

Whitepaper

Customer Targeting for Residential Energy Efficiency Programs: Enhancing Electricity Savings at the Meter

Adam M. Scheer¹, Sam Borgeson², Kali Rosendo³

¹*Pacific Gas and Electric Company*

²*Convergence Data Analytics*

³*Massachusetts Institute of Technology*

October 27, 2017



Table of Contents

Abstract

Executive Summary.....	i
I. Introduction.....	1
II. Background.....	1
III. Customer Targeting.....	2
IV. The Programs and Datasets.....	5
i. Air Conditioning Quality Care (AC/QC).....	5
ii. Advanced Home Upgrade.....	5
iii. Sample Characterization.....	5
V. Methodology.....	7
i. Computing Pre-Period Load Characteristics.....	7
ii. Compute Cooling Energy Savings.....	7
a. Model Cooling Energy Usage.....	8
b. Compute Cooling Savings.....	8
iii. Determine Peak Summer Demand Savings.....	8
iv. Filter Customers via Pre-Program Usage Metrics.....	9
v. Treatment of Outliers.....	9
VI. Results and Discussion.....	9
i. Conglomerate Filter.....	9
ii. Exploring the Impact of Negative Savers.....	12
iii. Testing Individual Criteria.....	17
iv. Correlation Between Savings and Individual Criteria.....	20
v. Geographic Considerations.....	20
vi. Potential for Advanced Targeting Schemes.....	22
vii. Additional Implications.....	25
VII. Customer Targeting to Enhance EE as a Grid Resource.....	25
VIII. Programmatic Challenges and Recommendations.....	27

IX. Study Limitations.....	28
IX. Conclusions.....	29
Appendix.....	A1
A1. Outliers.....	A1
A2. Null Dataset Results.....	A2

Customer Targeting for Residential Energy Efficiency Programs: Enhancing Electricity Savings at the Meter

Adam M. Scheer¹, Sam Borgeson², Kali Rosendo³

¹*Pacific Gas and Electric Company*

²*Convergence Data Analytics*

³*Massachusetts Institute of Technology*

Abstract

The adoption of smart meters in California has yielded a stream of hourly electric usage (AMI) data for nearly every residential utility customer. These data are being put to use in a variety of ways, including more rapid identification of outages and development of time-of-use rate structures. However, California's energy efficiency (EE) programs and evaluations are only beginning to explore the vast insights offered from AMI data. In this paper we investigate the potential for increased electricity and demand savings by targeting customers for EE intervention based on features derived from their AMI usage profiles. Using recent past program participants, these data-driven targeting strategies are developed and tested against observed savings outcomes. Our analysis focuses on two long-standing residential EE programs offered by Pacific Gas and Electric Company, the HVAC Quality Maintenance (AC/QC) program and the whole home retrofit (Advanced Home Upgrade, or AHU) program. Results show that effective targeting can yield significantly enhanced per-capita savings and peak demand reduction. We find that even straightforward targeting schemes, based on pre-intervention usage data alone, are effective at selecting high-saving customers while also limiting the fraction of customers who consume more after the program. The highest performing schemes have the potential to increase per-capita savings and demand reduction by 50 – 150% when applied at moderate levels. Special attention is given to potential pitfalls, which could lead both programs and evaluations astray, including the effects of outliers and other statistical artifacts, and recommendations are given to avoid such stumbling blocks. Finally, we discuss broader implications, motivations, and barriers to implementing customer targeting.

Executive Summary

With the advent of smart meters and advancing data analysis techniques, effective customer targeting presents a major opportunity to increase the savings and cost effectiveness of many energy efficiency (EE) programs. In this whitepaper we develop and test customer targeting schemes based on interval data analytics. If utilized in current and future EE programs, these and similar methods can enhance the value of EE to participating customers, yield higher returns on investments from the ratepayer base, and provide more benefits at lesser cost for utility program sponsors.

This research analyzes data from two longstanding Pacific Gas and Electric (PG&E) residential EE programs to investigate the potential for improved electricity savings and peak demand reduction; the HVAC quality maintenance (AC/QC) and whole home retrofit (Advanced Home Upgrade) programs. For both programs cooling energy savings¹ are estimated for each participating customer via weather-normalized pre/post billing analysis. Customer targeting schemes based only on customers' pre-intervention electricity usage data then select subsets of the full sample. Savings for these subsets are compared to the full sample to gauge the effectiveness of the particular scheme.

The customer targeting approaches rely on intuitive considerations and are built from characteristics readily derived from a customer's hourly interval electricity usage data. These features are shown schematically in Figure E1 and together are expected to be indicative of high consumption, inefficiency, and potential for evening peak demand reduction. Some of these characteristics (i – iii) could be determined from monthly billing data alone while others (iv – vi) require hourly data.




Figure E1: Schematic summer (July) and shoulder month (April) usage and load shape characteristics that would be expected for a customer with high savings potential from an HVAC or building shell energy efficiency program

¹ While not the exclusive benefit of these programs, we focus specifically on cooling for several reasons. First, cooling savings account for a high fraction of total savings within both programs. Second, cooling needs are a primary driver of peak demand and high electricity procurement costs. Third, studying cooling load provides for a straightforward focus of this research.

The features i – v in Figure 1 are rolled into a conglomerate targeting strategy based on threshold values for each element. As thresholds are made more stringent, the targeting scheme becomes more selective in an attempt to identify customers with higher savings potential. The conglomerate scheme is applied to pre-program usage data and the resulting subsets consist of customers who have passed each threshold. Figure E2 shows how average per-household savings change (y-axis) for the Advanced Home Upgrade program as the targeting criteria are enhanced (x-axis). Three timeframes are broken out and studied independently: annual, summer, and summer peak.




Figure E2: Advanced Home Upgrade - Per capita normalized mean energy savings (kWh/hr) for each level of targeting applied during the three time periods: annual, summer (June – September) and summer peak (June – September; 3 – 9 pm). At the Loose, Medium, and Strict targeting levels, approximately 27%, 62%, and 86% of customers are removed.

Over each timeframe Figure E2 shows that applying the targeting filters substantially increases per-capita savings. Key findings include,

- Average household savings are more than doubled at the Moderate targeting level compared to the full sample.
- When the Loose criteria are applied, eliminating more than a quarter of the participants, total program cooling savings *increases*, indicating the lower-level filtering is particularly effective at removing customers with a propensity for neutral and negative savings.
- The summer timeframe is the predominant driver of cooling savings. Though accounting for only a third of the annual hours, the summer accounts for 70% of the annual savings.

Similar trends are observed in PG&E’s residential HVAC quality maintenance (AC/QC) program.

While the conglomerate filter is clearly effective, additional insights are gained from studying the performance of the individual criteria of Figure E1. In Figure E3 we show how electricity savings change as several individual criteria are applied in isolation. Similar to the results for the conglomerate filter, one can see that most individual criteria are effective at selecting customers with higher average savings.




Figure E3: Advanced Home Upgrade. Per-capita summer electricity savings for remaining program participants after filtering based on the identified individual criteria.

In both Advanced Home Upgrade and AC/QC, the total summer usage filter (iii) is very effective in predicting electricity savings and peak demand reduction. Similarly, the daily average maximum-to-minimum demand ratio (v) underperforms in both programs. In general, the spread in performance of the individual criteria shown in Figure E3 illustrates the need for empirical evidence in developing specific targeting strategies.

Further insight is gained when focusing on specific geographic regions. The Advanced Home Upgrade sample contained a large proportion of participants from both the hot Central Valley and the more temperate climate zones. Results indicate that the neutral and negative savers eliminated by the lower targeting levels originate almost exclusively within temperate climate regions. Collectively, customers in these regions save very little electricity from the building shell and HVAC system EE measures provided through the programs. This is a consequential observation as it suggests that both program and targeting strategies can be made more effective if focused by climate region. In that vein, Table E1 gives Central Valley savings results and average percentage household savings for two high performing individual filters: total summer kWh, and the ratio between summer-to-shoulder month usage.

Table E1: Central Valley Subset – Comparison of Summer kWh to Summer-to-Shoulder Ratio Filters; Average Household Savings

% Customers Filtered Out	Total Summer kWh		Summer to Shoulder Ratio	
	Summer kWh Savings	% Household Savings	Summer kWh Savings	% Household Savings
0	712	19%	712	19%
25	883	20%	880	21%
50	1,055	20%	1,022	23%
75	1,308	21%	1,170	26%
90	1,457	19%	1,269	28%

The total summer usage filter isolates high users while the summer-to-shoulder ratio attempts to identify inefficiency. While both criteria yield substantially increased savings, monitoring savings as a

percentage of total usage shows the filters succeed for different reasons. When applying the summer kWh filter, the percentage household savings is relatively static across increasing levels of usage. This indicates that the higher per-capita savings is attributable to selecting larger households with higher total usage. In contrast, after the summer-to-shoulder filter is applied, we observe a substantial increase in savings as a percentage of household usage, which indicates this criterion is identifying customers who offer a higher savings opportunity due to inefficiency. Strategically combining these elements provides a more optimal and robust targeting scheme. Table E2 gives recommended targeting parameters for retrofit programs operating in the Central Valley along with the percentage of participating Advanced Home Upgrade customers filtered out at each step. At the 50% targeting level, per-household savings is observed to increase by more than a factor of 1.5.

Table E2: Recommended Central Valley Targeting Scheme; Threshold Filter Values

% Customers Filtered Out	Average Daily Summer kWh (iii)	Summer to Shoulder^a Usage Ratio (ii)
10	12.93	0.827
25	19.60	0.827
50	26.98	1.138
75	34.58	1.498
90	42.14	1.805

^aSummer = June, July, August; Shoulder = November, February, March

Because the customer targeting schemes developed here are straightforward to compute from and apply to a large population of customers, they can be directly implemented to support future programs. The quantification of enhanced electricity savings and peak demand reduction within multiple programs upon customer targeting also provides a starting point for the valuation of AMI data and analysis. Within existing downstream programs, identifying customers who stand to benefit the most from a particular intervention can motivate enhanced and individualized value propositions to encourage participation. The opportunity is particularly salient within program designs, including Pay-for-Performance, for which identifying customers at the outset with high savings potential has inherent value to program implementers who are incentivized based on savings observed directly at the meter. Especially if overlaid with demographic and/or customer segmentation information, powerful messaging can be developed and directed to customers most in need of EE support.

Beyond enhancing savings for traditional programs, customer targeting can help EE compete as a distributed energy resource (DER). Effective targeting ensures that EE is directed to customers best able to deliver measurable grid benefits. However, structural changes are needed within EE program design and measurement for this vision to become reality. Counter to many traditional program designs, in the DER landscape, we posit that EE programs are likely to compete most directly with other resources if designed specifically to achieve and reward reliable savings at the meter. When considered alongside technologies such as solar, storage, and demand response, all of which allow for definitive and immediate measurement with a high degree of certainty, EE will be most competitive if resources are directed to high potential customers with results that can be quantified via meter-based analysis.

I. Introduction

California accounts for nearly \$1 billion of the more than \$6 billion spent annually on publically funded energy efficiency (EE) programs in the United States.² Despite the impressive investment, if the residential portion of this EE funding were spread evenly across the state, each household would only receive about \$15.³ This is enough to replace a few light bulbs, but orders of magnitude short of the five-figure cost of a deep retrofit. The limited reach of EE funding compels a simple question: How can resources be directed to the customers who will save the most? In this paper we show how smart meter data can be used within existing programs to target high-saving customers.

More specifically, we explore how characteristics of a customer's usage data can inform filtering rules that select customers with greater energy savings achieved 'at the meter.' The analysis focuses on Pacific Gas and Electric's (PG&E) residential HVAC quality maintenance (AC/QC) and whole home retrofit (Advanced Home Upgrade) programs. Because the filter rules developed here are straightforward to compute from and apply to a large population of customers, they can be directly implemented in support of future program targeting. We observe that billing data alone can provide for effective targeting schemes, but that insights offered through interval (AMI) data are essential to understanding the impacts of targeting on key metrics such as peak demand reduction.

In addition to increasing savings at the meter, effective customer targeting can also enhance utility avoided costs and cost effectiveness of EE programs. In particular, as Pay-for-Performance⁴ programs take root, and as EE policy goals intensify,⁵ program administrators, implementers, and regulators all have a stake in optimizing savings with data-driven techniques. Additionally, by selecting customers with higher propensity to deliver savings at the meter, targeting has the potential to make EE a more competitive and reliable grid resource within the broader energy and distributed energy resources market.

II. Background

Recent billing analysis studies on several California residential EE programs reveal consistent patterns:

1. Impacts observed at the meter vary widely among program participants.
2. A small fraction of program participants accounts for a high fraction of the total metered savings.
3. A significant number of program participants display *negative* savings (i.e. consume more energy after the program than before) when assessed at the meter.

Two residential HVAC EE programs were investigated as part of the *Phase I AMI Billing Regression Study*⁶ conducted by Evergreen Economics: PG&E's Residential Quality Maintenance (AC/QC) program, which promotes HVAC system maintenance and is also a focus of this work, and Southern California Edison's

²https://library.cee1.org/sites/default/files/library/12670/CEE_2015_AIR_Tables_March_2015.pdf; Table 4

³<http://www.census.gov/quickfacts/table/HSG010215/06>; California holds nearly 14 million households. The residential sector accounts for roughly one-fifth of California EE spending (\$200 million). \$200M/14M households = \$13.30/household

⁴ Pay-for-Performance programs are designed to pay all or part of the rebate based on the savings observed at the meter.

⁵ For instance, in California recent legislation (Senate Bill 350) has doubled the energy efficiency goals within existing buildings.

⁶ *AMI Billing Regression Study (Phase I)*. Evergreen Economics, 2016. CALMAC ID: SCE0383.01

(SCE) Quality Installation program (RQI), which incentivizes HVAC system installation based on industry performance standards. Both of these programs achieve savings goals based on deemed⁷ engineering estimates relative to a counterfactual⁸ baseline.

Results on the AC/QC program indicated that participant average annual household electricity savings were on the order of 3.5%, with significantly higher savings during summertime and periods of peak demand. In aggregate these results are sensible. However, when customers were ranked and grouped by their pre-period baseload⁹ energy usage, total program savings were observed to originate almost exclusively from the top 25% of customers. Further, the bottom 50% of customers collectively accounted for negative savings. In other words, in aggregate, half of program participants increased energy usage after the program. The negative savers had the effect of discounting a quarter of the positive savings in this baseload binning scheme. Results of the SCE RQI program show similar trends.^{10,11}

Recent study of the whole home retrofit Energy Upgrade California (EUC) Program¹² provided a unique perspective via a combination of billing analysis and complementary customer surveys. The billing analysis results showed that 12% of customers increased energy usage by at least 5% in the year after program intervention, while an additional 18% were near neutral savers (-5% to +5% savings). Survey results of the high negative savers revealed that half of these customers noted a change in heating behavior indicative of takeback, compared with only 19% among the population of high positive savers. Additionally, two-thirds of the negative savers indicated a change in occupancy compared to only one quarter of high positive savers. These results suggest that when measuring savings at the meter, enrolling customers likely to save more than average while avoiding negative-saving customers may make or break the success of an EE program.

It is worth noting that any sample of customers will show a distribution of pre/post “savings” due to natural variability in usage. Especially for ‘light touch’ EE intervention, the program may not be the primary factor that causes an average customer to use more or less. Therefore, the term ‘negative saver’ does not necessarily indicate a customer who used more *because* of the program. Nevertheless, the goal of customer targeting – finding customers most likely to yield high positive savings – remains.

⁷ ‘Deemed’ savings refer to engineering estimates for the *average* savings that would be expected for a particular EE measure.

⁸ The ‘counterfactual’ refers to what would have happened in the absence of the program. Many EE programs are designed to produce savings, and evaluated to assess savings, relative to the counterfactual.

⁹ Here ‘baseload’ refers to a household’s average minimum electricity demand throughout the course of a day. Baseload may also describe the minimum energy demand for a given region that must be serviced by power providers.

¹⁰ Again customers were binned by pre-program baseload energy consumption. In this case average annual total household electricity savings were found to be 7.0%. Again the lower usage customers increased consumption after program intervention, while the 25% of highest usage customers accounted for approximately 65% of total savings.

¹¹ In both programs, the negative savers may exhibit a takeback effect. Consider the RQI program, in which customers have a new HVAC system installed. If these households did not have central AC before participating, or had a smaller system, even optimal installation of an energy efficient unit would result in increased electricity demand. In the HVAC maintenance (AC/QC) program, it could be that a high occurrence of repairs, returning partially failed systems to full service, occurred along with program measures for many of the negative saving customers. However, without further research, these possibilities will remain unconfirmed.

¹² *PG&E Whole House Program: Marketing and Targeting Analysis*. Opinion Dynamics Corporation, 2014. CALMAC ID: PGE0302.05

III. Customer Targeting

In developing EE targeting criteria, we note that high usage by itself does not guarantee system inefficiency. Some customers have high usage simply because they have higher than average needs for heating, cooling, maintenance, and lighting. Therefore, we posit that even a minimalist EE targeting strategy should include a specific criterion focused on end use efficiency in addition to a total usage threshold to best gauge savings potential. In Table 1 we give the total electricity usage for three hypothetical customers on two days. Day 1 has a high temperature of a moderate 70 °F and the high on day 2 is 90 °F, a hot day.

Table 1

Outside High Temperature	Daily Electricity Usage		
	Customer A	Customer B	Customer C
Day 1: 70 °F	3 kWh	10 kWh	20 kWh
Day 2: 90 °F	6 kWh	17 kWh	22 kWh

If this were the only available information, which customer should be recruited for an air conditioning or building shell EE program? On both the mild and hot days Customer C uses the most electricity. However her usage only increases by 10% after the sizable rise in outside temperature. It would seem that baseload measures may be more appropriate for this customer.¹³ By comparison, on the hotter day, Customer A's electricity usage doubled, a potential indication of inefficiency. But with a total consumption, even on the hot day, of only 6 kWh, Customer A has little potential for savings and therefore would not be a top choice for program intervention. Finally, Customer B has significant total electricity usage *and* shows a considerable fractional increase (70%) in total usage on the hot day. Based on these limited data, only Customer B meets both the threshold of high usage and high temperature sensitivity that would indicate high potential to save.

Building upon these considerations, Figure 1 is a schematic of the energy usage characteristics we hypothesize a customer with high savings potential for air conditioning or building shell measures would exhibit. The red curve indicates average daily summertime demand while the dark blue curve shows average daily demand during the milder shoulder months for the same customer.¹⁴

¹³ Another possibility for the small usage change could be the lack of an HVAC system, or a system in need of repair. In such a case, this customer would be likely to demonstrate significant takeback after program participation, which would lead to negative savings in a pre/post billing analysis.

¹⁴ Summertime is taken as June – September, typically the hottest months in central California climate zones, and the shoulder months are taken as February, March and November. Cooling needs are expected to be greatest in the summer months while electricity usage for combined heating and cooling is expected to be minimal during the shoulder months.

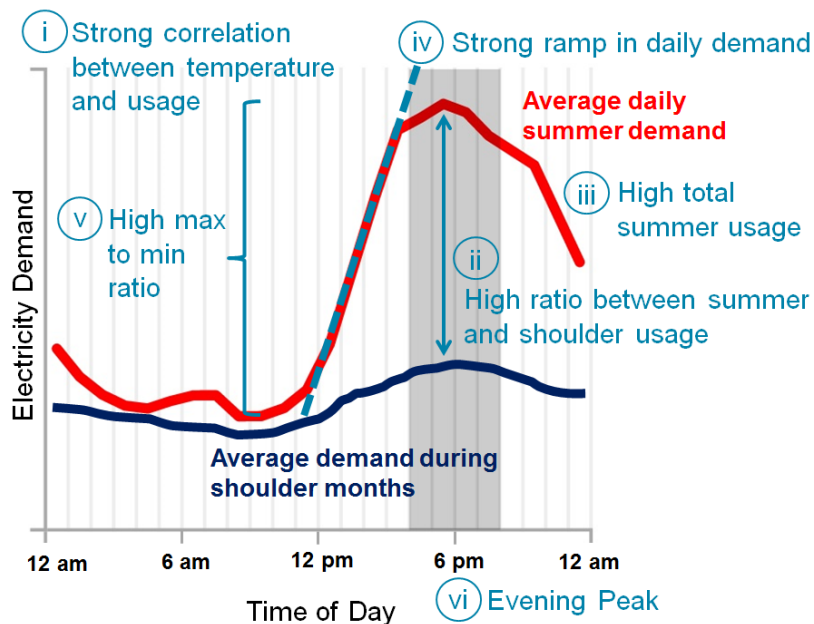


Figure 1: Summertime and shoulder month load shapes for an ideal residential customer for an HVAC or building shell energy efficiency program

The metrics i – iii are associated with total electricity usage while iv – vi are those related to evening peak demand. Below we develop and test targeting filters based on one or more of these criteria:

Expected predictors of electricity savings potential

- i. Statistical correlation¹⁵ between hourly readings of outside temperature and demand (ranges from -1 to 1)
- ii. Ratio of total summertime (June, July, August) electricity consumption to usage in the mild shoulder months (November, February, March)
- iii. Total summertime (June, July, August) electricity usage in kWh

Expected predictors of electric peak demand savings potential

- iv. Average hourly usage increase (kWh/hr) during the ramp up period to the evening peak (defined as 3 pm to 7 pm)
- v. Average of the ratio between daily maximum and minimum demand, as derived from hourly meter readings
- vi. The fraction of total summer load occurring during peak hours (defined as 3 pm to 9 pm)

The seasonal and total usage metrics i – iii, illustrated here with hourly interval data, can be discerned through analysis of monthly billing data if necessary, but the peak demand filters iv – vi can only be assessed with interval data. To optimize a residential building shell or HVAC EE program for *electricity* (kWh) savings, filters based on these *demand* (kW) metrics may be overly restrictive or even

¹⁵ Using the cor function in R; For an explanation on correlation, see <https://www.socialresearchmethods.net/kb/statcorr.php>

counterproductive. However, to target peak demand reductions along with total energy savings, combined criteria iv – vi are hypothesized to pay dividends.

In the next section, we develop and apply customer selection filters using threshold criteria for predictors i – v at varying levels. The threshold values for each filter are computed from customers' pre-intervention AMI data. We ascertain the ability of this targeting scheme to select customers with above average electricity savings and peak demand reduction and quantify the associated change in average per-capita¹⁶ savings.

IV. The Programs and Datasets

Analysis focuses on three datasets: 1. Recent participants in PG&E's residential HVAC quality maintenance program (Air Conditioning Quality Care; AC/QC); 2. Recent participants in PG&E's whole home retrofit (Advanced Home Upgrade) program; 3. A 'null' dataset of randomly selected residential customers who did not participate in either program. The datasets contain 1-hour interval electricity usage data for a minimum of one full year pre- and post-intervention for every customer.

i. Air Conditioning Quality Care (AC/QC)

The AC/QC dataset contains 1,216 participating customer records for program years 2012 – 2014 and was also used for the Evergreen Economics study⁶ discussed above. All data cleaning steps performed by Evergreen were retained for this analysis along with additional cleaning described in the Methodology section below. In the years studied, the AC/QC program incentivized contractors to perform several measures including system assessment and cleaning, refrigerant charge adjustment, blower motor retrofit, and duct sealing. Contractors were allowed to complete one or more of these measures for any job and received a "kicker" incentive for jobs with multiple measures.

ii. Advanced Home Upgrade

The Advanced Home Upgrade dataset contains records from 6,286 participating customers spanning 2013 – 2016 and was also used in the development of the CalTRACK¹⁷ billing analysis platform. Advanced Home Upgrade is part of the Energy Upgrade California (EUC) program and is designed to deliver whole home retrofits that reduce household energy usage by 20% or more through building shell and HVAC upgrades. Most often, lighting and plug loads are not addressed. Ex ante savings are calculated by contractors using approved modeling software and verified by quality control checks performed by the program implementation firm. The EUC program also hosts the Basic Home Upgrade pathway that delivers retrofits with deemed savings estimates. The Basic path is not studied here.

iii. Sample Characterization

Figure 2 illustrates California's climate zones¹⁸ and gives a breakdown of participants by zone. Only customers who passed all data cleaning steps (and were therefore included in the analysis) are shown.

¹⁶ Throughout this paper per-capita is used synonymously with per-household

¹⁷ <http://www.caltrack.org>

¹⁸ *The Pacific Energy Center's Guide to California Climate Zones and Bioclimatic Design*, 2006.

https://www.pge.com/includes/docs/pdfs/about/edusafety/training/pec/toolbox/arch/climate/california_climate_zones_01-16.pdf

Nearly all AC/QC participating customers live in climate zones 11, 12, and 13. Collectively these areas comprise the Central Valley, an inland region with hot and dry summers and relatively cold winters. Advanced Home Upgrade also realized significant participation in the Central Valley, but had many participants from across PG&E's service territory, including in the mild coastal climate zones.

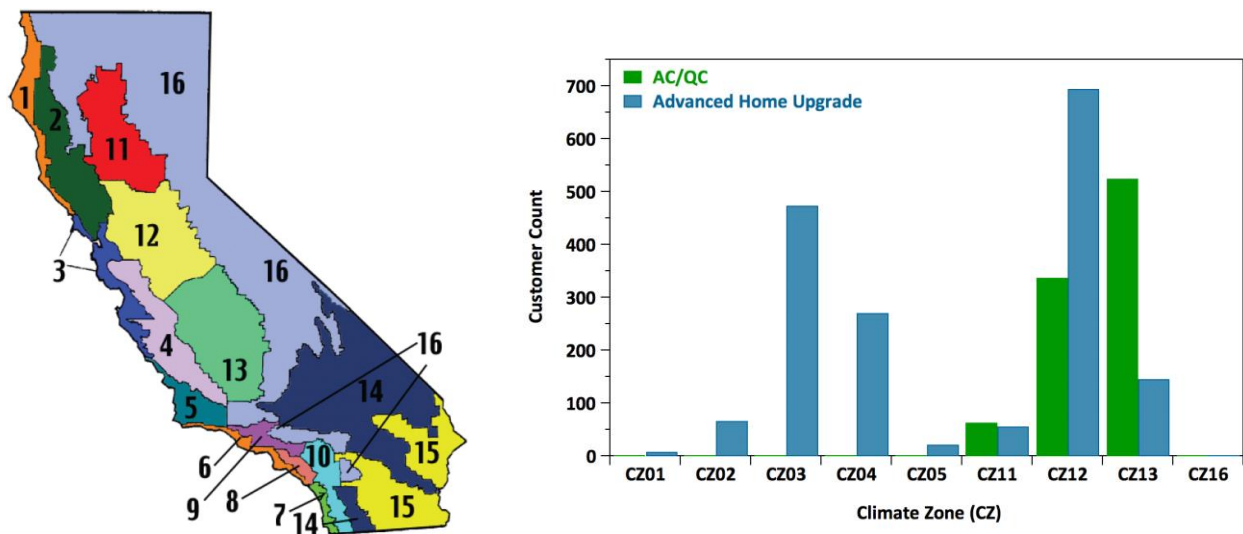


Figure 2: Left: The California Climate Zones, reproduced from ref. 18. Right: AC/QC and Advanced Home Upgrade program participants by climate zones. Only customers who passed all data cleaning steps are shown.

The distribution of total pre-program annual household energy usage for the AC/QC and Advanced Home Upgrade samples is given in Figure 3.

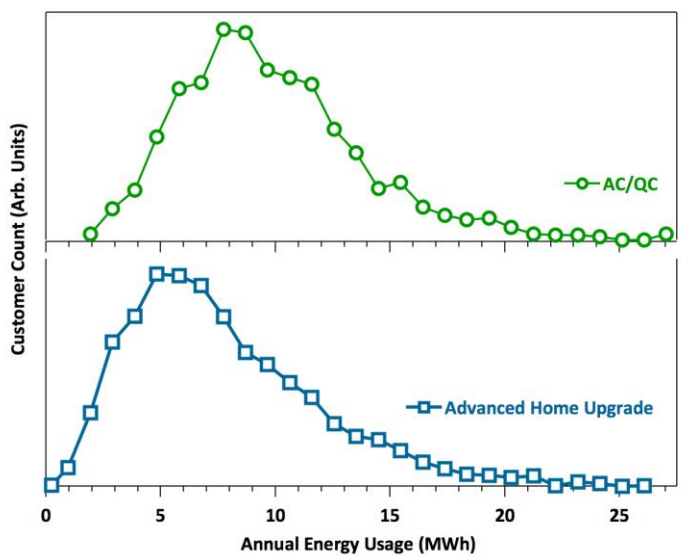


Figure 3: The distribution of annual electricity usage for participating AC/QC and Advanced Home Upgrade customers in the dataset

The AC/QC distribution peaks around 8 MWh compared to approximately 6 MWh for Advanced Home Upgrade. The average total annual energy usage for the AC/QC participants was also higher at 9.9 MWh, compared to 8.0 MWh for Advanced Home Upgrade. The peak in the AC/QC distribution occurs

approximately 30% higher than the average household consumption across PG&E’s service territory.¹⁹ The higher usage reflected in the AC/QC dataset is consistent with the Central Valley’s higher heating and cooling needs and larger home sizes compared to the coastal regions that are prevalent in the Advanced Home Upgrade sample. Both samples show a long tail with dozens of customers consuming more than 20 MWh annually.

V. Methodology and Experimental Design

Basic data quality filters were performed on each customer’s usage data, including (1) mean demand must be greater than 110W, (2) data must cover at least 180 days, and (3) fewer than 15% of readings are allowed to be zero. Table 2 shows how these quality filters impacted the total sample size and the application use for each subset of data.

Table 2

Data Cleaning Step	Customers Remaining in Sample		
	AC/QC	AHU	Null
Initial	1,216	5,981	6,544
Pre-period load characteristics	1,204	4,843	6,374
Pre-period CDD model estimation	1,118	4,179	6,342
Post-period CDD model estimation	930	1,786	5,785

A significant fraction of the Advanced Home Upgrade sample was eliminated, especially in the post-period data cleaning step. It is likely that more of the sample could be retained with more fastidious procedures, including interpolation schemes for absent meter-reads. However, sufficient sample remained without these measures for the purposes of this study and the authors chose to proceed with the smaller dataset. Consistency checks between retained and eliminated customers showed no issues.

The following steps were performed to compute a pre/post electricity savings assessment for each customer and apply targeting filters:

i. Computing Pre-Period Load Characteristics

Customer data is filtered to the year before program intervention. These data are then used to compute the usage metrics that are illustrated in Figure 1 and given in Table 3 below. We also compute a suite of “basic metrics” as defined by the open-source meter data analysis package VISDOM.²⁰

ii. Compute Cooling Energy Savings

In this study we focus specifically on cooling for several reasons. First, both programs service AC systems and/or provide building shell improvements. While the latter would also be expected to yield gas and potentially baseload savings, we take it as a prerequisite that a targeted customer should be expected to

¹⁹ PG&E’s Energy Efficiency Business Plan 2018 – 2025; Residential Appendix p. 17; https://media.wix.com/ugd/0c9650_cbeb1d9e14cf4575845e8d5cd6bce57f.pdf

²⁰ Sam Borgeson, Jungsuk Kwac and Ram Rajagopal (2016). visdom: R package for energy data analytics. R package version 0.7.0. <https://github.com/convergenceda/visdom>

deliver some degree of variable electricity savings after substantial building shell work. Second, meeting cooling needs is a primary driver of peak demand and high electricity procurement costs and associated savings are thus particularly valuable. Third, isolating variable load provides for a straightforward focus of this research. Often throughout this paper we refer to the cooling savings estimates simply as savings.

a. Model Cooling Energy Usage

To isolate cooling energy from total household load, 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. A day's CDH is the sum of the degrees the outside temperature is above 65°F (or 0 if cooler than 65°F) across all hours, h , in each day, d .

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

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

The regression coefficient α quantifies the cooling sensitivity of each household and can be used to predict daily cooling energy given a computed CDH for day d . The weather normalization model is run for each customer using data from the cooling season (May through September) of both pre- and post-intervention years. An additional modeling approach was tested in which unique change point temperatures were determined for each customer.²¹ Generally this model yielded similar results and has the advantage of allowing temperature responses to vary both above and below the change point. However, this method is more complex and not without drawbacks.²² For the purposes of this research we will report CDH model results for usage and kWh savings results.

b. Compute Cooling Savings

Now, using daily weather data from the post period, we compute post-period daily CDH and use the α coefficients from both the pre and post models to predict daily cooling energy. The predictions using the pre-intervention coefficient are used as the counterfactual for how much cooling energy would have been required if the efficiency intervention had not occurred. Thus the savings estimates can be made by computing the difference between model estimates. We use different subsets of post-period weather data to compute annual (365 days of data), summer (June through September), and summer peak period savings (June through September; 3 pm through 9 pm).

iii. Determine Peak Summer Demand Savings

Using data from each customer from the peak demand period (3 – 9 pm) for every day in the pre-intervention summer, we compute the overall maximum hour and 97th percentile hour of evening

²¹ The alternative model also applies a single change point, but the best change point is selected separately for each individual customer via a parametric search of candidate change points, accepting the one that explains the most variance (the max R^2). This model allows for temperature responses to vary both above and below the change point.

²² The added degrees of freedom can provide a better fit for some customers (for instance households with atypical balance point temperatures), but they can also overfit data from non-temperature responsive homes and produce higher variation and thermal response outliers.

demand. Next, we calculate the same metric using post-intervention summer days. By taking the difference between the pre- and post-intervention values, we obtain an estimate of peak demand reduction for each customer. The cooling energy and demand reduction metrics then become the values used to quantify the performance of usage features applied as customer filters.

iv. Filter Customers via Pre-Program Usage Metrics

Finally, with the usage metrics from the pre-intervention period (the same information available to program administrators at the time), we flag targeted customers. For example, a targeted customer might have overall consumption greater than a certain threshold value, show an evening-peaking load profile, have a higher than average increase in consumption during the summer, etc. Whatever the filter criteria, they include some customers and exclude others and thereby define a subset of customers. We compute the mean value of the energy and 97th percentile peak demand savings for the targeted subsets and compare the results to the untargeted program outcomes.

We start with a set of “conglomerate” filters chosen based on first principles and professional judgement (Figure 1) and define three levels, Loose, Moderate, and Strict, for their thresholds. At these levels approximately 70%, 36%, and 14% of customers remain in the sample. After applying these filters, we track the change in observed savings in the subsets of customers they isolate. To better understand the performance of individual targeting criteria and to compare to the conglomerate filter, the individual criteria are tuned to eliminate a progressing fraction of customers (10%, 25%, 50%, 75%, 90%) and the resulting performance is assessed via the mean summertime savings across the resulting subsets of customers. With this approach we can compare the performance of the individual metrics to the compound conglomerate filters during the time period of highest intensity savings.

v. Treatment of Outliers

Savings estimates are computed from both the full cleaned sample and a sample consisting of the middle 3 – 97% of savers. By removing the 3% of lowest and highest saving customers, we attempt to ensure that results are not overly dependent on a small number of customers who display a very high discrepancy in their year-to-year energy usage that is unlikely attributable to program intervention. Reasons for substantial outliers may include installation of rooftop solar PV, purchase of an electric vehicle, or other significant structural load changes, occupancy changes. This approach does risk exclusion of some customers with very substantial usage changes that are legitimately due to the program. In general, similar trends are observed with and without outliers, with savings generally being diminished somewhat upon their exclusion. Appendix A provides more detail on the quantitative effect of retaining or removing outliers in the sample.

VI. Results and Discussion

i. Conglomerate Filter

After removing outliers, the remaining participating customers were filtered into targeted and non-targeted samples based on their pre-program energy usage and load shapes using Loose, Moderate and Strict targeting criteria tied to variables i – v described above (Figure 1). Tables 3a and 3b give the

specific filters, associated threshold values, and the number of customers in each group for AC/QC and Advanced Home Upgrade, respectively. Note that with the Loose targeting criteria, the load shape requirements are not applied.

Table 3a: AC/QC Targeting Scheme

Level	Targeting Criteria					Fraction of customers remaining
	i	ii	iii	iv	v	
Loose	0.2	1.20	19	-	-	0.67
Moderate	0.3	1.70	26	0.25	2.0	0.34
Strict	0.3	1.85	36	0.40	5.0	0.13

Table 3b: Advanced Home Upgrade Targeting Scheme

Level	Targeting Criteria					Fraction of customers remaining
	i	ii	iii	iv	v	
Loose	-0.05	0.71	6.0	-	-	0.73
Moderate	0.07	1.12	12.5	0.13	2.0	0.38
Strict	0.07	1.30	26.4	0.31	5.1	0.14

i. Temperature to load correlation (ranges from -1 to 1)

ii. Ratio of summer-to-shoulder electricity consumption

iii. Total summer electricity usage (kWh)

iv. Evening ramp slope (kWh/hr)

v. Average daily max-to-min demand ratio

To be selected, a customer must pass the minimum thresholds for each filter. The thresholds were set such that all filters had approximately the same impact in removing customers. The last column shows the percentage of customers remaining in the sample after all the filters at a given level were applied. For instance, in AC/QC most (67%) customers passed the “Loose” criteria, while only 13% of customers passed the “Strict” filters.

Figures 4 and 5 show per-capita cooling savings for the four cases of no targeting, and Loose, Moderate and Strict targeting as detailed in Tables 3 and 4. Results are given in the normalized quantity of kWh saved per hour (kWh/hr) for three different timeframes: annual, which comprises the full year after program participation (8760 hours), summer, which isolates June – September (2928 hours), and summer peak, which assesses June – September during the hours of 3 – 9 pm (732 hours).²³ Peak hours were chosen to coincide with peak pricing in PG&E’s time-of-use rate structures. The summer and summer peak timeframes and savings are essentially subsets of annual.

²³ These hours are consistent with the summertime periods of peak demand in recent analyses, including in “*Revisiting the California Duck Curve*,” ScottMadden, 2016.

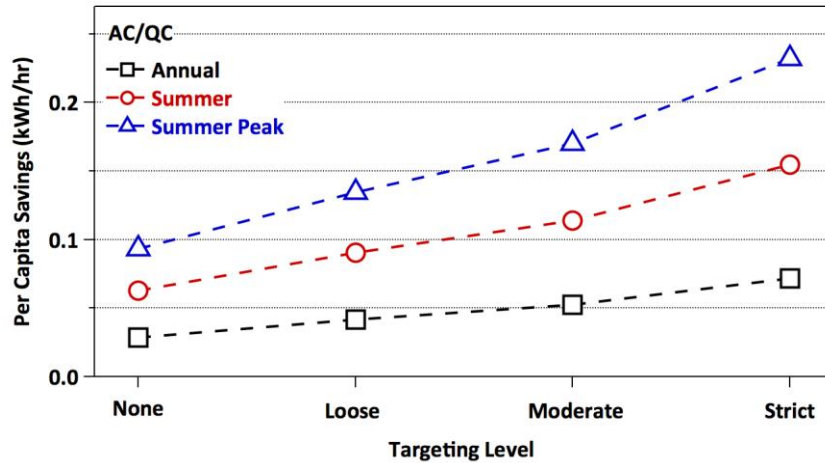


Figure 4: AC/QC - Per capita normalized mean cooling savings (kWh/hr) for each level of targeting applied during the three time periods: annual, summer (June – September) and summer peak (June – September; 3 – 9 pm)

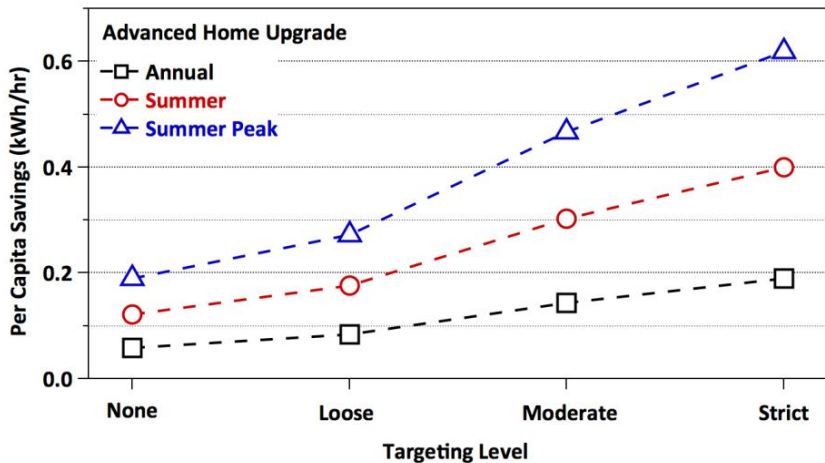


Figure 5: Advanced Home Upgrade - Per capita normalized mean cooling savings (kWh/hr) for each level of targeting applied during the three time periods: annual, summer (June – September) and summer peak (June – September; 3 – 9 pm)

Compared to annual, both programs show increased hourly savings (kWh/hr) during the summer and summer peak timeframes. For both programs and across each timeframe a significant increase in hourly savings is observed as the targeting criteria are enhanced. Applying the Moderate criteria, which eliminated two-thirds of the full sample, increased the per-capita savings by more than a factor of 1.8 in AC/QC and by more than a factor of 2.5 in Advanced Home Upgrade. Further enhancement in per-capita savings is observed when applying the Strict criteria, illustrating the continued effectiveness of augmenting the targeting rigor. While the results are informative, the Strict level of filtering, at which nearly 90% of potential participants are eliminated, may not be practical for most programs.

Repeating this analysis for the null dataset (Appendix A) shows a small positive differential, which is likely attributable to a regression to the mean effect. That the null data displays little pre/post differential, even at high targeting levels, lends confidence that the effects presented throughout this work are not due to a statistical issue or anomaly.

Along with electricity savings, we also computed the average peak demand reduction²⁴ as a function of targeting rigor. Similar trends are observed to those in Figures 4 and 5 and are detailed in Table 4.

Table 4: Average Peak Demand Reduction (kW)

Targeting Level	AC/QC	AHU
None	0.12	0.47
Loose	0.16	0.68
Moderate	0.17	1.07
Strict	0.19	1.30

One can see that as with electricity savings, targeting pays dividends in per-capita peak demand reduction.

As described in the Methodology section, independent models were optimized to best fit the data for the annual, summer and summer peak timeframes described in Figures 4 and 5. Table 5 shows that results calculated for these subsets are very highly correlated (0.991 – 0.998). For much of this paper, we focus on the summer timeframe, which provides both strong signal and sufficient hours.

Table 5: Savings Correlation	kWh Savings		
	Annual	Summer	Summer Peak
Annual kWh Savings	1.000		
Summer kWh Savings	0.998	1.000	
Summer Peak kWh Savings	0.991	0.997	1.000
Peak Demand Reduction (kW)	0.639	0.634	0.629

Table 5 also shows the correlation between electricity savings and peak demand reduction is significant but not as strong as between the usage subsets themselves.

ii. Exploring the Sample and the Impact of Negative Savers

Table 6 reproduces the per capita normalized kWh/hr data of Figure 5 along with the total MWh savings achieved by the specific subsets.²⁵

Table 6: Advanced Home Upgrade Cooling Electricity Savings by Targeted Subset and Timeframe

Targeting Level	Customer Count	Annual		Summer		Summer Peak	
		kWh/hr	Total MWh	kWh/hr	Total MWh	kWh/hr	Total MWh
None	1,565	0.061	832	0.128	585	0.197	226
Loose	1,140	0.086	854	0.179	597	0.275	230
Moderate	585	0.146	749	0.306	524	0.471	202
Strict	213	0.192	358	0.403	251	0.622	97

²⁴ Computed via the 97th percentile hour of demand as described in the Methodology section.

²⁵ The customer count for the full sample in Table 6 differs from Table 2 due to exclusion of outliers and a small number of customers for which one or more filters did not compute.

Several important insights are apparent in these data. First, the majority of annual savings occurs during the summer. The summer timeframe accounts for only a third of the 8,760 annual hours, but yields 70% of the annual electricity savings (556 MWh out of 789 MWh). Second, though one might expect that removing a large fraction of the program participants, even those who do not pass targeting criteria, would diminish total program savings, the data in Table 5 suggest otherwise. At the Moderate level, 90% of the total annual savings persist despite removing nearly two-thirds of the participants. Additionally, upon application of the Loose criteria, which removes more than a quarter of the participants, total cooling savings *increases* by 3%. This is in line with the observations of near neutral and negative savers from the previous studies discussed above and raises a key point of this research: The efficacy of targeting criteria hinges largely on the ability to predict customers who will not save or will consume more after program intervention.

Figure 6 shows the rank-ordered summer cooling savings for the full Advanced Home Upgrade sample (red) along with every customer eliminated by the Loose criteria (blue), and every customer retained by the Moderate filter (green).




Figure 6: The rank-ordered summer savings (kWh) for the full Advanced Home Upgrade sample (red), for every customer eliminated by the Loose criteria (blue), and for every customer retained by the Moderate filter (green).

While the Loose filter did eliminate a small number of significant positive savers, the large majority of households that were filtered out consisted of near neutral and negative savers. In contrast, customers retained by the Moderate filter were nearly all highly positive savers, with only a small fraction of negative savers passing the threshold criteria. As one would expect, the full sample displays characteristics of both subsets.

Figures 7 (summer electricity usage) and 8 (peak summer demand) detail how the fraction of negative savers in the sample changes upon targeting. The leftmost bars show that summertime electric cooling usage increased after program intervention for 39% and 38% of participants within AC/QC and Advanced Home Upgrade. Collectively, these negative savers eliminated 51% (AC/QC) and 21% (Advanced Home Upgrade) of the positive savings achieved by the remaining participants as indicated by

the light blue bars. Applying the Loose, Moderate and Strict criteria reduces the proportion of negative savers dramatically. At the Moderate level, the fraction of negative savers falls to 32% within AC/QC and to 10% in Advanced Home Upgrade. The fraction of positive savings eclipsed by these negative savers drops to 27% and 4%, respectively. These results, along with many others presented here, reflect the deeper savings delivered through more comprehensive measures via Advanced Home Upgrade compared to AC/QC. With the lighter touch of the latter, savings are generally smaller and harder to distinguish from natural fluctuations in consumption.




Figure 7: Top – AC/QC; Bottom – Advanced Home Upgrade. Solid red bars: The percentage of negative saving customers in the sample as a function of targeting level, as measured by summertime electricity usage. Light blue bars: The percentage of positive savings eclipsed by the negative saving participants at each level of rigor.

Similar trends are observed in Fig. 8 for peak demand reduction. Despite having relatively little effect on the total percentage of negative savers, significant impact is observed for AC/QC by applying Loose targeting, with the fraction of positive savings eliminated by negative savers falling from 66% to 48%. This is an indication that the Loose filtering has succeeded in removing some of the highest negative savers. Incremental improvements are observed thereafter. In contrast, results for Advanced Home Upgrade show definitive benefits at each progressing level of targeting rigor, with the percentage of positive savings eclipsed by negative savers falling from 40% to just 6% at the Moderate level.




Figure 8: Top – AC/QC; Bottom – Advanced Home Upgrade. Solid red bars: The percentage of negative saving customers in the sample as a function of targeting level, measured by 97th percentile peak demand reduction. Light blue bars: The percentage of positive savings eclipsed by the negative saving participants at each targeting level.

There are a number of reasons a customer may show negative savings in this type of pre/post billing analysis. Along with natural year-to-year fluctuation, these include addition of load, occupancy change, poor contractor performance, regression to the mean, and takeback. Some of these are indicative of a customer who was not a good match for the program, while others are more nuanced. For customers who have structural changes in their energy usage due to load additions, a deemed or modeled counterfactual is a more appropriate than a baseline determined through pre-program usage data alone. Both programs that are the topic of this paper serve a number of customers who fall into this category. However, the fact that the targeting criteria are effective at identifying negative savers through pre-program usage alone is an indication that such households are not in the majority.

Along with the assessment of negative savers, it is revealing to observe how the full distribution of cooling savings changes upon application of the targeting filters. Figure 9 shows density plots for summer electricity savings for the two programs. Each point represents the number of customers (y) in a bin centered at the respective x -value. The traces are scaled such that the total area under each curve is equal, thus providing a normalized comparison.

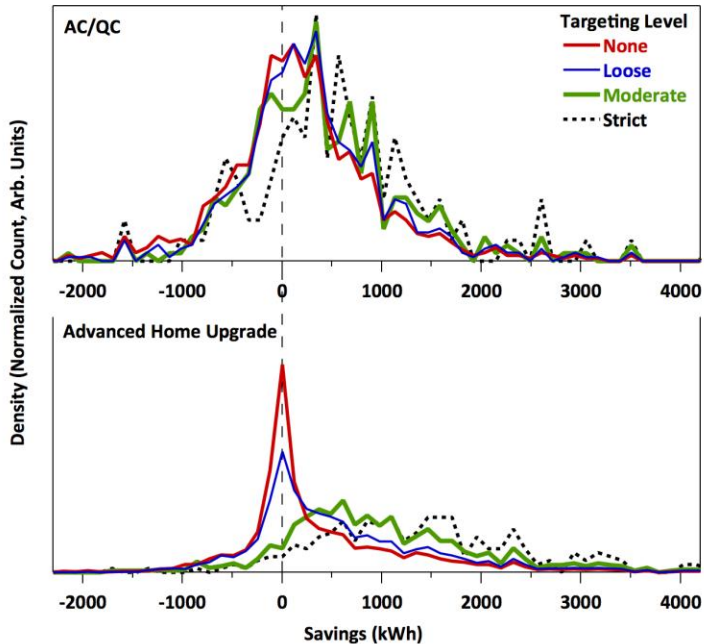


Figure 9: Top – AC/QC; Bottom – Advanced Home Upgrade. Density plots of the summer cooling savings observed as a function of targeting rigor. Positive savings are to the right of 0.

One can see that with no or Loose filtering, both the AC/QC and Advanced Home Upgrade distributions peak near zero savings. However, the Loose targeting broadens the distribution on the positive-savings side while repressing the neutral and negative savers. This effect is more readily visible in the Advanced Home Upgrade results. With the Moderate filters in place (green traces), peaks in both distributions shift off of zero, well into the range of positive savings. At the Moderate level in Advanced Home Upgrade nearly all substantial negative savers have been eliminated and only a small tail in negative territory remains, as would be expected from Figure 7. The Moderate filters eliminate 89% of the bottom quartile of the full sample and retain 87% of the top quartile. At this level, more than 55% of the Advanced Home Upgrade distribution is in the top quartile of the full sample. The shift in AC/QC is significant but small relative to Advanced Home Upgrade, which is again reflective of the lighter intervention of the former.

Analogously, Fig. 10 shows density plots for peak demand reduction.

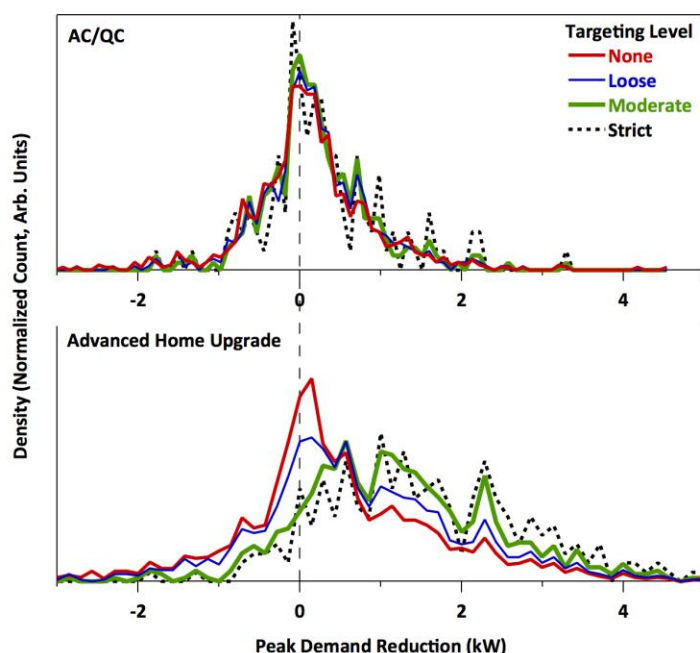


Figure 10: Top – AC/QC; Bottom – Advanced Home Upgrade. Density plots of the 97th percentile peak demand reduction observed as a function of targeting rigor. Positive savings are to the right of 0.

Similar to the electricity-savings, at the no and Loose targeting levels, distributions are peaked near zero. Again, the AC/QC program shows only minor visible shifts, while Advanced Home Upgrade displays a significant evolution to positive savings upon increasing targeting rigor. For the latter, at the Moderate level 47% of customers are in the top quartile of the full sample, 73% of the top quartile are retained, and 85% of the bottom quartile are eliminated.

iii. Testing Individual Criteria

Thus far, we have applied targeting criteria i – v as in Table 3 such that they all have significant influence over the final sample of customers. To gain more insight into the performance of each individual filter, we analyzed the per-capita electricity savings and peak demand reduction from independent application of criteria i – v, along with two additional elements:

- vi. The fraction of total summer load occurring during peak hours
- vii. The absolute range (kW) of average summertime daily minimum to maximum demand

Calculation of these additional filters requires access to interval data and both are hypothesized to correlate with peak demand reduction. The filters i – vii were individually applied and tuned such that 10%, 25%, 50%, 75%, and 90% of the total sample was eliminated, again based on pre-program usage data. In Figure 10 we present the results of this analysis for summertime electricity savings for AC/QC and Advanced Home Upgrade.

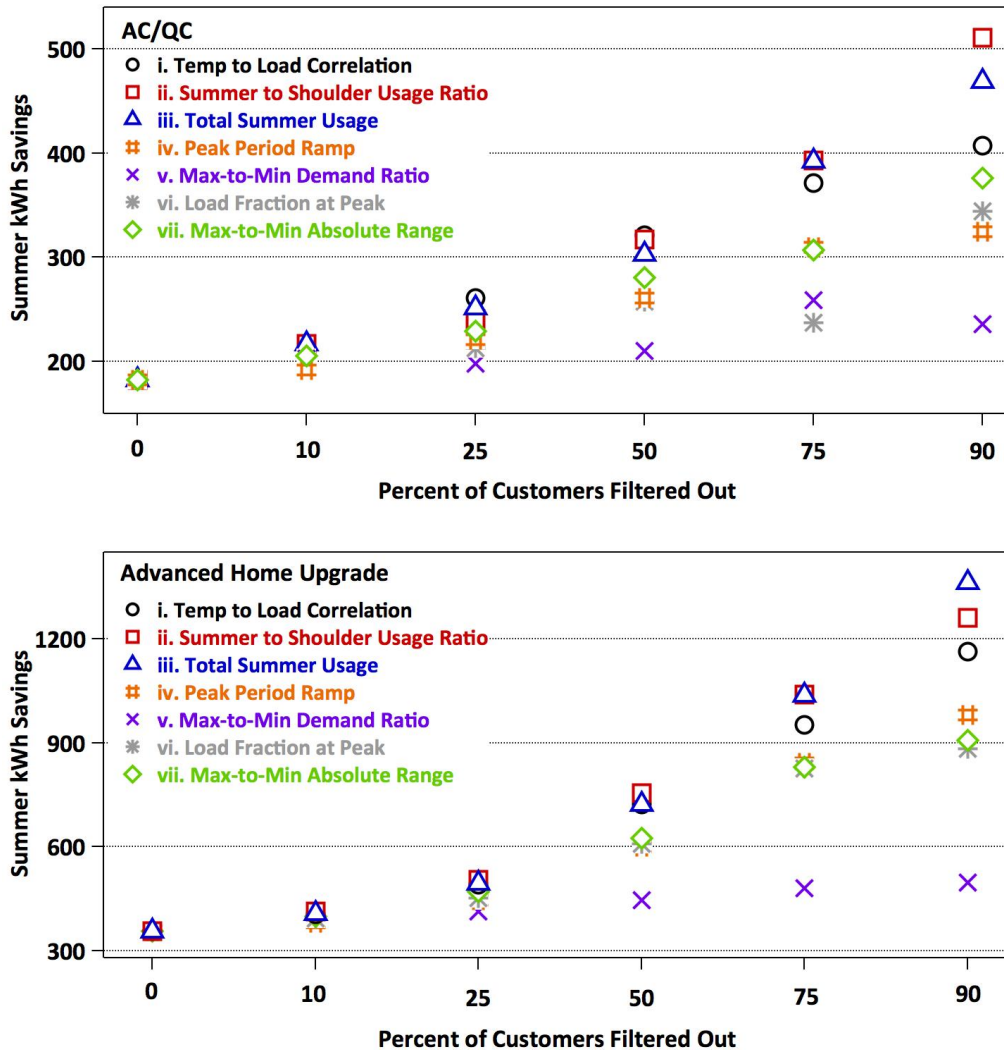


Figure 10: Top – AC/QC, Bottom – Advanced Home Upgrade. Per-capita summer electricity savings for remaining program participants after filtering based on the individual criteria i – vii.

Several interesting findings and trends emerge from these results. Generally, application of each filter has a positive effect on the per-capita savings. However, some filters perform very well, while results for others are underwhelming. In both the AC/QC and Advanced Home Upgrade samples, the three filters focused on electricity usage and efficiency (i – iii) all behave well, yielding comparable results. This is particularly notable because these are the three criteria that can be calculated from monthly billing data (without interval meter readings).

In both programs the least impactful filter is v, the ratio between maximum-to-minimum demand. This can be attributed to the propensity of this filter to retain customers with both low baseload and low total electricity consumption. Such customers could have a high max-to-min demand ratio simply because the denominator is small. However, in the conglomerate targeting scheme, the other filters would be expected to remove such customers. Therefore, it is unclear immediately if this criterion is a valuable element of the conglomerate scheme. Filters iv, vi, and vii all yield continuously enhanced per-

capita summer kWh savings as they are applied more aggressively, though not to the degree of i – iii. Considering iv, vi, and vii were designed primarily to identify customers with high potential for peak demand reduction, their underperformance in predicting kWh savings relative to the usage-focused criteria is not surprising. In general, these results showcase that selection of specific targeting criteria is critical to effectively identifying customers with high savings potential.

Figure 11 gives analogous individual filter results for peak demand reduction.

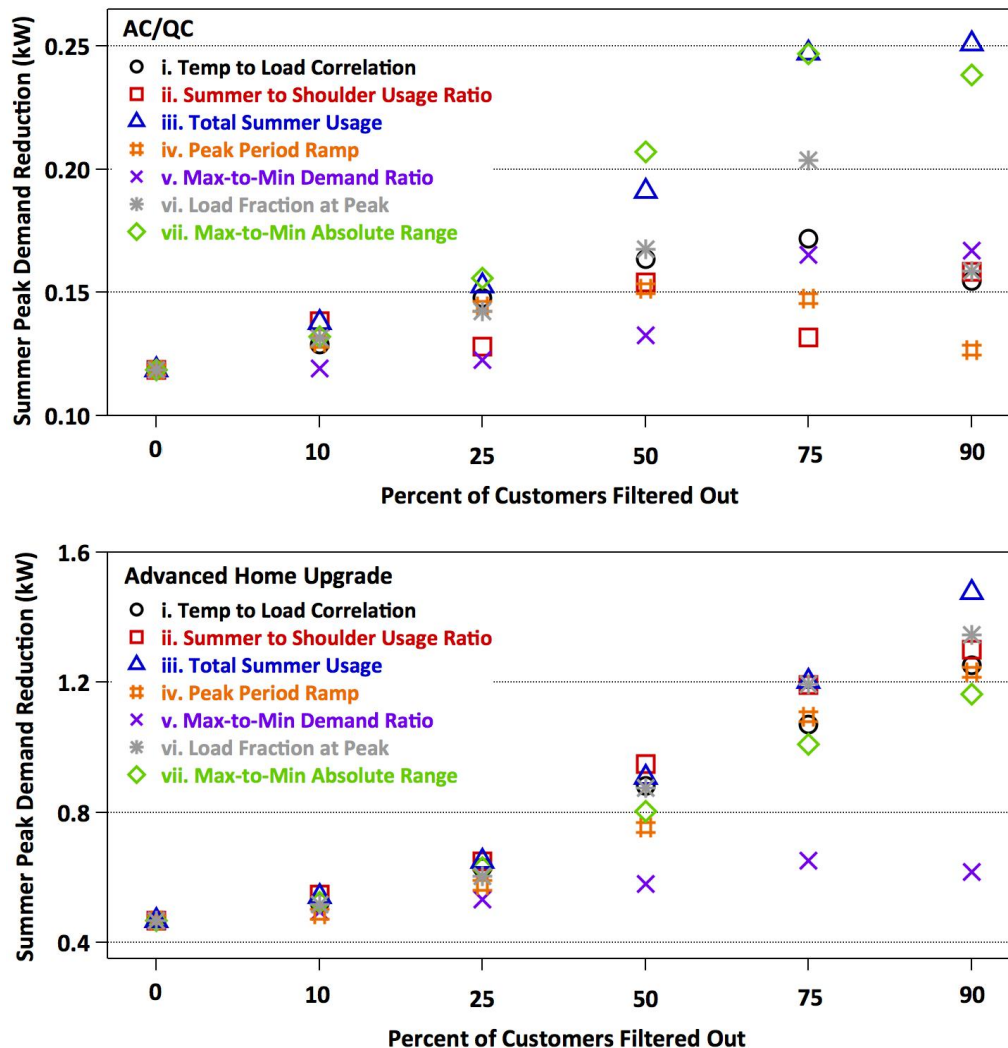


Figure 11: Top – AC/QC, Bottom – Advanced Home Upgrade. Summer per-capita 97th percentile peak demand reduction for remaining program participants after filtering based on the individual filters i – vii.

Compared to the results for electricity savings, different patterns are evident between AC/QC and Advanced Home Upgrade in assessing individual filter effects on peak demand reduction. However, in both electricity and peak demand results, the two consistent trends between programs are the high performance of the summer kWh filter (iii) and the continuing poor results from the daily maximum-to-minimum demand ratio filter (v). The performance of the other filters is shuffled between the two programs, with the absolute maximum-to-minimum range (vii) outperforming summer kWh (iii) at

certain levels of rigor in AC/QC, while the load fraction at peak (vi) and summer-to-shoulder usage ratio (ii) consistently show good results in Advanced Home Upgrade. As discussed earlier, the precision of the Advanced Home Upgrade results is greater than the AC/QC program due to the higher average savings and larger sample of the former. Indeed Fig. 11 shows steadier trends for Advanced Home Upgrade and more apparent scatter in the AC/QC results. The top performing individual filters for savings and peak demand reduction are observed here to very slightly outperform the conglomerate filter described above. We return below to the topic of potential for a more optimal filter.

iv. Correlation Between Savings and Individual Criteria

Table 7 is the correlation matrix for filters i – vii along with summer kWh savings and evening peak demand reduction for the Advanced Home Upgrade program.

Table 7: Advanced Home Upgrade Individual Targeting Criteria and Savings Correlation Matrix

Correlation	Summer kWh Savings	Peak Demand Reduction	i. Temp to Load Correlation	ii. Sumr/Shoulder kWh Ratio	iii. Total Summer Usage	iv. Peak Period Ramp	v. Max-to-Min Dmd Ratio	vi. Load Fraction at Peak
Summer kWh Savings	1.00							
Peak Demand Reduction	0.63	1.00						
i. Temp to Load Correlation	0.58	0.40	1.00					
ii. Sumr/Shoulder kWh Ratio	0.63	0.40	0.80	1.00				
iii. Total Summer Usage	0.64	0.44	0.64	0.57	1.00			
iv. Peak Period Ramp	0.38	0.28	0.17	0.29	0.37	1.00		
v. Max-to-Min Dmd Ratio	0.09	0.09	0.11	0.06	0.12	0.07	1.00	
vi. Load Fraction at Peak	0.41	0.33	0.57	0.55	0.25	0.53	0.17	1.00
vii. Max-to-Min Range	0.45	0.36	0.42	0.29	0.82	0.33	0.47	0.20

Correlation to savings provides an incomplete but meaningful measure of filter performance. Not surprisingly, the three top performing criteria of Figure 10, i – iii, show the highest degree of correlation with summer kWh savings (0.58 – 0.64). Intriguingly, though not to the same degree, these filters are also the most correlated with peak demand reduction (0.40 – 0.44). Filter v, maximum-to-minimum demand ratio, is the least correlated with savings and with other filters, confirming that by itself it is a poor criterion. Table 7 also shows that filters i and ii are very highly correlated (0.80). When considering that a customer with high temperature-to-load correlation would also be expected to have a high ratio between summer and shoulder period usage, the strong correlation between the two is an indication that they may not both be needed in a targeting scheme. Total summer usage (iii) and daily average maximum-to-minimum range (vii) are also very highly correlated (0.82). Yet Figure 10 shows iii significantly outperforms vii at predicting kWh savings, highlighting that correlation alone is an imperfect indicator.

v. Geographic Considerations

To this point we have treated the full AC/QC and Advanced Home Upgrade samples without parsing based on geographic, demographic or household characteristics. A detailed investigation of the latter two is outside the scope of this work. However, we can assess trends by climate region and we focus on

Advanced Home Upgrade, which has a much more varied geographic distribution of participants as shown in Fig. 2.

Given the targeting criteria that emphasize total and increased summertime usage, in the conglomerate scheme (Table 3b) one would expect that households in the hot climate zones with higher cooling needs would be favored. This is indeed observed to be the case. At the Loose level slightly more than half of participants in the mild climate regions²⁶ are eliminated, while nearly 90% of the Central Valley customers pass. At the Moderate level, only 7% of the temperate climate participants are retained compared to 63% from the hot climate regions.

Figure 12 shows the rank-ordered summer cooling savings for the full sample as well as participants in the mild and hot climate regions.




Figure 12: The rank-ordered summer savings (kWh/hr) for the full sample (red), Central Valley (CZ 11 – 13) participants (black), and all other climate zones (green).

Average per capita savings within the hot climate region are nearly double that of the full sample. In contrast, the large majority of mild climate participants saved very little or increased cooling usage after program intervention. Despite a tail of high savers, the average summer cooling savings for the mild climate customers was just 0.006 kWh/hr, essentially zero given statistical uncertainty.²⁷ The disparity between temperate and hot regions suggests both that distinct program designs should be developed for the different climate zones and that the independent programs would benefit from separate targeting schemes.

Table 8 shows how the household average summer kWh usage, savings, and savings as a percentage of usage compares for the full Advanced Home Upgrade sample and the Central Valley subset at different levels of targeting. By the summer kWh filter design, average summer usage of the remaining subsets

²⁶ For the purposes of this discussion, the “hot” climate regions are the Central Valley zones 11 – 13 and the “mild” climate regions are all others.

²⁷ With potentially greater fractional electric heating consumption, one may expect that these customers would save more on an annual basis. However, average annual household savings are even less, 0.002 kWh/hr.

increases at each step. Across levels, threshold filter values to achieve the same percentage of filtering are significantly higher for the Central Valley than for the full sample.

Table 8: Summer kWh (iii) Filter Results for the Full Sample and the Hot Climate Zone Subset

AHU	Full Sample			Central Valley Climate Zones 11 - 13			
	% Customers Filtered Out	Avg. Summer kWh	Avg. Summer kWh Savings	Avg. % Household Savings	Avg. Summer kWh	Avg. Summer kWh Savings	Avg. % Household Savings
0		2,854	374	13.1%	3,811	712	18.7%
10		3,088	415	13.4%	4,083	775	19.0%
25		3,460	502	14.5%	4,480	883	19.7%
50		4,225	727	17.2%	5,220	1,055	20.2%
75		5,401	1,042	19.3%	6,294	1,308	20.8%
90		6,761	1,350	20.0%	7,476	1,457	19.5%

As the targeting is made more strict, the percentage household savings for the full sample increases substantially while the percent household savings in the Central Valley subset shows very little change. This indicates two important points. First, percentage household savings is a misleading metric if not put into proper context. The increase in percentage household savings for the full sample upon targeting is due mainly to eliminating the mild climate zone participants who in aggregate offer very little savings. Second, the relatively static percentage household savings in the Central Valley indicates that increased savings upon enhanced targeting is primarily due to selecting larger households with higher total usage and therefore higher savings potential.

vi. Potential for Advanced Targeting Schemes

Table 9 gives a comparison between the savings achieved when utilizing the total summer kWh filter (iii) and the summer-to-shoulder usage ratio filter (ii) for participants in the hot climate regions.

Table 9: Comparison of Summer kWh (iii) to Summer-to-Shoulder (ii) Filters; Central Valley Subset

AHU	Total Summer kWh Filter (iii)			Summer to Shoulder Filter (ii)			
	% Customers Filtered Out	Avg. Summer kWh	Avg. Summer kWh Savings	Avg. % Household Savings	Avg. Summer kWh	Avg. Summer kWh Savings	Avg. % Household Savings
0		3,811	712	18.7%	3,811	712	18.7%
10		4,083	775	19.0%	3,964	782	19.7%
25		4,480	883	19.7%	4,180	880	21.1%
50		5,220	1,055	20.2%	4,433	1,022	23.0%
75		6,294	1,308	20.8%	4,491	1,170	26.1%
90		7,476	1,457	19.5%	4,490	1,269	28.3%

Counter to the summer kWh filter results, upon application of the summer-to-shoulder filter, a significant increase in savings as a percentage of household usage is observed. Thus both filters are very

effective but for different reasons. While the summer kWh criterion successfully selects customers with enhanced savings due to their larger initial demand, the summer-to-shoulder filter identifies customers with higher propensity to save due to their initial inefficiency.

This returns us to our original hypothesis: Targeting schemes can be more effective with components that address both threshold usage and efficiency. The combination of filters ii and iii provides both elements. Because the summer kWh filter somewhat outperformed the summer-to-shoulder ratio, we establish it as the ‘dominant’ criterion and use the summer-to-shoulder filter to ‘sweep’ for homes with high usage but that also display high efficiency, such as hypothetical Customer C in Table 1. The thresholds for both filters are determined at each level with the following logic:

1. At every level, both filters are set to eliminate a minimum of 5% of customers.
2. At the 50% level, the two filters are applied in such a way that maximizes savings.
3. Adhering to step 1 as default, both filters are applied in the same ratio as determined in step 2. The ratio is figured by rank ordering customers for each filter. After optimizing savings at the 50% targeting level, the ratio between the rank of ii and iii is kept constant for other levels.

Following these steps yields the threshold filter values of Table 10 and the results of Figure 13.

Table 10: Central Valley Subset: Threshold Filter Values

% Customers Filtered Out	<u>Individual Filters</u>		<u>Combined ii + iii</u>	
	Summer kWh (iii)	Summer to Shoulder (ii)	iii	ii
10	14.440	0.946	12.930	0.827
25	20.250	1.184	19.600	0.827
50	28.810	1.496	26.980	1.138
75	32.826	1.913	34.580	1.498
90	52.010	2.330	42.140	1.805




Figure 13: For Advanced Home Upgrade Central Valley Climate Zones 11 – 13: A comparison of the best performing individual filters for kWh savings (iii. Summer kWh usage and ii. Summer-to-Shoulder Usage Ratio) compared to the combination of these two filters as targeting rigor is increased. The black diamonds show results for a hypothetical “best possible” filter, which results from simply taking the highest saving customers at each level.

Application of the combined filter yields slight improvement in per-capita savings at every level with the exception of the most strict at which a more substantial enhancement is observed. A similar analysis shows that small additional gains can be achieved for peak demand reduction by combination of a threshold and efficiency filter.

Also included in Figure 13 is a “best possible” filter, which simply gives the upper bound for targeting performance based on the rank-ordered list of savings for each customer. Several real world considerations are expected to limit the ability to achieve a result close to this ideal hypothetical:

- A number of different contractors deliver an EE program in the field. The quality of work performed varies from one contractor to another.
- Contractors and customers must plan the EE investment on an individual basis. Despite a customer’s potential based on their pre-intervention usage data, he or she may or may not elect optional energy saving program offerings. No effort has been made here to differentiate the depth of work performed at individual households.
- Customer behavior and household occupancy may change within the observation period.

Future research should investigate the efficacy of more sophisticated modeling for customer targeting for programs with relatively small savings, such as AC/QC, as well as for programs designed to deliver deeper savings like Advanced Home Upgrade. Use of detailed load disaggregation modeling and/or machine learning algorithms might be expected to outperform the relatively straightforward criteria developed here. However, the high fidelity of the straightforward filters in eliminating near neutral and negative savers, while retaining the high positive savers, as described above, indicates that more advanced techniques may not be expected to yield dramatically improved results.

vii. Additional Implications

In the process of assessing the performance of different targeting filters, several additional insights of value arose. As has been found in previous research discussed above, here we confirm that a high portion of customers do not produce electricity savings at the meter. Figure 9 shows that for Advanced Home Upgrade, a building shell and HVAC retrofit program designed to achieve 20% household-level savings, the full sample distribution *peaks* at zero cooling savings. Knowing only this fact may lead to speculation on multiple fronts: Are contractors ‘gaming the system’ and/or delivering poor quality work? Is there a high fraction of projects with added load, a high propensity of occupancy changes, or a high degree of behavioral takeback? To date only marginal understanding has been gained on the underlying reasons for the negative and neutral-saver phenomenon.

While there are certainly some contractors who are not performing up to program standards and some customers using the programs in conjunction with significant load additions, our results indicate that the principal reason for the high fraction of neutral and negative savers is the low propensity of building shell and HVAC upgrades to yield substantial electric savings in the temperate climate zones. Among Advanced Home Upgrade projects completed in the hot Central Valley, Figure 12 shows that the large majority result in significant savings. In short, when the program serves a customer with a high opportunity for savings, the customer most often indeed saves. This implies that takeback and contractor underperformance issues are likely restricted to isolated incidences instead of being a pervasive feature of the program.

That targeting may eliminate subsets of customers with low propensity to benefit from a particular program may raise questions of fairness in resource allocation. We suggest that the best way to address this concern is through a more individualized approach and a robust suite of offerings. While in this study, we have identified customers with little savings potential for building shell and AC upgrades, they may have opportunity to save in other ways, notably gas, lighting, and baseload. Further, to save nothing on cooling from a program that emphasizes relatively expensive HVAC and building shell measures does not provide for a reasonable return on investment for either participating customers or the ratepayer base. In general, the results of this study showcase the need for more tailored design and delivery of EE programs to specific regions and for specific customers. Addressing baseload, lighting, and plug-loads among a wider customer population in temperate climates instead of pursuing deep building shell upgrades for a relatively small number of households may be a wiser use of limited resources.

Finally, the features upon which targeting can be built are often themselves instruments to gain important insights. As discussed in Section vi, customers with higher usage tend to save more, but do not appear to be using energy less efficiently. In contrast, customers with high summer-to-shoulder period usage ratios appear to offer higher savings because of their more inefficient homes. Through the lens of a program implementer, a high summer-to-shoulder usage ratio may be a particularly attractive feature as it signals the probability of higher savings for a smaller scope of work in comparison to focusing on large households with high total usage. It is likely that many more such insights are available with a careful study of these data and it is worth note that these implications are ascertained without many traditional evaluation tactics and only customers’ AMI data as inputs.

VII. Customer Targeting to Enhance EE as a Grid Resource

Many of the considerations that compel customer targeting for meter-based EE programs can also enhance EE as a grid resource. Traditionally EE programs have served to meet greenhouse gas reduction targets, support markets for efficient technologies and services, and help customers save on their bills. While EE programs will still be called upon to meet these important objectives, the role of EE is expanding. Both regulatory bodies and program administrators are pursuing EE as a tool to enhance grid reliability, defer investment in new generation, and mitigate expensive procurement during periods of peak electricity demand. In many cases EE is being considered alongside a broad suite of distributed energy resources (DERs) including demand response, solar photovoltaics and other renewables, storage, electric vehicles, and even time-of-use pricing schemes. Because the residential sector accounts for a high fraction of peak evening demand, EE delivered in this sector can serve an important role in alleviating stress on the grid and reducing usage during the times of highest procurement costs. Further, with the onset of time-of-use electricity rates, EE and demand response programs that curtail load demands during peak hours will save customers more on their bill than conventional programs. In short, from a customer benefit and utility avoided cost perspective, not every kWh is created equal.²⁸

This point is demonstrated in recent research from E2e²⁹ in which the avoided procurement costs for several EE programs were investigated. Pre/post billing analysis using 1-hour interval data was performed for approximately 10,000 participants in SCE's Residential Quality HVAC Installation (RQI) program between 2010 – 2015. The RQI savings were found to occur in high coincidence with periods of high procurement costs.³⁰ The authors quantified this overlap, then compared the resulting avoided costs to those estimated by a simple averaging of annual savings and procurement costs. The RQI program was estimated to deliver more than 50% higher avoided costs than would be expected with simple averaging over the annual timeframes. By more effectively targeting customers for program participation as described here, the avoided costs and benefits of EE to the grid can be enhanced further and made more reliable. These considerations will become even more important in the next decade, with the "Duck Curve"³¹ poised to become more pronounced due to increasing adoption of solar and residential plug loads exacerbating the mid-day trough and evening peak, respectively. Therefore, we anticipate even greater value for peak summer kW savings than currently prescribed in short term savings models.³¹

It is worth noting that EE faces a unique circumstance compared to traditional generation and many other DERs: An EE intervention results in avoided energy use, which cannot be measured instantly or

²⁸Electricity procurement prices in California consist of two principle components: costs of procuring energy (wholesale) and the cost to ensure adequate capacity (capacity costs), both of which are time-varying. Wholesale prices depend on the specific generation mix called upon to deliver electricity at any given time, while capacity fees are paid in advance to ensure sufficient supply-side resources are in reserve to guard against blackouts. California wholesale prices tend to range from \$20 - \$60 per MWh, peaking during periods of high demand, while capacity costs are usually near zero during off-peak hours, but can approach \$400 or more per MWh during peak hours in the summer months.

²⁹*Do Energy Efficiency Investments Deliver at the Right Time?* J. Boomhower and L.W. Davis (2016).

³⁰The Model showed that RQI savings occur almost exclusively during the hot summer months and predominantly from noon - 10 pm, when air conditioning needs are greatest.

³¹Unlike many states, the time valuation of savings is built into the avoided cost assessment utilized in California.

directly. While this may seem benign, the consequences are dramatic and pose significant challenges to successful deployment of EE as a DER. First, an accurate assessment of EE savings, i.e. accounting for seasonal variations, often requires a year or more of both pre- and post-intervention data and careful analysis to quantify uncertainties – major detriments in an industry with justifiably low risk tolerance. Second, while a grid operator may be most interested in the ability of EE to reduce demand, EE savings are often measured relative to baselines that do not relate to an observable change in metered consumption. These realities foster EE programs that are not designed to produce evaluable savings at the meter, with a major overhaul of program design, targeting, and measurement and verification methods required to better align EE outcomes with grid needs.

Ultimately, programs that will be evaluated by measuring the change in metered energy consumption must be specifically designed to deliver savings that are measurable at the meter. Such programs will find inherent value in the types of targeting methods described here.

VIII. Programmatic Challenges and Recommendations

PG&E has implemented customer targeting beta-tests in both its commercial and residential HVAC quality maintenance programs. In both cases analysis of AMI electricity usage data was performed to identify the ~20% of customers with the highest potential for HVAC savings. In both scenarios, “kicker” incentives were offered to contractors for serving a targeted customer. Both efforts have thus far proven unsuccessful due to similar challenges: Customers often have established contractor relationships. Targeted customers may be hesitant to sign on with a new contractor through an EE program. Similarly, it can be expensive and time consuming for participating contractors to recruit new customers. Targeted customers may also be dispersed geographically, which poses challenges for recruitment and delivering high volume interventions. To address these challenges the following approaches merit examination.

Programmatic Considerations and Recommendations

- While kicker incentives can motivate the extra work to recruit a targeted customer within a downstream deemed program, transitioning to Pay-for-Performance program models in which incentives are awarded based on savings observed at the meter would inherently reward serving customers with high propensity to save.
- Utilize the targeting analysis for more personalized messaging that conveys an enhanced value proposition for the identified customers. Results of automated calculations similar to those performed for this study have been delivered in the field,³² including strategies that provide graphical comparisons to similar customers’ energy usage patterns, a powerful social norming concept.
- When developing lists of targeted customers, cluster by location as much as possible.

³² See for example, *Review and Validation of 2014 Pacific Gas and Electric Home Energy Reports Program Impacts (Final Report)*, DNV GL, 2016.

- Utilities can play an important role in customer targeting by generating leads for participating contractors. Customers often trust their utility as an energy advisor, which positions the utility to communicate customer-specific usage details and provide recommendations on ways to manage usage, including EE programs. When third parties are implementing EE programs, customer data along with the utility brand and marketing apparatus can be valuable commodities that are woven into agreements.

Analytics and Evaluation Recommendations

- Pay close attention to potential pitfalls, including the impact of outliers, when testing a targeting scheme or determining aggregate program savings via billing analysis. Here we have taken precautions to remove the top and bottom 3% of savers and selected the 97th percentile hour of demand to determine peak demand and peak demand reduction. There are more precise ways to identify outliers and evaluators should determine an appropriate strategy to assess spurious data points for each study. The potential for regression to the mean should also be monitored as was done here with the null dataset analysis. If a significant effect is observed, it could be backed out for the determination of final results.
- Introduce grid-relevant criteria, such as time and space variations in the value of savings, into the measurement and verification process and clarify how the differentiated value of savings at different times and locations can be monetized in programs.
- Even where official evaluation criteria have not changed, methods of pre/post and treatment/control evaluation of savings at the meter should be developed and refined, including in retrospect, to aid in understanding the grid impacts of EE and to facilitate a smooth transition to programs designed for at-the-meter savings.
- Combine energy usage data with econometric data on high energy saving customers to further refine customers with a high propensity to save. Previous research¹² found that positive savers are more likely to be longstanding customers with higher incomes in larger, older homes, among other factors.

IX. Study Limitations

As with any study, this research was performed with limited time and budget. Here we identify known limitations, how they may have impacted uncertainty or conclusions, and where future investigation would be useful. Limitations include:

- **Development and testing of targeting criteria utilized data from past programs.** While the AC/QC and Advanced Home Upgrade data served as rich grounds for the development and testing of customer targeting schemes, the ultimate goal is to apply and evaluate such strategies in the field. Ideal for an experiment would be a program that serves both non-targeted and targeted populations.
- **Determination of savings and targeting results focused exclusively on cooling energy.** While this was a choice by design, contractors in the Advanced Home Upgrade program also service

gas appliances as well as pool pumps. (Most often, lighting and plug loads are not addressed.) Therefore the targeting schemes likely excluded some customers who saved substantially or included others who were relative underperformers. In particular, customers in the temperate climates may use little electricity for cooling but may save in these other ways. If the goal of a targeting scheme is to maximize all savings from a truly comprehensive program, a multipronged strategy may be more beneficial. On the other hand, a focus on cooling energy alone is probably sufficient to target electricity savings during times of high procurement costs and to maximize grid benefits. Phase II of this research will address gas. Within the residential sector, further research that assesses targeting for baseload and lighting needs would also be valuable.

- **This research did not attempt to quantify how customer and contractor choices, work quality, and home size affected performance.** Customers and contractors in both AC/QC and Advanced Home Upgrade are faced with choices on the scope of a retrofit. This research did not attempt to normalize savings for the scale of work performed. Some customers may elect a deep retrofit while others opt for a less expensive, scaled back scope. Some contractors may tend toward recommending larger scale projects while others gravitate toward a lighter touch. The quality of work among contractors is likely to vary as well as post-program customer behavior. It is also likely that smaller households required less program resources to serve and could therefore be undervalued in this analysis. In future work it would be valuable to assess savings on a kWh-per-dollar basis. Nevertheless, despite all these sources of potential noise in the data, our analysis yielded consistent, sensible trends along several lines of inquiry, which gives us confidence in the validity of results.
- **Analysis of a matched control group and adjustment of results accordingly was not performed.** While a deep control group analysis may be needed to verify the absolute savings results described throughout this paper, such a sophisticated analysis was tangential to the goals of this study. We would not expect the customer targeting trends to vary substantially if a control group was used to adjust the absolute value of savings. Further, the null dataset results showed only a small regression to the mean effect, which helps validate the merit of the pre/post billing analysis upon which savings are based.
- **A fixed setpoint temperature of 65 °F was used.** While a fixed setpoint temperature was deemed sufficient for the purposes of this research, a variable setpoint method with a high degree of quality control to mitigate spurious results and model overfitting may increase accuracy at an individual household level.
- **Free ridership implications are not addressed.** Though beyond the scope of this research, it would be valuable for future studies to assess the implications of customer targeting on free ridership. Within downstream programs like AC/QC and Advanced Home Upgrade, the responsibility for customer recruitment often rests primarily with participating contractors. The easiest customers to recruit are those who are already planning upgrades in line with program offerings. If successful in recruiting customers from targeted lists, especially with enhanced and individualized marketing outreach, such bias may be largely avoided.

X. Conclusions

In this work we have investigated how targeting criteria that are straightforward to compute from a customer's electricity usage data can impact the cooling savings expected from residential HVAC and building shell programs. For both the AC/QC and Advanced Home Upgrade programs investigated here, as the targeting criteria were made more rigorous, per-capita electricity and peak demand savings increased substantially. Specific filters were shown to be effective not only at identifying high-saving customers, but also eliminating underperforming households. These targeting strategies utilized datasets from past EE programs as testing grounds and the resulting strategies can be put to use immediately for future program targeting.

Several other important insights arose as part of this research, including the observation that a large fraction of program participants consisted of neutral and negative savers. Further analysis showed that the overwhelming majority of these customers were located in temperate climate zones. In contrast, participants from the hot Central Valley tended to save a great deal. Focusing targeting schemes on the Central Valley customers revealed that filtering based on usage alone selects customers who save more due primarily to their higher initial energy needs. Using a filter designed to identify inefficiency (the dimensionless ratio between summer and shoulder-period usage) also performed well by retaining customers with a substantial increase in their savings as a percentage of total household usage. The combination of these filters was shown to yield somewhat higher savings than either individually.

Though per capita savings may increase substantially upon targeting, ultimately selection of targeting rigor depends on many factors including the population of potential customers, the effectiveness of program marketing at driving uptake, program budget and goals, and any grid needs being addressed. Consider Figure 14, which shows how Advanced Home Upgrade per-capita savings relates to total summer savings for the combined summer kWh (iii) and summer-to-shoulder ratio (ii) filters for the Central Valley Climate zones (Table 10).

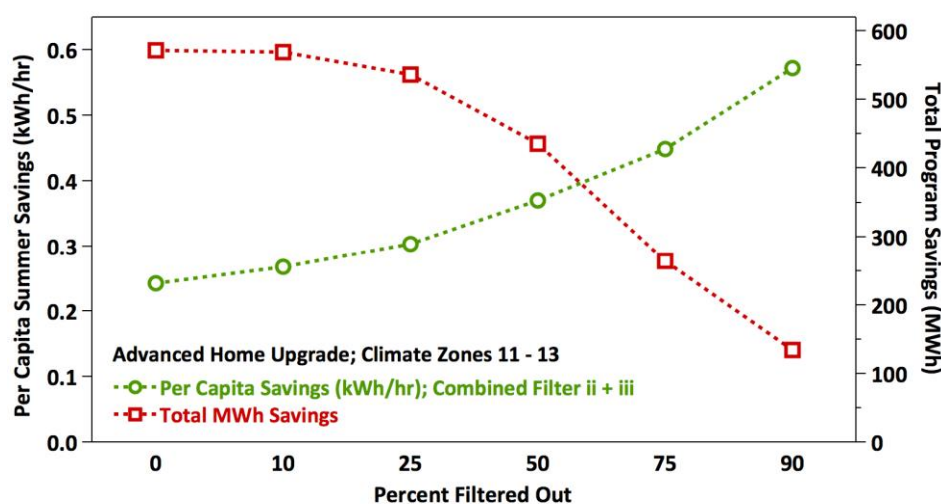


Figure 14: For Advanced Home Upgrade Central Valley Climate Zones: Juxtaposition of total program savings from the associated subsets (red squares) and corresponding per-capita savings (kWh/hr; green circles) upon application of the combined Summer kWh (iii) and Summer-to-Shoulder Usage Ratio (iii) targeting scheme of Table 10.

Here 94% of total program savings persist after targeting removes one quarter of the sample. Similarly, despite eliminating half of the participants, more than three-quarters of total savings are retained due to an enhancement of per-capita savings by 52%. These would appear to be good investments, especially if loss in uptake can be made up with a focused marketing effort. However, while improvements in per-capita savings continue, more than half of total savings are lost if targeting is set to remove 75% of the sample. In many cases, this may be too severe, especially considering the reduced potential population from which to draw participants. Utilities and their program implementers may wish to segment targeted customers into tranches based on their potential and pursue uptake in waves.

It is our hope that this paper will serve as an example for EE program administrators of the opportunity for data analytics to enhance performance. New program designs that prioritize savings at the meter, including Pay-for-Performance, can realize immediate and significant benefits by implementing the types of targeting strategies detailed here. Doing so has the potential to unlock the potential of EE as a competitive grid resource. While this work focused on cooling energy savings from residential HVAC and building shell programs, customer targeting schemes based on AMI data analytics could also be developed for programs designed to address baseload, lighting, and gas usage. Phase II of this research will focus specifically on gas.

Appendix: Outliers and Null Dataset Results

A1. Outliers

Figure A1 reproduces the results of Figure 5 for AC/QC with outliers included. The solid points show results for the entire sample with the analogous open points displaying results with the top and bottom 3% of savers excluded.




Figure A1: AC/QC - Per capita normalized mean energy savings for each level of targeting applied during the three time periods: Annual, Summer (June – September) and Summer Peak (June – September; 3 – 9 pm). Results for the entire sample are given as solid points while results for the middle 3 – 97% of customers (determined by a rank order of savings) are given by analogous open points.

Analogous to Figure A1, Figure A2 reproduces the results of Figure 6 for Advanced Home Upgrade with outliers included.




Figure A2: Advanced Home Upgrade - Per capita normalized mean energy savings for each level of targeting applied during the three time periods: Annual, Summer (June – September) and Summer Peak (June – September; 3 – 9 pm). Results for the entire sample are given as solid points while results for the middle 3 – 97% of customers (determined by a rank order of savings) are given by analogous open points.

While similar trends are observed with and without outliers, their inclusion had a substantial effect on the absolute values of each point. Generally In each case, excluding outliers has the effect of diminishing per capita savings, usually by approximately 20%.

Other methods may be used to determine outliers more precisely. For example, one normalize the determination of outliers by eliminating customers with the lowest and highest ratio of savings to usage. Ultimately assigning outliers does require professional judgement.

A2. Null Dataset Results

Figure A3 gives results from the null dataset analysis using the conglomerate filter set such that the Loose Moderate and Strict targeting levels eliminated 34%, 66% and 87% of participants, respectively.

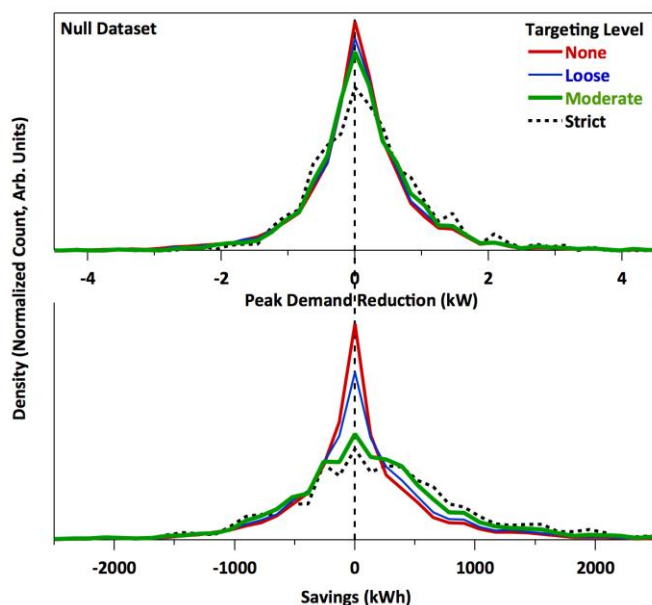


Figure A3: Null Dataset Results. Top – 97th percentile peak demand reduction; Bottom – Electricity savings observed as a function of targeting rigor.

A slight effect that we are assigning as regression to the mean is observed and is more prevalent in the kWh assessment shown in the bottom panel. For evaluation purposes, we recommend that the magnitude of these effects are quantified and backed out for the determination of final results.