Public Version. Redactions in “2018 Load Impact Evaluation for Pacific Gas & Electric Company’s SmartRate™ Program” and appendices.

2018 Load Impact Evaluation for Pacific Gas & Electric Company’s SmartRate™ Program

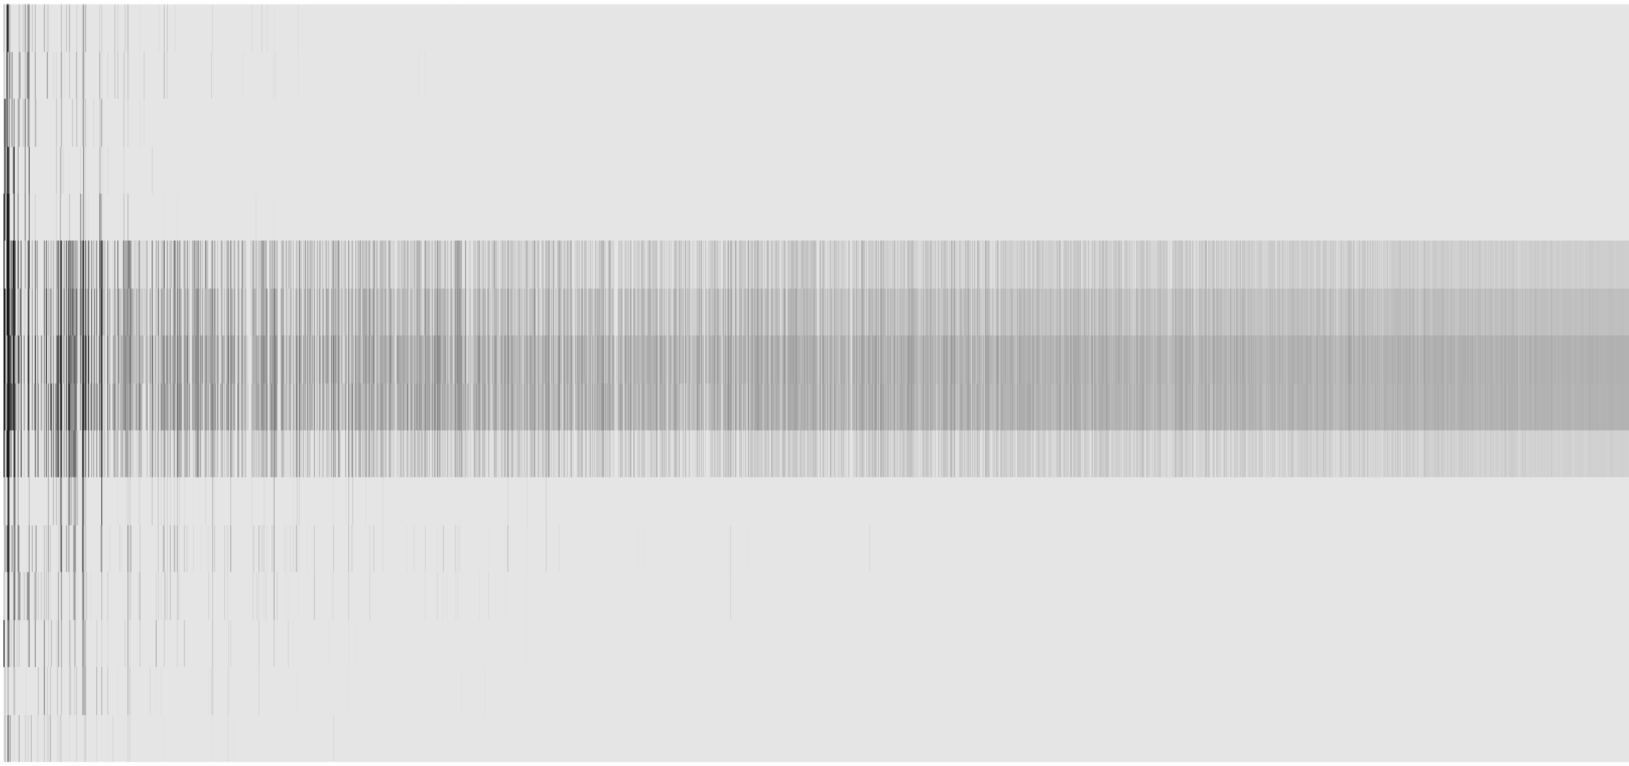

March 27, 2019

Prepared for Pacific Gas and Electric Company

CALMAC Study ID PGE0428.01

Prepared by
Sam Borgeson
Phil Price
David Noonan
Convergence Data Analytics, LLC

Jennifer Mitchell-Jackson
Grounded Research and Consulting, LLC

Convergence Data Analytics

Cover image: Drawing from a pool of 2500 event participants and 7500 matched controls, each column represents a separate difference-in-differences calculation for the same event using a randomly selected subset of participants, growing from 100 to 2500 from left to right. Rows are 8 am to midnight (ascending) on an event day. Color represents the different between average loads of participants and controls (3x per-participant) for each hour.

Contents

Figures

Tables

Executive Summary

This report presents results of a load impact analysis of PG&E's SmartRate™ program. SmartRate™ is a demand response (DR) program developed by PG&E that offers residential customers an economic incentive to shift consumption away from SmartRate™ events. The program uses a rate modifier that applies time varying costs on top of any given underlying rate. It reduces rates by just under $0.024/kWh from June 1st through September 30th, except between the hours of 2 pm and 7 pm on up to 15 event days per year, when rates are increased by $0.60/kWh over the underlying rate.[1] Customers can enroll online when doing a rate comparison in PG&E's Electric Rate Comparison Tool, which calculates electricity costs for customers under different rates using their metered consumption and enables them to select and enroll in any rate plan that best matches their lifestyle. They can also enroll by calling into PG&E and working with a customer service representative. The SmartRate™ price differences provide a financial incentive for customers to save or shift energy consumption away from periods of grid congestion while reducing overall costs for most participants. Customers are notified of SmartRate™ events by 2 pm a day in advance by their choice of phone, email, or text message. Customers in their first full season of the program are also offered a Bill Protection guarantee that credits any cost increases caused by SmartRate™ back to their account at the end of the summer.

In this document we present:

1. Ex post load impacts for the Smart Rate™ program for program year 2018 (PY2018)
2. Ex ante forecast of SmartRate™ for program years 2019-2029
3. A deeper dive into issues of particular relevance to the future of the program

Within these analyses, we examined impacts across geography and by customer segments. This included findings for:

- *Local capacity areas (LCAs).* There are seven California Independent Systems Operator (CAISO) LCAs[2] in PG&E territory, spanning a great deal of geographic/climatic variability: Humboldt, North Coast/North Bay, Greater Bay Area, Sierra, Stockton, Greater Fresno and Kern. The LCA analysis provides insights on the magnitude of available capacity from events in each geographic area.
- *Dually enrolled customers.* Up through the 2018 program year, SmartRate™ customers have been allowed to simultaneously enroll in PG&E's SmartAC™ program. These customers are described as "dually enrolled" or as "duals" for short in the graphics, with all others are described as "SmartRate™ only," abbreviated as "SR only" in figures. The SmartAC™ program installs hardware that can automatically curtail the operation of central air conditioners in response to event signals dispatched by program planners, so at least part of their response is directly controlled by PG&E. Dual customers are no longer being recruited by the SmartRate™ program (per D18-11-029), however, a significant minority of SmartRate™ customers are dually enrolled and understanding the differences between dually enrolled customers and SmartRate™ only customers allows for more accurate impact assessment and forecasts.
- *CARE customers.* PG&E provides discounted rates for low income customers under a program called the California Alternate Rates for Energy (CARE) program. It is a public policy priority to serve CARE customers through energy efficiency and DR programs, including SmartRate™. CARE customers will, on

[1] Participants also receive a participation credit of $.0075/kWh beyond the baseline (Tier 1) level of monthly consumption from June through September (and any days of May or October that fall within those June/Sept billing cycles). Customers on the opt-in E-TOU-B rate, receive an addition $.0050/kWh on all consumption from June through September.

[2] See http://www.caiso.com/informed/Pages/StakeholderProcesses/LocalCapacityRequirementsProcess.aspx for more details on the CAISO local capacity requirements process.

average, have different consumption patterns (i.e. reference loads) and respond differently to price signals than their non-CARE peers.

- *Time of use (TOU) customers.* Only a small fraction of PG&E customers is currently on a time of use (TOU) rate, but TOU rates are slated to become the default choice for customers in the fall of 2020. Since TOU rates also expose customers to costs that vary by hour of day, it has become important to empirically quantify the additional SmartRate™ impacts on customers enrolled in TOU rates. In anticipation of the transition to default-in TOU rates, this year's evaluation includes estimates of program impacts within sub-groups of TOU customers (i.e., default-in or opt-in TOU).
- *Enrollment cohorts.* The SmartRate™ program enrolled customers as early as 2008. This report examines impacts by year of enrollment to provide insights on how cohorts of customers enrolled under different marketing strategies and incentive structures over time use energy, and to understand the impacts recruitment strategies can have on the composition of load reductions during events.

The first three categories (LCAs, dually enrolled customers, and CARE) provide data points to inform our ex ante model, while the latter two (TOU customers and enrollment years) provide additional insights to help steer the program into the future. Load shed estimates based on a more inclusive range of customer characteristics could potentially enable more accurate ex ante forecasts in the future.

This report (delivered March 2019) marks some notable changes in the hours used to create the ex ante forecasts. While past reports have provided ex ante results for a 1 pm – 6 pm window in the summer months, D.18-06-030 modified the CAISO resource adequacy measurement hours (RA window) to 4 pm – 9 pm for each month of the year beginning in 2019. As such, our ex ante results for 2019-2029 are presented for a 4 pm – 9 pm RA window unless a specific comparison is being made.

The event period for PY2018 (2 pm – 7 pm) is consistent with the event period used in past years, but it is not consistent with peak periods defined in time of use (TOU) rates (either 4 pm – 8 pm or 5 pm – 9 pm depending on the rate). The event period is also not consistent with the RA window. It begins two hours earlier and ends two hours before the end of the new RA window. Proposed program changes currently under consideration would shift the SmartRate™ event window later in the day, increasing the overlap with the RA window and peak periods of increasingly popular (and soon mandatory) TOU rates. This information should be considered by the reader as context for this year's forecast.

Ex Post Load Impacts

PY2018 Aggregate Impacts: The aggregate ex post program load impact for a typical day in 2018 was 17.8 MW. Figure 1 presents aggregate impacts (stacked by LCA impacts) for each event day in PY2018, as well as for a typical event day. As shown in the figure below, there were nine event days in 2018.

Figure 1: Aggregate ex post event impact by event date and LCA.

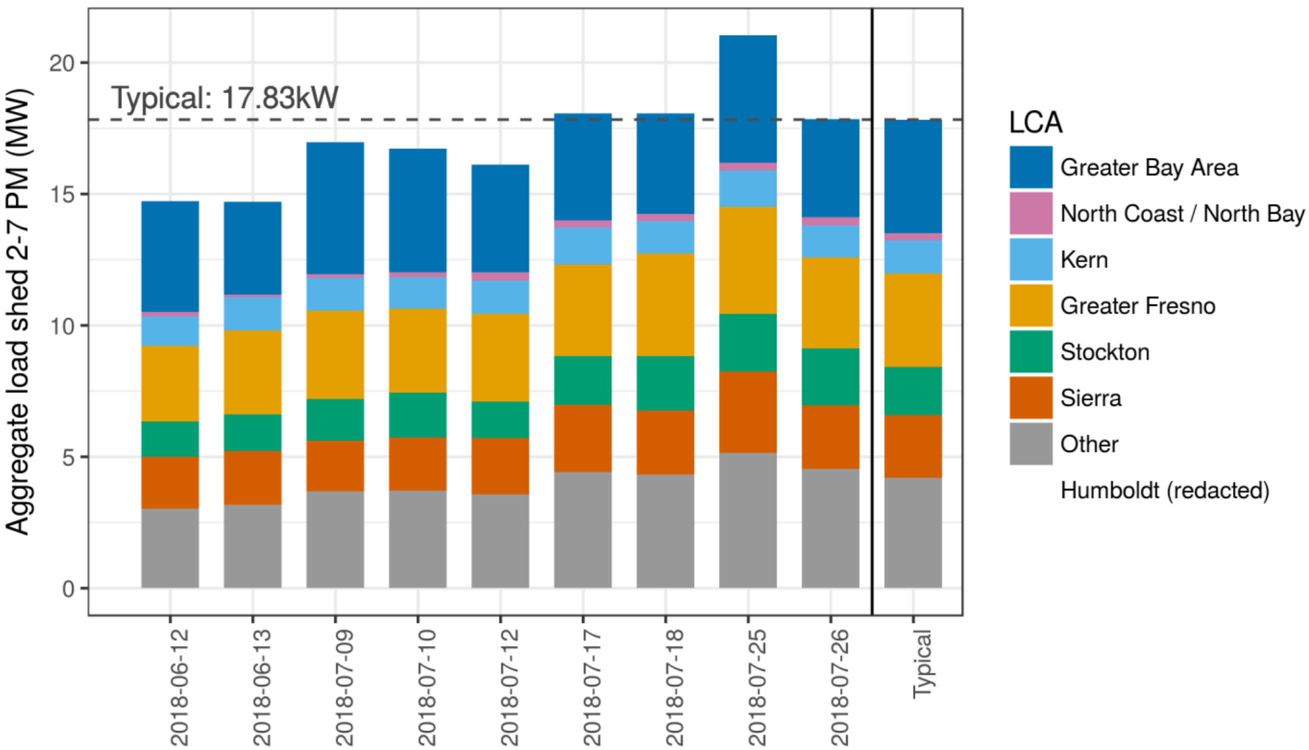

Note: Due to communication issues, some enrolled customers did not participate in the events on June 12 and 13, reducing the impact for those days compared to the others. Those dates are not included in *typical event day* impact calculations in this report.

Comparison to PY2017 Aggregate Impacts: Aggregate ex post results were significantly lower than in prior years, primarily due to lower overall participation numbers and lower temperatures on event days. On a typical event day, approximately 110,000 customers participated in each of the nine SmartRate™ events that occurred in 2018, a decrease of approximately 11% compared to 2017. Cooler event day temperatures led to a 32% reduction in the average number of cooling degree hours (CDH relative to 65˚F) across event days in 2018 compared to 2017, lowering air conditioning and therefore reference loads and responses for event participants. All events occurred in June or July. Unlike in prior years, there were no August or September events.

PY2018 Per-Customer Impacts: The average per-customer impact in 2018 was 0.16 kW with significant variation across local capacity areas (LCAs) and customer segments (see Figure 2). Figure 2 illustrates typical per-customer event average impacts for all nine events called in 2018 (i.e., each circle represents an event) for all categories of customers examined.

Figure 2: Typical per-customer average impacts by category for each of the 9 events in 2018

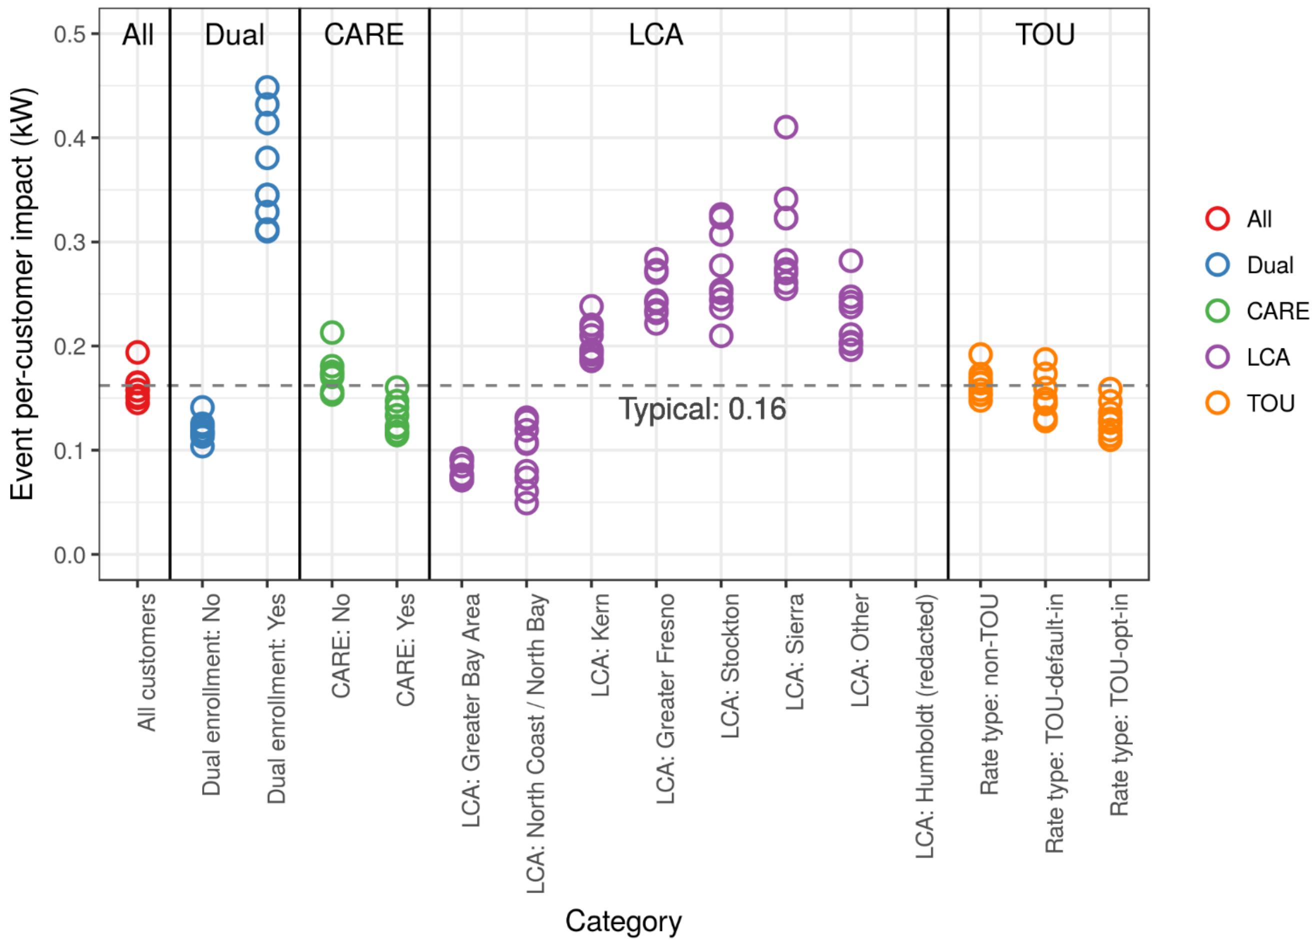

As shown in the figure:

- **Dual enrollment provides the largest impacts per customer**. Their automated event day air conditioner (AC) load curtailment is responsible for most of their performance boost, but dually enrolled customers also tend to live in hotter climates (shown elsewhere) with larger reference loads than SmartRate™ only customers.
- **LCA is the next biggest source of variability in load curtailment**, with the cooler coastal LCAs such as the Greater Bay Area and the North Coast underperforming relative to the typical program average, and the hotter inland LCAs such as Sierra and Stockton outperforming the average.
- **CARE customers tend to produce smaller impacts than customers not on CARE**. This is the case despite the fact that CARE customers tend to live in hotter climates and have larger reference loads than customers not enrolled in CARE. Their load shed as a percentage of their reference load is the smallest of any examined subgroup.
- **Customers on default-in TOU shed more load than customers on opt-in TOU rates.** However, the TOU opt-in customers are disproportionately found in cooler climates with lower reference loads compared to other program participants. Furthermore, the sample sizes of the TOU groups are relatively small. TOU opt-in customers make up 17% of the enrolled population; default-in TOU customers make up just

3% of the population. More data will be required to fully characterize the interaction of TOU rates and the SmartRate™ program.

Comparison to PY2017 Per-Customer Impacts: The 2018 average impacts are lower than in prior years primarily because of cooler weather but also because the proportion of dually enrolled customers, i.e. customers enrolled in both SmartRate™ and SmartAC™, decreased significantly between 2017 and 2018. On average, dually enrolled customers have substantially higher load impacts than SmartRate™ only customers (0.38 kW vs. 0.12 kW for a typical day, respectively). Thus, the decrease in the proportion of dually enrolled customers decreased the overall ('All') average impacts.

Ex Ante Load Impacts

CDA forecasts August peak aggregate impacts of 8.6 MW during an average hour within the 4-9 pm resource adequacy window using the 1-in-2 PG&E weather scenario, and an aggregate impact of 9.8 MW using a 1-in-10 PG&E weather scenario. (See Figure 3.) These estimates assume that SmartRate™ events run from 2 – 7 pm.

Figure 3. Mean and aggregate ex ante impacts for 4-9 RA window (4 scenarios, summer only)

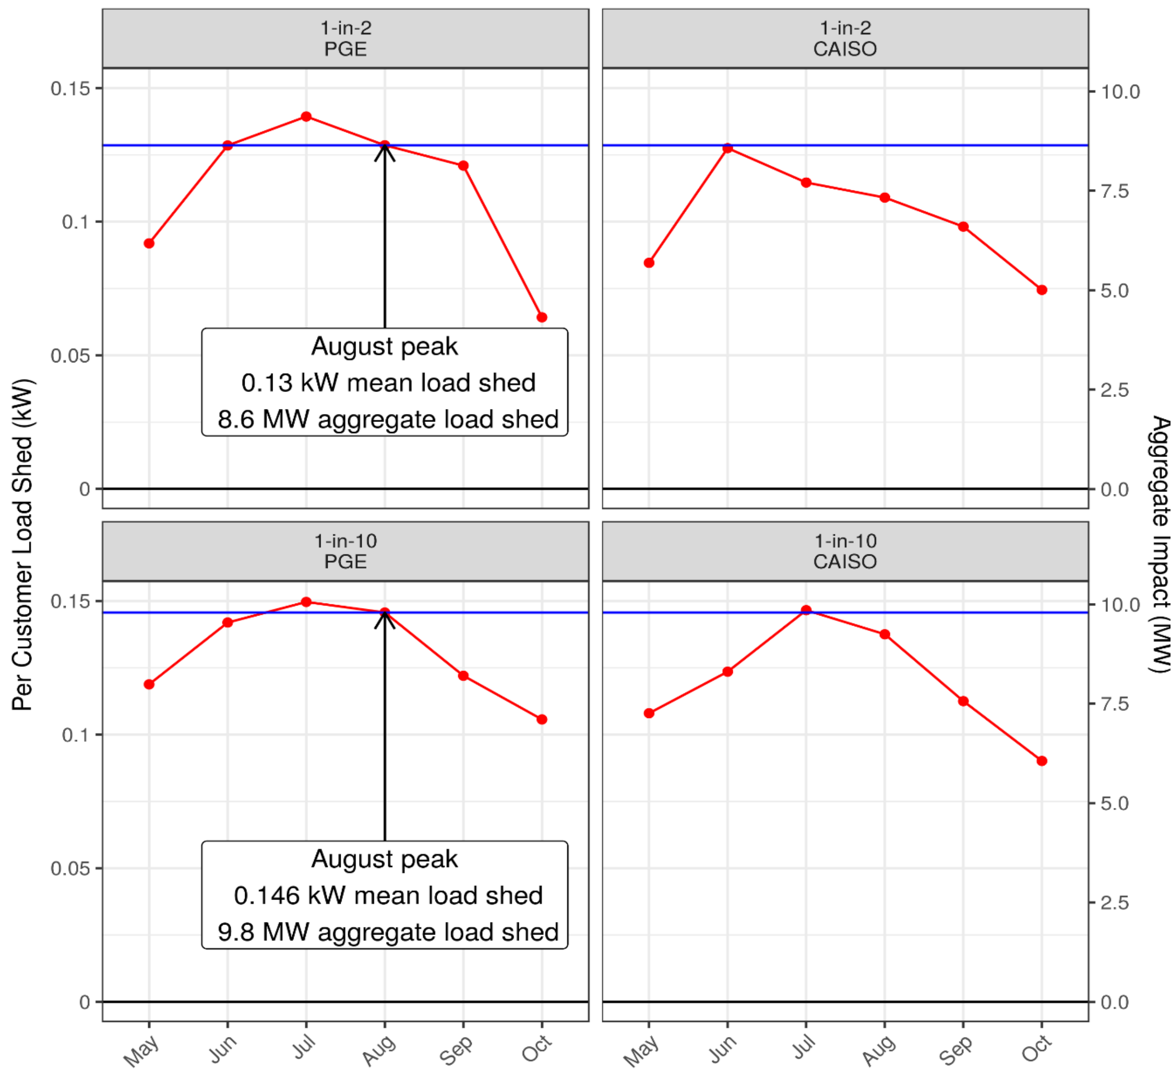

Figure note: The left and right axes apply to all four figures. Mean load shed is indicated on the left, and aggregate on the right, as demonstrated by the August examples in PG&E 1-in-2 and 1-in-10 and the corresponding blue line leading left and right.

Aggregate ex ante estimates for 2019 and beyond are significantly lower than ex ante estimates from prior years. For example, in the previous study's estimate, the August peak day using the 1-in-2 weather scenario was 13.5 MW, compared to 8.6 MW for the same scenario in this year's estimate. While there are several differences between last year and this year's estimates (discussed in the *Current Ex Ante to Prior Ex Ante* section) the factor that led to the largest reduction compared to prior ex ante estimates was changing the RA window. Starting in January 2019, Decision 18-06-030 modified the resource adequacy measurement hours to HE17-HE21 (4-9 pm) for each month of the year—a fundamental change from the prior summer RA window of 1-6 pm. Now, the event period overlaps with this RA window for only three event hours, rather than the four hours that were previously included. In addition, within the new RA window there are two hours of post-event 'snapback' during which electric load is typically *higher than it would have been in the absence of an event* (primarily due to air conditioners working harder to recondition spaces at the conclusion of each event). As such, the average hourly load impact over the window is significantly lower than in past years.

Ex ante estimates are provided for all weather scenarios and all months, but since SmartRate™ events take place exclusively during warm or hot weather in summer the predictions for cool weather and non-summer months are poorly supported by empirical data. The detailed ex ante estimates for all weather scenarios can be found in the *Detailed Summary of Ex Ante Impacts* section.

Conclusions and Recommendations

Based on our evaluation of the 2018 SmartRate™ program, we provide the following recommendations:

Recommendation	Description
Embrace behavioral savings from households with lower per-customer impacts to account for fewer dually enrolled customers	Because PG&E can no longer add new customers who are also enrolled in SmartAC™, the count and therefore impact of these dually enrolled customers will predictably decline through attrition over time. Increased enrollment and greater engagement will be needed to overcome reduced per-household impacts.
Leverage the transition to TOU rates to boost enrollment; prepare for a lot of novices	As the wave of defaulted-in TOU customers use the web-based tool to compare and choose rates, there may well be a boost in enrollment in SmartRate™. If this occurs, PG&E should be ready to recruit and work closely with many customers for whom this type of rate is very new.
Target those with higher potential for savings	Understanding differences among customers within the SmartRate™ only segment will offer the best insights as to how to proceed as the program becomes more dependent on customer behavior (i.e. as the direct load control savings from dually enrolled customers decline).
Shift event times to better align with TOU peak periods and the RA window	The SmartRate™ event window should be harmonized with the now later 4 – 9 pm RA window and the peak periods of emerging TOU rates. Changing the event times will provide larger RA window impacts while also minimizing confusion for future customers.
Explore effects among TOU customers	The geographic difference between the opt-in versus default-in TOU groups are striking and confound any easy comparisons. PG&E should review this study's findings in context of the parallel TOU impact findings and consider steps that could be undertaken to support learning and evaluation related to customers on TOU rates.

Introduction to the 2018 SmartRate™ Program

SmartRate™ is a voluntary critical peak pricing program that overlays a standard electric rate. SmartRate™ is designed to help reduce load on the electric grid on days when resources are constrained, also known as SmartDays™. During summer non-event hours, customers receive a credit of approximately $0.024 per kWh; and between 2 pm and 7 pm on SmartDays™ (i.e., SmartRate™ events) customers are charged a peak-price of $0.60 per kWh over their regular rate. These credits are adjusted slightly for customers on an E-TOU rate.

The program calls a minimum of 9 and a maximum of 15 SmartDays™ in a year. In 2018, PG&E called nine events, and all of these occurred in June and July.

PG&E provided customers with day-ahead notification of SmartDays™ via phone, text and/or email to allow customers to plan for reducing their energy use or shifting their load during event hours. During their first full summer season of program enrollment (and any preceding partial season), customers are covered by a rate protection guarantee that refunds any net costs associated with SmartRate™. Customers on bill protection are credited in their November bill cycle if they didn't save.

Participant Characteristics

Within the SmartRate™ program, we explored impacts by geography and customer segments. The categories that we examined include LCAs, dual enrollment status, CARE status, TOU customers and annual enrollment cohorts.

Local Capacity Areas (LCAs)

As a program designed to reduce loads on demand, SmartRate™ is tracked as a part of a larger framework developed by the CAISO to identify and mitigate grid congestion in areas where it is a persistent problem. That framework divides the congested portions of California's grid into Load Capacity Areas (LCAs), seven of which, Humboldt, North Coast/North Bay, Greater Bay Area, Sierra, Stockton, Greater Fresno and Kern, are in PG&Es service territory. Customers are in a named LCA if they are served by a portion of the grid that experiences congestion. **If their portion of the grid is not congested, customers are said to be in the "Other" LCA, which therefore contains customers from all over the service territory**. Although they are technically defined using the grid topology, LCAs span a great deal of geographic/climatic variability. To support CAISO's need to characterize local resources within each LCA, this evaluation quantifies and reports SmartRate™ program impacts for the sub-groups of customers in each named LCA in PG&E's territory and "Other".

Dual enrollment

Up through the 2018 program year, SmartRate™ customers have been allowed to simultaneously enroll in PG&E's SmartAC™ program.[3] These customers are described as "dually enrolled" or as "duals" for short, with all others described as "SmartRate™ only," abbreviated as "SR only" where necessary in this report. SmartAC™ installs hardware that can automatically curtail the operation of central air conditioners (AC) in response to event signals dispatched by program planners. Like SmartRate™, SmartAC™ is operated for a limited number of events during summer months to curtail grid congestion. This demand response (DR) strategy is known as direct load control. The AC control hardware of dually enrolled customers is activated during SmartRate™ events, providing an automated boost to their load curtailment.

[3] Per D18-11-029, after October 26th, 2018, no new customers will be allowed to be dually enrolled, but existing duals will be allowed to keep their dual enrollment.

CARE

PG&E provides discounted rates for low income customers under a program called California Alternate Rates for Energy (CARE) program. It is a public policy priority to serve CARE customers through energy efficiency and DR programs, including SmartRate™ and to ensure that such programs do not adversely impact them. CARE customers will, on average, have different consumption patterns (i.e. reference loads) and respond differently to price signals than their non-CARE peers. For these reasons, CARE customers are evaluated with impacts reported as a distinct sub-group of customers in this report.

Time of Use rates

The SmartRate™ program was originally enabled by the widespread deployment of SmartMeters™ that record electricity consumption on an hourly basis and allow for the tabulation of costs that vary as a function of time of day or for limited event periods. As a program based on a rate modifier that applies time varying costs on top of any given underlying rate, SmartRate™ is a natural complement to and can be applied on top of the time of use (TOU) rates that were similarly enabled by SmartMeters™. Only a small fraction of PG&E customers is currently on a TOU rate, but TOU rates are currently slated to become the default choice for customers in the fall of 2020. Since TOU rates already expose customers to (continuous and gentler) time varying costs, it has become important to empirically quantify the additional SmartRate™ impacts on customers already enrolled in TOU rates.

PG&E introduced new rates structures and rate changes in 2018 relevant to SmartRate™. Of particular interest is the default enrollment pilot, which defaults selected customers into a the TOU-C rate to learn from customer responses in anticipation of widespread deployment of TOU rates by default, scheduled for the fall of 2020. Just under 20% of all SmartRate™ participants in 2018 were on a TOU rate (this includes both opt-in and default-in TOU customers).

Enrollment Cohorts

Over the history of the program, program managers have recruited different types of customers to the program—partly because of the method of recruitment, and partly because of the types of customer that have been targeted over time.[4] At the same time, rules and the regulatory environment continue to evolve, as do the loads and interests of customers eligible to enroll. All these factors taken together shape the ever-changing composition of new customers enrolled for a given program year. Once enrolled, customers learn how to best participate in the program and make other choices, including leaving the program, over time. For all these reasons, it is useful to look at the differences and similarities of enrollment cohorts as they participate in the program over time. Because enrollments take place year-round but the program events primarily occur in the summer, we define enrollment cohorts using the year of the first event they participated in. For convenience and sample size, we limit the cohorts to pre-2015 and individual years spanning 2015-2018

Participation Trends

At about 108,000, the total enrollment the SmartRate™ at the end of the 2018 program season represents a decline of 13% from the end of the 2017 program year (typical event day enrollment declined by 11% year over year). Enrollment numbers, and thus the make-up of event participants, are being significantly changed as a result of new Community Choice Aggregators (CCAs). By default, all utility customers in the service territory of a CCA are enrolled in the CCA's retail service when they start operating.[5] CCA customers are no longer eligible for SmartRate™ and are automatically de-enrolled. This was the largest reason for "unenrolls" in 2018. Overall, enrollment decreased despite a large 2018 enrollment cohort. Over the same time period, dual enrollment declined by nearly 34% (decreasing from 23% of the population to 17%), CARE enrollment declined by 9% (while

[4] PG&E has not done any active acquisition to SmartRate™ program since 2015.

[5] These customers are given an opportunity to opt-out of the CCA.

increasing their proportion of total participants slightly from 28% of the population to 30% of a smaller population).

This trend is expected to continue. In total, 97,000 SmartRate™ participants have been unenrolled (or are anticipated to be unenrolled) as the CCAs expand.[6] The 2018 results reflect the unenrollment of 34,000 participants. An additional 63,000 "unenrollments" are anticipated to occur before the 2019 season. These changes either have, or will, affect forecasts of participants in this program.

- **2017:** In 2017, 34,000 SmartRate™ participants unenrolled due to CCAs, which impacted the 2018 participation numbers.
- **2018:** In November 2018, another 17,000 unenrolled when East Bay Community Energy started up, and another 23,000 were forecasted to be unenrolled from San Jose and San Francisco before the end of 2019.
- **Early 2019:** In February 2019, when San Jose Community Energy launched, an additional 12,000 unenrolled; and in April 2019, another 11,000 are forecasted to be unenrolled with as the Clean Power San Francisco completes residential expansion.

While 2018 participation numbers decreased because of the CCAs, over 20,000 new SmartRate™ customers enrolled in 2018. (See Figure 4 below). A few LCAs actually saw an increase in enrollment compared to 2017. This was due both to new enrollees, and the realignment of customer assignments to LCAs prior to 2018. The re-alignment had a particularly large impact on the "Other" LCA.

[6] This data was provided to the evaluation team by PG&E.

Figure 4: End of season PY2018 SmartRate™ participants by their year of enrollment

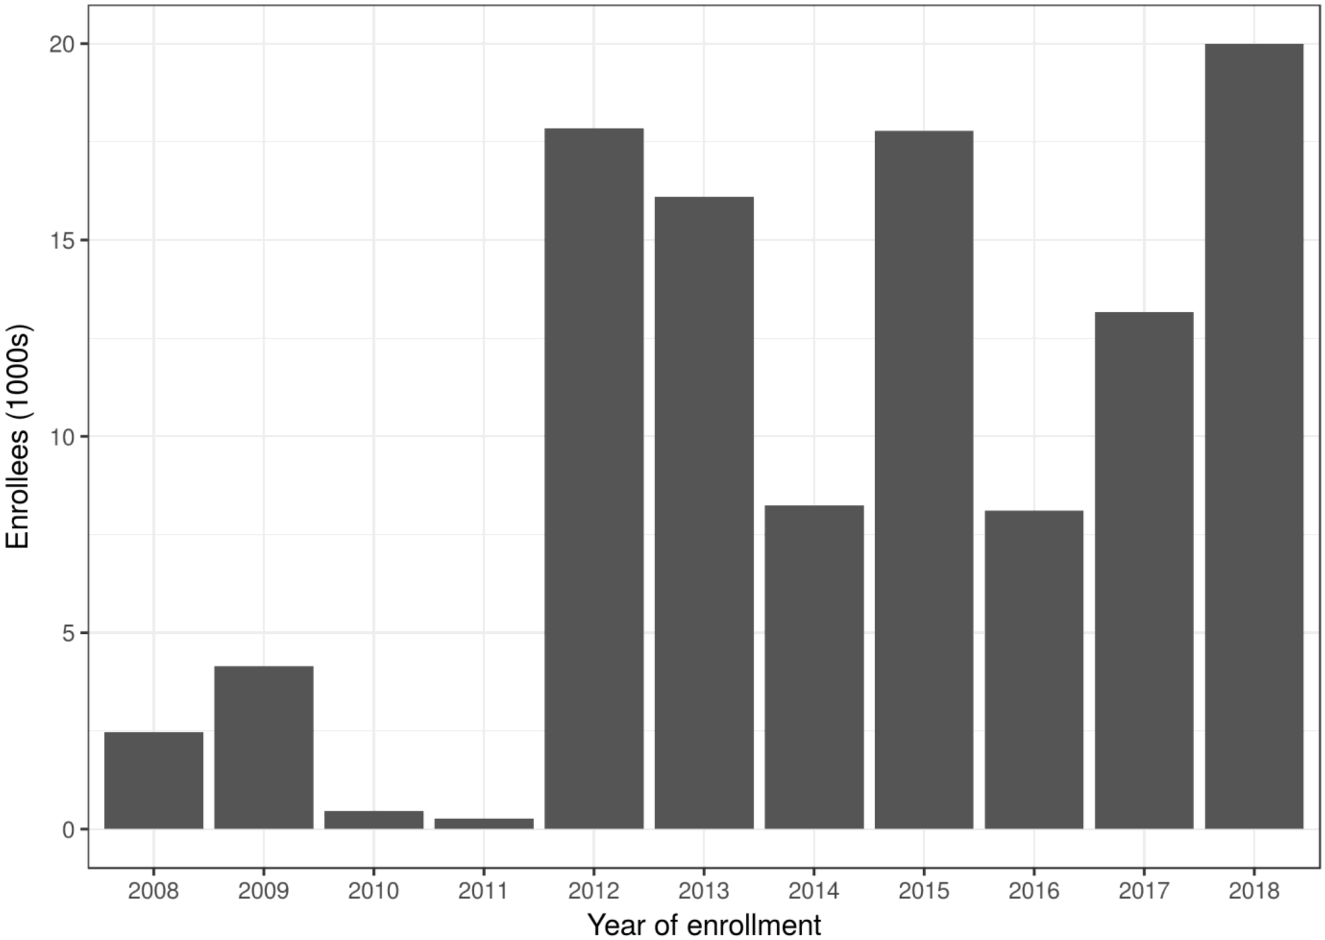

The breakdown by LCA and dually enrolled vs. SmartRate™ only for a typical event day in 2018 compared to 2017 is shown in Table 1, with the LCAs in order from the largest (Greater Bay Area) to the smallest (Humboldt).

Table 1. PY2018 Participants in the SmartRate™ Program

LCA	Dually enrolled	SmartRate™ only	Total	Change from typical 2017
Greater Bay Area	5,858	48,325	54,183	5% Decrease
Greater Fresno Area	3,236	11,121	14,357	4% Increase
Sierra	2,051	5,488	7,540	43% Decrease
Stockton	1,812	4,868	6,680	None
Kern	1,142	5,290	6,432	50% Increase
North Coast and North Bay	400	2,126	2,526	9% Increase
Humboldt	[redacted]	[redacted]	[redacted]	[redacted]
Other (not associated with an LCA)	3,956	14,297	18,254	31% Decrease
Grand Total	**18,455**	**91,515**	**109,972**	**11% decrease**

Table note: These are the numbers at the end of the season. On a typical event day participation was 109,972. Humboldt findings are not called out in this report for confidentiality reasons given the small number of participants in this LCA. The numbers in this table are the result of averaging across event days to get "typical" enrollment. Thus, small differences between totals and sub-groups are possible.

Largely driven by the previously discussed enrollment dynamics, event participation varied by event day. Figure 5 displays event participation by event for both 2017 (faded color) and 2018. The off season CCA reductions create the main discontinuity between years; however, there were communication problems on two dates, (June12 and June 13, 2018), which lead to a dip of about 21,000 participants during those events. As shown in the figure, the communication problem impacted all LCAs on those two dates.

Figure 5. Participation by LCA in 2017 and 2018

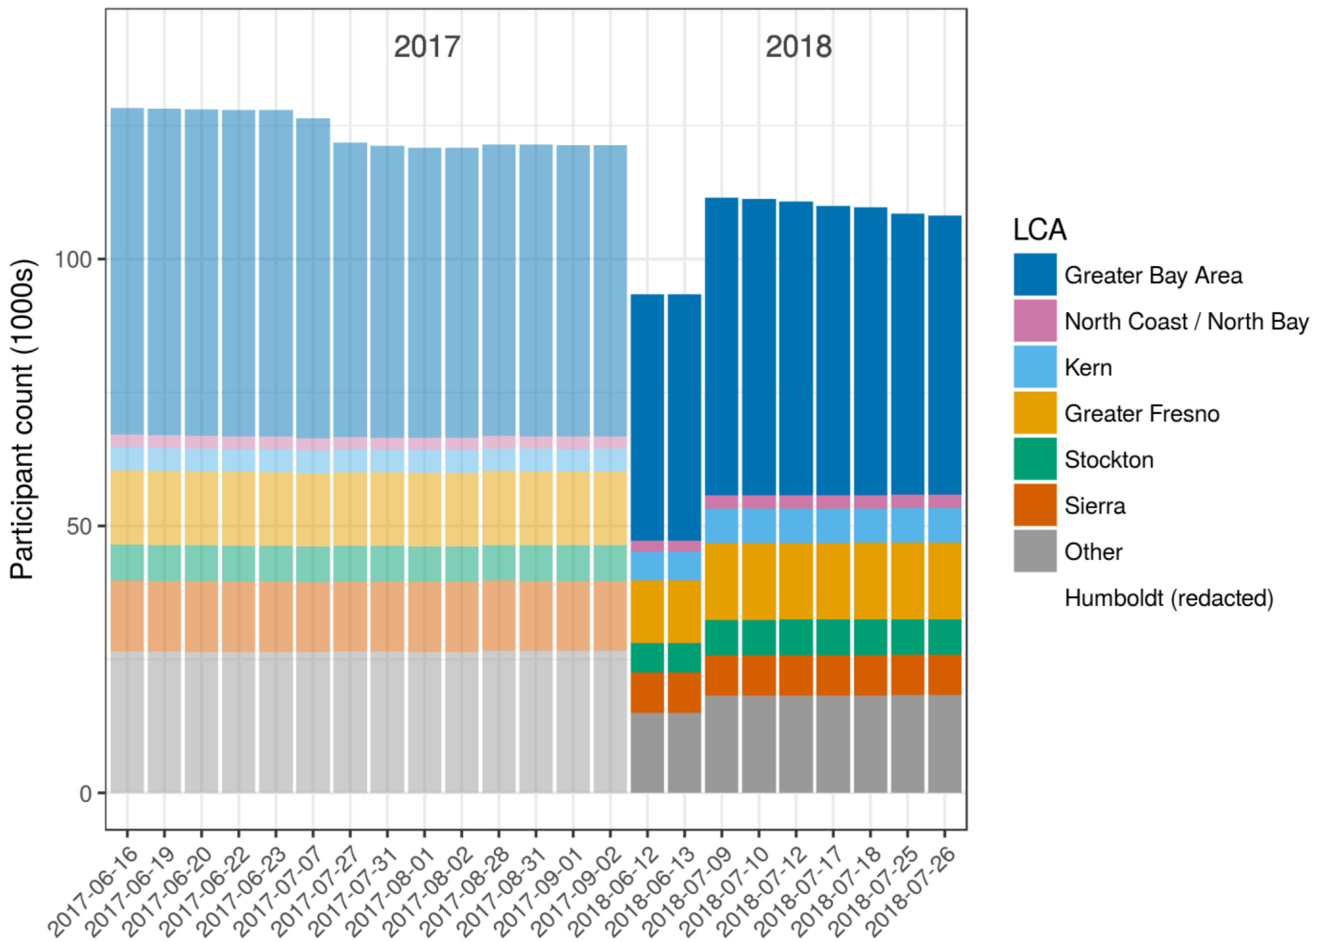

Temperature Trends – Mild Weather in 2018

On average, event temperatures were down 7° F in 2018 compared to the events in 2017. This resulted in a 32% year over year reduction in cooling degree hours (CDH relative to 65°F) in 2018 (See Figure 6). Customers agree to up to 15 events per year in the SmartRate™ program. However, the mild summer weather meant that the conditions on only nine days in June and July met the criteria for an event. As mentioned earlier, no events were called in August or September.

Figure 6. Comparison of cooling degree hours (CDH) 2017 vs 2018

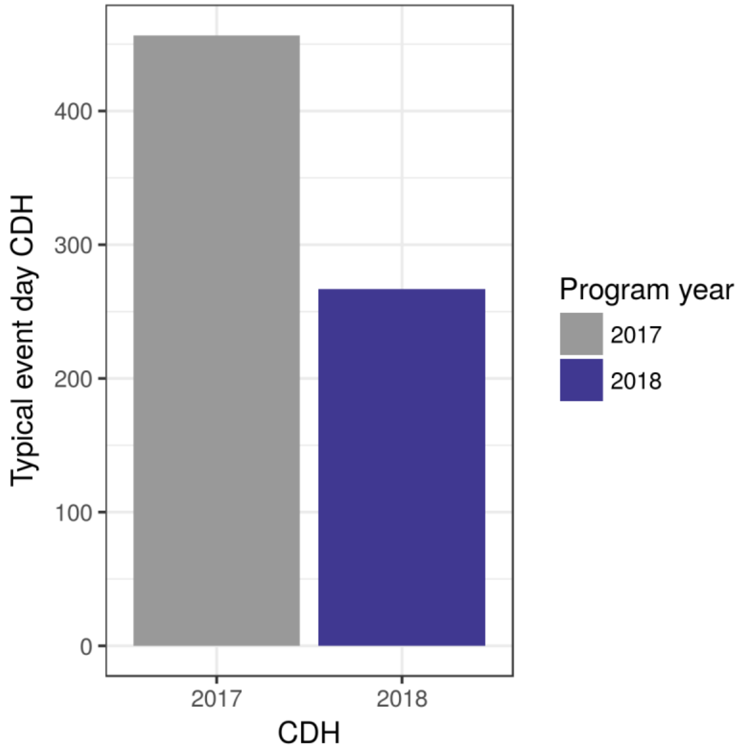

Programmatic Changes and Important Operational Details in 2018

There were no major programmatic or operational changes between 2017 and 2018. For example, events ran from 2 pm to 7 pm on event days in both years. However, it should be noted that the CAISO Resource Adequacy window, formerly set to 1 pm to 6 pm was shifted to 4 pm to 9 pm starting January 2019, meaning that the RA window no longer contains the first two hours of events but does include the two hours after each event.

It is also worth calling out that when SmartRate™ events are called, the AC loads of dually enrolled customers are controlled during the SmartRate™ hours and evaluated as part of the ex post event performance of SmartRate™ customers.[7]

Looking forward from the 2018 program year, there are other changes anticipated. The California Public Utilities Commission (CPUC) determined that no new customers will be allowed to be dually enrolled, but existing ones can remain in the program. The CPUC will also determine whether the dispatch hours of SmartRate™ events should be updated to more closely align with the new RA window and the peak periods of current and future TOU rates.

[7] Ex ante estimates use the specified number of dually enrolled customers, and present both program and portfolio values to avoid double counting, as described in the ex ante chapter.

Key Research Questions and Study Methods

Key Research Questions

The research:

1. Estimates the ex post load impacts for the SmartRate™ program for PY2018
2. Estimates the ex ante load impacts of SmartRate™ for program years 2019-2029
3. Looks at dual-enrollment-, CARE-, and LCA-specific effects
4. Also examines effects by TOU rate type and enrollment cohort

Challenges to our analysis included:

- **Signal to noise ratio:** We note that the effects of the SmartRate™ programs are small with respect to the natural variability within the population. In cases of smaller samples of customers, estimation errors may, at times, exceed the estimated impacts.
- **Segment size:** Some LCAs have a very small numbers of participants. For example, Humboldt county has fewer than 100 SmartRate™ participants. Further, specific combinations of relevant customer attributes can produce small samples. We want to learn from the variability across those segments without making estimates using too few customers.

Ex Post Impact Analysis Methods

For the **ex post analysis**, we estimated load impacts and reference loads for participants on event days compared to similar non-event days (comparison days) and also compared to their matched controls using panel regressions that effectively implement difference in difference estimations. The analysis produced estimates of hourly and event average impacts and reference loads, with errors, for the entire population and relevant sub-group.

We also estimated load impacts for sub-groups, such as PG&E's Local Capacity Areas (LCAs), Dual enrollment status, CARE, and for the first time this year, rate types. A side analysis also introduced the year of enrollment as a sub-group defining customer attribute.

More specifically, we estimated average reference loads and load impacts (both with uncertainties) and tabulated participant enrollment and weighted temperatures for each hour of each event day for every customer sub-group modeled for or reported in this report.

These are the steps we took to arrive at our ex post estimates. Subsequent sections describe each in more detail:

(1) **Identify comparison days**: Match temperature shapes for event days to similar non-event days using weather data.
(2) **Identify potential controls**: Using customer attribute data, including LCA, CARE status, and rate type, we identified a pool of customers from which we would draw matched controls for program participants.
(3) **Match potential controls to participants using meter data**: Compute average load shape for each customer and potential control across comparison days (from 24 hr. a day load data) and match some or all of the resulting load shape features using a similarity distance metric, keeping the 5 "closest" controls for every participant.
(4) **Identify impact model input data**: The model requires participant and control meter data and the local weather data associated with the sub-group(s) of customers whose impacts are to be modeled, with data pulled for the event day and all comparison days.

(5) **Run event models**: Process model data into panel data sets and run separate “difference in difference” regressions for each hour of the day. These regressions do the mathematical analog of computing the difference between participants on comparison and event days (a rough estimate of event impacts) and correcting them with the difference between comparison and event data for the controls, but also allow for normalization against structural drivers of consumption like outside temperature and hour of the day.
(6) **Estimate errors**: Not trusting the error metrics reported by the regression models as under-estimates for various statistical reasons, we estimated model errors by running our event models on comparison days. The correct answer for these non-events is zero impact, so any deviations from zero were taken as empirical model errors.
(7) **Run and store estimates for every customer sub-group**: The basic prescription of steps 1-6 was repeated over and over for every combination of customer attributes defining each sub-group, and for event day and comparison day (i.e. to compute the errors), with 24 hourly estimates of reference loads and load impacts returned with empirical errors alongside of participant counts and population weighted hourly average temperatures.

Identify comparison days

The goal for comparison day matching was to find non-event days with 24hr temperature profiles that were closest by distance metric to event days across all the weather stations associated with program participants. Figure 7 shows the event day and non-event day temperature profiles from those stations for the full summer of 2018. The event day profiles are blue, and the non-event day profiles are gray. It can be verified that the event days are among the hottest days of the summer, but there are good candidate non-event days to match them.

Figure 7. Candidate temperature profiles and event day temperature profiles for all weather stations associated with program participants

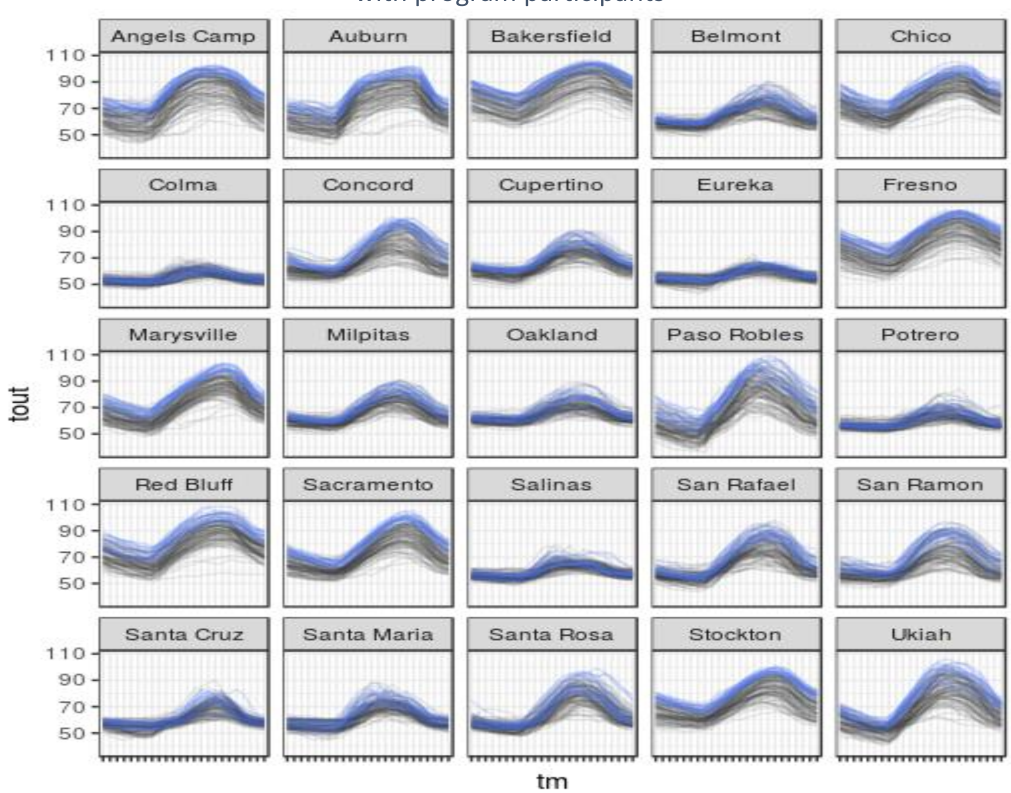

Comparison days were selected based on three criteria. First, day comparisons were made by computing the Euclidean distance between 24-hour temperature profiles for each SmartRate™ event day in the 2018 season and the available non-event days (excluding SmartAC™ events as well) within each weather station. Dates that were matched 3 or more times were kept (N=19). Second, Euclidean distances were computed for the event and non-event average temperature profiles across all weather stations, keeping all matches (N=7). Finally, dates that fell just short of the actual event trigger logic, which is based on the max temperature of the average day ahead forecast for stations Red Bluff, Sacramento, Fresno, Concord, and San Jose exceeding 98F were added (N=6). Many dates overlapped between these methods and the final list had 20 unique comparison dates in it.

There were just 9 event days, but 20 comparison days (see Figure 8). The reason for this generous inclusion of comparison days was to improve the observation of correlation between customers across a wider range of environmental conditions.

Figure 8. Territory wide average temperature profiles for the 9 event days and the 20 comparison dates.

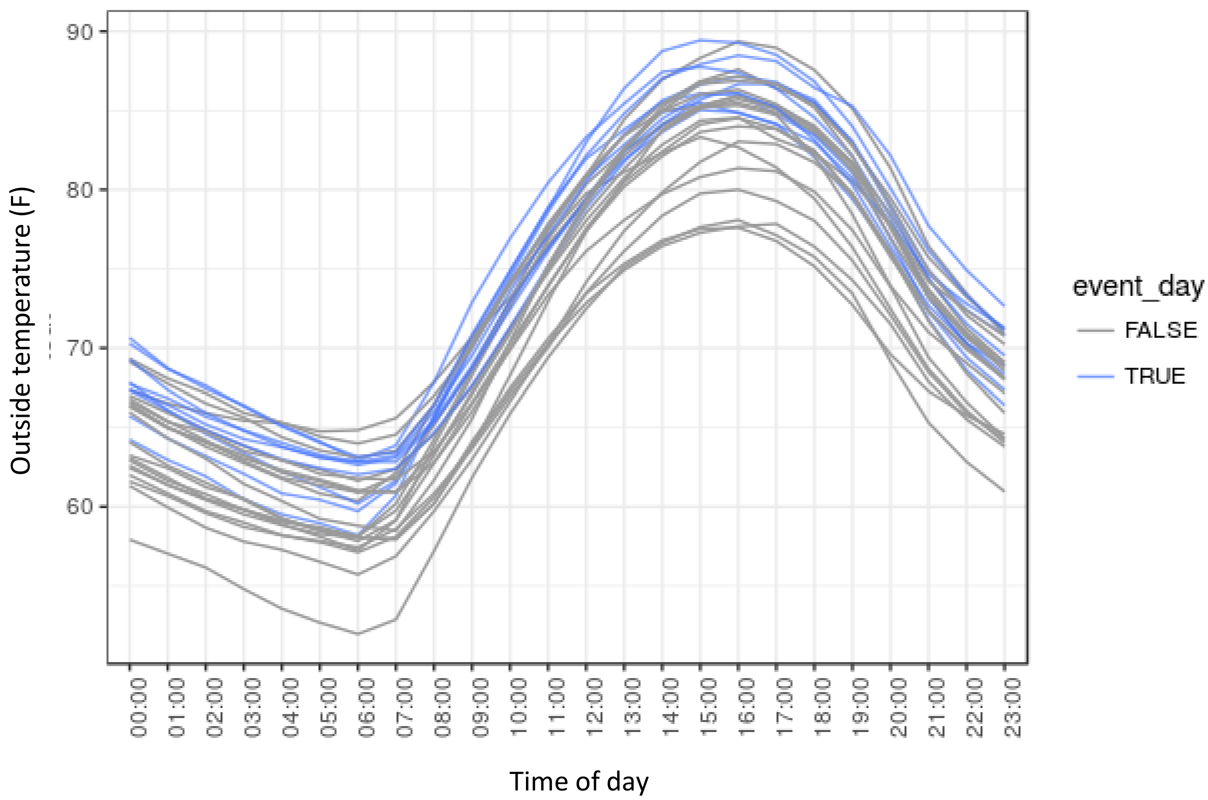

Identify potential controls

Potential controls were identified in two rounds: sampling and matching (and then controls were identified in a third round described below). The goal was to find controls with load patterns and locations similar to the program-enrolled-customers without requesting more sensitive customer data than required to make good fits.

First round: stratified sample

For the first round of potential control matching, the set of enrolled customer account information was used to develop stratified sample counts of potential controls with the strata being LCA, CARE enrollment, and TOU rate type. These criteria were to be run against PG&E's internal account and billing databases. The goal was to keep requested controls roughly in proportion to enrolled sub-category counts while keeping the overall request for billing data under 1M customers. The minimum category request was fixed to about 300 potential controls, boosting counts in the sparser categories, especially to ensure good TOU rate type representation, and the larger categories that would have good statistical power anyway were trimmed to make room. In cases where the request exceeded the total count of customers in a sub-category, it was lower to the actual count and the extra samples were allocated to the smallest sub-categories that could provide more samples. The result was a request for summer 2018 billing and account data for 953,000 potential control customers.

Second round: bill matching

For the second round of matching, our goal was to begin to narrow in on customers with similar consumption patterns as program participants who were also in the same LCA. We computed Euclidean distances between participants and potential controls returned from the round 1 data request on total kWh consumption for the six bills between 2018-04-01 and 2018-10-01 and kept the 10 closest potential controls for each participant.

The result was a request for meter data for a list of 500K customers on all comparison and event days. These customers were the final sample of "potential controls", some of whom became actual controls based on matching patterns in meter data with program participants.

Match potential controls to participants using meter data

Starting with the average comparison day load shape for all participants and potential controls (each hour of the load shape is the average of all the meter readings for that hour across all comparison days), we computed a distance metric between all participants and potential controls *from the same weather station coverage area*. To ensure a robust control sample, we kept the 5 potential controls most similar to each participant. Although a larger number of control customers is good in that it decreases the statistical noise compared to a smaller sample, there is a point of diminishing returns because the fifth-best match to a given participant will not be as good a match as the best match (and the estimation errors are also determined by the number of cases, which cannot be boosted). We selected five controls per customer as side modeling demonstrated that that number balance the tradeoffs between reduction of small-sample noise and the bias introduced by choosing control customers who are less-perfect matches.

Using the R package MatchIt, we modeled different characteristic "features" of the load data using different distance metrics. Based on comparisons of aggregate load shapes of participants and controls, a single "morning load" feature averaging the load from midnight to noon and the loads from individual hours from noon to midnight were identified as reliable matching features. The distance metric that matched well for the customers studies is called Mahalanobis distance. It is similar in concept to adding up the absolute difference between every feature (or taking the sqrt of the sum of the squared difference to find Euclidean distance), but it includes variance corrections that down-weight "distance" derived from highly variable features and up weight stable features so good fits on the more stable features are not automatically canceled out by the noisy ones. This matters here because load variance is often at its highest when consumption is at its highest. Also, more subtly, the morning load metric was already stabilized by averaging several hours together.

The same potential control can match any number of participants. Also, the matches made using this approach are fixed – the same SmartRate™-enrolled customer always has the same matches. As a practical matter, after a long running batch process that did the matching concluded, looking up and using matches was very efficient.

Identify model input data

Customer sub-groups were defined by combinations of customer attributes called LCA, Dual, CARE, TOU-type, and enrollment cohort (the latter covering first program year of enrollment). A sub-group is defined by values for each of these attributes, where the value ‘All’ indicates that the attributes shouldn’t be used to filter the sub-group and a specific value indicates that only customers with that value for that attribute should be included in the group. For example, Dual=’Yes’, CARE=’Yes’, LCA=’All’ would identify all dually enrolled CARE customers from any LCA.

Based on the list of customers in the sub-group identified, we loaded all individual meter data for comparison and event days and weather data from the weather station assigned to each customer (PG&E assigns weather stations to each customer to support analysis involving weather data).

Run event models

Impact estimates and reference loads are based on running 2 different models for each of the 24 hours hour of event days. The input data (the same for both models) is designed to be run as a panel regression with data for the sub-group of customers whose load impact is to be estimated. These data columns serve as the variables:

Table 2. Data columns involved in panel estimation of hourly load impacts

Data column	Explanation
kWh	The total kWh consumption for the hour. This is the variable on the left-hand side of the regression equation to be explained by all the other factors.
date	Date of the electricity consumption. The dates in one panel will include one event day and all comparison days.
customer_id	The unique identifier of the customer each meter reading belongs to. The customer_ids in a panel will include all the event participants (the cases) and all of their controls.
hour	Hour of day of the electricity consumption indexed to 1 through 24, with 1 spanning midnight to 1am and so on. The panel will have been filtered to a single hour of day prior to the estimation of the load impacts for that hour.
case	Indicator that is T if the reading belongs to an enrolled customer, or F if the reading belongs to a control customer.
event	Indicator that is T if the reading is from an event day or F if the date is a comparison day.
case_event	Indicator that is T if the reading in question falls during an event day and belongs to an enrolled customer or F otherwise.
cdh	The total number of degrees by which the average temperature for the date and hour exceeds 65F at the nearest weather station to each customer. This is used to quantify the air conditioning (and other temperature sensitive load) contribution to the load data.
case_cdh	Additional temperature sensitivity term only estimated for enrolled customers (i.e. `cdh` for cases, 0 for controls)
early_aft_kWh	Average consumption from 11am to 1pm, added to the model to allow it to perform a same day correction that recalibrates impacts to near zero just before an event.
case_early_aft_kWh	Additional same day correction term only estimated for enrolled customers (i.e. `early_aft_kWh` for cases and 0 for controls)

In R’s formula notation, the hourly ordinary least squares load impact model, which assumes that the data has already been filtered to a single hour of the day, is:

```
kWh ~ case + event + case_event + cdh + case_cdh + early_aft_kWh
```

This means that the total energy consumed for a given hour (kWh) is explained by a regression model fit that estimates multiplicative coefficients for each of the listed variables. Because the data is a panel data set spanning readings from many customers days the fits apply to the average outcome across all modeled customers. This particular model formulation is an extension of the difference in differences evaluation methods used with treatment and control groups. Under difference in differences, the difference in loads for participants between non-event-days and the event day (an approximate estimate of load impacts) is corrected for outside influences on consumption using the difference in loads for controls between non-event-days and event days. If there are no outside influences impacts consumption in the control group, the correction term will be zero and the change in participant loads will stand.

With our regression model, we are effectively implementing a difference in difference calculation that can also control for outside temperature and make a mid-day impact correction to ensure that event impacts start near zero prior to each event. The load impact term is the case_event coefficient, corresponding to the difference in consumption seen in cases during event days on the hour in question, after controlling for event day consumption for controls and comparison day consumption of both participants and controls, but with weather normalization and mid-day-correction built into the estimates.

"Reference load" is the load the participants would have experienced on event day if there were no event. In R’s formula notation, the hourly reference load model, which assumes that the data has already been filtered to a single hour of the day, is:

```
kWh ~ case + event + early_aft_kWh + case_early_aft_kWh
```

This model is run on the same data as the load impact model, but the resulting fit is then used to predict kWh when the input data is forced to be `case`=T and `event`=T. In other words, it predicts the load that the participants would have experienced if there hadn’t been an event.

The load impact and reference load models are run 24 times each for every customer sub-group and event day combination, with their outputs stored as 24-hour load shapes. Although the regression models provide error estimates, we calculate reported errors empirically as described below.

Estimate errors

There are two important sources of differences (aka errors) between a model and the phenomena it is trying to capture: measurement error and model error.

Measurement error comes from the fact that data, especially whole home electricity data in this work, can be imperfectly measured and noisy and full of ups and downs unrelated to the driving forces that determine event savings. This variability of the data itself around its central tendency is what regression standard errors capture. Regression model standard errors are confidence ranges on the average value of the effect being studied. The more data available, the more confident the model gets. Additionally, for their standard error estimates to be accurate, regression models require the data supplied to them to be uncorrelated observations. Time series data is full of correlated observations, a phenomenon called serial correlation. This means, that the model will be over-confident in its estimates and the regression errors will be systematically too small.

Model errors have to do with the structure of the model and the choices and assumptions of the modeler. A simple form of model error associated with regressions is over-fitting. The way a regression model works, the more parameters you add to it, the lower its standard errors get, regardless of whether it is actually being improved. More generally, modelers can get the structure of their model wrong[8]. For example, if a modeler observes that loads decrease when the outside temperature drops from 90F to 80F, i.e. due to reduced air conditioning loads, the modeler might not know that the AC is on a timer and cannot run at certain times regardless of the temperature or a model might be inadvertently structured to expect the loads to keep decreasing as the temperature drop from 50F to 40F, even though most AC units would be off at that point with nothing left to reduce. There is no way for a model to report its model error and there may not be a way for a modeler to know what the structural errors of their model are.

If you know the correct answer, there is an alternative to model-reported errors when quantifying the uncertainty of estimates. You can simply quantify the difference between the model and the answer. We do not know the true event impact from events, but we do know the true impact from non-events: zero. There can be no impact if there is no event[9]. Our approach to errors estimates was to run the event models on non-event (aka comparison days) days using data from the same participants and controls used to evaluate event days. By definitions, all deviations from zero impact on comparison days are errors so each comparison day gives us a different estimate of errors separately for every hour of the day. All the comparison non-event days can be averaged to provide a good estimate of the errors associated with the particular sub-group and its controls. These errors are the same across all events because the same comparison days apply to all events.

Run and store estimates for every customer sub-group

As a practical matter, much of the work done to provide ex post estimates revolved around the data management, memory consumption, and CPU utilization associated with the estimation process. There was also work done to configure and pre-compute all the permutations of customer sub-groups associated with each round of estimation and the necessary data structures and storage formats required to store and retrieve the results. All of this infrastructure was encapsulated in a “batch run” framework that allowed us to run and store estimates on any number of arbitrarily defined customer sub-groups. The final step of the ex post estimation process was running the estimates after putting all the other pieces in place.

There were three sets of batch runs made in support of this evaluation.

(1) The **“official run”** created and modeled impacts for sub-groups for every combination of LCA, Dual enrollment status, and CARE enrollment status, including interactions between Dual and CARE not previously modeled. This run forms the basis of all official data products, including the table generator, and many of the tables and figures in this report. The event outputs of the official run were also used as the inputs to the ex ante modeling, which quantified the effects of outside temperature, hour of day,

[8] In fact, there is no such thing as a perfect model. On this topic, the statistician George Box famously wrote “… all models are approximations. Essentially, all models are wrong, but some are useful.” He was talking about model error.

[9] We are aware that the program can and very likely does have spillover effects on no-event days, but our job as evaluators of ex post impacts is to assess the impacts of calling an event vs. not calling an event because that is the dispatchable resource.

dual program enrollment, CARE enrollment, and LCA for the purposes of forecasting future event outcomes.

(2) The **"TOU-type run"** modeled impacts for sub-groups for every combination of LCA, TOU-type, and Dual enrollment status. While the number of TOU customers (especially TOU-default-in) is small overall, they are anticipated to be the leading edge of a territory-wide roll out of default-in TOU rates and are therefore of particular strategic interest to program planners and (one assumes) future program evaluators. This run's outputs support all the analysis presented in the "Deeper Dive" TOU analysis found in this report.

(3) The **"enrollment cohort run"** modeled impacts for sub groups defined as SmartRate™ only, non-TOU, and every enrollment cohort from 2015 through 2018 plus a cohort for all customers enrolled since before 2015. This run's outputs support all the analysis presented in the "Deeper Dive" Enrollment Cohort Analysis found in this report.

Validity assessment of the study findings

Errors in model estimates generally come from two sources. The first source is sampling error related to the precision or representativeness of the underlying data. The meter data is not a perfect record of just the loads that are in play for DR events. For example, whole home meter data contains lots of random excursions caused by occupant behavior and device operations whose timing and magnitude are not correlated with any externally observable information. Sampling error tends to decline with sample size, but with diminishing returns such that it never drops to zero. The second source is model error, which is when assumptions or model structure don't match the true underlying conditions. There is error associated with all models. None are perfect. Moreover, model errors are not always observable or verifiable. The job of a modeler is to use their experience and professional judgement to match their methods and models to the process being studied until testable sources of error are minimized and well characterized and the model is useful, without claiming to have perfected their approach.

Our modeling task was inherently about comparing two different time periods: comparison days and event days. A potential model error would be the assumption that event days will look exactly like the average of comparison days absent events. Of course, there are factors that mean each day is not like the next. To address this concern, our approach employed a matched control group to help quantify the degree to which even and comparison days don't match for otherwise similar customers.

When models are given too little data to work with, their estimates are inherently noisy. For this reason, all of our estimates involve panels of data drawn from sub-categories of customers, with average effects estimated across all of them at once. Furthermore, individual estimates are themselves probabilistic in nature and do not just add or average in a simple manner. For this reason, we estimate and report group averages, not individual estimates. We also draw upon 5 control customers for every participant so that the control side of the estimate is a stable as possible and the errors are dominated by the properties of the participant data.

Modeling and data decisions are about reducing the model errors, but they still need to be accurately tabulated and reported. The regression models used to make our ex post impact estimates come with a set of assumptions about the modeled data and the validity of the model structure that have to hold for their coefficient standard errors and prediction error estimates to be valid. For example, error estimates decline when applied to time series data. Error estimates also decline with the number of model parameters and by definition assume the model structure is valid. The upshot is that in this context, the regression model errors should be viewed as underestimates.

Perhaps the most important step we took to report accurate errors was to make empirical estimates of errors based on running the models on non-event days when the correct answer, error load impact, was known. See the Estimate Errors section above for more details).

In the end models will have errors that need to be studied and reported. As a start, the estimate percentiles reported in the table generators are derived from the empirical error estimates described above and meet the reporting requirements of Protocol 24. As a metrics of error performance, we looked at the absolute value of hourly standard errors and the ratio of those standard errors to the underlying impact estimates.

Our main ex post analysis is made up of 16,700 hourly event estimates. Of those 3,585 are during event hours with impacts greater than 10 Watts (W). For that group, tabulated the value of standard errors and the ration of standard errors to estimates.

Table 3. Counts of hourly estimates by error characteristics

Estimated standard error (in Watts)	**Count**
<10 W	664
10-20 W	702
20-50 W	1878
50-100 W	323
>100 W	18
Standard error to impact ratio	**Count**
<0.1	1381
0.1-0.2	1138
0.2-0.5	746
0.5-1	215
>1	105

There are just 341 hours with standard errors over 50W and 320 whose magnitude is half of the impact estimate or more. In most of those cases, the sample size is to blame for large errors. Due to the small sample sizes these poor estimates have minimal effects on larger averages or aggregates. Overall, we are satisfied that errors are small enough to evaluate the program accurately. The exceptions tend to re-enforce the larger conclusion that only isolated hours and mostly small samples experience significant errors.

Ex Ante Impact Analysis Methods

The ex ante load impact predictions extrapolate the ex post predictions to future years, given assumptions about future enrollment and event-day weather. Ex ante load impact estimates represent the expected per-customer average and system-wide aggregate load impacts that would occur during a SmartAC™ event under normal (1-in-2 year, i.e. temperatures that have a 50% chance of occurring in a given year) and extreme (1-in-10 year) weather conditions. Ex ante results serve two purposes: 1) they assist PG&E and the State with long-term resource planning, and 2) they allow PG&E to assess year-to-year changes in the program's effectiveness.

We used the same basic approach taken in past evaluations: analyze data from the 2018 SmartRate™ events to determine the load impact for each event and each customer segment; develop a statistical model based on the ex-post load impacts to allow load impact to be predicted based on outdoor temperature and other explanatory variables; use this model to predict the load impact for future events under standardized weather conditions; and perform a similar procedure to predict reference loads (that is, what will the average loads be under those conditions if no SmartRate™ event is called).

We performed three key steps.

1. Find the ex post load impact for each event using the approach described above in the section on the 'Ex-Post Impact Analysis' for each hour of the day, each LCA, and separately for each combination of (SmartRate™-only and dually enrolled customers) and (CARE and non-CARE customers). For each SmartRate™ event this yields one number per hour of the day, in each LCA, for each combination of CARE status and Dual enrollment status.
2. Fit a linear regression model to the ex post load impact. Separate models were used for pre-event, during-event, and post-event ("snapback") periods.
    a. We assume that during the period more than two hours prior to the event there are no event impacts: load shed is set to zero during the early hours. The two hours immediately prior to the event are fit with a regression model.
    b. The during-event model predicts the load impact as a function of temperature, with a different relationship between temperature and load shed for each combination of CARE and dually enrolled status. Within each customer category, the model assumes a linear relationship between temperature and load shed for temperatures above 70˚ F, and assumes the load shed for temperatures below 70˚ F is independent of temperature. The slope of the shed-vs-load relationship is different in the first and fifth hour of an event than in the intermediate hours. Additional terms in the modeling equation adjust for differences in behavior between LCAs where the ex-post data show that those differences are large enough to be important; for example, dual-enrolled customers in Stockton, Fresno, and Sierra have different shed-vs-load relationships than dual-enrolled customers in other LCAs.
    c. The period after the event is likely to experience "snapback", in which air conditioning has to work harder than normal in order to cool the residences back to their desired temperature once the event has ended. This results in increased electric load, i.e. negative load shed. Snapback may also include increased load due to people delaying activities such as laundry or dishwashing until after the event. The model used for predicting load after the event ends is conceptually similar to the one for predicting load shed during the event: it is a linear model that predicts load impact separately in different LCAs and customer categories. One of the inputs to the snapback model is the predicted load shed in the final hour of the event: empirically, higher load shed leads to larger snapback, and the model incorporates this.

In order to provide estimates of the relative (percentage) load shed due to SmartRate™, reference loads are also needed for the weather scenarios. This is provided by a model that is also similar to the event-period model.

3. Use the predictive model to make forecasts for the hourly load impact in the required weather scenarios. For each scenario and each month of the year, summarize the ex ante predictions with the mean and maximum load sheds that occur during the RA window. In this report the ex ante predictions assume future SmartRate™ events will take place during the same 2-7 pm event window that has been used in the past, including 2018, since that is still the official event window for future events. The events are evaluated over the new RA window of 4-9 pm.

Customers in the LCA called "Other" required special treatment. Most LCAs are fairly compact spatially, but the "Other" LCA includes customers from all around PG&E's service territories: Any customer who is not in a spatially defined LCA is included in Other, so Other customers experience a much wider variety of weather than customers in the named LCAs. For modeling purposes most customers in the Other LCA were assigned to a spatially-defined LCA based on the customer's weather station: if a customer in Other is in the same weather zone as many SmartAC™ customers in a spatially defined LCA, that customer was assigned to that spatially defined LCA fitting the model; that is, the model was fit as if that customer were just like any other customer in that LCA. The statistical model was fit and used to make ex ante predictions for each LCA, and then the Other customers were effectively reassigned from the LCA in which they were modeled back into the "Other" LCA to make the ex ante predictions. As a simple example, if 40% of the customers in the Other LCA had been assigned to LCA 1, and 60% had been assigned to LCA 2, then the predicted load shed for Other would be a weighted average of the load shed in LCA 1 and LCA 2, with weights of 0.4 and 0.6 respectively.

In brief, the predicted mean load shed per customer for a given customer segment, for a given event hour, in a given LCA, is:

Load Shed (kW per customer) = (sum of applicable additive coefficients) +
*(sum of applicable multiplicative coefficients) * (T – 70˚ F)*

All of the terms and the details of the model are described in Appendix B.

Ex Post Results

2018 SmartRate™ ex post results on a typical 2018 event day were 17.83 MW in aggregate or 0.16 kW per customer. Table 4 summarizes the outcome of every SmartRate™ event of the 2018 season.

Impacts were observed to vary by date, LCA, CARE status, and dual enrollment status but are generally lower than in prior years due to cooler weather, the de-enrollment of customers who were transitioned to new community choice aggregators (CCAs), and a decline in the number and percentage of dually enrolled customers in the population. The Humboldt LCA contains fewer than 100 customers, too few to report their results publicly while still protecting the privacy of those enrolled customers. There are also too few customers to materially influence the average or aggregate event results reported in this document. For these reasons, Humboldt results are excluded or redacted from all tables and figures in this report.

Table 4 summarizes the outcomes for all nine 2018 event days and the representative typical event day. At this highly aggregated level (i.e. the totals/averages across all LCAs), reference loads roughly track temperatures and per-customer impacts track reference loads.

Table 4. Summary of 2018 events, conditions, enrollment, and impacts.

Date	Number enrolled	Per-cust. ref. load (kW)	Per-cust. load impact (kW)	Agg. ref. load (MW)	Agg. load impact (MW)	% Impact (% of ref.)	Average temp. (F)
2018-06-12*	93,340	1.25	0.16	116.36	14.73	13%	87
2018-06-13*	93,359	1.35	0.16	126.08	14.70	12%	87
2018-07-09	111,536	1.35	0.15	150.98	16.96	11%	89
2018-07-10	111,296	1.47	0.15	163.56	16.72	10%	88
2018-07-12	110,803	1.43	0.15	158.24	16.12	10%	86
2018-07-17	109,878	1.52	0.16	167.03	18.07	11%	88
2018-07-18	109,640	1.58	0.16	172.84	18.07	10%	88
2018-07-25	108,463	1.63	0.19	177.33	21.04	12%	90
2018-07-26	108,185	1.59	0.16	172.14	17.85	10%	87
Typical*	109,972	1.51	0.16	166.02	17.83	11%	88

Table note: *Due to signaling problems on 6/12 and 6/13, not all customers were notified of the events. This explains the lower enrolled participation numbers on those days. Due to their unusual conditions, those events are excluded from the typical event day calculations in this report.

Aggregate Ex Post Summary

The ex post aggregate load impact on a typical 2018 event day was 17.83 MW (down from 28.1 MW in 2017), with the largest aggregate load shed (21.04 MW) occurring on July 25th, the hottest day. Figure 9 depicts the breakdown of aggregate impacts across LCAs for all 2018 events.

Figure 9. Aggregate 2018 ex post load shed (overall, by date and LCA)

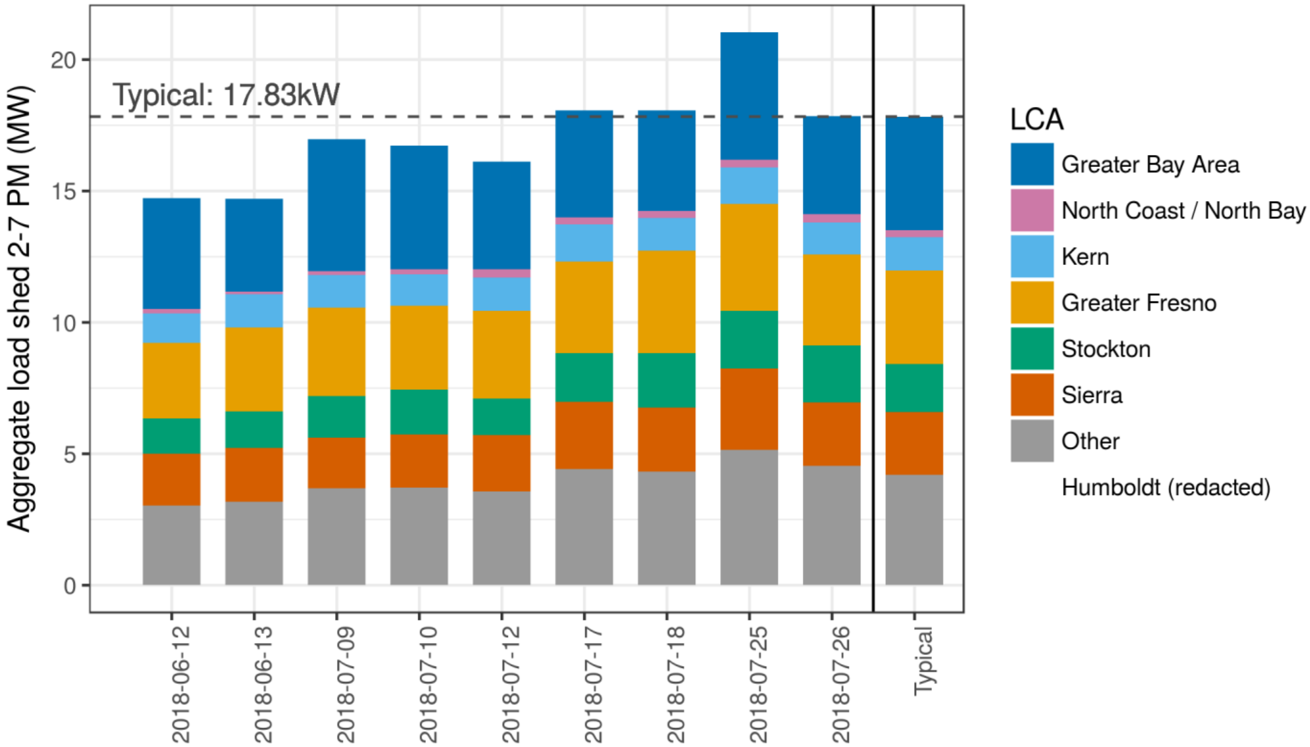

The events of June 12th and 13th suffered from decreased participation due to technical problems and are therefore not directly comparable to other event days. As such, they are not part of the typical day calculation.

Figure 10 illustrates the aggregate load impact contribution of every sub-group relevant to program performance for each of the nine event days. For each event day, the sum across each set of categorical sub groups totals the “All customers” aggregate impact.

Figure 10. Aggregate load impacts by customer sub-groups

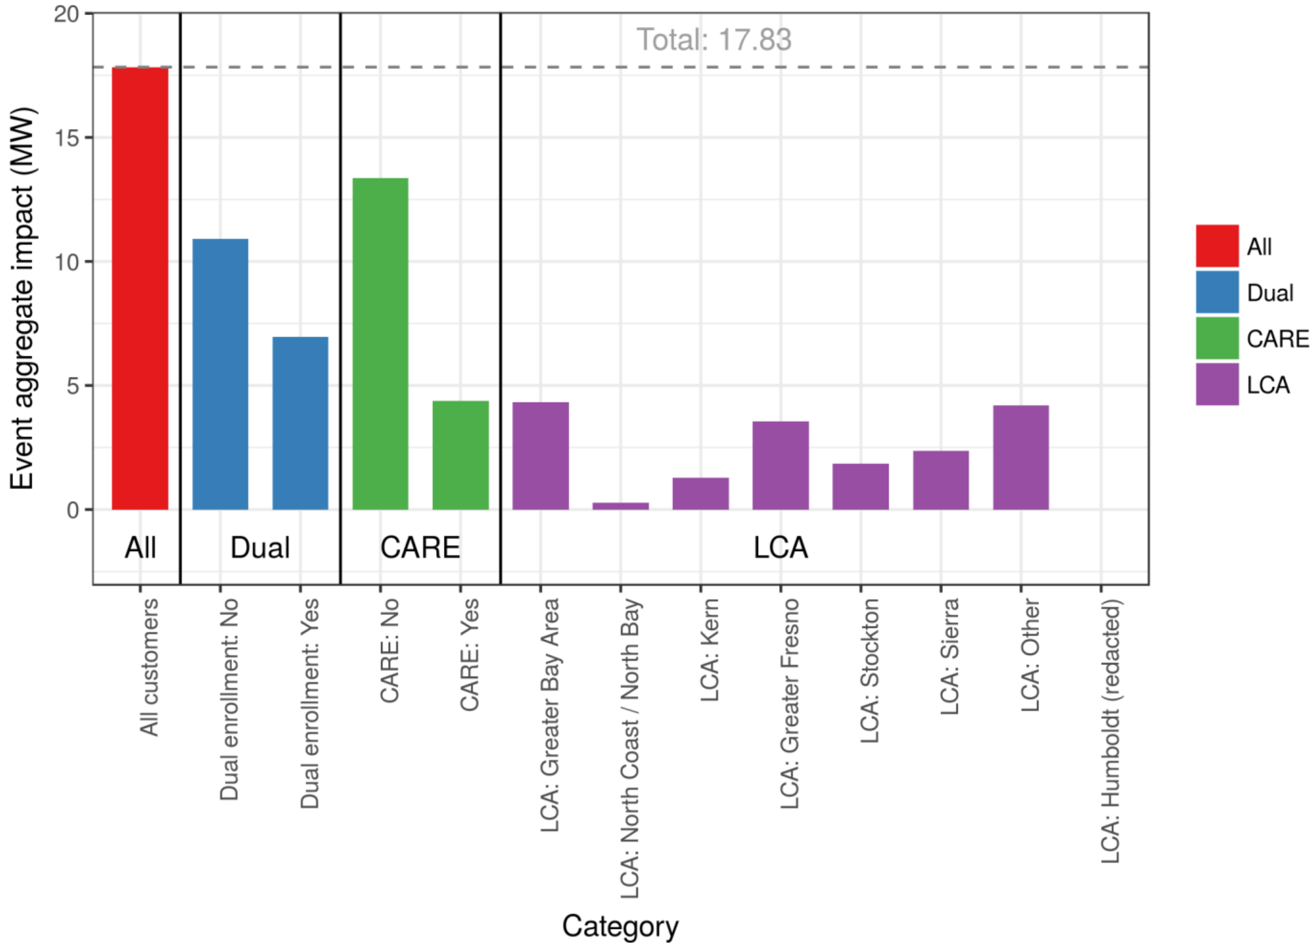

Figure note: The categories with multiple groups add to All.

From this view, it can be verified that:

-
- Dually enrolled customers contributed less than SmartRate™ only customers, in aggregate. The 34% decline in the number of dually enrolled customers leaves them lagging behind SmartRate™ only customers in aggregate impact for the first time.
- CARE customers contribute just a small fraction of aggregate impacts.
- Greater Bay Area and Fresno are the named LCAs with the largest aggregate contributions to the total. For context, the large number of customers from the Greater Bay Area makes it the named LCA with the largest aggregate contribution to the total despite cooler weather and the lowest reference loads and per-customer impacts of any LCA.

Average Ex Post Load Impacts by Event

On a typical event day, the average per customer load shed was 0.16 kW in 2018, compared to 0.23 kW in 2017. This varied slightly across the nine 2018 events (i.e., some of the black dots in Figure 11 are above, and some are below the typical event day line) but the results are generally consistent across event days, with the greatest load impacts occurring on July 25th when the average temperature during the event hours (across all LCAs) was just above 90 degrees—the hottest event during the 2018 season.

Figure 11. Mean 2018 ex post event load shed (overall, by date and LCA)

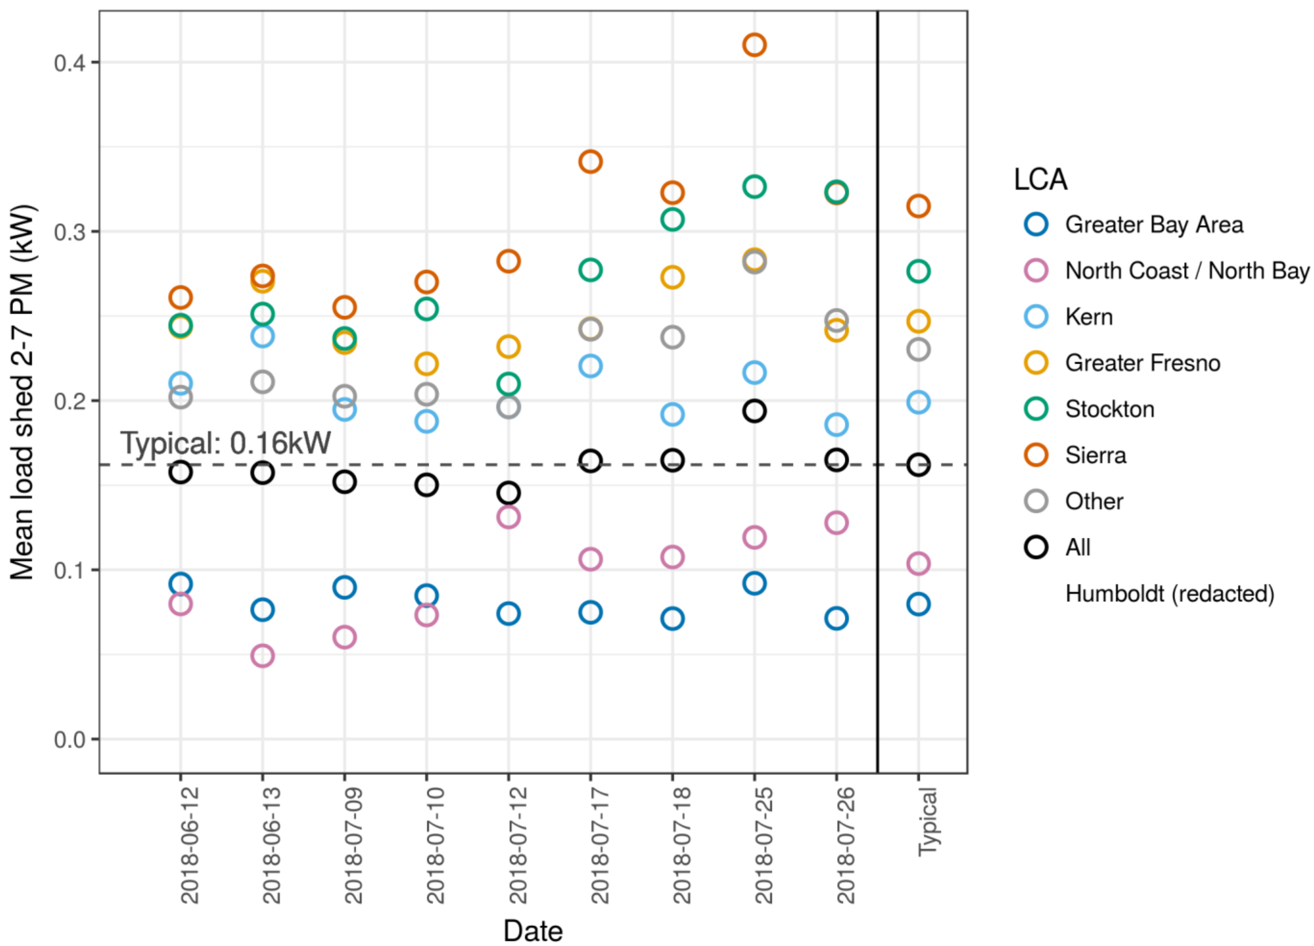

While the load shed across event days varies somewhat, the mean load shed varies much more by LCA, as represented by the vertical spread of the points on each day. There is frequently more than a 4x difference between the LCAs with the lowest and highest per-customer impacts. Results by LCA are discussed in more detail in the LCA-specific results section below.

Load Impacts by Customer Sub-group

This section presents load impacts for dually enrolled vs. SmartRate™ only, CARE customers, and LCA. Table 5 summarizes typical event day outcomes for all the relevant sub-group categories.

Table 5. Typical event day outcomes for all evaluated customer sub-groups

	Category	Enrollment	Temp. (F)	Per-cust. Ref. (kW)	Per-cust. Impact (kW)	Agg. Impact (MW)	Agg. Ref. (MW)	% Impact per cust. (% of ref.)
All	**All customers**	109,972	88	1.51	0.16	17.83	166.02	10.7%
LCA	**Humboldt**	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
	Greater Bay Area	54,183	79	0.83	0.08	4.32	44.85	9.6%
	Greater Fresno	14,357	102	2.53	0.25	3.54	36.32	9.8%
	Kern	6,432	101	2.58	0.20	1.28	16.58	7.7%
	North Coast / North Bay	2,526	88	1.17	0.10	0.26	2.94	8.9%
	Other	18,254	92	1.87	0.23	4.20	34.07	12.3%
	Sierra	7,540	98	2.15	0.32	2.37	16.24	14.6%
	Stockton	6,680	97	2.24	0.28	1.85	14.98	12.3%
Dual	**SmartRate™ only**	91,515	87	1.43	0.12	10.90	130.61	8.3%
	Dually enrolled	18,456	94	1.92	0.38	6.97	35.37	19.7%
CARE	**Non-CARE**	77,573	86	1.37	0.17	13.35	106.37	12.6%
	CARE	32,397	93	1.85	0.14	4.37	59.97	7.3%

Table note: Enrollment numbers are based on the typical event day for 2018, which averages enrollments across events. The total enrollment for each major category of customers above adds to the same total, apart from small differences due to rounding.

A review of the LCA rows reveals that they experience a wide variety of temperatures, with the Greater Bay Area notably cooler than all others, and Greater Fresno and Kern notably hotter. The per-customer reference loads track up and down with temperatures fairly well. The per-customer impacts, on the other hand, are not just fixed percentages of the reference loads. Sierra was the third hottest LCA but featured the largest per-customer impacts by a significant margin. This is reflected in their position as the top impact as a percentage of reference loads (at 14.6%). Looking at aggregate impact, the large enrollment numbers in the Greater Bay Area outweigh the lowest per-customer impacts of any LCA to produce the top entry. A combination of fairly high enrollment numbers, very hot weather, and decent savings as a % of reference loads make Greater Fresno the second largest named LCA in terms of aggregate impacts. The “Other” LCA is comprised of customers from all over the service territory that happen to not live in areas of grid stress, so it is difficult to generalize about what drives their savings, but they are notably different from All customers.

A review of the dually enrolled vs. SmartRate™-only rows reveals that more dually enrolled customers are in hotter climates than SmartRate™ only customers. The population-weighted typical event temperature for dually enrolled customers is 7 degrees hotter. That difference carries into the reference loads. The per-customer impacts for dually enrolled customers are further enhanced by the nearly 20% impact as a % of reference load response they have. This surely bears the imprint of the automated air conditioning controls deployed for dually enrolled participants. When tabulating aggregate impacts, the outsized per-customer response can’t make up for the fact that just 17% of enrolled customers are dually enrolled, so their impacts come in at about 7 MW compared to nearly 11 MW from SmartRate™ only customers. This is the first year that dually enrolled customers contributed less to aggregate impacts than SmartRate™ only customers.

A review of the CARE vs. Non-CARE rows reveals that more CARE customers are in hotter climates than non-CARE customers. The population weighted typical event temperature for CARE customers is 7 degrees hotter. The temperature difference carries into the reference loads, but CARE customers have significantly lower impacts as a percentage of reference loads. They contribute impacts of just 4.4 MW compared to 13.3 MW from non-CARE customers, despite the fact that they comprise over 29% of the enrolled population. See the *CARE Enrollment Status Results* below for more details and discussion of CARE customers.

LCA-specific Results

The 2018 ex post results indicate that the mean load shed for customers varies strongly across LCAs. Figure 12 provides a summary of ex post results by LCA, with panels from left to right that show aggregate load impacts for typical event days, per-customer event impacts for all event days, and impacts as a % of reference loads with one open circle per category per event day. The LCAs with highest temperatures (i.e., Sierra, Stockton, Fresno) yielded the largest load impacts per customer, as has been the case historically. Overall, Sierra and Stockton participants contributed the most *per customer*, and the Greater Bay Area participants contributed the least *per customer* but the most in *aggregate*. The North Coast / North Bay LCA experienced the greatest variability in impacts as a % of reference loads. This may have to do with the variable weather that comes from straddling cool coastal and hot inland climates. Notably, Kern, which is also among the hottest LCAs, underperformed in 2018. The potentially anomalous Kern results are discussed further in Appendix C.

Figure 12. Summary of aggregate, per-customer load impact and load impact as a % of reference loads for LCAs

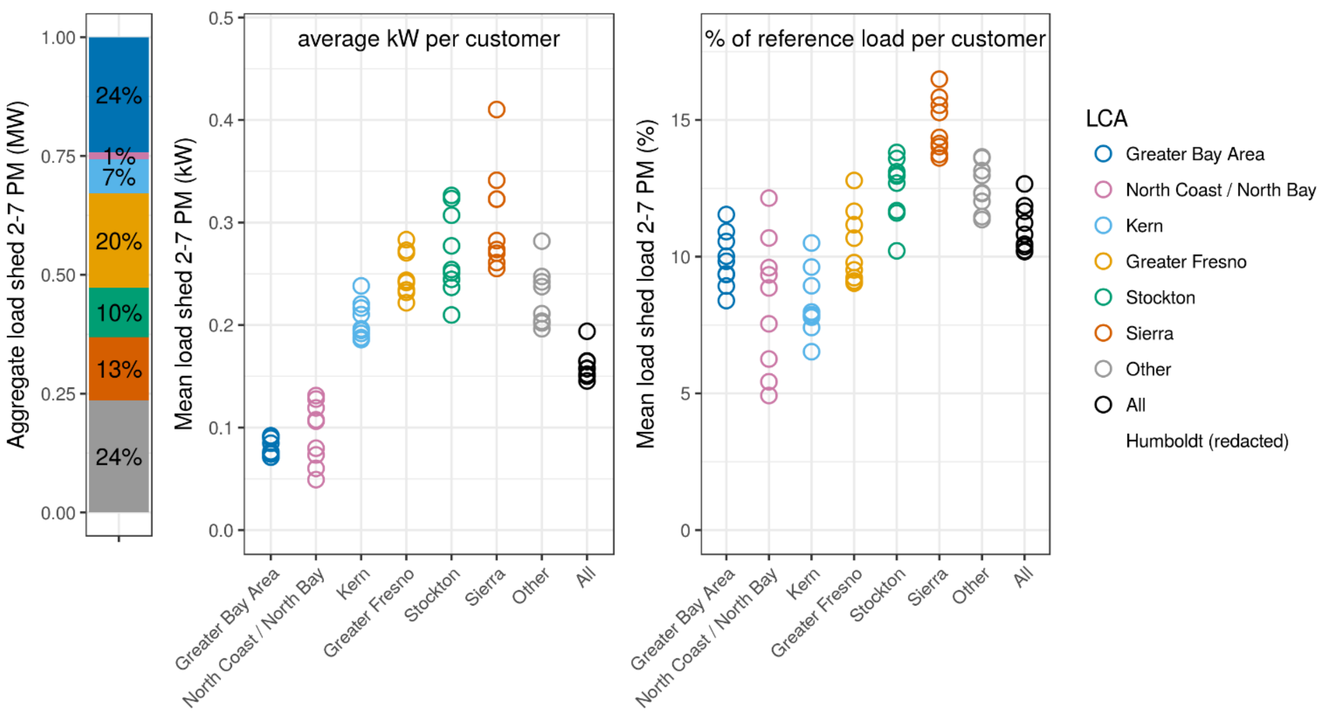

To better examine relative performance by LCA, Figure 13 magnifies the middle panel above to provide a closer look at per-customer event impacts by LCA, with each of the nine events from 2018 plotted as a colored open circle for each LCA. The 'All' LCA represents every enrolled customer in the program and shows what the high-level results look like. The dashed gray line represents the typical event day per-customer impact for the All LCA. It cleanly bisects the results into the 4 hot LCAs (+ Other) that consistently save above the average and the cooler LCAs that consistently save less.

Figure 13. Average participant load shed by LCA (each dot is an event)

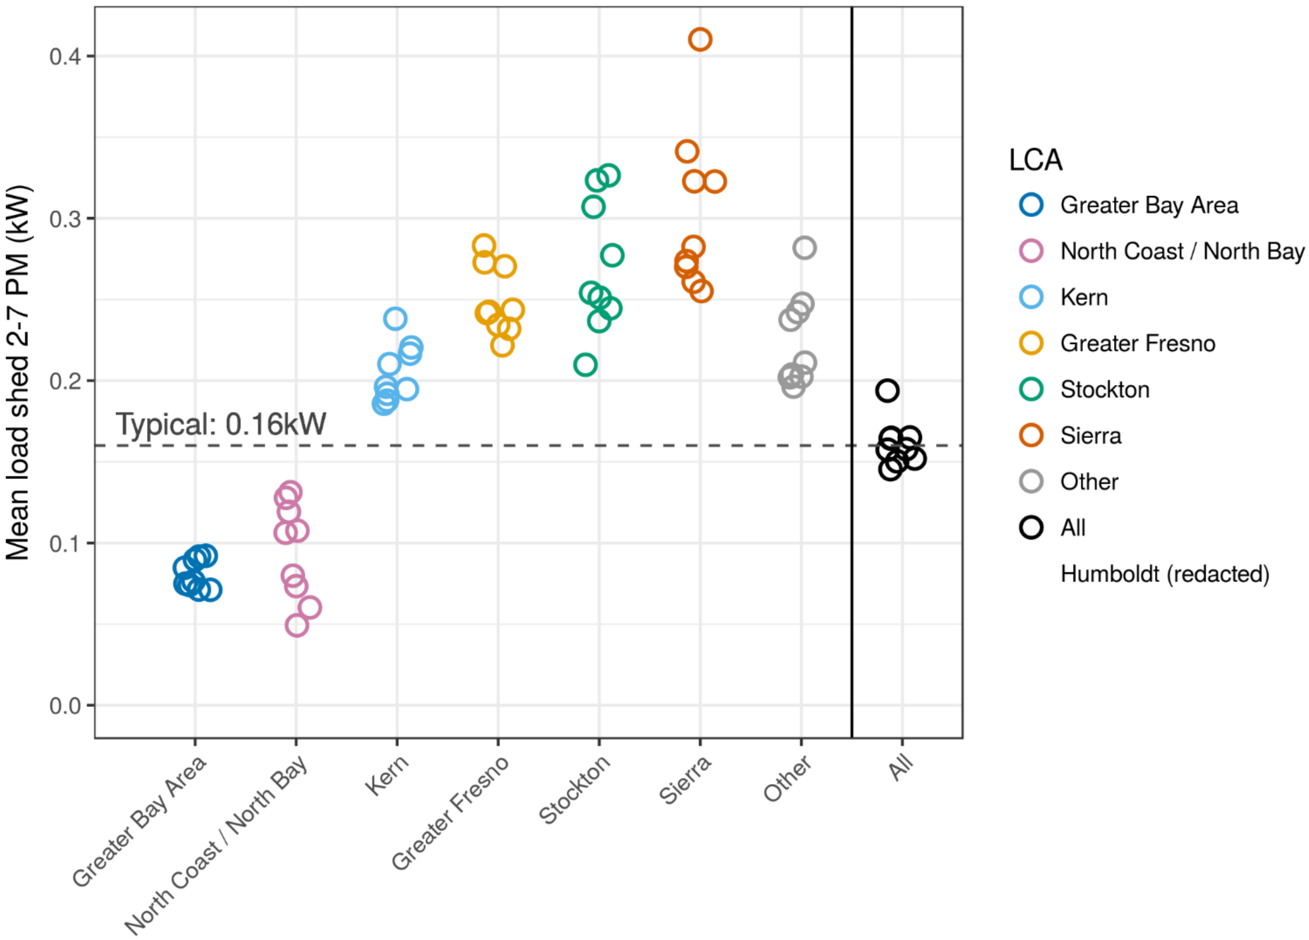

The contribution of the LCA to the total, however, shows a different pattern than the pattern above (see Figure 14). Despite the fact that the Greater Bay Area has the lowest mean load shed per customer, it contributes the most to the aggregate savings due to the large number of participants in the LCA, followed closely by the contribution of customers who are not in any LCA (i.e., “Other”) and Greater Fresno.

Figure 14: Contribution of LCA to Aggregate Load Shed

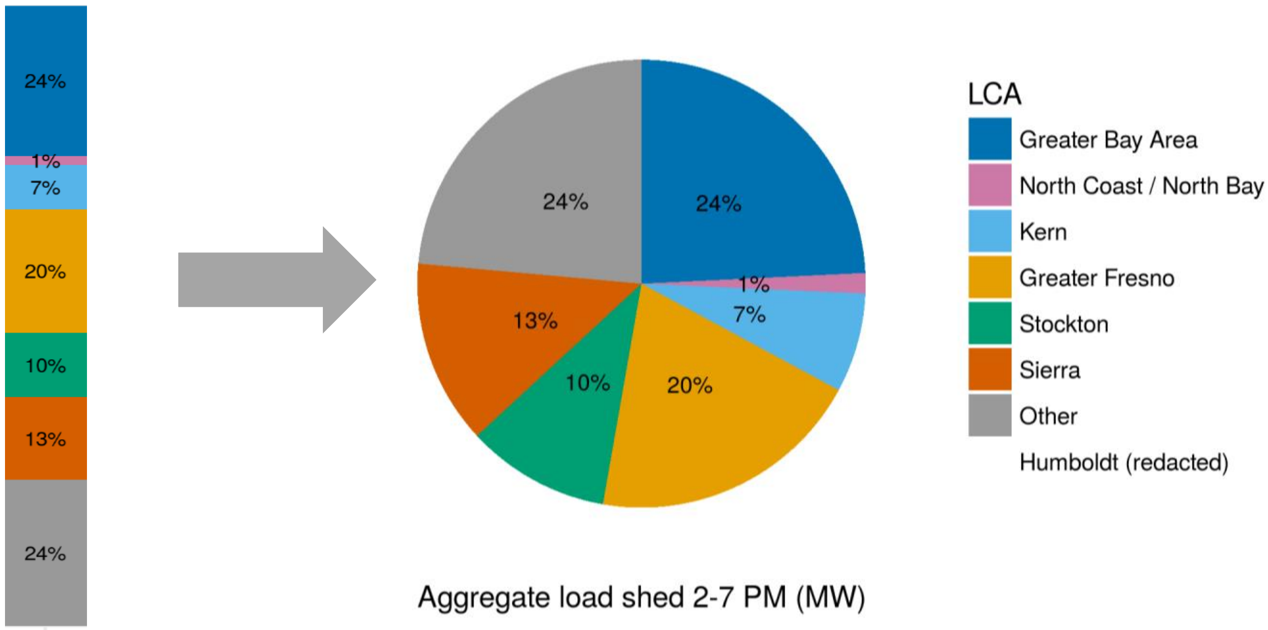

When we look at the contribution of SmartRate™ only and dually enrolled across the LCAs, Greater Bay Area SmartRate™ only customers can be seen contributing the most to the overall aggregate load shed in 2018 (despite the fact that they have lowest per-customer savings) due to the large number of participants. See Figure 15, which depicts the range of outcomes across the nine event days as a box with whiskers, where the box spans the 25th to 75th percentile of the results. There is no LCA where dually enrolled customers consistently contribute aggregate savings greater than SmartRate™ only customers.

Figure 15. Contribution to the Aggregate Load Shed by LCA and Enrollment Status

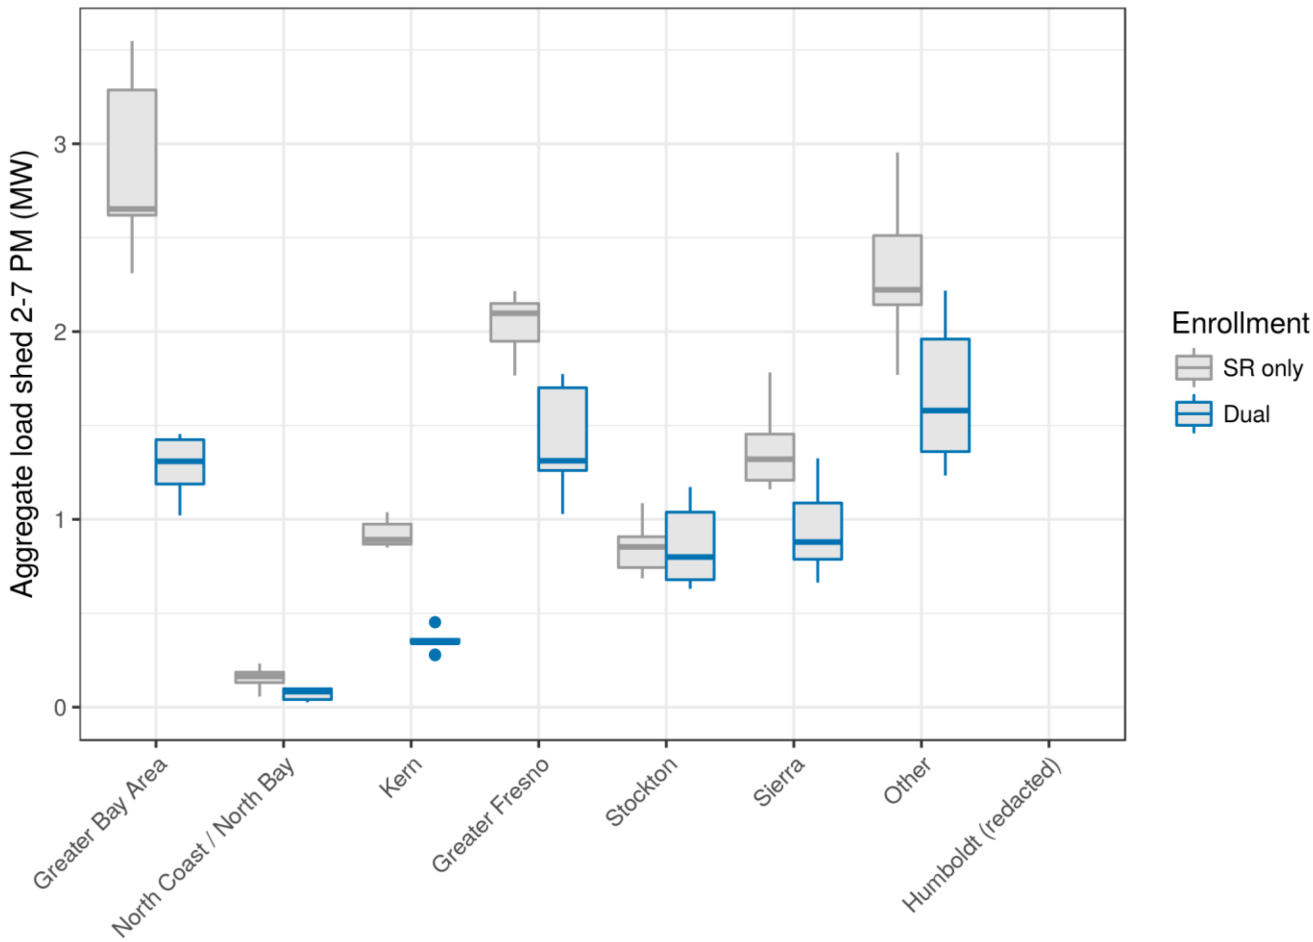

Dual Enrollment Status Results

Figure 16 provides a summary of ex post results by dually enrolled vs. SmartRate™ only enrollment, with panels from left to right that show aggregate load impacts for typical event days, per-customer event impacts for all event days, and impacts as a % of reference loads with one open circle per category per event day.

Figure 16. Event load impact aggregate MW, per-customer kW and % of reference load by dual enrollment status

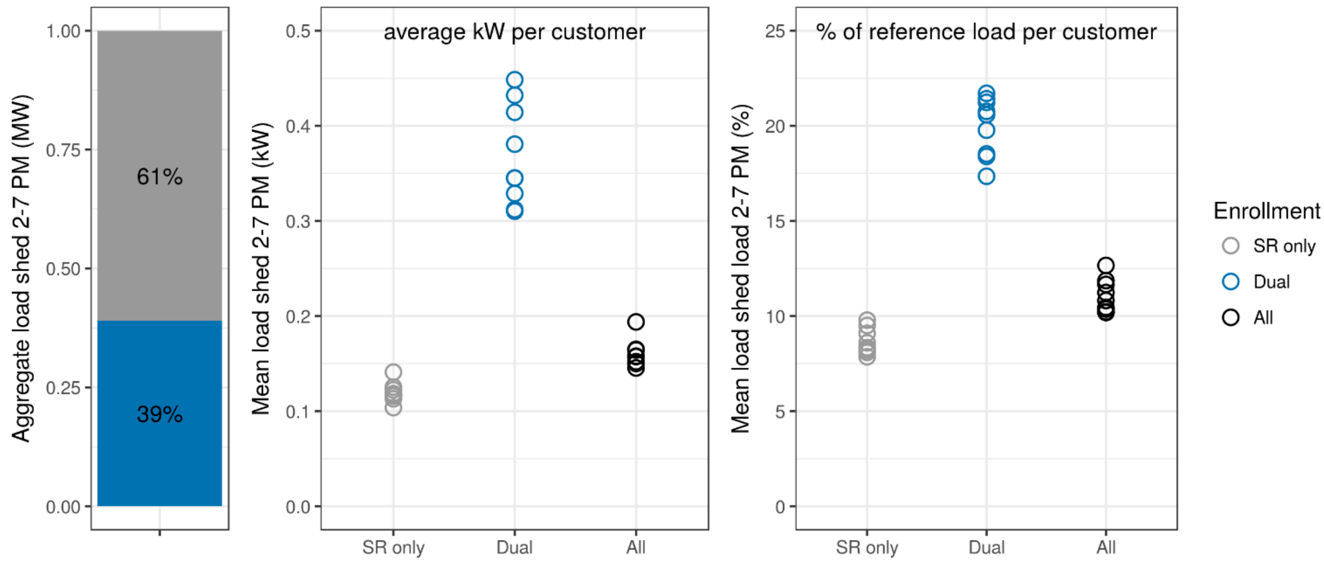

Dually enrolled participants typically save more than 3x more than SmartRate™ only participants. Dually enrolled participants also save a much larger percentage of their load – averaging and often exceeding 20% (compared to an average of less than 9% by the SmartRate™ only customers), which demonstrates that direct load controls have the ability to save much more than customer behaviors alone. These results are consistent with past results.

However, this is the first year that the aggregate impacts are less than 50% from dually enrolled customers. Across the season, dually enrolled customers contribute about 39% of the total aggregate load. This is primarily because the makeup of the customers enrolled has shifted. The decrease in dually enrolled participants was particularly high (36% decrease compared to under 5% decrease of Smart-Rate™ only customers). As a result, dually enrolled customers represent just under 17% of all participants in 2018 (compared to 23% in 2017).

One issue with comparing dually enrolled and SmartRate™ only customers is that dually enrolled customers tend to be in hotter places. If we want to know how dually enrolled and SmartRate™ only customers compare *under the same conditions*, we need to control for reference loads (which are driven by temperatures). Figure 17 presents a scatter plot of per-customer load impacts (y-axis) vs. reference loads (x-axis) using 9 events x 8 LCAs, separated by enrollment type (SR only and Dual). There are guide lines for 1.5 and 2 kW reference loads.

Figure 17. Event load shed vs reference load – dually enrolled and SmartRate™ only (highlights at 1.5 kW and 2 kW reference loads)

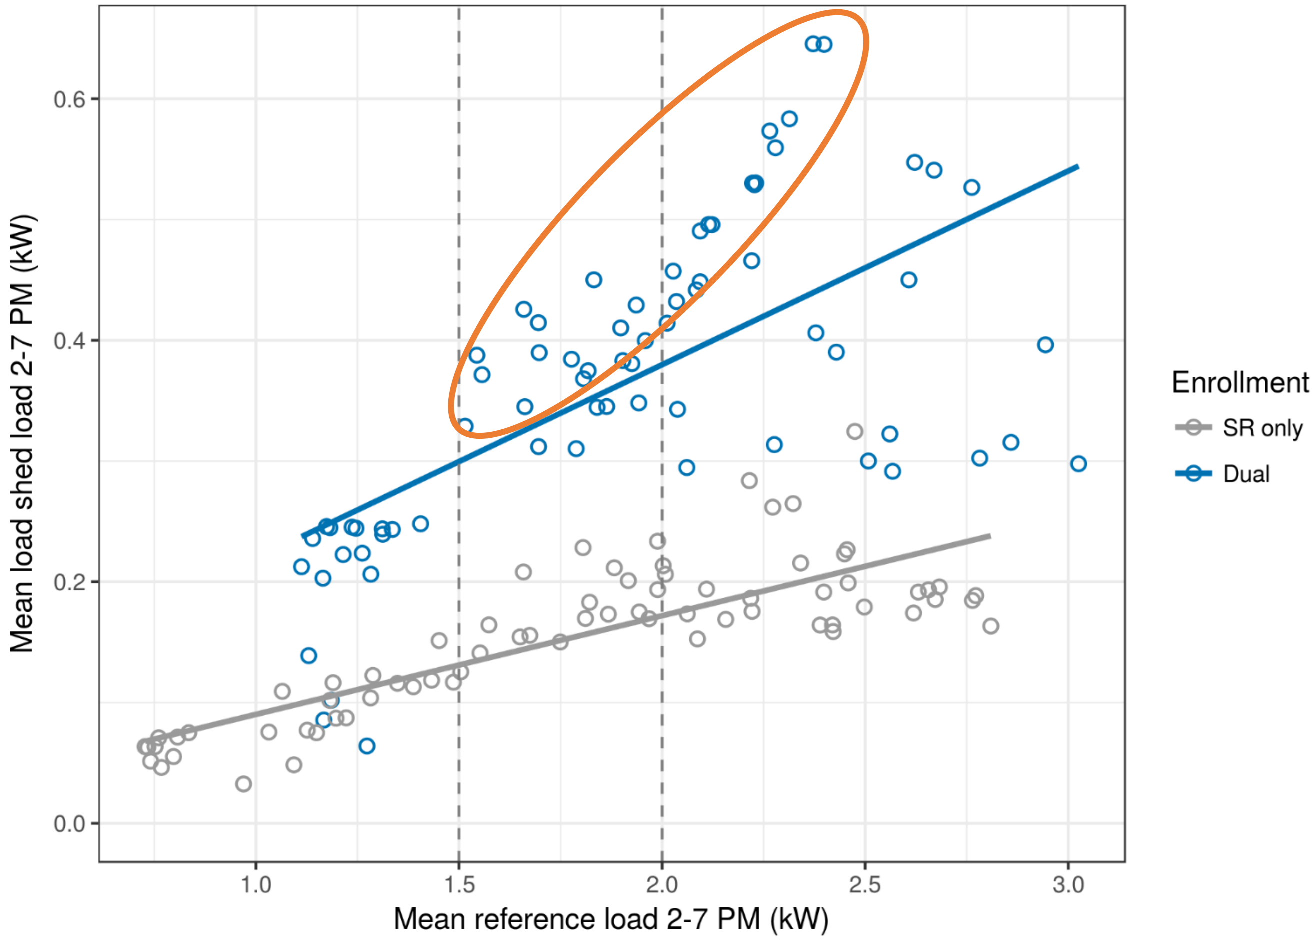

From this plot, it can be verified that there is a systematic difference between reference loads of the two groups (SR only extends lower and Dual extends higher). At the same reference loads, dually enrolled customers have impacts roughly double their SmartRate™ only peers. There is also cohort of 'Dual' points that outperform the linear fit highlighted in an oval. They start above the line near the 1.5 kW reference load line and diverge more and more above the line as reference loads increase. Those points come from Sierra, Stockton, and Other. The lowest impact 'Dual' points at the hottest temperatures are from Kern.

CARE Enrollment Status Results

Figure 18 provides a summary of ex post results by CARE enrollment status, with panels from left to right that show aggregate load impacts for typical event days, per-customer event impacts for all event days, and impacts as a % of reference loads with one open circle per category per event day.

Figure 18. Event load impact aggregate MW, per-customer kW and % of reference load by CARE status

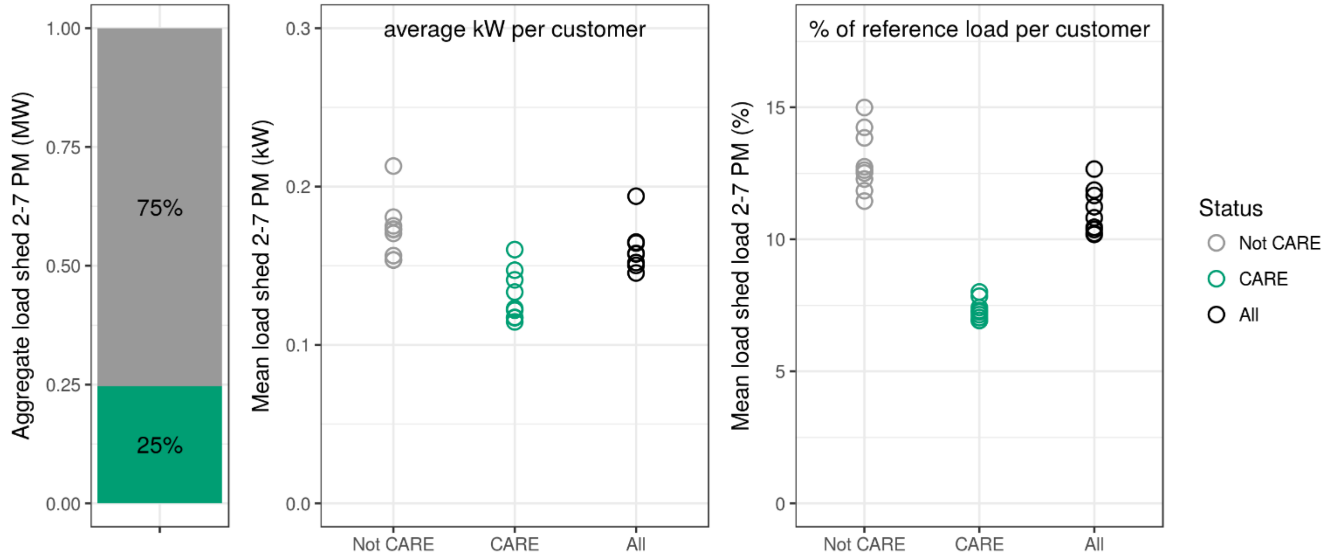

Consistent with past findings, across the events, CARE customers tend to save less than individuals who are not on the CARE program, despite higher reference loads, due to significantly lower load impacts as a % of reference loads. For comparison: they make up a smaller portion of the aggregate impacts than dually enrolled customers (a little over 60% of what dually enrolled customers contribute) despite out numbering them nearly 2 to 1. CARE customers experience similar weather and only slightly lower reference loads compared to dually enrolled customers. The primary explanation for the smaller impacts for CARE customers is that their load shed is the smallest percentage of reference loads of any customer group examined.

Why CARE customers would shed less is an open question. On the one hand, the effect flies in the face of economic incentives. They are known to be low income and have gone out of their way to apply for lower rates, so we might expect them to be more responsive to the price signals of SmartRate™. On the other hand, the effect could be structural. The portion of CARE customers in apartments, which is higher than the general population, may have less direct control over their energy consuming devices. Finally, the effect could also be behavioral. Lower shed is consistent with longer and higher occupancy in homes, which is correlated with lower incomes.

Because typical event temperatures are higher for CARE customers than non-CARE, it is important to understand how their reference loads relate to those of other customers. Figure 19 presents a scatter plot of per-customer load impacts (y-axis) vs. reference loads (x-axis) using 9 events x 8 LCAs worth of open circle points, separated by CARE status. There are reference lines at 1.5 kW and 2 kW reference loads.

Figure 19. Event load impact vs. reference load across all events in all LCAs by CARE enrollment status

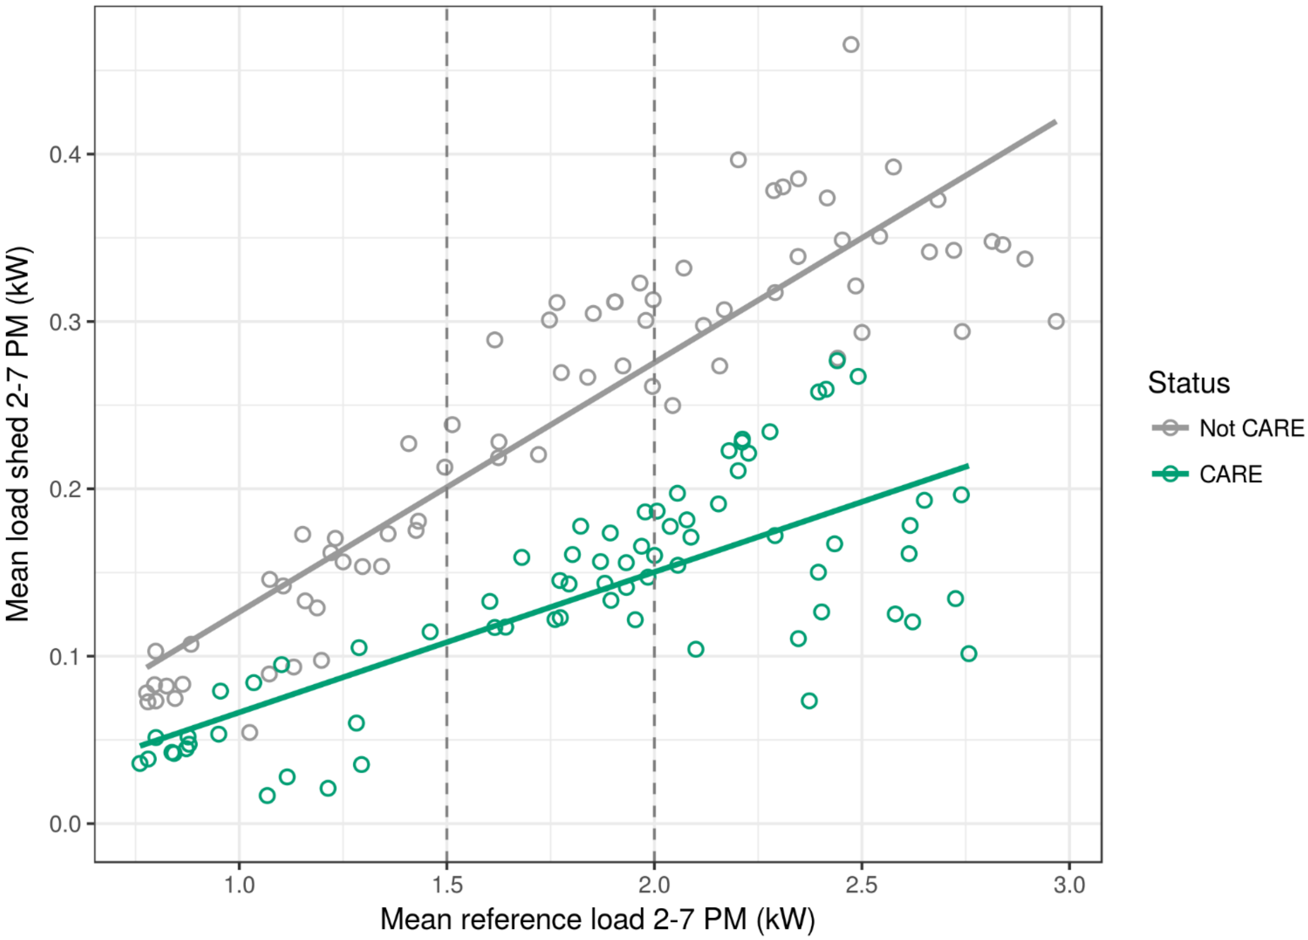

From this plot, it can be verified that there is a systematic difference between reference loads of the two groups. In this case, CARE customers have a lower maximal extent of reference loads, but they have a higher density of points at elevated reference loads and their average is, in fact higher, than non-CARE customers. This somewhat confusing state of affairs may be attributable to CARE customers living in smaller homes with lower loads than other customers in the same areas, but more concentrated in the hotter climates. CARE customers respond with impacts relative to the same reference loads that are about one half of non-CARE customers.

Interestingly, the dually enrolled customers on CARE actually perform as well or perhaps even a little better on average than dually enrolled non-CARE customers since the utility is controlling their loads during this time (see Figure 20). It is the SmartRate™ only CARE customers (shown in blue below) that appear to be the lowest performing segment. CARE customers, therefore, only tend to underperform without direct load controls.

Figure 20. Event load shed by CARE status – dually enrolled vs SmartRate™ only

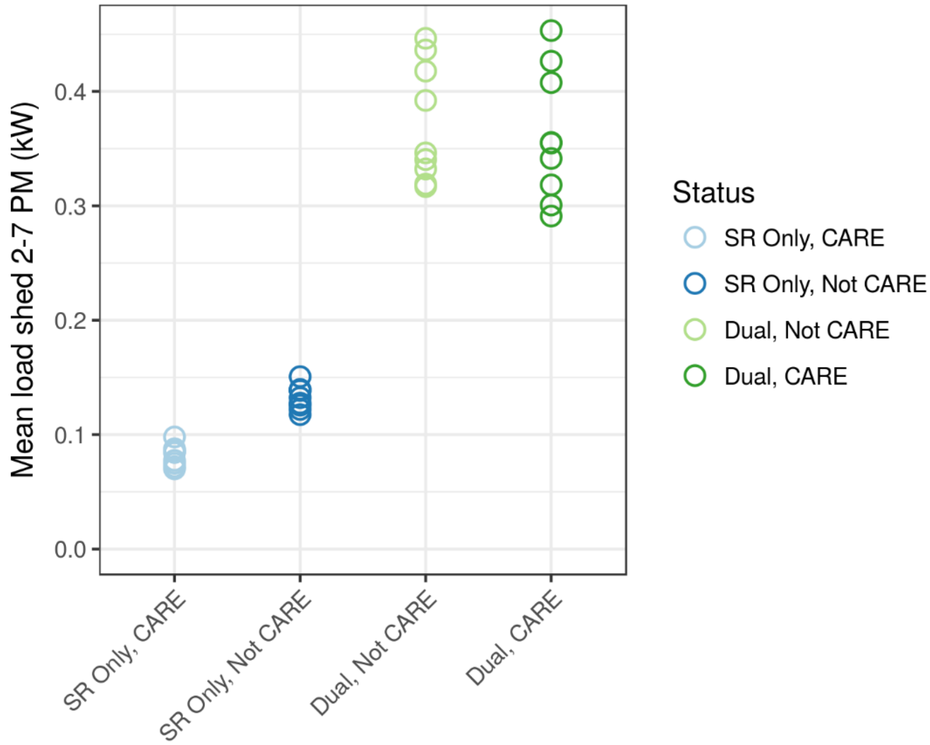

Current Ex Post to Prior Ex Post

We compared CDA's ex post results to the prior year's ex post results for their respective typical event days. (See Table 6.) Overall, our 2018 ex post estimates are lower because of cooler weather, decreasing numbers of participants, and the shifting make-up or composition of participants.

Table 6. Comparison of 2018 ex post load impacts on a typical event day compared to 2017 ex post impacts on a typical event day for hours 2pm to 7 pm

	Estimate	Enrolled	Agg. load impact (MW)	Agg. Ref. load (MW)	Per-cust. load impact (kW)	Per-cust. Ref. load (kW)	% Load Impact	Avg. Event Temp (F)
All	2017 Ex Post	124,049	28.1	226.0	0.23	1.82	12.4%	98
	2018 Ex Post	109,972	17.8	166.0	0.16	1.51	10.7%	88
Dually Enrolled	2017 Ex Post	28,923	14.3	69.5	0.50	2.40	20.6%	100
	2018 Ex Post	18,456	7.0	35.4	0.38	1.92	19.7%	94
SR only	2017 Ex Post	95,126	13.8	156.5	0.14	1.64	8.8%	97
	2018 Ex Post	91,515	10.9	130.6	0.12	1.43	8.3%	87

Table note: Humboldt numbers have been redacted.

As described earlier in this report, cooler weather led to a reduction in air-conditioning demand as quantified via cooling degree hours during the 2018 season, which reduced the mean load shed during 2018 events.

Figure 21 displays the average event impacts (y-axis) for all events in 2017 (gray) and 2018 (purple) against the average outside temperature for each event. The figure illustrates that impacts are sensitive to temperature, that the temperatures in 2018 were significantly cooler than 2017, and that after accounting for temperature with all else being equal, our 2018 per-customer impact estimates are roughly in line with the results from 2017.

Figure 21. Mean load shed by temperature (2017 vs 2018)

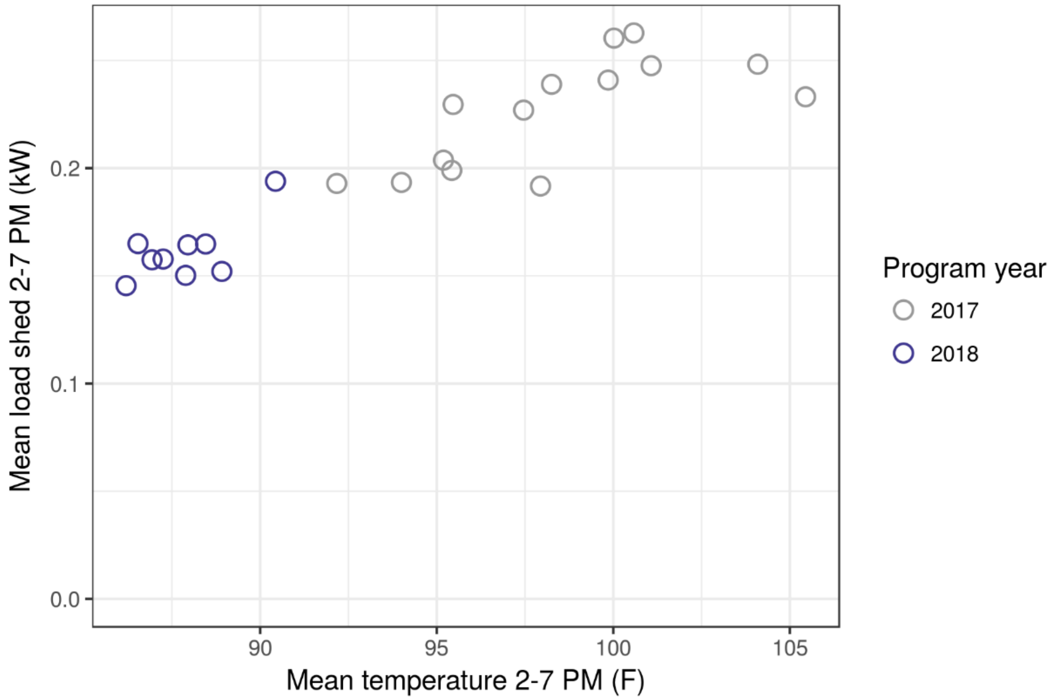

Along with temperatures, LCA and dual enrollment status are the most important determinants of load impacts, and they all interact. Figure 22 extends the use of scatter plots to present load impacts vs. temperatures for all events in 2017 and 2018, broken out by dual enrollment status. The green/blue points are from dually enrolled customers and the red/orange points are from SmartRate™ only customers.

Figure 22: Event load shed vs. outside temperature by LCA (2017 vs. 2018); Humboldt redacted

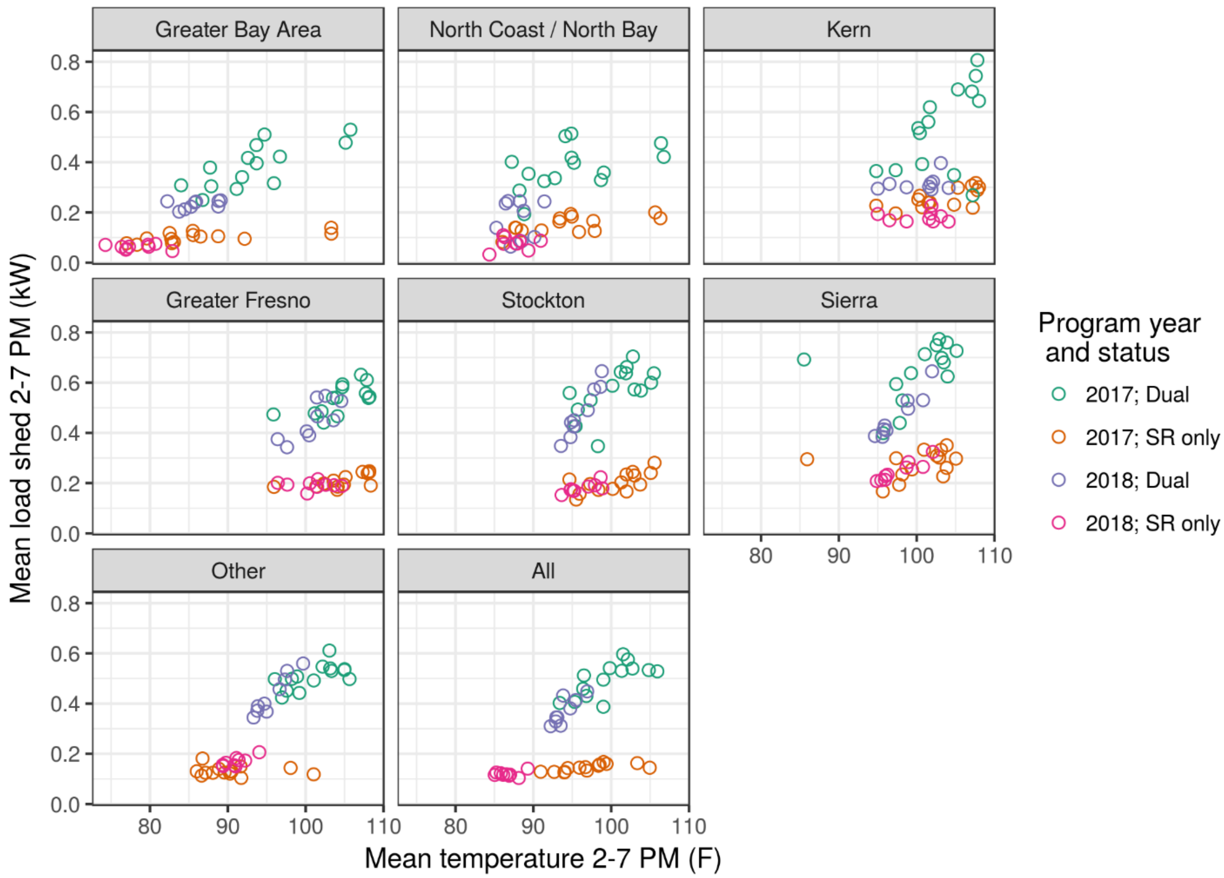

From the figure it can be seen that the 2018 ex post results are generally consistent with (i.e. grouped with and extending the trend of) the 2017 results after taking into account LCA, dual enrollment status, and event temperatures. The figure also reveals more subtle patterns, including the tendency of duals in some LCAs to be in the hotter parts of the LCA (Greater Bay Area and Other in particular); that the temperature response of customer impacts is different by dual enrollment status and LCA; and that the normally clean separation of dually enrolled and SmartRate™ only customers broke down in Kern in 2018. We explore the drivers of event impacts in the Ex Ante Results section and the Kern numbers (shown in the figure above) further in Appendix C.

Table 7 quantifies the changes in key contributing factors and the associated changes in per-customer and aggregate load impacts.

Table 7. Year over year changes from 2017 to 2018 leading to reduced impacts

	Enrollment	Event avg. Temp (F)	CDH	Reference load (kW)	Impact (kW)	Impact (%)	Aggregate impact (MW)
2017	124,049	95	390	1.82	0.23	13%	28.10
2018	109,972	88	267	1.51	0.16	10%	17.83
Difference	-14,077	-7	-123	-0.31	-0.06	-2%	-10.26
% Difference	-11%	-7%	-32%	-17%	-28%	-	-37%

Table note: Humboldt numbers are not included.

Some of these changes, like the weather, will change back again in the future. Most of the 17% drop in reference loads can be explained by the decrease in event temperatures and CDH (i.e. less air conditioning). All else being equal, reference loads in future years could easily match those of 2017. The effect of de-enrollment, on the other hand, represents a permanent shift in the composition of the SmartRate™ population.

Because CCAs serve specific geographies and have been primarily in coastal or near coastal communities, de-enrollments affected some segments of the populations more than others. For example, customers in the Greater Bay Area LCA saw much greater losses in the number of participants than their counterparts. There was a >50% drop in Bay Area aggregate impact on high level of population churn (but not much overall loss of head count). Sierra experienced a large (43%) decrease in participants. See Table 8.

Table 8. 2017 to 2018 change in aggregate impacts and end of summer enrollment by LCAs

		Aggregate impact (MW)							
	Program year	**Bay Area**	**N. Coast/ N. Bay**	**Kern**	**Fresno**	**Stockton**	**Sierra**	**Other**	**All**
Aggregate impact	2018	4.30	0.30	1.30	3.50	1.80	2.40	4.20	17.8
	2017	8.70	0.40	1.40	4.30	2.20	5.40	5.70	28.1
	Difference	-4.40	-0.10	-0.10	-0.80	-0.40	-3.00	-1.50	-10.3
	% Difference	-51%	-25%	-7%	-19%	-18%	-56%	-26%	-37%
Enrollment on typical event day	2018	54,183	2,526	6,432	14,357	6,680	7,540	18,254	109,972
	2017	57,309	2,320	4,290	13,816	6,693	13,146	26,476	124,049
	Difference	-3,126	206	2,142	541	-13	-5,606	-8,222	-14,077
	% Difference	-5%	9%	50%	4%	0%	-43%	-31%	-11%

Table note: Humboldt numbers are not included.

These population changes within specific LCAs can have very different effects on the overall total aggregate impact since the mean load shed of some LCAs is much higher than others (see Figure 21). Sierra mean load sheds, for example, are consistently higher than Greater Bay Area.

Dually enrolled customers also saw a disproportionate drop. As shown in Table 9, the number of dually enrolled customers *on typical event days* dropped by 36%. However, the aggregate impacts of the group of dually enrolled customers dropped by 51%, with very roughly 36% is from the population change and therefore 15% from cooler weather.

Table 9. 2017 to 2018 change in aggregate impacts and enrollment by Dual enrollment status

	Aggregate impact (MW)			Enrollment		
Program year	**SR Only**	**Dual**	**Total**	**SR Only**	**Dual**	**Total**
2018	10.9	7.0	17.8	91,515	18,456	109,972
2017	13.8	14.3	28.1	95,126	28,923	124,049
Difference	-2.9	-7.3	-10.3	-3,611	-10,467	-14,077
% Difference	-21%	-51%	-37%	-4%	-36%	-11%

CARE enrollments and aggregate impacts *on a typical event day* dropped less than the population average (see Table 10). This could be due to the tendency for CARE customers to live inland way from CCA defections and in hotter climates.

Table 10. 2017 to 2018 change in aggregate impacts and enrollment by CARE enrollment status

	Aggregate impact (MW)			Enrollment		
Program year	**CARE**	**Not CARE**	**Total**	**CARE**	**Not CARE**	**Total**
2018	4.4	13.4	17.8	32,397	77,573	109,972
2017	5.5	22.9	28.1	35,310	88,782	124,049
Difference	-1.1	-9.5	-10.3	-2,913	-11,209	-14,077
% Difference	-20%	-41%	-37%	-8%	-13%	-11%

Factor		Changes between 2017 and 2018	Individual average results expected differences (Mean Load Shed)		Total aggregate result expected difference (Aggregate Load Shed)	
Total participation		Overall 11% decrease in typical event day participation-- from 124K in 2017 to approximately 110K in 2018	No expected change in per customer results due to this factor alone	●	Overall decrease in aggregate results	↓
Shifting composition of participants	Dually enrolled	Dually enrolled customers, who tend to be high savers, make up a smaller percentage of the total customers (17% in 2018 compared to 23% in 2017).	Decrease in average per customer results	↓	Same as individual	●
	CARE	In 2018, a slightly larger percentage of participants were CARE participants 30% v 28% in 2017. These CARE participants tend to save less than non-CARE participants.	Small decrease in average per customer results	↓	Same as individual	●
	Overall make-up by LCA	Greater Bay Area participants, who tend to have the lowest per customer savings, made up a larger percentage of the overall total participants (48% in 2018 compared to 46% in 2017)	Slight decrease in average per customer results	↓	Same as individual	●
	TOU	TOU customers make up just under 20% of customers in 2018. The % of TOU customers in 2017 was undocumented. See section on TOU.	The price effects of TOU rates could decrease what we expect at the individual level, but less clear since not reported in 2017	↓?	Same as individual	●
Temperature and CDH		Average event temperatures were roughly 7°F lower in 2018 than in 2017. As a result, there were significantly fewer CDHs in 2018. Impacts have been shown to be lower on cooler days.	Will decrease reference loads and load impacts at the individual level	↓	Overall decrease in aggregate results	↓

Current Ex Post to Prior Ex Ante

Referring to Table 11, one can see that the current ex post load impacts are lower than previously forecasted. The overall participant population was greater (109,972 on the typical event day compared to 89,829 in last year's ex ante), but the temperatures, reference loads, and differences in the composition of participants, e.g., dually enrolled customers who have higher impacts make up a smaller percentage of the total population than in prior years, resulted in lower savings than predicted. For comparison purposes, we looked at the differences during event periods (2 pm -7 pm), rather than using the 2017 RA window which is no longer relevant. The new 4-9 p.m. RA window was not evaluated in the report on the 2017 program year.

Table 11. Comparison of ex post load impacts for the typical event day in 2018 compared to prior ex ante forecasts (August peak day, 1-in-2 PG&E weather scenario for SmartRate™) during the 2 pm to 7 pm event window

	Estimate	Enrolled	Agg. load impact (MW)	Agg. Ref. load (MW)	Per-cust. load impact (kW)	Per-cust. Ref. load (kW)	% Load Impact	Avg. Event Temp (F)
All	2017 Ex Ante	89,829	19.9	166.2	0.22	1.85	11.9%	93
	2018 Ex Post	109,972	17.8	166.0	0.16	1.51	10.7%	88
Dually Enrolled	2017 Ex Ante	20,257	9.7	48.1	0.48	2.38	20.2%	99
	2018 Ex Post	18,456	7.0	35.4	0.38	1.92	19.7%	94
SR only	2017 Ex Ante	69,572	10.1	118.1	0.15	1.70	8.6%	91
	2018 Ex Post	91,515	10.9	130.6	0.12	1.43	8.3%	87

Table note: The 2017 Ex Ante forecasts shown here are specific to PY2018. The table compares the forecast for 2018 to the actual numbers from 2018.

Ex-Ante Results

This section presents the results of CDA's ex ante forecast for the period 2019-2029. For context, we first present key information on policy and enrollment changes, followed by the aggregate and mean ex ante results. We also compare the current forecast to 2018 ex post results and findings from prior years.

A companion ex ante "table generator" spreadsheet, which is a supplement to this report, provides hourly load shed predictions, with uncertainties, for all of the standardized weather conditions, using projected enrollment rates for the next eleven years. That granularity of data is not reproduced in this report.

Ex Ante Background

PG&E has historically evaluated the ex ante effects as the impact of events over the "resource adequacy (RA) window." In past evaluations, the RA window was 1 pm - 6 pm from April to October and historically all SmartRate™ events occurred from 2pm - 7 pm. However, in 2019 the RA window shifted to 4 pm - 9 pm to align with the timing of the California Independent System Operator Corporation's (CAISO's) annual availability assessment hours (AAH). The CPUC's Decision 18-06-030 made several changes to the RA program rules and implementation, including:

- The resource adequacy measurement hours were modified to HE17-HE21 (4 pm – 9 pm) for each month of the year beginning in 2019.
- The CAISO annual availability assessment hour analysis was to be submitted into the resource adequacy proceeding for consideration as to whether the Commission should adjust its resource adequacy measurement hours.

The new resource adequacy window has a large effect on the evaluation of the program, since (1) only three hours of the five-hour-long events fall within the RA window; and (2) two hours of the 'snapback' period are also included in the window. (See Figure 23 below.)

The ex ante model for the event hours is described at length in the methods section of this report (see *Ex Ante Impact Analysis Methods*) but one of the key features is that the model is fit to the observed outcomes of various customer segments, including: 1) Dually enrolled CARE, 2) Dually enrolled non-CARE, 3) SmartRate™ only CARE, and 4) SmartRate™ only non-CARE.[10] The model is used to predict the mean load shed for customers in each of these segments in the ex ante weather conditions. The mean load for *all* customers is the enrollment-count-weighted mean of the mean of each segment. An example of the hourly predictions for each segment, for the weather conditions corresponding to the PG&E 1-in-2 day in August, is shown in Figure 23 below. We discuss the effect of this three-hour overlap between the RA window and the event window throughout this chapter.

[10] "SmartRate™ only" and "Dually enrolled" refer to whether customers are simultaneously enrolled in the SmartAC™ program.

Figure 23. Predicted load shed by customers on the August peak (4-9 pm window, PG&E 1-in-2)

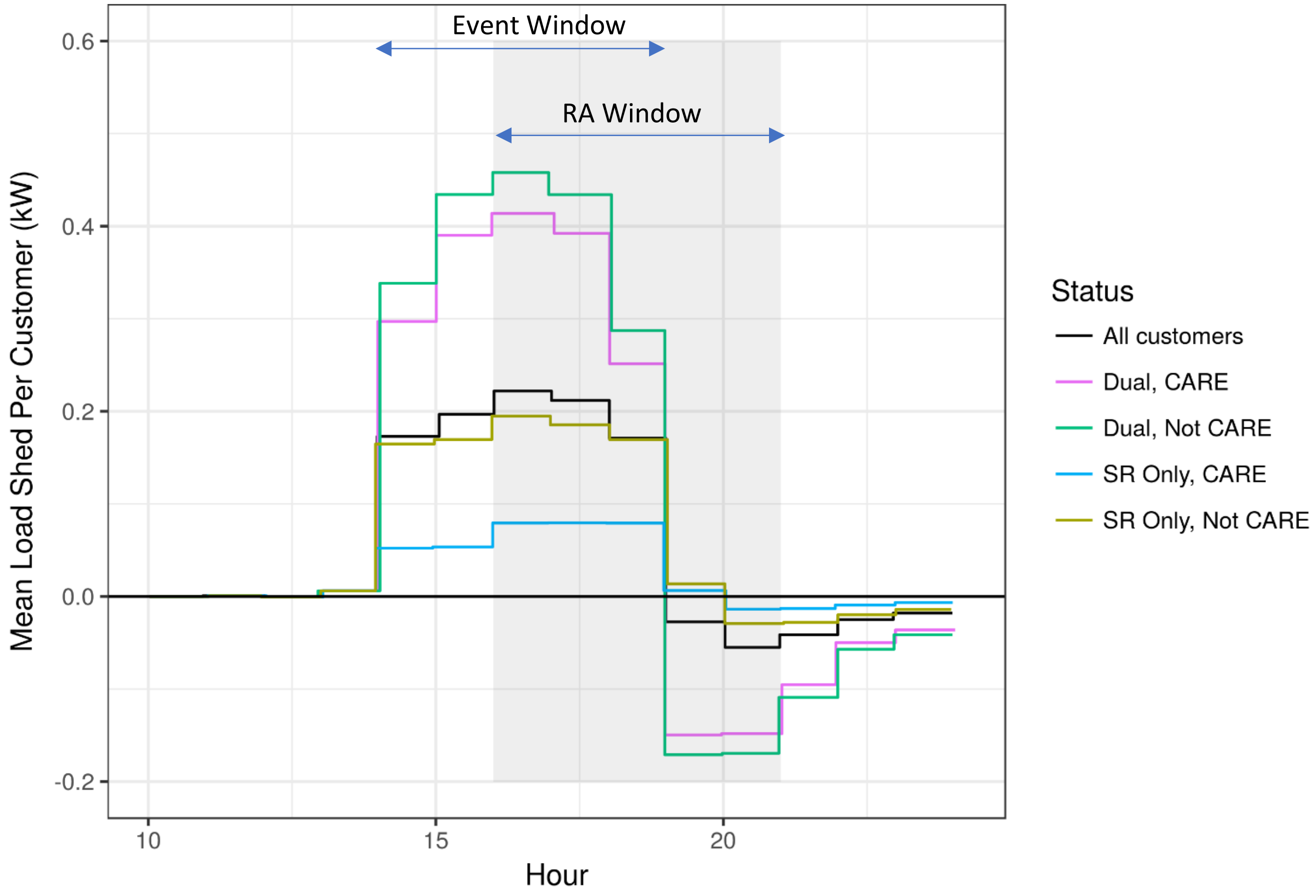

Figure note: The figure is for events that occurred 2-7 pm. The resource adequacy window is 4-9 pm, shown in gray.

Ex Ante Enrollment

Predicting the impact of the SmartRate™ program requires predictions about the number of customers in the program, which are then multiplied by the estimated impact per customer.

The enrollment forecast for 2019 and beyond increased from prior years. Because of the changes in participation (discussed in the ex post results section) the composition or makeup of future participants is assumed to be different than in prior forecasts. That is, the proportion by enrollment status, CARE status, and LCA has shifted.

Over the 11-year forecast, PG&E assumes that participation levels change during 2019 to 67,206 participants by the end of the year, approximately 15,000 (22%) of which are also in the SmartAC™ program (these customers are referred to as dually enrolled). The distribution across the LCAs and proportion of SmartRate™ only to dually enrolled customers were provided by PG&E enrollment forecasts and are assumed to be constant starting in 2020 (see table and figures below).

Table 12. Forecasted SmartRate™ enrollments 2019-2029 (provided by PG&E)

	SmartRate™ only		Dually enrolled		Total	
LCA	**2019**	**2020-2029**	**2019**	**2020-2029**	**2019**	**2020-2029**
All	63,548	52,380	14,826	14,826	78,374	67,206
Greater Bay Area	19,652	8,543	2,306	2,306	21,958	10,849
North Coast / North Bay	2,147	2,147	329	329	2,476	2,476
Kern	4,712	4,712	1,040	1,040	5,752	5,752
Humboldt	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
Greater Fresno	10,945	10,945	3,300	3,300	14,245	14,245
Stockton	4,946	4,946	1,823	1,823	6,769	6,769
Sierra	5,604	5,604	2,018	2,018	7,622	7,622
Other	15,495	15,436	4,010	4,010	19,505	19,446

Table note: 2019 is for March 2018 through March 2019. By summer of 2019 the enrollment numbers stop changing, so 2019 looks just like 2020 etc.

Humboldt is not shown in the table above because there are too few participants from Humboldt to disclose their outcomes publicly. They do not have a significant effect on average or aggregate results, so they are not shown in any of the ex ante tables and figures. **Humboldt numbers (in their entirety) have been excluded from this section.**

PG&E did not provide estimates of CARE to non-CARE customers so in the model we assumed that within each load capacity area and category of SmartRate™-only or dually-enrolled, the proportion of CARE customers in the future will be the same as in the current population.

Figure 24. Projected enrollment by (a) LCA 2019-2029 and (b) customer segment

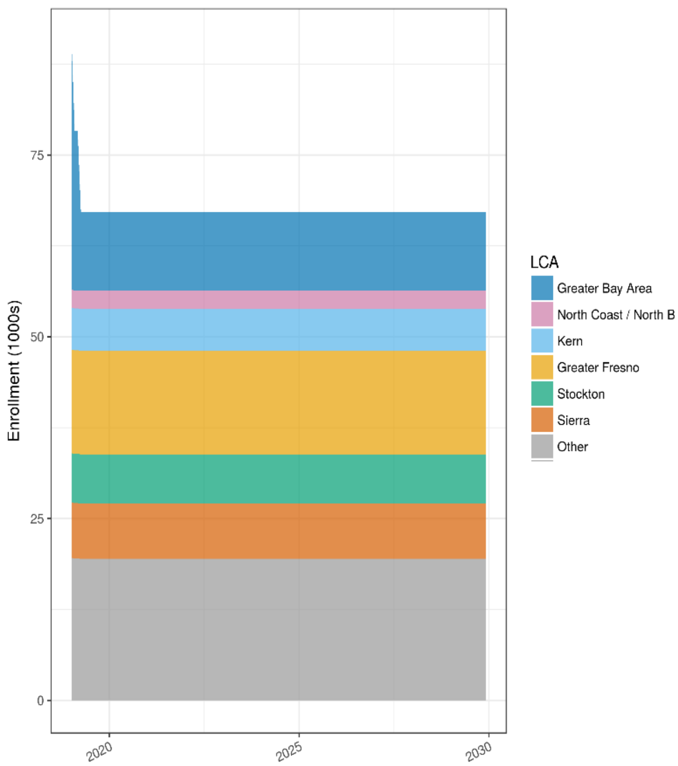

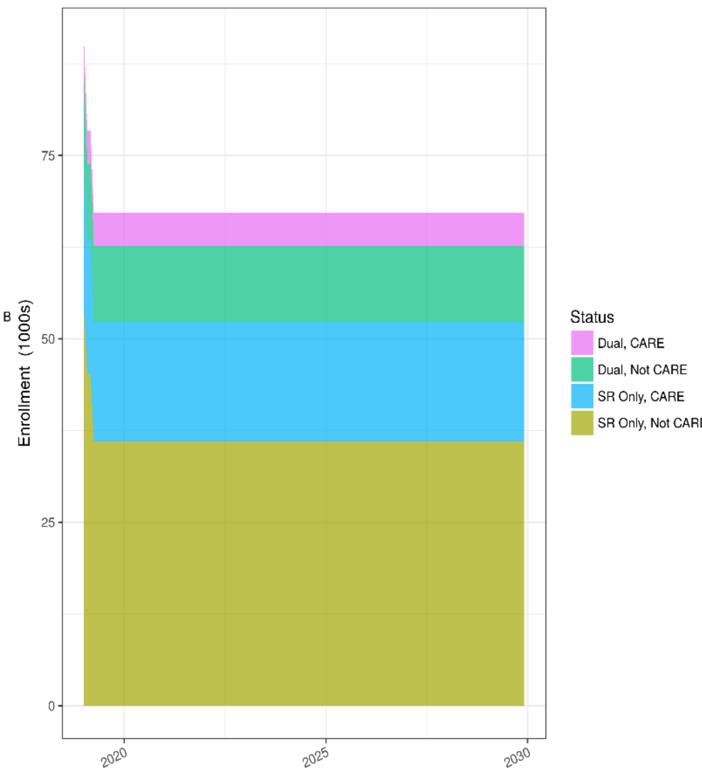

Our model does not explicitly incorporate TOU status. CDA assumed that future participants will behave in a manner similar to existing participants, even if more are on TOU rates. CDA was unable to incorporate TOU segments in this model due to limitations in the current data (i.e., the small number of TOU participants in SmartRate™ in some LCAs and customer categories and the absence of TOU rate types in the enrollment forecast).

Ex Ante Load Reduction

CDA forecasts aggregate impacts of 8.6 MW during the average hour of the 4-9 pm resource adequacy window assuming SmartRate™ events that run from 2 – 7 pm for an August peak day using the 1-in-2 PG&E weather scenario. Figure 25 shows both the mean impact per customer (axis on left) and the aggregate impact (axis on right) for PG&E and CAISO (1-in-2 and 1-in-10 weather years) summer peak weather, for the population of customers assumed to be enrolled in future years. The weather during the CAISO peak is different from the PG&E peak, producing differently shaped (and generally lower) predicted load shed than those for PG&E. The aggregate impact calculation is equal to the mean impact per customer (values from the axis on the left side of the graphic) times the count of approximately 67,000 customers forecast for future program years.

In aggregate, ex ante estimates for the aggregate load shed (averaged over the five-hour-long RA window) reach over 10 MW in *July* of PG&E's 1-in-10 weather year, as shown in the bottom left figure below (using the right-hand scale).

Figure 25. Mean and aggregate ex ante impacts during 4-9 RA event window (4 scenarios, summer only)

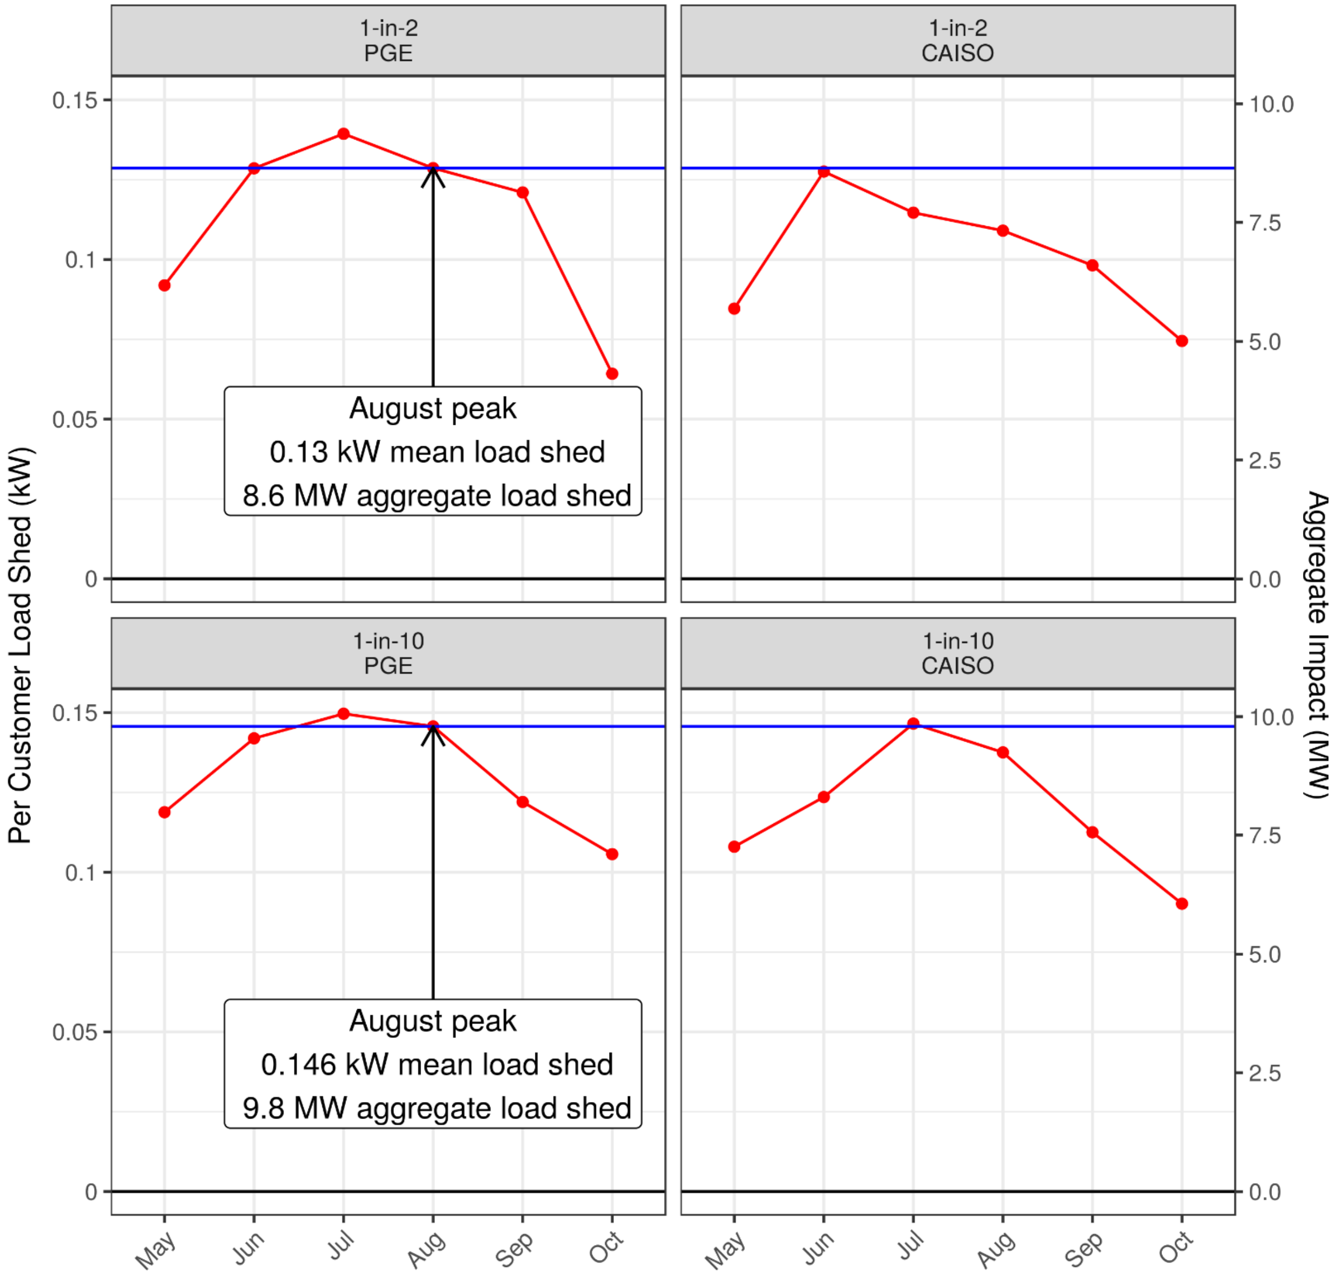

Figure note: The left and right axes apply to all four figures. Mean load shed is indicated on the left, and aggregate on the right, as demonstrated by the August examples in PG&E 1-in-2 and 1-in-10 with the blue line leading left and right.

The table below shows program and portfolio-adjusted impacts for the PG&E 1-in-2 scenario by program enrollment type. Portfolio-adjusted results assume all of the forecasted impacts from SmartRate™ only customers, and 20% of the forecasted impacts from the dually enrolled participants. A detailed table of results is in the summary section at the end of the chapter.

Table 13. Ex ante program and portfolio-adjusted load aggregate impacts – PG&E 1-in-2, 4-9 RA window

	Program			**Portfolio-Adjusted**		
	SmartRate™ only	Dually enrolled	Total	SmartRate™ only	Dually enrolled	Total
January Peak	***0.66***	***0.93***	***1.58***	***0.66***	***0.93***	***1.58***
February Peak	***0.66***	***0.93***	***1.58***	***0.66***	***0.93***	***1.58***
March Peak	***0.66***	***0.93***	***1.58***	***0.66***	***0.93***	***1.58***
April Peak	2.40	1.10	3.54	2.40	1.10	3.54
May Peak	4.26	1.86	6.18	4.26	0.37	4.64
June Peak	5.63	3.00	8.64	5.63	0.60	6.22
July Peak	6.01	3.31	9.37	6.01	0.66	6.67
August Peak	5.61	2.99	8.64	5.61	0.60	6.21
September Peak	5.29	2.80	8.13	5.29	0.56	5.85
October Peak	***3.00***	***1.28***	***4.32***	***3.00***	***0.26***	***3.25***
November Peak	***0.66***	***0.93***	***1.58***	***0.66***	***0.93***	***1.58***
December Peak	***0.66***	***0.93***	***1.58***	***0.66***	***0.93***	***1.58***

Table notes: Bold italics indicates less certainty because there were no events during this timeframe.

Ex Ante Load Reduction by Enrollment Status

While dually enrolled participants are expected to continue to contribute the highest per customer load shed (see blue in left-hand figure), in *aggregate* (see blue in right-hand figure below) they are not anticipated to contribute as much as SmartRate™ only customers because there are many fewer of them.

Figure 26. Mean (a) and aggregate (b) load shed by enrollment status (PG&E 1-in-2)

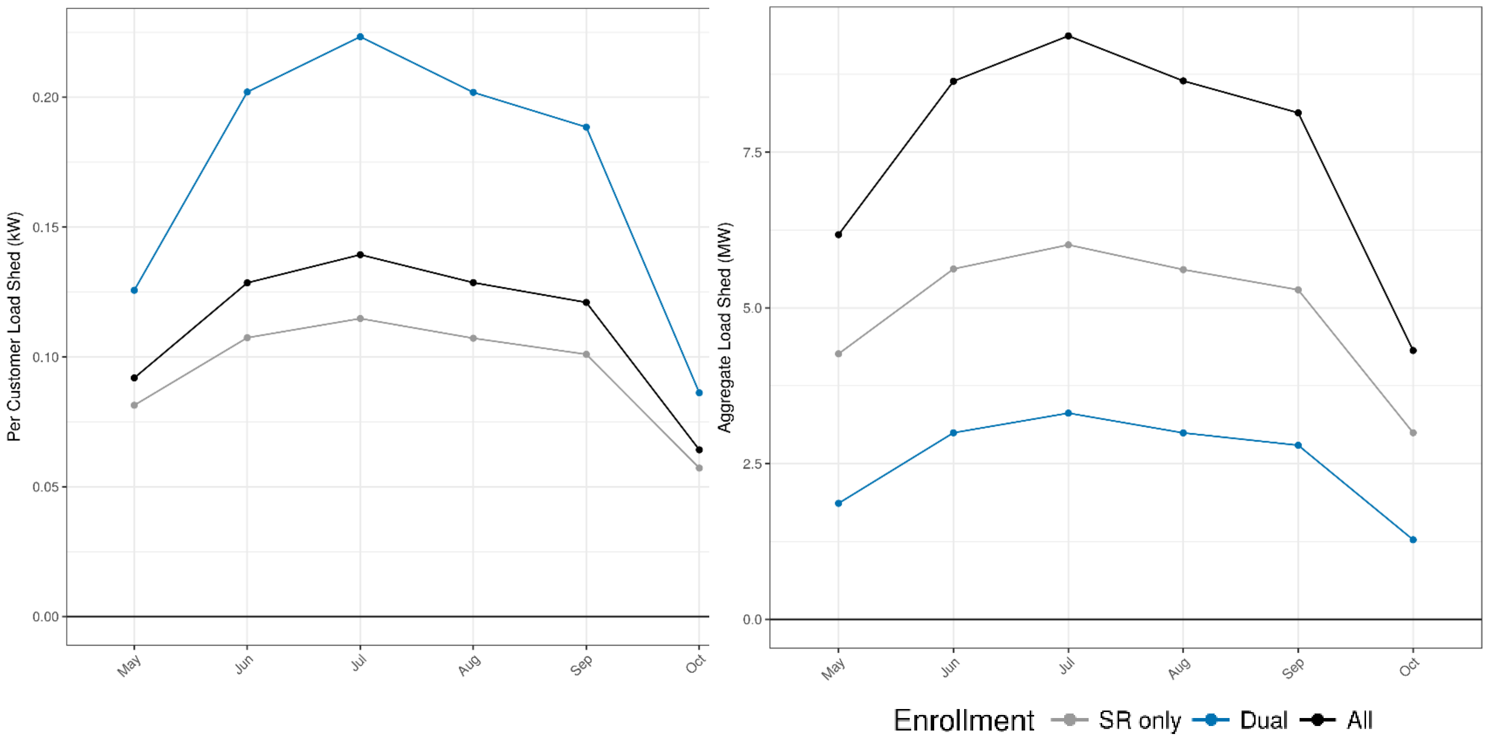

A more detailed discussion of the effects of dually enrolled customers can be found in the section comparing across years, see the *Current Ex Ante to Prior Ex Ante* section at the end of this chapter.

Ex Ante Load Reduction by LCA

The per-customer ex ante impacts vary across the local capacity areas. Figure 27 shows mean ex ante load shed by LCA for each of the weather years using the 4-9 pm RA window.

Figure 27. Mean impact per customer by LCA (4-9 RA window)

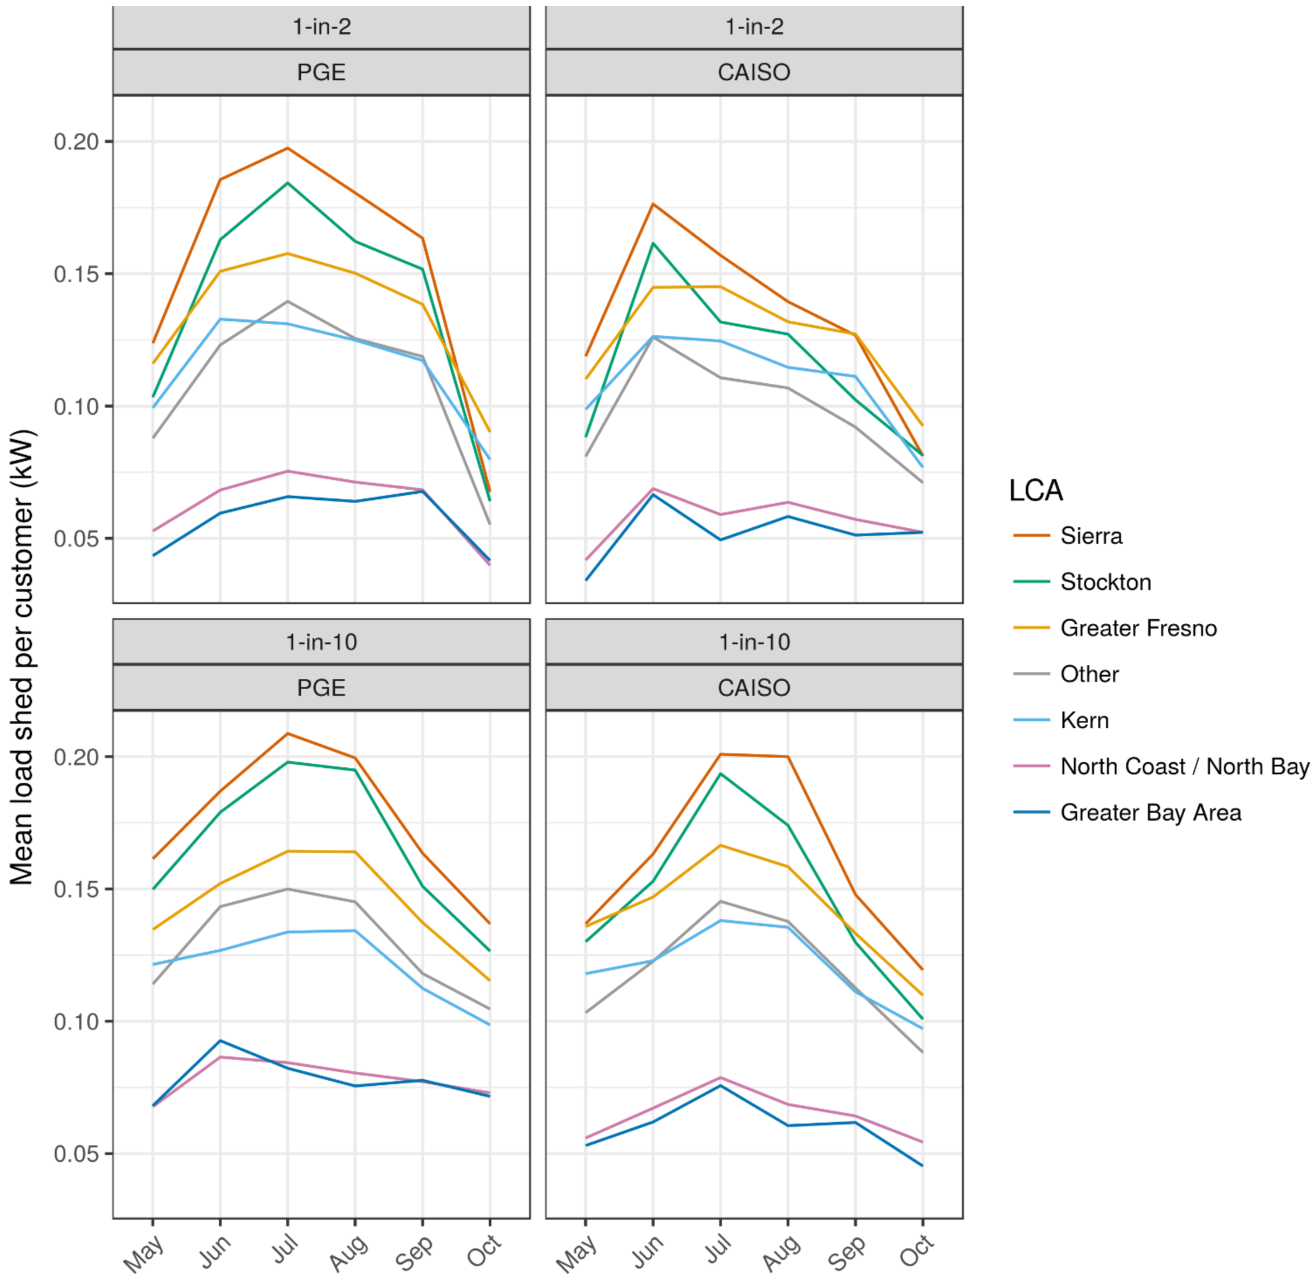

In the figure above, Sierra, Stockton and Greater Fresno have the highest load shed per customer while the Greater Bay Area and the Northern Coast are among the lowest.

Kern does not change much across the scenarios. In the 2018 events, Kern had very low temperature sensitivity – load shed was about the same on very hot days as on more moderate days – and the statistical model assumes that will also be true in future years—so Kern's load shed projections are about the same in moderate years (1-in-2) as in hotter years (1-in-10).

Like the mean load shed per customer, the aggregate impacts also vary greatly between LCAs, due in part to the difference in weather, but even more because of the differences in enrollment among LCAs, as shown in Figure

28 (a) which presents the forecasted ex ante enrollment by LCA. The LCAs are in order from highest load shed per customer at the top to lowest at the bottom. Figure 28 (b) shows the relative contribution of each LCA to the overall aggregate load shed. This is the product of the load shed per customer times the number of customers. The Greater Fresno LCA (middle gold-color band) contributes the most to the overall load shed because that LCA has a high mean load shed per customer due to the weather that they experience and a large number of participants. The Greater Bay Area contributes little to the program compared to most other LCAs due to very significant declines in enrollment from CCA defections and relatively poor load shed per customer (primarily due to mild weather).

Figure 28. (a) Forecasted enrollment and (b) Aggregate ex ante impacts by LCA (4-9 RA window)

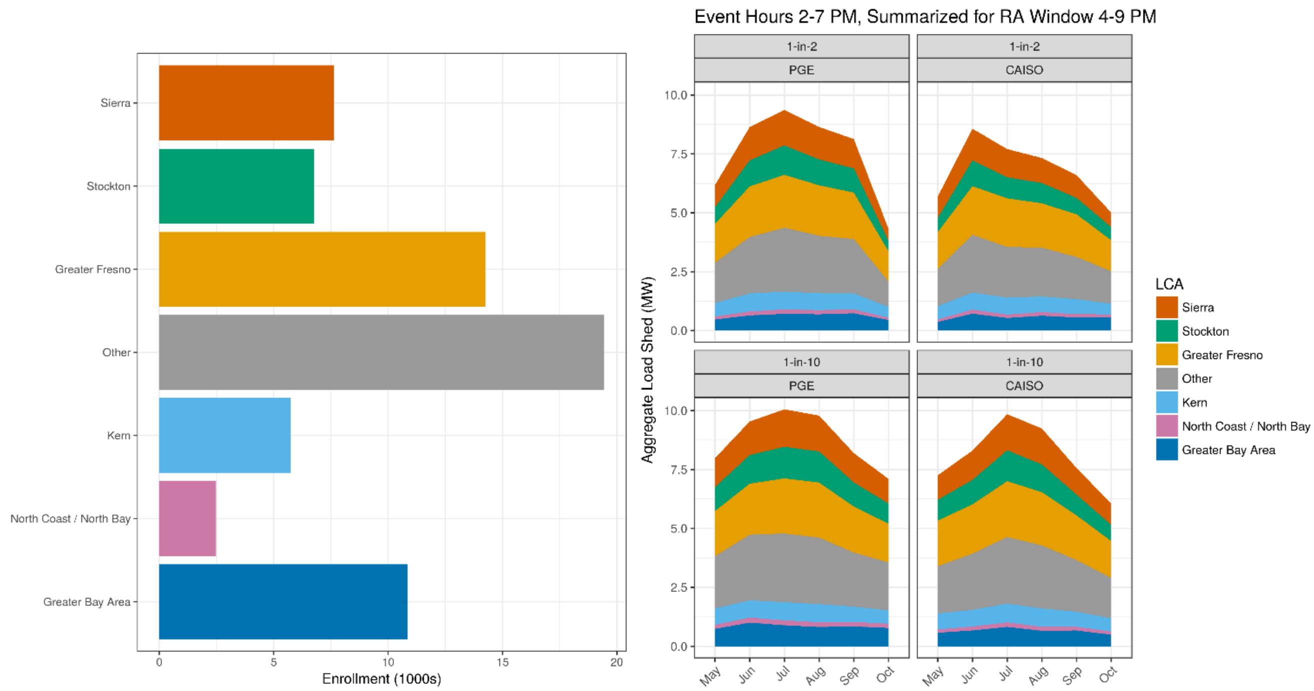

Table 14. Aggregate ex ante impacts by LCA in the PG&E 1-in-2 Weather

LCA	SmartRate™ only	Dually enrolled	Total
All	5.61	2.99	8.64
Sierra	0.91	0.47	1.38
Stockton	0.61	0.49	1.10
Greater Fresno	1.44	0.70	2.14
Other	1.52	0.89	2.44
Kern	0.56	0.16	0.72
North Coast / North Bay	0.14	0.04	0.18
Greater Bay Area	0.44	0.25	0.69
Table note: Shown in order to match the figure above.			

The figures and table above apply to the ex ante predictions for all future years 2020 and beyond.

Current Ex Ante to Current Ex Post

The current ex post and ex ante aggregate load impacts differ for a variety of reasons. Table 15 provides the numbers for comparison, and Table 16 summarizes the factors that influence the relationship between ex post observations and ex ante predictions.

Table 15. Comparison of current ex ante to current ex post load impacts for the typical event day in 2018 (load forecast for 1-in-2 PG&E weather scenario for SmartRate™ for event hours 2 to 7, August peak day)

			Aggregate		Per Participant			
		Enrollment	**Load Impact (MW)**	**Ref. Load (MW)**	**Load Impact (kW)**	**Ref. Load (kW)**	**% Load Impact**	**Average Event Temp**
All	2018 ex post	109,972	17.83	166.02	0.16	1.51	10.7%	88
	Current ex ante	67,206	16.26	132.89	0.24	1.98	12.2%	97
SmartRate™ only	2018 ex post	91,515	10.90	130.61	0.12	1.43	8.4%	87
	Current ex ante	52,380	9.18	101.63	0.18	1.94	9.0%	96
Dually enrolled	2018 ex post	18,456	6.97	35.37	0.38	1.92	19.7%	94
	Current ex ante	14,826	6.98	30.93	0.47	2.09	22.6%	99

Table note: The 'enrollment' figures for 2018 ex post show the mean number of participants, not the mean numb er of enrolled customers. These differ because signaling errors caused lack of participation by some customers in the first two SmartRate™ events of 2018. Humboldt numbers are not included.

Table 16: Important factors that relate ex post observations to ex ante predictions

Factor	Ex Post	Ex Ante	Magnitude of Impact of Assumptions
Measurement window	2018 SmartRate™ events took place from 2-7 pm	Resource adequacy credit only applies from 4-9 pm, which covers three hours of the event	Large impact: Only three hours of the five-hour event are credited, and the resource adequacy period also includes two hours of 'snapback'. Mean load shed per hour is 7.9 MW higher in the event window than the RA window.
Total enrollment	The typical event population in 2018 was 109,972	The population is estimated to decrease to 67,206 in future years.	The decrease in ex ante enrollments decreases the total load impact relative to ex post by about 1.8 MW. Most de-enrollment is projected to take place in the Greater Bay Area.
Loss of most Greater Bay Area participants	The program had more than 33,000 Bay Area participants in 2018	Enrollment is impacted by CCA roll out and projected to fall to under 11,000 participants by summer 2019	Bay Area aggregate load impact decreases by about 2/3 compared to 2018. This is a large part of the total enrollment decrease.
Dually enrolled in SmartAC™ and SmartRate™	The program had 18,522 dually enrolled customers at the end of the 2018 season, representing 17% of the overall population.	The forecast assumes 14,826 dually enrolled participants, which make up 22% of the overall population.	While the proportion of dually enrolled customers is increasing, the overall savings will decrease because there are fewer participants.
Weather	Event days in 2018 were cooler than normal in the Bay Area and Northern Coast. Other LCAs close to normal. (Here "normal" means the PG&E 1-in-2 typical event day).	Predictions are made for typical and for hot weather.	The Ex Ante predictions for 1-in-10 weather years are higher than ex post. Other ex ante predictions can be higher or lower depending on scenario and month.

Step-by-Step from Ex Post to Ex Ante

The figure below summarizes all of the changes between ex post and ex ante.

Figure 29: Relationship between Ex Post and Ex Ante Load Shed during the 4-9 PM Resource Adequacy Window

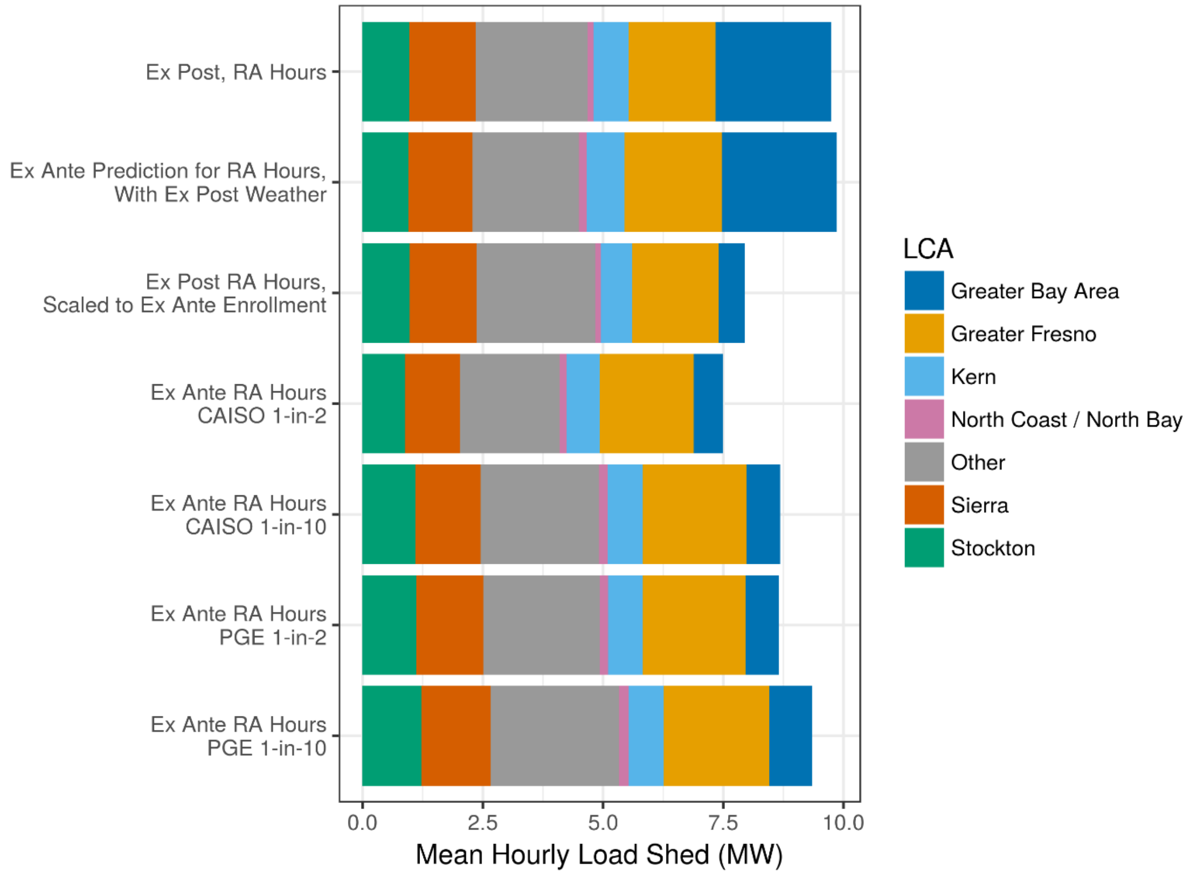

Figure 29 shows the relationship between ex post and ex ante predictions. The top bar shows the mean ex post load impact for the 2018 events (excluding the first two events, for which there were technical problems so that not everyone in the program was informed of the event day). The ex ante model was used to predict the load shed by LCA, given the weather the same event days summarized in the first bar; the resulting predictions are summarized in the second bar.

Summertime program enrollment in 2019 and beyond is projected to differ substantially from 2018, in particular in the Greater Bay Area. To generate the third bar, the ex post load impact from 2018 was scaled to the 2019 numbers, separately within each LCA and for each dual-enrollment status. For example, dual-enrolled customers in the Greater Bay Area are projected to drop from 6,267 to 2,306, so the third bar scales the contribution from such customers to 2306/6276 = 37% of its 2018 value. This bar therefore provides an estimate of what the program would be expected to contribute in 2019 (or future years) if the weather is the same as 2018.

The lower four bars show the mean hourly load shed for the four standard Ex Ante scenarios. Contributions from all LCAs are higher in extreme weather years (the 1-in-10 scenarios) than in more typical years (1-in-2), but details depend on the mix of dual-enrollees versus SmartRate-Only customers, the temperature sensitivity of each group of customers, and the temperature assumptions of each scenario.

Current Ex Ante to Prior Ex Ante

This year's ex ante estimates are lower than the prior year's ex ante estimates (see Table 17 and Figure 30).

Table 17. Comparison of current ex ante load forecasts to prior ex ante
(1-in-2 PG&E weather scenario for SmartRate™ for event hours 2 to 7, August peak day)

			Aggregate		Per Participant			
		Enrolled	**Load Impact (MW)**	**Ref. Load (MW)**	**Load Impact (kW)**	**Ref. Load (kW)**	**% Load Impact**	**Avg. Event Temp**
All (Program)	Prior ex ante 2019+	59,697	13.00	104.93	0.22	1.76	12.39	94
	Current ex ante	67,206	8.64	135.17	0.13	2.01	6.39	93
All (Portfolio)	Prior ex ante 2019+	59,697	7.60	104.93	0.13	1.76	7.24	94
	Current ex ante	67,206	6.21	135.17	0.09	2.01	4.60	93
SmartRate™ only (Program/ Portfolio)	Prior ex ante 2019+	44,011	6.30	70.42	0.14	1.60	8.95	92
	Current ex ante	52,380	5.61	103.57	0.11	1.98	5.42	92
Dually enrolled (Program)	Prior ex ante 2019+	15,686	6.70	34.51	0.43	2.20	19.42	98
	Current ex ante	14,826	2.99	31.38	0.20	2.12	9.54	95
Dually enrolled (Portfolio)	Prior ex ante 2019+	15,686	1.30	34.51	0.08	2.20	3.77	98
	Current ex ante	14,826	0.60	31.38	0.04	2.12	1.91	95

Table note: References loads are taken as reported for the 2018 group in the report. No other information was available at the time of this report. Humboldt is not included.

Figure 30 shows ex ante predicted load shed for dually enrolled, SmartRate™ only and All customers in 2019 that were made as part of evaluating the 2017 SmartRate™ program (i.e., "prior ex ante estimates"), along with the ex ante predictions made for 2019 in the present report. In both cases the predictions are for events that take place from 2-7 pm, but the Prior Ex Ante results summarize the behavior during the previous 1-6 pm. Resource Adequacy period, whereas the current ex ante predictions are for the new RA period of 4-9 pm. The prior ex ante results thus included four event hours and no snapback hours, while the current ex ante results include only three event hours and two snapback hours; these factors significantly reduce the current ex ante load impact compared to prior ex ante. The difference for SmartRate™ only customers is smaller than for dually enrolled customers. The current ex ante numbers do assume more than 7,000 additional participants in the program than were assumed last year, but this increase is not enough to overcome the substantial decrease in aggregate load impact that is due to the changed RA window.

Figure 30. Aggregate Ex Ante Impact Estimates: Comparison Across Report Years (PG&E 1-in-2)

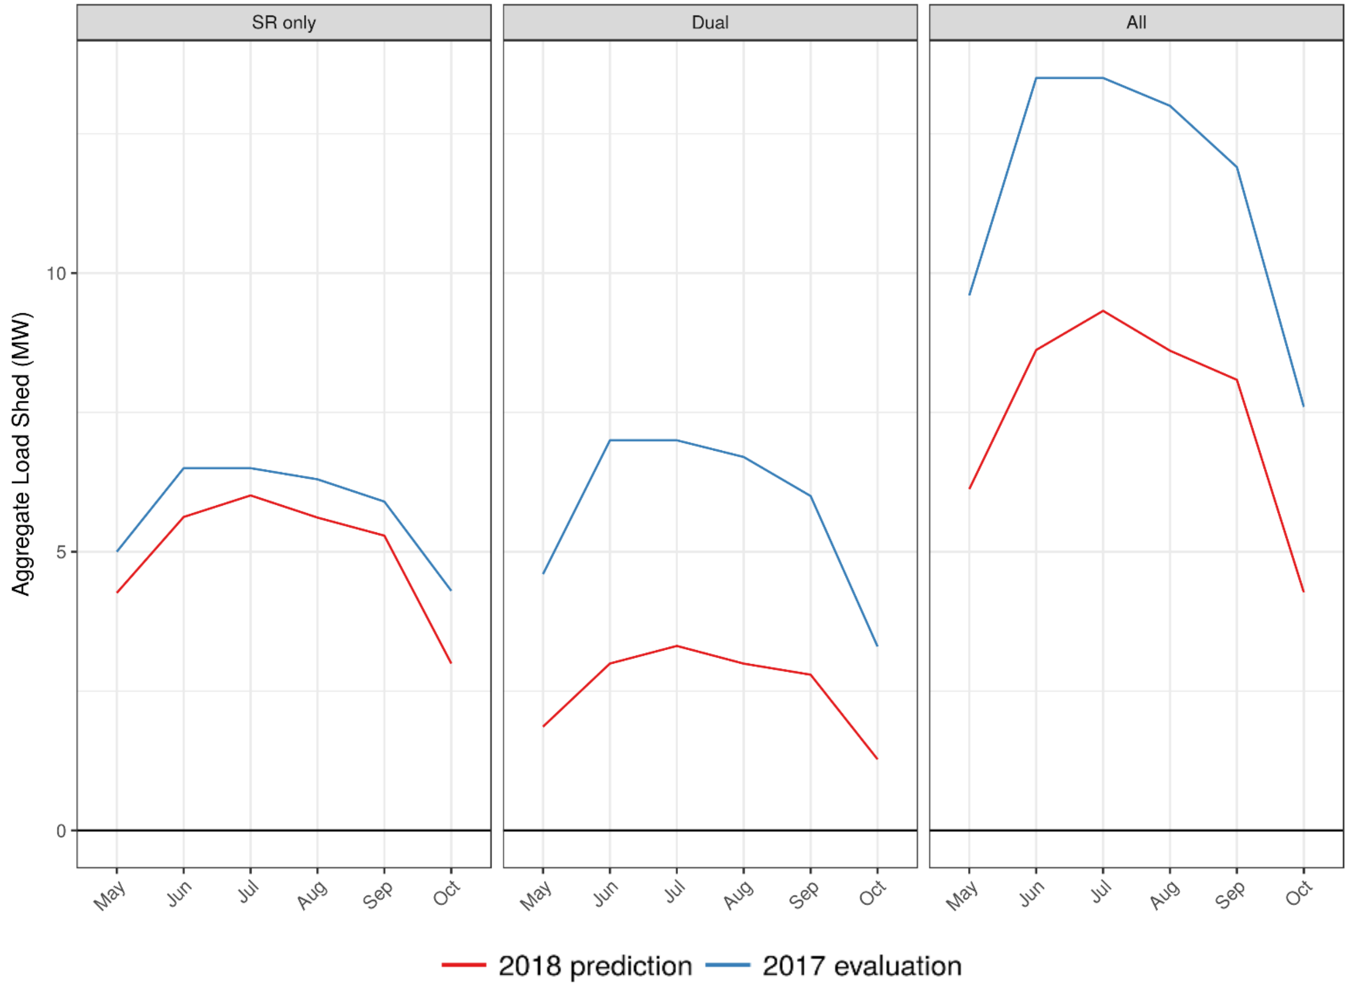

There are four reasons this year's ex ante predictions are different from last year's:

1. **Changes in RA window -** This year's evaluation is based on a resource adequacy window from 4-9 pm rather than the 1-6 pm period that was used in previous years, while events for both years were called from 2-7 pm.
2. **Changes in performance -** There were changes in the load shed performance of customers in the program – such as decreased load shed and decreased temperature sensitivity in Kern compared to previous years – and these changes are projected to persist into the future.
3. **Changes in enrollment -** The number of enrolled customers in the Greater Bay Area is projected to decrease substantially compared to last year's assumptions, and there are also changes in projected enrollment in other LCAs and in the mix of SmartRate™ only and dually enrolled customers.
4. **Changes in the modeling -** The statistical model used this year is different from that used in previous years; for example, it takes account of the significant difference between CARE and non-CARE customers in load shed as a function of temperature.

Below we discuss each of the four individual factors that lead to differences between the predictions from the 2017 evaluation and the predictions from the current evaluation.

Changes in RA Window

The biggest quantitative change is caused by the change in the RA window; had the current evaluation been done with the same 1-6 pm RA window that was used for the 2017 evaluation the predicted aggregate load shed would be roughly 4 to 5 MW higher in the peak months.

The figure below shows the current ex ante estimates using both the older RA window (1-6 pm, in gold) and the newer RA window 4-9 pm in red). Had the RA window remained 1-6 pm as in prior evaluations, the August peak would be 13.3 MW.

Figure 31. Effects of Changing RA Window

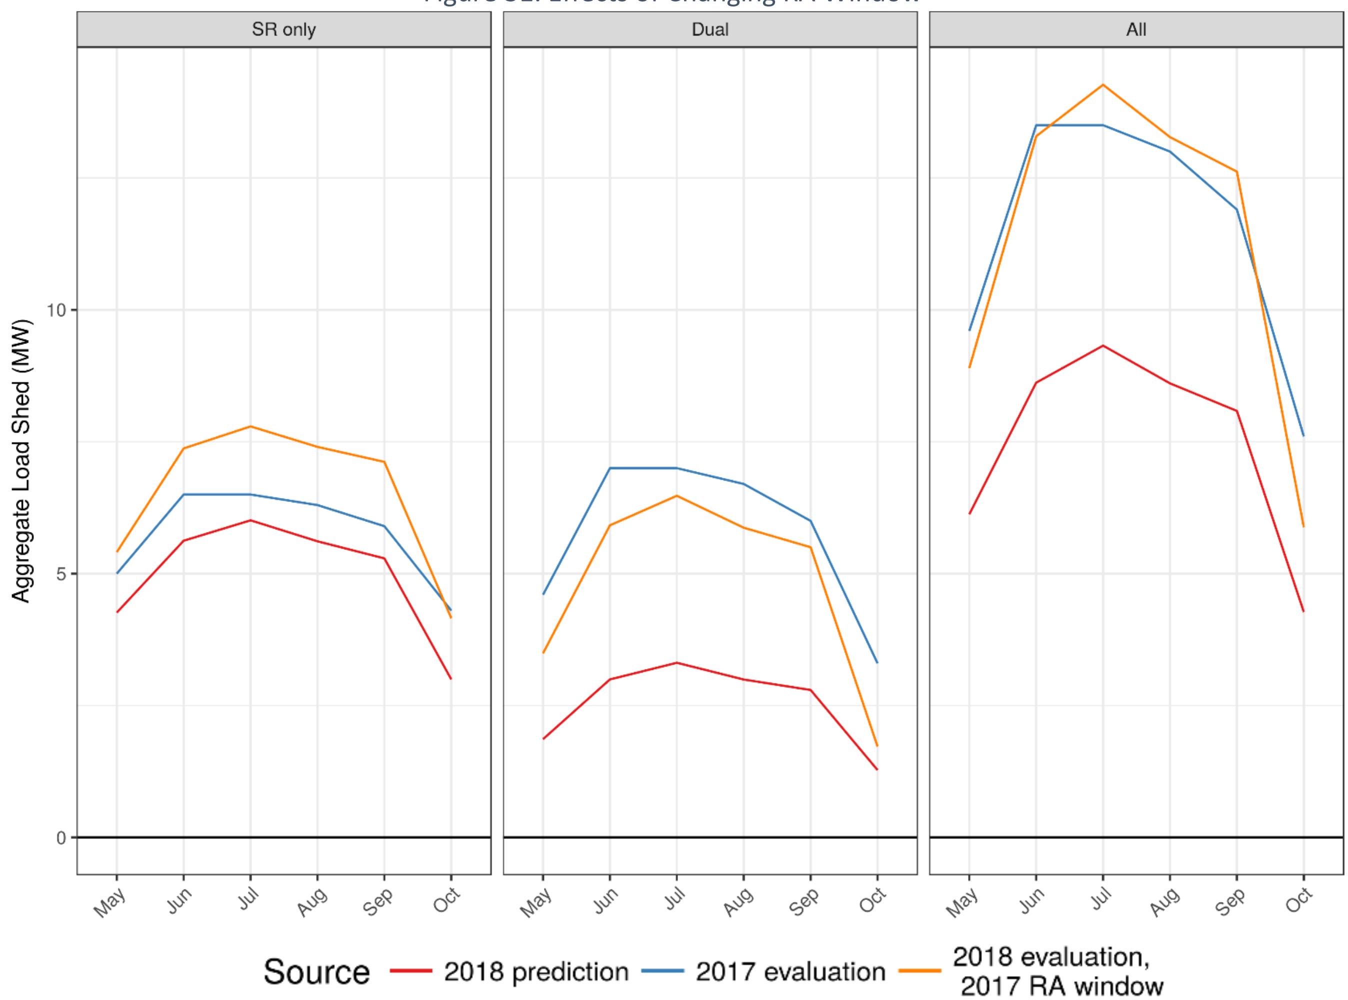

Noticeably in the figure above, changing the RA window for SmartRate™-only customers (left panel) would lead to results that are higher than last year's ex ante estimates: the gold line is above the blue line. However, for dual-enrolled customers (center panel), the changing RA window is not the only reason for the difference in predicted load shed, so there is still a gap between last year's ex ante estimates (in blue) and this year's ex ante estimates (in gold).

Changes in Customer Performance

This year's model reflects a lower load shed from dually enrolled customers in 2018 compared to 2017. Part of this difference is due to the unusually low load impact of dually-enrolled customers in the Kern LCA in the 2018 SmartRate™ events.

Changes in Enrollment Numbers

Forecasted enrollment also has an effect on these numbers. Forecasted enrollment increased from the prior year's estimate of 59,731 to 67,206 (12.5% increase). This new enrollment forecast leads to about a 0.75 MW increase in the estimate, over what it would have been with the 2017 enrollment forecast, however, the forecasted number of dually enrolled customers declined between the 2017 and 2018 enrollment forecasts.

Figure 32. Effect of Changing Enrollment

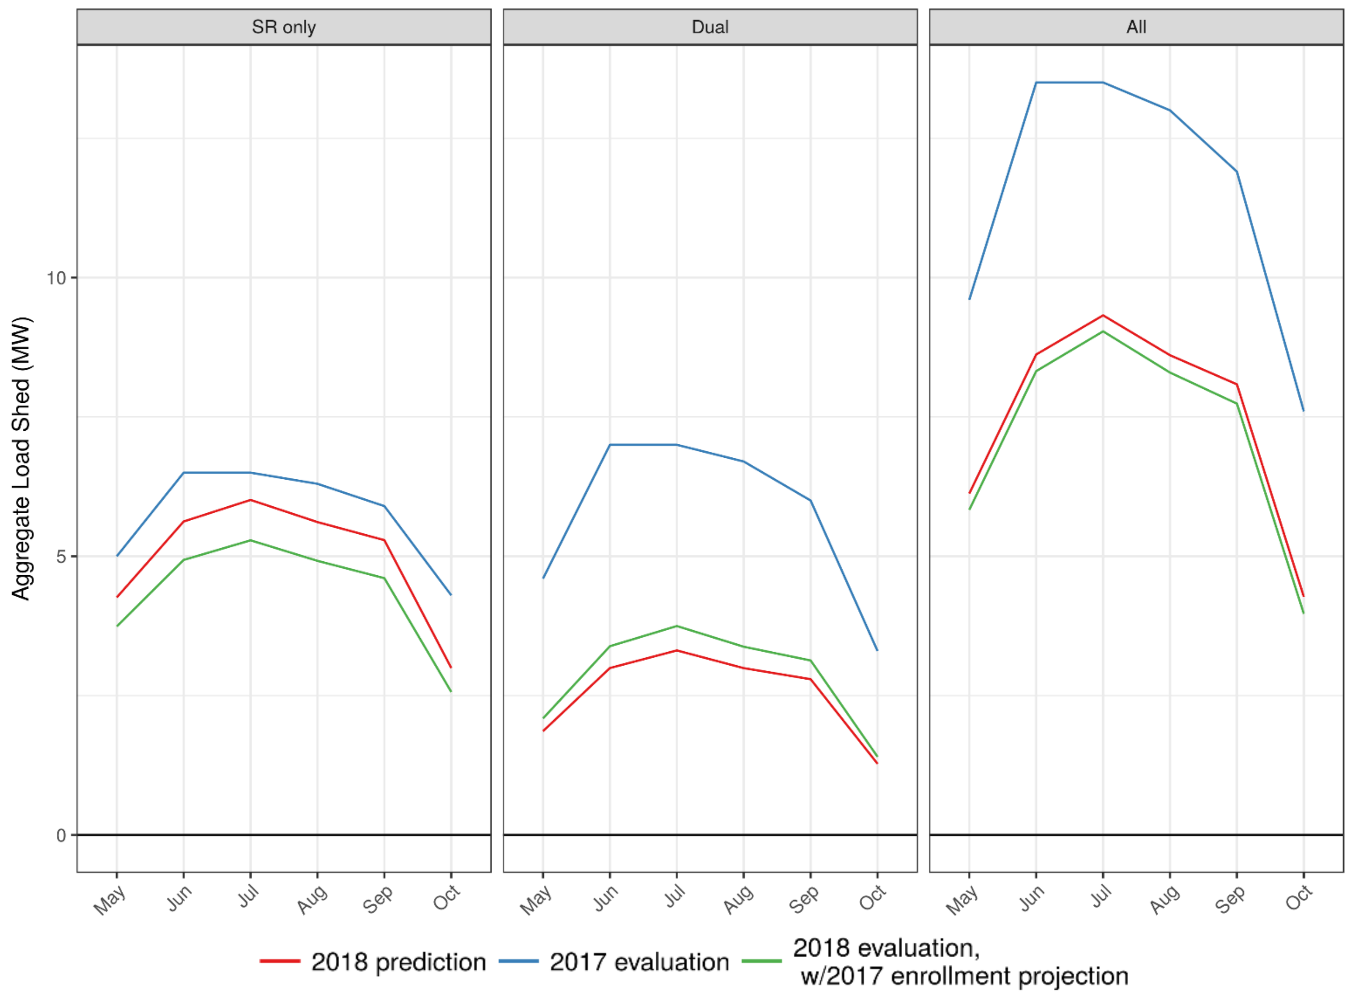

Changes in Model

Finally, some of the difference between years may be due to modeling changes that were implemented to make the ex ante statistical model more accurate. This year's model included assumptions for CARE customers because they are very different. Specifically, the model was fit to the observed outcomes of the interaction between CARE and dually enrolled segments, i.e., 1) Dually enrolled CARE, 2) Dually enrolled non-CARE, 3) SmartRate™ only CARE, and 4) SmartRate™ only non-CARE. The model was used to predict the mean load shed for customers in each of these segments in the ex ante weather conditions. While the magnitude of modeling changes is unknown (since there is no point of comparison from prior years), we do not expect the changes in the model to have been a large source of the differences between the years. This year's model is discussed in detail in Appendix B.

Detailed Summary of Ex Ante Impacts

The detailed ex ante impact estimates for all customers, are shown in table form below (Table 18). These, as well as additional ex ante details in the appendices, are provided to help with long-term resource planning. As described earlier, the highest estimated aggregate impact occurs on the July peak day in every weather scenario.

Table 18: SmartRate™ ex ante load impact estimates by weather year and day type (event period 4-9 pm)

			Per Customer (kW)		Aggregate (MW)	
		Day type	Mean Hourly Impact	Max Hourly Impact	Mean Hourly Impact	Max Hourly Impact
PG&E	**1-in-2**	**Typical Event Day**	**0.13**	**0.27**	**8.69**	**18.33**
		January Peak	0.02	0.03	1.58	2.00
		February Peak	0.02	0.03	1.58	2.00
		March Peak	0.02	0.03	1.58	2.00
		April Peak	0.05	0.10	3.54	6.39
		May Peak	0.09	0.19	6.18	12.69
		June Peak	0.13	0.27	8.64	18.16
		July Peak	0.14	0.29	9.37	19.55
		August Peak	0.13	0.27	8.64	18.29
		September Peak	0.12	0.26	8.13	17.38
		October Peak	0.06	0.13	4.32	8.61
		November Peak	0.02	0.03	1.58	2.00
		December Peak	0.02	0.03	1.58	2.00
PG&E	**1-in-10**	**Typical Event Day**	**0.14**	**0.29**	**9.39**	**19.80**
		January Peak	0.02	0.03	1.58	2.00
		February Peak	0.02	0.03	1.58	2.00
		March Peak	0.02	0.03	1.58	2.00
		April Peak	0.09	0.19	6.19	12.92
		May Peak	0.12	0.25	7.98	16.85
		June Peak	0.14	0.29	9.54	19.80
		July Peak	0.15	0.32	10.06	21.25
		August Peak	0.15	0.31	9.79	20.67
		September Peak	0.12	0.26	8.20	17.50
		October Peak	0.11	0.23	7.10	15.41
		November Peak	0.02	0.03	1.58	2.00
		December Peak	0.02	0.03	1.58	2.00
		Continued next pg.				…

			Per Customer (kW)		Aggregate (MW)	
		Day type	**Mean Hourly Impact**	**Max Hourly Impact**	**Mean Hourly Impact**	**Max Hourly Impact**
CAISO	**1-in-2**	**Typical Event Day**	**0.11**	**0.24**	**7.55**	**15.83**
		January Peak	0.02	0.03	1.58	2.00
		February Peak	0.02	0.03	1.58	2.00
		March Peak	0.02	0.03	1.58	2.00
		April Peak	0.06	0.11	3.80	7.22
		May Peak	0.08	0.17	5.69	11.73
		June Peak	0.13	0.27	8.57	17.96
		July Peak	0.11	0.24	7.71	16.15
		August Peak	0.11	0.23	7.33	15.32
		September Peak	0.10	0.21	6.60	13.89
		October Peak	0.07	0.16	5.01	10.64
		November Peak	0.02	0.03	1.59	2.00
		December Peak	0.02	0.03	1.58	2.00
CAISO	**1-in-10**	**Typical Event Day**	**0.13**	**0.27**	**8.73**	**18.44**
		January Peak	0.02	0.03	1.58	2.00
		February Peak	0.02	0.03	1.58	2.00
		March Peak	0.03	0.04	2.09	2.83
		April Peak	0.09	0.18	5.91	12.27
		May Peak	0.11	0.23	7.26	15.23
		June Peak	0.12	0.26	8.30	17.32
		July Peak	0.15	0.31	9.85	20.70
		August Peak	0.14	0.29	9.24	19.53
		September Peak	0.11	0.24	7.56	16.20
		October Peak	0.09	0.20	6.06	13.17
		November Peak	0.02	0.03	1.64	2.00
		December Peak	0.02	0.03	1.58	2.00

A Deeper Dive to Inform the Future

In this section, we provide two investigations that can help inform the program, as well as future forecasts. This section includes an assessment of enrollment cohort impacts and TOU customer impacts. Below we present the following:

- Enrollment Cohort Analysis - Understanding structural and behavioral changes
- TOU Analysis – Exploration of TOU impacts

Enrollment Cohort Analysis – Understanding Structural and Behavioral Changes

As a pricing program, SmartRate™ provides incentives for customers to shift and curtail their loads, but it does not prescribe how customers should respond. Apart from the automated dispatch of AC controls for dually enrolled customers, all of the impacts of SmartRate™ can be understood as changes in customer behavior. As dually enrolled customers decrease, SmartRate™ impacts will be increasingly based on purely behavioral changes. Having a deeper understanding of behaviors and their determinants is important to understanding the future of the program.

This section examines cohorts of customers that enrolled at different times to understand if there are structural differences in the groups, and how cohort impacts differ. Because enrollments take place year-round but the program events primarily occur in the summer, cohorts are identified using the year of the first event they participated in.

Analysis by enrollment year

Examining the number of dually enrolled customers versus SmartRate™ only customers enrolled over time gives us our first insight that the types of customers that enrolled year to year are very different. The program recruited a large number of dually enrolled customers during the 2012-2014 period. As shown in Figure 33 (note the difference in the y-axis scale of the two panels), dual enrollments surged in 2012-2014, with many of those customers still present in the program. SmartRate™ only enrollments, on the other hand, have accelerated in the last two years, with a large fraction of customers coming from the Bay Area in particular. The figure below provides a breakdown, by enrollment year cohort of customers enrolled for the 2018 program year, with dually enrolled and SmartRate™ only customers viewed separately. Note the scales of the y-axis for each are different.

Figure 33. First year of enrollment for all customers enrolled as of the beginning of the 2018 program year, by enrollment status

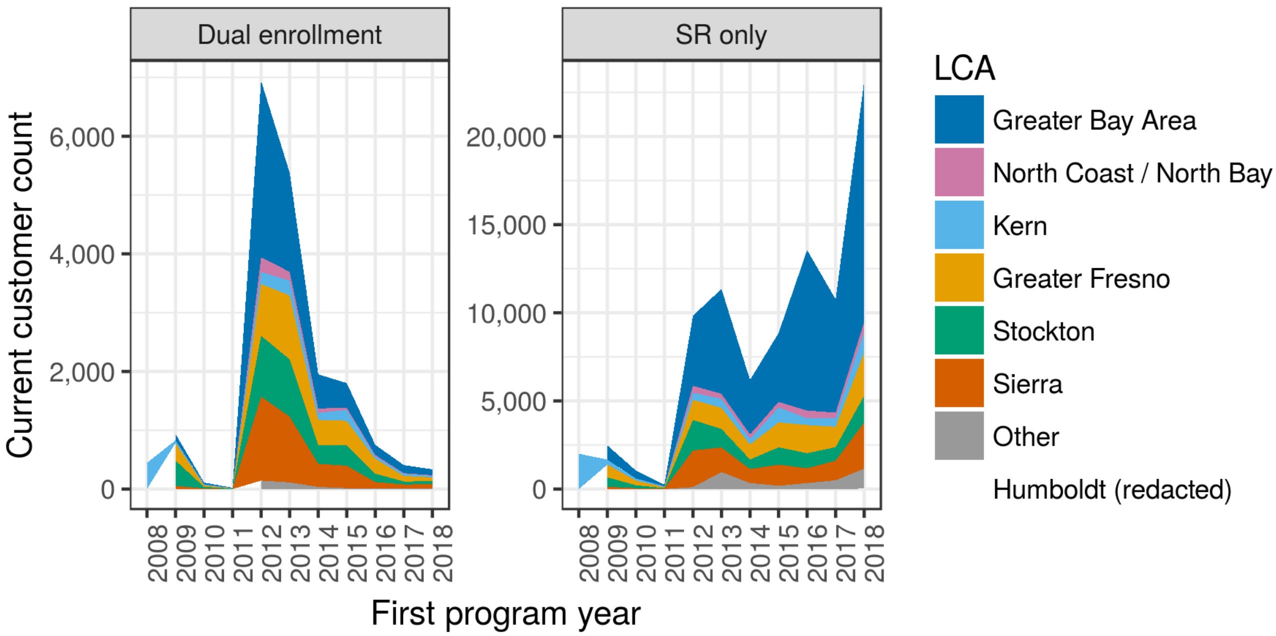

Because dually enrolled customers are enrolled in the SmartAC™ program, which utilizes direct load controls to reduce electricity demand from central air conditioning units (AC) in the customer's home, the impacts are larger (as discussed in this report) and they are partially outside of the control of the customer. As such, they do not tell us as much about customer-controlled behavioral differences between cohorts. Below we describe our method for getting a deeper look at the variation in customer, and customer behavior, that is not due to the automated dispatch of AC controls.

Enrollment cohort impact estimation methods

To concentrate on behavioral differences, dually enrolled customers were removed from the sample that we analyzed. To further isolate program effects as distinct from other programmatic effects, customers on TOU rates were also eliminated. Figure 34 provides a total count of the SmartRate™ only non-TOU program year 2018 participants by first program year cohort and LCA. This subset of customers was used for our analysis.

Figure 34. First year of enrollment for SmartRate™ only customers not on TOU rates enrolled as of the beginning of the 2018 program year

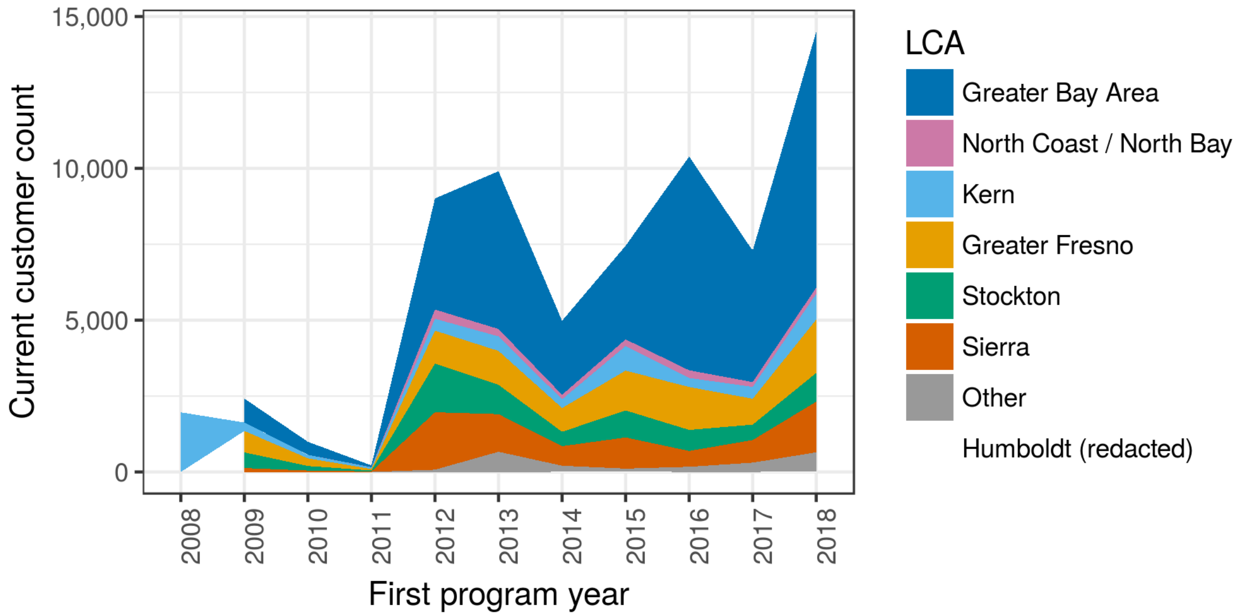

The participant counts in the sample selected are dominated by the Greater Bay Area, with fluctuating, but significant rates of enrollment dating back to 2012 and a handful of founding participants reaching as far back as 2008.

To estimate the reference loads, event outcomes, and other program statistics for each enrollment cohort, the matched control comparison process used to estimate outcomes based on dual enrollment status, CARE participation, TOU rate types, and other sub-groups was *repeated using customer sub-groups based on the cohorts defined using their first year of program participation*. To avoid examining too many cohorts with too few customers, enrollment cohorts were defined annually from 2015 through 2018, with a single cohort from all years prior to 2015. As described above, dually enrolled customers were excluded to focus the analysis on behavioral, rather than automated control, effects and customer on TOU rates were eliminated because they have only recently become available and to avoid complications of attribution across pricing mechanisms.

Because the modeling and post-processing approach was the same, the estimation methods details can be found main ex post methods section of this report*, Ex Post Impact Analysis Methods.*

Results: Event reference loads

We analyzed reference loads for each cohort across all nine events and all LCAs to understand if there are structural difference in consumption across the enrollment cohorts. Over time, the groups may have come from different geographies or have systematically larger or smaller loads.

We use reference loads as a metric of cohort similarity across years and LCAs because they will be different if the loads of each year's new customers are structurally different. Since reference loads are temperature sensitive, the relationship between reference loads and temperature for each enrollment cohort within and across LCAs provides the clearest view of structural differences in consumption. Figure 35 illustrates reference loads vs. outside temperature for each of the enrollment cohorts and LCAs, for each event. Linear fit regression lines help to illustrate the magnitude and temperature sensitivity of each group's reference loads.

Figure 35. 2018 event reference loads vs. outside temperature for each enrollment cohort in each LCA. 2018 events are the dots, each with a regression line illustrating temperature trends for each cohort.

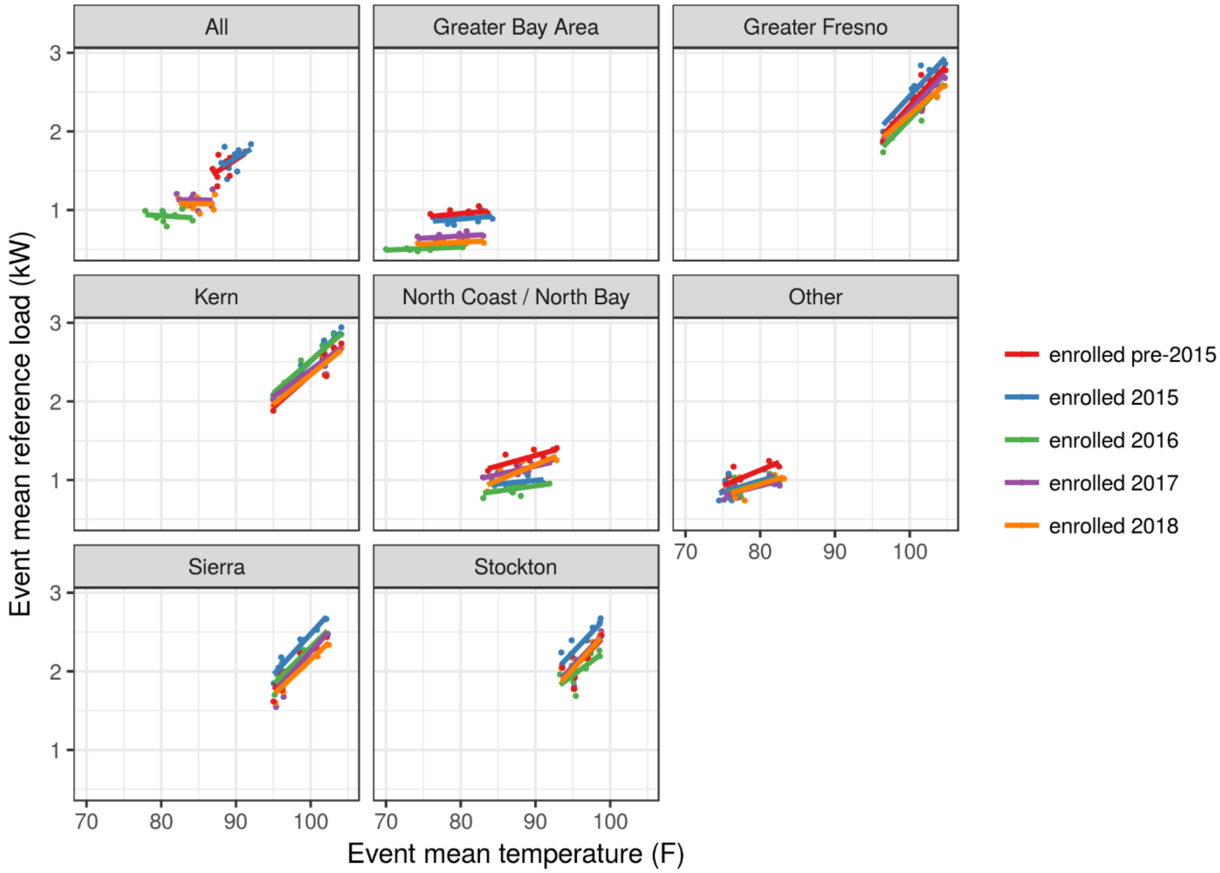

As a result of differential enrollment across LCAs over time, the "All" LCA can be seen to have the highest event temperatures experienced by the earlier cohorts and the lowest temperatures experienced by the 2016 cohort. That same qualitative description applies to the results from the Greater Bay Area, so different temperature regimes can occur within as well as across LCAs. These differences appear to indicate that recruitment and retention has changed the loads of enrolling customers over time, apparently with the earlier cohorts having been recruited from different (hotter) places and with different loads than more recent ones.

If recruitment and retention had been consistent over time, these reference loads should all be overlapping. Since they are not, we conclude that differences in the composition of new enrollment cohorts from year to year contribute substantially to the differences observed across them. In other words, earlier cohorts are not simply the same types of customers as the more recent enrollees with more program experience, where the difference can be explained by the effects of that experience. These differences are likely larger than and would partially obscure any behavioral learning or habituation effects (spillover) that might be able to be seen by looking across enrollment cohorts over time. However, the data do provide evidence that these cohorts are behaving differently than one another. We discuss this in the next section.

Results: Impacts from Behavioral Changes

Figure 36 illustrates impacts vs. outside temperature for each of the enrollment cohorts and LCAs, for each event. Linear fit regression lines help to illustrate the magnitude and temperature sensitivity of each group's load impacts, i.e. behavior during events.

Figure 36. 2018 event impacts vs. outside temperature for each enrollment cohort in each LCA. 2018 events are the dots, each with a regression line illustrating temperature trends for each cohort.

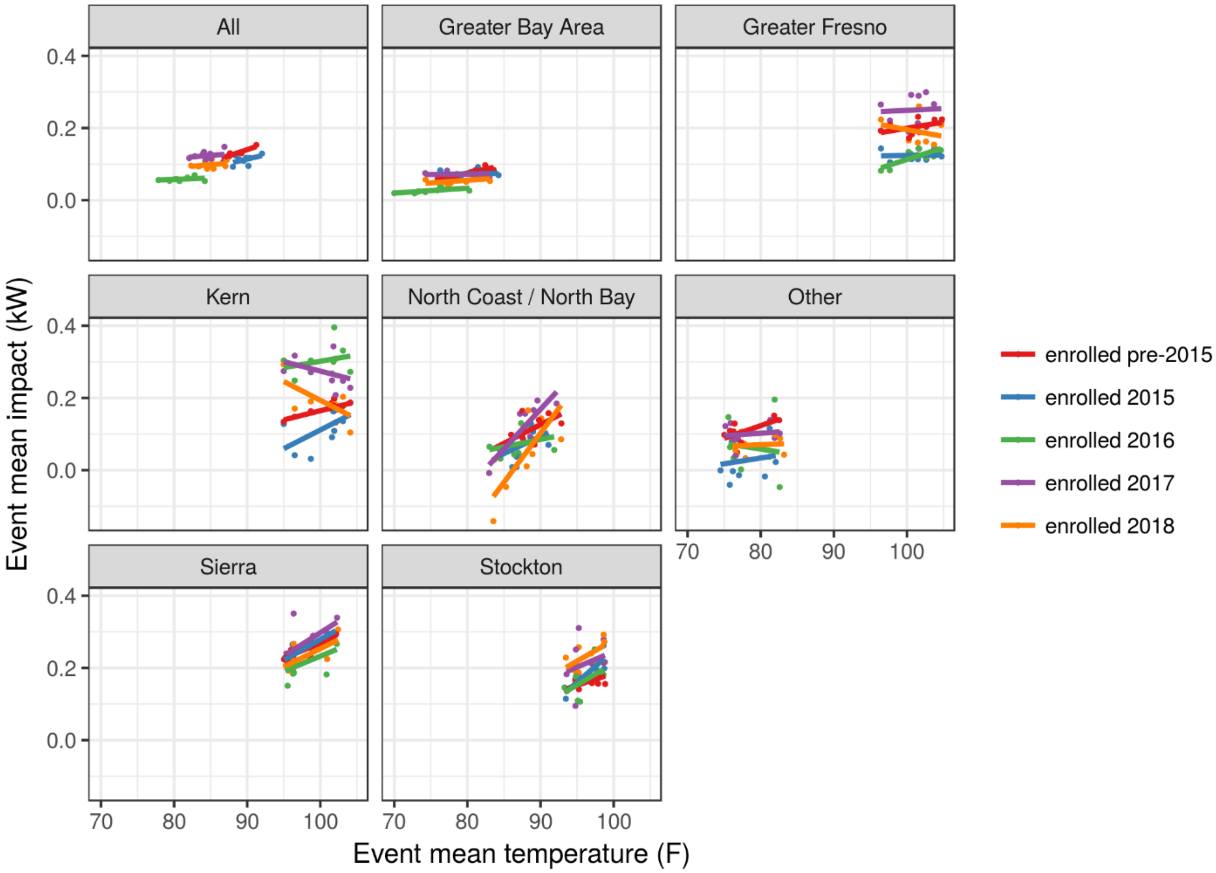

While the reference load slopes (indicating amount of energy that customers need in response to changing temperatures) were relatively consistent within each individual LCA, the heterogeneous slopes in some of the LCAs below indicate that the behaviors (and the magnitude of the impacts that result from those behaviors) are likely different across the enrollment cohorts.

We also see signs that the irregularities in Kern and North Coast LCAs outcomes documented elsewhere. These appear to be driven by specific cohorts of customers. (Kern is also discussed in Appendix C.). There is some indication of a pattern in the magnitude of impacts across cohorts within each LCA. To isolate that pattern, Figure 37 illustrates the typical day impact for each enrollment cohort for each LCA.

Figure 37. Typical event impact by cohort and LCA

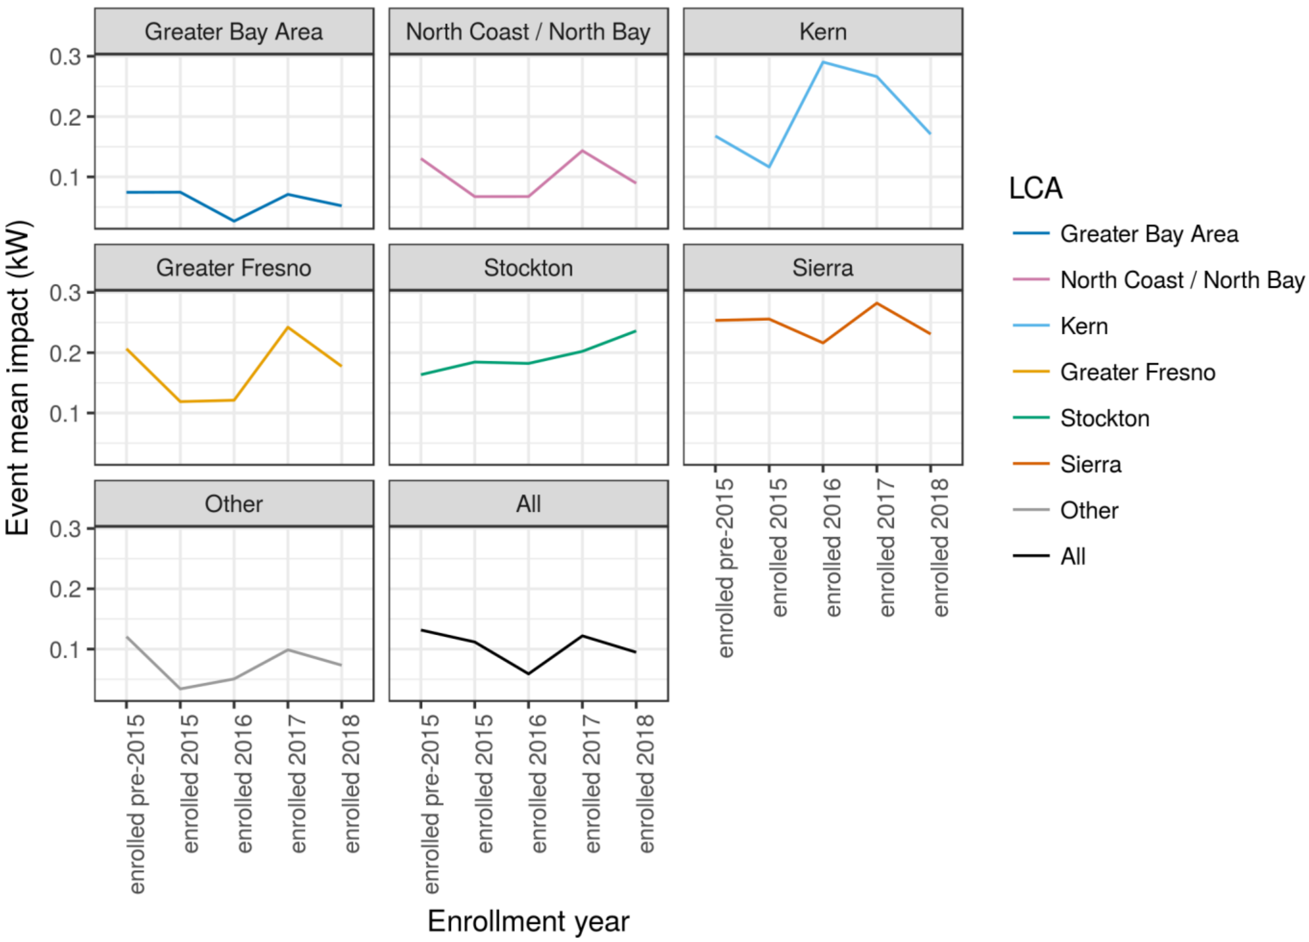

The cohorts are producing different impacts within each LCA, but the impacts are not uniform. There is not a single consistent pattern, like steadily increasing or decreasing savings with program experience. However, within almost all LCAs, the 2018 cohort under performs relative to 2017, suggesting the possibility that they are still learning how to respond to the program, the 2016 cohort tends to under-perform and the pre-2015 cohort tends to out-perform compared to other cohorts.

Figure 38 presents load impact as a percentage of reference load by enrollment cohort and LCA. The qualitative patterns seen in Figure 37 are more or less preserved, suggesting that the differences between cohorts should be interpreted more as behavioral impacts rather than structural differences in consumption (i.e. the differences are not explained by reference loads).

Figure 38. Typical event impact as a percentage of reference loads by enrollment cohort and LCA.

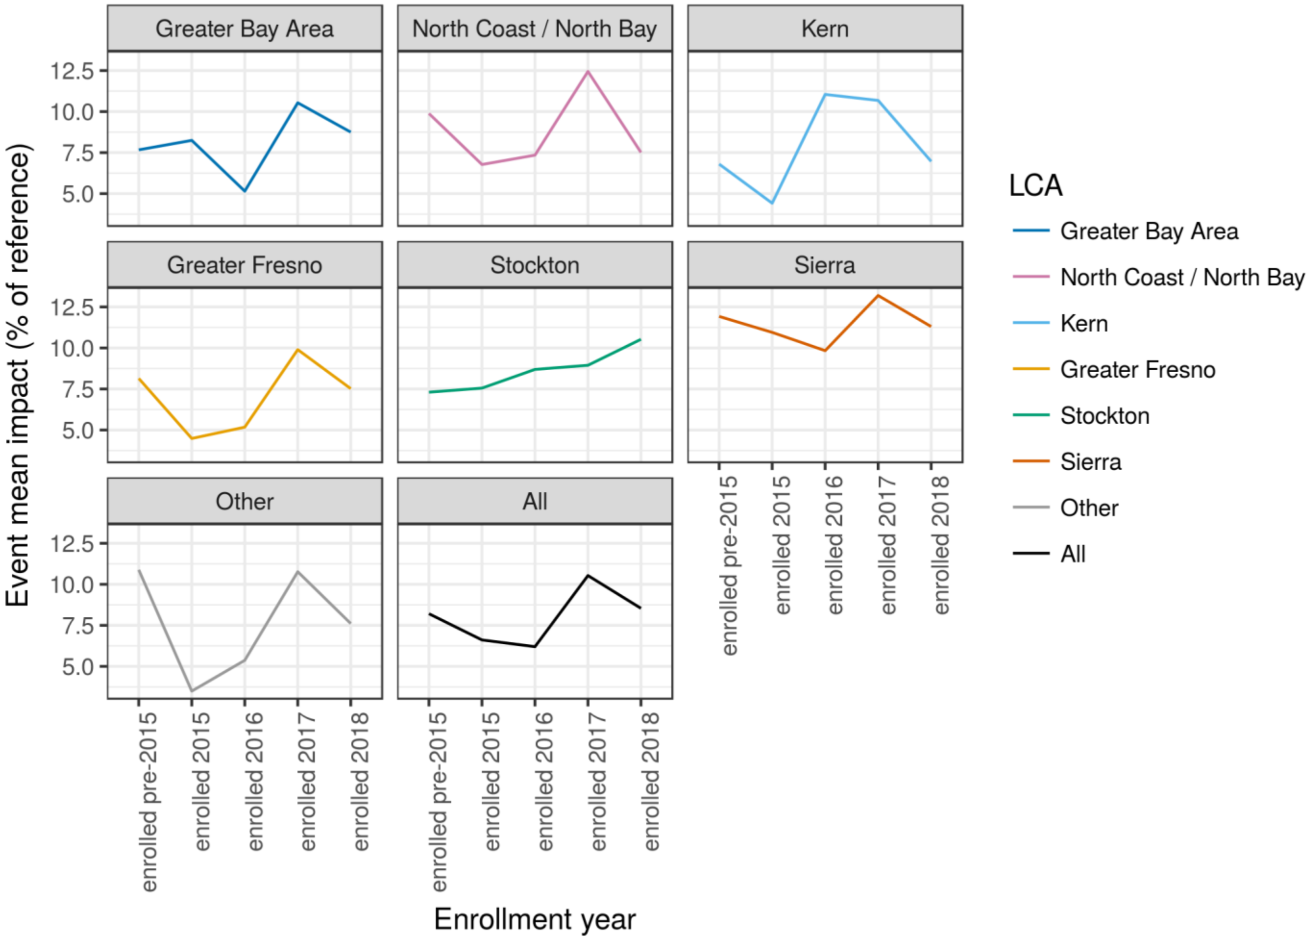

Enrollment Cohort Summary – A Look Towards the Future

We have documented significant structural differences in reference loads and temperature experienced by different enrollment cohorts. We have also shown that the cohorts are likely taking different behaviors in response to events. However, given the large differences across cohorts and time, these results do not provide conclusive insights about habituation and learning from year-over-year participation. They do, however, demonstrate that past marketing, targeting strategies, participant restrictions, and other factors that shape each enrollment cohort have an enduring effect on program outcomes. They also suggest that future impact assessments may want to consider building a cohort analysis into the model to improve estimates.

TOU Analysis – Exploration of TOU Impacts

The CDA team also looked at event impacts for customers on TOU rates to try to understand the variation in impacts across customers who are not on a TOU rate, and those on either opt-in or default-in rates. For this analysis, the rates TOU-A, TOU-B, and E6 are classified as “TOU-opt-in” rates because they are optional, and customers must self-select to be enrolled in them. TOU-C is classified as a “TOU-default-in” rate because it is the default rate (with bill protection) associated with the 2018 default pilot. All other rates are considered to be “non-TOU”.

The SmartRate™ program was originally enabled by the widespread deployment of SmartMeters™ that record electricity consumption on an hourly basis and allow for the tabulation of costs that vary as a function of time of day or for limited event periods. As a program based on a rate modifier that applies time varying costs on top of any given underlying rate, SmartRate™ is a natural complement to and can be applied on top of the time of use (TOU) rates that were similarly enabled by SmartMeters™. Only a small fraction of PG&E customers is currently on a TOU rate, but TOU rates are currently slated to become the default choice for customers in the fall of 2020. Since TOU rates already expose customers to (continuous and gentler) time varying costs, it has become important to empirically quantify the additional SmartRate™ impacts on customers already enrolled in TOU rates.

TOU SmartRate™ Participants

As of the last event of the 2018 program season on 2018-07-26, just under 20% of SmartRate™ participants were on TOU rates, with 17% on a TOU-opt-in rate and 3% on a TOU-default-in rate. Table 19 provides the count of customers enrolled in each LCA and TOU rate type category and the percentage of the rate type category each represents. The TOU rate enrollments are not evenly distributed across the LCAs. For example, the Greater Bay Area represents 50% of program enrollment, but just 21% of TOU-default customers and 63% of TOU-opt-in customers. In contrast, Sierra represents 13% of all enrolled customers, but 24% of TOU-default customers. Note that the Humboldt LCA is not shown due to its low enrollment.

Table 19. Summary of enrollment count and percentage in each LCA by TOU type, as of the last event of the year on 2018-07-26.

LCA	All	All %	non-TOU	non-TOU %	TOU-default-in	TOU default-in %	TOU-opt-in	TOU opt-in %
Other	3,747	3	2,344	3%	437	15%	966	5%
North Coast / North Bay	3,050	3	1,962	2%	275	9%	813	4%
Kern	7,934	7	7,217	8%	167	6%	550	3%
Humboldt	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
Stockton	11,528	11	9,879	11%	421	14%	1,228	7%
Sierra	13,770	13	11,052	13%	724	24%	1,994	11%
Greater Fresno	14,367	13	12,699	15%	302	10%	1,366	7%
Greater Bay Area	53,789	50	41,567	48%	635	21%	11,587	63%
All	108,185	100	86,733	100%	2,961	100%	18,532	100%

Figure 39 visualizes the number of customers in each TOU rate category for each LCA (panel a) and the relative contribution each LCA makes to the total count of customers by TOU category.

Figure 39. (a) TOU enrollment category breakdown for each LCA and (b) Enrollment fraction by LCA within each TOU category

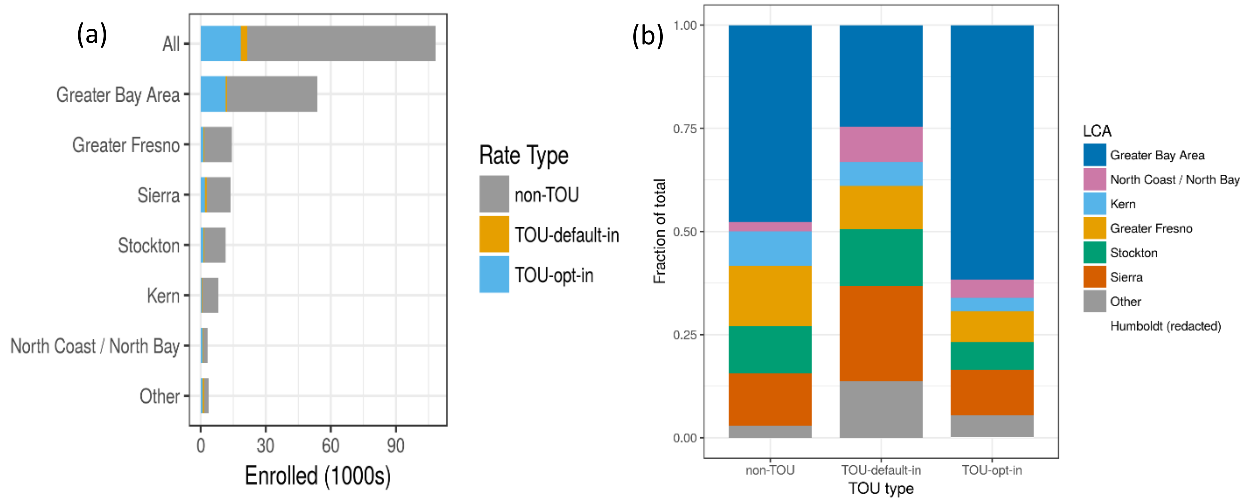

Panel (a) illustrates that the fraction of TOU enrolled customers is small (about 20%) compared to total program enrollment and TOU-default-in enrollment is a small portion of all the TOU rates (about 3% of total enrollment). Panel (b) illustrates that the breakdown by LCA of each TOU rate category varies. This means that part of the difference between their impacts with be attributable to the different weather they experience, rather than the circumstances of their rate enrollment. TOU opt-in customers are quite concentrated in the Greater Bay Area, while TOU-default-in customers are under-represented in the Greater Bay Area and more common in Sierra, Other and North Coast / North Bay LCAs.

While the number of TOU-default-in customers in the program is small overall—just 3% —they are anticipated to be the leading edge of a territory-wide roll out of TOU-default-in rates (scheduled for October 2020) and are therefore of particular strategic interest to program planners and (one assumes) future program evaluators. As a part of the pilot, customers were defaulted in but also presented costs associated with their current rate and the rate they would save the most on – event if it wasn't TOU-C. When a customer uses PG&E's rate comparison tool (and they are eligible), they are automatically presented a cost analysis with and without SmartRate™, with enrollment as easy as selecting their preferred rate. Based on this information, many pilot participants selected a different TOU rate, but still enrolled in SmartRate™.

Figure 40 breaks down the first program year customers currently assigned the different rate types participated in.

Figure 40. SmartRate™ enrollment year by TOU group

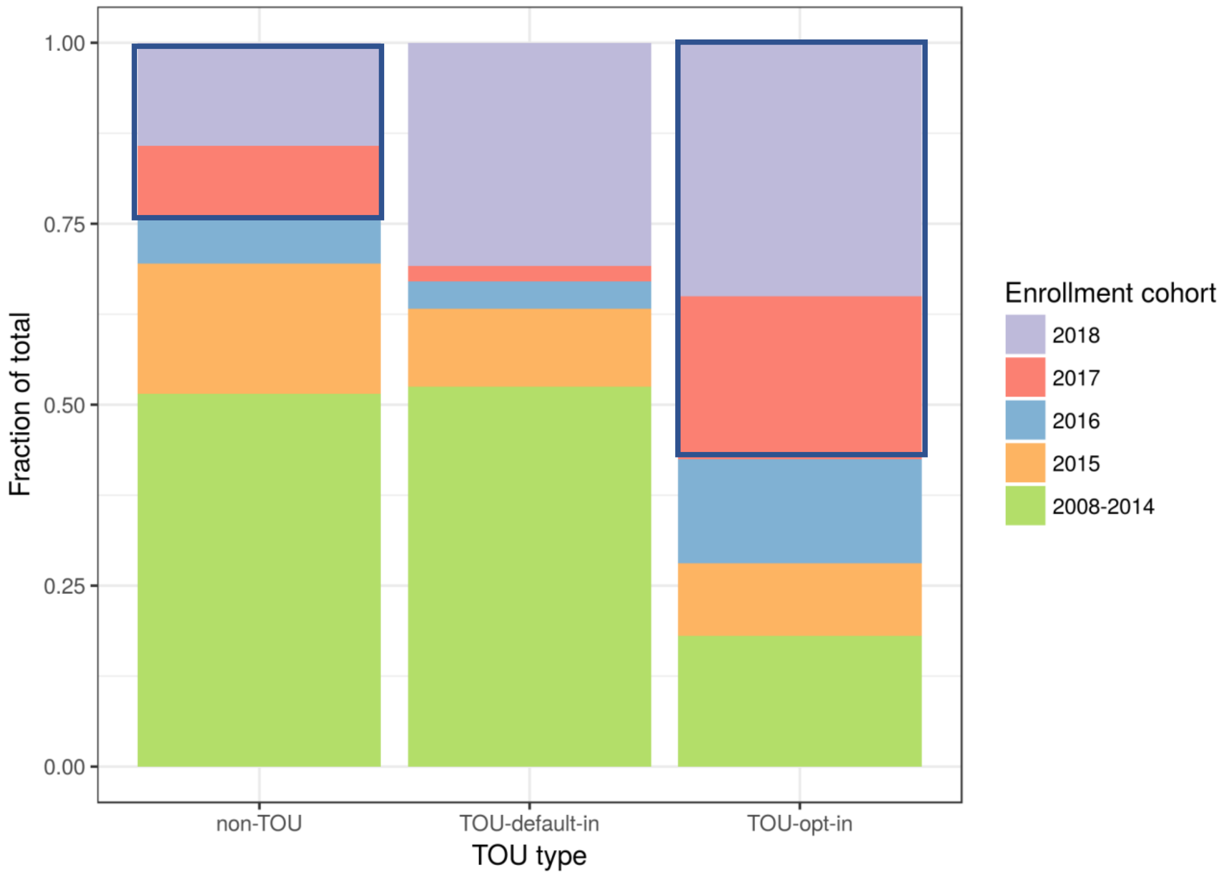

Both types of TOU rates are weighted towards recent enrollment. More than 50% of TOU-opt-in customers enrolled in the past two years, while just 25% of non-TOU customers enrolled during the same period. About 30% of the TOU-default-in customers enrolled in 2018 (i.e. the year of the pilot). The rest were already SmartRate™ customers before the pilot started.

TOU Impacts

Using the same estimation machinery employed to determine impacts by LCA, Dual enrollment status, and CARE enrollment status, we made ex post load impact estimates for each of the three TOU rate type categories for each of the nine event days. Figure 41 presents TOU rate type breakdowns for the typical day aggregate impact (left), average per-customer impact (middle) and impact as a percentage of reference load (right) for each of the nine events from 2018.

Figure 41. Typical event day aggregate, per-customer, and % of reference load shed by TOU rate category.

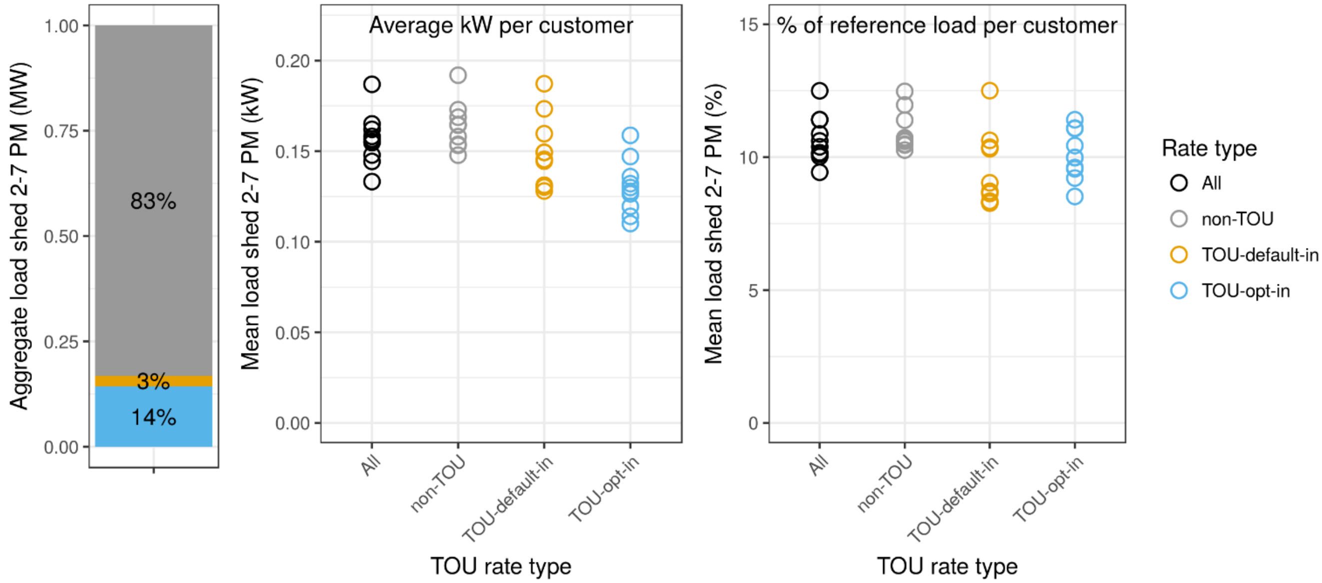

The figure shows that most aggregate program impacts are currently coming from non-TOU customers; however, in the coming years the share of customers on TOU rates is expected to grow quickly. The center and right hand panel show that the range of outcomes across event days overlaps for each rate type, but per-customer load impacts for TOU-default-in are somewhat smaller than non-TOU and larger than TOU-opt-in. Figure 42 plots per-customer reference load vs. load impact (i.e. same y-axis as the middle panel above) for each of the 9 events and each of the rate categories, with a linear fit for each set. From this view, it appears that the load impacts of TOU-opt-in customers should increase fastest with increases in reference loads (i.e. greatest slope of the linear fit), but they have systematically lower reference loads than the other rate categories. The TOU-default-in customers have the flattest load impact response to increasing reference loads, but both the reference loads and the saving vary widely. This could easily be an artifact of their smaller sample size.

Figure 42. Event load shed vs reference load by TOU group

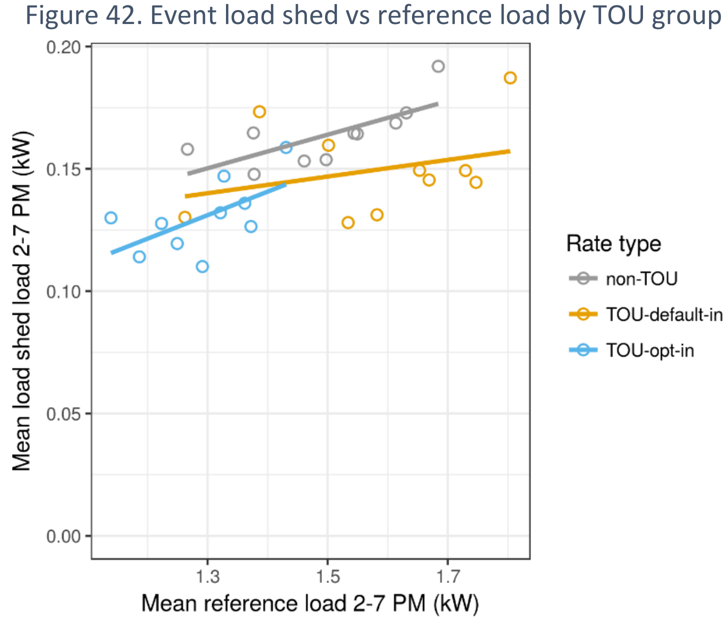

Table 20 summarizes the results for the typical 2018 event day to better quantify the differences observed.

Table 20. Typical event day outcomes for all evaluated TOU rate types

	Category	Enrollment	Temp. (F)	Per-cust. Ref. (kW)	Per-cust. Impact (kW)	Agg. Impact (MW)	Agg. Ref. (MW)	Impact (%)
All	**All customers**	109,972	88	1.51	0.16	17.8	166.0	10.7%
TOU	**non-TOU**	88,094	89	1.54	0.16	14.5	136.0	10.7%
	TOU-default-in	2,924	91	1.65	0.15	0.4	4.8	9.0%
	TOU-opt-in	18,953	84	1.32	0.13	2.5	25.1	10.0%

Due to the geographic effects mentioned earlier, the average event temperature varies by rate type. The TOU-opt-in customers, concentrated in the Greater Bay Area, experienced the coolest temperatures while TOU-default-in customers were the hottest. Per-customer impacts are lower among TOU customers, with savings among the TOU-opt-in customers lagging savings from TOU-default-in customers. However, TOU-default-in customers have the lowest value of load impacts as a percentage of reference loads (Impact %).

It is tempting to speculate about why that might be. To what extent do the geographic patterns in TOU rate enrollment indicate a recognition of opportunities to saving money due to their prevailing load patterns? To what degree is self-selection influencing TOU opt-in savings? Are the new enrollees still learning about the program? To what extent do TOU rate customers re-arrange their consumption away from peak periods, thus reducing their reference loads? Are TOU-default-in customers on the default rate because they are more passive? Is the effect attributable to the timing or cost of the peak period of the TOU-C rate? In 2018, the SmartRate™ event hours were misaligned with the TOU peak hours, which results in the TOU customers having to reduce load for a longer period of time, which can be difficult on hot days.[11] Unfortunately, it is too early in the deployment and there are too few enrolled default-in customers to know for sure. These unresolved questions suggest areas for future inquiry by program planners and evaluators.

TOU Summary – A Look Towards the Future

As described above, there are more opt-ins TOU customers in Greater Bay Area, and generally, Greater Bay Area households contribute lower impacts on an individual household (or meter) level. Given changes with the CCAs—particularly the expectation that even more Greater Bay Area participants will move to CCAs and thus leave the program—the composition of TOU customers within the SmartRate™ program is expected to shift in the future. As such, the make-up of SmartRate™ TOU customers will include a lower proportion of Greater Bay Area households. This could increase the per household savings that PG&E will see from the TOU groups.

Other factors may also change the savings that PG&E sees. If the future event hours are aligned (i.e., the SmartRate™ hours fall into all TOU peak as is proposed for 2020), PG&E would likely see greater savings (due to the reduced burden/reduced number of hours).

Furthermore, it seems likely that the 2018 cohort of TOU-default-in customers arrived at the SmartRate™ program via the web-based rate comparison tool, which automatically runs cost calculations with and without SmartRate™ and recommends the least expensive option (with first year bill protection). It would be logical for

[11] PG&E notes that when hours are aligned and SmartRate™ hours fall into all TOU peak periods, they anticipate seeing a greater load reduction from these customers.

such customers to use the tool upon understanding that they have been enrolled in a new rate and also logical for them to adopt SmartRate™ when the tool recommends it. However, the relatively passive process leading to their enrollment could come with lower attention and below average event responses.

Given the upcoming widespread residential TOU-default-in roll out, it may be necessary to prepare for a large increase in customers enrolling in SmartRate™ and the need to educate them more about the program and their options for responding and offer sufficient motivation for them to participate. On the other hand, assuming administrative costs for the program only increase modestly with enrollment, it might also be rational to embrace a much larger pool of customers even if they have low per-customer impacts.

Conclusions and Recommendations

The 2018 SmartRate™ program delivered 17.83 MW of capacity for a typical event, down from prior year's reduction of 28.1 MW. This change was due to decreases in total participation, the shifting composition of participants, and cooler temperatures.

Ex ante forecasts have also significantly changed since last year. While the factors above all affect the ex ante forecast, one of the largest reasons for the decrease in the ex ante forecast is due to the change in the RA window, which overlaps for only three hours with the event window (and includes two hours of snapback). The PG&E 1-in-2 August peak load reduction forecast is 8.6 MW in 2018 over the new 4-9 pm RA window, down from 13.5 MW in the prior year's ex ante forecast for a 1-6 pm RA window.

Based on our evaluation of the 2018 SmartRate™ program, we provide the following recommendations:

- **Embrace behavioral savings with lower per-customer impacts to account for fewer dually enrolled customers – Attrition of dually enrolled customers will reduce impacts further in the future.** With the December 2018 decision that prohibits new dual enrollment, the number (and likely proportion of) dually enrolled customers will continue to decline. While customers who had been enrolled prior to 10/26/18 are grandfathered in, no new dually enrolled customers will be added. The count and therefore impact of dually enrolled customers will predictably decline through attrition over time. As such, SmartRate™ impacts will be increasingly based on purely behavioral outcomes, which will tend to be lower per-customer. The decline in per-customers savings could be offset somewhat by the recruitment, education, and encouragement of customers most capable of providing behavioral savings. Matching past aggregate impacts without dually enrolled customers will require significantly higher program enrollment.
- **Leverage the transition to TOU rates to boost enrollment; prepare for a lot of novices.** SmartRate™ enrollment takes place through a web-based tool that allows customers to use their own meter data to compare what their costs would be under different rates and rate features, including SmartRate™. The advent of default TOU enrollment will very likely drive significantly more traffic to that tool and may cause a spike in interest and enrollments in SmartRate™. Many of these customers will recognize it as a way to potentially save money but not have much context on how to optimally participate in the program. PG&E will want to prepare to walk this new cohort through their orientation and options in the context of the confusion of their rate being changed. PG&E should also monitor the enrolling cohort carefully since it is unclear which segments will be inclined to sign up for the program, although it is a safe bet that the rate comparison tool recommendations (knowable by PG&E in advance) will shape the process. PG&E may also want to actively recruit segments with higher potential for load reductions; beyond focusing on populations in hotter areas this could be done using individual load patterns and attributes.
- **Target those with higher potential for reductions – Segment-specific results provide insights on the best individual targets**. As the program becomes more purely based on customer behaviors, recruitment targeting should be shaped by analysis of which customers have the technical capacity to shed loads and the tendency to do so. This could be done segment by segment or with metrics unique to each customer. The historical performance of the SmartRate™ only segment of current customers will offer the best insights as to how to proceed here.
- **Shift event times to better align with TOU peak periods and RA window – The RA window has shifted to 4-9 pm, TOU rate peak times have shifted, but events are still 2-7 pm.** PG&E has proposed to change SmartRate™ event hours to 5 - 8 pm. This change, if approved in July 2019, would apply to the 2020

season. Something very much like this change will be needed to ensure that the program continues to be relevant to the CAISO and manageable by customers. The SmartRate™ event window should be harmonized with the now later 4 – 9 pm RA window and the peak periods of emerging TOU rates. Although it will require more work on the back end to send the right messages and tabulate the right costs, it might be useful to make the smart rate event window match the peak period of the TOU rates customers are on (if applicable). Changing the event times will provide larger RA window impacts while also minimizing confusion for TOU customers who are enrolled in SmartRate™.

- **Explore effects among TOU customers – Look for patterns among TOU customers, and in particular those offered the default TOU rate.** TOU rates are set to become the default for residential customers in 2020. This year's report began to explore the effects of TOU rates. Despite the very limited number of defaulted-in TOU customers in the data set, they appear to have higher impacts than their peers who opted into TOU rates. However, the geographic difference between the two groups are striking and confound any easy conclusions. The findings on TOU customers in this report should also be looked at in the context of TOU impact evaluation findings, which are being conducted in parallel with this report. There are several steps that could be undertaken to support learning and evaluation related to TOU rate customers. Notably, the current ex ante model does not explicitly incorporate TOU status and while the forecasts do predict TOU enrollment, the limited spillover modeling done to date has not been rigorous or deep enough to detect subtle price effects.

In addition to the recommendations above, future evaluations may want to consider building time of use rates and enrollment cohorts into future models to better account for behavioral effects. Time of use rates are scheduled to become the default choice in the fall of 2020. Marketing and targeting strategies (and participant restrictions) most likely influenced who participated in each cohort. Our findings show significant structural differences in customer attributes, reference loads, and temperature experienced by different enrollment cohorts. They also indicate that cohorts are likely taking different types of behaviors in response to events. Future impact assessments may want to consider building a cohort analysis into the model to have better have explanatory power because they are so different.

Appendix A: Event by event results: SmartRate™ only and Dually enrolled

SmartRate™ only

Date	Number enrolled	Per-customer ref. load (kW)	Per-customer load impact (kW)	Aggregate ref. load (MW)	Aggregate load impact (MW)	Impact as % of ref.	Average temperature (F)
2018-06-12	77,961	1.19	0.12	92.81	9.08	10	86
2018-06-13	77,980	1.29	0.12	100.49	9.55	10	86
2018-07-09	93,103	1.28	0.10	119.45	9.67	8	88
2018-07-10	92,862	1.39	0.11	128.91	10.50	8	87
2018-07-12	92,366	1.35	0.12	124.61	10.71	9	85
2018-07-17	91,438	1.43	0.12	131.04	10.84	8	87
2018-07-18	91,199	1.49	0.12	135.61	10.66	8	87
2018-07-25	89,977	1.55	0.14	139.68	12.70	9	89
2018-07-26	89,663	1.50	0.13	134.94	11.25	8	85
Typical	91,515	1.43	0.12	130.61	10.90	8	87

Dually enrolled

Date	Number enrolled	Per-customer ref. load (kW)	Per-customer load impact (kW)	Aggregate ref. load (MW)	Aggregate load impact (MW)	Impact as % of ref.	Average temperature (F)
2018-06-12	15,379	1.51	0.33	23.30	5.06	22	93
2018-06-13	15,379	1.66	0.34	25.56	5.31	21	93
2018-07-09	18,433	1.70	0.31	31.26	5.75	18	94
2018-07-10	18,434	1.86	0.35	34.35	6.36	19	93
2018-07-12	18,437	1.79	0.31	32.97	5.72	17	92
2018-07-17	18,440	1.93	0.38	35.51	7.02	20	95
2018-07-18	18,441	2.01	0.41	37.11	7.64	21	95
2018-07-25	18,486	2.09	0.45	38.70	8.29	21	97
2018-07-26	18,522	2.04	0.43	37.70	8.00	21	94
Typical	18,456	1.92	0.38	35.37	6.97	20	94

Appendix B: Ex ante model

The statistical model is a linear regression model whose coefficients are fit to the ex post event data and are used to predict the load impact for future events. For each event, an ex post estimate of load shed is produced for each combination of (LCA, CARE status, dual status, event hour), as described in the section on Ex Post methods. The ex post estimates have varying degrees of statistical uncertainty, and when fitting the model each point is given statistical weight inversely proportional to the standard error of the uncertainty.

The load impact is strongly dependent on outdoor air temperature, since adjusting the thermostat setpoint (or otherwise reducing the use of air conditioning) is one of the principal ways customers participate in SmartRate™ events. Of course, load shed due to reduced air conditioning usage is only possible if air conditioners are on, so we follow standard practice of using "cooling degree-hours" in the model rather than temperature. We used a baseline temperature for Cooling Degree Hours of 70˚ F, so 'cdh70' is equal to zero for outdoor temperature below 70˚ F, and is equal to (Temperature – 70˚ F) for temperatures above 70˚ F.

The regression model includes a large number of terms, so it is cumbersome to write out as a formula. In brief, the predicted mean load shed per customer for a given customer segment, for a given event hour, in a given LCA, is:

Load Shed (kW per customer) = (sum of applicable additive coefficients) + (sum of applicable multiplicative coefficients)/10 * (T – 70˚ F)

We multiplied the multiplicative coefficients by 10 in the table below to avoid having to show more digits; that's why the formula above requires dividing those coefficients by 10.

'Applicable' coefficients are a subset of those tabulated below. See the example immediately below the table. A '**:**' indicates multiplication in this table. For instance, all of the multiplicative coefficients are denoted by 'cdh70' in the table below, to indicate that they multiply cdh70. The dual:Sierra term means multiply the dual indicator variable times the Sierra indicator variable; since these are both either 0 or 1, that coefficient only applies to dual-enrolled customers in the Sierra LCA.

Table 21. Coefficient estimates of the Ex Ante model during event hours.

Parameter	Estimate	Standard Error
first_hours	-0.031	0.003
cool_lca	0.031	0.008
care	0.041	0.009
dual:Other:noncare	0.072	0.010
dual:Sierra	-0.373	0.066
dual:Stockton	-0.701	0.077
cdh70	0.097	0.001
cdh70:dual	0.069	0.003
cdh70:care	-0.070	0.004
cdh70:care:dual	0.062	0.003
cool_lca:cdh70	-0.046	0.005

cdh70:dual:hour1	-0.029	0.003
cdh70:dual:hour5	-0.046	0.003
cdh70:care:dual:Sierra	-0.046	0.006
cdh70:care:dual:Fresno	-0.044	0.006
cdh70:care:dual:Kern	-0.085	0.004
cdh70:Sierra:SRonly	0.015	0.002
cdh70:dual:Sierra	0.168	0.024
cdh70:dual:Stockton	0.301	0.029
cdh70:dual:Fresno	0.017	0.005

The parameter 'first_hours' is an indicator variable (either 0 or 1) indicating whether the data point is from the first two hours of an event. 'cool_lca' indicates whether it is from either Bay Area or Northern Coast. 'care' indicates whether the data point is from customers in the CARE program, 'dual' indicates whether it is from dually-enrolled customers, 'hour1' and 'hour5' indicate whether the data point is from the first or fifth hour of the event, respectively.

To illustrate, consider predicting the load shed for the third hour of an event, for care customers who are not dual-enrolled, in Stockton, when the temperature at that hour is 90 F. Look up all of the coefficient estimates that apply to this customer segment in this hour:

Additive terms:

first_hours: -0.031

CARE: 0.041

Coefficients that multiply cooling degree-hours:

All: 0.097 (the coefficient denoted 'cdh70' applies to all customer segments)

CARE: -0.07

Thus, using the formula for predicted load shed, the estimate for non-dual-enrolled CARE customers in Stockton in event hour 2, with a temperature of 90 F, is:

Load Shed = (-0.031 + 0.041) + (0.097 – 0.07)/10*(90˚ F – 70˚ F) = 0.064 kW per customer.

The model differs from that used in previous years in several important ways:

1. For each of the four combinations of (CARE status, dual-enrollment status) the model fits a different slope in the load impact as a function of temperature. In previous years only the dual status was taken into account in fitting the temperature slope.
2. The present model uses temperature data for each hour of the day, so, for instance, the predicted load impact at 2 PM depends on the temperature at 2 PM while the predicted load impact at 6 PM depends on the temperature at 6 PM. Previous years used a single temperature value that summarized the entire day. This difference becomes important in the evening when temperatures can drop many degrees per hour.

3. The present model has additive terms only for specific LCAs, or specific customer segments within each LCA, whose load impact differs substantially from the rest. In previous years, each customer segment and each LCA had a separate additive term, risking over-fitting.
4. The present model assumes one hour differs from another only through an additive "event hour" term, and that the temperature is different in each hour. In previous years a completely separate set of coefficients was determined for each event hour.

A similar approach was used to model the snapback period immediately following an event. One of the inputs to this model is the predicted load impact in the final hour of the event, from the model described above. That input value, called 'pred_final_hour' in the table below, is multiplied by coefficients that apply in different hours after the event.

Table 22. Coefficient estimates of the ex ante model in the post-event ("snapback") period.

Parameter	Estimate	Standard Error
pred_final_hour:after1	0.087	0.023
pred_final_hour:after2	-0.169	0.024
pred_final_hour:after3	-0.160	0.022
pred_final_hour:after4	-0.113	0.021
pred_final_hour:after5	-0.080	0.019
pred_final_hour:after1:dual	-0.658	0.030
pred_final_hour:after2:dual	-0.389	0.030
pred_final_hour:after3:dual	-0.203	0.027
pred_final_hour:after4:dual	-0.076	0.026
pred_final_hour:after5:dual	-0.044	0.024

For example, to calculate the predicted snapback for dual customers in a given LCA and customer category, in the third hour after an event has ended, first use the 'event model' tabulated above to calculate the predicted load shed in that LCA and customer category in the final hour of the event. Then multiply by the coefficients appropriate to the 3rd hour after the event for duals, and the result is:

Predicted load shed = -0.160*(final hour load shed) -0.203*(final hour load shed)

We also produced a pre-event model but the estimated load shed in pre-event hours is close to zero, and no pre-event ex ante forecasts are relevant to this report, so we do not report the coefficients here.

To generate predictions for the RA period, the event-hours model was used to generate predictions for the event hours of 2-7 PM, and the post-event model was used to generate predictions for hours after 7 PM. From these sets of predictions, the RA hours of 4-9 PM were selected and the mean hourly load shed per customer was calculated.

Model Selection

The model choices described above strike a balance between "underfitting" and "overfitting." A model is "underfit" if there are predictive variables that would improve the predictive accuracy of the model but that are not included in the model. A model is "overfit" if many predictive variables are included that the model fits the observed data very well but performs badly when predicting future data. The model used in previous years has 90 adjustable parameters, and would likely overfit when applied to data from only 9 event days from 2018.

To avoid overfitting while still providing an accurate model, model parameters were included only if they are clearly real effects rather than apparent effects that could be due to random variability – they are 'statistically significant' -- and if their estimated effect is large enough to change the predictions noticeably.

Figure 43. Ex post load impact from three LCAs

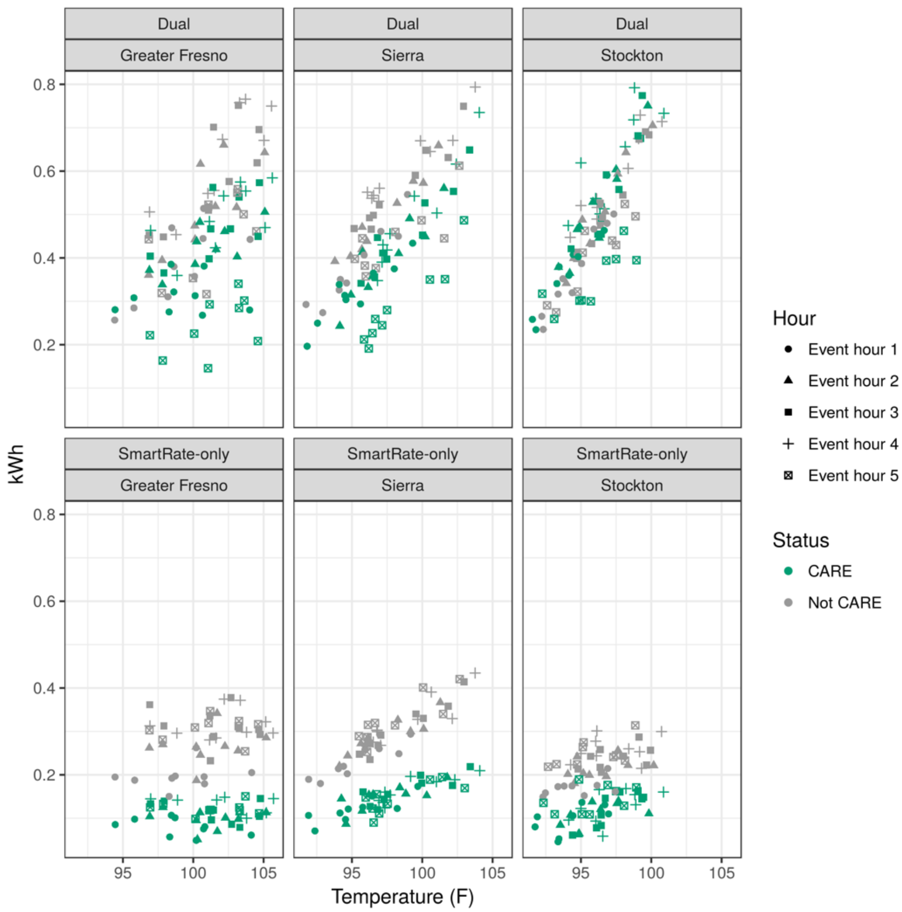

Figure 43 shows ex-post load impact from three LCAs and demonstrates the need for a model that treats customers differently by CARE and dual categories: CARE customers (green) have substantially different load shed at given temperature than non-care customers (gray); Dual-enrolled customers (top row) have a different relationship between temperature and load shed than SmartRate™-only customers (bottom row), and the difference in load shed between CARE and non-CARE customers is different for SmartRate™-only customers than for dual-enrolled customers (that is, the vertical difference between gray and green differs between the top and bottom row of plots). Imagine drawing a best-fit line through the data points for dual-enrolled non-CARE customers in Sierra. That line would not provide good predictions if applied to any other LCA or customer category.

Figure 44: Comparison of ex post load shed to predictions from the ex ante model.

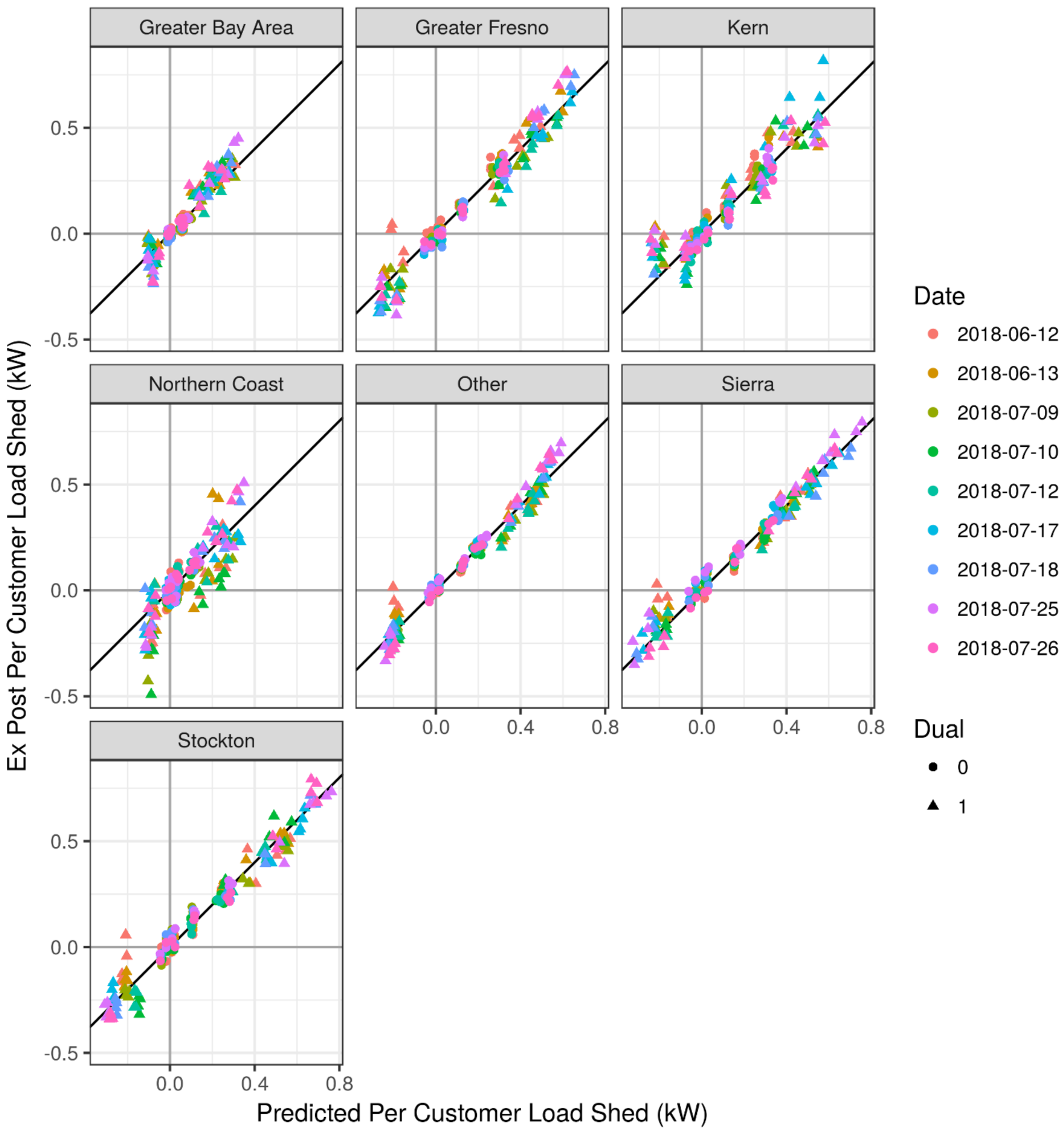

Figure 44 shows the ex post load shed vs the load shed that is predicted by the ex ante model, for the hours of the RA window (4-9 PM), for each event. Values below zero are from the snapback period (7-9 PM).

Bayesian Multi-Level Model

Fitting a linear model in which there is a different temperature coefficient for each combination of (LCA, dual-enrollment status, CARE status, event hour) would lead to overfitting, since there were only nine events, and since some combinations of have rather low numbers of customers – hundreds, not thousands – and thus have ex post load impact estimates that are subject to substantial uncertainty. As discussed above, rather than fitting a parameter for each combination of factors we selected a subset of parameters that are both well-supported by the data and whose values are large enough to be practically significant. The effect of using a subset is to treat some combinations of factors as if they have exactly the same coefficient. For instance, the temperature effect for dual-enrolled customers in hour 1 of the event differs from that for the other hours by the same amount in every LCA. Lumping dual-enrolled customers from all LCAs together when estimating a coefficient is called ‘complete pooling’. In ordinary linear regression the only alternative is ‘no pooling’: dual-enrolled customers in hour 1 of the event could have a coefficient in, say, the Sierra LCA that is fit completely independently from the coefficient for other LCAs.

Complete pooling treats different categories of data as if they are the same. But in fact, it isn't possible for Sierra to have *exactly* the same temperature relationship as the other LCAs. Complete pooling guarantees that predictions for some categories will be too high and for other categories will be too low. But 'no pooling' leads to overfitting, in which both signal and noise influence the predictions.

There is an alternative, called 'partial pooling.' Suppose we somehow knew the correct temperature coefficient for dual-enrolled CARE customers in six of the LCAs, and they all fell within a fairly tight range of values but are not identical. We would know a fair amount about the value of the coefficient in the seventh LCA even before seeing any data from that LCA. Further suppose that LCA has only a small number of dual-enrolled CARE customers, so that the ex post estimates from those customers are subject to substantial uncertainty. Rather than estimate the temperature coefficient for those customers based solely on those noisy data, one could obtain a better estimate by taking a weighted average between the best fit to the noisy data, and the mean estimate from the other LCAs. This is the basic idea (but not the mathematical method) behind a statistical approach called Bayesian Multi-Level Modeling. In a Bayesian Multi-Level Model, parameters (such as temperature coefficients for different categories of customers) are assumed to be drawn from a common distribution. The values of the individual parameters, as well as parameters that describe the distribution, are all estimated from the data. Details can be found in many textbooks, such as "Bayesian Data Analysis", by Gelman, Carlin, Stern, Dunson, Vehtari, and Rubin (CRC Press, 2014).

A Bayesian multi-level model allows many more parameters to be included without causing a problem with over-fitting. We fit a Bayesian model that includes a different temperature coefficient for each combination of (LCA, dual-enrollment status, CARE status), with a modified temperature slope in each LCA for each event hour, and used the results – which include estimated coefficients and uncertainties – but ultimately did not use that model for the ex ante predictions: Switching completely to a Bayesian model for the ex ante predictions would be a substantial change in methods, and would have complicated the comparison of last year's ex ante predictions to this year's. Bayesian multi-level models are also harder to summarize than the ordinary regression model we ultimately used, since they have many more parameters and since the model structure is more complicated. However, we did use the Bayesian model results to help select what parameters to include in the ex ante model.

The Bayesian approach has many advantages and should be considered for future evaluations.

Appendix C: A closer look at the Kern LCA

Within the LCA-specific results, one in particular—Kern—underperformed compared to other LCAs and its own past performance. While Kern had event temperatures ranging from 95 to 104°F, and reference load magnitudes and temperature responses similar to reference loads in other hot LCAs, Kern participants did not respond as expected based on past performance. Figure 45 illustrates Kern's reference loads and event impacts for 2017 and 2018 in comparison to Sierra, which represents the other hot LCAs well and was largely aligned with Kern in 2017. The top row is event reference loads scatter plotted against outside temperatures for both Kern and Sierra, with events from 2017 in gray and 2018 in purple. The bottom row is event impacts scatter plotted against outside temperatures for the same LCAs, but with the performance for dually enrolled (triangles) and SmartRate™ only customers (circles) separated. The results for dually enrolled customers are highlighted by orange ovals for both Kern and Sierra.

Figure 45. Kern and Sierra reference loads (top row) and mean load shed vs temperatures (bottom row) for 2018 and 2017

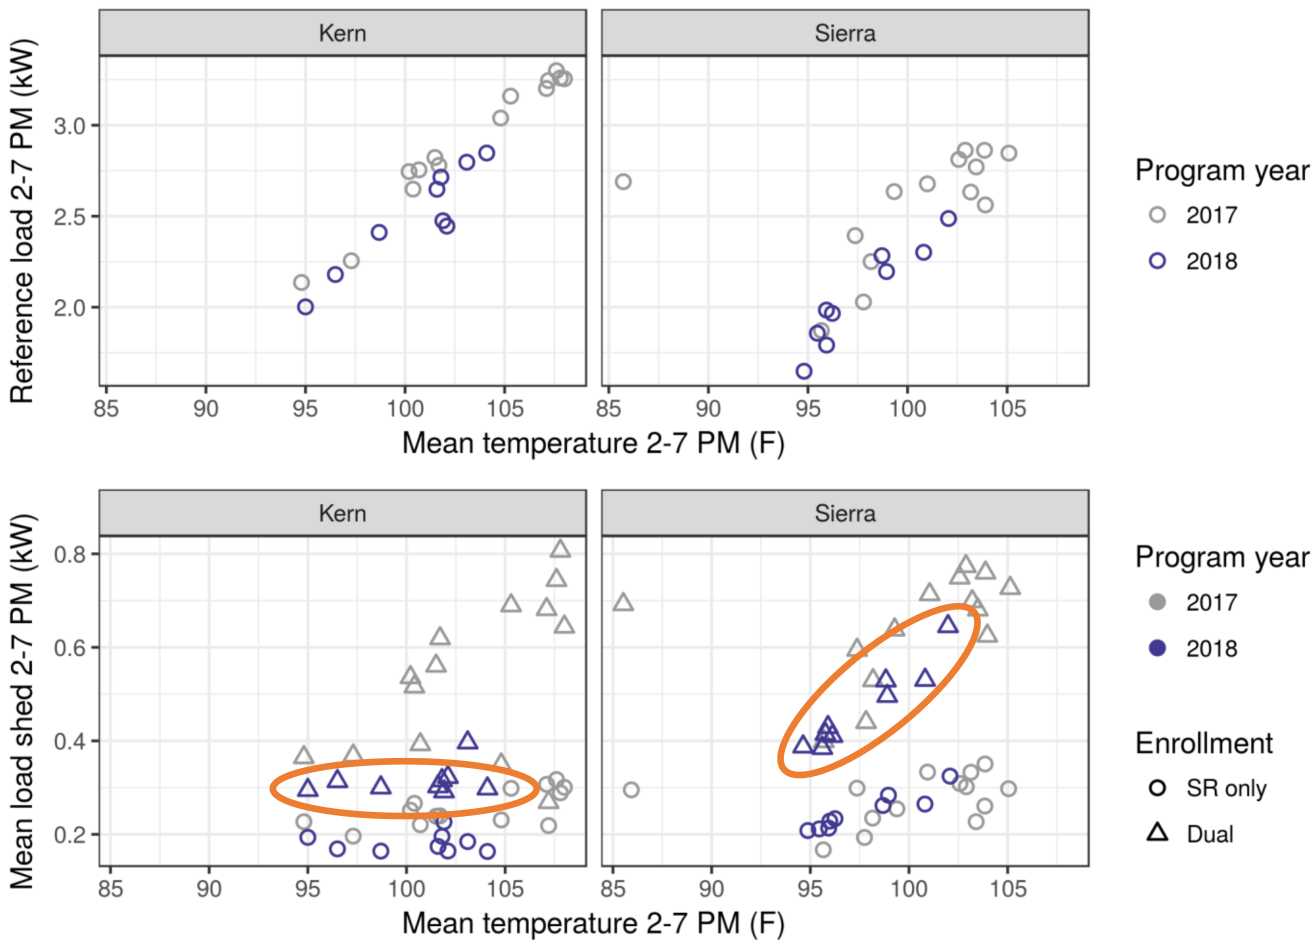

The underperformance of Kern is particularly noticeable among the *dually enrolled* customers, who usually provide the highest impacts. Under normal conditions, event impacts increase with outside temperature (and reference loads), especially for dually enrolled customers. Referencing the figure, it is clear that Kern's dually enrolled customers (highlighted in the oval) were far less temperature responsive than Sierra's (other oval) and less responsive in 2018 than 2017 (see Kern's gray vs. purple triangles). It appears that there was also some underperformance in 2017 (a couple of low gray triangles in Kern). These occurred during the final two events of 2017.

Based on the information available, the most likely explanation for the well below expectation response from Kern's dually enrolled customers is some sort of data, configuration, or signaling issue concentrated in Kern's geography that prevents automated curtailment of AC equipment. As of the time of this writing, the PG&E program team is investigating this issue. These findings also occurred in the context of a 50% increase in the number of participants in Kern (to 6,449 households) compared to 2017. We believe this change is the result of an official remapping of which customers are in which LCAs as well as conventional enrollments and may be contributing to signaling discrepancies because LCAs are input into each SmartAC™ device in the field.

Appendix D: Bill Protection Analysis

PG&E provided billing data with program specific credits and charges for all SmartRate™ participants enrolled as of 2018-05-01. This included billing impacts for the period 2018-05-01 through 2018-11-01.

We analyzed the data to understand the impact on customer bills and which customers qualify for bill protection. Bill protection means that participants with net total costs associated with the program during their first full summer on the program receive a credit at the end of the summer period (i.e., on their November bill). This credit is for the difference between what they would have been charged if they had not been on the SmartRate™ and what they were charged while enrolled in the program. After the first full summer, participants are charged and credited for their electric usage based on SmartRate™ without bill protection.

CDA developed the billing impacts data set using two sources: (1) the bill data with SmartRate™ charges provided by PG&E, and (2) the SmartRate™ customer lists as of 2018-05-01 that had all of the needed flags required for our analysis (i.e., CARE, LCA and flag for enrollment status). After CDA completed the data cleaning steps, there were a total of 107,195 customers for the bill protection analysis. (See Table 23.)

Table 23. Data cleaning steps for bill protection status analysis

Filter type	Customer count
SmartRate™ customers	149,364
SmartRate™ active during 2018 program year	128,700
Billing customers	128,005
Billing and SmartRate™ data merged	128,005
CARE, LCA, and Dual values present	126,147
3+ months of billing May - September 2018	107,195
Working sample	**107,195**

Average Savings and Refunds Across all Customers

Across all 107,195 customers in our analysis, the average participant saved $23.80. Dually enrolled participants were more likely to see a bigger reduction in their bill (and a smaller increase) than SmartRate™ only participants. See Table 24.

Table 24. Bill impacts for all customers. Dollar values are the total for the summer of 2018. Numbers discussed in the text are highlighted.

Enrollment Status	Impact	Count of Participants	% of Participants	Avg. Bill Change
SR only	Increased Bill	4,628	5%	$11.43
	Decreased Bill	83,559	95%	-$33.97
	All	88,187	100%	-$22.54
Dual	Increased Bill	1,374	7%	$9.58
	Decreased Bill	17,634	93%	-$38.80

	All	19,008	100%	-$29.22
All	Increased Bill	6,002	6%	$11.01
	Decreased Bill	101,193	94%	-$34.81
	All	107,195	100%	-$23.80

Bill Protection

Based on the definition of bill protection as applying to any enrolled customer during their first full summer season, we flagged all customers whose enroll date was after 2017-05-01, i.e. after the start of the prior year's summer season, as protected.

Overall 25% of all participants fell into the bill protection status during the 2018 SmartRate™ season. (See Table 25Table 26.) This is higher than the number on bill protection in 2017 (25% compared to 17%, but similar to prior years, almost all of those on bill protection status were *SmartRate™ only* customers.

Table 25. Share of participants with and without bill protection status (i.e., "protected")

Enrollment Status	Protection Status	Count of Participants	% Protected	% Population
SR only	Unprotected	62,204	71%	58%
	Protected	25,983	29%	24%
	All	88,187	100%	82%
Dual	Unprotected	18,555	98%	17%
	Protected	453	2%	0%
	All	19,008	100%	18%
All	Unprotected	80,759	75%	75%
	Protected	26,436	25%	25%
	All	107,195	100%	100%

Of those who were eligible for bill protection (26,436), 96% of bill protected participants experienced bill reductions (shown as a negative bill change in Table 26Table 26). The average *bill reduction total across the summer of 2018* was $34.36, but the small number of dually enrolled customers who experienced a bill reduction saw an even larger reduction in their bill ($55.48 for dually enrolled customers compared to $33.99 for SmartRate™ only customers).

Among the 4% of participants that saw their bill increase, it increased by an average of $10.16. We discuss the average refund to these customers towards the end of this appendix.

Table 26. Bill impacts for protected customers. Dollar values are the total for the summer of 2018. Numbers discussed in the text are highlighted.

Enrollment Status	Impact	Count of Participants	% of Participants	Avg. Bill Change
SR only	Increased Bill	950	4%	$10.18
	Decreased Bill	25,033	96%	-$33.99

	All	25,983	100%	-$23.81
Dual	Increased Bill	15	3%	$8.99
	Decreased Bill	438	97%	-$55.48
	All	453	100%	-$46.50
All	Increased Bill	965	4%	$10.16
	Decreased Bill	25,471	96%	-$34.36
	All	26,436	100%	-$24.20

Participants without Bill Protection

Table 27 presents the proportion of program-enrolled participants who had been enrolled longer than a full summer and therefore did not have bill protection status (80,759). Among this group, 94% experienced a decrease, and the average decrease was roughly the same as those on bill protection ($34.96 compared to $34.36).

Table 27. Bill impacts for unprotected customers. Dollar values are the total for the summer of 2018. Numbers discussed in the text are highlighted.

Enrollment Status	Impact	Count of Participants	% of Participants	Avg. Bill Change
SR only	Increased Bill	3,678	6%	$11.76
	Decreased Bill	58,526	94%	-$33.96
	All	62,204	100%	-$22.21
Dual	Increased Bill	1,359	7%	$9.59
	Decreased Bill	17,196	93%	-$38.38
	All	18,555	100%	-$28.79
All	Increased Bill	5,037	6%	$11.17
	Decreased Bill	75,722	94%	-$34.96
	All	80,759	100%	-$23.79

Bill Protection Refunds

To ensure representative results, the preceding sections examined bill impacts for customers who participated in the program for 3 or more summer months. However, bill protection refunds are paid to all customers. *For the analysis of bill refunds only*, we analyzed all 126,147 customers with billing data available for 2018. Of the participants who were bill protected and participated in any portion of the summer of 2018 (35,320 SmartRate™ only and 537 dually enrolled participants), we calculate that a total of 1,451 or 4% experienced bill increases and were therefore eligible to receive refunds (Table 28).[12] The average refund amount for SmartRate™ only customers was $8.44. The 15 dually enrolled customers with a refund received slightly more ($8.99).

[12] These numbers are based on our interpretation of the eligibility criteria and were not derived from actual rebate payment data.

Table 28. Bill protection refund amounts. Dollar values are the total for the summer of 2018. Numbers discussed in the text are highlighted.

Enrollment Status	Received Refund	Count of Participants	Avg. Refund
SR only	Yes	1,436	$8.44
	No	33,884	-
	Either	35,320	-
Dual	Yes	15	$8.99
	No	522	-
	Either	537	-
Table note: Unlike all other tables presented in this section, this table was derived from data for all customers participating in the 2018 program year, not just those with at least 3 months of participation.			

Bill Impacts by Customer Segment

The tables below present the average effects on customer bills by LCA and for participants who are on CARE.

We calculated the share of participants who experienced bill reductions in each LCA and the average bill reduction among these participants. (See Table 29.) In every LCA, the average bill was reduced, with the largest reductions experienced by customers in Kern and Fresno Area LCAs and the smallest reductions in the Greater Bay Area.

Table 29. SmartRate™ bill impacts by dual enrollment and LCA

Enrollment Status	LCA	% of Population*	Avg. Bill Change
SR only	Stockton	6%	-$25.39
	Sierra	6%	-$29.96
	Other (blank)	16%	-$26.33
	North Coast / North Bay	2%	-$25.87
	Kern	6%	-$44.71
	Humboldt	[redacted]	[redacted]
	Greater Fresno Area	12%	-$37.63
	Greater Bay Area	51%	-$15.61
	All	100%	-$22.54
Dual	Stockton	10%	-$30.29
	Sierra	11%	-$30.18
	Other (blank)	21%	-$25.88
	North Coast / North Bay	2%	-$20.98
	Kern	6%	-$48.40
	Humboldt	[redacted]	[redacted]
	Greater Fresno Area	19%	-$40.76

	Greater Bay Area	31%	-$21.78
	All	100%	-$29.22

*Does not add to 100% due to rounding.

Customers on CARE experience similar (just slightly smaller) reductions in their average bill than non-CARE customers. The average bill change for CARE customers ranged from a reduction of $21.66 to a reduction of $28.42 depending on the category. See Table 30.

Table 30. SmartRate™ bill impacts by dual and CARE enrollment statuses

Enrollment Status	CARE status	Count of Participants	% of Participants	Avg. Bill Change
SR only	Non-CARE	62,556	58%	-$22.91
	CARE	25,631	24%	-$21.66
	All	88,187	82%	-$22.54
Dual	Non-CARE	13,164	12%	-$29.63
	CARE	5,844	5%	-$28.42
	All	19,008	18%	-$29.22
All	Non-CARE	75,720	71%	-$24.20
	CARE	31,475	29%	-$22.95
	All	107,195	100%	-$23.80