perfect filter outcome where customers are eliminated from lowest to highest savings – the equivalent of a consumption feature with perfect prediction of outcomes.

Figure 10: Filter performance as magnitude of mean daily savings(left) or depth of mean daily savings (right) vs. depth of filter % for the example HVAC feature filters. Dotted horizontal line is the unfiltered population average savings. Dashed line is the perfect filter outcome that truncates the savings distribution from least to most to achieve the filter % desired.

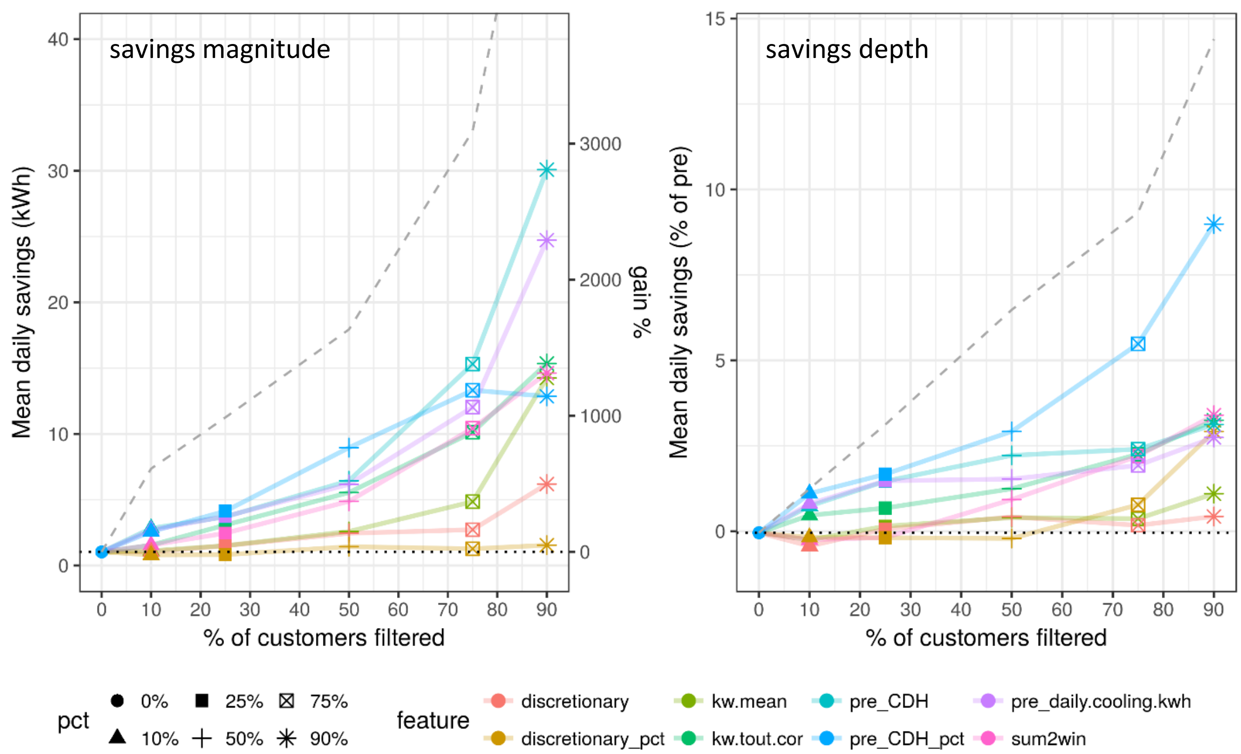

The unfiltered average savings are very small for the HVAC program, but the temperature correlated filters are able successfully predict premises with significant savings - 10-15 kWh/day with 75% of customers filtered out. In contrast to the results from DI program filters, features that amplify the magnitude of HVAC program savings also amplify the depth of savings. This suggests there is a wider range of options for achieving deeper savings through HVAC interventions than addressing larger HVAC loads and that HVAC loads can vary to a significant degree independent of the magnitude of other loads. For example, restaurants have high ventilation requirements in their kitchens, resulting in HVAC loads much larger than a typical office building of similar size. Updates to controls and scheduling, better system zoning, and more efficient equipment can all contribute to savings.

The *pre_CDH* feature based on modeled temperature sensitivity performs best in amplifying the savings magnitude and the same feature as a percentage of pre-intervention consumption performs best in amplifying the depth of savings. The regression approach that produced those features is a more precise way to isolate the temperature effects from other sources of variability than the other thermal filters.[19]

Even though general efficiency savings potential is often observed to increase with the magnitude of consumption, *kw.mean* is a much weaker predictor of savings outcomes than the more specialized thermal features for the HVAC program. Even though HVAC loads might be considered as primarily discretionary loads that are above and beyond the baseload, the *discretionary* and *discretionary_pct* features are poor performers.

The DI results compared to the HVAC results are a good example of different programs with different intervention types being best targeted by different features. In the case of DI, the end-uses correlate well enough with total consumption that total consumption turns out to be one of the best feature filters. In the case of the HVAC program, features more finely tuned to pick out cooling loads from meter data are a better choice.

7.4 General patterns in the results

The results presented in this paper are drawn from two specific programs, but the methods are designed to be broadly applied. This section provides analysis of the general patterns present in the DI and HVAC program intended to building intuition for how the methods and insights might generalize to other programs and situations.

With a set of just over 100 consumption features, each designed to quantify a different aspect of consumption, it is important to contemplate which features are relevant to which aspect of program performance. To study how well different types of consumption features isolate different EE-program-relevant load characteristics, they are categorized into four “feature families”:

- **baseload** – Features that isolate always-on loads. For example, daily minimum consumption.
- **consumption –** Features related to the total magnitude of consumption. For example average consumption.
- **thermal** – Features related to the correlation between outsode temperature and consumption. For example, the ratio of summer to shoulder season consumption.
- **variability –** Features related to how variable consumption is. For example, the average daily range from minumum to maximum consumption or the ratio of overnight to mid-day consumption.

See Appendix A for a comprehensive list of the features cateloged in each family for the discussion that follows.

[19] The regression model is also the basis for the evaluation of savings in the first place, so the pre_CDH feature has a bit of a structural advantage over the others.

7.4.1 Cross comparison of results

Figure 11 depicts the best consumption feature filter performance at each filter depth for all of the categories of DI participants studied, with the premise count of each sub-group along the x-axis and the multiplier for the gains over the unfiltered average savings on the y-axis. In other words, the plotted points correspond to the single best feature filter at the given filter depth. Ideal filters will maximally concentrate average savings while eliminating as few customers as possible. By these criteria, the consumption feature filters applied to all DI participants (i.e. the black points illustrating performance without narrowing to within specific customer categories) are the highest performing filters.

Figure 11: Summary of consumption feature filter performance for 50%, 75%, and 90% depths (keeping 1 in 2, 1 in 4 and 1 in 10 customers respectively) for all categories of DI customers studied. The x-axis is the number of premises remaining after the filter application, running from high to low values. The y-axis is the savings multiplier achieved over the average savings across all DI participants. Ideal filters will maximally concentrate average savings while eliminating as few customers as possible.

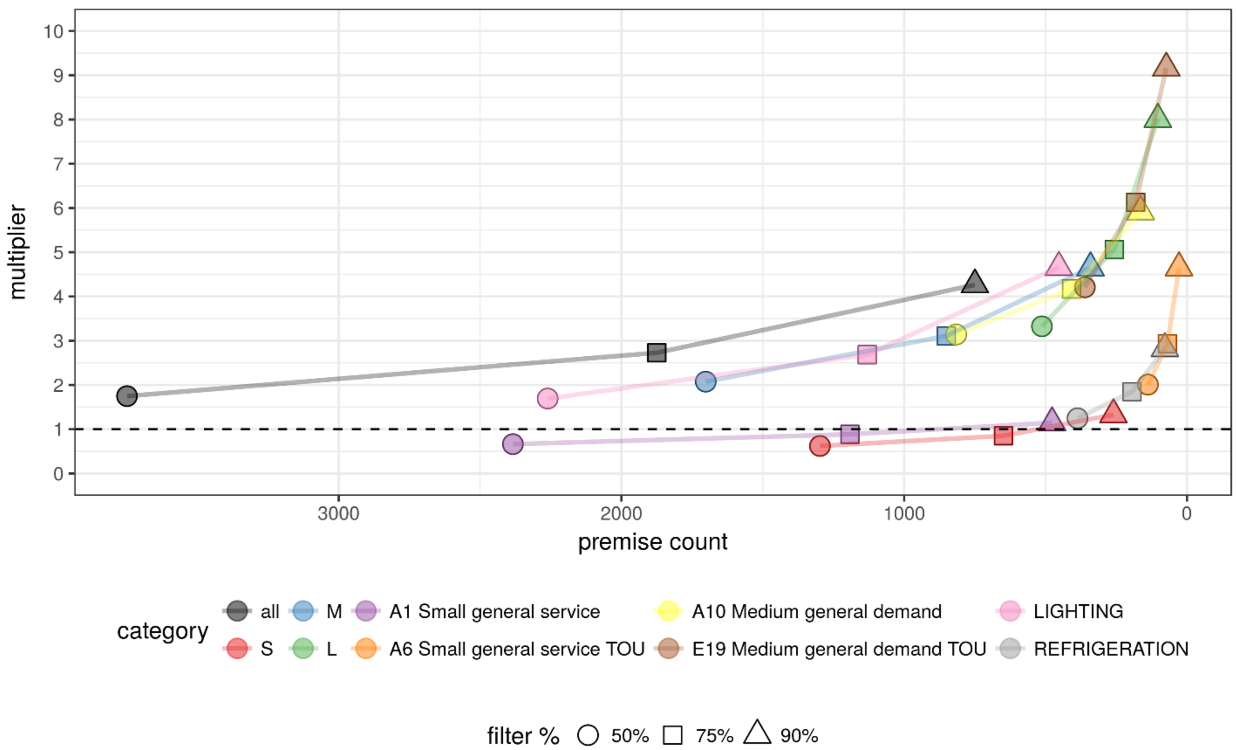

This figure shows that outcomes that double average customer savings can be achieved with targeting logic that eliminates just over 50% of participants and can triple savings by eliminating 75% of participants. The large savings multipliers (over 4x of the unfiltered average), come at 90% filter depth or within specific customer categories, particularly size L or E19 rates (medium general demand with TOU pricing). For programs that can afford to focus their offerings so narrowly, the gains can be considerable (4-9x unfiltered savings), but with such low premise counts, the group average savings can be significantly impacted by outliers. Overly aggressive filtering will tend to produce results here that over-estimate the savings gains achievable outside this group of customers.

Figure 12 illustrates the relationship between filter depth and savings multipliers for the HVAC program. The unfiltered average savings from the HVAC program evaluated to just over 1 kWh/day. This modest starting place is the reason the multipliers are so large.

Figure 12: HVAC program best filter performance at filter depths of 25%, 50%, 75% and 90%. The x-axis is the descending premise could for the filtered groups. The y-axis is the multiplier over the unfiltered average savings that each filtered group achieves.

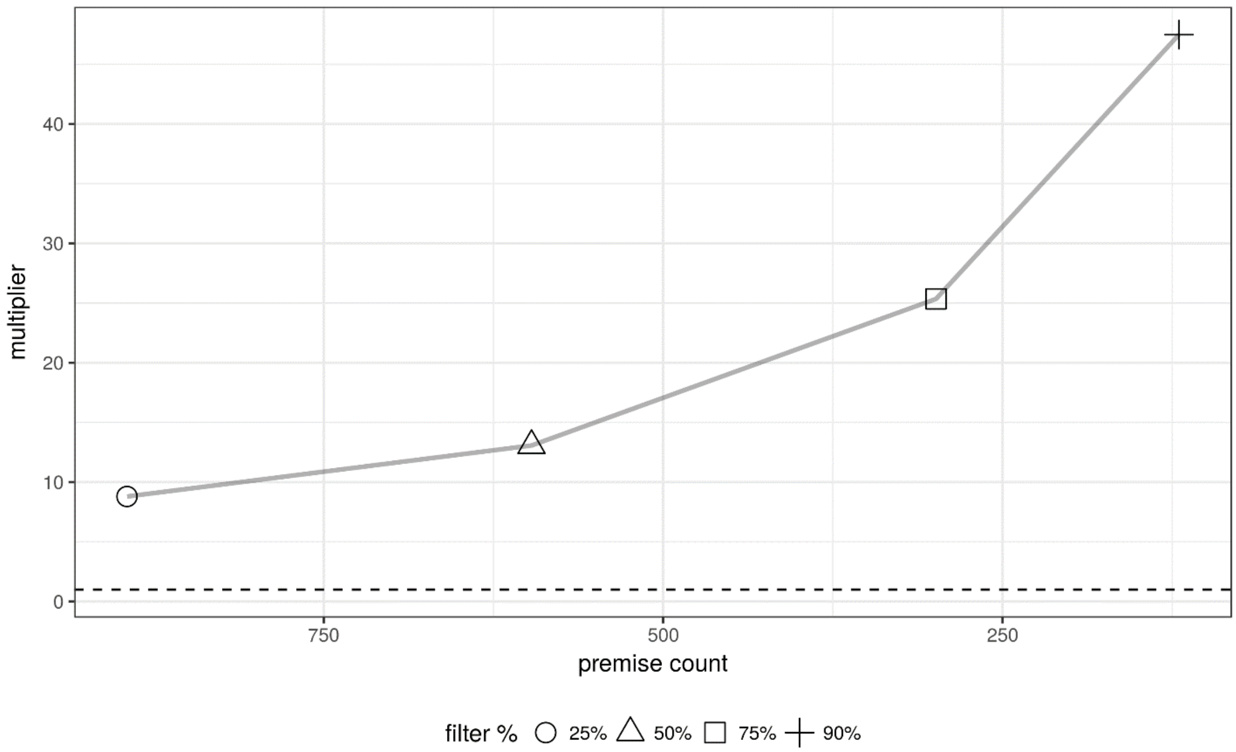

7.4.2 Magnitude vs. depth for feature families

Figure 13 shows the evolution of DI program savings magnitude (x-axis) vs. savings depth - i.e. as a % of pre-intervention consumption - (y-axis) for all features across all filter depths, colored by feature family. At 0% depth, filters eliminate no one and all sub-groups match the full population averages. As the filter depth grows more stringent, DI savings are amplified most reliably by the baseload and total consumption families. In contrast, the thermal and variability families scatter, improving savings magnitude and/or depth in some cases, performing worse than no filter at all in other cases. These results support the understanding that the DI program addresses end uses that are part of the baseload or corrlate strongly with total consumption. Features from the same families should perform well in other programs that address end-uses that are similarly corrleated with total/baseload consumption.

Figure 13: Evolution of DI savings (x-axis) vs. savings as % of average pre-intervention consumption (y-axis) across all filter depths (panel titles) for all feature filters, color coded by feature family. Dashed lines are the unfiltered average values for each axis and each panel represents a different filter depth.

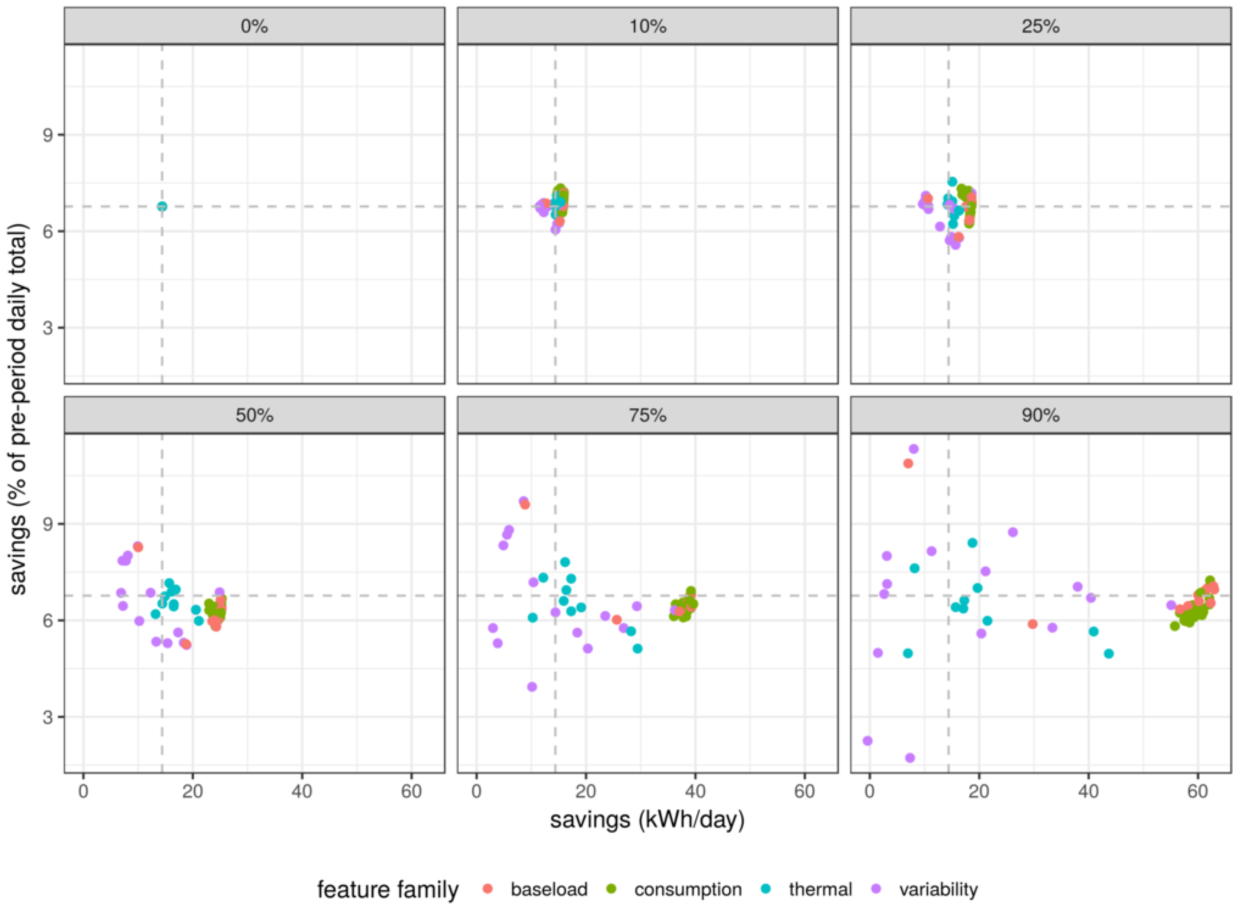

Figure 14 shows the evolution of HVAC program savings magnitude (x-axis) vs. depth (y-axis) for all features across all filter depths, colored by feature group. As the filter depth grows more stringent, HVAC savings magnitude and depth are amplified by baseload and total consumption families, but not as consistently as DI savings were. Baseload/total features that impact total savings also tend to increase % savings. These features are more differentiated, resulting in a stretched out diagonal line of points at higher filter depths. Similar to DI, several variability features improve the savings depth, while others in the same family do not. The most notable difference between the DI and HVAC versions of this plot is the performance of the thermal features, several of which deliver far and away the biggest gains for both savings magnitude and depth and validate the hypothesis that different feature perfom best in predicting the outcomes of different programs. Features from the thermal family should perform well in other AC and heating programs, but the specific features that do best are likely to be adapted to the specific types of interventions offered.

Figure 14: Evolution of HVAC savings (x-axis) vs. savings as % of average pre-intervention consumption (y-axis) across all filter depths (panel titles) for all feature filters, color coded by feature category. Dashed lines are the unfiltered average values for each axis and each panel represents a different filter depth.

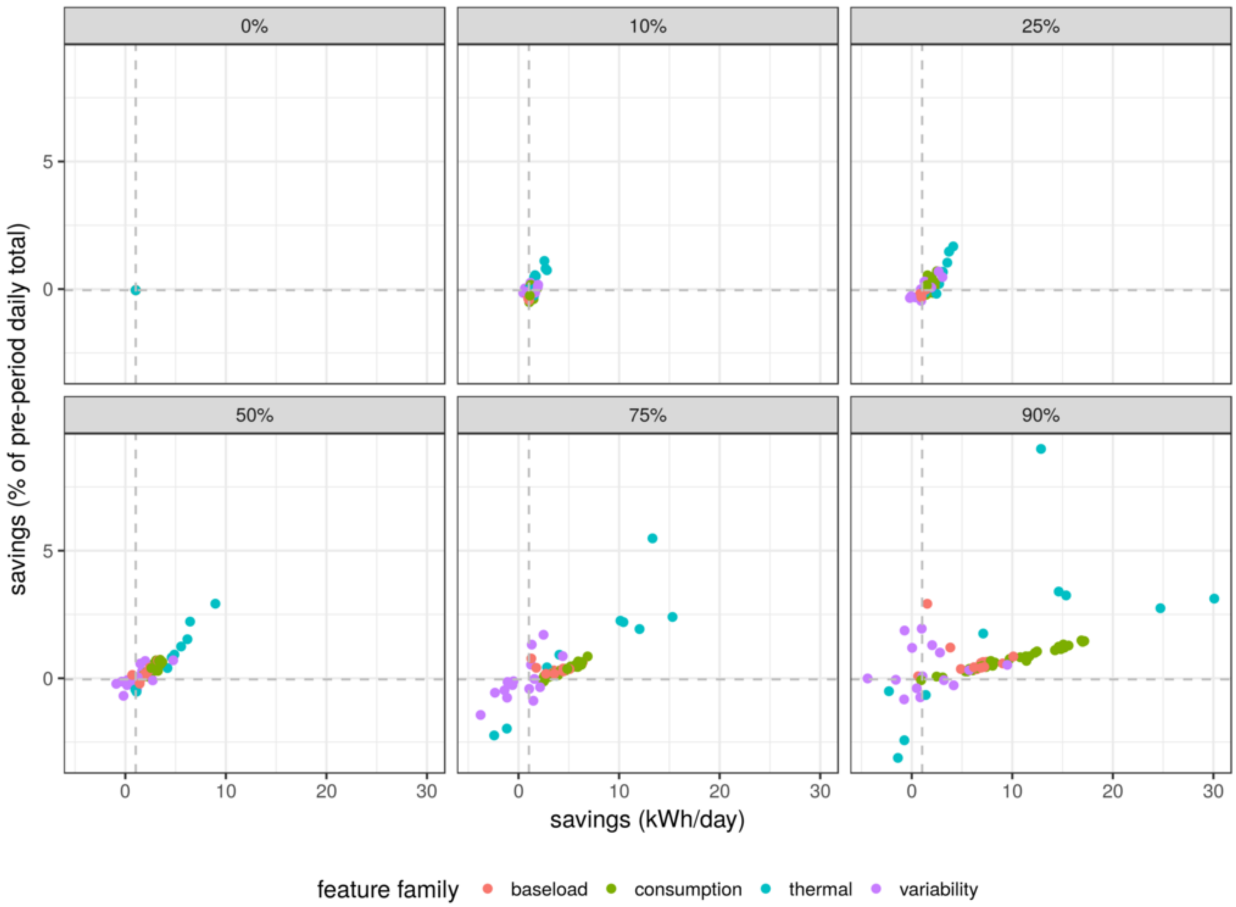

Figure 15 illustrates the progression of magnitude (x-axis) and depth (y-axis) of savings for several key DI program filters as they progress from low to high stringency, connected as paths. In this view, one can see the relationship, or lack thereof between filter outcomes in magnitude and depth of savings. For DI, the filters that improve the magnitude of savings (baseload, average demand, and discretionary load in particular) do little to change the depth of savings. This implies that the corresponding customer premises are larger than the program-wide average, but that interventions performed there do not go any deeper than in the unfiltered population of participants. The % of loads that turn on and off each day, known as the discretionary %, on the other hand, handily increases the depth of savings, but at the expense of magnitude of savings. This filter demonstrates that it is possible to amplify the depth of savings with a filter, but that greater depth of savings for the DI program comes mostly from smaller premises.

Figure 15: Magnitude (x-axis) vs. depth (y-axis) for key DI program-relevant features.

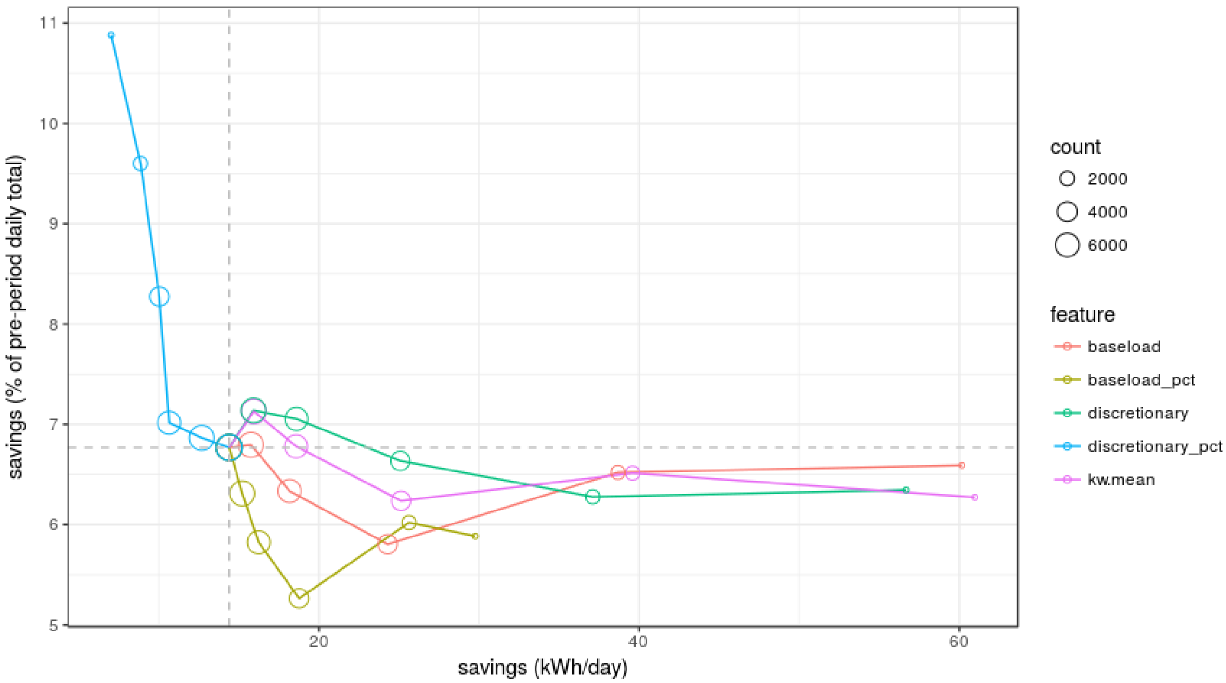

Figure 16 is the HVAC program version of Figure 15. It illustrates the progression of magnitude (x-axis) and depth (y-axis) of savings for several key HVAC program filters as they progress from low to high stringency, connected as paths. Unlike the DI program, the best performing filters in terms of magnitude also increase the depth of savings and the best performing filter in terms of depth of savings also improves the magnitude. These results suggest that larger HVAC loads tend to represent a larger fraction of their site total consumption. Put another way, HVAC loads tend to vary more independently of total load than lighting and refrigeration loads addressed by DI.

Figure 16: Magnitude (x-axis) vs. depth (y-axis) for key HVAC program-relevant features.

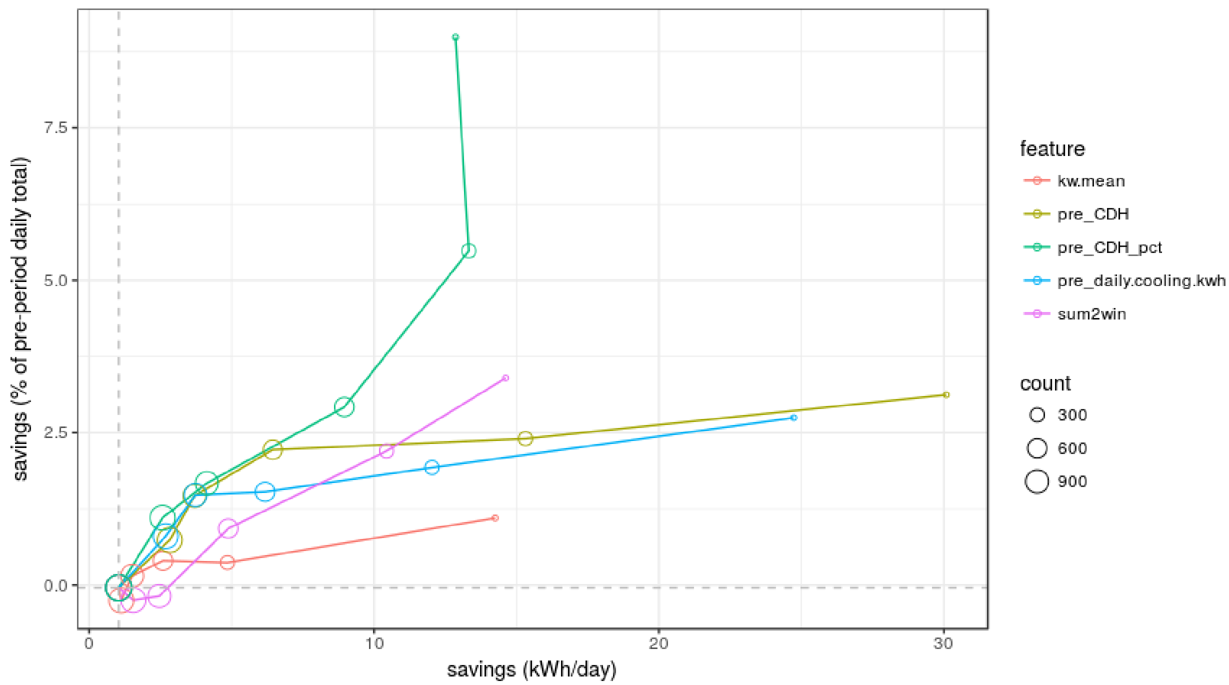

Figure 17 presents savings results for the DI program grouped by feature group and presented as box and whisker plots of the gain % across all feature families, with separate panels for 25%, 50%, and 75% filter depths. The center line of the box corresponds to the 50th percentile value associated with the groups of filters represented therein. The lower and upper bounds correspond to the 25th and 75th percentile, respectively and the whiskers beyond them extend 1.5 inter-quartile distances above and below the boxes, with dots for values beyond 1.5 inter-quartile distances. The baseload and total consumption features are the best choices, with gains approaching 175% at the deepest filter depth, with remarkably tight ranges of outcomes across all features in those categories. The retroactive targeting gains are robust across a wide range of consumption metrics.

Figure 17: DI Savings gain % by feature filter group type for filters eliminating 25%, 50%, and 75% of customers. The center line of the box corresponds to the 50th percentile value. The lower and upper bounds correspond to the 25th and 75th percentile, respectively and the whiskers (extending up to 1.5x the interquartile distance) and dots beyond them correspond to the lower and upper quartiles of performance.

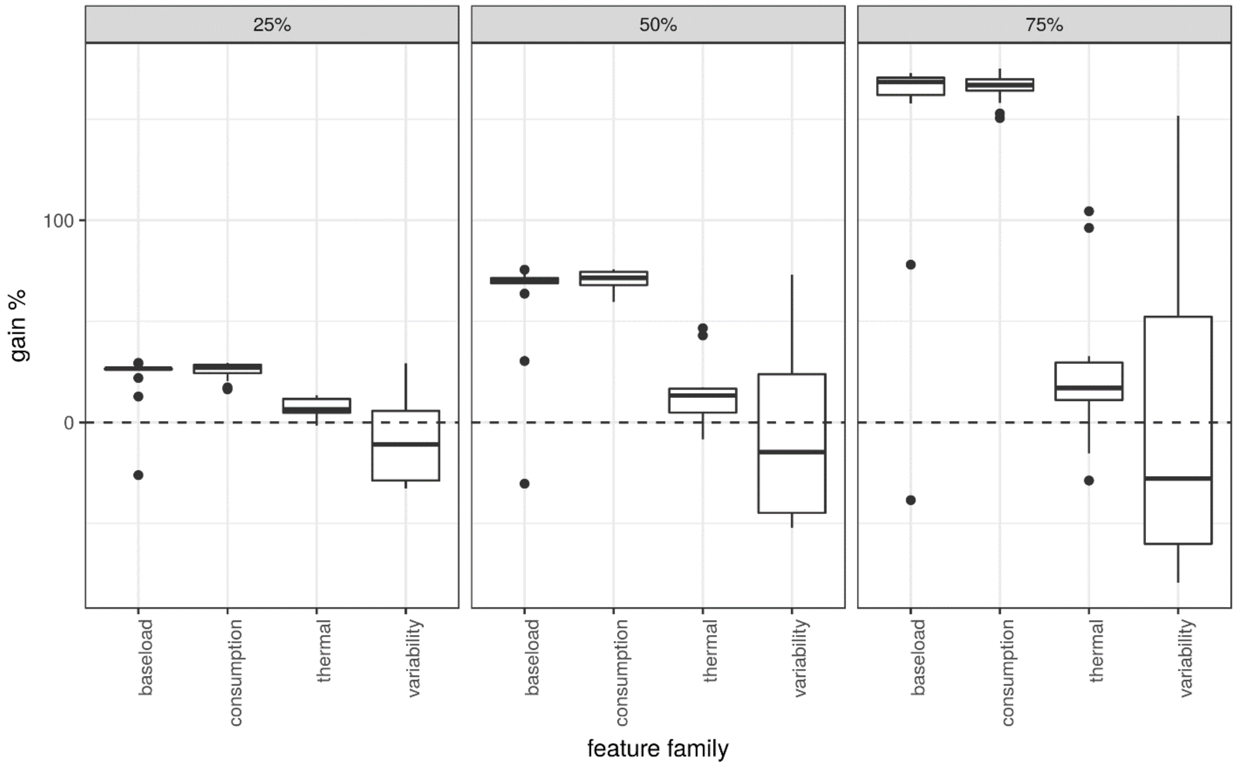

Figure 18 presents HVAC program savings associated with the feature families. The "thermal" feature group includes the best performers, but that group is more hit and miss with its performance. What is clear is that features from different groups are performing relatively better and worse for different programs. The drivers of DI end-uses (primarily lighting and refrigeration) are largely different from the drivers of HVAC end-uses.

Figure 18: HVAC program savings by feature filter group with 25%, 50%, and 75% of customers filtered out. The center line of the box corresponds to the 50th percentile value. The lower and upper bounds correspond to the 25th and 75th percentile, respectively and the whiskers (up to 1.5x the interquartile range) and dots beyond them correspond to the lower and upper quartiles of performance.

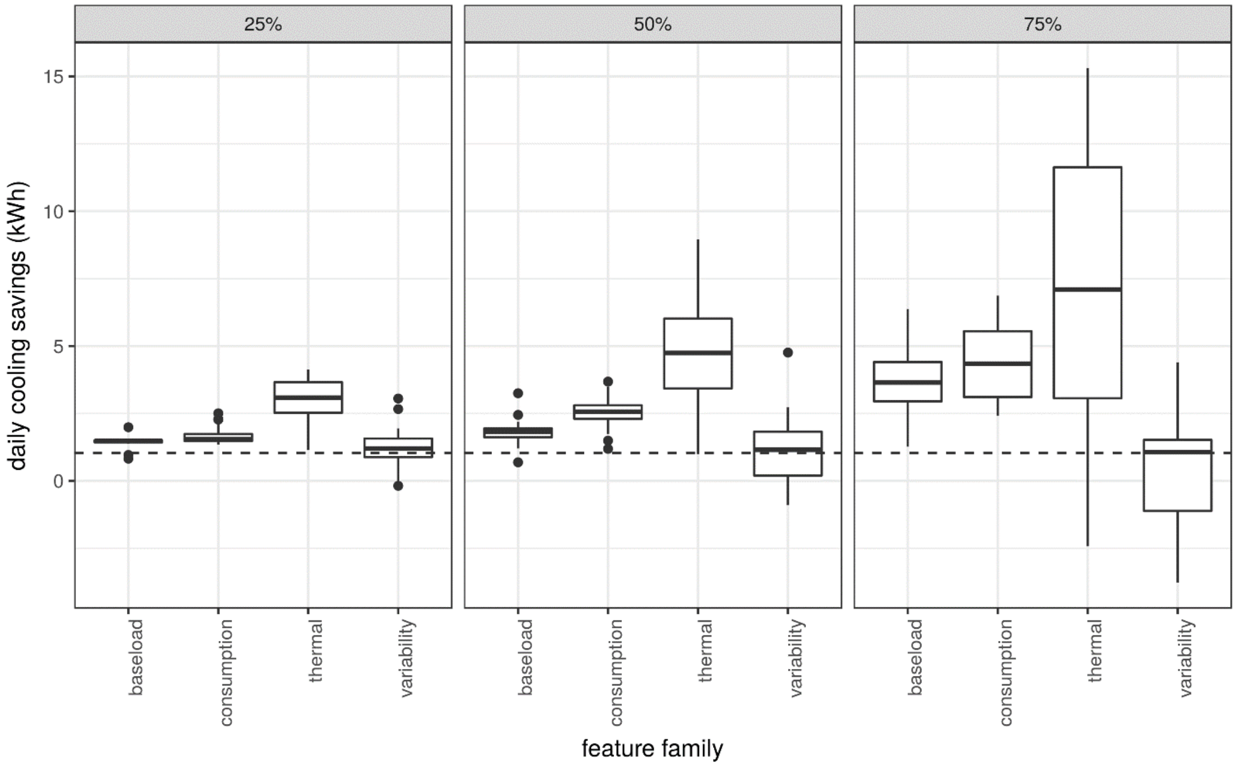

7.4.3 Customer characteristic sub-group performance

Figure 19 shows the average savings (y-axis) of every DI customer characteristics sub-group (circles whose size corresponds to the count of premises in the group) grouped by characteristic type across the x-axis. Note that "year" corresponds to the program year and "more than one" indicates that customers in the sub-group in question participated in more than one program year. Some sub-groups are defined using size or rate types in addition to a primary characteristic type. In these cases, the outcomes are plotted in the column associated with their primary characteristics. Table 9 below the figure provides details on the top 10 saving groups depicted in the figure.

- The top performing sub-groups typically have small premise counts, indicating a possible outsized impact due to outliers.
- The highest savings come from the year and NAICS categories when they are crossed with the S, M, L customer size categories.
- The highest performing groups with largest premise counts come from the customer size and rate categories, indicating these categories are likely reliable for future targeting strategies.
- It can be confirmed that size L customers, Medium and TOU rates, and retail stores are all correlated with elevated savings delivered to the meter.

Figure 19: DI program sub-group average savings (y-axis) by customer characteristic type (across x-axis), sized by the number of premises in each group.

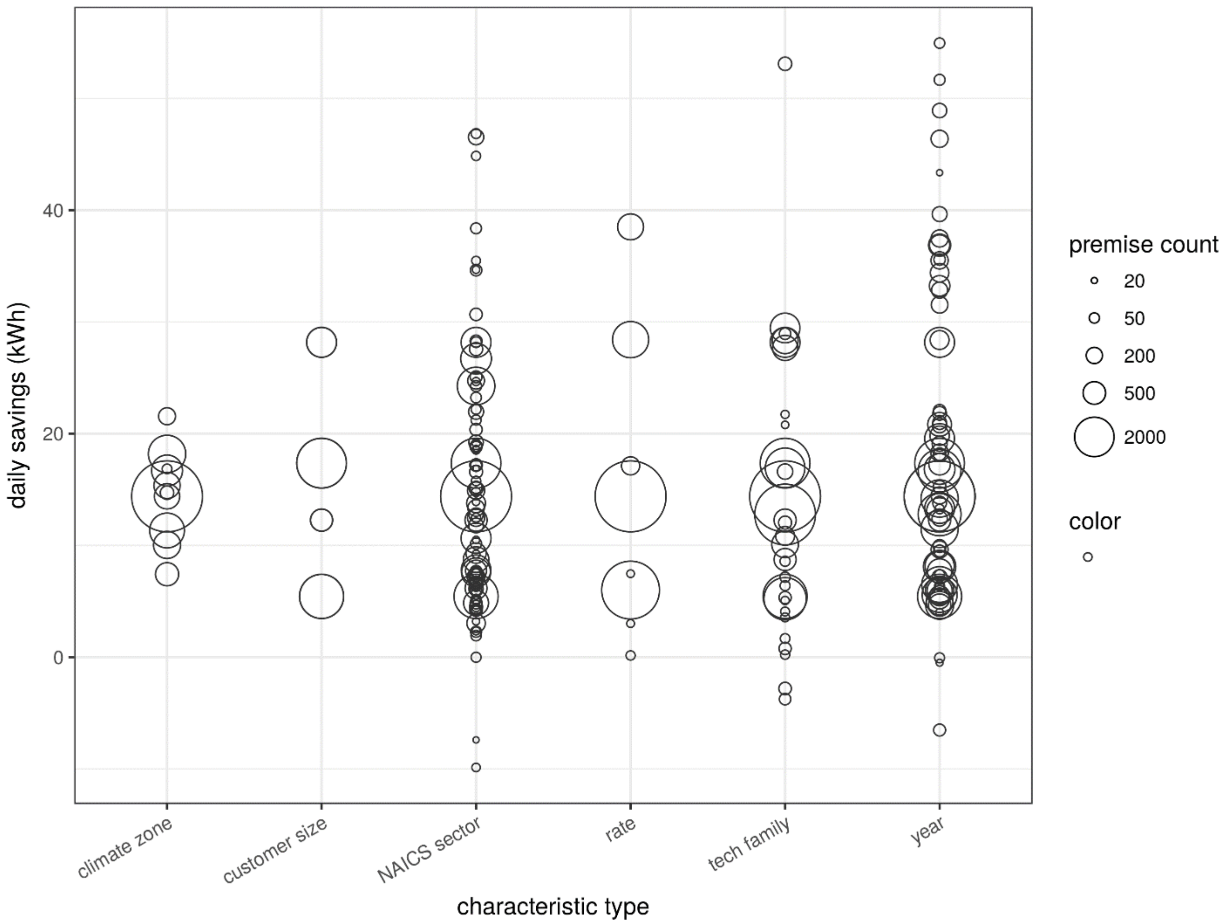

Table 9: Top 10 sub-groups, by savings for the plot above, where the labels describe the categories that define the sub-groups and n is the number of premises in the sub-group.

daily savings (kWh)	label
54.95	year=more than one; customer size=L; n=51
53.1	tech family=LIGHTING and REFRIGERATION; customer size=L; n=112
51.66	year=more than one; rate=E19 Medium general demand TOU; n=60
48.91	year=more than one; rate=A10 Medium general demand; n=138
46.86	NAICS sector=Real Estate and Rental and Lea; customer size=L; n=45
46.54	NAICS sector=RETAIL TRADE - 1; customer size=L; n=169
46.38	year=2015; rate=E19 Medium general demand TOU; n=232
44.85	NAICS sector=Uncategorized; customer size=M; n=36
43.35	year=more than one; customer size=N; n=20
39.65	year=more than one; tech family=LIGHTING and REFRIGERATION; n=155

Figure 20 depicts the *depth of saivngs, or "% of pre"* (y-axis) of every DI sub-group (circles whose size corresponds to the number of premises in the group) defined by all customer characteristics (Across the x-axis). Table 10 under the figure provides more details on the top 10 sub-groups depicted. For % of pre savings, NAICS sector includes the highest performing sub-categories, in other words % of pre savings are concentrated in specific program types and areas of focus. Unlike the absolute savings results, normalized savings accrue to size S customers, and especially to retailers. These results point toward deeper % savings being achievable for smaller customers with consumption dominated by a program-eligible end-use. The electricity consumption of retail businesses, for example, is often driven by lighting, a DI program specialty.

Figure 20: Sub-group average savings (y-axis) by customer characteristic type (across x-axis), sized by the number of premises in each group.

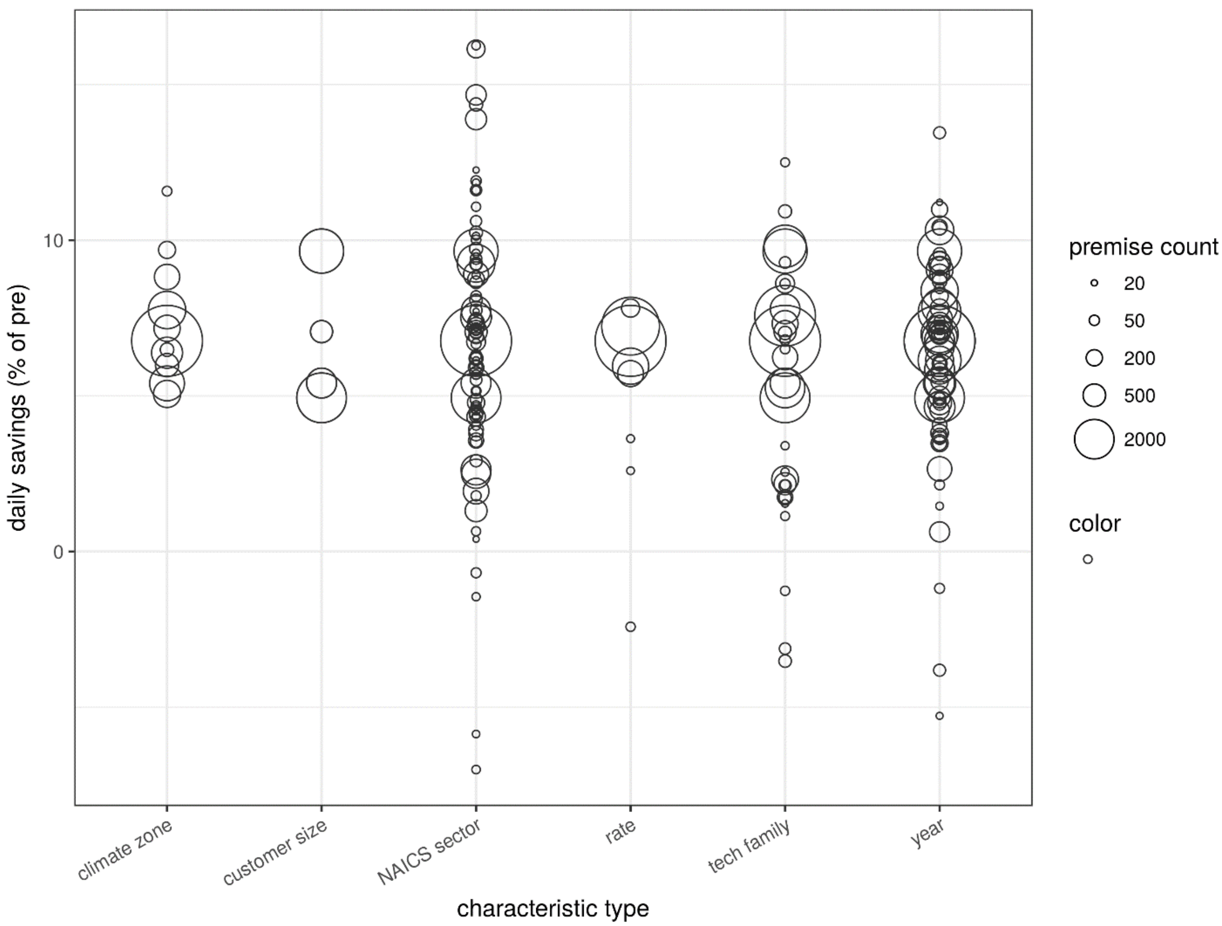

Table 10: Top 10 groups, by savings for the plot above, where the labels describe the categories that define the groups and n is the number of premises in the group.

daily savings (% of pre)	label
16.25	NAICS sector=RETAIL TRADE - 2; customer size=N; n=31
16.14	NAICS sector=RETAIL TRADE - 2; customer size=S; n=261
14.67	NAICS sector=RETAIL TRADE - 2; n=375
14.36	NAICS sector=RETAIL TRADE - 1; customer size=N; n=109
13.88	NAICS sector=RETAIL TRADE - 1; customer size=S; n=412
13.45	year=more than one; customer size=S; n=78
12.5	tech family=ELECTRONICS AND IT and LIGHTING; customer size=S; n=33
12.25	NAICS sector=Management of Companies and En; customer size=L; n=20
11.9	NAICS sector=Administrative and Support and; customer size=S; n=51
11.85	NAICS sector=TRANSPORTATION; customer size=M; n=23

Figure 21 shows the average savings (y-axis) of every HVAC sub-group (circles whose size corresponds to the number of premises in the group) defined by all customer characteristics (Across the x-axis). Some sub-groups are defined using size in addition to a primary characteristic type. In these cases, the outcomes are plotted in the column associated with their primary characteristics. Table 11 below it provides more details on the top 10 sub-groups depicted in the figure. The customer characteristics with the greatest variability in outcomes, including the highest savings, are NAICS sector and program year. Finding the right type of customer and getting the intervention implementation details right appear to be crucial differentiators for HVAC program savings. Accommodation and food service (note that they tend to have large ventilation/HVAC requirements), retail, and public administration are all higher than average savers.

Figure 21: HVAC program sub-group average savings (y-axis) by customer characteristic type (across x-axis), sized by the number of premises in each group.

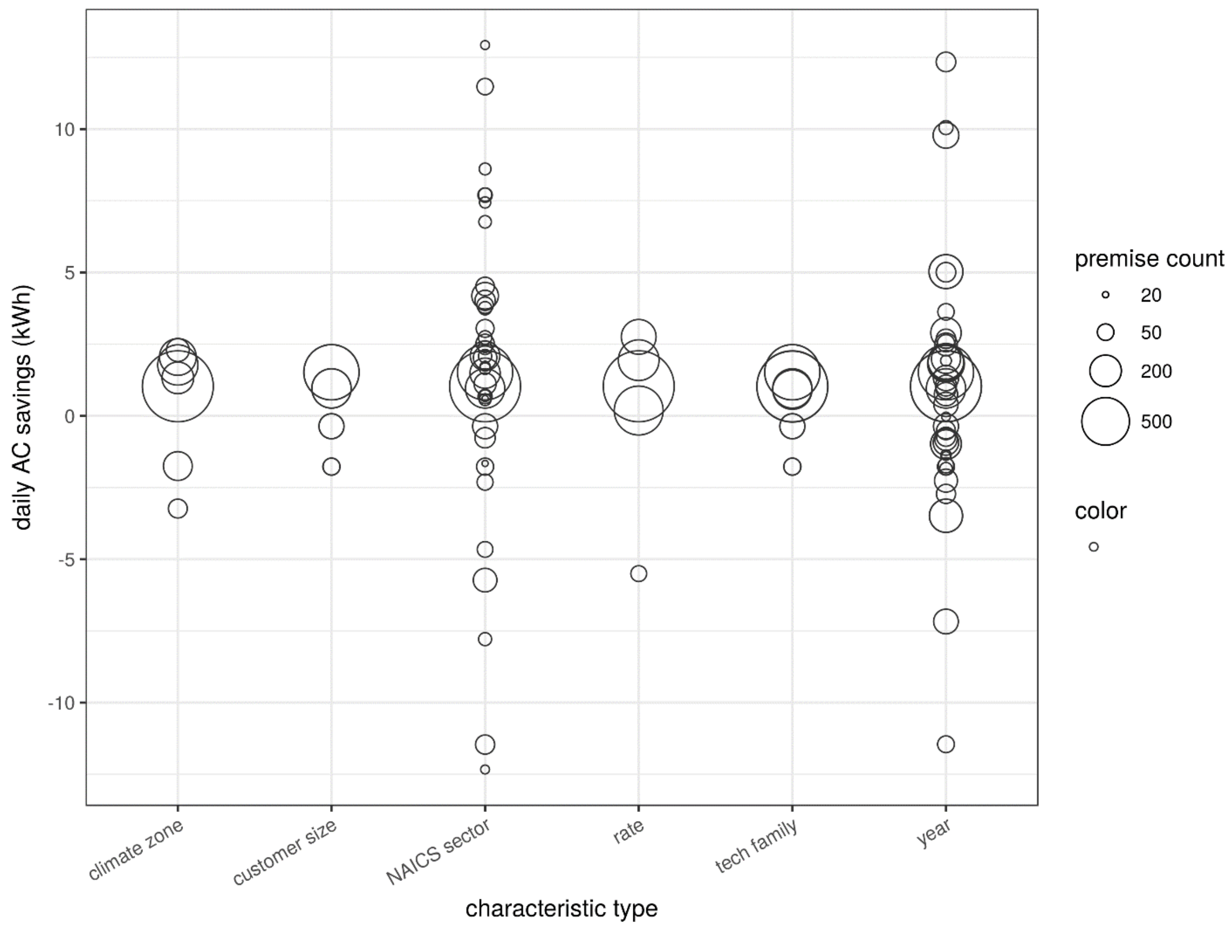

Table 11: Top 10 groups, by savings for the plot above, where the labels describe the categories that define the groups and n is the number of premises in the group.

daily AC savings (kWh)	Label
12.93	NAICS sector=Accommodation and Food Service; customer size=M; n=22
12.35	year=2015; rate=A10 Medium general demand; n=70
11.49	NAICS sector=Accommodation and Food Service; n=49
10.06	year=2015; rate=E19 Medium general demand TOU; n=37
9.79	year=2015; customer size=L; n=127
8.61	NAICS sector=RETAIL TRADE - 2; customer size=L; n=29
7.72	NAICS sector=Public Administration; n=34
7.7	NAICS sector=RETAIL TRADE - 2; n=39
7.44	NAICS sector=Public Administration; customer size=L; n=27
6.77	NAICS sector=Arts, Entertainment, and Recr; n=31

7.4.4 Customer category vs. feature filter performance

Figure 22 depicts DI program savings results (y-axis) for all customer categories (x-axis) in black (as seen in Figure 20), and savings results for all feature filters (at 10%, 25%, 50%, 75%, and 90% filter depths), grouped by feature group, in blue. Circle sizes correspond to the number of premises in each sub-group. Table 12 beneath the figure provides details on some of the top performing sub-groups. As previously discussed, the consumption and baseload feature groups are the best at improving sub-group savings. **While there are category-based groups that perform as well or better than the better usage-based filters, the count of customers in those sub-groups is quite low** – typically well under 100 customers, raising concerns about the influence of a small number of outliers on the sub-group average performance as well as the scalability of recruitment strategies that rely on small minorities of customers. By design, the feature filter groups have 750 (at 90% filter depth) or more (at less aggressive filter depths) customers. **The fact that such large sub-groups achieve average savings comparable to or better than nearly all customer-category-defined sub-groups is affirmation of the value of individual consumption data in identifying potential savers** and suggests that even targeting strategies based on aggressive feature filters have the potential to recruit enough program participants to keep implementation pipelines full.

Figure 22: DI customer (sub)category savings (black) compared to feature filter savings (blue) for all examined customer categories and features. Circle size corresponds to the number of premises in each group. The filters range across 10%, 25%, 50%, 75%, and 90% filter depth.

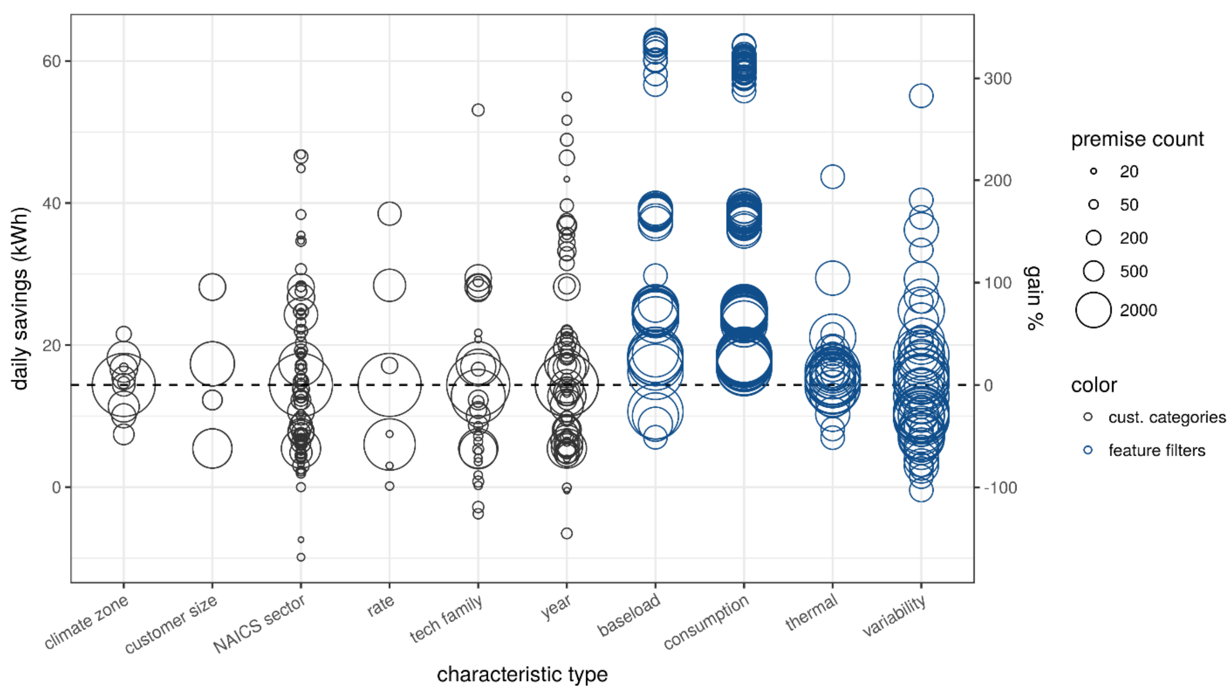

Table 12: Sample of 6 top DI customer category and feature filter savings results.

	daily savings (kWh)	label
Category filters	54.95	year=more than one customer size=L n=51
	53.1	tech family=LIGHTING and REFRIGERATION customer size=L n=112
	51.66	year=more than one rate=E19 Medium general demand TOU n=60
Feature filters	62.48	feature: 1st hour of day mean demand (90%) n=750
	62.31	feature: pre-period modeled non-HVAC weekday daily kWh (90%) n=750
	62.06	feature: mean demand (winter) (90%) n=750

Figure 23 depicts HVAC program savings results (y-axis) for all customer categories (x-axis) in black (as seen in Figure 21), and savings results for all feature filters (at 10%, 25%, 50%, 75%, and 90% filter depths), grouped by feature group, in blue. Circle sizes correspond to the number of premises in each sub-group. Table 13 beneath the figure provides details on some of the top performing sub-groups. As previously discussed, the “thermal” feature group, with features focused on isolating temperature correlation and AC loads are the best at improving HVAC sub-group savings, with just a handful performing particularly well. While there are category-based groups, particularly specific NAICS sectors, that perform as well or better than the best feature filters, the count of customers in those sub-groups is quite low – typically well under 100 customers, raising concerns about the influence of one or two outliers on the sub-group average performance and also the scalability of recruitment strategies that rely on small minorities of customers. By design, the feature filter groups have 120 (at 90% filter depth) or more (at less aggressive filter depths) customers.

Figure 23: HVAC customer (sub)category savings (black) compared to feature filter savings (blue) for all examined customer categories and features. Circle size corresponds to the number of premises in each group. The filters range across 10%, 25%, 50%, 75%, and 90% filter depth.

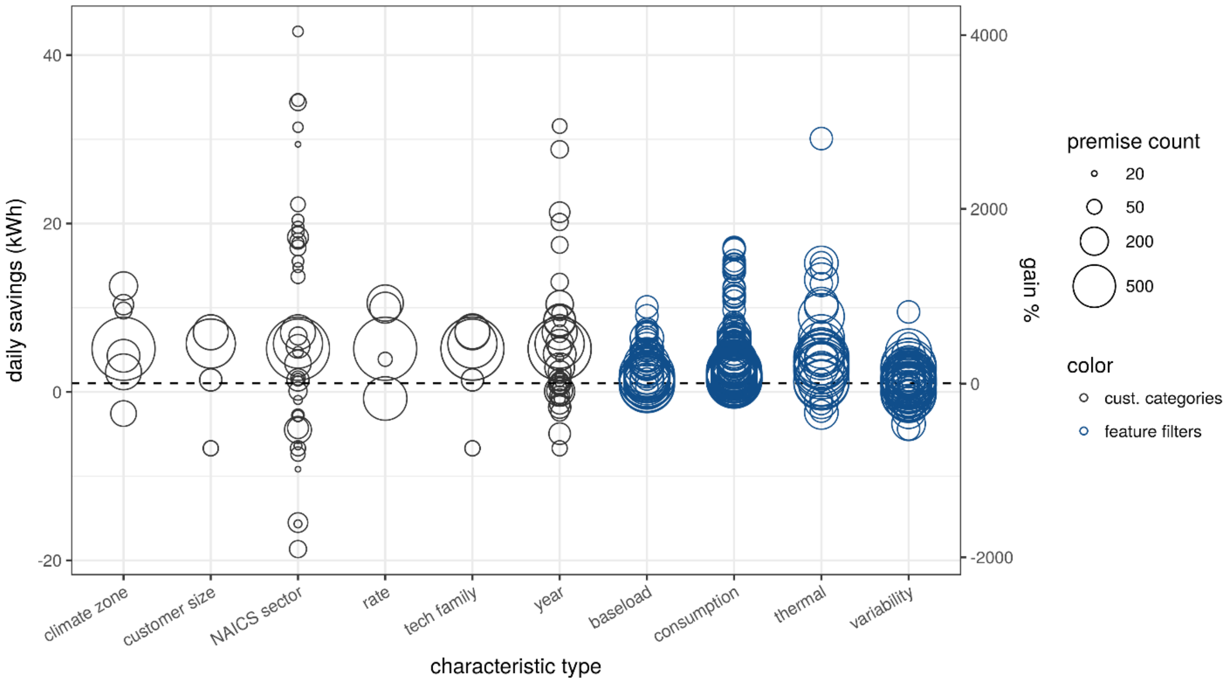

Table 13: Sample of top HVAC savings results for sub-groups defined by customer categories (rows 1-6) and feature filters (rows 7-12).

	daily savings (kWh)	label
Category filters	42.80	NAICS sector=RETAIL TRADE - 2 customer size=L n=29
	34.67	NAICS sector=RETAIL TRADE - 2 n=39
	34.38	NAICS sector=N/A customer size=L n=62
	31.57	year=2014 rate=E19 Medium general demand TOU n=49
	31.43	NAICS sector=Professional, Scientific, and … n=28
	29.40	NAICS sector=MANUFACTURING - 4 n=20
Feature filters	30.09	feature: pre-period modeled temperature sensitivity (kWh/day/degree F) (90%) n=120
	24.73	feature: pre-period modeled daily cooling load (kWh/day) (90%) n=120
	17.16	feature: Aug. total electric energy (90%) n=120
	16.88	feature: total kWh consumed during summer months (90%) n=120
	15.35	feature: correlation between electric consumption and outside temperature (90%) n=120
	15.31	feature: pre-period modeled temperature sensitivity (kWh/day/degree F) (75%) n=299

7.4.5 DI all filter comparison

A final category that is relevant to the DI program is the end-use affected. DI interventions are primarily lighting and refrigeration and results can be examined independently for each. Figure 24 depicts feature filter results, by feature group for all of DI (black), DI lighting (blue), and DI refrigeration (green). Table 14, which follows the figure, provides details for a sample of the top-performing sub-groups in each of the DI technology families.

The savings from the lighting part of DI are higher on average than the savings from refrigeration. The top performing features for lighting and refrigeration are generally from the consumption and baseload categories, but the specific features that perform best are different for each.

Figure 24: DI program feature filter performance (y-axis) by feature group (x-axis) for the entire program (black), just lighting (blue), and just refrigeration (green). Circles are sized by the premise count of the sub-categories.

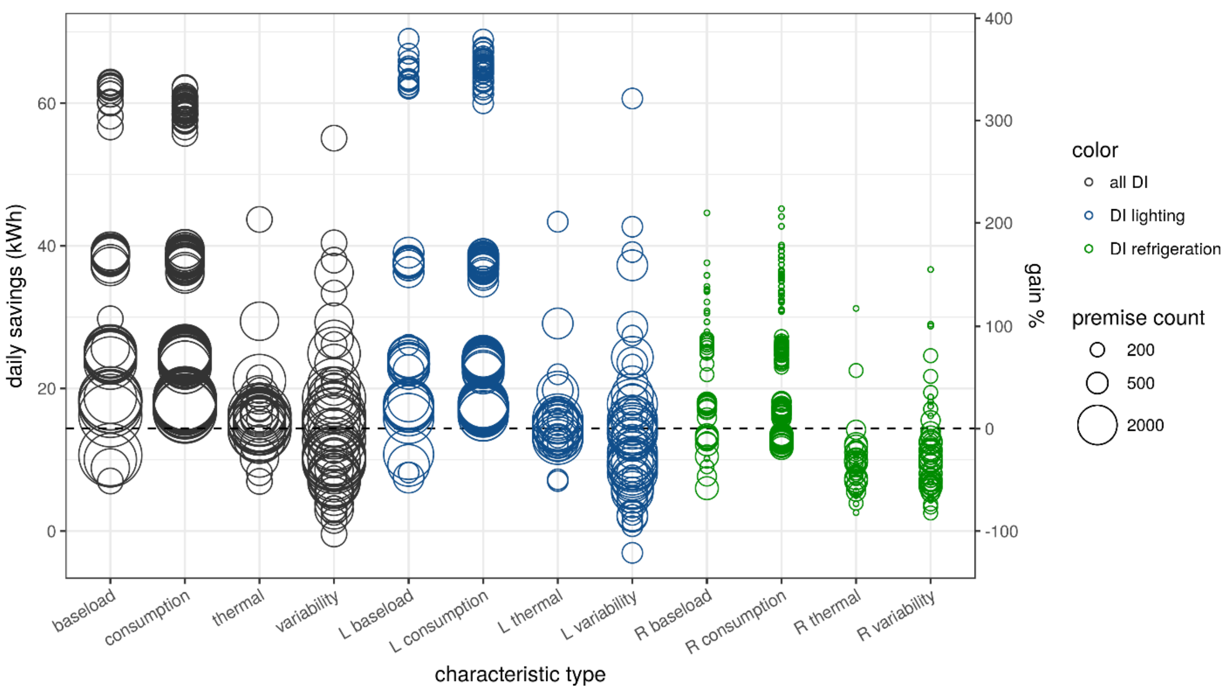

Table 14: Sample of top DI savings results for feature filters for all of DI, just DI lighting, and just DI refrigeration.

	daily savings (kWh)	label
DI all	62.48	feature: 1st hour of day mean demand (90%) n=750
	62.31	feature: pre-period modeled non-HVAC weekday daily kWh (90%) n=750
	62.20	feature: 6th hour of day mean demand (90%) n=750
	62.06	feature: mean demand (winter) (90%) n=750
DI lighting	69.04	feature: pre-period modeled non-HVAC weekday daily kWh (90%) n=453
	68.94	feature: mean demand (winter) (90%) n=453

	daily savings (kWh)	label
	67.83	feature: 19th hour of day mean demand (90%) n=453
	67.20	feature: pre-period modeled daily total load (kWh/day) (90%) n=453
	67.20	feature: mean of daily mean demand (90%) n=453
	67.20	feature: mean demand (all obs) (90%) n=453
DI refrigeration	45.21	feature: 22nd hour of day mean demand (90%) n=78
	40.11	feature: Nov. total electric energy (90%) n=78
	39.29	feature: Feb. total electric energy (90%) n=78
	38.59	feature: mean of daily max demand (90%) n=78

8 Discussion

8.1 Different ways of evaluating savings

Savings are achieved one project at a time, and they can be tabulated in different ways, depending on the context. There are three primary tabulations of program savings relevant to discussion of targeting and program performance.

8.1.1 Average per-customer savings (magnitude)

These are reported in energy units and are useful when contemplating the energy benefits of data-driven targeting and their customer-level impacts. These are particularly relevant because typical program benefits are computed, and rewarded, in energy terms. This is especially true in the context of pay for performance programs.

8.1.2 Percentage per-customer savings (depth)

These are the percentage of pre-program energy consumption estimated to be saved *per-customer*. This metric helps to get at the question of how significant program impacts are on a relative basis per site. A smaller customer will need to work harder to achieve similar energy savings as a much larger customer, so the savings as a percentage of original consumption is also a proxy for on-site effort.

8.1.3 Total program savings

Total savings are a multiple of per-customer savings and the number of customers. The true value of many EE programs is best expressed as their total saving. In the context of targeting, where some fraction of potential participants are ignored, it is important to understand at what point the filter depth is interfering with a program's ability to keep its project pipeline full. See "How deep can filters go" for more discussion of this issue.

8.1.4 Savings over time and location

The grid is a dynamic and complex system. The value of load shifting and reduction varies over time and space in ways that shape strategies for grid operations that better accommodate rapidly growing sources of renewable energy and distributed energy resources of all kinds. It is not hard to imagine a time when EE resources will be valued according to the timing and location of their savings. Indeed, this is an active area of policy discussion. The interval meter data used to derive feature filters for this study will also be integral to the analysis that credits the timing and location of EE savings. Moreover, consumption features should be able to predict the timing and location of savings in addition to magnitude and depth. The advent of time and location varying resource valuation in EE programs will make the case for data-driven targeting even stronger.

8.2 How deep can filters go

Policy makers and implementers might be concerned that data-driven targeting filters eliminate half or more of potential program candidates. Is there a point where targeting will run into practical limits? The obvious answer is that yes, this is entirely possible and the solution is to ease filter restrictions to the point where there are enough candidates to run programs at full capacity. However, the kinds of EE programs amenable to at-the-meter evaluation typically touch thousands of customers per year out of a pool of hundreds of thousands or millions. Those programs are constrained by implementation capacity, not by the pool of potential participants.

On the other hand, the more aggressive filtering becomes, the more likely that high or low “savings” outliers caused by non-program changes in consumption will dominate the results. At 50% of customers eliminated, the influence of remaining outliers is doubled. At 90% it is multiplied by 10. For many programs, concerns about the influence of outliers will manifest before there is slack in program capacity.

Here are some real-world numbers: there are approximately 400,000 SMB customers in PG&E’s territory (account + premise) that are eligible to participate in the programs studied.

Regional DI serves approximately 5,000 customers per year. Based on conversations with implementers, there is an approximately 50% conversion rate (higher or lower for particular implementers). Therefore, we can assume that 10,000 customers are approached each year.

For HVAC CQM approximately 2,400 customers participated in 2016 and another 2400 in 2017. Assuming a 25% conversion rate, approximately 9,600 customers need to be approached to get 2,400 participants.

Thus, HVAC requires outreach to just 2.4% and DI requires outreach to 2.5% of all SMB customers a year. We assume there are also practical limitations on outreach that impact who implementers pursue - geography, business type, etc. It is probable that meeting those requirements would require a larger pool of candidates to draw upon. But it looks like programs like the ones examined could be pretty aggressive with their filtering (absent other constraints).

Certainly filters that eliminate 50% of candidates would be a good starting point, but it may be possible to go to 75% or higher in the real world (with the caution that the influence of outliers on estimated savings is expected to be non-negligible at those high filtering levels, leading to higher on-paper mean savings improvements than should be expected in the real world).

8.3 Didn’t we already know that customer size is an important determinant of savings?

Logically, savings from EE interventions will scale with the degree of waste or inefficiency on site prior to their implementation, along with the ability of implementers to identify and address that waste and efficiency.

As a broad rule of thumb, savings have been observed to scale with total energy consumption, but there are many cases of large consumers with efficient and well operated systems and there are many cases of smaller customers who have significant waste and/or inefficiency or with usage dominated by a program-relevant end-use. A more accurate rule of thumb would be that savings tend to scale with the magnitude of the relevant end-use(s).

In reality, beyond rules of thumb, savings scale with the magnitude of correctable waste and inefficiency in the relevant end-uses and their interactions with one another and the premise occupants. But that presents a practical problem: end uses, waste, inefficiency, and occupants are not directly observed. The best resource widely available for estimating these core drivers of efficiency potential is interval meter data and the question of how well that data can support program improvements is an open question – and the subject of this research.

Furthermore, savings are not the only, or even the most important metric of EE program success. Cost effectiveness calculations require understanding the costs associated with implementation in opposition to the savings achieved. Social, policy, or corporate goals will often dictate that a broad mix of customers should be served by programs. Real-world targeting should never be one dimensional and even the most aggressive profit maximizing approach to efficiency would want to target the least expensive savings, not just the largest ones.

8.4 The limitations of NMEC savings calculations

Another potentially problematic aspect of at-the-meter targeting and evaluation is the difficulty in confidently attributing savings to program actions as opposed to other factors. Standard pre/post at-the-meter NMEC "savings" calculations are actually calculations of all changes that occurred between the pre and post periods regardless of cause. As we have seen in our control groups, there can be very broad distributions of "natural variability," presumably caused by a combination of mean-zero fluctuations in consumption from year to year, site-specific changes in energy use intensity (i.e. business expansion or contraction), and long-term trends in consumption (i.e. LED lighting adoption or more efficient computing). On top of those actual fluctuations, NMEC also adds the uncertainties associated with the weather normalization.

Potential sources of NMEC savings bias:

1. In large samples, mean-zero fluctuations and site-specific changes in consumption are often assumed to cancel out across premises (for every site with an increase, there is a corresponding site with a decrease). However, shared factors like droughts, prevailing economic conditions, etc. can cause shifts in consumption that do not cancel out. Further, these exogenous factors can impact certain customer segments more than others.
2. Similarly, a weather normalization model that is overly temperature sensitive or was trained using relatively cool (or hot) weather data, could create systematic biases when trying to normalize consumption for a relatively hot (or cold) year.
3. Trends in energy consumption (i.e. organic LED adoption or plug load growth) can also undermine the assumption that models trained on pre-period data can provide unbiased estimates of the counterfactual conditions for the post-period.

8.4.1 Synthetic controls for NMEC savings

A potential solution to the biases listed above is to identify control groups that experienced the same fluctuations, weather, and energy trends as the participants. To serve as controls, customers would also need to respond to all those conditions in the same manner as participants would have. For example, if participants are all higher than average AC users, their controls would need to be as well.

The best controls come from randomly controlled trials (RCT) that randomly assign participants and controls out of a single, large pool of customers. In those cases, there is no risk that some unobserved characteristic (politics, energy awareness, past participation, etc.) has influenced the enrollment choices of participants. In the real world, there are limited opportunities to run RCTs.

When RCTs are impractical or impossible, there are several promising methods for generating synthetic control groups using methods that match between participants and controls using various metrics of similarity. This matching can be done at any time and can therefore be applied to a program in retrospect. Unfortunately, there is no way to rigorously *prove* that a matched control is going to behave as the participants would have, but there are some practical tests that can be run to evaluate both the matching methods and the groups they select.

With all this as context, it does appear based on this and related research that NMEC savings without control groups runs a significant a risk of counting saving that were not caused by the program interventions or ignoring savings that were. More work should be done on methods for forming synthetic control groups for at-the-meter savings calculations and methods for evaluating the resulting matched groups.

As a corollary, NMEC savings should be expressed with error bars. Those error bars should shrink with larger samples sizes and controls, but will never be zero, and the errors on individual projects will likely be so large as to often (accurately) make individual assessments extremely difficult.

8.4.2 Gaming the system?

Theoretically, the imperfections in NMEC savings estimates could be exploited by an implementer who recognizes a situation where at-the-meter estimates will have a bias from some condition outside the program effects. A standard example is an implementer who only works with households with children about to graduate from high school and move out of the house. This implementer would get NMEC reductions from the occupancy changes for free unless they were compared to control group of other households with children about to graduate.

All of these concerns are reason to proceed with care, but they are not cause for despair! There are certainly technical fixes for a lot of the issues raised. We recommend more research in this area and practical steps taken to ensure best practices are understood and used by evaluators, but more importantly, there are many non-technical fixes as well. Issues like long term trends can be spotted at a high level and corrected. Implementers that game the rules will stand out against a background of honest peers and can be policed accordingly. Effective program administrators and regulators can ensure that such issues are well contained.

The most appropriate metric of success for at-the-meter evaluation will be whether new strategies for delivering savings become viable and mature and scale up because evaluation rules that reward performance (even imperfectly) are better for innovation that rules that ignore performance.

8.5 Additional considerations

Although we assume programs will continue to have open enrollment for eligible customers, targeting logic selects some people for focused recruitment and leaves others out. As a matter of policy, we should be concerned about ensuring that the inclusions and exclusions continue to serve the greater public purpose of efficiency programs. At first blush, one might assume that targeting will result in programs that just focus on the biggest customers, who represent a modest fraction of customers. However, targeting based on individual consumption spans a variety of customers far better than targeting based on broad categories. For example, if your HVAC program targets hot climate zones, you are likely to pick up some serious AC users, but you will also be including some customers without much AC at all and excluding the higher AC users from other climates.

It is important to note that untargeted programs are not evenly applied. Customers typically self-select to enroll, introducing a bias towards the profile of the most motivated participants. To the extent that targeting is merit based and results in contact with atypical program participants, it would naturally improve the diversity of characteristics of participants.

One category of customers that can't be well served by data-driven methods is anyone who just started a new account. Without the data to analyze, data-driven programs could have a blind spot for customers that move frequently. However, many aspects of efficiency are more accurately tied to the premise than the customer, so it may be possible to overcome this blind spot to some extent by looking at the location's consumption history rather than the specific customer.

In the case of targeting based on retrospective analysis, it is also important to recognize cases where the targeting could amplify one or more aspects of current practice while ignoring the potential gains from a strategy that hasn't yet been tried. This means that it is important that data-driven programs are designed to ensure that they meet the varied needs and savings potential of all customers. In that setting, the targeting logic that excludes customers from one program would help to identify the other program(s) that is (are) the best fit for each customer.

Whatever biases and blind spots targeting rules have, the best response is to improve the rules. Targeting can be used to recruit a population of program participants with any number of constraints, including income, business type, etc. and, using the methods of retrospective analysis described in this report, the mix of program participants can be monitored and adjusted as needed with more precision than untargeted programs are currently able to provide.

9 Conclusions and next steps

The research presented in this report has demonstrated that very significant gains in at-the-meter EE program savings per-customer can be achieved using targeting based on pre-program-implementation consumption features and customer characteristics. We estimate that well-executed targeting can improve per-customer average savings by a factor of 2-3x by pre-screening potential participants using data-driven targeting methods described here and focusing recruitment efforts on the most attractive 25-50% of potential customers.

These results are based on practical calculations that can be undertaken by a wide variety of program planners and implementers, provided they have access to meter data, and the lessons already learned can be applied even without the benefit of meter data.

Because this research was performed retrospectively using data from actual programs, we can state with confidence that the targeting gains we have observed are additional to real world best practices already in place.

Because the savings calculations were based on NMEC pre/post comparisons, our methods and findings are particularly relevant to the planning and evaluation of pay-for-performance programs.

Based on the strength of these findings, we recommend follow-up in three areas:

First, every program is different, so **more program data should be retrospectively analyzed** using these or similar methods to expand the body of knowledge on generalizable and program-specific findings on the drivers of program savings. These methods are well suited to **applications in both research and evaluation settings**.

Second, program designers, planners, and implementers should **incorporate data-driven targeting into their program recruitment strategies**. There are many details, potential pitfalls, and synergies that can only be worked out in the field. For example, is data-driven targeting best suited to help improve the performance of existing programs or should programs be designed with individualized information in mind from the ground up? However, based on the savings improvements we've documented, efforts along the lines of this work should out-perform current practice, with improvements to savings and cost effectiveness that make the effort well worth it.

More broadly, the end goal of data-driven programs should be the **tailoring of efficiency services to the specific needs, waste, and inefficiency of individual customers**. Properly done, individualized targeting, diagnosis, education, and support has the potential to significantly increase realized benefits of program interventions, unlock program specialties that are not cost effective when prescriptively applied, and improve customer satisfaction. Further, the insights and implications derived from customer targeting research have the potential to drive the frontier of our collective understanding of how customers use energy and what can be done to make that use more flexible at a time when grid operators are looking to demand side resources for greater flexibility in time and location.

There is valuable information about the drivers of program savings and the wide range of "energy behaviors" of customers locked away in the customer and intervention data of past programs. We now have the tools to unlock that information, with a focus on increasing benefits for both programs and customers.

Appendix A: Definitions

attribution – The process of attributing apparent savings to program effects rather than from unrelated customer activity. In EE programs, attribution is closely associated with distinguishing net from gross savings. Our usage of the term is meant to be broader than the standard set of net to gross concerns.

consumption feature – A metric of consumption computed using meter data that characterize some aspect of consumption and can be used to define filtering criteria. See the table of features below in this appendix.

control group – A group of non-participating customers with characteristics comparable to a set of participants, used to quantify the "natural variability" in savings estimates produced by at-the-meter methods in the absence of program impacts.

customer attributes – Known characteristics of customers that can be used to define sub-groups of customers.

depth of savings – daily savings as a % of pre-period daily total consumption.

filter percentage – The percentage of customers eliminated by a given level of a targeting filter or filters.

filtered savings – The average at-the-meter savings values for a sub-group of customers selected by a filter.

gain – The mean daily kWh saved in a filtered sub-group of participants subtracted from the average program savings across all participants. If a program saved 6 kWh/day on average, but a filtered subset of participants saved 15 kWh/day on average, their gain would be 9 kWh/day.

% gain – The percentage of average savings experienced across all program participants represented by the gain of a filtered sub-group of participants (i.e. gain / all participant average x 100). If an average participantsaved 6 kWh/day on average, but a filtered subset of participants saved an average of 15 kWh/day, their gain would be 9 kWh/day, for a % gain of 150%.

impact estimates – See savings estimates.

magnitude of savings – estimated difference in daily average consumption (kWh) between the weather normalized pre- and post-periods.

natural variability – By definition, non-participants in programs derive zero savings from them, but at-the-meter evaluation methods still measure changes in their energy usage over time. These non-program changes are defined as "natural variability" but could be the product of specific trends in consumption and technology adoption.

NMEC / Normalized Metered Energy Consumption – Meter data modified to control for known sources of variability that are different between the period when the data was gathered and the period to which it should be compared. In the context of at-the-meter savings, pre-period meter data can be normalized against outside temperature, day of week, time of day, and other known determinants of consumption. This most typically involves running a weather normalization regression model, with roots traceable to PRriSM and IPMVP Option 3, on pre-period data and using it to forecast "baseline" loads during the post period.

outliers – The most extreme instances of natural variability, like doubling or halving of pre-program consumption that cannot be explained as program impacts but make an outsized contribution to participant means. Such outliers need to be trimmed from the results to improve the likelihood that observed savings can be attributed to the program.

P4P / Pay for performance – A type of efficiency program that pays implementers based on at-the-meter evaluation of savings, typically based on NMEC.

perfect filter – A perfect filter would systematically eliminate customers from the lowest savings to the highest, in order. It is not possible for a real-world filtering approach based solely on pre-program data to match the gains achievable through perfect filtering, but it can be useful to compare the performance of a given filter to the perfect one.

proportional savings - Savings that are proportional to some property of consumption observed in the controls, for example, proportional to the pre-intervention average daily consumption. These produce a positively skewed savings distribution. Our filtering approach assumes that savings are roughly proportional to some consumption features computable prior to the program implementation.

savings estimates - Savings in kWh/day = (the consumption from the post-period – the NMEC from the pre-period with post-period weather data applied), averaged across all post period days. These estimates can be made for both *total* energy savings and *HVAC*, or *cooling* energy savings.

unfiltered savings – The program wide average at-the-meter savings across all participants.

Feature families and feature definitions

The table below lists all the features used in this analysis and provides a family assignment and simple definition for each.

family	feature	label	units
baseload	Aug_min	Aug. mean of daily min demand	kW
baseload	base_pct	baseload as a percentage of total consumption	%
baseload	discretionary	Non-baseload consumption (kWh/day)	kWh
baseload	discretionary_pct	discretionary load as a percentage of total consumption	%
baseload	HOD_mean_24	24th hour of day mean demand	kW
baseload	HOD_mean_4	4th hour of day mean demand	kW
baseload	min	mean of daily minimum demand	kW
baseload	min_3	3rd percentile electric demand	kW
baseload	pre_intercept	pre-period modeled non-HVAC weekday daily kWh	kWh
consumption	Aug_mean	Aug. mean of daily mean demand	kW
consumption	HOD_mean_12	12th hour of day mean demand	kW
consumption	HOD_mean_16	16th hour of day mean demand	kW
consumption	HOD_mean_20	20th hour of day mean demand	kW
consumption	HOD_mean_8	8th hour of day mean demand	kW
consumption	kw_mean	mean demand (all obs)	kW

family	feature	label	units
consumption	kw_mean_summer	mean demand (summer months)	kW
consumption	kw_mean_winter	mean demand (winter months)	kW
consumption	kw_total_Apr	Apr. total electric energy	kWh
consumption	kw_total_Aug	Aug. total electric energy	kWh
consumption	kw_total_Dec	Dec. total electric energy	kWh
consumption	kw_total_Feb	Feb. total electric energy	kWh
consumption	kw_total_Jan	Jan. total electric energy	kWh
consumption	kw_total_Jul	Jul. total electric energy	kWh
consumption	kw_total_Jun	Jun. total electric energy	kWh
consumption	kw_total_Mar	Mar. total electric energy	kWh
consumption	kw_total_May	May total electric energy	kWh
consumption	kw_total_Nov	Nov. total electric energy	kWh
consumption	kw_total_Oct	Oct. total electric energy	kWh
consumption	kw_total_Sep	Sep. total electric energy	kWh
consumption	max	mean of daily maximum demand	kW
consumption	max_day_kw	total electric energy on max day	kWh
consumption	mean	mean of daily mean demand	kW
consumption	pre_daily_kwh	pre-period modeled daily total load (kWh/day)	kWh
consumption	summer_kwh	total kWh consumed during summer months	kWh
thermal	kw_tout_cor	correlation between electric consumption and outside temperature	cor
thermal	max_day_pct	percentile of temperature for max demand day	%
thermal	max_day_tout	mean temperature of max demand day	F
thermal	max_hr_tout	temperature at max demand	F
thermal	pre_CDH	pre-period modeled temperature sensitivity (kWh/day/degree F)	kWh/day/F
thermal	pre_CDH_pct	pre-period temperature sensitivity as a % of daily kWh (%/degree F)	%
thermal	pre_daily_cooling_kwh	pre-period modeled daily cooling load (kWh/day)	kWh
thermal	sum2win	ratio of total consumption during summer months to winter months	
thermal	summer_tout	Summer mean outside temperature	deg F
thermal	tout	Annual mean outside temperature (F)	deg F
variability	daily_kw_max_var	variance of daily max electricity demand	kW
variability	daily_kw_min_var	variance of daily min electricity demand	kW
variability	daily_kw_var	variance of daily mean electricity demand	kW
variability	kw_var	electric demand variance	kW
variability	kw_var_summer	electric demand variance (summer)	kW
variability	kw_var_winter	electric demand variance (winter)	kW
variability	max_ramp	maximum rate of ramp	kW

family	feature	label	units
variability	mean_ramp	the average load shape ramp rate across early evening hours	kW
variability	mn2mx	ratio of *min* to *max* average daily demand	
variability	morning_ramp	average ramp rate change in consumption between sequential hours from 7am to 10am (kWh/hr)	kW
variability	mx2mn	ratio of *max* to *min* features	
variability	n2d	ratio of night (1-5am) to daytime (3-7pm) consumption	
variability	nv2dv	ratio of night (1-5am) to daytime (3-7pm) variances	
variability	peak_frac	the fraction of all energy consumed during peak period hours	
variability	range	mean of daily range of demand	kW

Appendix B: Methods details

Data cleaning and validation

This section details the steps taken to clean and validate the customer data used for this study.

Isolating occupied buildings

To ensure that our analysis ran against customers with occupied buildings with correctly operating meters, we required that all customers have an average consumption greater than 180W and that no more than 15% of meter readings are missing or zero.

Eliminating Spurious Outliers

The central measure of program savings in our analysis is the average daily kWh savings in the post-intervention period. Averaged values can be highly sensitive to the presence of large outliers in the data set being considered, so it is important to examine any extreme values and potentially reject them from the sample if they do not represent the effect being measured. In this analysis, evaluation of all participant sites yielded a significant proportion with very extreme savings values relative to the mean. Standard naive approaches to trimming outliers (e.g. excluding the highest and lowest N percent of values or excluding outliers beyond a certain number of standard deviations) demonstrated that the average savings were strongly dependent on the exact outlier-rejection strategy used. At the same time, it is also possible that some fraction of the extreme values is derived real savings (e.g., for very large sites that were operating very inefficiently prior to program intervention), and it would be preferable to retain such sites in our analysis. Therefore, it is desirable to use other observed characteristics of each site, besides the measured savings, to determine which sites should be excluded as potentially spurious outliers. We developed a two-step approach for this analysis, as detailed in the following sections.

Step 1: Requiring sufficient time-series data for evaluation

Because the measured savings depends on a model of site-level energy consumption fitted to meter data from the pre- and post-intervention period, sites having limited pre-period and post-period data are likely to yield badly constrained models and, thus, spurious savings estimates. We can explore this potential source of error by looking at the CDH model parameter, which represents the scaling of cooling load with daily cooling degree hours, for each site in the pre- and post-intervention period. If both the pre-period and the post-period model have sufficient data to constrain the response of cooling load to weather, we would expect these two coefficients to be strongly correlated with one another, whereas they would not be expected to have any relation to one another if one of the periods has insufficient data to constrain the model.

The left panel of Figure 25 shows a scatter plot of these two parameters for each customer evaluated, with color-coding indicating the minimum number of pre-period or post-period days for each meter-data time series. For sites with less than 120 days' worth of hourly meter data in either the pre-period or the post-period, there is little to no correlation between the cooling coefficients for the two periods. Above this threshold, the correlation is visually clear, with some outliers remaining, and the relation tightens as we increase the pre/post-period threshold. In the right panel, we see the fraction of the sample having a pre-period or post-period shorter than a given threshold level. There is a clear tradeoff between outlier fraction and sample size: eliminating sites with pre- or post-intervention periods shorter than 300 days would yield a very clean sample, but it would also eliminate more than 50% and 70% of DI and HVAC sites, respectively. To strike a balance between outlier fraction and sample size, in this analysis we chose to require a minimum of 120 days of data in both the pre-intervention and post-intervention period.

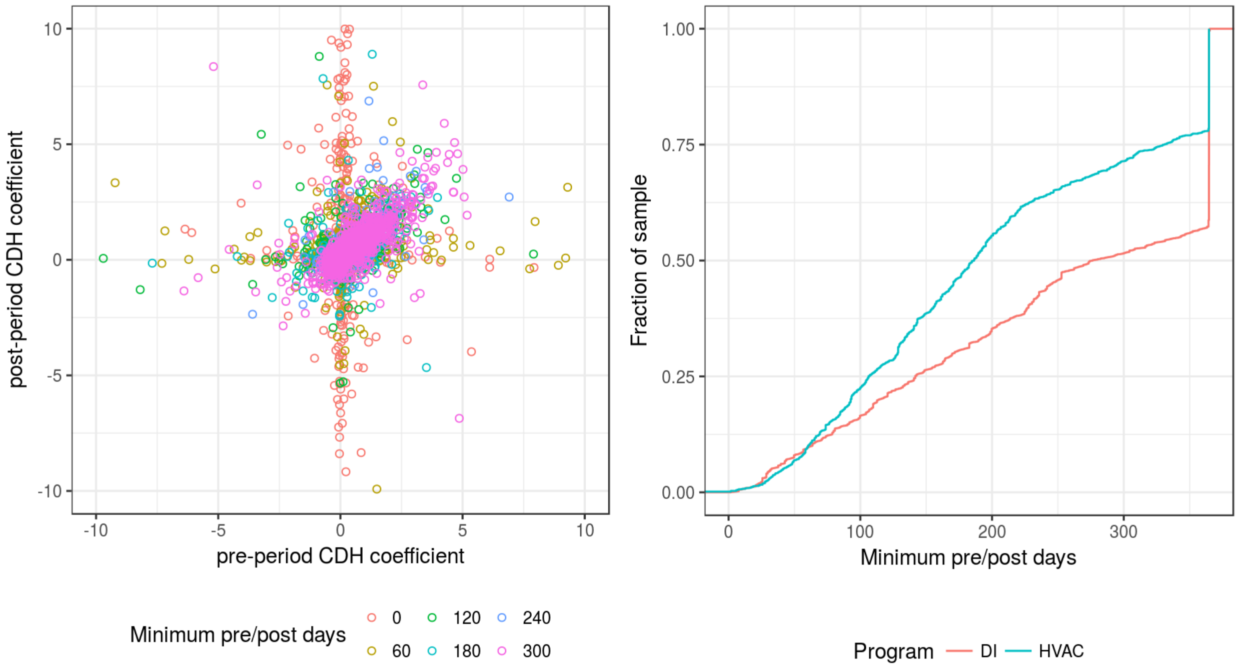

Figure 25. Impact of pre-period and post-period duration on evaluation quality and sample size. The left panel shows the model parameter that represents the response of cooling load to weather, which would be expected to show a strong correlation between the pre-period and the post-period. The expected relation is evident, but large outliers are present for sites with pre-periods or post-periods shorter than 120 days. The right panel shows the fraction of the sample that falls below a given threshold in pre/post period duration. As we increase the threshold, the fraction of the sample excluded rises rapidly, especially for the HVAC participants.

Step 2: Eliminating observed changes that are impossible through efficiency

As shown in Figure 25, although restricting to sites having a long pre-post period eliminates a large number of outliers, there is a small number of points that remain far off of the main relation, regardless of how high we set the threshold. In examining these sites, we found that a significant number had extremely large negative measured savings–sometimes well in excess of −1000 kWh/day–both in terms of total energy consumption and cooling energy consumption. Upon further examination, we found that many of these large outliers exhibited a qualitative change in usage pattern between the pre-intervention and post-intervention period. For example, we found numerous sites that appeared to have been largely idle prior to the intervention but fully operational afterward (see the ‘Computing and visualizing gains from filtering’ section that follows for an example). In these cases, there was an attendant large increase in energy consumption that was clearly not the result of the program intervention.

Indeed, the program intervention may have saved significant energy compared to what would have been consumed otherwise, but it is not possible to measure this accurately using the pre-intervention data as a baseline. There could also be sites in which the opposite situation holds: occupied during the pre-intervention period, but idle afterward. We anticipate that the former situation is much more common among customers who choose to participate in efficiency programs, since there would be little point in participating for customers who expect to shut down operations, but a large potential upside for customers who expect their consumption to ramp up. This has the potential to introduce a substantial (and spurious) negative bias in the average apparent savings.

To guard against this, we chose to exclude from our analysis any site where the daily energy consumption in the post-intervention period was different from the pre-period value by more than a factor of two in either direction. For instance, if a particular site had a pre-period energy consumption of 50 kWh/day, we would reject it from the sample if its post-period energy consumption was smaller than 25 kWh/day or larger than 100 kWh/day. This corresponds to a simple assumption that neither the DI or HVAC program should have a large enough impact to halve or double the total site-level energy consumption.

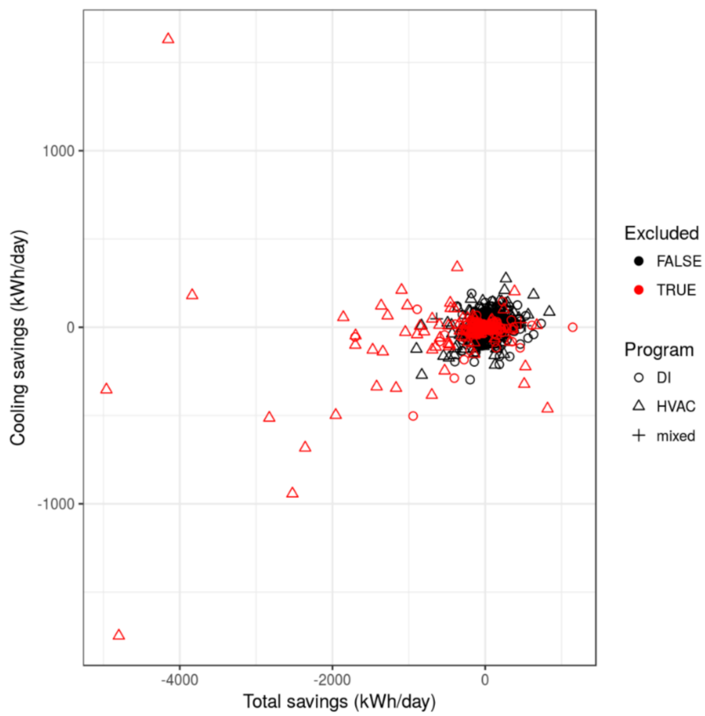

Figure 26. Scatter plot of daily total energy savings and cooling energy savings, for all evaluated program-participant sites. Extreme outliers are evident in both dimensions. Red points indicate sites whose total energy consumption changed by more than a factor of 2 between the pre-intervention sample (the central cluster of points is visually saturated by heavy overplotting); nevertheless, this simple filter effectively excludes all of the most extreme outliers.

This simple criterion eliminates only 4% of the program participants from the sample that passes the 120-day threshold, but it is very effective in excluding outliers in the savings distribution. Figure 26 shows the impact on outliers, both for total energy savings and cooling energy savings. There is a tight cluster of values near zero, and a smaller number of extremely large outliers are evident, both in the total-savings and cooling-savings dimensions, with a strong bias toward negative values in both dimensions. Red points indicate customers who are excluded by our simple factor-of-two criterion on the change in total energy consumption. Remarkably, this simple trim excludes all of the most visually evident outliers. It is also notable that most of the largest outliers occur among the HVAC program participants, which may not be surprising since this program skews heavily toward larger customers. As discussed in the results section below, the outlier rejection has an especially significant impact on the measured average savings of the HVAC program.

It is worth pausing briefly to discuss our reasons for choosing a factor of two as our threshold in pre- versus-post energy consumption, rather than some other multiple, since there is a balance that must be struck between stringently excluding outliers and retaining a sufficiently large and representative sample for further analysis. Figure 27 shows the measured total savings and cooling savings (post-period energy less baseline energy) for all program participants, compared to the logarithm of the ratio of post-period to pre-period total energy consumption. Horizontal lines indicate a factor of two (solid) or three (dashed) between the pre and post periods. It is clear that most of the extremely large outliers occur at sites with a post-to-pre-ratio exceeding a factor of three, and essentially all of them are associated with changes larger than a factor of two. Although the latter threshold excludes somewhat more data, the total amount of data excluded is still relatively small, at just over 4% of the sample. Further, we do not expect that these programs are likely to yield savings more than a few tens of percentage points in general. For these reasons, we selected the more restrictive factor-of-two threshold for our main analysis, to more effectively exclude the bulk of very large outliers.

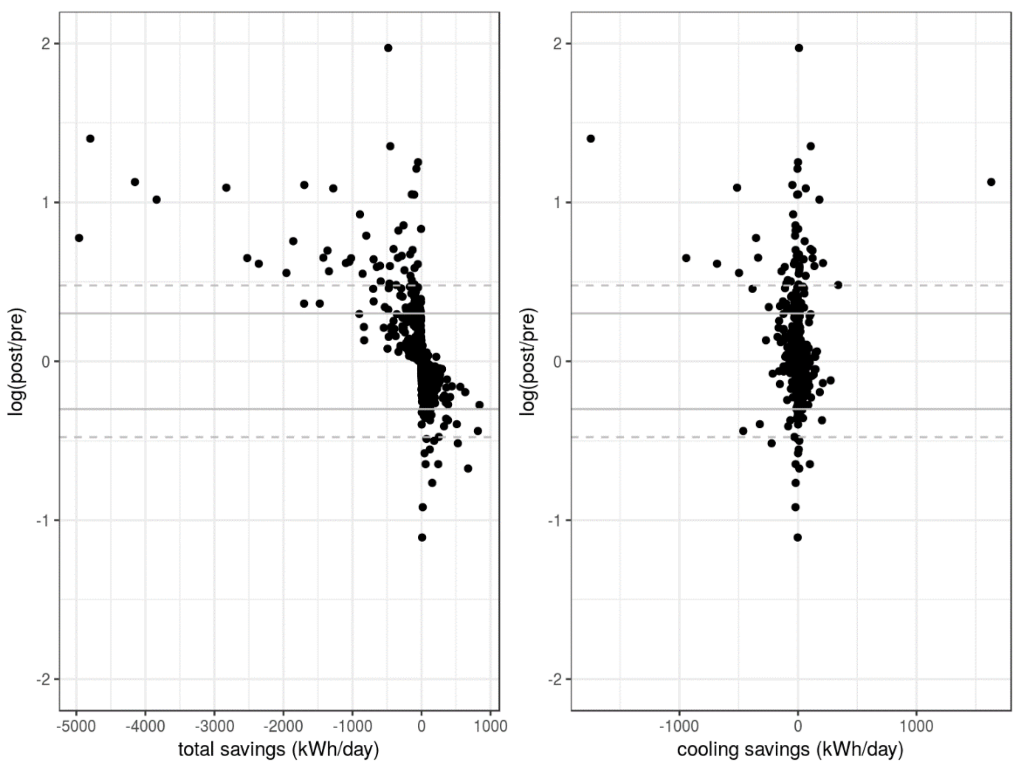

Figure 27. Total savings and cooling savings, compared to the logarithm of the post-period to pre-period energy consumption ratio. Horizontal lines indicate a factor of 2 (solid) and 3 (dashed). It is evident that most of the very large savings outliers lie at a ratio greater than 3, and essentially all of them fall at a ratio greater than 2.

Control group comparison

Simply observing a reduction in energy consumption following a program intervention does not definitively demonstrate that the intervention *caused* the savings. In addition to program impacts, there may be other external forces or trends (aside from weather) that drive energy savings across a broad range of customers. For instance, an increase in participation in other efficiency programs, or natural adoption of efficient technologies like flat screen TVs or LED lighting might reduce average energy consumption across all customers, regardless of their program participation. Conversely, wider uptake of consumptive devices can lead to greater consumption. If such non-program effects are not accounted for in our savings assessment, they might lead to a spuriously large or small measurement of the program impact. To guard against this, we compare our measured energy savings for each of the program participant groups against a control group of customers who did not participate in either the HVAC or the DI program, but are otherwise similar to program participants.

Control group resampling

Simple exploration of the control and participant data reveals that the control sample has a very different distribution of energy consumption than do the samples of DI or HVAC program participants. This is shown in Figure 28 for the HVAC program: the participants (blue) skew toward much higher energy consumption than the control group (gray). In the same figure, the distribution of site sizes reveals that the control group has a much greater proportion of small (S) and un-computed (N) site sizes than does the participant group. This will naturally lead to the observed differences in the energy consumption distributions.

As a simple thought experiment, let us suppose there is a non-weather, non-program effect that increases energy consumption by 1%, on average, across all SMB customers. In this situation, the HVAC participant group (blue) would show a much larger average increase than the control group (gray), simply by virtue of the fact that the participant sites are much larger on average. Such an effect could easily swamp any savings from the program interventions and lead to the (incorrect) conclusion that the program has had no impact, or even a negative impact.

To guard against this and ensure a fair comparison between the control and participant groups, we re-sample the control group so that it more closely resembles the relevant participant group. For each customer in the DI and HVAC program participant samples, we draw a customer from the control sample that is the nearest-neighbor to the participant of interest, within some space of parameters for which we would like to match the two samples. This selection is done *with replacement*, so that a given control customer might be chosen multiple times (in order to better populate the tails of the distribution being matched, for instance). To minimize the amount of duplication in our resampled control group, it is important not to include too many parameters in the resampling.

In this analysis, we choose to match on the total daily site energy consumption in the pre-intervention period, and on the site size category (S, M, L, or N). Figure 28 shows the results of this resampling strategy for the HVAC program: the resampled control group (red) has a nearly identical distribution in energy consumption to the participant group, and a much more similar distribution in site size. Broadly similar improvements in the control-participant sample resemblance occur when we resample for the DI program.

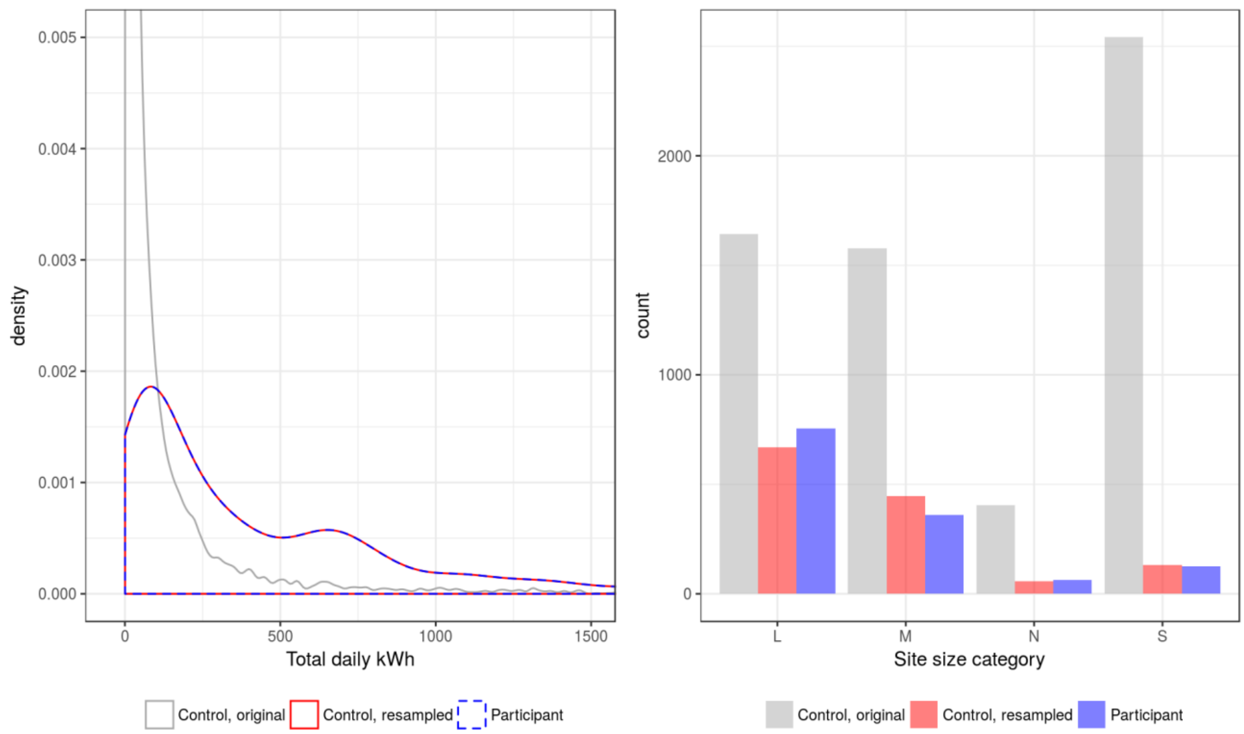

Figure 28. Distributions by total daily energy consumption and site size, for the original control group, the control group as resampled to match the HVAC participant group, and the HVAC participant group itself. Note that the original control group is skewed toward much smaller site sizes and much lower energy consumption than the participant group, while the resampled control group matches the participant distributions much more closely.

It is interesting to note that, even though we have not matched on the CDH-dependent component of the energy consumption (i.e., the cooling energy consumption), the above resampling strategy also yields a much-improved match between the control and participant distributions for this parameter. Because of this effect, we chose not to explicitly match on the HVAC energy consumption in order to minimize the amount of duplication in the resampled control group.

Control group evaluation

To compare the energy savings of the program participants to the control group, we subject the control group members' meter data to the same evaluation procedures that we used to measure weather-normalized energy savings for the program participants. Because the control group has (by definition) no program interventions, we developed a method for dividing their data into synthetic "pre-intervention" and "post-intervention" periods for the purposes of evaluation. First, to ensure that the evaluation code has access to realistic intervention metadata, we assigned each customer in the control group a random intervention drawn from the pool of all program interventions. We then assigned this intervention a randomly selected date between the beginning of 2014 and the end of 2015. This date delineates the pre- and post-intervention period for each control-group customer.

A simpler approach would be to choose a single date in the middle of the intervention period and split all control-group sites at that date. However, in the participant sample, the pre-intervention and post-intervention periods can be longer or shorter depending on whether the program intervention occurred early or late in the intervention period, and this variation in the length of the evaluation periods is a major source of noise in the program evaluation results. By contrast, simply splitting the control group at a fixed date would introduce significantly less noise of this type, and the control group evaluation could end up being unfairly advantaged relative to the program participants. Our random selection of dates allows us to better simulate the stochasticity in the program participants' evaluation periods.

Impact on average savings

Table 15 shows the impact of our outlier-rejection approach on the measured savings for each program. There are negligible changes in the measured savings of the DI program, but the impact on the HVAC savings is quite substantial, changing both the total and the cooling savings from negative to positive. This is consistent with Figure 26, which shows that the overwhelming majority of very large outliers are associated with the HVAC program.

Table 15. Savings before and after trimming sites having a >2x change in post/pre-consumption

Program type	Total savings	Cooling savings	Trimmed total savings	Trimmed cooling savings
DI	14.11	-0.17	14.39	0.04
HVAC	-37.58	-3.02	5.13	1.03

Computing and visualizing gains from filtering

Figure 29a is a histogram of the natural variability distribution for the DI control group. Figure 29b illustrates the NMEC pre/post savings calculated for all DI program participants. And Figure 29c illustrates the savings of the DI customers in the top quartile of pre-intervention average daily consumption. We define filter % as the percentage of customers eliminated by a given filter criterion - 75% in the case illustrated. There is more on the methods for making the NMEC savings estimates for controls and program participants and the filtering logic in the sections that follow. For now, our focus is simply on the shape and means of the distributions.

Figure 29: Computed outcomes for the direct install (DI) program studied. a – natural variability for reference; b – DI program NMEC savingss – note the skew; c – DI program participants filtered using the upper quartile of average pre-program energy consumption – note the skew and the mean.

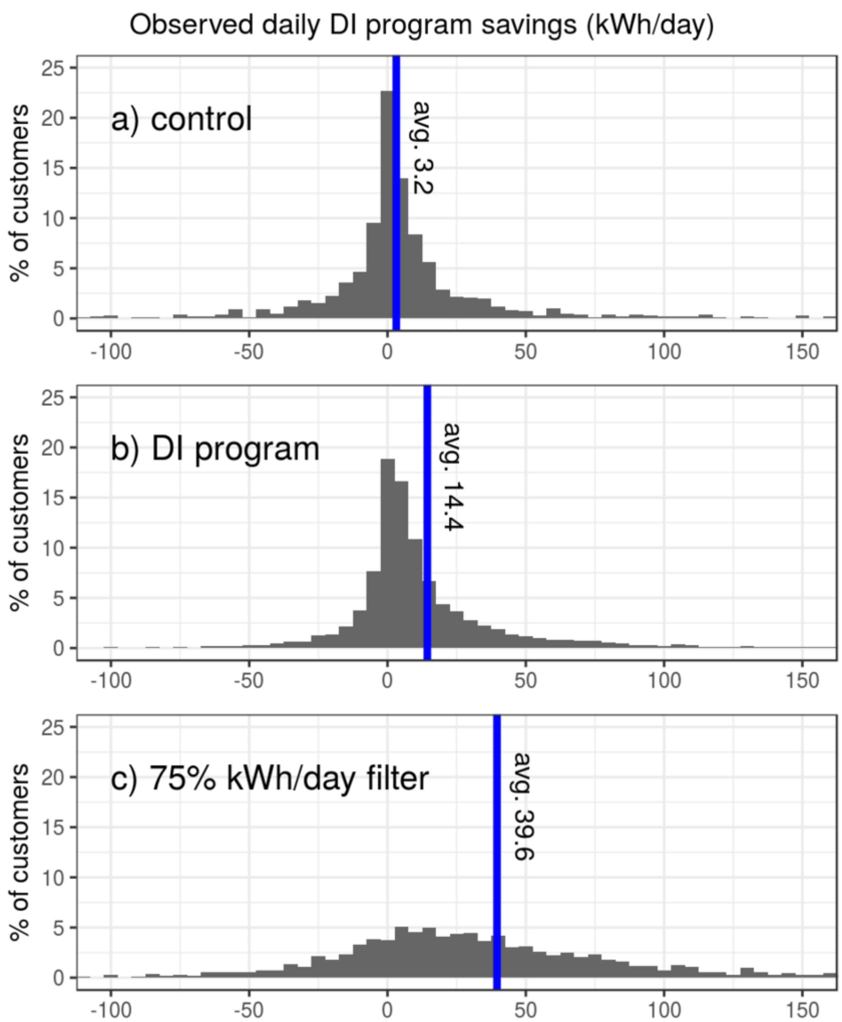

Referring to Figure 29, we observe the following:

- The controls have an average "savings" value close to zero and roughly symmetrical in the positive and nagative directions. But it is not exactly zero. There may be a long term trend of

energy savings or even changing business conditions playing out or perhaps the weather normalization is imperfect.
- The variability in "natural" outcomes (the width or standard deviation of a) is very large, suggesting that customer consumption is a moving target from year to year.
- The histogram of DI savings, seen in b, illustrates savings averaging over 14 kWh/day, but with deviations as large as found in the control group in a. This distribution is skewed, with more customers in positive territory than negative indicating that the program has a positive savings impact on average with heterogeneous results per-customer.
- The mean of the filtered participants found in c is nearly triple the mean of all participants. This helps quantify the potential value of targeting for this program.
- There is also a lot of variability in the outcomes of the filtered group, but the "skew" that provided evidence in b that the program worked is much more pronounced, with just a small minority of customers showing negative outcomes.

While these histograms are useful for recording proportions of customers in the studied data, it is useful at times to smooth those histograms into "density distributions" that more readily support overlaid comparisons and show the relative proportion of customers rather than absolute numbers. Figure 30a visualizes the same data from above as overlapping density distributions.

Figure 30: Panel a) provides a more compact visualization of all three panels from Figure 29. Savings are drawn as density curves rather than histograms and can now be more directly compared. In many cases, just the average savings for each distribution will be the subject of our analysis. Here the averages are displayed as dotted vertical lines. Panels b, c, and d extend the dotted lines with savings, gains about the unfiltered average, and gains as a percentage of the unfiltered average, respectively.

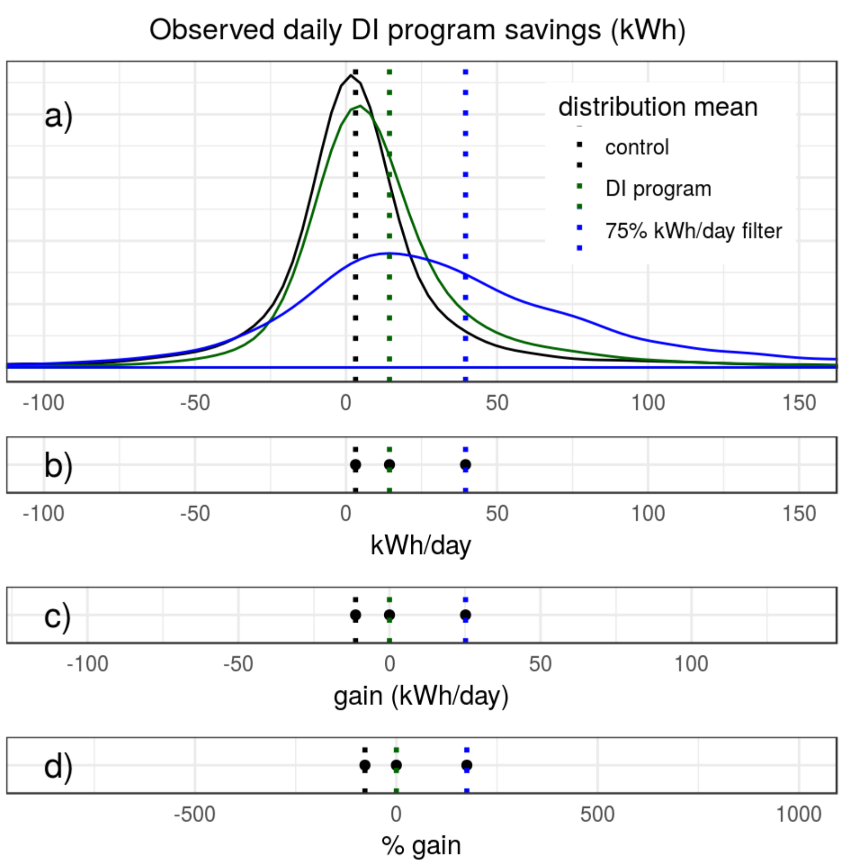

Figure 31 displays the density curve version of the HVAC program outcomes.

Figure 31: HVAC program savings are drawn as density curves for comparison between the HVAC controls, all program participants, and program participants filtered using the upper quartile of the ratio between summer and winter consumption, one of the top performing filters for HVAC, with corresponding mean values as dotted vertical lines.

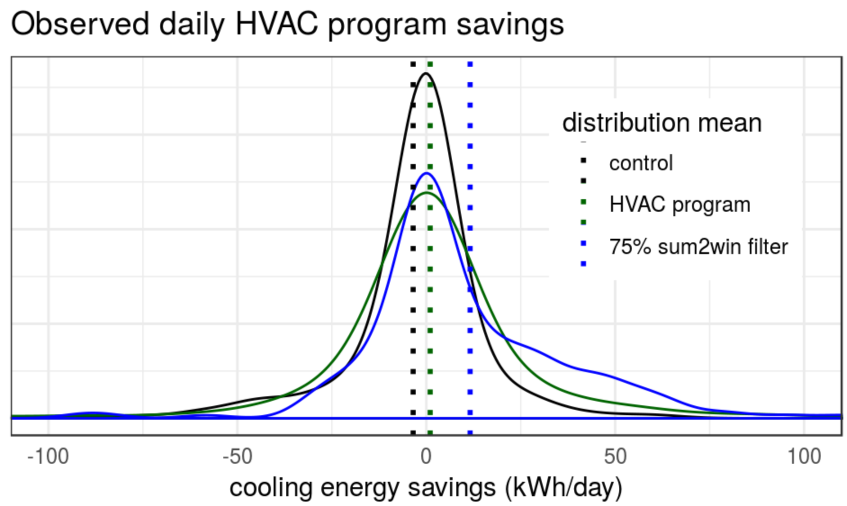

We will also be adopting a more compact visualization of the mean savings for different sets of customers, i.e. without the underlying distributions, that keeps the a kWh/day savings axis (although typically as a y-axis), but shows only a single point at the mean value of each relevant distribution. In some cases, we are more interested in how well targeting can work in general than what it would have accomplished in the specific programs studied, we will also present results as *% gain relative to all program participants*.

Who are the outliers?

It is instructive to look at the site-level data for the sites we have excluded as outliers due to unrealistic savings values, i.e., those whose total energy consumption has changed by more than a factor of two. Referring to Table 1, we note that removing such outliers reduced the premise count from 7767 to 7467 for DI and from 1304 to 1193 from HVAC.

Figure 32, Figure 33, and Figure 34

Figure 34 show a four-panel plot summarizing basic energy characteristics for each of three selected sites eliminated due to large swings in energy consumption.

In Figure 32, we see the energy consumption was effectively zero prior to 2015. This is evidently a site that was unoccupied during the pre-intervention period. Presumably the HVAC intervention was performed as part of an overall renovation before re-occupying the site. Broadly speaking this scenario may have the desired impact, but it is impossible to measure the savings for this site from a pre/post comparison. The pre-period data will yield an irrelevant baseline in the post period, and we will estimate an extremely large *negative* savings that will skew the average savings substantially. Examples like this one are quite common among sites whose energy consumption changes by more than a factor of two.

Figure 32. Summary plots of energy consumption characteristics for a site that has a large change in energy consumption between the pre and post periods. From upper left to lower right, the panels show the hourly demand for each day recorded as a

color-coded heat map, an example ten-day demand curve, the total kWh per day for each day recorded (black) overlaid on top of the average daily outside temperature (gray), and the relation between energy consumption (y-axis) and average outside temperature (x-axis) for each day recorded.

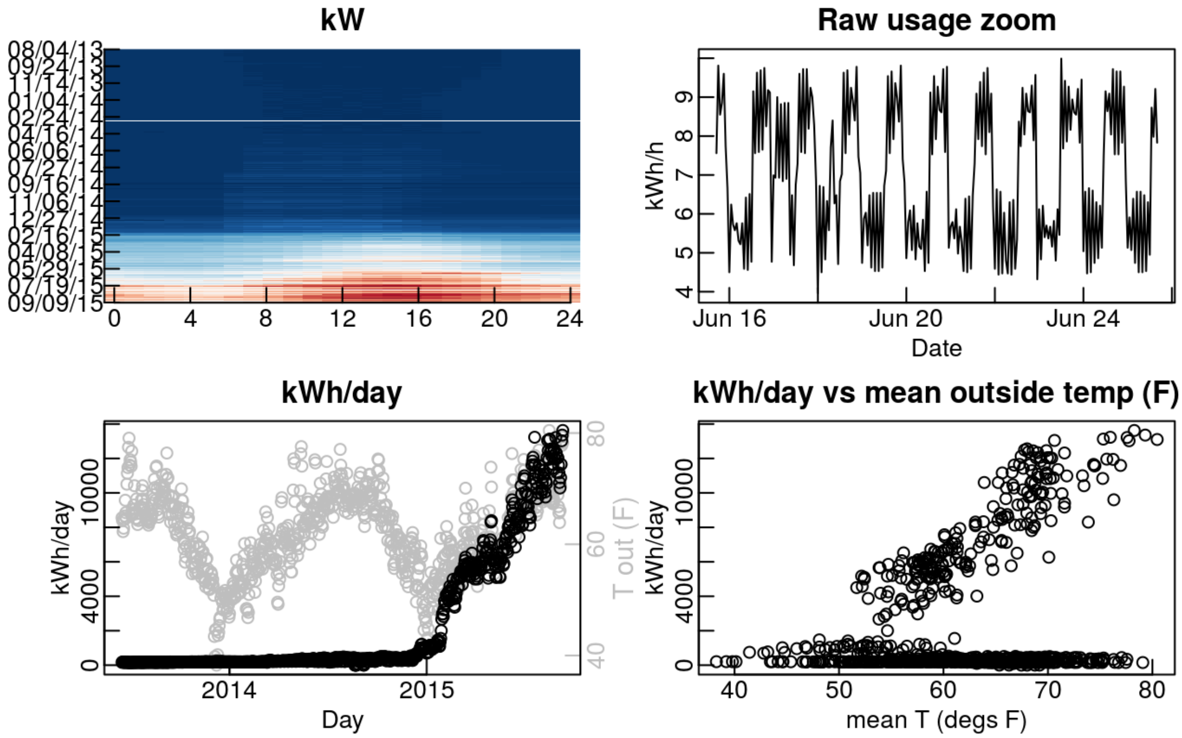

In Figure 33, we see the opposite case, where the site appears to have been shuttered late in the period of measurement. This scenario will yield a spurious positive outlier in our savings distribution. Finally, in

Figure 34, we see a more complicated case, where a large transient spike in consumption appears in the middle of the period. The impact on the measured savings will depend on whether the intervention occurred before or after this spike, but in either case, the pre-intervention period is unlikely to yield a relevant baseline against which to compare the post-intervention data.

Figure 33: Similar to Figure 32 for a different site.

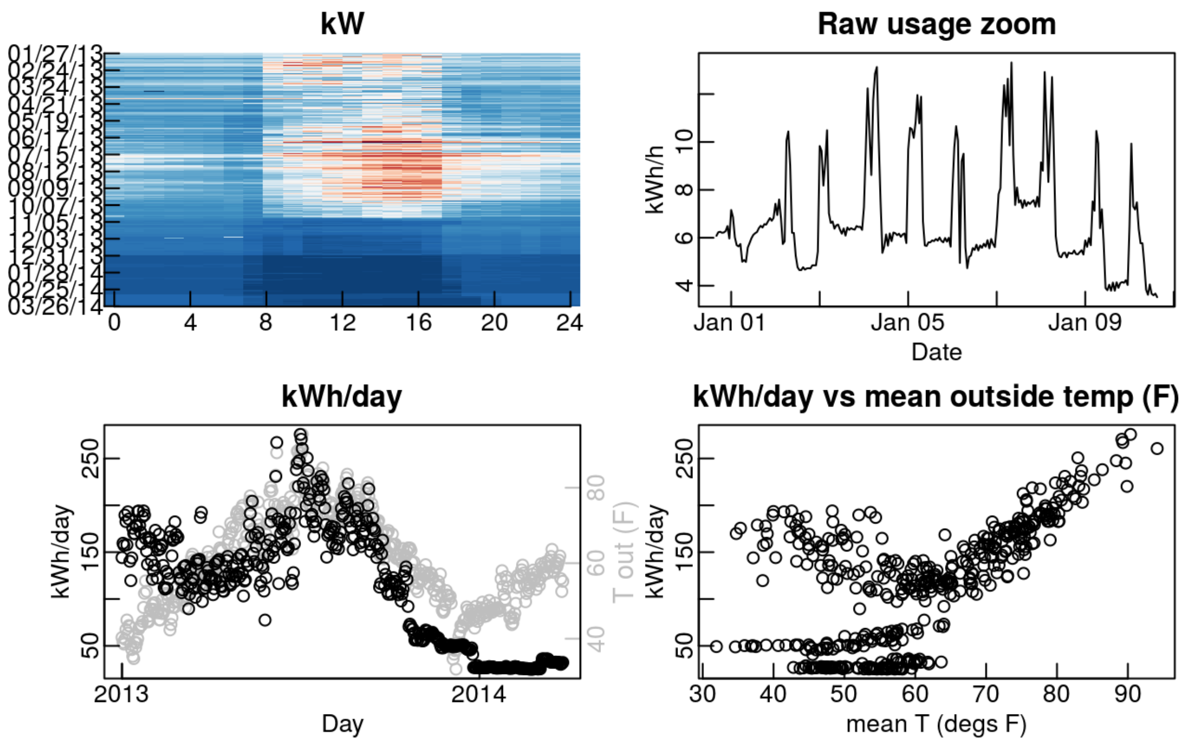

Figure 34. Similar to Figure 32, for a different example site.

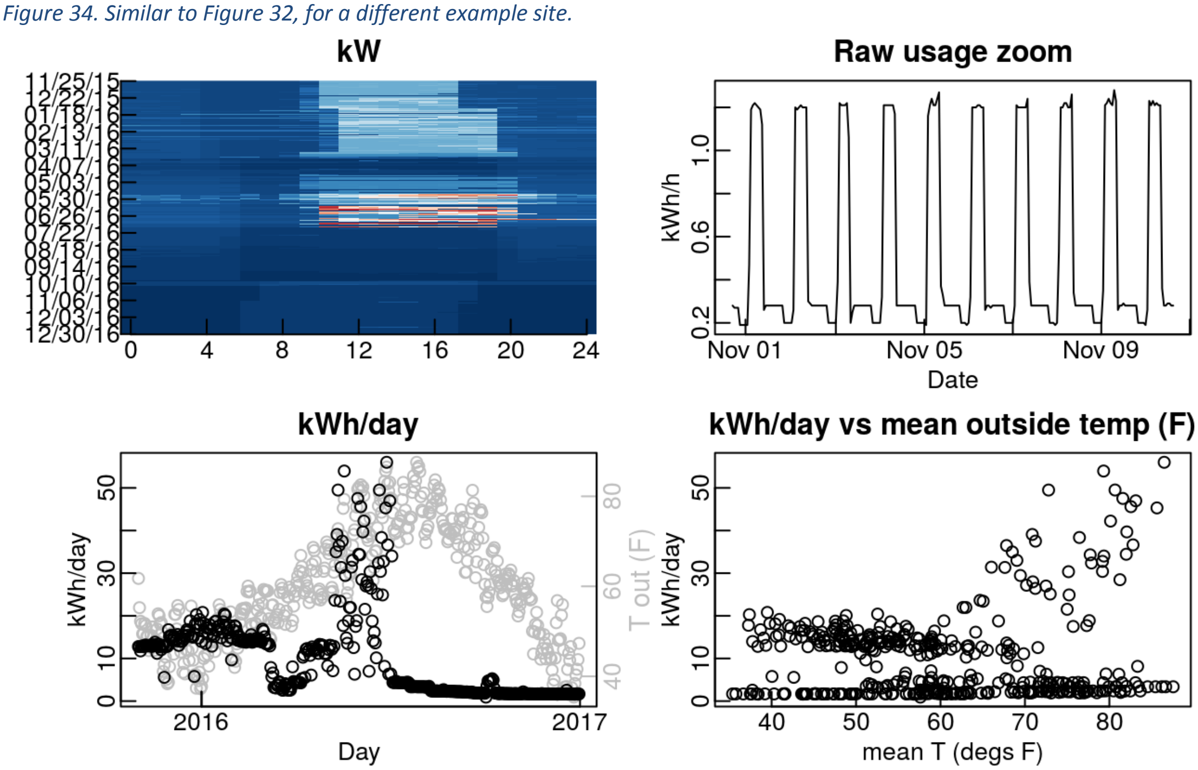

Appendix C: Tabulation of savings by customer characteristics

In this appendix, savings results for all participants are compared to sub-groups defined by customer characteristics for both programs.

How to read these tables

The format for each table in the section is as follows:

The heading describes the one or two categories of characteristics being examined. Then from left to right, the columns are:

<group characteristics> - The first one or two columns provide the specific category values that define the sub-group on each row. For example, in the first table the sub-groups are defined by their NAICS sector business type classification, whereas the second table splits each NAICS sector further into S, M, L, N total consumption 'customer size' categories.

Premise count – The count of premises that are in the sub-category group for the row. Groups with fewer than 20 premises have been excluded as groups with the fewest premises are the most likely to have fluke results. Note that the tables have an 'All participants' group row at the top. This row provides the count and other metrics for all premises used in the tabulation of the table, with all the sub-groups together. In the case of table that use two separate group characteristics, the 'All participants' entries are split across the secondary group, so the second table in this appendix has such entries for S, M, L, and N total consumption categories.

daily savings – The at-the-meter average estimated daily savings in kWh for the sub-group.

daily savings (% of pre) - The at-the-meter average estimated daily savings for the sub-group as a percentage of pre-intervention average daily total consumption.

daily savings std dev – The standard deviation of the daily savings.

gain - The daily savings improvement, in daily kWh, over the corresponding 'All participants' savings entry that a sub-group exhibits.

% gain – The percentage improvement over the corresponding 'All participants' daily savings entry that a sub-group exhibits.

We define savings "gain" as the mean daily kWh saved in the filtered group subtracted from the mean of the unfiltered program savings. The purpose of gain calculations is to quantify the per-customer savings a filter provides *in addition to* what was already being achieved. The gain visualized in Figure 29c is 39.6 – 14.4 = 25.2 kWh/day of additional mean savings for the filtered group compared to the average savings of all participants. Similarly, **we define "% gain" as the gain divided by the average savings of all participants** 25.2 / 14.4 = a % gain of 175. In other words, the filtered group of participants has an average savings that is 175% above and beyond what the average participant saves.[20] Similarly, the unfiltered group of all participants, by definition has a % gain of 0.

Where customer size is part of the tabulation, % gain numbers are relative to all customers that share the same size designation, rather than relative to the larger group of all participants. Note that none of the sub-groups were filtered using feature data – all results in this appendix are based solely on

[20] Note that in this example, the filtered savings mean is 2.75x the unfiltered savings but only 2.75 – 1 = 1.75x is additional and counted as the gain.

customer characteristics, not feature filtering. Appendix D presents feature filtering and mixed customer characteristics and feature filtering results.

Tabulation of DI savings by customer characteristics

For DI savings, we are reporting the gain %, or the percentage of additional savings on top of the average for all customers each sub-group achieves. In other words, a gain % of 10 means that a sub-group saved an average of 10% more than all customers together.

DI tabulation by size

For DI, savings track customer "size," meaning magnitude of consumption, categories very well. Recall that size "N" indicates customers with too little data to compute the official size metric. These newer customers tend to be smaller businesses, which are more likely to start up and move than larger ones. Note also that smaller customers tend to save a greater percentage of their pre-intervention consumption than larger customers, indicating deeper savings are being achieved.

customer size	premise count	daily savings (% of pre)	daily savings (kWh)	daily savings std dev	% gain
L	1025	5.41	28.17	83.86	96
M	3403	4.94	17.36	45.94	21
N	474	7.06	12.27	42.23	-15
S	2595	9.65	5.44	13.38	-62

DI tabulation by rate

This tabulation shows strong performance for both "Medium" rates (E19 and A10), which makes sense given the correlation between customer size and savings, but the TOU version of the rate performs significantly better than the standard one. Possibly more striking, the group on the TOU rate for "Small" customers (A6) out-saves the general population by 19%, while the standard "Small" rate (A1) saves 58% less than the general population.

rate	premise count	daily savings (% of pre)	daily savings (kWh)	daily savings std dev	% gain
All participants	7497	6.77	14.39	46.38	0
E19 Medium general demand TOU	720	5.72	38.51	85.36	168
A10 Medium general demand	1631	5.95	28.41	66.47	97
A6 Small general service TOU	275	7.82	17.12	38.64	19
AG5 Large TOU agriculture	24	2.60	7.47	140.60	-48
A1 Small general service	4766	7.24	6.01	18.98	-58
GNR1 Small commercial gas	25	3.63	3.01	11.80	-79
TC traffic control	37	-2.41	0.15	3.14	-99

DI tabulation by technology family

The DI program is dominated by lighting and refrigeration installs, at least one of which is performed at nearly every premise served by the program. Lighting alone out-saves refrigeration alone, but they perform 11% and 30% (respectively) worse than the DI program average. When lighting and refrigeration are done together, the savings are a remarkable 105% greater than the DI program average. Lighting interventions tend to save a little over 7% of pre-intervention total consumption, but refrigeration saves just over 2%.

tech family	premise count	daily savings (% of pre)	daily savings (kWh)	daily savings std dev	% gain
All participants	7497	6.77	14.39	46.38	0
LIGHTING and REFRIGERATION	1005	7.80	29.46	56.60	105
HVAC	27	3.40	21.73	56.92	51
APPLIANCES and LIGHTING and REFRIGERATION	23	1.55	20.79	47.35	44
LIGHTING	5331	7.57	12.77	44.30	-11
REFRIGERATION	772	2.32	10.10	47.03	-30
ELECTRONICS AND IT and LIGHTING	50	8.60	7.16	17.23	-50
NA and REFRIGERATION	45	2.14	6.40	23.14	-56
APPLIANCES and REFRIGERATION	33	1.14	3.56	34.83	-75
APPLIANCES and LIGHTING	85	2.12	0.78	14.65	-95
APPLIANCES	91	-3.51	-2.80	36.05	-119

DI tabulation by technology family, size

The general technology family pattern observed above holds across size categories, with savings from “lighting and refrigeration” interventions for size L customers a remarkable 269% larger than the DI program average and 88% larger than the size L average. Refrigeration savings as a % of pre-intervention total consumption tend to decrease with size (the refrigeration loads are a smaller % of the total for large customers). The pattern is a little less clear for lighting. At nearly 10%, lighting savings as a % of pre-intervention total consumption are greatest for size S customers, but lowest at 5% for size M.

tech family	customer size	premise count	daily savings (% of pre)	daily savings (kWh)	daily savings std dev	% gain
All participants	L	1025	5.41	28.17	83.86	0
All participants	M	3403	4.94	17.36	45.94	0
All participants	N	474	7.06	12.27	42.23	0
All participants	S	2595	9.65	5.44	13.38	0
LIGHTING and REFRIGERATION	L	112	7.02	53.10	89.25	88
LIGHTING	L	662	6.24	27.67	86.18	-2
REFRIGERATION	L	171	1.74	16.61	79.46	-41
LIGHTING and REFRIGERATION	M	720	7.32	28.32	52.78	63
LIGHTING	M	2013	5.25	16.94	46.62	-2

tech family	customer size	premise count	daily savings (% of pre)	daily savings (kWh)	daily savings std dev	% gain
REFRIGERATION	M	468	2.18	8.74	31.00	-50
NA and REFRIGERATION	M	28	2.56	5.07	24.25	-71
APPLIANCES and LIGHTING	M	37	-1.26	0.23	20.70	-99
APPLIANCES	M	70	-3.12	-3.75	33.16	-122
LIGHTING and REFRIGERATION	N	70	9.28	28.94	54.21	136
LIGHTING	N	311	8.61	10.85	40.25	-12
REFRIGERATION	N	85	1.73	5.33	36.36	-57
LIGHTING and REFRIGERATION	S	103	10.93	12.05	17.64	121
REFRIGERATION	S	48	6.80	8.59	36.24	58
LIGHTING	S	2345	9.81	5.23	12.34	-4
ELECTRONICS AND IT and LIGHTING	S	33	12.50	4.09	9.16	-25
APPLIANCES and LIGHTING	S	41	6.50	1.67	6.33	-69

DI tabulation of technology for LIGHTING family only

Within the lighting category, LEDs (24% gain) out-save linear fluorescents (-49% gain), and CFLs (-85% gain). Projects involving both LEDs and linear fluorescents (55% gain) have the best savings performance. CFLs have little savings impact on their own or paired with other technologies.

technology	premise count	daily savings (% of pre)	daily savings (kWh)	daily savings std dev	% gain
All participants	5331	7.6	12.8	44.3	0
LED and LINEAR FLUORESCENT	645	9.7	19.8	56.4	55
LED and LIGHTING CONTROLS AND SENSORS and LINEAR FLUORESCENT	32	12.1	17.8	23.1	40
LED	2065	9.0	15.8	46.9	24
COMPACT FLUORESCENT and LED	124	8.3	15.8	42.7	24
COMPACT FLUORESCENT and LED and LINEAR FLUORESCENT	207	9.8	15.6	53.9	22
COMPACT FLUORESCENT and LED and LIGHTING CONTROLS AND SENSORS and LINEAR FLUORESCENT	26	5.9	15.1	37.4	19
LIGHTING CONTROLS AND SENSORS and LINEAR FLUORESCENT	90	7.3	12.4	80.3	-3
LED and NA	772	6.8	7.7	25.4	-40
COMPACT FLUORESCENT and LIGHTING CONTROLS AND SENSORS and LINEAR FLUORESCENT	31	4.3	7.6	17.4	-40
LINEAR FLUORESCENT	995	3.3	6.5	36.3	-49
COMPACT FLUORESCENT and LINEAR FLUORESCENT	204	8.2	6.3	35.3	-50
LED and LINEAR FLUORESCENT and NA	33	5.9	3.6	36.6	-72
COMPACT FLUORESCENT	40	-0.6	1.9	19.0	-85

DI tabulation of technology for REFRIGERATION family only

Refrigeration interventions are dominated by walk-in cooler projects. Savings from walk-in cooler interventions on their own are 5% less than all refrigeration projects (33% less than all DI), which are lower performers than DI projects in general. Savings from controls alone are 75% less than all refrigeration projects. However, savings from walk-in coolers and controls together are nearly 140% greater than all refrigeration projects (67% greater than all DI). Controls appear to be key drivers of refrigeration savings or perhaps controls problems are key drivers of inefficiency in refrigeration.

technology	premise count	daily savings (% of pre)	daily savings (kWh)	daily savings std dev	% gain
All participants	772	2.3	10.1	47.0	0
REFRIGERATION CONTROL and WALK-IN COOLER	42	9.0	24.0	25.2	138
WALK-IN COOLER	687	1.9	9.6	47.1	-5
REFRIGERATION CONTROL	40	1.5	2.5	60.7	-75

DI tabulation by intervention year

Customers participating in program year 2015 were likely to save a bit more than other recent years. Program year 2016 had disappointing savings. However, the most notable result here is that customers that participated in programs to the extent that their interventions spanned more than one year dramatically out-performed the general population. This is likely partially due to the cumulative impact of multiple interventions and partially to the self-selection effect of highly motivated customers participating multiple times in programs.

Intervention year	premise count	daily savings (% of pre)	daily savings (kWh)	daily savings std dev	% gain
All participants	7497	6.77	14.39	46.38	0
more than one	386	9.32	33.23	74.66	131
2015	2409	7.69	16.70	49.01	16
2013	1026	6.78	13.31	48.28	-7
2014	2431	6.14	12.79	41.96	-11
2016	1245	5.41	8.11	32.55	-44

DI tabulation by intervention year, size

Intervention year	customer size	premise count	daily savings (% of pre)	daily savings (kWh)	daily savings std dev	% gain
All participants	L	1025	5.41	28.17	83.86	0
All participants	M	3403	4.94	17.36	45.94	0
All participants	N	474	7.06	12.27	42.23	0
All participants	S	2595	9.65	5.44	13.38	0
more than one	L	51	7.42	54.95	113.87	95

Intervention year	customer size	premise count	daily savings (% of pre)	daily savings (kWh)	daily savings std dev	% gain
2015	L	375	6.99	36.88	91.02	31
2013	L	120	3.66	21.82	95.33	-23
2016	L	147	4.06	20.73	64.12	-26
2014	L	332	4.55	19.82	71.06	-30
more than one	M	237	8.21	35.51	74.90	105
2015	M	1019	5.86	19.56	45.95	13
2013	M	439	4.97	17.27	47.13	-1
2014	M	1094	4.64	16.84	42.22	-3
2016	M	614	2.66	7.71	32.42	-56
more than one	N	20	11.22	43.35	64.79	253
2013	N	120	10.39	13.80	35.82	12
2015	N	101	4.90	9.95	32.21	-19
2014	N	230	5.78	9.64	45.89	-21
more than one	S	78	13.45	9.51	17.32	75
2015	S	914	10.32	5.97	13.34	10
2013	S	347	8.91	5.21	20.39	-4
2014	S	775	9.05	4.99	10.92	-8
2016	S	481	9.27	4.66	9.23	-14

DI tabulation by intervention year, technology family

Intervention year	tech family	premise count	daily savings (% of pre)	daily savings (kWh)	daily savings std dev	% gain
All participants	All participants	7497	6.77	14.39	46.38	0
more than one	LIGHTING and REFRIGERATION	155	8.80	39.65	76.31	175
2015	LIGHTING and REFRIGERATION	291	9.16	34.40	59.12	139
more than one	LIGHTING	194	10.99	32.85	78.17	128
2014	LIGHTING and REFRIGERATION	308	7.81	28.38	51.18	97
2013	LIGHTING and REFRIGERATION	159	6.05	19.10	39.25	33
2015	REFRIGERATION	249	3.81	18.29	59.66	27
2016	LIGHTING and REFRIGERATION	92	4.75	18.16	47.61	26
2015	LIGHTING	1744	8.37	14.20	45.80	-1
2013	LIGHTING	700	7.47	13.12	50.95	-9
2014	LIGHTING	1731	6.98	11.49	40.49	-20
2016	LIGHTING	962	6.58	8.16	30.04	-43
2015	ELECTRONICS AND IT and LIGHTING	46	8.45	7.38	17.92	-49
2013	REFRIGERATION	151	3.72	7.18	45.51	-50
2016	NA and REFRIGERATION	45	2.14	6.40	23.14	-56
2014	REFRIGERATION	363	0.63	5.50	36.42	-62

2016	APPLIANCES and REFRIGERATION	24	1.46	4.02	32.47	-72
2015	APPLIANCES and LIGHTING	45	-1.18	-0.05	8.19	-100
2016	APPLIANCES	79	-3.81	-6.52	32.77	-145

DI tabulation by NAICS SECTOR abbreviation

Savings from sports, entertainment, and recreation venues, non-department stores, and more technical manufacturing are all around 70% greater than DI projects in general.

NAICS sector	premise count	daily savings (% of pre)	daily savings (kWh)	daily savings std dev	% gain
All participants	7497	6.77	14.39	46.38	0
Uncategorized	81	7.48	28.28	54.14	97
Arts, Entertainment, and Recreation	218	8.75	24.74	70.66	72
RETAIL TRADE - 1	1786	9.26	24.27	55.37	69
MANUFACTURING - 3	62	4.60	24.26	110.88	69
TRANSPORTATION	56	8.21	23.23	50.00	61
Educational Services	146	7.10	19.23	38.45	34
MANUFACTURING - 2	85	2.92	18.88	57.43	31
Management of Companies …	63	10.61	17.16	36.07	19
TRANSPORTATION & WAREHOUSING	103	4.37	16.59	86.76	15
Wholesale Trade	237	7.24	14.91	63.67	4
Information	35	0.65	14.78	78.82	3
Real Estate and Rental and Leasing…	308	6.74	13.78	64.56	-4
Health Care and Social Assistance	211	7.36	13.35	38.53	-7
MANUFACTURING - 4	151	6.92	11.79	49.11	-18
Accommodation and Food Service	1037	2.62	10.65	35.68	-26
Agriculture, Forestry, Fishing …	62	7.00	10.17	58.65	-29
RETAIL TRADE - 2	375	14.67	9.31	18.60	-35
Other Services	993	7.73	7.83	25.51	-46
Construction	74	5.52	7.69	32.21	-47
N/A	957	2.50	7.58	29.66	-47
Administrative and Support …	93	9.25	7.49	21.20	-48
Public Administration	149	3.92	7.43	33.17	-48
Professional, Scientific, …	126	6.20	6.63	30.01	-54
Finance and Insurance	68	9.41	6.04	13.43	-58

DI tabulation by NAICS SECTOR abbreviation, size

NAICS sector	customer size	premise count	daily savings (% of pre)	daily savings (kWh)	daily savings std dev	% gain
All participants	L	1025	5.41	28.17	83.86	0
All participants	M	3403	4.94	17.36	45.94	0
All participants	N	474	7.06	12.27	42.23	0
All participants	S	2595	9.65	5.44	13.38	0
Real Estate and Rental …	L	45	6.67	46.86	117.70	66
RETAIL TRADE - 1	L	169	5.91	46.54	102.75	65
Wholesale Trade	L	32	7.08	35.48	143.88	26
Accommodation and Food Service	L	71	3.53	34.63	78.12	23
Arts, Entertainment, and Recreation …	L	88	8.07	30.67	93.23	9
Other Services	L	56	7.14	28.19	70.69	0
Educational Services	L	79	6.20	25.09	45.02	-11
N/A	L	160	3.57	21.97	50.09	-22
TRANSPORTATION & WAREHOUSING	L	27	-1.45	19.18	141.18	-32
Health Care and Social Assistance	L	27	5.17	18.53	52.48	-34
Management of Companies and En	L	20	12.25	16.89	40.07	-40
Public Administration	L	121	4.46	7.37	35.71	-74
Uncategorized	M	36	11.07	44.85	58.97	158
Arts, Entertainment, and Recreation …	M	59	6.31	38.38	69.41	121
TRANSPORTATION	M	23	11.85	34.68	54.62	100
RETAIL TRADE - 1	M	1096	7.53	26.71	52.63	54
Management of Companies …	M	27	9.58	24.67	41.76	42
TRANSPORTATION & WAREHOUSING	M	44	5.16	22.17	73.27	28
MANUFACTURING - 2	M	48	1.79	21.22	61.72	22
Health Care and Social Assistance	M	86	5.94	20.37	48.84	17
MANUFACTURING - 4	M	73	4.88	18.79	51.11	8
Construction	M	20	7.71	18.69	45.52	8
Wholesale Trade	M	93	5.75	17.20	52.20	-1
Educational Services	M	39	5.68	15.76	34.37	-9
RETAIL TRADE - 2	M	81	9.22	15.20	30.26	-12
Real Estate and Rental …	M	129	7.43	15.11	65.57	-13
MANUFACTURING - 3	M	23	-5.86	13.90	82.83	-20
Other Services	M	236	4.79	12.48	31.90	-28
Professional, Scientific, …	M	37	4.05	12.48	41.17	-28
Public Administration	M	22	6.14	10.50	19.88	-40
Accommodation and Food Service	M	694	1.94	8.73	29.23	-50
N/A	M	464	1.31	6.17	27.21	-64
Administrative and Support …	M	28	4.34	5.15	27.28	-70

NAICS sector	customer size	premise count	daily savings (% of pre)	daily savings (kWh)	daily savings std dev	% gain
"Agriculture, Forestry, Fishing …	M	20	0.40	-7.40	50.58	-143
RETAIL TRADE - 1	N	109	14.36	27.57	61.03	125
Accommodation and Food Service	N	130	4.48	12.73	37.74	4
RETAIL TRADE - 2	N	31	16.25	6.27	12.25	-49
Other Services	N	52	7.11	1.93	25.73	-84
N/A	N	46	-0.68	-0.01	18.64	-100
Real Estate and Rental …	N	29	-7.00	-9.87	35.16	-180
MANUFACTURING - 3	S	25	10.13	9.23	21.78	70
RETAIL TRADE - 1	S	412	13.88	7.75	12.07	42
Educational Services	S	28	11.61	7.54	12.77	39
RETAIL TRADE - 2	S	261	16.14	7.17	8.80	32
Arts, Entertainment, and Recreation …	S	65	11.61	6.63	10.93	22
Wholesale Trade	S	98	10.25	6.49	13.64	19
Accommodation and Food Service	S	142	3.79	6.12	22.13	13
Administrative and Support …	S	51	11.90	5.16	8.56	-5
TRANSPORTATION & WAREHOUSING	S	27	8.88	5.02	21.18	-8
"Professional, Scientific, …	S	80	7.75	4.86	10.09	-11
Other Services	S	649	8.91	4.86	9.79	-11
Real Estate and Rental …	S	105	9.72	4.49	10.08	-17
Uncategorized	S	27	7.22	4.44	8.57	-18
Health Care and Social Assistance	S	93	8.65	4.22	8.39	-22
Construction	S	49	4.68	4.19	12.57	-23
Finance and Insurance	S	40	10.01	4.15	7.14	-24
TRANSPORTATION	S	22	4.33	3.21	8.41	-41
N/A	S	287	4.33	3.04	11.99	-44
MANUFACTURING - 2	S	26	4.51	2.35	5.21	-57
MANUFACTURING - 4	S	61	9.21	2.27	40.89	-58

DI tabulation by climate zone

The coastal climate zones cz03 and cz05 significantly out-perform others.

climate zone	premise count	daily savings (% of pre)	daily savings (kWh)	daily savings std dev	% gain
All participants	7497	6.77	14.39	46.38	0
cz05	228	9.69	21.56	46.01	50
cz03	1756	7.76	18.19	57.98	26
N/A	42	11.57	16.85	30.94	17
cz13	1120	6.39	16.70	49.75	16
cz11	737	7.17	15.37	35.44	7

cz16	116	6.49	14.72	26.54	2
cz04	673	8.82	14.40	45.99	0
cz12	1485	5.40	11.34	39.61	-21
cz02	805	5.06	10.05	42.66	-30
cz01	535	5.99	7.42	34.01	-48

Tabulation of HVAC savings by customer characteristics

Because the average for all customers for the HVAC program is quite small (on the order of 1 kWh/day) and gain percentages would therefore be potentially misleadingly large, we are reporting the gain (not the gain %). As a reminder, the gain for a sub-group is the daily kWh of additional savings on top of the average for all customers it achieves. In other words, a gain of 5 means that a sub-group saved 5 kWh/day more than the average of all customers.

Savings are computed as the change in modeled temperature responsive loads between the pre- and post- periods. Due to varying occupancy or non-AC loads that are systematically lower during warmer periods, some temperature response coefficients can be negative (less energy used during hotter weather). Those negative values can factor into the % of pre or kWh values of group savings and will occasionally produce a negative value for group averages.

HVAC tabulation by size

HVAC program savings are loosely correlated with customer size, but the effect is not nearly as strong as it was for the DI program.

customer size	premise count	daily AC savings (% of pre)	daily AC savings (kWh)	daily AC savings std dev	% AC gain
L	694	0.27	1.52	38.31	48
M	327	0.66	0.96	26.65	-7
N	53	-2.97	-1.77	34.70	-272
S	119	-2.47	-0.36	6.37	-135

HVAC tabulation by climate zone

One might expect HVAC saving to strictly correlate with hotter climate zones. However here we see that the hottest climate zones, cz12, and cz13 are not the strongest performers. The northern central valley, cz11, and northern coast, cz03 perform best. There are clearly some location-specific customer attributes overcoming climate effects. It is also worth noting in this context that the HVAC program primarily performs quality and maintenance actions that impact pumps, fans, and other aspects of air and water distribution as much as the AC unit systems themselves.

climate zone	premise count	daily AC savings (% of pre)	daily AC savings (kWh)	daily AC savings std dev	% AC gain
All participants	1193	-0.04	1.03	33.24	0
cz11	98	1.51	2.30	25.24	122

cz03	282	0.55	2.06	32.13	99
cz12	345	-0.59	1.78	35.77	72
cz04	202	0.41	1.33	24.61	28
cz13	163	-2.09	-1.75	44.23	-269
cz02	65	1.43	-3.23	24.52	-413

HVAC tabulation by rate

rate	premise count	daily AC savings (% of pre)	daily AC savings (kWh)	daily AC savings std dev	% AC gain
All participants	1193	-0.04	1.03	33.24	0
E19 Medium general demand TOU	249	0.36	2.76	35.93	167
A10 Medium general demand	354	0.16	1.94	48.23	88
A1 Small general service	530	-0.28	0.19	16.66	-82
A6 Small general service TOU	46	-1.10	-5.50	22.70	-632

HVAC savings by technology

Chiller projects dramatically out-perform other types, with unitary AC projects associated with above average savings. Notably, the quality maintenance interventions are associated with below average savings – the actual average is negative, but this is likely just a symptom of the variability being so much higher than the savings that outliers can dominate average outcomes.

technology	premise count	daily AC savings (% of pre)	daily AC savings (kWh)	daily AC savings std dev	% AC gain
All participants	1193	0.0	1.0	33.2	0
CHILLER	45	0.7	7.2	43.1	593
AIR DISTRIBUTION and QUALITY MAINTENANCE and UNITARY AC/HP	20	0.6	5.0	13.2	384
UNITARY AC/HP	721	0.3	1.9	37.6	83
AIR DISTRIBUTION and MOTORS PUMPS AND FANS and QUALITY MAINTENANCE	20	0.2	1.9	6.1	80
AIR DISTRIBUTION and MOTORS PUMPS AND FANS and QUALITY MAINTENANCE and UNITARY AC/HP	29	-0.2	1.5	10.3	48
HVAC CONTROL and QUALITY MAINTENANCE	96	-1.6	0.0	18.2	-100
QUALITY MAINTENANCE	90	-0.4	-2.3	22.5	-323
HVAC CONTROL and QUALITY MAINTENANCE and UNITARY AC/HP	127	-1.1	-4.4	24.9	-524

HVAC tabulation by intervention year

As with DI, program year 2015 had noticeably better results than others, with program year 2013 returning noticeably worse results than others. Unlike DI, premises with interventions spanning more than one year did not perform better than their peers. The overall savings numbers are pretty small and the sample of participants have high variance of estimates savings and fairly low count of premises with data for a given year. The annual differences could be artifacts of imperfect weather normalization as weather changes from one year to the next.

year	premise count	daily AC savings (% of pre)	daily AC savings (kWh)	daily AC savings std dev	% AC gain
All participants	1193	-0.04	1.03	33.24	0
2015	235	-0.11	5.03	44.63	387
2016	264	1.31	1.90	34.66	84
2014	281	-0.05	1.82	27.23	76
more than one	189	-0.43	-0.98	29.87	-194
2013	224	-1.22	-3.48	25.94	-437

HVAC tabulation by intervention year, size

year	customer size	premise count	daily AC savings (% of pre)	daily AC savings (kWh)	daily AC savings std dev	% AC gain
All participants	L	694	0.27	1.52	38.31	0
All participants	M	327	0.66	0.96	26.65	0
All participants	N	53	-2.97	-1.77	34.70	0
All participants	S	119	-2.47	-0.36	6.37	0
2015	L	127	1.75	9.79	56.94	543
2016	L	185	0.43	2.90	34.53	90
2014	L	167	0.32	1.99	28.93	31
more than one	L	99	-0.32	-2.26	33.45	-248
2013	L	116	-1.20	-7.17	32.50	-571
2014	M	69	-0.50	2.69	21.81	180
2013	M	55	0.91	2.53	16.65	164
2015	M	73	-0.80	0.94	23.40	-2
2016	M	64	4.38	-0.50	39.00	-152
more than one	M	66	-0.31	-0.73	27.28	-176
2013	N	22	-2.01	-1.40	21.01	-21
2014	S	26	0.01	1.93	9.14	-636
more than one	S	22	-1.79	-0.02	5.90	-93
2015	S	25	-3.97	-1.36	6.47	279
2013	S	31	-4.54	-1.84	4.46	412

HVAC tabulation by intervention year, rate

year	rate	premise count	daily AC savings (% of pre)	daily AC savings (kWh)	daily AC savings std dev	% AC gain
All participants	All participants	1193	-0.04	1.03	33.24	0
2015	A10 Medium general demand	70	1.84	12.35	59.78	1094
2015	E19 Medium general demand TOU	37	1.40	10.06	65.62	873
2016	A10 Medium general demand	70	-0.29	5.01	57.90	385
2014	E19 Medium general demand TOU	49	0.40	3.63	36.56	251
2014	A10 Medium general demand	96	-0.04	2.50	35.18	141
2014	A1 Small general service	121	-0.12	1.32	12.21	28
2016	A1 Small general service	81	4.28	1.06	24.48	3
2016	E19 Medium general demand TOU	102	0.26	0.76	18.74	-26
2015	A1 Small general service	111	-1.68	0.40	18.89	-61
more than one	A1 Small general service	93	-1.24	-0.82	14.21	-179
2013	A1 Small general service	124	-1.42	-0.93	13.58	-190
2013	E19 Medium general demand TOU	43	-0.51	-1.72	28.12	-266
more than one	A10 Medium general demand	67	0.49	-2.72	42.21	-363
2013	A10 Medium general demand	51	-1.58	-11.45	42.22	-1208

HVAC tabulation by NAICS SECTOR abbreviation

Accommodation and Food Service has the largest average savings by far, followed by Public Administration and Retail Trade – 2. None of these has a very large premise count, however.

NAICS sector	premise count	daily AC savings (% of pre)	daily AC savings (kWh)	daily AC savings std dev	% AC gain
All participants	1193	-0.04	1.03	33.24	0
Accommodation and Food Service	49	2.23	11.49	29.30	1011
Public Administration	34	4.20	7.72	65.50	647
RETAIL TRADE - 2	39	-1.39	7.70	35.97	645
Arts, Entertainment, and Recreation …	31	-2.21	6.77	38.49	555

NAICS sector	premise count	daily AC savings (% of pre)	daily AC savings (kWh)	daily AC savings std dev	% AC gain
Other Services (except Public	75	-1.12	4.03	23.00	290
Finance and Insurance	58	1.24	3.05	13.66	195
Information	36	0.25	2.38	26.97	130
RETAIL TRADE - 1	173	0.27	2.10	32.94	103
N/A	103	2.47	2.03	26.03	97
Professional, Scientific …	28	-0.43	1.65	20.97	59
Real Estate and Rental and Leasing …	182	0.83	1.49	33.09	44
Construction	38	-4.31	0.68	14.80	-34
MANUFACTURING - 4	20	0.11	0.61	41.56	-41
Health Care and Social Assistance …	76	-2.20	-0.76	25.02	-174
Management of Companies …	45	-0.61	-4.66	30.65	-550
Educational Services	105	-0.73	-5.73	40.48	-654

HVAC tabulation by NAICS SECTOR abbreviation, size

Limited conclusions can be drawn from such small premise counts.

NAICS sector	customer size	premise count	daily AC savings (% of pre)	daily AC savings (kWh)	daily AC savings std dev	% AC gain
All participants	L	694	0.27	1.52	38.31	0
All participants	M	327	0.66	0.96	26.65	0
All participants	N	53	-2.97	-1.77	34.70	0
All participants	S	119	-2.47	-0.36	6.37	0
RETAIL TRADE - 2	L	29	-1.11	8.61	41.05	466
Public Administration	L	27	2.24	7.44	73.06	389
N/A	L	62	0.11	4.52	25.53	197
RETAIL TRADE - 1	L	136	0.76	4.19	33.99	175
Finance and Insurance	L	49	1.11	3.86	14.14	154
Information	L	35	0.65	2.74	27.28	80
Real Estate and Rental and Leasing	L	90	2.37	1.12	44.46	-27
Arts, Entertainment, and Recreation …	L	20	-2.58	-1.66	38.18	-209
Health Care and Social Assistance	L	47	-2.00	-2.31	30.53	-252
Management of Companies …	L	33	-0.48	-7.79	34.17	-612
Educational Services	L	67	-1.50	-11.46	44.57	-853

Accommodation and Food Service	M	22	4.11	12.93	28.34	1249
Educational Services	M	34	-0.39	3.76	31.47	292
Real Estate and Rental and Leasing …	M	60	-0.09	2.53	14.10	164
Other Services	M	34	-1.46	2.09	16.72	118
Construction	M	24	-3.31	1.73	16.32	81
N/A	M	28	16.05	0.71	31.30	-26
RETAIL TRADE - 1	M	22	-3.39	-12.33	33.19	-1386
Other Services	S	29	-0.13	0.57	5.16	-257

Appendix D: Tabulation of savings from consumption feature filters

In this appendix, savings results for all participants are compared to sub-groups defined by consumption feature filters (some applied within specific customer characteristics sub-groups) for both programs. Note that features were computed using the open source meter data analysis R package, VISDOM.[21]

How to read these tables

The format for each table in the section is as follows:

filter % - the percentage of customers eliminated by the feature filter criteria. For example, 75% indicates that the premises associated with the lower 3/4 of values for the feature in question were eliminated, leaving the top 25%.

<group characteristics> - An optional second column provides specific customer characteristic values that define a sub-group that filters were applied within. For example, the sub-groups could be based on S, M, L total consumption 'customer size' categories or rate types.

feature – The name of the feature whose values were used as criteria for filtering out premises.

premise count – The count of premises that are in the filtered sub-group for the row. Groups with fewer than 20 premises have been excluded as groups with the fewest premises are the most likely to have fluke results.

daily savings (% of pre) - The at-the-meter average estimated daily savings for the sub-group as a percentage of pre-intervention average daily total consumption.

daily savings – The at-the-meter average estimated daily savings in kWh for the sub-group.

% gain – The percentage improvement over the corresponding 'All participants' daily savings entry that a sub-group exhibits.

DI program filtering results

All DI filtering results presented here are in % gain terms. So a value of 10 means that the filtered group saved an average of 10% more than the whole population.

DI 90% filter depth top 25 filters

The top 25 feature filters at 90% filter depth in the DI program are all consumption and baseload related features, with the average consumption from 11pm to midnight delivering the strongest gains. % gains ranging from 306-336% are equivalent to multiplying the average savings across all participating premises by 4.06 - 4.36x.

filter %	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
90%	HOD_mean_24	750	7.06	62.80	336
90%	pre_intercept	750	6.54	62.31	333
90%	kw_mean_winter	750	6.49	62.06	331
90%	HOD_mean_4	750	6.93	61.24	326

[21] http://github.com/convergenceda/visdom

filter %	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
90%	kw_mean	750	6.27	60.97	324
90%	mean	750	6.27	60.97	324
90%	pre_daily_kwh	750	6.27	60.97	324
90%	min_day_kw	750	6.98	60.89	323
90%	kw_mean_summer	750	6.16	60.71	322
90%	kw_total_Jul	750	6.39	60.65	321
90%	kw_total_Apr	750	6.32	60.31	319
90%	kw_total_Aug	750	6.33	60.20	318
90%	min	750	6.59	60.16	318
90%	min_3	750	6.79	60.02	317
90%	Aug_mean	750	6.29	59.99	317
90%	HOD_mean_8	750	6.48	59.67	315
90%	kw_total_Mar	750	6.43	59.50	313
90%	kw_total_Jun	750	6.10	59.35	312
90%	kw_total_May	750	6.10	59.19	311
90%	kw_total_Feb	750	6.47	59.07	310
90%	summer_kwh	750	6.21	58.86	309
90%	HOD_mean_20	750	6.05	58.62	307
90%	max	750	6.12	58.53	307
90%	HOD_mean_12	750	6.20	58.53	307
90%	max_day_kw	750	5.93	58.47	306

DI 75% filter depth top 25 filters

The top 25 feature filters at 75% filter depth in the DI program are all consumption and baseload related features, with metrics of average/total consumption delivering the strongest gains. % gains ranging from 162-175% are equivalent to multiplying the average savings across all participating premises by 2.62 - 2.75x.

filter %	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
75%	kw_mean	1875	6.51	39.59	175
75%	mean	1875	6.51	39.59	175
75%	pre_daily_kwh	1875	6.51	39.59	175
75%	kw_mean_summer	1875	6.49	39.54	175
75%	pre_intercept	1875	6.42	39.26	173
75%	kw_mean_winter	1875	6.38	39.15	172
75%	min_day_kw	1875	6.91	39.13	172

filter %	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
75%	HOD_mean_4	1875	6.76	38.95	171
75%	Aug_mean	1875	6.62	38.91	170
75%	HOD_mean_24	1875	6.63	38.88	170
75%	HOD_mean_8	1875	6.67	38.86	170
75%	kw_total_Aug	1875	6.58	38.84	170
75%	kw_total_Jul	1875	6.56	38.80	170
75%	kw_total_Jun	1875	6.51	38.72	169
75%	min	1875	6.52	38.68	169
75%	HOD_mean_16	1875	6.36	38.60	168
75%	HOD_mean_20	1875	6.39	38.43	167
75%	min_3	1875	6.59	38.40	167
75%	max	1875	6.25	38.39	167
75%	max_day_kw	1875	6.13	38.30	166
75%	kw_total_May	1875	6.36	38.19	165
75%	HOD_mean_12	1875	6.31	38.17	165
75%	kw_total_Apr	1875	6.30	38.02	164
75%	Aug_max	1875	6.46	38.01	164
75%	kw_total_Jan	1875	6.56	37.71	162

DI 50% filter depth top 25 filters

The top 25 feature filters at 50% filter depth in the DI program are all consumption and baseload related features, with metrics related to non-weather-responsive loads and maximum, mid-day, and average daily consumption delivering the strongest gains. % gains ranging from 69-76% are equivalent to multiplying the average savings across all participating premises by 1.69 – 1.76x.

filter %	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
50%	pre_intercept	3749	6.37	25.27	76
50%	max	3749	6.46	25.24	75
50%	HOD_mean_12	3749	6.56	25.20	75
50%	kw_mean_winter	3749	6.29	25.16	75
50%	kw_mean	3749	6.24	25.13	75
50%	mean	3749	6.24	25.13	75
50%	pre_daily_kwh	3749	6.24	25.13	75
50%	discretionary	3749	6.64	25.08	74
50%	HOD_mean_16	3749	6.48	25.04	74

filter %	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
50%	kw_mean_summer	3749	6.09	25.02	74
50%	max_97	3749	6.42	24.99	74
50%	min_day_kw	3749	6.28	24.95	73
50%	max_day_kw	3749	6.17	24.92	73
50%	range	3749	6.87	24.90	73
50%	min_3	3749	6.04	24.78	72
50%	HOD_mean_8	3749	6.24	24.72	72
50%	HOD_mean_4	3749	6.16	24.63	71
50%	max_hr_kw	3749	6.32	24.55	71
50%	HOD_mean_24	3749	6.07	24.49	70
50%	HOD_mean_20	3749	6.01	24.47	70
50%	Aug_max	3749	6.69	24.41	70
50%	kw_total_Jun	3749	6.31	24.40	70
50%	kw_total_Jul	3749	6.34	24.32	69
50%	kw_total_May	3749	6.18	24.29	69
50%	min	3749	5.80	24.29	69

DI lighting 90% filter depth top 25 filters

The top 25 feature filters at 90% filter depth in the lighting only projects within the DI program deliver % gains ranging from 338-380% are equivalent to multiplying the average savings across all participating premises by 4.38 – 4.80x.

filter %	tech family	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
90%	LIGHTING	pre_intercept	453	8.20	69.04	380
90%	LIGHTING	kw_mean_winter	453	8.21	68.94	379
90%	LIGHTING	kw_mean	453	7.82	67.20	367
90%	LIGHTING	mean	453	7.82	67.20	367
90%	LIGHTING	pre_daily_kwh	453	7.82	67.20	367
90%	LIGHTING	HOD_mean_24	453	8.23	66.95	365
90%	LIGHTING	kw_mean_summer	453	7.35	66.13	360
90%	LIGHTING	min_day_kw	453	8.33	66.03	359
90%	LIGHTING	kw_total_Apr	453	7.81	65.99	359
90%	LIGHTING	Aug_mean	453	7.48	65.43	355
90%	LIGHTING	kw_total_Aug	453	7.45	65.32	354
90%	LIGHTING	max	453	7.78	65.19	353
90%	LIGHTING	kw_total_Mar	453	7.95	65.13	353

filter %	tech family	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
90%	LIGHTING	HOD_mean_20	453	7.67	65.07	352
90%	LIGHTING	kw_total_Dec	453	7.97	65.03	352
90%	LIGHTING	kw_total_Jun	453	7.46	64.91	351
90%	LIGHTING	kw_total_Jan	453	8.01	64.72	350
90%	LIGHTING	kw_total_Jul	453	7.48	64.58	349
90%	LIGHTING	kw_total_May	453	7.42	64.30	347
90%	LIGHTING	HOD_mean_12	453	7.61	64.27	347
90%	LIGHTING	kw_total_Feb	453	7.96	64.22	346
90%	LIGHTING	HOD_mean_16	453	7.60	64.20	346
90%	LIGHTING	summer_kwh	453	7.38	63.62	342
90%	LIGHTING	HOD_mean_4	453	7.67	63.24	339
90%	LIGHTING	min_3	453	7.67	63.06	338

DI lighting 75% filter depth top 25 filters

The top 25 feature filters at 75% filter depth in the lighting only projects within the DI program deliver % gains ranging from 160-171% are equivalent to multiplying the average savings across all participating premises by 2.60 – 2.71x.

filter %	tech family	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
75%	LIGHTING	pre_intercept	1131	7.36	39.05	171
75%	LIGHTING	kw_mean_winter	1131	7.28	38.87	170
75%	LIGHTING	kw_mean_summer	1131	7.09	38.85	170
75%	LIGHTING	kw_mean	1131	7.17	38.78	169
75%	LIGHTING	mean	1131	7.17	38.78	169
75%	LIGHTING	pre_daily_kwh	1131	7.17	38.78	169
75%	LIGHTING	max	1131	7.29	38.52	168
75%	LIGHTING	HOD_mean_16	1131	7.42	38.39	167
75%	LIGHTING	HOD_mean_12	1131	7.39	38.35	167
75%	LIGHTING	kw_total_Aug	1131	7.26	38.25	166
75%	LIGHTING	Aug_mean	1131	7.21	38.14	165
75%	LIGHTING	min_day_kw	1131	7.49	38.05	164
75%	LIGHTING	kw_total_Jun	1131	7.29	38.04	164
75%	LIGHTING	discretionary	1131	7.61	37.94	164
75%	LIGHTING	max_97	1131	7.10	37.82	163
75%	LIGHTING	HOD_mean_4	1131	7.41	37.71	162
75%	LIGHTING	Aug_max	1131	7.08	37.67	162

filter %	tech family	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
75%	LIGHTING	HOD_mean_24	1131	7.31	37.66	162
75%	LIGHTING	kw_total_May	1131	7.17	37.63	161
75%	LIGHTING	kw_total_Jul	1131	7.10	37.62	161
75%	LIGHTING	min_3	1131	7.10	37.53	161
75%	LIGHTING	kw_total_Apr	1131	7.06	37.53	161
75%	LIGHTING	HOD_mean_20	1131	7.02	37.51	161
75%	LIGHTING	HOD_mean_8	1131	7.09	37.48	160
75%	LIGHTING	max_day_kw	1131	6.68	37.45	160

DI lighting 50% filter depth top 25 filters

The top 25 feature filters at 50% filter depth in the lighting only projects within the DI program deliver % gains ranging from 64-70% are equivalent to multiplying the average savings across all participating premises by 1.64 – 1.70x.

filter %	tech family	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
50%	LIGHTING	kw_mean_summer	2261	7.79	24.52	70
50%	LIGHTING	pre_intercept	2261	7.91	24.48	70
50%	LIGHTING	kw_mean_winter	2261	7.83	24.39	70
50%	LIGHTING	max	2261	7.98	24.36	69
50%	LIGHTING	range	2261	8.38	24.36	69
50%	LIGHTING	discretionary	2261	8.21	24.30	69
50%	LIGHTING	min_day_kw	2261	7.96	24.27	69
50%	LIGHTING	kw_mean	2261	7.63	24.26	69
50%	LIGHTING	mean	2261	7.63	24.26	69
50%	LIGHTING	pre_daily_kwh	2261	7.63	24.26	69
50%	LIGHTING	HOD_mean_12	2261	7.97	24.16	68
50%	LIGHTING	max_97	2261	7.79	24.11	68
50%	LIGHTING	kw_total_Jun	2261	8.18	24.00	67
50%	LIGHTING	HOD_mean_16	2261	7.93	23.99	67
50%	LIGHTING	max_day_kw	2261	7.40	23.96	67
50%	LIGHTING	HOD_mean_20	2261	7.84	23.90	66
50%	LIGHTING	min_3	2261	7.45	23.86	66
50%	LIGHTING	kw_total_May	2261	8.09	23.83	66
50%	LIGHTING	kw_total_Aug	2261	7.94	23.83	66
50%	LIGHTING	kw_total_Jul	2261	7.99	23.82	66
50%	LIGHTING	Aug_mean	2261	7.91	23.81	65

filter %	tech family	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
50%	LIGHTING	Aug_max	2261	8.14	23.80	65
50%	LIGHTING	kw_total_Apr	2261	8.06	23.69	65
50%	LIGHTING	Aug_range	2261	8.54	23.64	64
50%	LIGHTING	max_hr_kw	2261	7.17	23.55	64

DI refrigeration 90% filter depth top 25 filters

The top 25 feature filters at 90% filter depth in the refrigeration only projects within the DI program deliver % gains ranging from 144-212% are equivalent to multiplying the average savings across all participating premises by 2.44 – 3.12x.

filter %	tech family	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
90%	REFRIGERATION	max_hr_kw	78	4.03	44.85	212
90%	REFRIGERATION	HOD_mean_20	78	3.53	42.74	197
90%	REFRIGERATION	max_97	78	3.73	42.26	194
90%	REFRIGERATION	Aug_max	78	3.56	41.15	186
90%	REFRIGERATION	kw_total_Nov	78	3.25	40.11	179
90%	REFRIGERATION	HOD_mean_8	78	3.36	39.84	177
90%	REFRIGERATION	kw_total_Feb	78	3.12	39.29	173
90%	REFRIGERATION	max	78	3.09	38.59	168
90%	REFRIGERATION	max_day_kw	78	3.26	38.57	168
90%	REFRIGERATION	kw_total_Dec	78	3.08	38.52	168
90%	REFRIGERATION	kw_total_Mar	78	2.99	38.03	164
90%	REFRIGERATION	pre_intercept	78	2.90	37.62	161
90%	REFRIGERATION	kw_mean_winter	78	2.87	37.37	160
90%	REFRIGERATION	range	78	3.76	36.68	155
90%	REFRIGERATION	kw_total_Sep	78	2.63	36.57	154
90%	REFRIGERATION	kw_total_Oct	78	2.88	36.54	154
90%	REFRIGERATION	summer_kwh	78	2.79	36.39	153
90%	REFRIGERATION	HOD_mean_16	78	2.74	36.06	151
90%	REFRIGERATION	kw_mean	78	2.68	35.89	149
90%	REFRIGERATION	mean	78	2.68	35.89	149
90%	REFRIGERATION	pre_daily_kwh	78	2.68	35.89	149
90%	REFRIGERATION	min	78	2.83	35.86	149
90%	REFRIGERATION	kw_total_Jun	78	2.67	35.61	147
90%	REFRIGERATION	kw_total_Apr	78	2.66	35.46	146
90%	REFRIGERATION	kw_total_Aug	78	2.59	35.09	144

DI refrigeration 75% filter depth top 25 filters

The top 25 feature filters at 75% filter depth in the refrigeration only projects within the DI program deliver % gains ranging from 77-89% are equivalent to multiplying the average savings across all participating premises by 1.77 – 1.89x.

filter %	tech family	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
75%	REFRIGERATION	kw_total_Sep	194	2.95	27.27	89
75%	REFRIGERATION	Aug_min	194	3.27	26.96	87
75%	REFRIGERATION	pre_intercept	194	3.06	26.93	87
75%	REFRIGERATION	HOD_mean_24	194	3.23	26.68	85
75%	REFRIGERATION	kw_mean_winter	194	3.01	26.65	85
75%	REFRIGERATION	Aug_max	194	2.99	26.48	84
75%	REFRIGERATION	kw_total_Aug	194	2.95	26.42	84
75%	REFRIGERATION	kw_total_Mar	194	2.97	26.31	83
75%	REFRIGERATION	Aug_range	194	3.13	26.13	82
75%	REFRIGERATION	min	194	3.13	26.12	81
75%	REFRIGERATION	kw_total_Jun	194	2.87	26.10	81
75%	REFRIGERATION	kw_total_Jul	194	2.89	26.10	81
75%	REFRIGERATION	min_3	194	3.19	26.09	81
75%	REFRIGERATION	kw_total_Oct	194	2.87	25.99	81
75%	REFRIGERATION	summer_kwh	194	2.89	25.90	80
75%	REFRIGERATION	max	194	2.91	25.85	80
75%	REFRIGERATION	kw_mean_summer	194	2.81	25.77	79
75%	REFRIGERATION	kw_mean	194	2.79	25.74	79
75%	REFRIGERATION	mean	194	2.79	25.74	79
75%	REFRIGERATION	pre_daily_kwh	194	2.79	25.74	79
75%	REFRIGERATION	kw_total_Nov	194	2.76	25.73	79
75%	REFRIGERATION	Aug_mean	194	2.81	25.73	79
75%	REFRIGERATION	max_97	194	2.90	25.66	78
75%	REFRIGERATION	kw_total_Feb	194	2.86	25.62	78
75%	REFRIGERATION	kw_total_Dec	194	2.96	25.48	77

DI refrigeration 50% filter depth top 25 filters

The top 25 feature filters at 50% filter depth in the refrigeration only projects within the DI program deliver % gains ranging from 17-27% are equivalent to multiplying the average savings across all participating premises by 1.17 – 1.27x.

filter %	tech family	feature	premise count	daily savings (% of pre)	daily savings(kWh)	% gain
50%	REFRIGERATION	summer_kwh	387	3.11	18.24	27
50%	REFRIGERATION	kw_total_Jul	387	2.98	18.17	26
50%	REFRIGERATION	HOD_mean_4	387	3.08	18.14	26
50%	REFRIGERATION	Aug_min	387	3.09	18.13	26
50%	REFRIGERATION	kw_total_Aug	387	2.92	18.04	25
50%	REFRIGERATION	Aug_mean	387	2.92	18.04	25
50%	REFRIGERATION	min	387	3.00	17.85	24
50%	REFRIGERATION	Aug_max	387	2.94	17.85	24
50%	REFRIGERATION	pre_intercept	387	2.79	17.67	23
50%	REFRIGERATION	kw_total_Jun	387	2.67	17.56	22
50%	REFRIGERATION	min_3	387	2.85	17.54	22
50%	REFRIGERATION	kw_mean_winter	387	2.72	17.52	22
50%	REFRIGERATION	kw_mean	387	2.68	17.41	21
50%	REFRIGERATION	mean	387	2.68	17.41	21
50%	REFRIGERATION	pre_daily_kwh	387	2.68	17.41	21
50%	REFRIGERATION	kw_total_Mar	387	2.81	17.34	21
50%	REFRIGERATION	max	387	2.67	17.24	20
50%	REFRIGERATION	kw_total_May	387	2.77	17.23	20
50%	REFRIGERATION	kw_total_Sep	387	2.58	17.23	20
50%	REFRIGERATION	HOD_mean_24	387	2.65	17.22	20
50%	REFRIGERATION	HOD_mean_20	387	2.56	17.13	19
50%	REFRIGERATION	min_day_kw	387	2.72	16.99	18
50%	REFRIGERATION	Aug_range	387	2.67	16.92	18
50%	REFRIGERATION	kw_total_Apr	387	2.60	16.89	17
50%	REFRIGERATION	HOD_mean_16	387	2.48	16.87	17

DI size 90% filter depth top 10 filters per category

The top 10 feature filters per customer size category at 90% filter depth for the DI program exhibit % gains

For L ranging from 689-715% or 7.89 – 8.15x average program savings.
For M ranging from 349-378% or 4.49 – 4.78x average program savings, number that eclipse the average savings of size L customers.
For S ranging from 29-35% or 0.84 – 1.35x average program savings. It is an impressive feat for the smallest customers to outperform the program wide average savings.

filter %	customer size	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
90%	L	Aug_mean	103	6.62	117.29	715
90%	L	pre_intercept	103	6.46	116.64	710
90%	L	HOD_mean_20	103	6.51	116.40	709
90%	L	kw_mean_summer	103	6.53	116.32	708
90%	L	kw_total_Aug	103	6.48	115.76	704
90%	L	kw_total_May	103	6.61	115.49	702
90%	L	HOD_mean_12	103	6.57	113.89	691
90%	L	kw_mean	103	6.23	113.62	689
90%	L	mean	103	6.23	113.62	689
90%	L	pre_daily_kwh	103	6.23	113.62	689
90%	M	pre_intercept	341	8.46	68.74	378
90%	M	kw_mean_winter	341	8.41	68.38	375
90%	M	HOD_mean_24	341	8.87	67.68	370
90%	M	kw_mean	341	8.13	67.20	367
90%	M	mean	341	8.13	67.20	367
90%	M	pre_daily_kwh	341	8.13	67.20	367
90%	M	HOD_mean_4	341	9.04	67.05	366
90%	M	kw_total_Mar	341	8.24	66.23	360
90%	M	min_day_kw	341	8.63	64.64	349
90%	M	kw_total_Feb	341	8.19	64.64	349
90%	S	discretionary	260	16.30	19.41	35
90%	S	kw_total_Jan	260	15.12	19.39	35
90%	S	pre_intercept	260	15.21	19.28	34
90%	S	kw_mean	260	14.94	19.25	34
90%	S	mean	260	14.94	19.25	34
90%	S	pre_daily_kwh	260	14.94	19.25	34
90%	S	kw_mean_winter	260	15.07	19.12	33
90%	S	kw_total_Feb	260	14.37	18.74	30
90%	S	kw_total_Mar	260	14.10	18.69	30
90%	S	HOD_mean_16	260	16.34	18.51	29

DI size 75% filter depth top filters 10 per category

The top 10 feature filters per customer size category at 75% filter depth for the DI program exhibit % gains

For L ranging from 397-422% or 4.97 – 5.22x average program savings.
For M ranging from 203-219% or 3.03 – 3.19x average program savings.

For S ranging from -16 to -14% or 0.84 – 0.86x average program savings.

filter %	customer size	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
75%	L	kw_total_Jun	257	5.79	75.08	422
75%	L	kw_mean_summer	257	5.49	73.80	413
75%	L	kw_total_Jul	257	5.71	73.07	408
75%	L	kw_total_Aug	257	5.66	73.05	408
75%	L	Aug_mean	257	5.66	73.05	408
75%	L	max	257	5.45	72.20	402
75%	L	kw_mean	257	5.34	72.12	401
75%	L	mean	257	5.34	72.12	401
75%	L	pre_daily_kwh	257	5.34	72.12	401
75%	L	Aug_max	257	5.59	71.48	397
75%	M	kw_mean_winter	851	7.73	45.87	219
75%	M	pre_intercept	851	7.68	45.59	217
75%	M	kw_mean	851	7.56	45.21	214
75%	M	mean	851	7.56	45.21	214
75%	M	pre_daily_kwh	851	7.56	45.21	214
75%	M	HOD_mean_4	851	8.10	44.85	212
75%	M	kw_mean_summer	851	7.21	43.97	206
75%	M	min	851	7.71	43.97	206
75%	M	max	851	7.28	43.69	204
75%	M	min_day_kw	851	7.98	43.60	203
75%	S	kw_total_May	649	12.22	12.42	-14
75%	S	max	649	12.54	12.40	-14
75%	S	kw_mean	649	11.81	12.34	-14
75%	S	mean	649	11.81	12.34	-14
75%	S	pre_daily_kwh	649	11.81	12.34	-14
75%	S	HOD_mean_12	649	12.81	12.24	-15
75%	S	discretionary	649	12.61	12.13	-16
75%	S	summer_kwh	649	12.20	12.13	-16
75%	S	kw_mean_summer	649	11.78	12.13	-16
75%	S	pre_intercept	649	11.55	12.10	-16

DI size 50% filter depth top 10 filters per category

The top 10 feature filters per customer size category at 50% filter depth for the DI program exhibit % gains

For L ranging from 229-245% or 3.29 – 3.45x average program savings.

For M ranging from 104-110% or 2.04 – 2.10x average program savings.
For S ranging from -38 to -37% or 0.62 – 0.63x average program savings.

filter %	customer size	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
50%	L	kw_total_Jun	513	5.53	49.69	245
50%	L	kw_total_Jul	513	5.34	48.38	236
50%	L	min_day_kw	513	5.46	47.90	233
50%	L	Aug_mean	513	5.36	47.80	232
50%	L	kw_total_Aug	513	5.31	47.64	231
50%	L	Aug_max	513	5.43	47.63	231
50%	L	kw_mean_summer	513	4.91	47.59	231
50%	L	summer_kwh	513	5.38	47.36	229
50%	L	max	513	5.04	47.34	229
50%	L	Aug_min	513	5.27	47.33	229
50%	M	pre_intercept	1702	6.78	30.29	110
50%	M	kw_mean	1702	6.64	30.05	109
50%	M	mean	1702	6.64	30.05	109
50%	M	pre_daily_kwh	1702	6.64	30.05	109
50%	M	kw_mean_winter	1702	6.66	30.03	109
50%	M	kw_mean_summer	1702	6.61	29.90	108
50%	M	min_day_kw	1702	7.10	29.87	108
50%	M	min_3	1702	6.88	29.49	105
50%	M	HOD_mean_24	1702	6.62	29.39	104
50%	M	max	1702	6.37	29.33	104
50%	S	discretionary	1298	12.24	9.05	-37
50%	S	HOD_mean_16	1298	12.42	8.97	-38
50%	S	kw_total_May	1298	11.64	8.94	-38
50%	S	max	1298	11.78	8.94	-38
50%	S	kw_mean	1298	11.45	8.94	-38
50%	S	mean	1298	11.45	8.94	-38
50%	S	pre_daily_kwh	1298	11.45	8.94	-38
50%	S	kw_mean_summer	1298	11.50	8.90	-38
50%	S	range	1298	12.15	8.89	-38
50%	S	HOD_mean_12	1298	12.08	8.89	-38

DI rate 90% filter depth top 10 filters per category

The top 10 feature filters per customer rate type at 90% filter depth for the DI program exhibit % gains

For E19 ranging from 795-855% or 8.95-9.55x average program savings.

For A10 ranging from 482-505% or 5.82-6.05x average program savings.
For A6 ranging from 354-377% or 4.54-4.77x average program savings.
For A1 ranging from 13-16% or 1.13 – 1.16x average program savings.

filter %	rate	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
90%	E19 Medium general demand TOU	HOD_mean_4	73	7.76	137.42	855
90%	E19 Medium general demand TOU	max_day_kw	73	7.55	137.35	854
90%	E19 Medium general demand TOU	pre_intercept	73	7.15	133.03	824
90%	E19 Medium general demand TOU	HOD_mean_24	73	7.53	132.65	822
90%	E19 Medium general demand TOU	Aug_max	73	7.29	130.19	805
90%	E19 Medium general demand TOU	max_97	73	7.09	130.10	804
90%	E19 Medium general demand TOU	kw_mean	73	6.84	129.64	801
90%	E19 Medium general demand TOU	mean	73	6.84	129.64	801
90%	E19 Medium general demand TOU	pre_daily_kwh	73	6.84	129.64	801
90%	E19 Medium general demand TOU	kw_total_Jan	73	6.89	128.85	795
90%	A10 Medium general demand	Aug_max	164	7.96	87.12	505
90%	A10 Medium general demand	kw_total_Jun	164	7.77	86.71	503
90%	A10 Medium general demand	summer_kwh	164	7.74	86.51	501
90%	A10 Medium general demand	kw_total_May	164	7.53	85.72	496
90%	A10 Medium general demand	Aug_mean	164	7.41	84.81	489
90%	A10 Medium general demand	pre_intercept	164	7.37	84.43	487
90%	A10 Medium general demand	kw_total_Aug	164	7.32	84.28	486
90%	A10 Medium general demand	kw_mean_summer	164	7.22	84.11	484
90%	A10 Medium general demand	kw_total_Jul	164	7.43	84.04	484
90%	A10 Medium general demand	kw_mean	164	7.23	83.80	482
90%	A6 Small general service TOU	kw_total_Dec	28	10.61	68.72	377
90%	A6 Small general service TOU	min_3	28	11.55	68.25	374
90%	A6 Small general service TOU	min_day_kw	28	11.16	67.71	370
90%	A6 Small general service TOU	HOD_mean_24	28	11.22	67.60	370
90%	A6 Small general service TOU	kw_mean_winter	28	10.21	67.31	368
90%	A6 Small general service TOU	kw_total_Jan	28	10.14	66.73	364
90%	A6 Small general service TOU	pre_intercept	28	9.91	65.99	359
90%	A6 Small general service TOU	HOD_mean_4	28	10.05	65.62	356
90%	A6 Small general service TOU	kw_total_Feb	28	10.23	65.60	356
90%	A6 Small general service TOU	kw_total_Nov	28	9.98	65.39	354
90%	A1 Small general service	pre_intercept	477	6.28	16.71	16
90%	A1 Small general service	HOD_mean_16	477	6.64	16.56	15
90%	A1 Small general service	kw_mean_winter	477	6.14	16.52	15

filter %	rate	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
90%	A1 Small general service	HOD_mean_12	477	6.64	16.52	15
90%	A1 Small general service	kw_mean	477	6.05	16.51	15
90%	A1 Small general service	mean	477	6.05	16.51	15
90%	A1 Small general service	pre_daily_kwh	477	6.05	16.51	15
90%	A1 Small general service	kw_total_Jan	477	6.64	16.35	14
90%	A1 Small general service	kw_total_Jun	477	6.22	16.28	13
90%	A1 Small general service	HOD_mean_8	477	6.80	16.25	13

DI rate 75% filter depth top 10 filters per category

The top 10 feature filters per customer rate type at 75% filter depth for the DI program exhibit % gains

For E19 ranging from 500-526% or 6.00-6.26x average program savings.
For A10 ranging from 310-329% or 4.10-4.29x average program savings.
For A6 ranging from 176-217% or 2.76-3.17x average program savings.
For A1 ranging from -13 to -10% or 0.87 – 0.90x average program savings.

filter %	rate	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
75%	E19 Medium general demand TOU	kw_mean	181	6.66	90.11	526
75%	E19 Medium general demand TOU	mean	181	6.66	90.11	526
75%	E19 Medium general demand TOU	pre_daily_kwh	181	6.66	90.11	526
75%	E19 Medium general demand TOU	pre_intercept	181	6.50	88.62	516
75%	E19 Medium general demand TOU	kw_mean_summer	181	6.43	88.11	512
75%	E19 Medium general demand TOU	HOD_mean_24	181	6.59	87.52	508
75%	E19 Medium general demand TOU	kw_mean_winter	181	6.30	87.17	506
75%	E19 Medium general demand TOU	kw_total_Jan	181	6.59	87.07	505
75%	E19 Medium general demand TOU	HOD_mean_20	181	6.35	86.71	503
75%	E19 Medium general demand TOU	kw_total_Dec	181	6.67	86.39	500
75%	A10 Medium general demand	HOD_mean_24	408	7.78	61.79	329
75%	A10 Medium general demand	kw_total_Apr	408	7.22	61.51	327
75%	A10 Medium general demand	min_day_kw	408	7.90	60.26	319
75%	A10 Medium general demand	pre_intercept	408	7.10	60.16	318
75%	A10 Medium general demand	kw_total_Jul	408	7.14	59.84	316
75%	A10 Medium general demand	kw_mean_winter	408	6.97	59.73	315
75%	A10 Medium general demand	kw_total_Aug	408	7.05	59.32	312
75%	A10 Medium general demand	Aug_mean	408	7.03	59.25	312
75%	A10 Medium general demand	kw_total_Mar	408	7.06	59.08	311

filter %	rate	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
75%	A10 Medium general demand	Aug_min	408	7.57	58.98	310
75%	A6 Small general service TOU	HOD_mean_4	69	11.24	45.58	217
75%	A6 Small general service TOU	kw_total_Dec	69	10.39	44.57	210
75%	A6 Small general service TOU	HOD_mean_24	69	10.68	44.51	209
75%	A6 Small general service TOU	min_day_kw	69	9.88	44.15	207
75%	A6 Small general service TOU	min_3	69	9.60	41.40	188
75%	A6 Small general service TOU	kw_total_Nov	69	9.56	41.18	186
75%	A6 Small general service TOU	min	69	9.09	40.37	181
75%	A6 Small general service TOU	kw_mean_winter	69	8.75	40.15	179
75%	A6 Small general service TOU	HOD_mean_8	69	8.97	40.08	179
75%	A6 Small general service TOU	HOD_mean_20	69	9.03	39.79	176
75%	A1 Small general service	pre_intercept	1192	6.73	12.94	-10
75%	A1 Small general service	HOD_mean_12	1192	7.12	12.91	-10
75%	A1 Small general service	max	1192	6.94	12.84	-11
75%	A1 Small general service	kw_mean_winter	1192	6.65	12.82	-11
75%	A1 Small general service	kw_total_Jan	1192	7.04	12.70	-12
75%	A1 Small general service	HOD_mean_16	1192	7.02	12.68	-12
75%	A1 Small general service	kw_mean	1192	6.43	12.61	-12
75%	A1 Small general service	mean	1192	6.43	12.61	-12
75%	A1 Small general service	pre_daily_kwh	1192	6.43	12.61	-12
75%	A1 Small general service	discretionary	1192	7.01	12.47	-13

DI rate 50% filter depth top 10 filters per category

The top 10 feature filters per customer rate type at 50% filter depth for the DI program exhibit % gains

For E19 ranging from 319-328% or 4.19 – 4.28x average program savings.
For A10 ranging from 211-218% or 3.11-3.18x average program savings.
For A6 ranging from 98-103% or 1.98-2.03x average program savings.
For A1 ranging from -35 to -33% or 0.65 - 0.67x average program savings.

filter %	rate	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
50%	E19 Medium general demand TOU	HOD_mean_24	361	6.11	61.63	328
50%	E19 Medium general demand TOU	pre_intercept	361	5.79	60.68	322
50%	E19 Medium general demand TOU	HOD_mean_20	361	5.81	60.68	322
50%	E19 Medium general demand TOU	kw_mean_summer	361	5.75	60.47	320

filter %	rate	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
50%	E19 Medium general demand TOU	max_day_kw	361	5.74	60.41	320
50%	E19 Medium general demand TOU	kw_mean_winter	361	5.76	60.41	320
50%	E19 Medium general demand TOU	max	361	5.73	60.27	319
50%	E19 Medium general demand TOU	kw_mean	361	5.70	60.25	319
50%	E19 Medium general demand TOU	mean	361	5.70	60.25	319
50%	E19 Medium general demand TOU	pre_daily_kwh	361	5.70	60.25	319
50%	A10 Medium general demand	Aug_mean	816	7.40	45.70	218
50%	A10 Medium general demand	kw_total_Aug	816	7.39	45.70	218
50%	A10 Medium general demand	pre_intercept	816	7.24	45.63	217
50%	A10 Medium general demand	kw_mean_summer	816	6.96	45.22	214
50%	A10 Medium general demand	kw_mean	816	7.06	45.11	213
50%	A10 Medium general demand	mean	816	7.06	45.11	213
50%	A10 Medium general demand	pre_daily_kwh	816	7.06	45.11	213
50%	A10 Medium general demand	kw_mean_winter	816	7.07	45.10	213
50%	A10 Medium general demand	kw_total_May	816	6.99	44.92	212
50%	A10 Medium general demand	kw_total_Jun	816	6.99	44.73	211
50%	A6 Small general service TOU	HOD_mean_4	138	9.11	29.18	103
50%	A6 Small general service TOU	pre_intercept	138	8.80	29.15	103
50%	A6 Small general service TOU	kw_mean_winter	138	8.74	29.08	102
50%	A6 Small general service TOU	min_day_kw	138	8.51	28.94	101
50%	A6 Small general service TOU	kw_mean	138	8.63	28.91	101
50%	A6 Small general service TOU	mean	138	8.63	28.91	101
50%	A6 Small general service TOU	pre_daily_kwh	138	8.63	28.91	101
50%	A6 Small general service TOU	max	138	8.86	28.76	100
50%	A6 Small general service TOU	max_97	138	8.87	28.48	98
50%	A6 Small general service TOU	kw_total_Mar	138	8.90	28.47	98
50%	A1 Small general service	discretionary	2384	7.79	9.65	-33
50%	A1 Small general service	max	2384	7.48	9.65	-33
50%	A1 Small general service	range	2384	8.04	9.58	-33
50%	A1 Small general service	HOD_mean_12	2384	7.63	9.57	-33
50%	A1 Small general service	kw_mean_winter	2384	7.12	9.54	-34
50%	A1 Small general service	kw_total_Apr	2384	7.45	9.50	-34
50%	A1 Small general service	pre_intercept	2384	7.05	9.48	-34
50%	A1 Small general service	kw_total_Jul	2384	7.39	9.45	-34
50%	A1 Small general service	max_97	2384	7.47	9.45	-34

filter %	rate	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
50%	A1 Small general service	HOD_mean_16	2384	7.54	9.42	-35

HVAC program filtering results

Note that because the HVAC program was relatively light touch, the unfiltered population savings were difficult to pull out from the background noise of other changes in consumption. The point estimate for the program mean is highly influenced by outliers. This does not mean that the program didn't save energy. It means that the savings were too small compared to the background noise to be accurately quantified using NMEC savings methods. **As a result of the near zero NMEC savings, the % gain values presented are very large, up to 52x the average savings of all participating premises.**

HVAC 90% filter depth top 25 filters

The top 25 feature filters at 90% filter depth in the HVAC program deliver daily savings ranging from 38.40-53.69 kWh/day are equivalent to multiplying the average savings across all participating premises by 37.2 – 51.9x.

filter %	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
90%	pre_CDH	120	3.12	53.69	5094
90%	Aug_mean	120	1.45	53.04	5031
90%	HOD_mean_4	120	0.64	50.86	4821
90%	kw_total_Aug	120	1.46	49.98	4735
90%	HOD_mean_20	120	1.43	49.83	4720
90%	kw_total_Sep	120	1.28	48.89	4630
90%	max	120	0.76	47.44	4489
90%	kw_mean	120	1.10	45.73	4324
90%	mean	120	1.10	45.73	4324
90%	pre_daily_kwh	120	1.10	45.73	4324
90%	kw_total_Apr	120	0.56	45.58	4310
90%	pre_intercept	120	0.44	45.20	4273
90%	min	120	0.36	44.63	4218
90%	pre_daily_cooling_kwh	120	2.75	44.13	4170
90%	Aug_min	120	0.85	43.24	4084
90%	Aug_max	120	0.81	42.06	3969
90%	summer_kwh	120	1.48	41.60	3924
90%	kw_total_Oct	120	1.14	41.39	3904
90%	max_day_kw	120	0.68	39.82	3753

filter %	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
90%	kw_total_Jul	120	1.24	39.46	3718
90%	HOD_mean_12	120	0.44	39.08	3681
90%	kw_total_Nov	120	0.80	39.07	3680
90%	kw_total_Jun	120	1.32	38.97	3670
90%	kw_mean_summer	120	1.22	38.51	3625
90%	kw_total_May	120	1.05	38.40	3615

HVAC 75% filter depth top 25 filters

The top 25 feature filters at 75% filter depth in the HVAC program deliver daily savings ranging from 21.84-28.66kWh/day are equivalent to multiplying the average savings across all participating premises by 21.1 – 27.7x.

filter %	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
75%	pre_CDH	299	2.40	28.66	2672
75%	Aug_mean	299	0.54	27.49	2559
75%	kw_total_Aug	299	0.56	27.35	2546
75%	pre_daily_cooling_kwh	299	1.93	27.01	2513
75%	kw_total_Oct	299	0.44	26.03	2418
75%	kw_mean_summer	299	0.46	25.68	2385
75%	HOD_mean_8	299	0.01	25.21	2339
75%	kw_mean	299	0.37	24.87	2306
75%	mean	299	0.37	24.87	2306
75%	pre_daily_kwh	299	0.37	24.87	2306
75%	summer_kwh	299	0.64	24.46	2266
75%	kw_total_Jun	299	0.66	24.37	2258
75%	kw_total_May	299	0.26	23.91	2213
75%	pre_intercept	299	0.18	23.87	2209
75%	HOD_mean_16	299	0.33	23.39	2163
75%	kw_total_Nov	299	0.45	23.33	2157
75%	HOD_mean_20	299	0.34	23.29	2153
75%	Aug_max	299	0.45	23.12	2137
75%	kw_total_Sep	299	0.53	23.06	2130
75%	HOD_mean_24	299	0.36	23.01	2126
75%	Aug_min	299	0.63	22.74	2100
75%	kw_mean_winter	299	0.20	22.52	2079
75%	max	299	0.05	22.33	2060

filter %	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
75%	kw_total_Jul	299	0.86	22.05	2033
75%	kw_tout_cor	299	2.25	21.84	2012

HVAC 50% filter depth top 25 filters

The top 25 feature filters at 50% filter depth in the HVAC program deliver daily savings ranging from 11.93-13.83kWh/day are equivalent to multiplying the average savings across all participating premises by 11.5 –13.4x.

filter %	feature	premise count	daily savings (% of pre)	daily savings (kWh)	% gain
50%	kw_total_Mar	597	0.38	13.83	1238
50%	Aug_range	597	0.44	13.81	1236
50%	pre_daily_cooling_kwh	597	1.53	13.79	1234
50%	pre_CDH	597	2.22	13.64	1219
50%	kw_total_Oct	597	0.55	13.52	1208
50%	Aug_max	597	0.45	13.50	1206
50%	HOD_mean_16	597	0.40	13.43	1200
50%	min_day_kw	597	0.15	13.37	1194
50%	kw_total_Nov	597	0.29	13.15	1172
50%	pre_CDH_pct	597	2.92	13.06	1163
50%	Aug_mean	597	0.65	13.03	1160
50%	kw_total_Aug	597	0.65	12.91	1149
50%	peak_frac	597	0.71	12.54	1113
50%	kw_mean	597	0.40	12.54	1113
50%	mean	597	0.40	12.54	1113
50%	pre_daily_kwh	597	0.40	12.54	1113
50%	max	597	0.30	12.49	1108
50%	min_3	597	-0.02	12.43	1102
50%	HOD_mean_12	597	0.33	12.42	1102
50%	kw_mean_winter	597	0.22	12.24	1084
50%	kw_total_Apr	597	0.36	12.15	1076
50%	HOD_mean_20	597	0.18	12.13	1074
50%	kw_total_Jun	597	0.70	12.09	1070
50%	sum2win	597	0.93	12.04	1064
50%	discretionary	597	0.42	11.93	1054