2019 Load Impact Evaluation for Pacific Gas & Electric Company's SmartRate™ Program

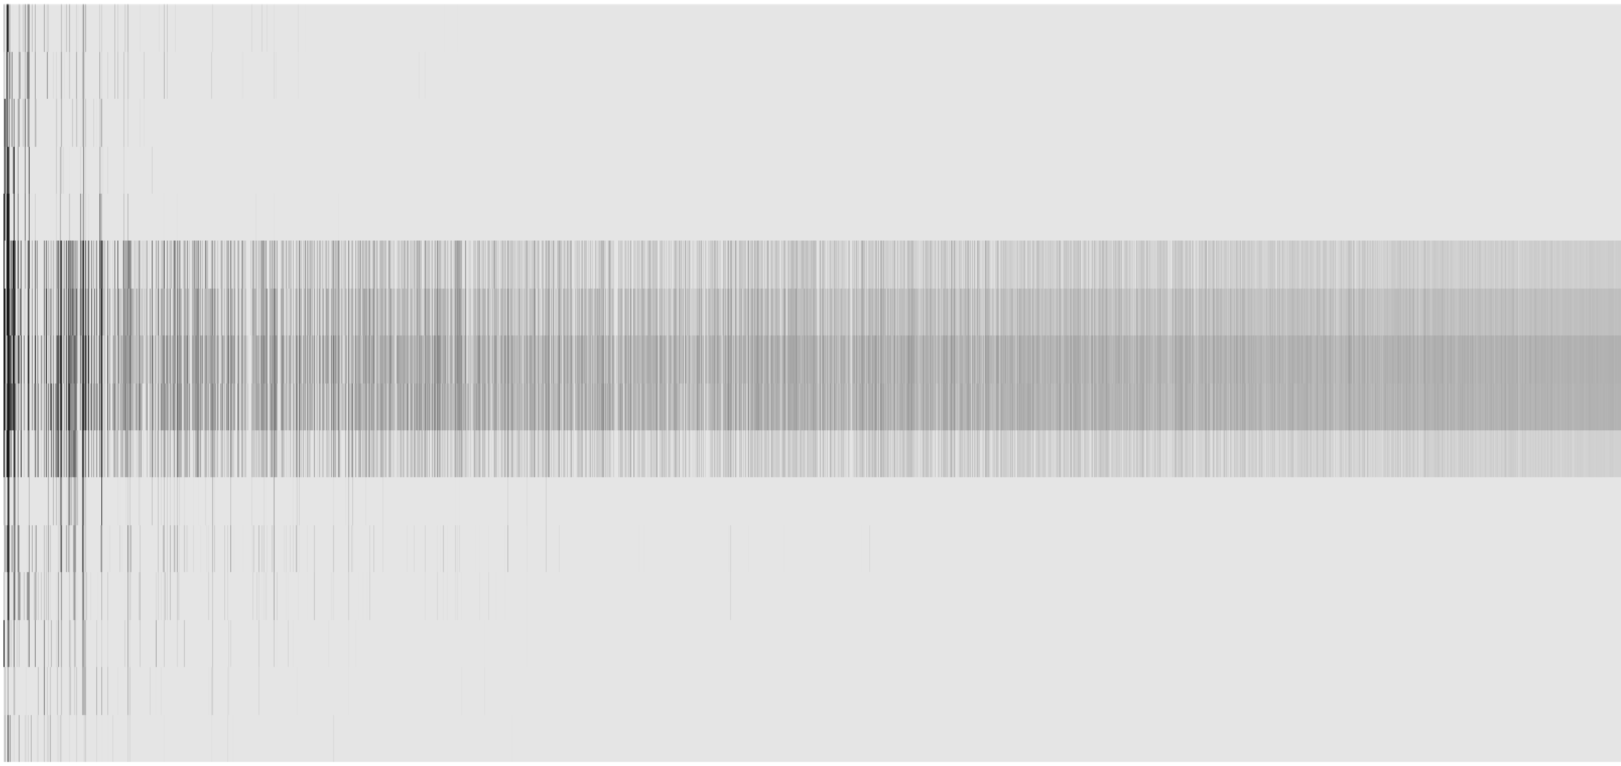

March 25, 2020 FINAL

Prepared for Pacific Gas and Electric Company

CALMAC Study ID PGE0445

Prepared by
Sam Borgeson
Phil Price
Convergence Data Analytics, LLC

Convergence Data Analytics

Stefanie Wayland
Jennifer Mitchell-Jackson
Grounded Analytics, LLC

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 an opt-in demand response (DR) program developed by PG&E based on a rate plan that provides residential customers an economic incentive to reduce or shift consumption during SmartRate™ events. The program uses a rate modifier that applies time varying costs on top of any given underlying rate. Participants receive a credit of approximately $0.024 per kilowatt hour (kWh) for all usage from June 1 to September 30, except for usage between 2 pm and 7 pm on between 9 and 15 SmartDays™ per year, when rates are increased by $0.60/kWh over the underlying rate. For all rates except E-TOU-B time-of-use, participants receive an extra credit of $.0075 for all usage above baseline (Tier 1) from June 1 to September 30. E-TOU-B participants receive an extra credit of $.0050 for all usage from June 1 to September 30.[1] Customers can enroll online from 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 patterns of usage. 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 still reducing overall costs for most participants. Customers are notified of SmartRate™ events by 2 pm the day before 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 in their November bill cycle.

In this document we present:

1. Ex post load impacts for the Smart Rate™ program for program year 2019 (PY2019)
2. Ex ante forecast of SmartRate™ for program years 2020-2030
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 named 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. Customers who do not live in one of the named LCAs are considered part of the "Other" LCA for the purposes of this evaluation. The LCA analysis provides insights on the magnitude of available capacity from events in each geographic area.
- *Dually enrolled customers.* Through the 2018 program year, SmartRate™ customers were allowed to simultaneously enroll in PG&E's SmartAC™ program. These customers are described as "dually enrolled" or as "duals" for short in some figures, 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. New dually participating customers are no longer allowed (per D18-11-029); however, a significant minority of SmartRate™ customers are still dually enrolled and understanding the differences between dually enrolled customers and SmartRate™ only customers allows for more accurate impact assessment and forecasts.

[1] https://www.pge.com/en_US/residential/rate-plans/rate-plan-options/smart-rate-add-on/discover-smart-rate/smart-rate-faq.page (viewed 03/01/2020).
[2] See http://www.caiso.com/informed/Pages/StakeholderProcesses/LocalCapacityRequirementsProcess.aspx for more details on the CAISO local capacity requirements process.

- *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, DR, and pricing products like SmartRate™. CARE customers will, on 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.* A growing minority of PG&E customers is on a time of use (TOU) rate. At the beginning of 2020, 23% of SmartRate customers were on a TOU rate (21% without E6), but TOU rates are slated to become the default choice for customers. This is currently expected to occur in Q4 of 2020. Since TOU rates also expose customers to costs that vary by hour of day, it has become important to empirically quantify SmartRate™ impacts on customers enrolled in TOU rates. In anticipation of the transition to default TOU rates, the PY2018 and PY2019 evaluations include estimates of program impacts within sub-groups of TOU customers (i.e., default or opt-in TOU).

The first three categories (LCAs, dually enrolled customers, and CARE) provide data points to inform our ex ante model, while the latter (TOU customers) provides 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.

In this report, we use a 4 pm – 9 pm RA window to create the ex ante forecasts. While past reports (PY2017 and earlier) 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 2020-2030 are presented for a 4 pm – 9 pm RA window unless a specific comparison is being made.

The event period for PY2019 (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 3 pm – 8 pm or 4 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. CPUC-approved changes to the SmartRate™ will update event hours to 5-8 pm to align within the high price periods of current and upcoming TOU rates.[3] These SmartRate™ event hour changes are currently slated for the 2021 season.

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 ex ante forecast.

Ex Post Load Impacts

PY2019 Aggregate Impacts: The aggregate ex post program load impact for a typical day in 2019 was 14.89 MW. Figure 1 presents aggregate impacts (stacked by LCA impacts) for each event day in PY2019, as well as for a typical event day. As shown in the figure below, there were nine event days in 2019. Aggregate impacts were at their highest on the first event of the season (June 11, 2019) and lowest during the final event of the season (September 13, 2019). Events are dispatched based on high temperatures and other factors that included Public Safety Power Shutoff (PSPS) activity for the first time this year.[4]

[3] There will be a new TOU-D rate and those peak hours will be 5-8 pm.

[4] SmartRate events tend not to be called on the same day as a PSPS event to minimize demands on the customer service center, web, and meteorology teams – and to avoid additional messaging to impacted customers.

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

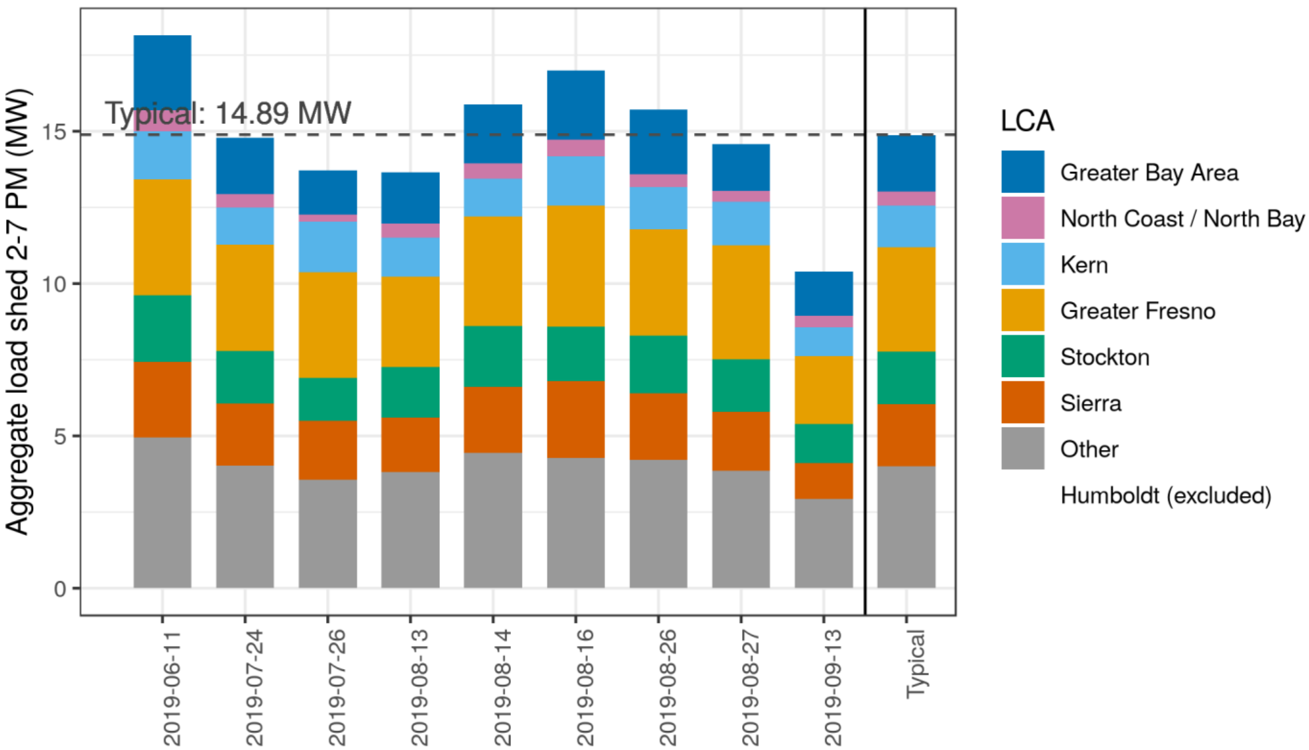

Comparison to PY2018 Aggregate Impacts: Aggregate ex post results were significantly lower than in prior years, primarily due to lower overall participation numbers on event days. Approximately 67,000 customers participated in each of the nine SmartRate™ events that occurred in 2019 (down from approximately 110,000 in PY2018). Most of the net decrease in enrollment comes from customers being automatically de-enrolled when they begin service under local Community Choice Aggregators (CCAs). Despite the fact that the typical event day cooling degree hours (CDH relative to 65˚F) was higher this year, meaning more air conditioning and therefore higher reference loads and responses for event participants, the roughly 40% reduction in participation brought the average aggregate impact down roughly 16% (from 17.83 MW to 14.89 MW).

PY2019 Per-Customer Impacts: The average per-customer impact in 2019 was 0.22 kW with 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 2019 (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 2019

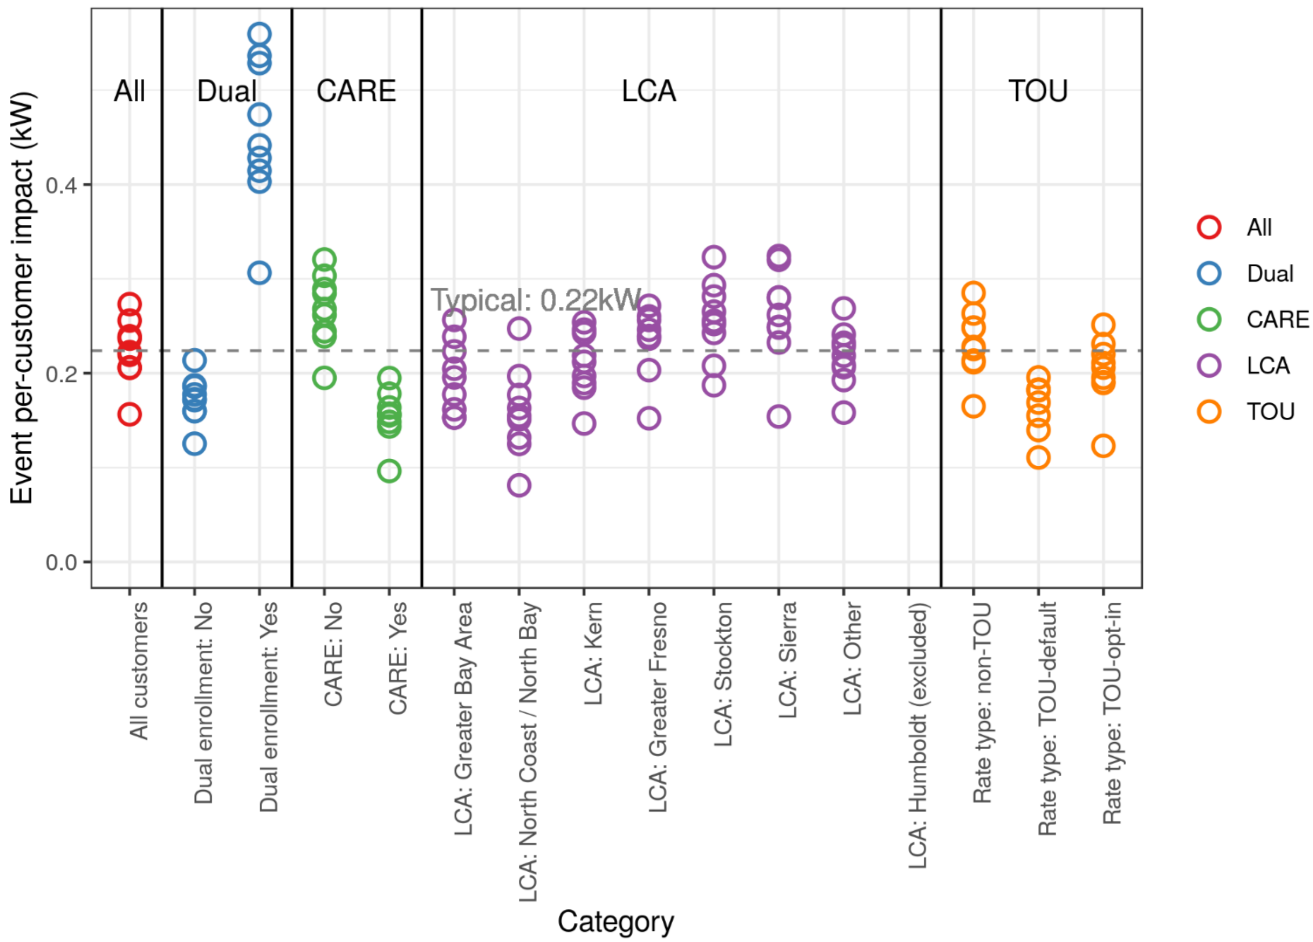

As shown in the figure:

- **Dual enrollment continues to provide the largest impacts per customer**. While 2019 saw significantly fewer dually-enrolled customers (due to de-enrollment of existing customers and the fact that new customers can no longer be dually enrolled), they still offer 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 with larger reference loads than SmartRate™ only customers.
- **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.
- **LCA continues to be a source of variability in load curtailment**. The cooler coastal LCAs such as the Greater Bay Area and the North Coast tend to offer lower per customer impacts, and the hotter inland LCAs such as Sierra and Stockton offer higher per customer impacts; however, there was less variability across LCAs in PY2019 than in PY2018, most likely due to the shift in the geography of program enrollment caused by CCA de-enrollments.

- **Customers not on TOU rates shed more load than those on TOU, with opt-in TOU customers shedding more than those who defaulted in.** TOU opt-in customers made up 18% of the enrolled population in PY2019; default TOU customers made up just under 5% of the population.

Comparison to PY2018 Per-Customer Impacts: The 2019 average impacts are higher than in PY2018 due to weather and CCA de-enrollments. The CCAs that came online between PY2081 and PY2019 were largely in the cooler portions of the Greater Bay Area. The resulting de-enrollments left a greater proportion of enrolled customers in warmer climates. Also, PY2019 was warmer than PY2018. The changes in enrolled customers and warmer weather overall resulted in an increase in cooling degree-hours (CDHs) experienced by enrolled customers. The de-enrollment also caused the proportion of dually enrolled customers increased slightly while their per customer savings increased significantly. On average, dually enrolled customers had substantially higher load impacts than SmartRate™ only customers (0.46 kW vs. 0.17 kW for a typical day, respectively – up from 0.38 kW and 0.12 kW in 2018).

Ex Ante Load Impacts

Ex Ante load impacts are forecasts of what would happen in a future event that takes place in standardized weather conditions. CDA forecasts an August peak aggregate impact of 8.2 MW during an average hour within the 4-9 pm resource adequacy window[5] using the 1-in-2 PG&E weather scenario, and an aggregate impact of 9.1 MW using a 1-in-10 PG&E weather scenario. (See Figure 3.) These estimates are down from prior years.

[5] Note that the timing of the SmartRate™ window is 2 to 7 pm while the RA window is 4 to 9 pm, so the reported mean aggregate RA impacts are the result of averaging 4 to 7 pm impacts with near-zero or even negative impacts after events from 7 to 9 pm.

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

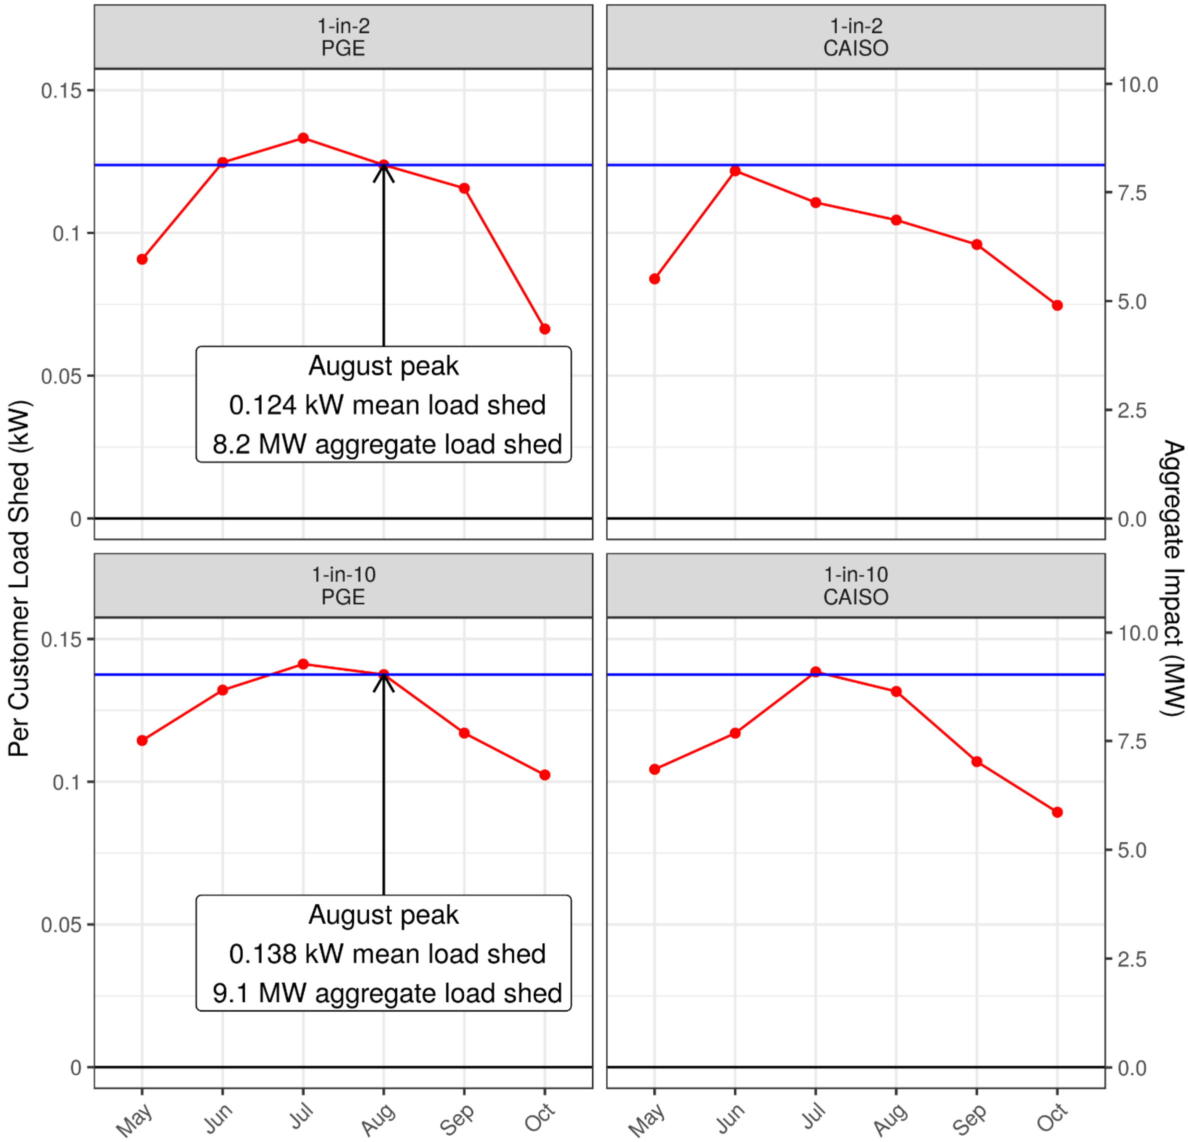

Figure note: The left and right axes apply to all four figures. Mean load shed is indicated on the left, and aggregate on the right, as demonstrated by the August examples in PG&E 1-in-2 and 1-in-10 and the corresponding blue line leading left and right.

Aggregate ex ante estimates for 2020 and beyond are 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 8.6 MW, compared to 8.2 MW for the same scenario in this year's estimate. The PY2018 forecast anticipated the CCA de-enrollments fairly accurately, so the factor that led to the largest reduction compared to prior ex ante estimates was the change in the number of dually enrolled customers who are forecast to be in the program in August 2020: The forecast from the PY2018 analysis had 14,826 customers in this group, compared to 11,885 for the program year 2019 analysis.

Note that the ex ante estimates use an RA window of 4-9 PM, even though the events take place from 2-7 pm. Decision 18-06-030 modified the resource adequacy measurement hours to HE17-HE21 (4-9 pm) for each month of the year—a change from the summer RA window of 1-6 pm that was in place when the current SmartRate™ event window was set. The event period overlaps with this RA window for only three event hours, rather than the four hours from earlier years (2017 and earlier). 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).

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 2019 SmartRate™ program, we provide the following recommendations:

Recommendation	Description
Plan the future of SmartRate in accordance with projections for new CCAs	When a CCA comes online in a customer's location, they are enrolled in the CCA by default and de-enrolled from SmartRate as a result. Those de-enrollments have reduced the customer base of SmartRate from 124k customers in 2017 to 66k customers in 2019. Where CCAs are not triggering de-enrollments, the number of customers enrolled in SmartRate is actually increasing each year, so the future of the SmartRate resource is likely to be determined by future CCA activity. At some point, if enrollment continues its dramatic decline, it will be necessary to determine whether and how the program can continue operating.
Try to mitigate the resource impacts of CCA de-enrollments	If customers de-enrolled through their transition to CCAs are not enrolled in a similar critical peak pricing program under their CCA, these self-identified customers willing to contribute to grid stability are being lost as a grid resource at a time of increasing value for load flexibility. This policy gap with unintended consequences for the DR resource should be addressed.
Embrace behavioral savings from households with lower per-customer impacts to account for fewer dually enrolled customers and target those with higher potential for savings	With the discontinuation of dual enrollment between SmartAC and SmartRate, SmartRate's event impacts will shift towards the largely behavioral impact of the group of customers currently classified as "SmartRate only". The ultimate value of a DR program lies in its aggregate impacts. To preserve and grow its aggregate impacts, SmartRate may need to pursue increased enrollments to offset smaller per-customer impacts and those enrollments may need to be targeted at customers likely to have higher potential for savings.
Leverage the transition to TOU rates to boost enrollment; prepare for a lot of novices	The only mechanism of outreach to enroll SmartRate customers is PG&E's rate comparison tool. On its own and without marketing expenditures, the rate comparison tool recruited more than enough customers between the end of PY2018 and PY2019 in every LCA other than the Greater Bay Area (i.e. without CCA de-enrollments) to produce a net increase in enrollment. It stands to reason that the advent of default TOU rates will drive traffic to the rate comparison tool and that surge in interest represents an opportunity to grow and re-configure the enrolled customer base of the program.

Recommendation	Description
Shift event times to better align with TOU peak periods and the RA window	Because SmartRate is implemented through rate structures, its details need to be hashed out and approved via the rate approval process. At the moment, the timing of the program event window is out of sync with both TOU rates and the current RA window. An update to the timing has been approved by the CPUC, but has not yet taken effect. The ex ante RA impacts are based on an average of observed impacts across the RA window, so the update to the SmartRate event window will substantially increase the RA average and bring the program incentives into alignment with current grid flexibility needs.
Prepare to evaluate impacts among TOU customers and account for rate type in choosing potential controls	With default TOU rates on the horizon, it will become critical to the evaluation of the SmartRate program that the ambient impacts of TOU rates are separable from the impacts of the peak rates on SmartDays. The synthetic control approach to event evaluation should be capable of differentiating those impacts, but perhaps only if TOU rate status is taking into account in the sampling and matching criteria.

Introduction to the 2019 SmartRate™ Program

SmartRate™ is a voluntary critical peak pricing program that overlays a standard electric rate. SmartRate™ is a rate-based tool designed to lower electric bills for customers that opt-in 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 2019, PG&E called 9 events, and these occurred in June through September. Events are dispatched based on high temperatures and other factors, including Public Safety Power Shutoff (PSPS) activity – to minimize demands on the customer service center, web, and meteorology teams – and to avoid unnecessary communications with impacted customers.

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, and TOU customers.

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 contains customers from a range of locations within 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 were allowed to simultaneously enroll in PG&E's SmartAC™ program.[6] 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.

[6] 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 the TOU-C rate is slated to become the default choice for customers in Q4 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 2019 relevant to SmartRate™. The TOU-A rate stopped enrolling new customers on the first day of 2020 and will be discontinued after 9/30/2020. The TOU-B rate will stop enrolling new customers sometime in 2020 and the TOU-C rate will become the default rate for customers around q4 of 2020. Around 23% of all SmartRate™ participants in 2019 were on a TOU rate (this includes both opt-in and default TOU customers).

Participation Trends

The total enrollment in SmartRate™ following the 2019 program season was 66,528 (roughly equivalent to 66,448 on a typical event day) –representing a substantial decline of nearly 40% from the end of the 2018 program year. Enrollment numbers, and thus the make-up of event participants, are significantly changing as a result of new Community Choice Aggregators (CCAs) beginning operations or expanding. By default, all utility customers in the service territory of a CCA are enrolled in the CCA's retail service when they start operating.[7] CCA customers are no longer eligible for SmartRate™ and are automatically de-enrolled. This was the largest reason for "unenrolls" in both 2018 and 2019. Areas not impacted by de-enrollments due to CCAs saw modest increases in total enrollment driven by PG&E's web-based rate comparison tool. Due to the net change in enrollment, dual enrollment declined by approximately 34% (yet increased just slightly as a percentage of the population from just under 17% to roughly 18.5%), CARE enrollment declined by 25% (while increasing their proportion of total participants from 29% of the population to nearly 38% of a smaller population).

Over the past three years, numerous SmartRate™ participants have been unenrolled as the CCAs expand.[8] The 2019 results reflect an additional 23,000 de-enrollments that occurred in 2019.

- **2017:** 22,000 SmartRate™ participants de-enrolled due to CCAs.
- **2018:** 34,000 SmartRate™ participants unenrolled due to CCAs. About half of these impacted the 2018 participation numbers and another 17,000 de-enrolled after the PY2018 season when East Bay Community Energy started up.
- **Early 2019:** In February 2019, when San Jose Community Energy launched, an additional 12,000 de-enrolled; and in April 2019, an estimated 11,000 customers were de-enrolled with as the Clean Power

[7] These customers are given an opportunity to opt-out of the CCA.
[8] This data was provided to the evaluation team by PG&E.

San Francisco completed its residential expansion. As such, the program was down 23,000 customers in calendar year 2019. This significantly reduced the number of dual customers, while also leading to fewer TOU customers in the program. Overall, the program was down approximately 44,000 customers at the end of the 2019 season compared to the end of the 2018 season.

While 2019 participation numbers continued to decrease because of the CCAs, over 10,000 new SmartRate™ customers enrolled in 2019. (See Figure 4 below.) Note that there was no marketing of the program in PY2019 but enrollment outside of areas with CCAs increased even as overall enrollment declined.

Figure 4: End of season PY2019 SmartRate™ participants by their year of enrollment (this reflects only participants that were still in the program during the 2019 season – not all participation over time)

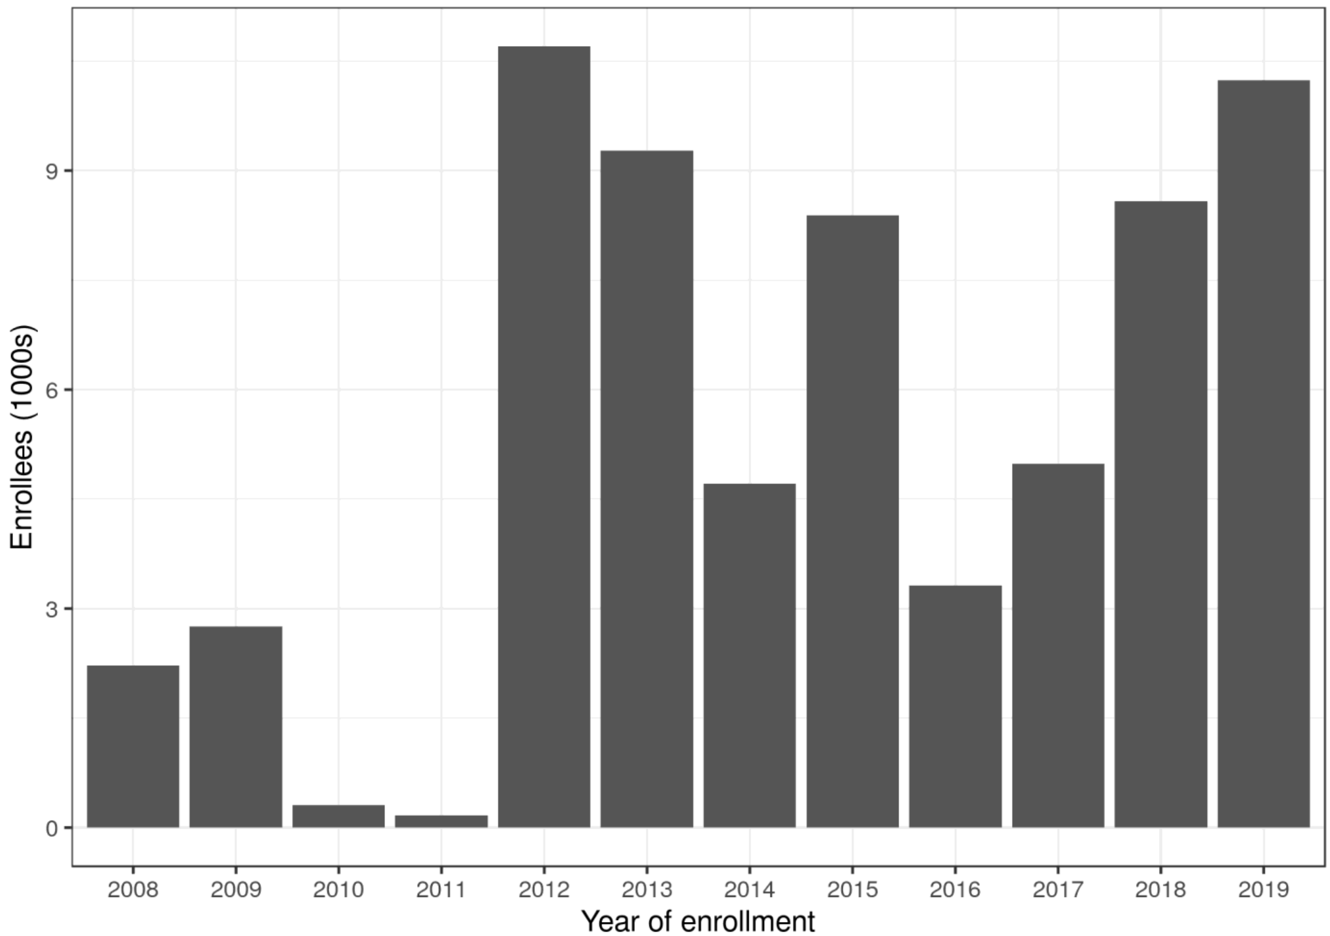

The breakdown by LCA and dually enrolled vs. SmartRate™ only for a typical event day in 2019 compared to 2018 is shown in Table 1, with the only LCA with a net *decrease* (Greater Bay Area) shown at the top. Largely due to de-enrollment of new CCA customers, the Greater Bay Area participation dropped 83% between 2018 and 2019, while participation in most other LCAs remained about the same. Overall, the program saw a roughly 40% drop in participation (primarily due to the CCAs in the Greater Bay Area).

Table 1. PY2019 Participants in the SmartRate™ Program

LCA	Dually enrolled	SmartRate™ only	Total	Change from typical 2018
Greater Bay Area	2,031	7,451	9,482	83% Decrease
Greater Fresno Area	2,369	12,267	14,635	2% Increase
Sierra	1,711	6,049	7,760	3% Increase
Stockton	1,667	5,111	6,778	1% Increase
Kern	705	5,818	6,523	1% Increase
North Coast and North Bay	307	2,446	2,816	11% Increase
Humboldt	-	-	-	-
Other (not associated with an LCA)	3,379	15,075	18,454	1% Increase
Grand Total	**12,169**	**54,217**	**66,448**	**~40% Decrease**

Table note: These are the numbers are on a typical event day. 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. Aggregate results exclude Humboldt due to confidentiality and too few customers to materially influence the average or aggregate event results.

Figure 5 displays event participation by event for past years (2017-2018 in faded color) and 2019. The off season CCA reductions create the main discontinuity between years.

Figure 5. Participation by LCA 2017 - 2019

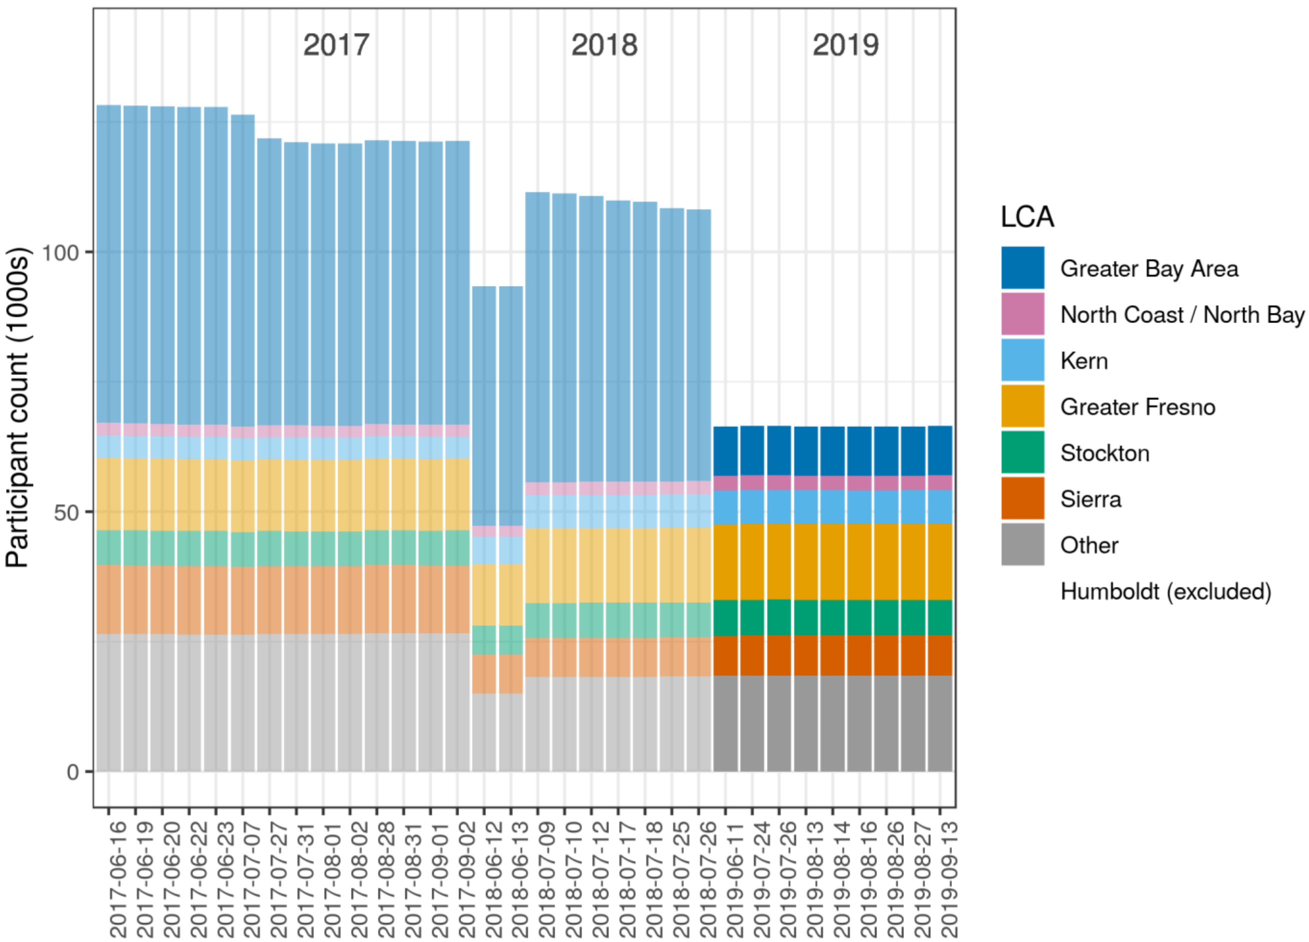

Temperature Trends – Higher temperatures in 2019 than in prior year

On average, event temperatures were up in 2019 compared to the events in 2018. Although the summer of 2019 was not as hot as the summer of 2017, temperatures experienced by typical event participants in 2019 were comparable to those experienced in 2017. (See Figure 6.) This is because a greater fraction of participants are now found in warmer climates - de-enrollments from CCAs in the Greater Bay Area remove customers from one of the coolest LCAs.

Figure 6. Comparison of cooling degree hours (CDH) 2017 - 2019

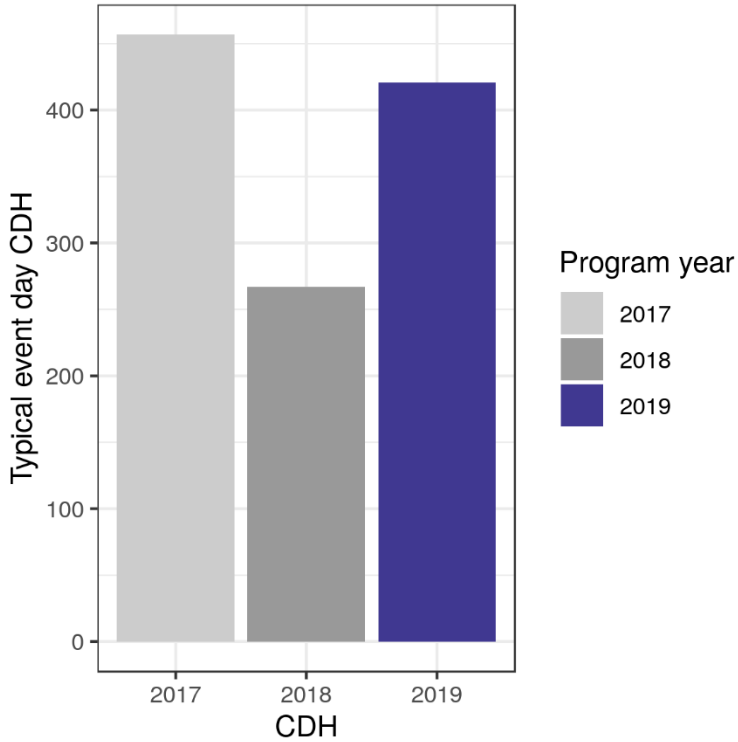

Programmatic Changes and Important Operational Details in 2019

There were no major programmatic or operational changes between 2018 and 2019. 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, 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.

The California Public Utilities Commission (CPUC) did determine that no new customers will be allowed to be dually enroll in this program and SmartAC™, but existing ones (enrolled prior to October 26, 2018) can remain in the program.

It is also worth calling out that when SmartRate™ events are called, the AC loads of dually enrolled customers are still controlled during the SmartRate™ hours and evaluated as part of the ex post event performance of SmartRate™ customers.[9]

We note that the program hours will shift from 2-7 pm to 5-8 pm by 2021.

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

Key Research Questions and Study Methods

Key Research Questions

The research:

1. Estimates the ex post load impacts for the SmartRate™ program for PY2019
2. Estimates the ex ante load impacts of SmartRate™ for program years 2020-2030
3. Looks at dual-enrollment-, CARE-, and LCA-specific effects
4. Also examines effects by TOU rate type

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 60 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.
- **Changing mix of customers:** the de-enrollment of customers due to CCAs were not randomly distributed across the population. Rather, they took place primarily in the Greater Bay Area. This resulted in shifts in the typical weather experienced by enrolled customers as well as changes in the fraction of dually enrolled customers in those covered by CARE rates.

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 rate types.

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. In PY2019 there were 15 days with Public Safety Power Shutoffs that precluded calling DR events and whose unusual grid activity make them unsuitable for use as non-event days as well. Figure 7 shows the event day, non-event day, and PSPS day temperature profiles from those stations for the full summer of 2019. The event day profiles are blue, and the non-event day profiles are gray and the PSPS profiles are red. It can be verified that the event days are among the hottest days of the summer, but there are some 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 for event days, non-event days, and PSPS days.

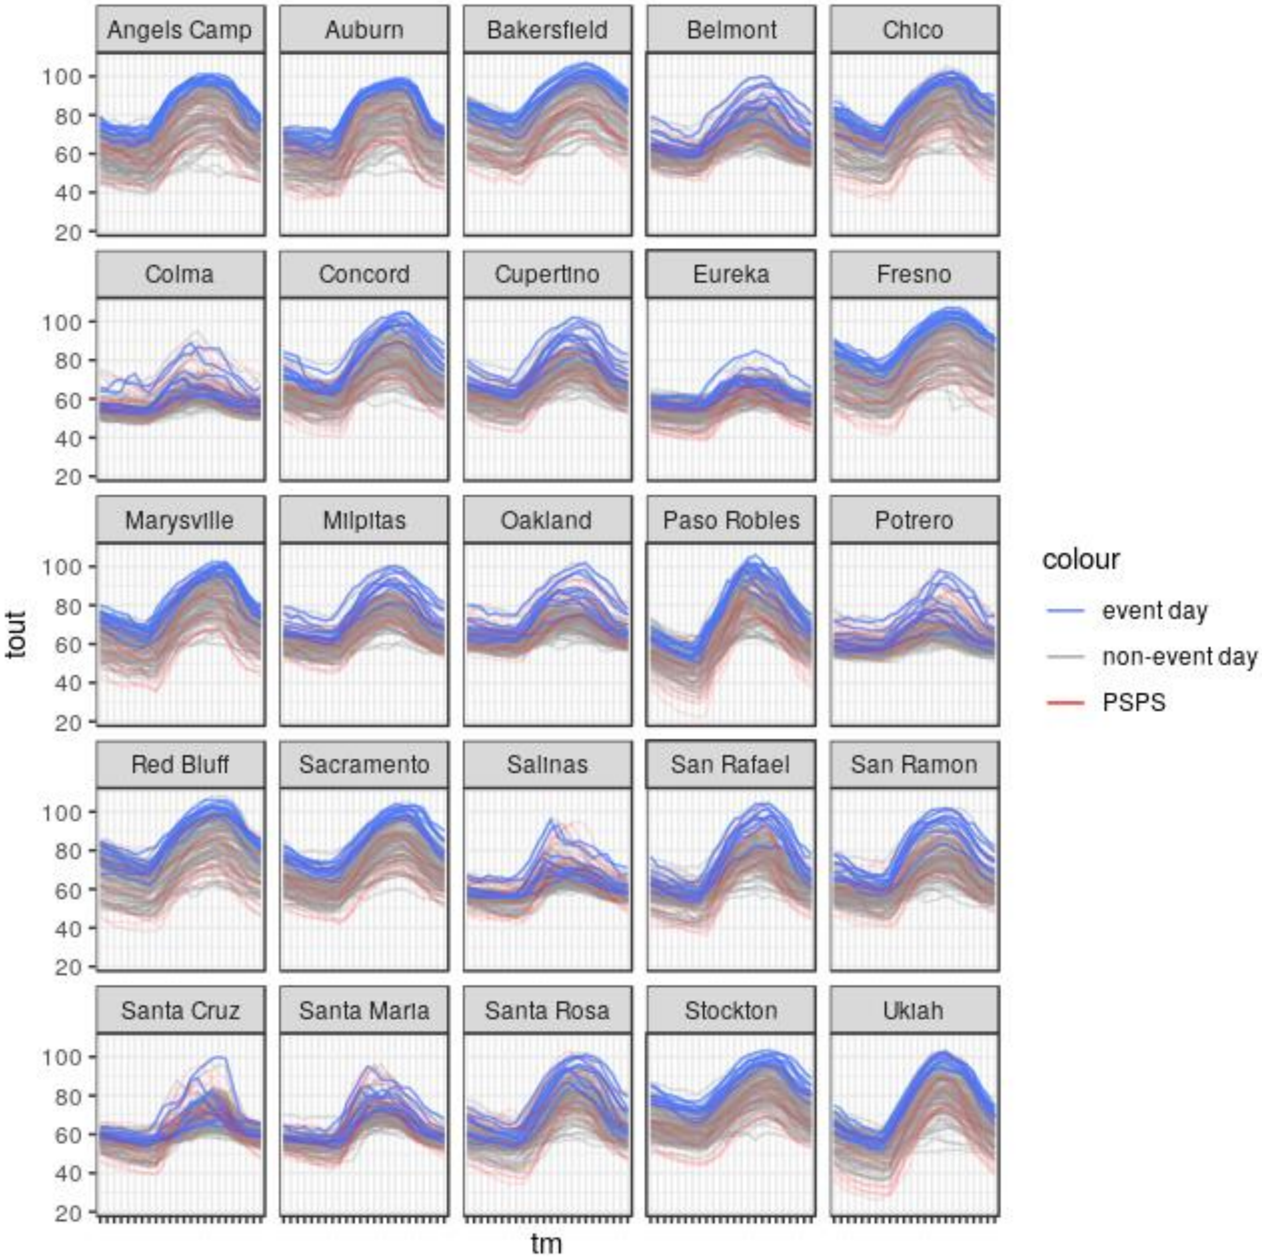

Comparison days were selected based on three criteria:

(1) Days whose average maximum temperature across 'Red Bluff', 'Sacramento', 'Fresno', 'Concord', and 'San Jose' were above 98F were selected based on that “trigger rule”. There were 8 such days.
(2) The enrollment-weighted average temperature profiles for each of the 9 event days were matched to the closest non-event day profiles (by Euclidean distance across all 24 hours). There were 5 such days.
(3) Day comparisons were made by computing the Euclidean distance between 24-hour temperature profiles for each SmartRate™ event day in the 2019 season and the available non-event days (excluding SmartAC™ events as well) within each weather station. There were 22 days that matched for three or

more weather station. Many dates overlapped between these methods and the unique list of dates across these three methods spanned 24 days, which were adopted as our comparison days.

There were just 9 event days, but 24 comparison days (see Figure 8). The reason for this generous inclusion of comparison days was to improve statistical support for the difference-in-difference baseline calculations.

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

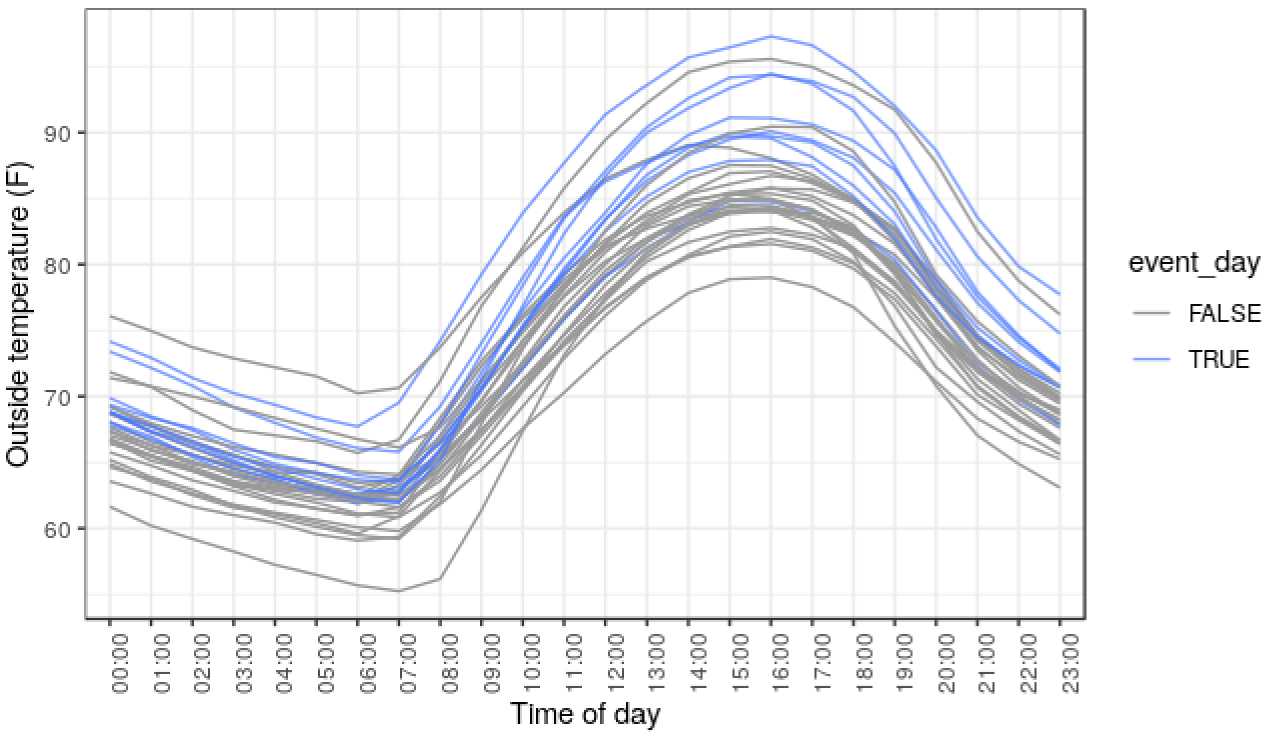

Identify potential controls

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 2019 billing and account data for 790,000 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 2019-04-01 and 2019-10-01 and kept the 10 closest potential controls for each participant. We also added 10,000 random NEM customers to our request because matching alone didn't provide a large enough sample of that relatively rare customer type.

The result was a request for meter data for a list of just under 360K 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 balanced 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, and TOU-type. 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 calculated by a statistical model that is run for each of the 24 hours 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 this regression model we are effectively implementing a difference-in-difference calculation that can also adjust for differences in temperature response between cases and controls, 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; this corresponds to the difference in consumption between event days and non-event days on the hour in question, for cases, after adjusting for differences in outdoor temperature.

In some LCAs and for some subsets of customers, the number of cooling-degree-hours can be 0 for some or most of the event days. If that occurs, the coefficients of the cdh terms in the equation above will be subject to very large uncertainty. For this reason, we require at least 5 event days with non-zero values of cdh in order to use the model given above; if that condition does not hold then the cdh terms are left out of the model.

"Reference load" is the load the participants would have experienced on an event day if there were no event. This is predicted, using the same model used to calculate load shed, as follows:

1. Calculate cdh, case_cdh, and early_aft_kWh for the cases on event day. (For the cases, case_cdh is equal to cdh, so those two variables have identical values).
2. Set the “event” and “case_event” variables to zero, and the “case” variable to 1.
3. Using those inputs, calculate the prediction from the event model described above.

The model is fit separately for each hour of the day, for each customer sub-group and event day combination, with their outputs stored as 24-hour load shapes. Although the regression models provide error estimates these estimates are too optimistic because they are based on assumptions (such as statistical independence of errors) that do not apply to this situation. Instead, 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[10]. 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 (note that we use cooling degree hours – the number of degrees above a threshold temperature the outside temperature is for each hour of the day to restrict cooling estimates to warmer temperatures). 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 cannot 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[11]. 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 two 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 ex post 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, dual program enrollment, CARE enrollment, and LCA for the purposes of forecasting future event outcomes.

---

[10] 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.

[11] 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.

(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) is small overall, they are anticipated to be the leading edge of a territory-wide roll out of default 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.

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 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 17,500 hourly event estimates. Of those 3,641 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	405
10-20 W	666
20-50 W	1,797
50-100 W	683
>100 W	90
Standard error to impact ratio	**Count**
<0.1	1,475
0.1-0.2	1,502
0.2-0.5	577
0.5-1	69
>1	18

There are 773 hours with standard errors over 50W but just 87 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. Figure 9 illustrates the relationship between standard errors (y-axis in kW) and the count of enrolled event participants (x-axis). 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.

Figure 9: Relationship between number of enrolled event participants and computed standard errors (in kW), colored by LCA, with vertical line at 1000 enrolled customers for reference.

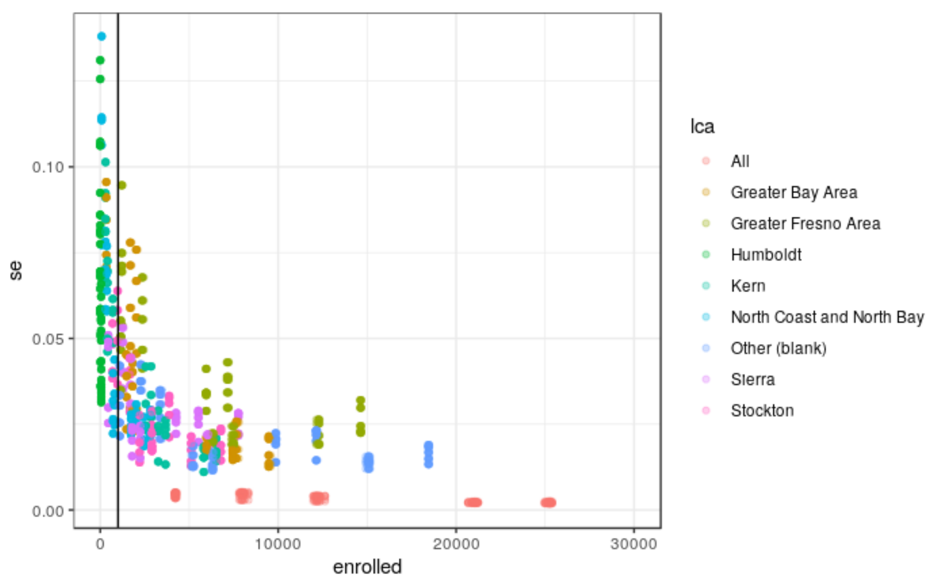

Ex Ante Impact Analysis Methods

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 SmartRate™ 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 2019 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. These are calculated using the same model used for the load shed estimates, as described in the "Event model" section above.

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 2019, since that is still the official event window for future events. The events are evaluated over the RA window of 4-9 pm, which was also the RA window last year.

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 SmartRate™ 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

2019 SmartRate™ ex post results on a typical 2019 event day were 14.89 MW in aggregate or 0.22 kW per customer. Table 4 summarizes the outcome of every SmartRate™ event of the 2019 season.

While aggregate impacts were lower than in 2018 (primarily due to the lower participation numbers), per customer impacts are higher than in 2018. Note that the Humboldt LCA contains 56 customers - too few to materially influence the average or aggregate event results reported in this document. For these reasons, Humboldt results will be excluded from results and figures in this report. Humboldt results are available in the private ex post table generator (for those with appropriate access) and the ex ante table generator (because forecasts do not disclose actual usage).

Table 4 summarizes the outcomes for all nine 2019 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 2019 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)
2019-06-11	66,508	2.35	0.27	156.05	18.17	12	99
2019-07-24	66,555	2.21	0.22	146.77	14.80	10	97
2019-07-26	66,584	2.14	0.21	142.71	13.73	10	94
2019-08-13	66,484	2.00	0.21	132.66	13.67	10	95
2019-08-14	66,489	2.27	0.24	150.60	15.90	11	99
2019-08-16	66,474	2.41	0.26	160.49	17.00	11	98
2019-08-26	66,445	2.16	0.24	143.67	15.72	11	96
2019-08-27	66,467	2.14	0.22	142.06	14.59	10	96
2019-09-13	66,528	1.78	0.16	118.72	10.41	9	95
Typical event	66,503	2.16	0.22	143.75	14.89	10	97

Aggregate Ex Post Summary

The ex post aggregate load impact on a typical 2019 event day was 14.89 MW (down from 17.83 MW in 2018), with the largest aggregate load shed (18.17 MW) occurring on June 11th. Figure 10 depicts the breakdown of aggregate impacts across LCAs for all 2019 events.

Figure 10. Aggregate 2019 ex post load shed (overall, by date and LCA)

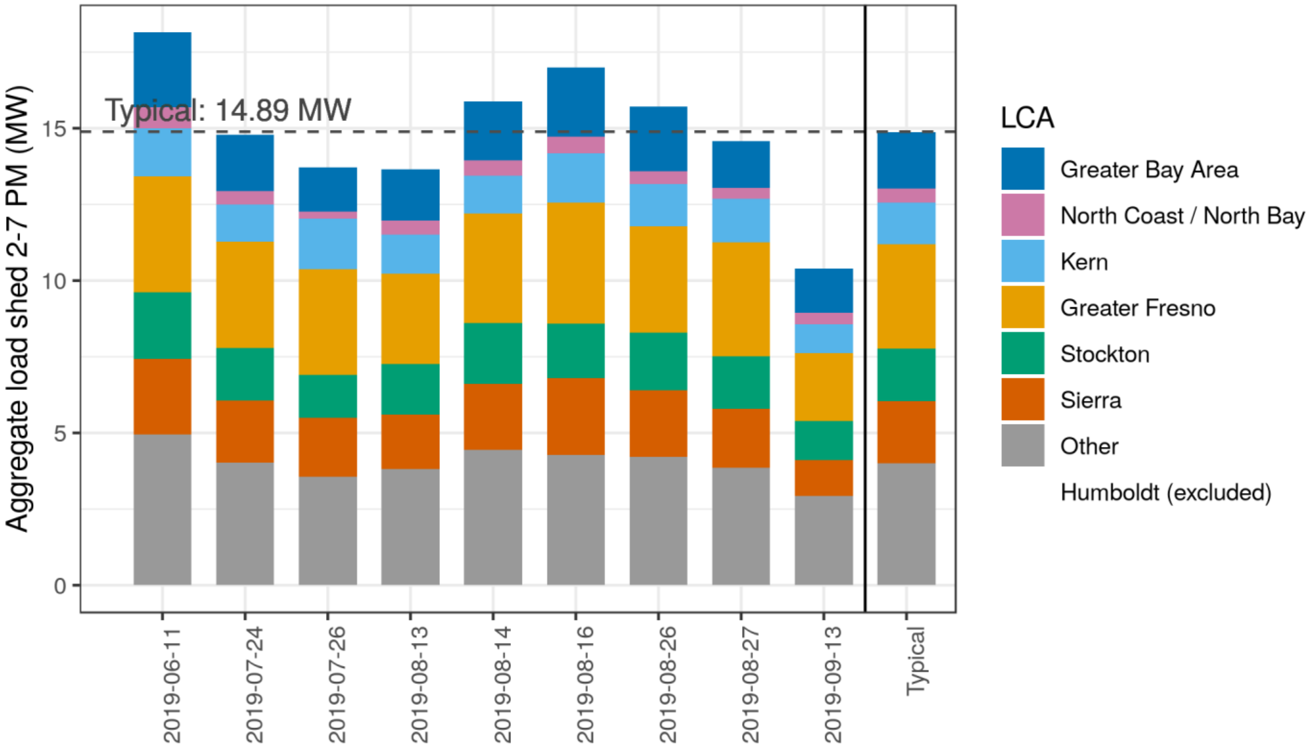

Figure 11 illustrates the aggregate load impact contribution of every sub-group relevant to program performance for each of the nine event days. For each event day, the sum across each set of categorical sub-groups totals the "All customers" aggregate impact.

Figure 11. Aggregate load impacts by customer sub-groups

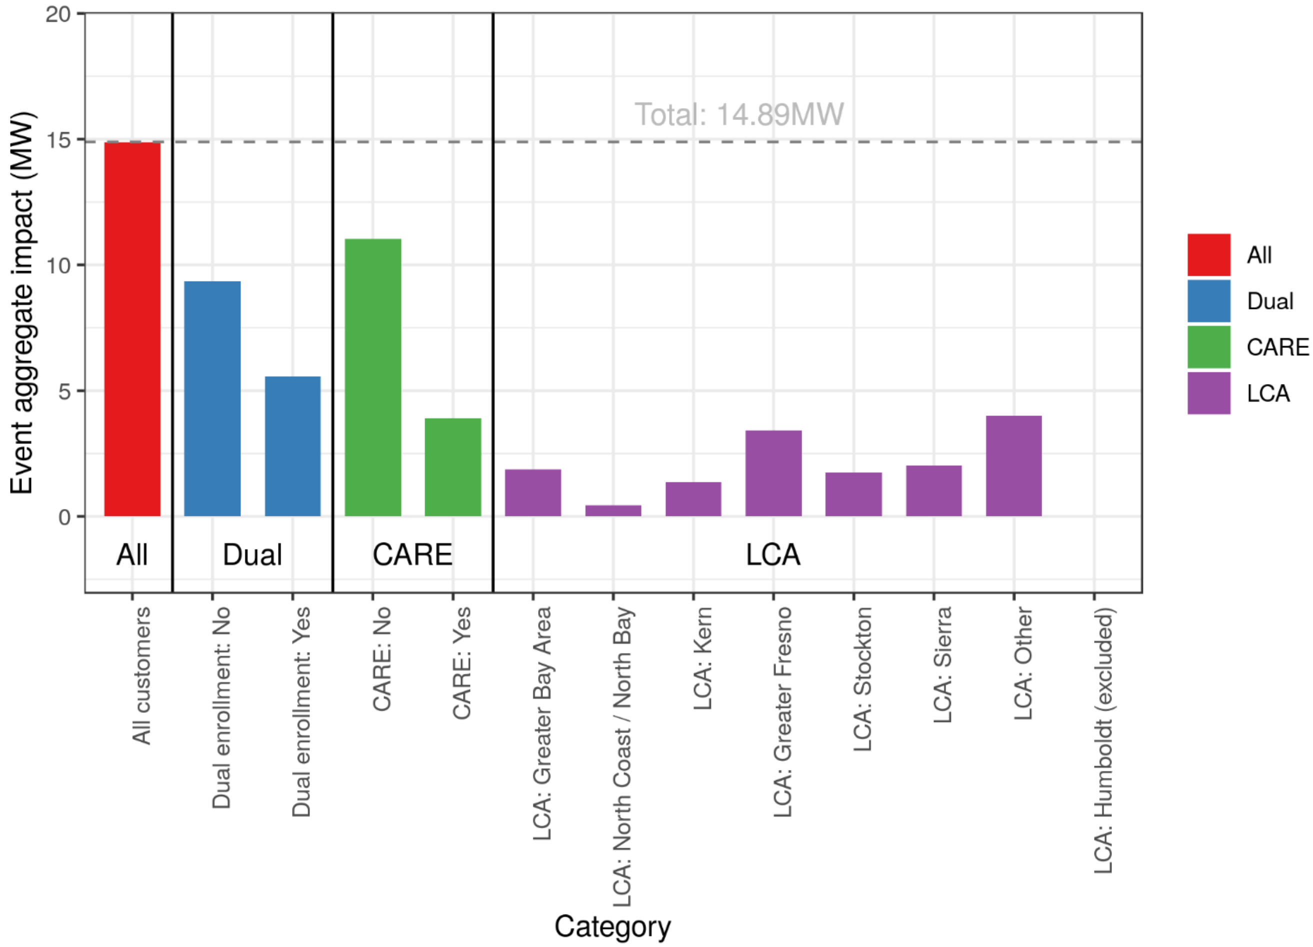

Figure note: The categories with multiple groups add to All.

From this view, it can be verified that:

- Dually enrolled customers contributed less than SmartRate™ only customers, *in aggregate*. The continued decline in the number of dually enrolled customers leaves them lagging behind SmartRate™ only customers in aggregate impacts.
- CARE customers contribute a smaller fraction of aggregate impacts than customers not on CARE.
- Fresno is the named LCA with the largest aggregate contribution to the total, but the total shed of "Other" customers outside of LCAs is slightly larger than Fresno's contribution. Greater Bay Area, which was among the largest contributors in PY2018, is no longer a standout contributor due to the loss of enrolled customers to CCAs.

Average Ex Post Load Impacts by Event

On a typical event day, the average per customer load shed was 0.22 kW in 2019—much higher than the 0.16 kW per customer in 2018 and close to the 0.23 kW level in 2017. This varied slightly across the nine 2019 events (i.e., some of the black dots in Figure 12 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 June 11th when the average temperature during the event hours (across all LCAs) was 99 degrees.

Figure 12. Mean 2019 ex post event load shed (overall, by date and LCA)

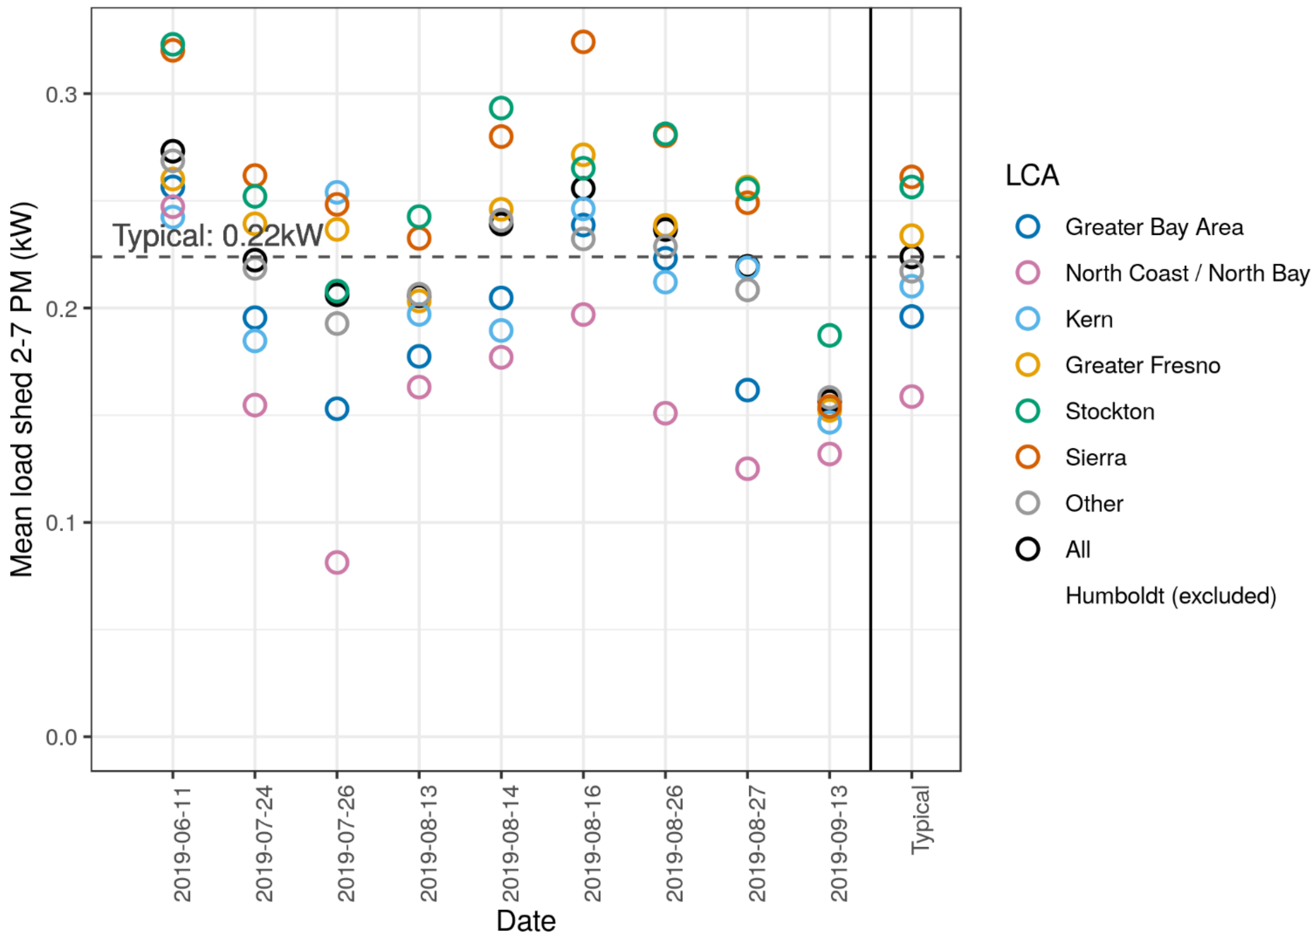

The load shed varies across event days, as well as across LCAs, as represented by the vertical spread of the points on each day. Results by LCA are discussed in more detail in the LCA-specific results section below.

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**	66,504	97	2.09	0.22	14.89	143.75	10.4%
LCA	**Humboldt (excluded)**							
	Greater Bay Area	9,482	91	1.53	0.20	1.86	15.55	**12.0%**
	Greater Fresno	14,635	102	2.55	0.23	3.42	**38.36**	8.9%
	Kern	6,523	101	2.60	0.21	1.37	17.41	7.9%
	North Coast / North Bay	2,816	94	1.31	0.16	0.45	3.94	11.3%
	Other	18,454	94	1.90	0.22	**4.01**	36.27	11.0%
	Sierra	7,760	97	2.06	**0.26**	2.03	16.49	12.3%
	Stockton	6,778	98	2.23	**0.26**	1.74	15.59	11.1%
Dual	**SmartRate™ only**	54,272	97	2.08	0.17	9.36	113.08	8.0%
	Dually enrolled	12,232	98	2.10	**0.46**	5.57	25.64	**20.5%**
CARE	**Non-CARE**	41,281	96	2.00	**0.27**	11.04	82.55	**12.8%**
	CARE	25,223	99	2.23	0.16	3.90	56.26	6.8%

Table note: Enrollment numbers are based on the typical event day for 2019, 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. Aggregate results exclude Humboldt due to confidentiality and too few customers to materially influence the average or aggregate event results.

A review of the LCA rows reveals that they experience a 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. Stockton and Sierra were the third and fourth hottest LCAs but featured the largest per-customer impacts at 0.26 kW per customer.

A review of the dually enrolled vs. SmartRate™-only rows reveals that more dually enrolled customers are in hotter climates than SmartRate™ only customers, leading to a slightly higher (1 degree) average temperatures. That difference carries into the reference loads.

The per-customer impacts for dually enrolled customers show a 20.5% impact as a % of reference load response. This surely bears the imprint of the automated air conditioning controls deployed for dually enrolled participants. However, 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 5.57 MW compared to nearly 9.36 MW from 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 3 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 3.9 MW compared to 11 MW from non-CARE customers. CARE customers comprise about 38% of the enrolled population. See the *CARE Enrollment Status Results* below for more details and discussion of CARE customers.

LCA-specific Results

The 2019 ex post results indicate that the mean load shed for customers varies across LCAs. Figure 13 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 and Stockton) 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 North Coast / North Bay LCA participants contributed the least *per customer* as well as the least *in aggregate* despite reasonable impacts as a % of reference loads.

Figure 13. Summary of aggregate, per-customer load impact and load impact as a % of reference loads for LCAs

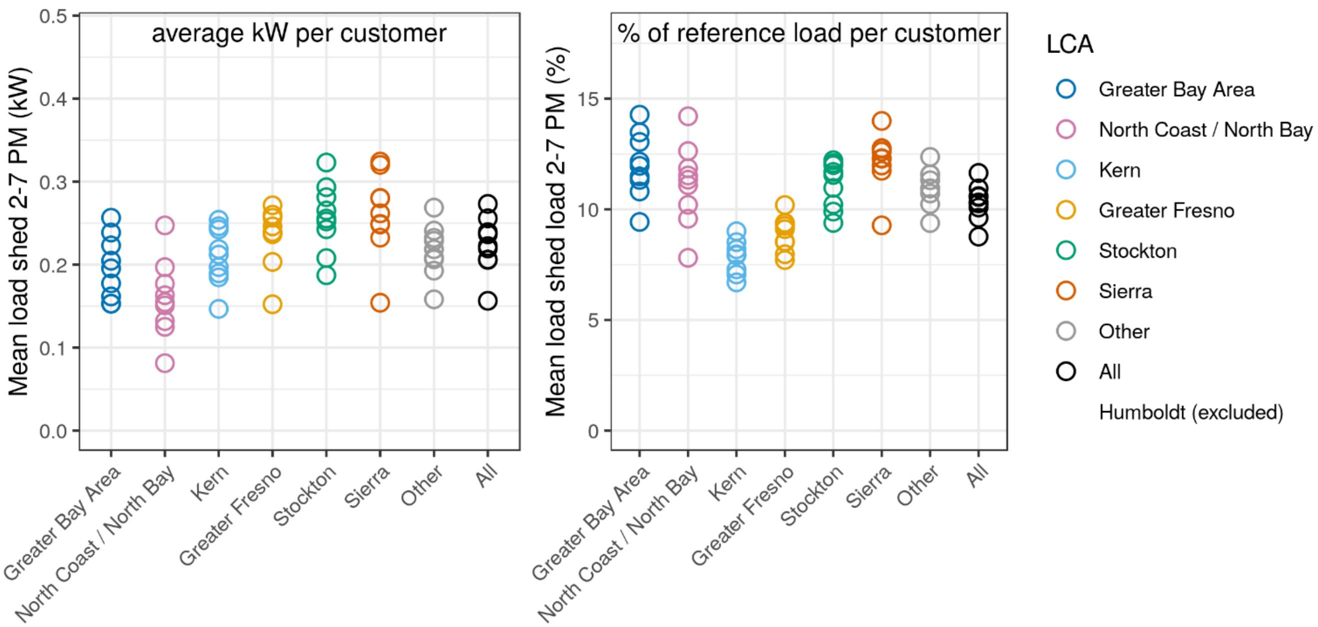

To better examine relative performance by LCA, Figure 14 magnifies the middle panel above to provide a closer look at per-customer event impacts by LCA, with each of the nine events from 2019 plotted as a colored open circle for each LCA. The 'All' LCA represents every enrolled customer in the program 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 somewhat bisects the results in the Greater Bay Areas and Kern LCAs (+ Other), while Greater Fresno, Stockton and Sierra LCAs were generally above the average. North Coast/North Bay is a cooler LCA that consistently saves less than the average.

Figure 14. Average participant load shed by LCA (each dot is an event)

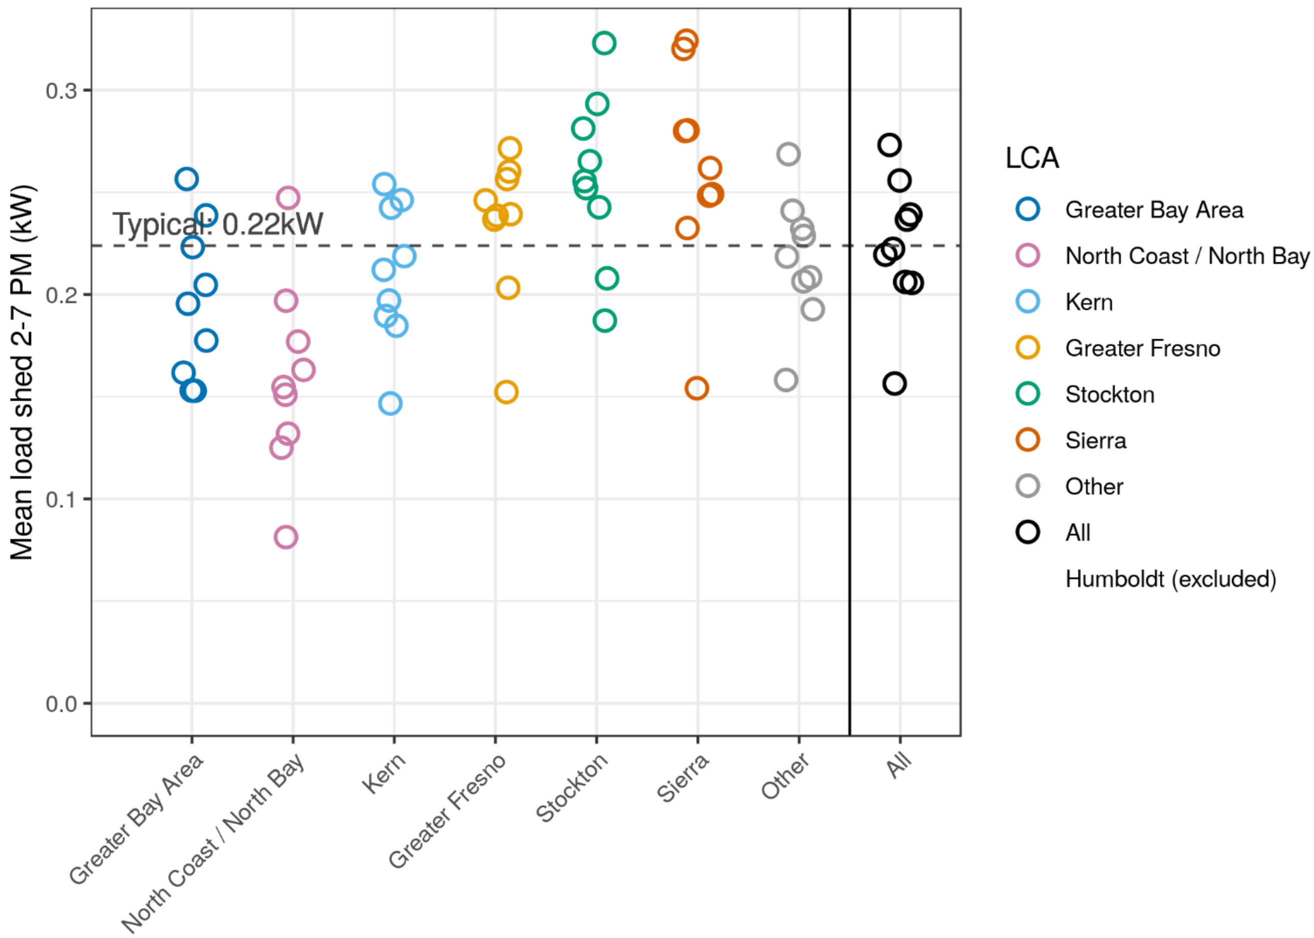

In looking at the contribution to the aggregate impacts, Greater Fresno contributes the largest percentage among the named LCAs (23%, see Figure 15). "Other" also plays an important role.

Figure 15: Contribution of LCA to Aggregate Load Shed

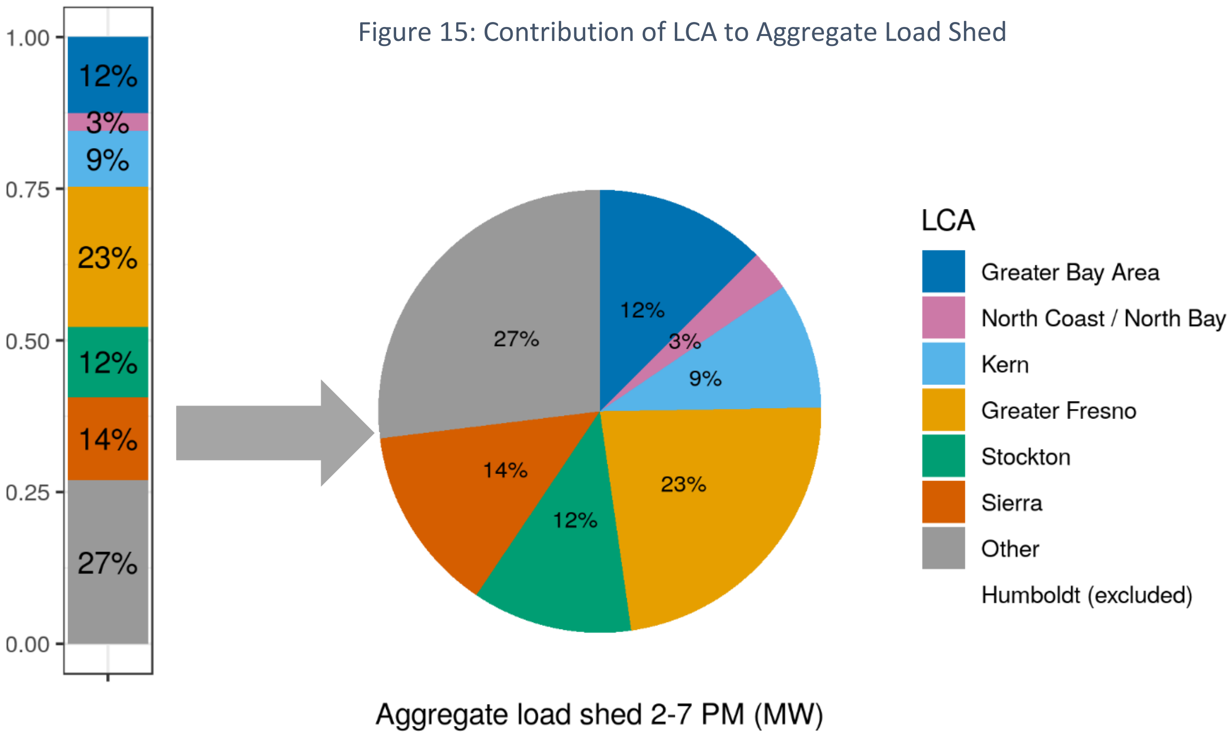

When we look at the contribution of SmartRate™ only and dually enrolled across the LCAs, SmartRate™ only customers contribute more to the overall aggregate load shed than dually enrolled customers. See Figure 16, which depicts the range of outcomes across the nine event days as a box with whiskers, where the box spans the 25^{th} to 75^{th} percentile of the results. There is no LCA where dually enrolled customers contribute aggregate savings greater than SmartRate™ only customers.

Figure 16. Contribution to the Aggregate Load Shed by LCA and Enrollment Status

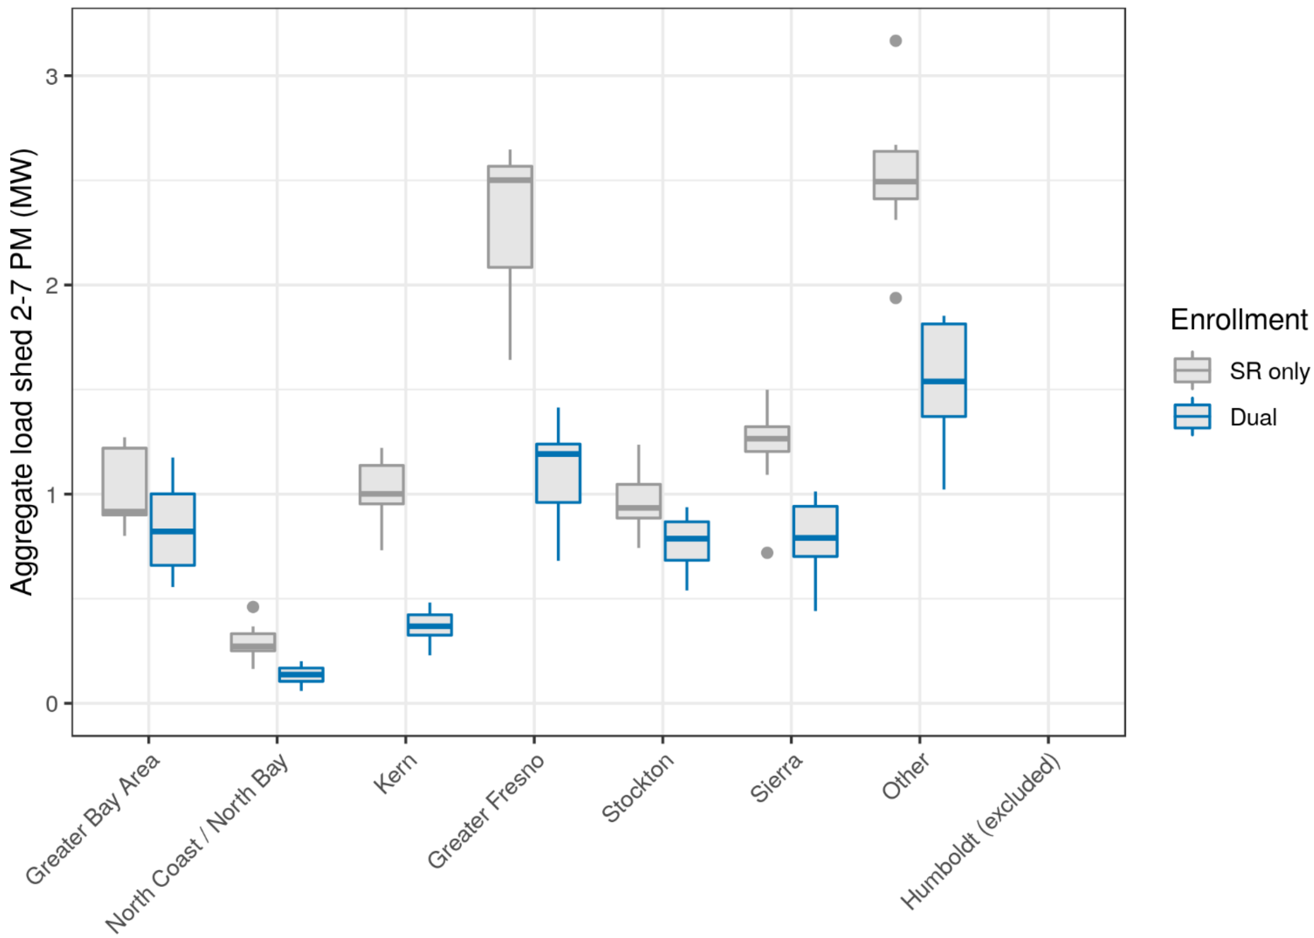

Dual Enrollment Status Results

Figure 17 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 17. Event load impact aggregate MW, per-customer kW and % of reference load by dual enrollment status

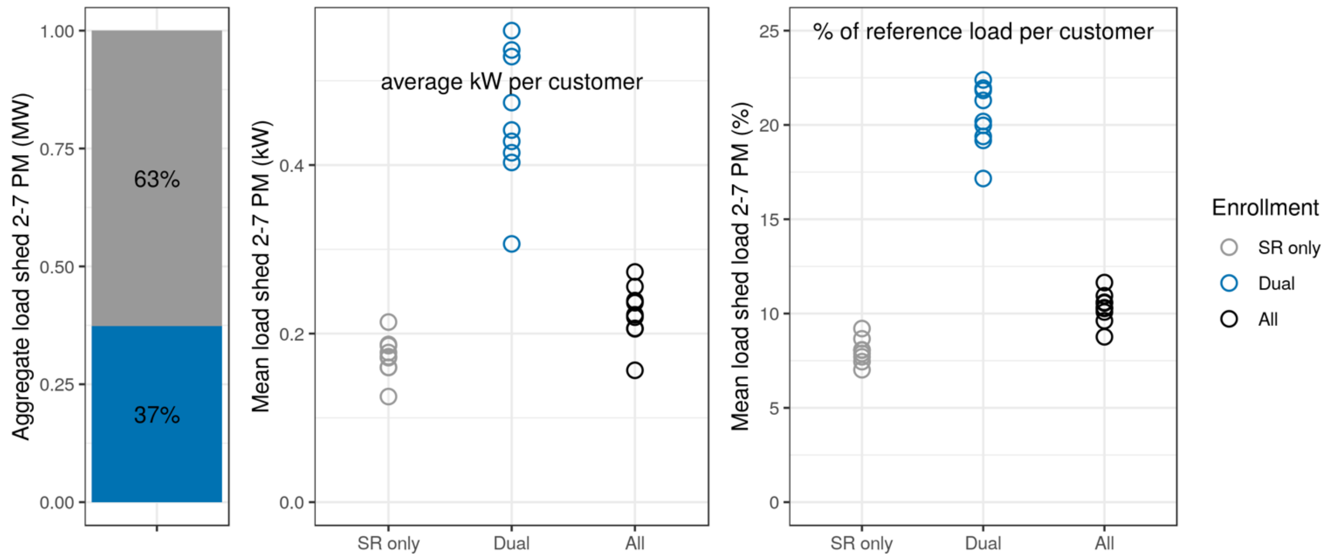

Dually enrolled participants typically save much more than SmartRate™ only participants. Dually enrolled participants also save a much larger percentage of their load – averaging and often exceeding 20% (compared to less than 10% by the SmartRate™ only participants), which demonstrates that direct load controls have the ability to save much more than what is typical of customer behaviors alone. These results are consistent with past results.

Across the season, dually enrolled customers contribute about 37% of the total aggregate load. As mentioned above, the majority of the aggregate impacts (63%) come from SmartRate™ only participants.

One issue with comparing dually enrolled and SmartRate™ only participants 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 18 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). These are highlighted at 1.5 and 2 kW reference loads to allow for comparisons.

Figure 18. Event load shed vs reference load – dually enrolled and SmartRate™ only (highlights at 1.5 kW and 2 kW reference loads)

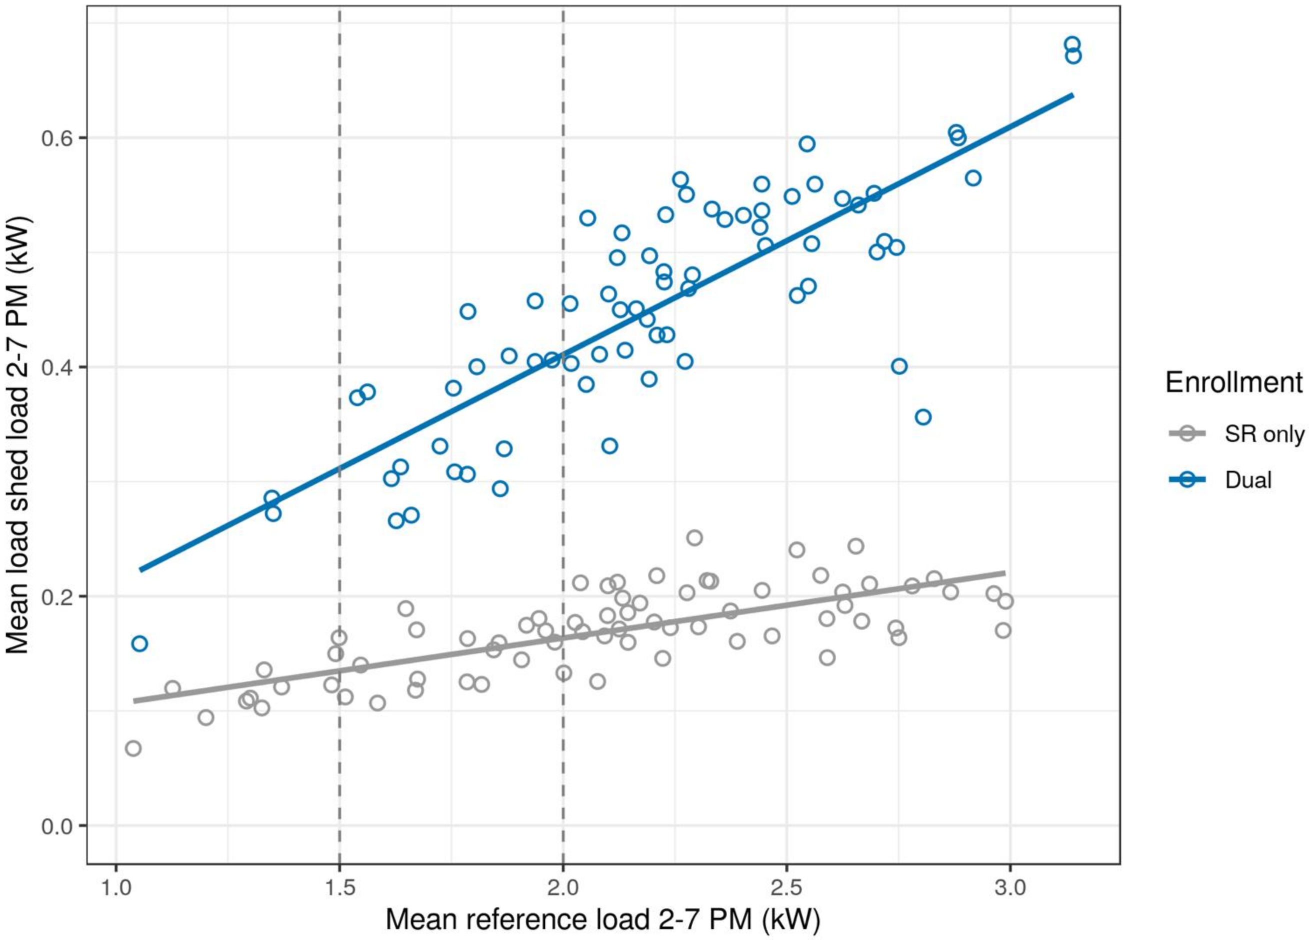

From this plot, it can be verified that there is not a systematic difference between reference loads (i.e. the range of values along the x-axis) of the two groups (SR only extends just slightly lower and Dual extends only slightly higher). At the same reference loads, dually enrolled customers have impacts roughly double their SmartRate™ only peers.

CARE Enrollment Status Results

Figure 19 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 19. Event load impact aggregate MW, per-customer kW and % of reference load by CARE status

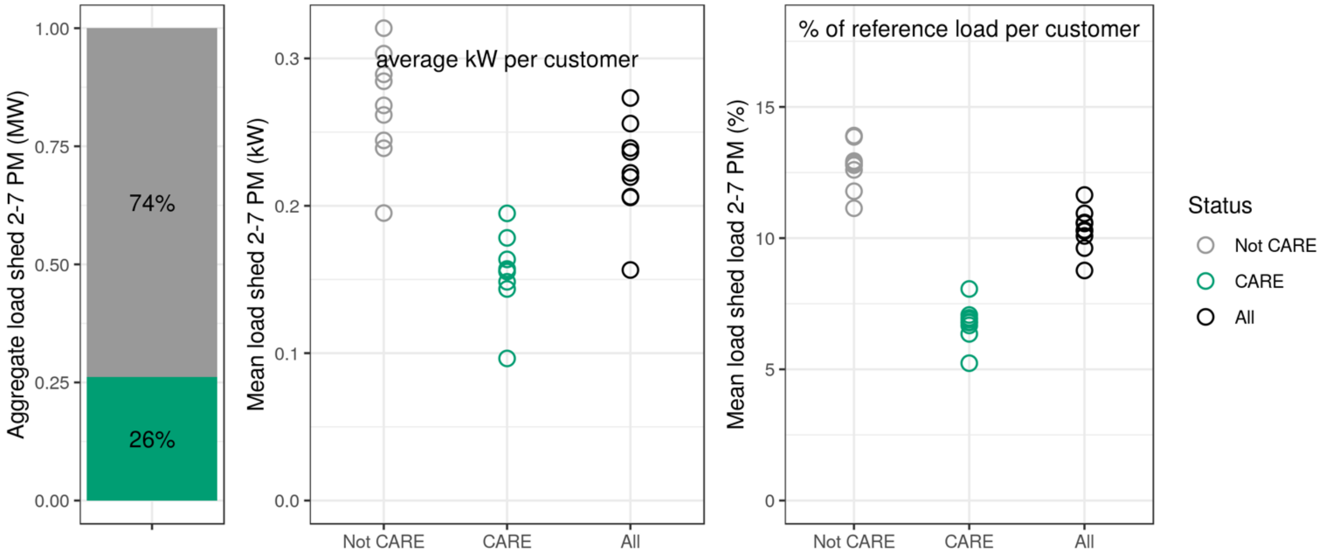

Consistent with past findings, across the events, CARE customers tend to save less than individuals who are not on the CARE program, despite higher reference loads, due to significantly lower load impacts as a % of reference loads. For comparison: they make up a smaller portion of the aggregate impacts than dually enrolled customers despite out numbering them nearly 2 to 1 (there were 25,000 CARE customers compared to 12,000 dually enrolled customers in 2019). CARE customers experience similar weather and larger 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 (6.8% among CARE customers compared to 10.4% among all customers, or 20.5% among dually enrolled customers).

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. We do not have the data to tease those effects apart from one another for this evaluation.

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

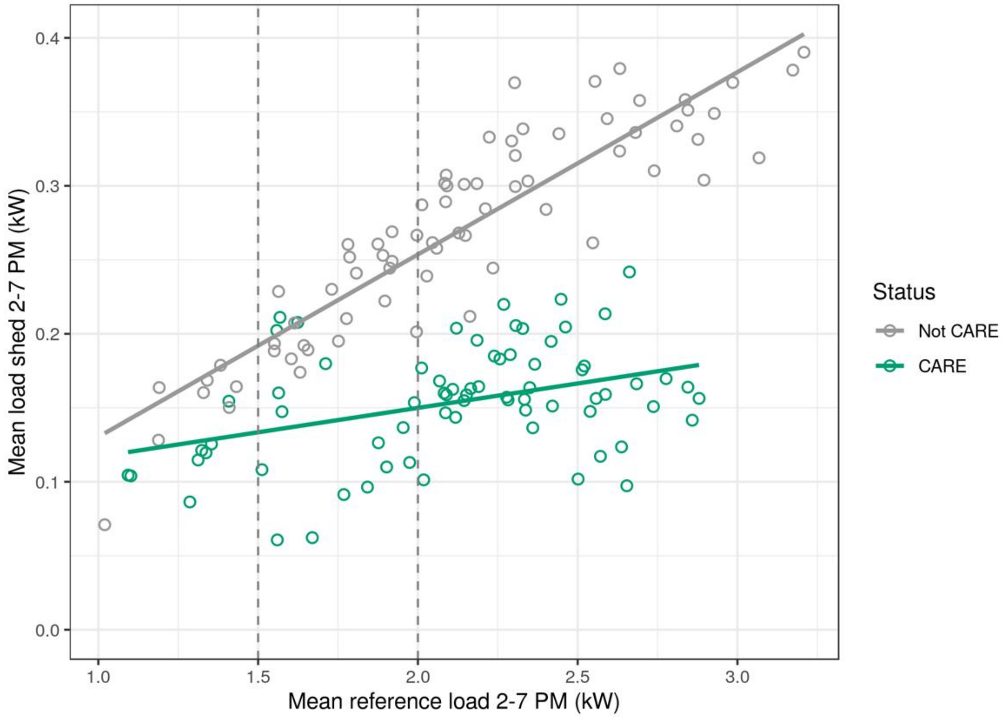

From this plot, it can be verified that there is a systematic difference between reference loads of the two groups. In this case, CARE customers have a lower maximal extent of reference loads, but they have a higher density of points at elevated reference loads and their average is, in fact higher, than 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 two-thirds of non-CARE customers.

Interestingly, the dually enrolled customers on CARE perform almost as well on average as dually enrolled non-CARE customers, most likely because the utility is controlling their loads during this time (see Figure 21). It is the SmartRate™ only CARE customers (shown in blue below) that appear to be the lowest performing segment. CARE customers most notably underperform when they are without direct load controls.

Figure 21. Event load shed by CARE status – dually enrolled vs SmartRate™ only

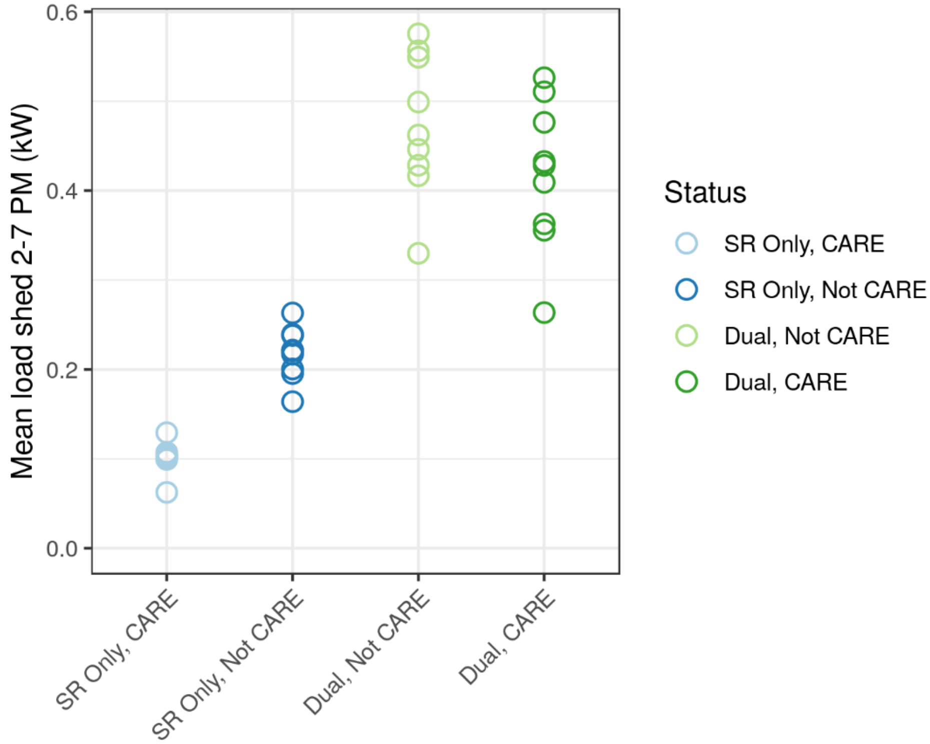

Curent Ex Post to Prior Ex Post

We compared 2019 ex post results to the prior year's ex post results for their respective typical event days. (See Table 6.) Overall, our 2019 ex post estimates are lower than in 2018 because of the decreasing number of participants, and the shifting make-up or composition of participants. Notably, however, the per customer values in 2019 were similar to 2017, which also had a similar average event temperature.

Table 6. Comparison of 2019 ex post load impacts on a typical event day compared to 2018 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	2019 Ex Post	66,504	14.9	143.7	0.22	2.16	10.4%	97
	2018 Ex Post	109,972	17.8	166.0	0.16	1.51	10.7%	88
	2017 Ex Post	124,049	28.1	226.0	0.23	1.82	12.4%	98
Dually Enrolled	2019 Ex Post	12,232	5.6	27.1	0.46	2.22	20.5%	98
	2018 Ex Post	18,456	7.0	35.4	0.38	1.92	19.7%	94
	2017 Ex Post	28,923	14.3	69.5	0.50	2.40	20.6%	100
SR only	2019 Ex Post	54,272	9.4	116.5	0.17	2.15	8.0%	97
	2018 Ex Post	91,515	10.9	130.6	0.12	1.43	8.3%	87
	2017 Ex Post	95,126	13.8	156.5	0.14	1.64	8.8%	97

Table note: Humboldt numbers have been excluded.

Figure 22 displays the average event impacts (y-axis) for all events in 2017-2018 (gray) and 2019 (purple) against the average outside temperature for each event. The figure illustrates that impacts are sensitive to temperature, that the temperatures in 2019 were significantly warmer than 2018, and roughly equivalent to those in 2017.

Figure 22. Mean load shed by temperature (2017 - 2019)

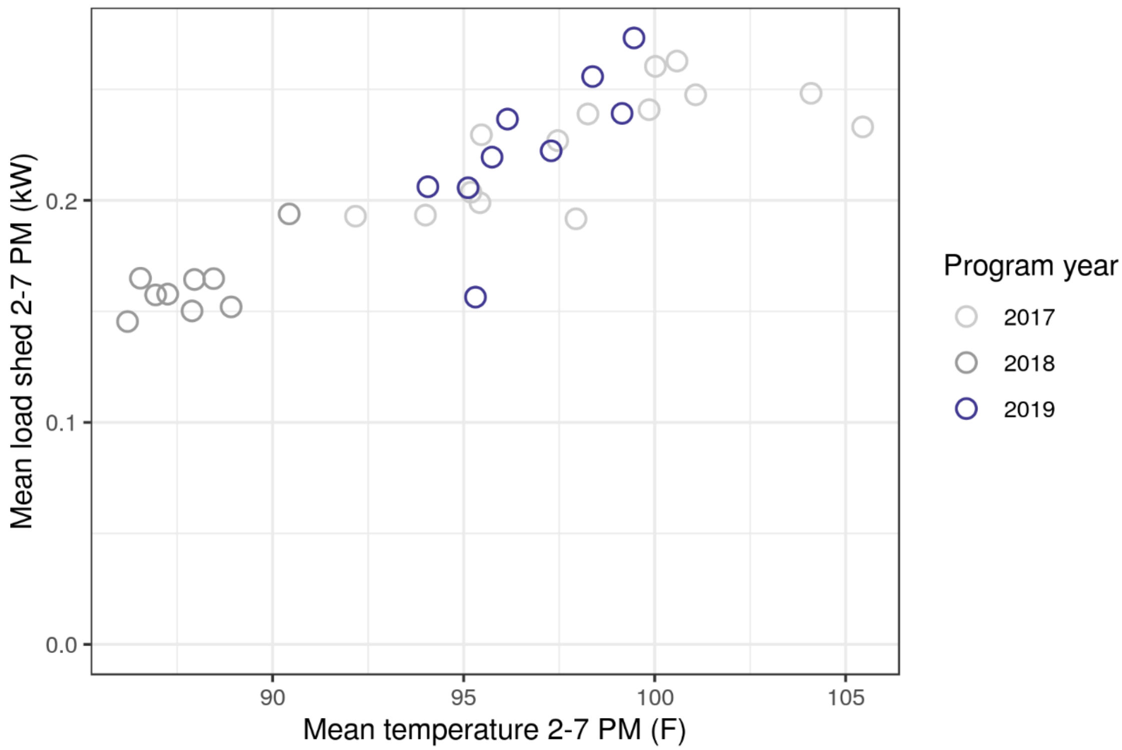

Along with temperatures, LCA and dual enrollment status are the most important determinants of load impacts, and they all interact. Figure 23 extends the use of scatter plots to present load impacts vs. temperatures for all events in 2017- 2019, broken out by dual enrollment status. The green/blue points are from dually enrolled customers and the red/orange points are from SmartRate™ only customers.

Figure 23: Event load shed vs. outside temperature by LCA (2017 -2019); Humboldt excluded

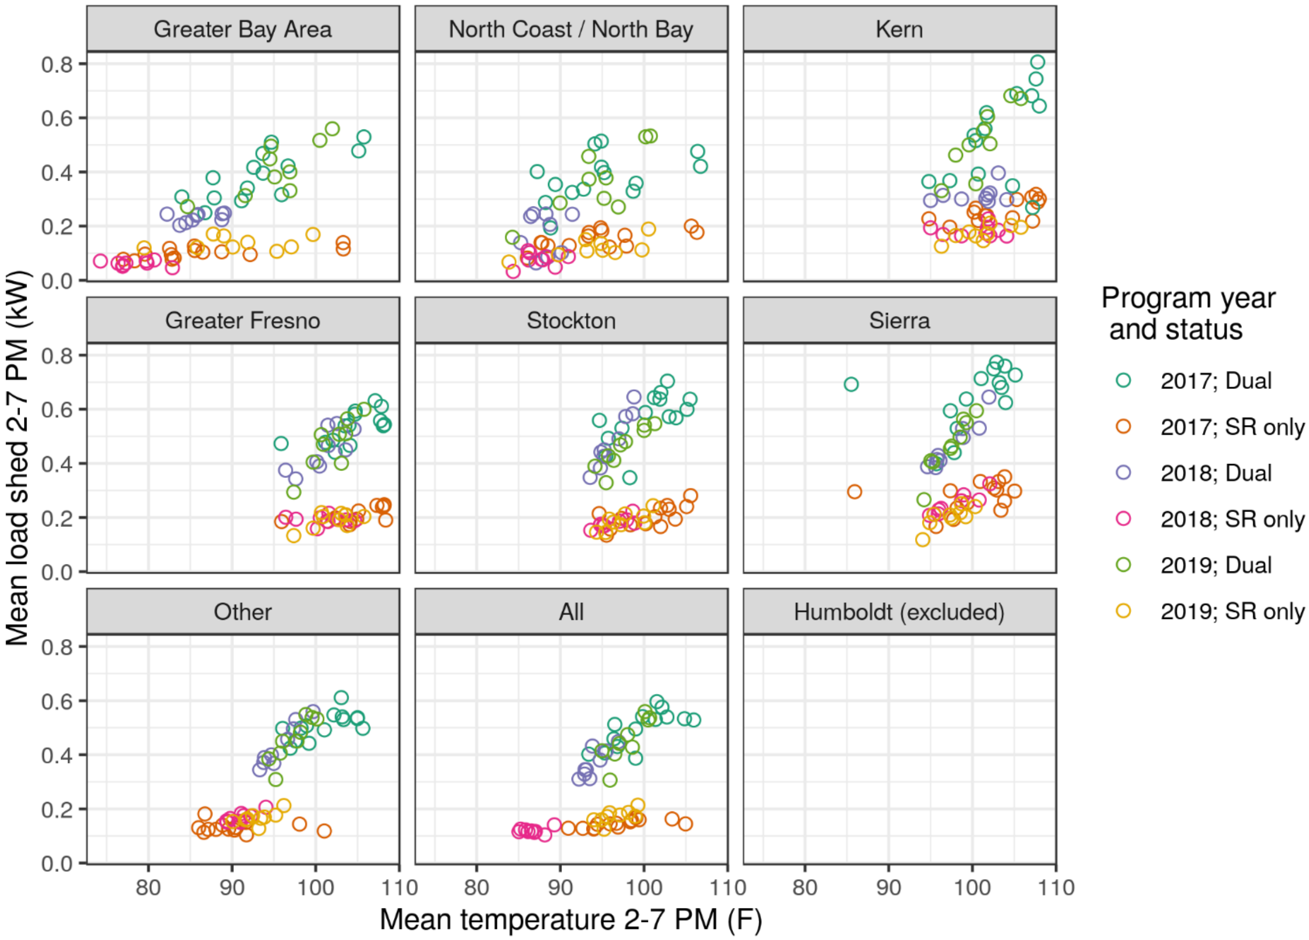

From the figure it can be seen that the 2019 ex post results are generally consistent with (i.e. grouped with and extending the trend of) the 2018 results after taking into account LCA, dual enrollment status, and event temperatures. The figure also reveals more subtle patterns, including the tendency 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.

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 2018 to 2019 leading to reduced impacts

	Enrollment	Event avg. Temp (F)	CDH	Reference load (kW)	Impact (kW)	Impact (%)	Aggregate impact (MW)
2019	66,447	97	421	2.16	0.22	11	14.9
2018	109,972	88	267	1.51	0.16	10	17.8
Difference ('18-'19)	-43,524	9	154	0.58	0.06	1	-2.9
% Difference	-40%	10%	58%	38%	38%	10%	-16%

Table note: Humboldt numbers are not included.

When comparing 2019 to the prior year, aggregate impacts dropped 16%, while enrollment dropped 40%. The reduction in the overall aggregate impacts was less than the drop in enrollment due to the increase in the other factors (temperatures, CDH, reference loads, etc.) The 38% increase in reference loads can be explained by the increase in event temperatures in 2019 – partly due to the large loss of customers in the relatively cool Bay Area, leading to higher mean temperature and thus higher air conditioner use among those remaining.

Because CCAs serve specific geographies, de-enrollments affected some segments of the populations more than others. As previously noted, the Greater Bay Area LCA saw much greater losses in the number of participants than their counterparts, with Sierra seeing the second largest drops. These LCAs also saw the largest drop in their aggregate impacts. See Table 8.

Table 8. 2018 to 2019 change in aggregate impacts and end of summer enrollment by LCAs

		Aggregate impact (MW)							
	Program year	**Greater Bay Area**	**North Coast / North Bay**	**Kern**	**Greater Fresno**	**Stockton**	**Sierra**	**Other**	**All**
Enrollment on typical event day	2019	9,482	2,816	6,523	14,635	6,778	7,760	18,454	66,504
	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
Aggregate impact	2019	1.9	0.4	1.4	3.4	1.7	2	4	14.9
	2018	4.3	0.3	1.3	3.5	1.8	2.4	4.2	17.8
	2017	8.7	0.4	1.4	4.3	2.2	5.4	5.7	28.1

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 22). 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 34%. However, the aggregate impacts of the group of dually enrolled customers dropped by 20% because of the warmer event temperatures.

Table 9. 2018 to 2019 change in aggregate impacts and enrollment by Dual enrollment status

	Aggregate impact (MW)			Enrollment		
Program year	**SR Only**	**Dual**	**Total**	**SR Only**	**Dual**	**Total**
2019	9.4	5.6	14.9	54,272	12,232	66,504
2018	10.9	7	17.8	91,515	18,456	109,972
Difference	-1.5	-1.4	-2.9	-37,243	-6,224	-43,468
% Difference	-14%	-20%	-16%	-41%	-34%	-40%

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. 2018 to 2019 change in aggregate impacts and enrollment by CARE enrollment status

	Aggregate impact (MW)			Enrollment		
Program year	**Total**	**Not CARE**	**CARE**	**CARE**	**Not CARE**	**Total**
2019	14.9	11	3.9	25,223	41,281	66,504
2018	17.8	13.4	4.4	32,397	77,573	109,972
Difference	-2.9	-2.4	-0.5	-7,174	-36,292	-43,468
% Difference	-16%	-18%	-11%	-22%	-47%	-40%

Factor		Changes between 2018 and 2019	Individual average results expected differences (Mean Load Shed)		Total aggregate result expected difference (Aggregate Load Shed)	
Total participation		Participation dropped from nearly 110,000 to closer to 66,000, with de-enrollments centered around the Greater Bay Area, causing the largest changes in aggregate impact and composition of enrolled customer characteristics.	Average customer now lives in a warmer climate and experiences more temperature sensitive impacts, resulting in an increase in per-customer impacts.	•	Overall decrease in aggregate impacts	↓
Shifting composition of participants	Dually enrolled	While the count of dually enrolled customers dropped 34%, these customers make up a slightly larger percentage of the population (18.5% rather than 17%), increasing average results. Dually enrolled per customer impacts went up significantly (due to weather).	Increase in average per customer results	↑	Overall fewer dually enrolled outweighs increased percentages and per customer impacts from this group	↓
	CARE	The number of CARE customers dropped 22%. However, they are now a larger percentage of the population (38% rather than 29%) –both decrease impacts.	Decrease in average per customer results	↓	Overall lower number greater percentage of CARE decrease impacts	↓
	Overall make-up by LCA	Enrollment dropped significantly (83%) in the Greater Bay Area, which has among the lowest per customer impacts, while the other areas netted about the same number of participants.	Increase in average per customer results	↑	While shifting percentages can increase per-customer impacts, the decrease in the Greater Bay Area led to decreased impacts	↓
	TOU	TOU made up a larger percentage of participants in 2019 (23%), but overall there were fewer TOU customers (down roughly 6,300 or 29%)	The price effects of TOU rates could decrease what we expect at the individual level compared to non-TOU	↓	Overall lower numbers and more TOU customers decrease impacts	↓
Temperature and CDH		CDH increased in 2019 service territory wide and the CDH experienced by enrolled customers also increased due to Greater Bay Area de-enrollments.	Increases reference loads and load impacts at the individual level	↑	Increased impacts	↑

Curent Ex Post to Prior Ex Ante

As part of the 2018 model evaluation, ex ante predictions were made and reported for what the average load impact would be for events called during the 4-9 PM RA window under specified weather conditions. But events in 2019 were actually called from 2-7 PM, not 4-9 PM. Comparing the predictions for the 4-9 PM *RA* window with the actual results from the 207 PM *event* window would not be useful for either looking at changes in programmatic effectiveness or for evaluating statistical methods used for making the ex ante predictions. Instead, we compare prior ex ante load impact predictions for the event window to current ex post load impacts: As part of the 2018 model evaluation, ex ante predictions were also made for the case of events being called during the 2-7 PM event window, although they were not reported in the report on the 2018 program year in order to avoid confusion, since the RA window is the official time period for ex ante results.

In Table 11 we show the ex ante predictions that were made last year for the August peak day under the 1-in-2 PG&E weather scenario during the *2-7 PM event window*, compared to this year's actual event averages in that same time window.

Table 11. Comparison of ex post load impacts for the typical event day in 2019 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	2018 Ex Ante (for 2019+)	67,206	16.1	132.3	0.24	1.97	12.1%	97
	2019 Ex Post	66,504	14.9	138.6	0.22	2.16	10.4%	97
Dually Enrolled	2018 Ex Ante (for 2019+)	14,826	6.90	30.8	0.47	2.08	22.4%	99
	2019 Ex Post	12,232	5.6	27.1	0.46	2.22	20.5%	98
SR only	2018 Ex Ante (for 2019+)	52,380	9.16	101.5	0.17	1.94	9.0%	96
	2019 Ex Post	54,272	9.4	112.9	0.17	2.08	8.3%	97

Table note: The 2018 Ex Ante forecasts shown here are specific to PY2019 (and beyond). The table compares the forecast for 2019 to the actual numbers from 2019.

As Table 11 shows, the average event temperature in 2019 was remarkably close to the temperatures of the PG&E 1-in-2 weather scenario for August, and the per customer load impacts predicted for that scenario are quite close to the actual per customer load impacts from 2019. The program-wide aggregate load shed was somewhat lower than the ex ante prediction but that is almost entirely due to the fact that the number of dually-enrolled customers was somewhat lower in 2019 than had been forecast.

Ex-Ante Results

This section presents the results of CDA's ex ante forecast for the period 2020-2030. 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 2019 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 evaluations prior to 2018, 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 24 below.) This is consistent with the reporting for PY2018.

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.[12] 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 24 below. We discuss the effect of this three-hour overlap between the RA window and the event window throughout this chapter.

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

Figure 24. Predicted load shed by customers on the August peak (4-9 pm window, PG&E 1-in-2)

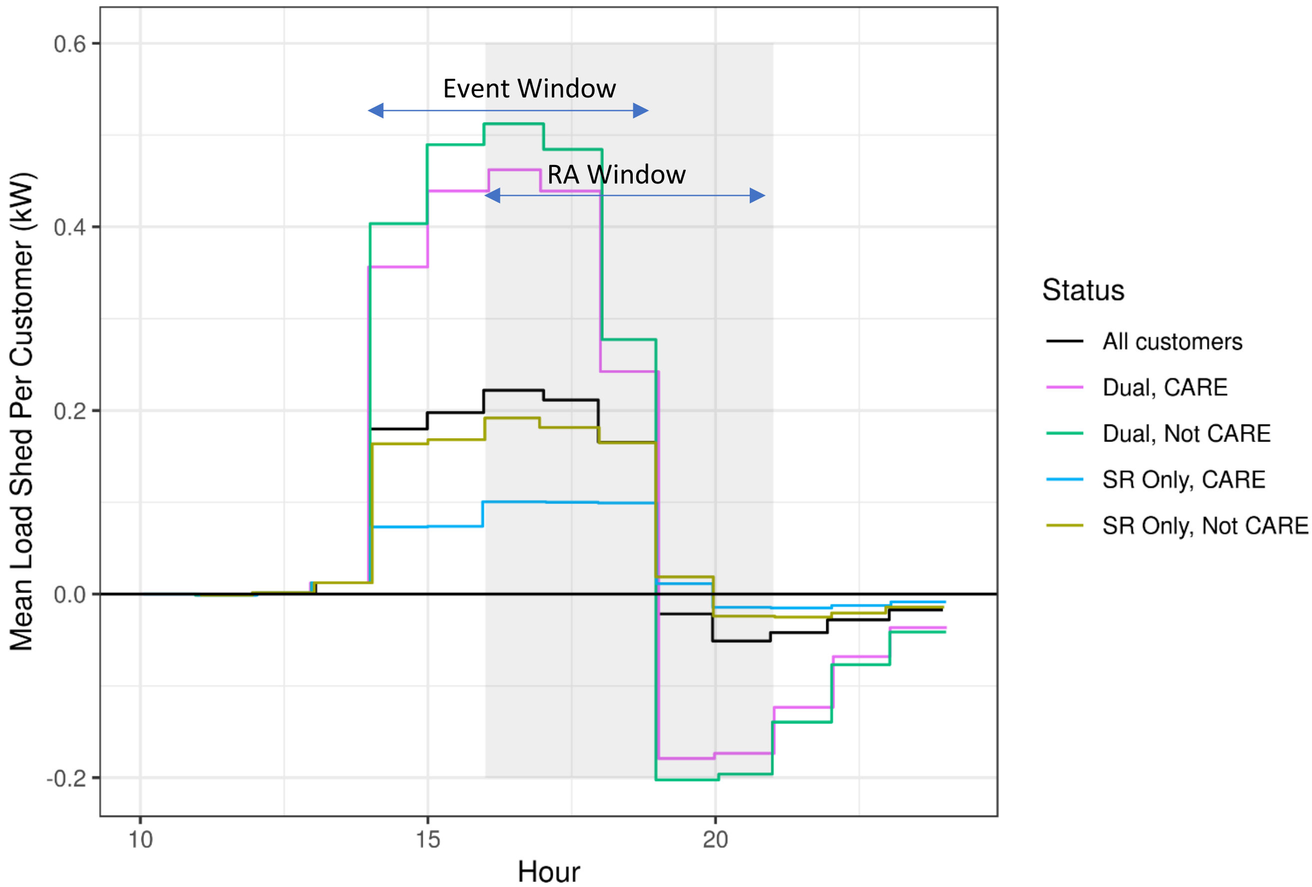

Figure note: The figure is for events that occurred 2-7 pm. The resource adequacy window is 4-9 pm, shown in gray.

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.

Over the 11-year forecast, PG&E assumes that participation levels stay approximately at 2019 levels of roughly 66,500[13] participants, approximately 12,000 (18%) 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).

[13] 66,519 at the end of 2019 and then steady at 65,673 for the remaining years.

Table 12. Forecasted SmartRate™ enrollments 2020-2030 (provided by PG&E)

	SmartRate™ only		**Dually enrolled**		**Total**	
LCA	**2020**	**2021-2030**	**2020**	**2021-2030**	**2020**	**2021-2030**
All	54,634	53,792	11,885	11,881	66,519	65,673
Greater Bay Area	7,494	7,494	1,970	1,970	9,464	9,464
North Coast / North Bay	2,342	2,342	296	296	2,638	2,638
Kern	5,061	5,061	630	630	5,691	5,691
Humboldt (excluded)						
Greater Fresno	12,055	12,055	2,251	2,251	14,306	14,306
Stockton	5,117	5,117	1,627	1,627	6,744	6,744
Sierra	6,103	6,103	1,628	1,628	7,731	7,731
Other	16,407	15,565	3,483	3,479	19,890	19,044

Humboldt contains 56 customers - too few participants to disclose their outcomes publicly. They do not have a meaningful effect on average or aggregate results, so they are not shown in any of the ex ante figures and tables. 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 25. Projected enrollment by (a) LCA 2019-2029 and (b) customer segment

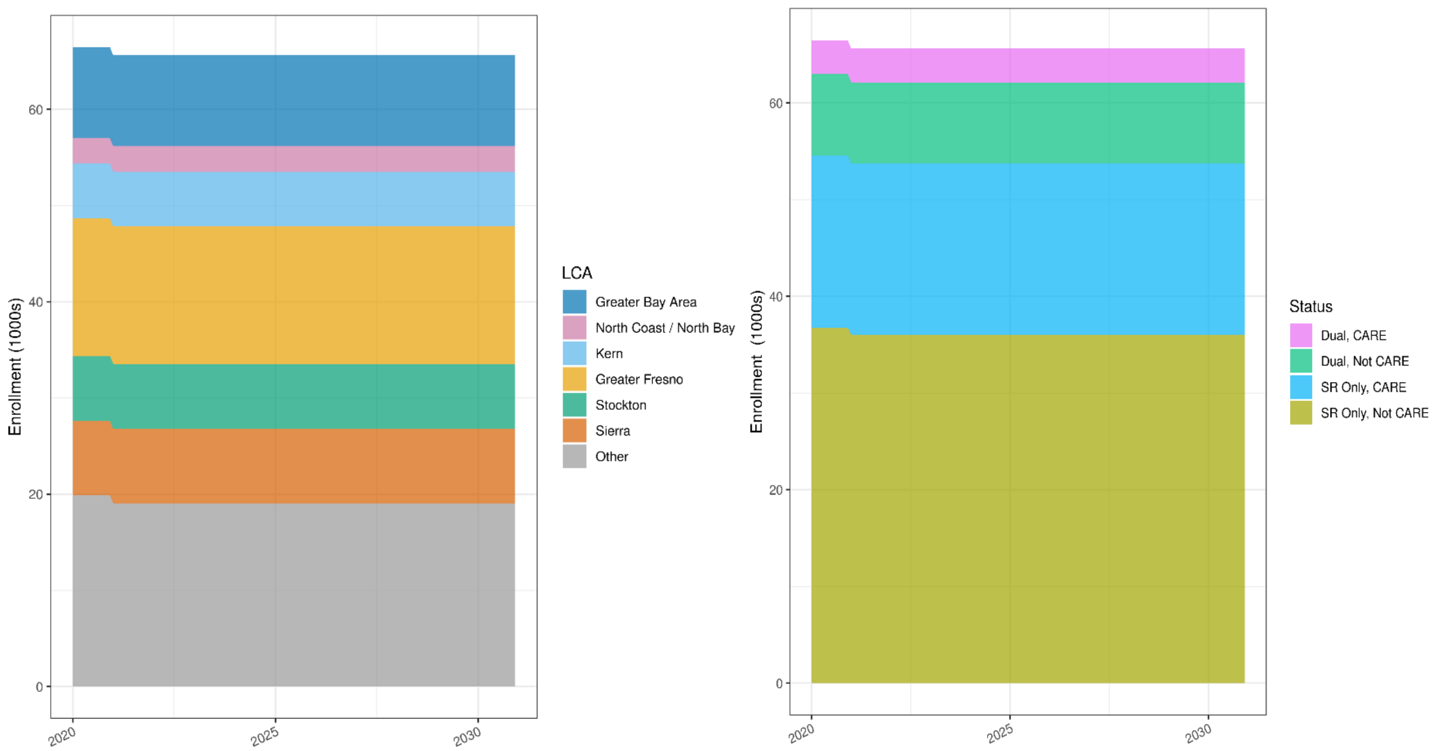

Our model does not explicitly incorporate TOU status. CDA assumed that future participants will behave in a manner similar to existing participants, even if more are on TOU rates. CDA was unable to incorporate TOU segments in this model due to limitations in the current data (i.e., the small number of 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 8.2 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 26 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 66,500 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 nearly 9.4 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 26. Mean and aggregate ex ante impacts during 4-9 RA event window (4 scenarios, summer only)

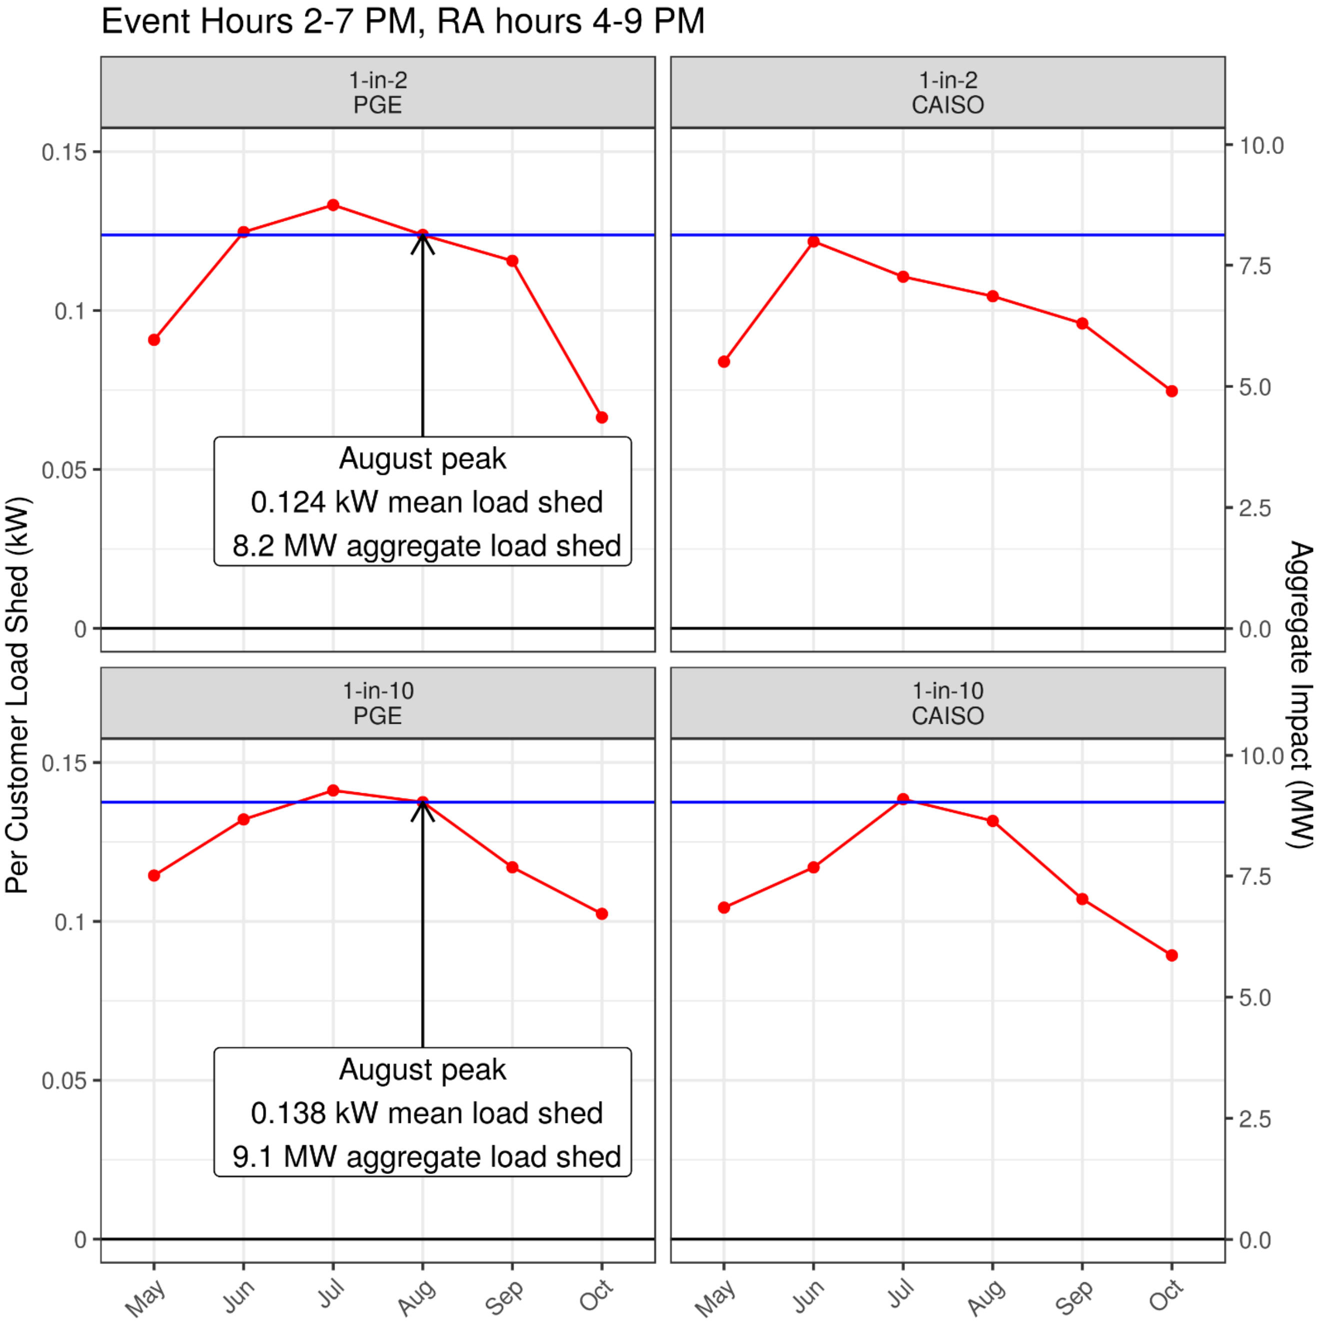

Figure note: The left and right axes apply to all four figures. Mean load shed is indicated on the left, and aggregate on the right, as demonstrated by the August examples in PG&E 1-in-2 and 1-in-10 with the blue line leading left and right.

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 18% 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	***0.92***	***0.84***	***1.74***	***0.92***	***0.84***	***1.74***
February Peak	***0.92***	***0.84***	***1.74***	***0.92***	***0.84***	***1.74***
March Peak	***0.92***	***0.84***	***1.74***	***0.92***	***0.84***	***1.74***
April Peak	***2.64***	***0.95***	***3.60***	***2.64***	***0.95***	***3.60***
May Peak	***4.55***	***1.47***	***6.04***	***4.55***	***0.29***	***4.85***
June Peak	5.96	2.33	8.30	5.96	0.47	6.43
July Peak	6.32	2.54	8.86	6.32	0.51	6.83
August Peak	5.91	2.32	8.24	5.91	0.46	6.37
September Peak	5.52	2.17	7.69	5.52	0.43	5.96
October Peak	***3.25***	***1.16***	***4.41***	***3.25***	***0.23***	***3.48***
November Peak	***0.92***	***0.84***	***1.74***	***0.92***	***0.84***	***1.74***
December Peak	***0.92***	***0.84***	***1.74***	***0.92***	***0.84***	***1.74***

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

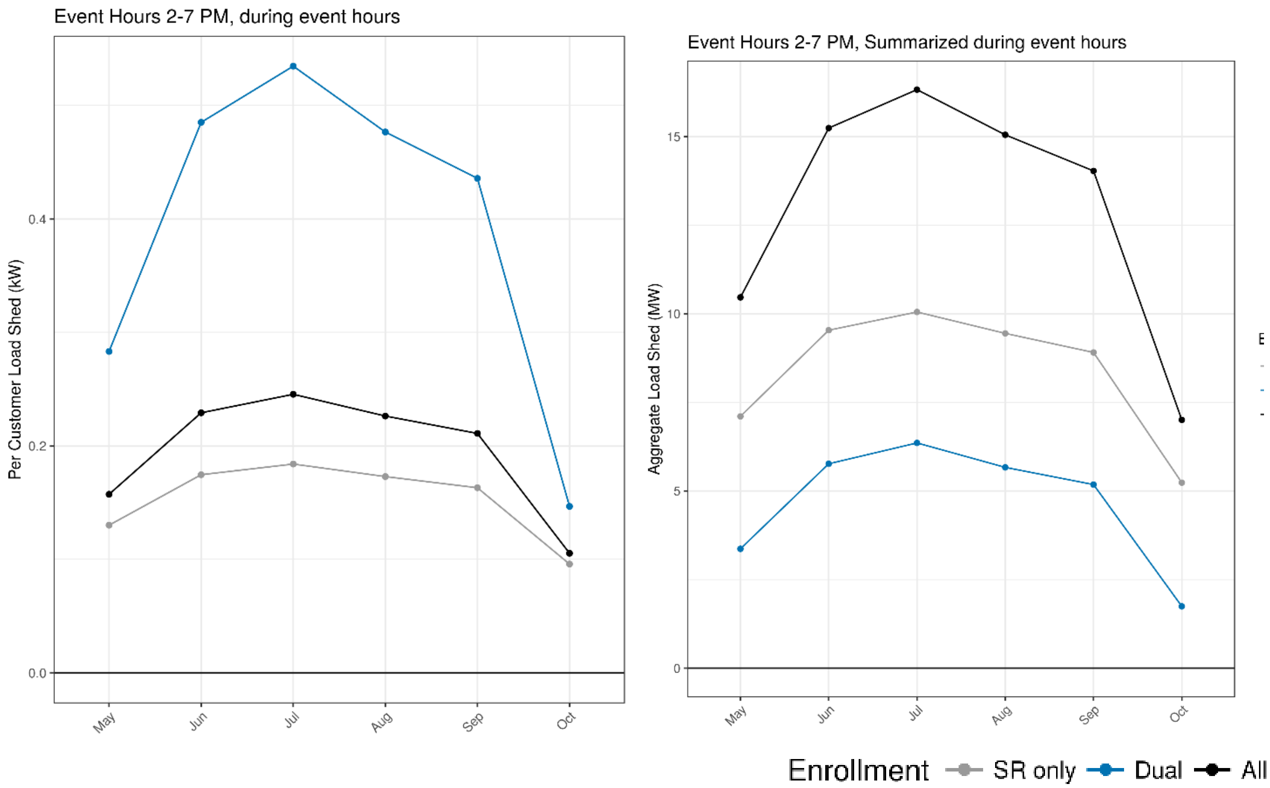

A more detailed discussion of the effects of dually enrolled customers can be found in the section comparing across years, see the section at the end of this chapter.

Ex Ante Load Reduction by LCA

The per-customer ex ante impacts vary across the local capacity areas. Figure 28 shows mean ex ante load shed by LCA for each of the weather years using the 4-9 pm RA window.

Figure 28. Mean impact per customer by LCA (4-9 RA window)

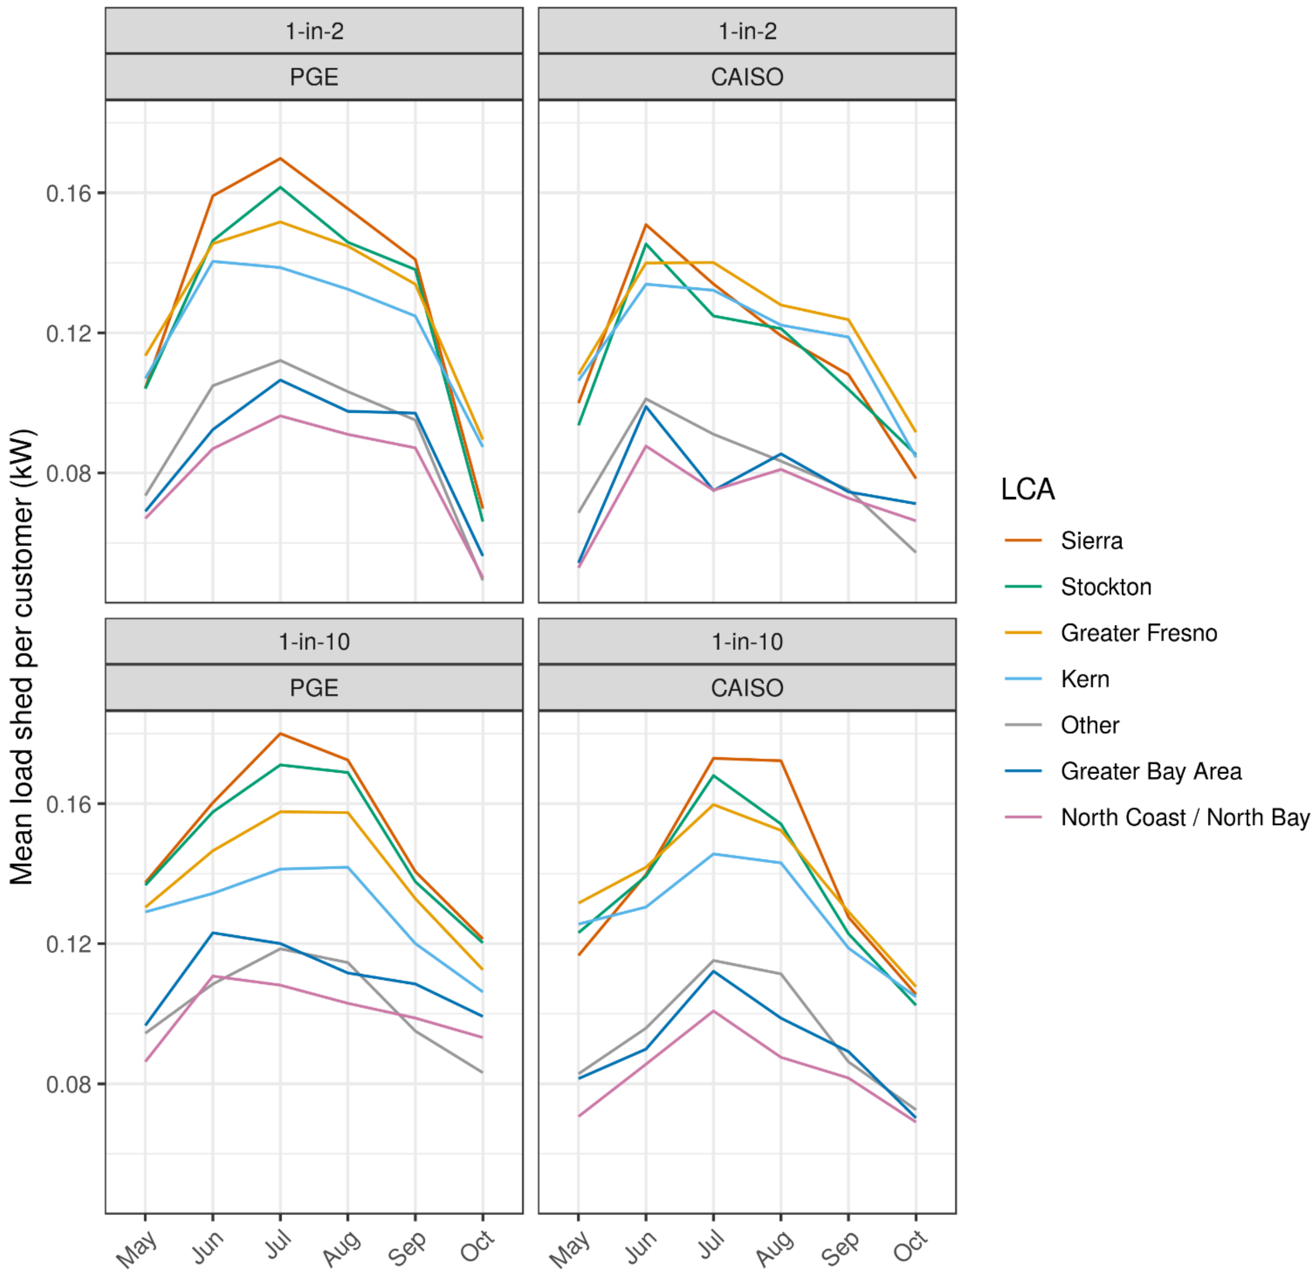

In the figure above, Sierra, Stockton and Greater Fresno have the highest load shed per customer while the Greater Bay Area and the Northern Coast are among the lowest.

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 29 (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 29 (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 North Coast/North Bay contributes little to the program compared to most other LCAs due to low enrollment and relatively poor load shed per customer (primarily due to mild weather).

Figure 29. (a) Forecasted enrollment and (b) Aggregate ex ante impacts by LCA (4-9 RA window)

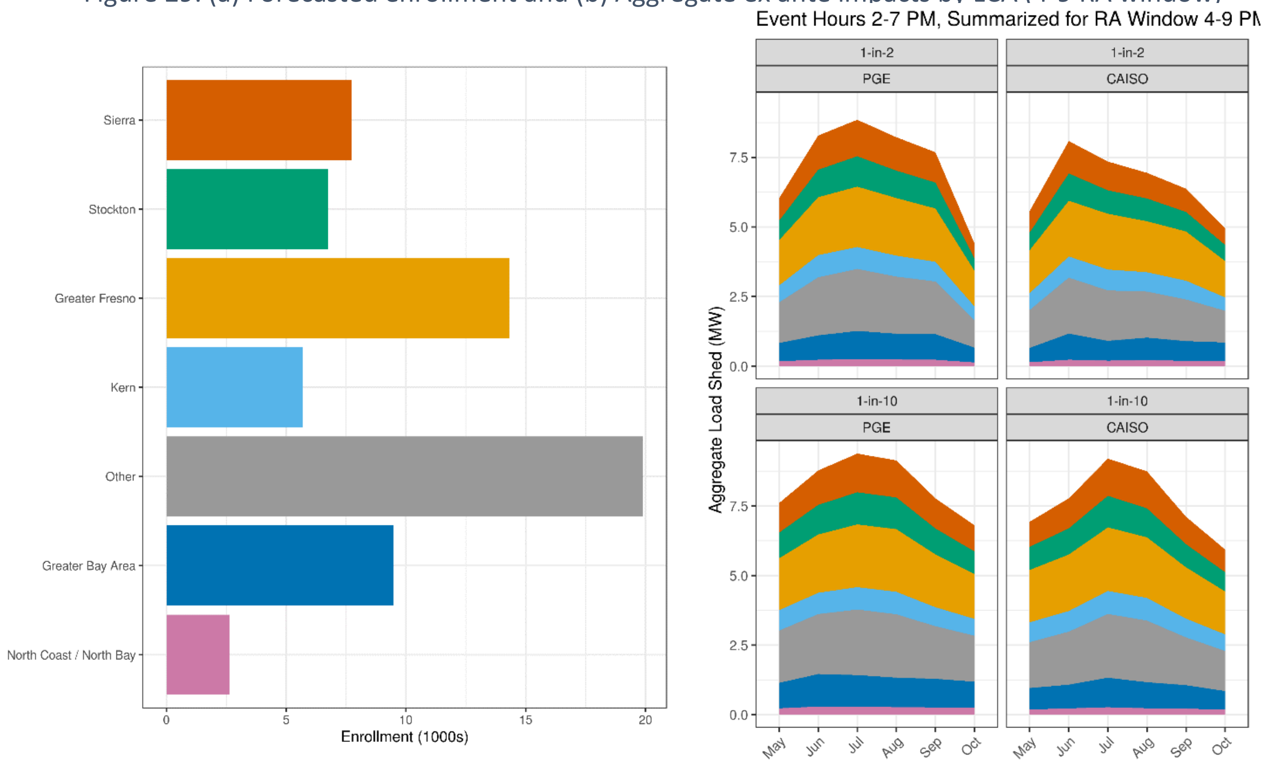

Table 14. Aggregate ex ante impacts by LCA in the PG&E 1-in-2 Weather

LCA	SmartRate™ only	Dually enrolled	Total
All	5.91	2.32	8.24
Sierra	0.83	0.37	1.20
Stockton	0.64	0.34	0.98
Greater Fresno	1.61	0.46	2.07
Kern	0.62	0.14	0.75
Other	1.41	0.64	2.05
Greater Bay Area	0.61	0.32	0.92
North Coast / North Bay	0.19	0.05	0.2
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 2021 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 2019 (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	2019 ex post	66,504	14.88	138.58	0.22	2.08	10.74	96.7
	Current ex ante	66,519	15.11	134.06	0.23	2.02	11.27	99.2
SmartRate ™ only	2019 ex post	54,272	9.36	112.87	0.17	2.08	8.29	96.5
	Current ex ante	54,634	9.45	108.93	0.17	1.99	8.67	99.0
Dually enrolled	2019 ex post	12,232	5.57	25.61	0.46	2.09	21.73	98.1
	Current ex ante	11,885	5.66	25.13	0.48	2.11	22.54	100.0

Table note: The 'enrollment' figures for 2019 ex post show the mean number of participants, not the mean numb-er of enrolled customers. 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	2019 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 20-70% higher in the event window than the RA window, depending on month. This applies to both ex post and ex ante calculations.
Weather	Event days in 2019 were 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 the per-LCA aggregate impacts for each of the modeling steps between ex post and ex ante.

Figure 30: Relationship between Ex Post and Ex Ante Load Shed during the 4-9 PM Resource Adequacy Window

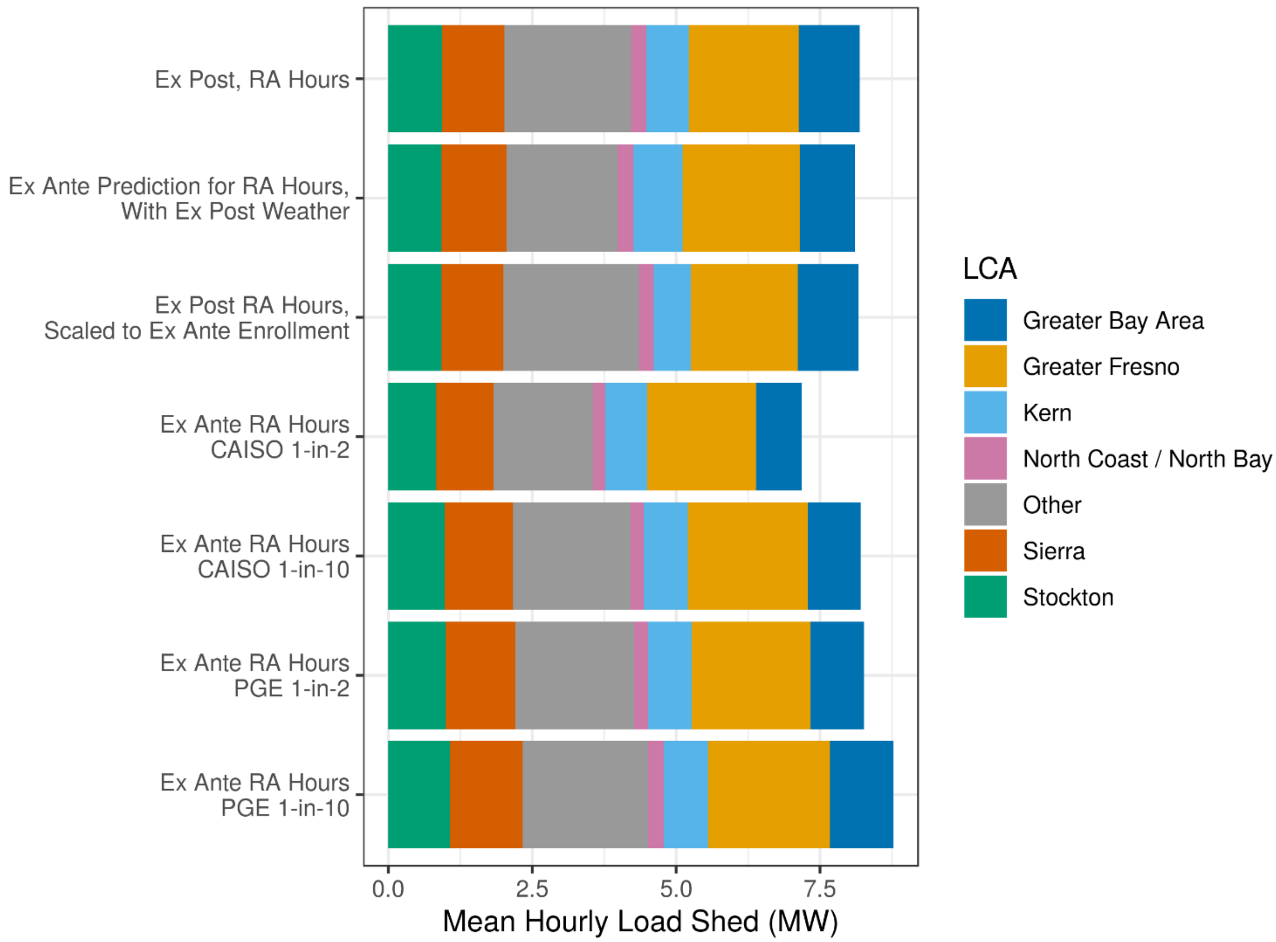

Figure 30 shows the relationship between ex post and ex ante predictions, with nearly all of the variation coming from the weather assumptions from the four scenarios. The top bar shows the mean ex post load impact for the 2019 events. The ex ante model was used to predict the load shed by LCA, given the weather from the same event days summarized in the first bar; the resulting in a just slightly lower prediction summarized in the second bar.

Summertime program enrollment in 2020 and beyond is projected to be relatively even with 2019 participation. To generate the third bar, the ex post load impact from 2019 was scaled to the 2020 numbers, separately within each LCA and for each dual-enrollment status.

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 estimate of 8.23 MW is lower than the prior year's ex ante estimate of 8.61 MW (see Table 17 and Figure 31.) Almost the entire reason for the difference between this year's ex ante predictions and last year's is around the changing mix of SmartRate™ only and dually enrolled customers -- and the performance of these groups. As shown in the table below, the prior estimates assumed that there would be 14,826 dually enrolled customers, which have higher impacts, while this year's enrollment estimate for that group is lower (11,885). These dually enrolled customers are expected to perform better than SmartRate™ only customers (0.20 compared to 0.11 kW per participant.)

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 2020+	67,206	8.61	134.95	0.13	2.01	6.38	93
	Current ex ante	66,519	8.23	135.78	0.12	2.04	6.06	96
All (Portfolio)	Prior ex ante 2020+	67,206	6.21	134.95	0.09	2.01	4.60	93
	Current ex ante	66519	6.37	135.78	0.10	2.04	4.69	96
SmartRate™ only (Program/ Portfolio)	Prior ex ante 2020+	52,380	5.61	103.57	0.11	1.98	5.42	92
	Current ex ante	54,634	5.91	110.45	0.11	2.02	5.35	96
Dually enrolled (Program)	Prior ex ante 2020+	14,826	2.99	31.38	0.20	2.12	9.54	95
	Current ex ante	11,885	2.32	25.33	0.20	2.13	9.17	96
Dually enrolled (Portfolio)	Prior ex ante 2020+	14,826	0.60	31.38	0.04	2.12	1.91	95
	Current ex ante	11,885	0.46	25.33	0.04	2.13	1.83	96

Table note: References loads are taken as reported for the 2019 group in the report. No other information was available at the time of this report. Humboldt is not included.

Figure 31 shows aggregate ex ante predicted load shed for dually enrolled, SmartRate™ only and All customers in 2020 that were made as part of evaluating the 2018 SmartRate™ program (i.e., "prior ex ante estimates" shown in red), along with the ex ante predictions made for 2020 in the present report (PY2019 in blue). SmartRate™ only provided slightly more aggregate impacts, while the dually enrolled participants in the 2019 ex ante estimates provided less aggregate impact than expected. The third panel below shows that overall, 2019 estimates are lower than past estimates.

Figure 31. Aggregate Ex Ante Impact Estimates: Comparison Across Report Years (PG&E 1-in-2)

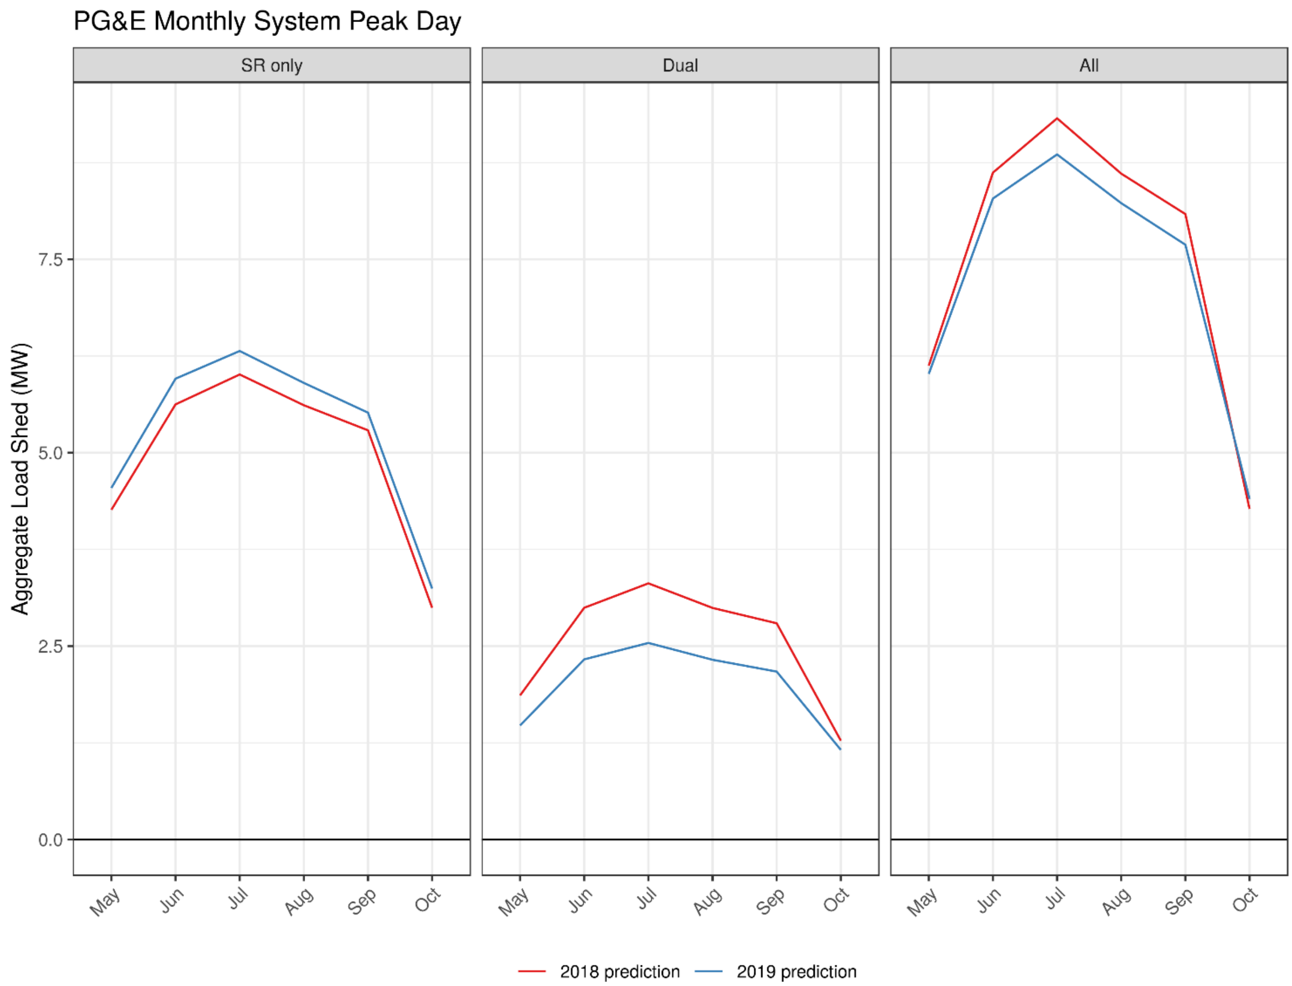

The forecasted aggregate load shed from the 2019 program year is somewhat lower than that from the 2018 program year, but this is almost entirely due to changes in the forecasted numbers of SR-only and dual customers: as discussed above, the total number of customers in the forecast is almost the same (66,619, compared to the 2018 estimate of 67,206) but the number of these customers who are dual-enrolled is now projected to be lower, 11,885 instead of 14,826. Meanwhile, the forecasted number of SR-only customers is expected to increase. Figure 32 shows (in green) the ex ante prediction that would have been obtained if the program year 2019 analysis were to use the 2018 enrollment forecast rather than the actual 2019 enrollment forecast. In that case the 2018 and 2019 ex ante predictions would have been almost exactly the same. That is, at this program-wide level of summarization, essentially all of the difference between ex ante aggregate load shed forecasts is due to changes in the enrollment projections rather than changes in per-participant performance.

Figure 32. Effect of Changing Enrollment

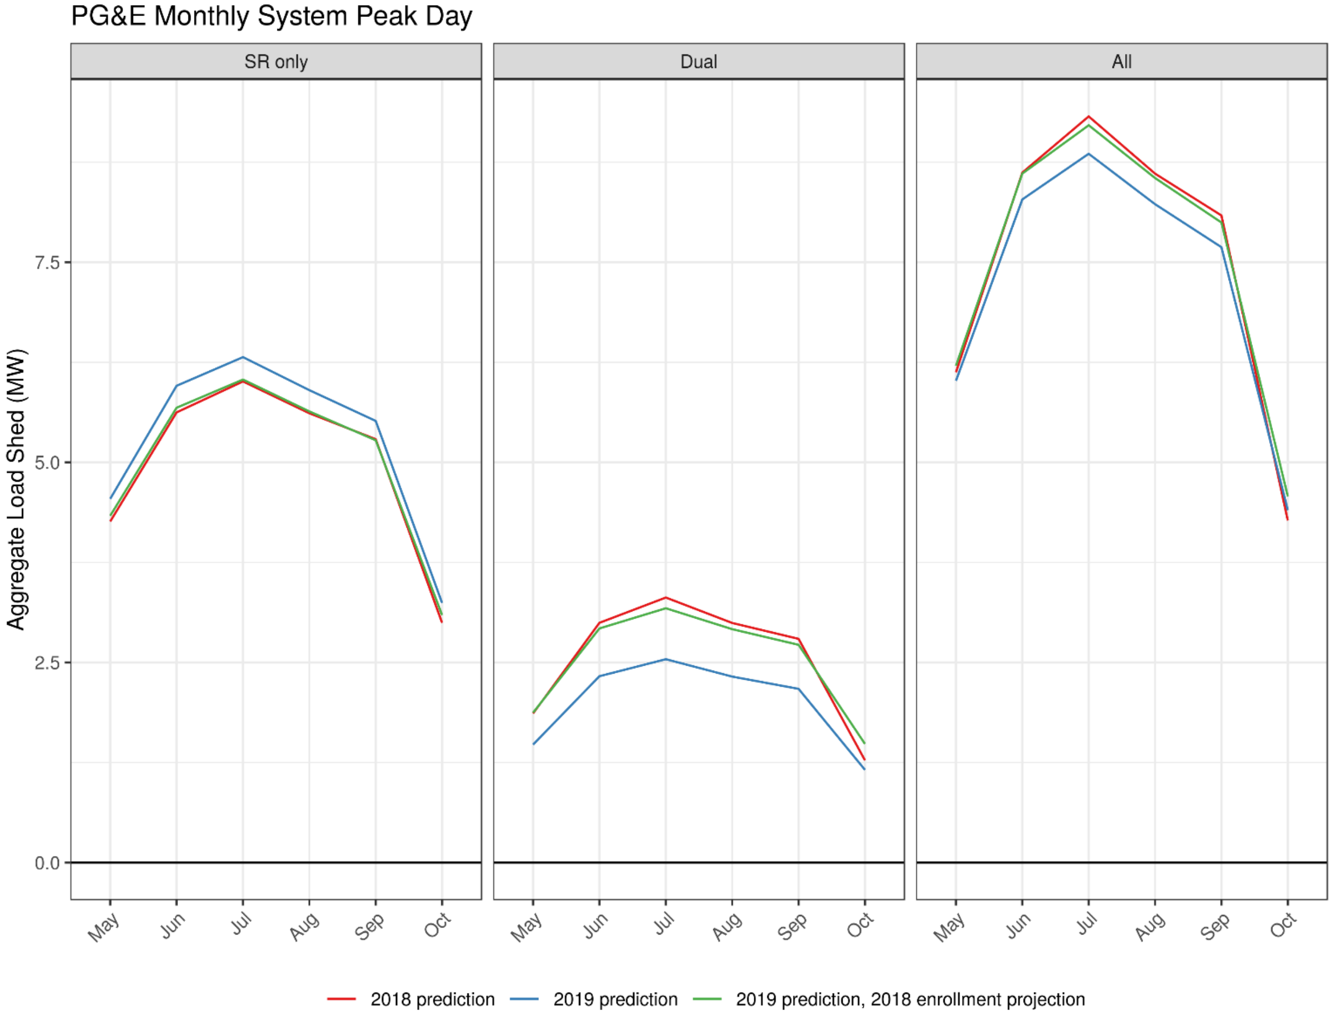

This year's model is discussed in detail in Appendix B.

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. The highest estimated aggregate impact occurs on the July peak day in every weather scenario except the CAISO 1-in-2 where there is a June peak.

Table 18: SmartRate™ ex ante load impact estimates by weather year and day type (RA 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**	0.12	0.26	8.27	17.03
		January Peak	0.03	0.03	1.74	1.94
		February Peak	0.03	0.03	1.74	1.94
		March Peak	0.03	0.03	1.74	1.94
		April Peak	0.05	0.10	3.60	6.35
		May Peak	0.09	0.18	6.04	12.05
		June Peak	0.12	0.26	8.30	17.05
		July Peak	0.13	0.27	8.86	18.10
		August Peak	0.12	0.26	8.24	16.99
		September Peak	0.12	0.24	7.69	15.97
		October Peak	0.07	0.13	4.41	8.52
		November Peak	0.03	0.03	1.74	1.94
		December Peak	0.03	0.03	1.74	1.94
PG&E	**1-in-10**	**Typical Event Day**	0.13	0.27	8.78	18.08
		January Peak	0.03	0.03	1.74	1.94
		February Peak	0.03	0.03	1.74	1.94
		March Peak	0.03	0.03	1.74	1.94
		April Peak	0.09	0.18	5.90	11.90
		May Peak	0.11	0.24	7.61	15.68
		June Peak	0.13	0.27	8.79	17.88
		July Peak	0.14	0.29	9.39	19.42
		August Peak	0.14	0.28	9.15	18.86
		September Peak	0.12	0.24	7.78	16.15
		October Peak	0.10	0.21	6.81	14.18
		November Peak	0.03	0.03	1.74	1.94
		December Peak	0.03	0.03	1.74	1.94
		Continued next pg.				

			Per Customer (kW)		Aggregate (MW)	
		Day type	**Mean Hourly Impact**	**Max Hourly I mpact**	**Mean Hourly Impact**	**Max Hourly Impact**
CAISO	**1-in-2**	**Typical Event Day**	0.11	0.22	7.20	14.66
		January Peak	0.03	0.03	1.74	1.94
		February Peak	0.03	0.03	1.74	1.94
		March Peak	0.03	0.03	1.74	1.94
		April Peak	0.06	0.11	3.93	7.21
		May Peak	0.08	0.17	5.58	11.19
		June Peak	0.12	0.25	8.10	16.54
		July Peak	0.11	0.23	7.36	15.02
		August Peak	0.10	0.21	6.95	14.11
		September Peak	0.10	0.20	6.38	12.98
		October Peak	0.07	0.15	4.96	9.90
		November Peak	0.03	0.03	1.75	1.94
		December Peak	0.03	0.03	1.74	1.94
CAISO	**1-in-10**	**Typical Event Day**	0.12	0.25	8.22	16.90
		January Peak	0.03	0.03	1.74	1.94
		February Peak	0.03	0.03	1.74	1.94
		March Peak	0.03	0.05	2.20	3.03
		April Peak	0.09	0.17	5.71	11.44
		May Peak	0.10	0.21	6.94	14.10
		June Peak	0.12	0.24	7.78	15.87
		July Peak	0.14	0.28	9.21	18.90
		August Peak	0.13	0.27	8.75	18.07
		September Peak	0.11	0.22	7.12	14.76
		October Peak	0.09	0.19	5.94	12.32
		November Peak	0.03	0.03	1.78	1.97
		December Peak	0.03	0.03	1.74	1.94

A Deeper Dive to Inform the Future

In this section, we provide an investigation that can help inform the program, as well as future forecasts. This section includes an assessment of TOU customer impacts.

TOU Analysis – Exploration of TOU Impacts

The CDA team 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 rates. For this analysis, the rates TOU-A, TOU-B, and E6 are classified as "TOU-opt-in" rates because they are all time of use rates that customers have opted into. Note that TOU-A has halted enrollment and will be discontinued and TOU-B will stop accepting new enrollment – both before the end of 2020. TOU-C is classified as a "TOU-default" rate because it will become the default residential rate. TOU-C was first associated with a 2018 pilot study where customers were defaulted into the rate, but it was available for voluntary enrollment in 2019. 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-C is currently slated to become the default choice for customers in Q4 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

In 2019, about 23% of SmartRate™ participants were on TOU rates, with 18% on a TOU-opt-in rate and just under 5% on the TOU-default 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 Fresno represents 22% of program enrollment, but just 14% of TOU-default customers and 15% of TOU-opt-in 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 2019-09-13.

LCA	All	All %	non-TOU	non-TOU %	TOU-default	TOU default %	TOU-opt-in	TOU opt-in %
Other	18,480	28	13785	27	1170	37	3583	30
North Coast / North Bay	2,808	4	1609	3	213	7	992	8
Kern	6,516	10	5743	11	197	6	602	5
Stockton	6,776	10	5542	11	290	9	969	8
Sierra	7,754	12	5625	11	510	16	1651	14
Greater Fresno	14,654	22	12436	24	454	14	1801	15
Greater Bay Area	9,480	14	6723	13	303	10	2471	20
All	66,528	100	51483	100	3138	100	12108	100

Figure 33 displays 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 33. (a) TOU enrollment category breakdown for each LCA and (b) Enrollment fraction by LCA within each TOU category

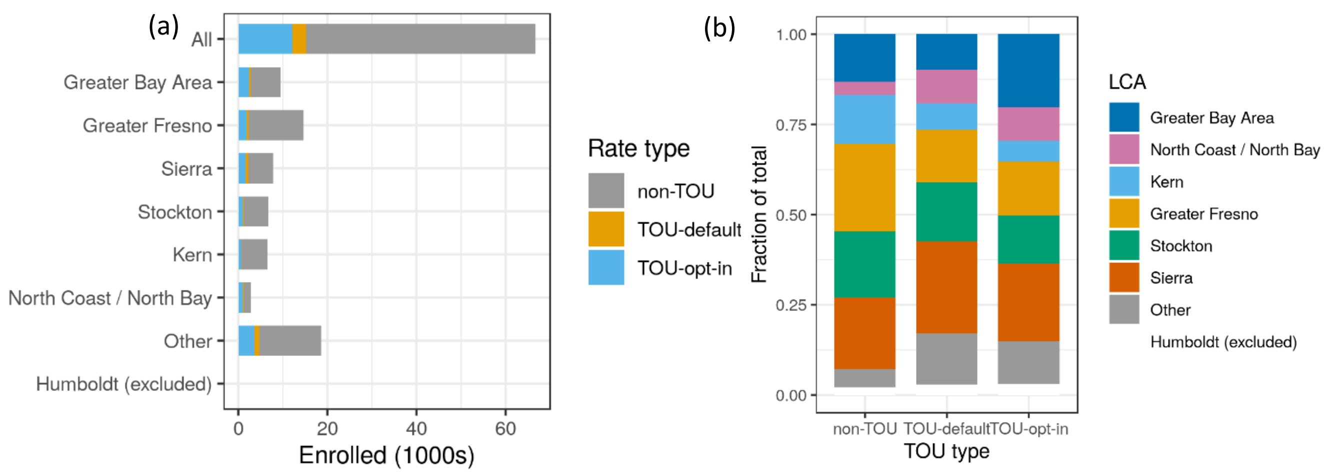

Panel (a) illustrates that the fraction of TOU enrolled customers is small (about 23%) compared to total program enrollment and TOU-default enrollment is a small portion of all the TOU rates (about 5% of total enrollment). Panel (b) illustrates that the breakdown by LCA of each TOU rate category varies. 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 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 customers in the program is small overall—roughly 5% —they are the leading edge of a territory-wide roll out of TOU-default 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™.

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 34 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 2019.

Figure 34. Typical event day aggregate, per-customer, and % of reference load shed by TOU rate category.

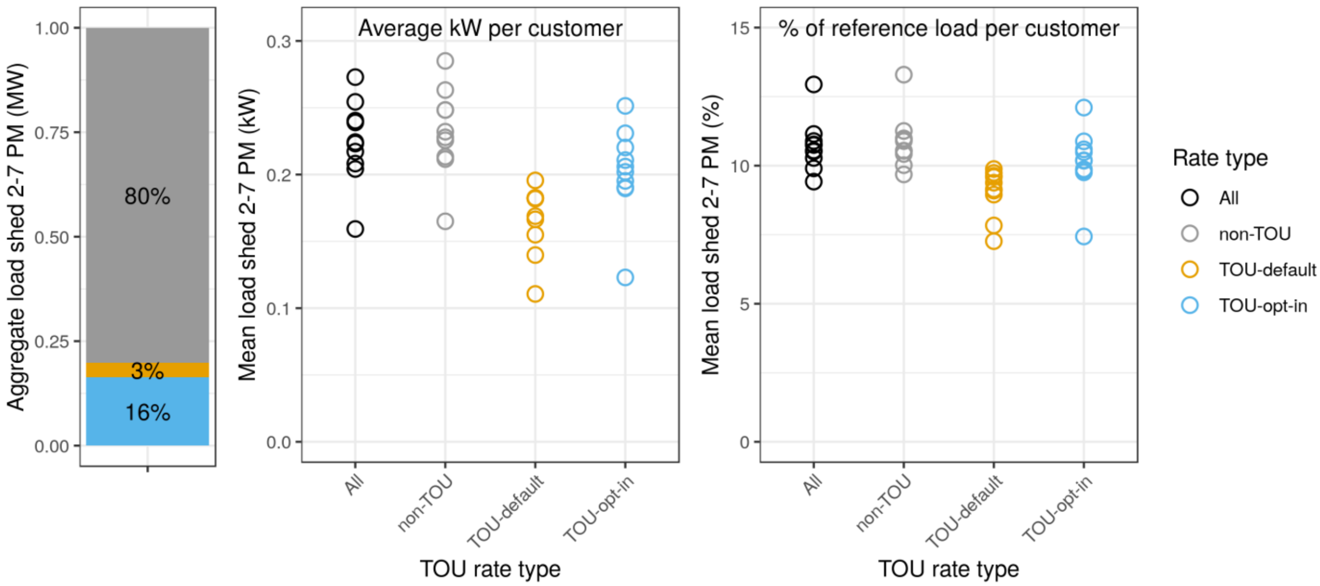

The figure shows that most aggregate program impacts are currently coming from non-TOU customers (80%); however, in the coming years the share of customers on TOU rates is expected to grow quickly. The center and right-hand panel show that the range of outcomes across event days overlaps across some of the rate types, but per-customer load impacts for TOU-default are somewhat smaller than non-TOU. They are also smaller than TOU-opt-in.

Table 20 summarizes the results for the typical 2019 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**	66,504	97	2.16	0.22	14.89	143.75	10.4%
TOU	**non-TOU**	51,664	97	2.13	0.23	11.99	110.11	10.9%
	TOU-default	3,107	95	1.83	0.17	0.52	5.69	9.1%
	TOU-opt-in	12,112	95	1.98	0.20	2.45	24.00	10.2%

Due to the geographic effects, the average event temperature varies between those on TOU rates and those not on TOU rates; however, in PY2019, both groups of TOU customers experienced the same average event temperature (95 degrees).Per-customer impacts are lower among TOU customers than those not on TOU rates, with savings among the TOU-default customers lagging savings from TOU-opt-in customers. TOU-default customers also have the lowest value of load impacts as a percentage of reference loads (Impact %). Note that the relative magnitude of the per customer impacts among the TOU customers changed in PY2019 (compared to PY2018 when default per customer impacts were higher). This is the result of shifting populations (i.e., fewer Greater Bay Area customers near coastal areas), higher temperatures, and the resulting higher reference loads.

Figure 35 plots per-customer reference load vs. load impact (i.e. same y-axis as the middle panel from the figure 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 increase fastest with increases in reference loads (i.e. greatest slope of the linear fit). The TOU-default customers have a flatter load impact response to increasing reference loads. Notably, there is a point between 1.75 kW and 2.0 kW where the two lines cross and opt-in per customer load shed is expected to be higher than default per customer impacts. The average temperatures and reference loads were higher in PY2019 than in PY2018, leading to higher per customer impacts for opt-in TOU customers than default TOU customers.

Figure 35. Event load shed vs reference load by TOU group

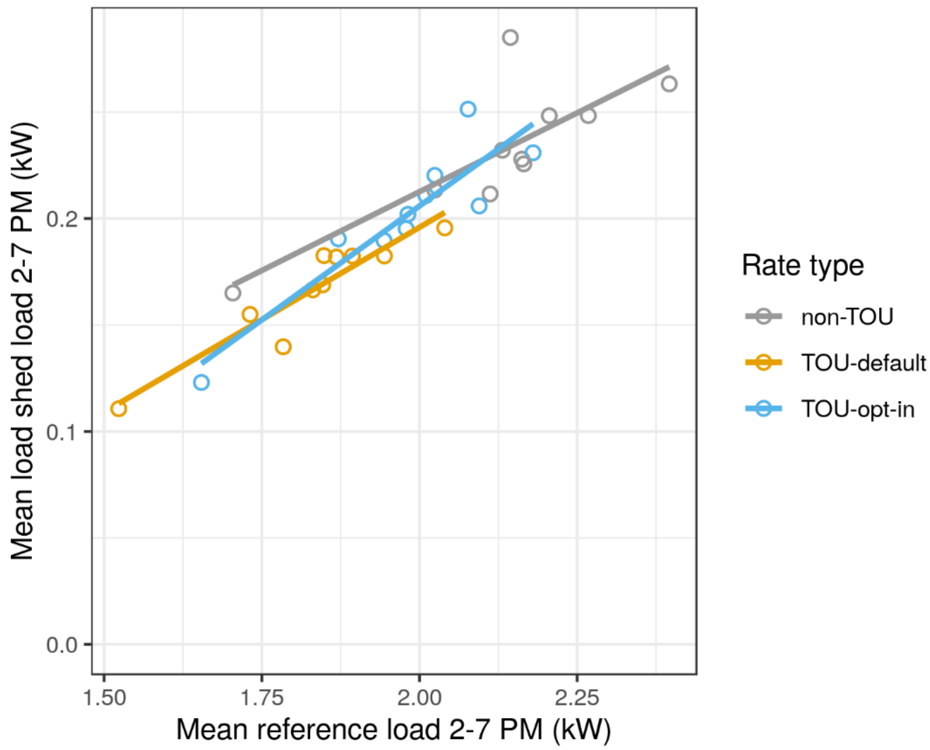

TOU Summary – A Look Towards the Future

The composition of TOU customers within the SmartRate™ program is expected to continue to shift in the future. As such, this could increase or decrease the per household savings that PG&E will see from the TOU groups.

Other factors may also change the savings that PG&E sees. As future event hours become more aligned (i.e., the CPUC-approved updated SmartRate™ hours slated to take effect in 2021 fall into all TOU rate peak periods), PG&E would likely see greater savings (due to the reduced burden/reduced number of hours).

Furthermore, it seems likely that the newer cohorts of TOU-default customers will continue to arrive 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 to eligible customers (with first year bill protection). It would be logical for 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 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 2019 SmartRate™ program delivered 14.89 MW of capacity for a typical event, down from prior year's 17.83 MW. This change was due primarily to the shifting composition of participants, especially the reduction in enrolled customers in the Greater Bay Area caused by customer enrollment in CCAs.

Ex ante forecasts have also decreased since last year, reflecting the same shifting composition of participants that influenced the 2019 load shed. The PG&E 1-in-2 August peak load reduction forecast is 8.2 MW in 2019 over the RA window, down from 8.6 MW in the prior year's ex ante forecast.

Based on our evaluation of the 2019 SmartRate™ program, we provide the following recommendations:

Recommendation	Description
Plan the future of SmartRate in accordance with projections for new CCAs	When a CCA comes online in a customer's location, they are enrolled in the CCA by default and de-enrolled from SmartRate as a result. Those de-enrollments have reduced the customer base of SmartRate from 124k customers in 2017 to 66k customers in 2019. Where CCAs are not triggering de-enrollments, the number of customers enrolled in SmartRate is actually increasing each year, so the future of the SmartRate resource is likely to be determined by future CCA activity. At some point, if enrollment continues its dramatic decline, it will be necessary to determine whether and how the program can continue operating.
Try to mitigate the resource impacts of CCA de-enrollments	If customers de-enrolled through their transition to CCAs are not enrolled in a similar critical peak pricing program under their CCA, these self-identified customers willing to contribute to grid stability are being lost as a grid resource at a time of increasing value for load flexibility. This policy gap with unintended consequences for the DR resource should be addressed.
Embrace behavioral savings from households with lower per-customer impacts to account for fewer dually enrolled customers and target those with higher potential for savings	With the discontinuation of dual enrollment between SmartAC and SmartRate, SmartRate's event impacts will shift towards the largely behavioral impact of the group of customers currently classified as "SmartRate only". The ultimate value of a DR program lies in its aggregate impacts. To preserve and grow its aggregate impacts, SmartRate may need to pursue increased enrollments to offset smaller per-customer impacts and those enrollments may need to be targeted at customers likely to have higher potential for savings.
Leverage the transition to TOU rates to boost enrollment; prepare for a lot of novices	The only mechanism of outreach to enroll SmartRate customers is PG&E's rate comparison tool. On its own and without marketing expenditures, the rate comparison tool recruited more than enough customers between the end of PY2018 and PY2019 in every LCA other than the Greater Bay Area (i.e. without CCA de-enrollments) to produce a net increase in enrollment. It stands to reason that the advent of default TOU rates will drive traffic to the rate comparison tool and that surge in interest represents an opportunity to grow and re-configure the enrolled customer base of the program.
Shift event times to better align with TOU peak periods and the RA window	Because SmartRate is implemented through rate structures, its details need to be hashed out and approved via the rate approval process. At the moment, the timing of the program event window is out of sync with both TOU rates and the current RA window. An update to the timing has

Recommendation	Description
	been approved by the CPUC, but has not yet taken effect. The ex ante RA impacts are based on an average of observed impacts across the RA window, so the update to the SmartRate event window will substantially increase the RA average and bring the program incentives into alignment with current grid flexibility needs.
Prepare to evaluate impacts among TOU customers and account for rate type in choosing potential controls	With default TOU rates on the horizon, it will become critical to the evaluation of the SmartRate program that the ambient impacts of TOU rates are separable from the impacts of the peak rates on SmartDays. The synthetic control approach to event evaluation should be capable of differentiating those impacts, but perhaps only if TOU rate status is taking into account in the sampling and matching criteria.

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 beginning in Q4 of 2020. Marketing and targeting strategies (and participant restrictions) influence who enrolls in each cohort. Our PY2018 findings showed 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. All of the cohort-level differences strongly suggest that recruitment strategies strongly influence the character of the resource over time.

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

SmartRate™ only

Date	Number enrolled	Per-cust. ref. load (kW)	Per-cust. load impact (kW)	Aggregate ref. load (MW)	Aggregate load impact (MW)	Impact as % of ref.	Average temperature (F)
2019-06-11	53,875	2.32	0.21	125.08	11.51	9	99
2019-07-24	54,196	2.20	0.18	119.42	9.62	8	97
2019-07-26	54,256	2.15	0.16	116.38	8.67	7	94
2019-08-13	54,270	1.98	0.16	107.51	8.68	8	95
2019-08-14	54,283	2.24	0.17	121.57	9.37	8	99
2019-08-16	54,302	2.37	0.19	128.92	10.16	8	98
2019-08-26	54,339	2.14	0.19	116.53	10.09	9	96
2019-08-27	54,376	2.13	0.17	115.56	9.32	8	96
2019-09-13	54,553	1.78	0.13	97.37	6.83	7	95
Typical event	54,272	2.15	0.17	116.48	9.36	8	97

Dually enrolled

Date	Number enrolled	Per-cust. ref. load (kW)	Per-cust. load impact (kW)	Aggregate ref. load (MW)	Aggregate load impact (MW)	Impact as % of ref.	Average temperature (F)
2019-06-11	12,633	2.44	0.54	30.88	6.78	22	101
2019-07-24	12,359	2.23	0.43	27.58	5.29	19	99
2019-07-26	12,328	2.14	0.41	26.36	5.11	19	95
2019-08-13	12,214	2.02	0.40	24.64	4.92	20	96
2019-08-14	12,206	2.36	0.53	28.82	6.45	22	100
2019-08-16	12,172	2.56	0.56	31.19	6.81	22	100
2019-08-26	12,106	2.23	0.47	26.95	5.74	21	98
2019-08-27	12,091	2.19	0.44	26.45	5.34	20	97
2019-09-13	11,975	1.79	0.31	21.38	3.67	17	96
Typical event	12,231	2.22	0.46	27.14	5.57	21	98

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 express 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)

where the (T-70 F) term is set to zero if 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, depending on what combination of LCA, dual, and CARE status is being predicted. See the example below the table.

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.

In this table, '**:**' indicates multiplication. 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 the value of the 'dual' indicator variable times the 'Sierra' indicator variable has a value of 1 only for dual-enrolled customers in Sierra, and 0 otherwise, the dual:Sierra coefficient only applies to dual-enrolled customers in the Sierra LCA.

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

Parameter	**2019 Estimate**	**2019 Standard Error**	**2018 Estimate**	**2018 Standard Error**

first_hours	-0.027	0.004	-0.031	0.003
cool_lca	0.031	0.015	0.031	0.008
care	0.063	0.011	0.041	0.009
cdh70	0.093	0.001	0.097	0.001
cdh70:dual	0.108	0.003	0.069	0.003
cdh70:care	-0.071	0.004	-0.070	0.004
cdh70:dual:hour5	-0.068	0.003	-0.046	0.003
cdh70:cool_lca	-0.031	0.006	-0.046	0.005
dual:Other:noncare	-0.082	0.013	0.072	0.010
dual:Stockton	-0.322	0.080	-0.701	0.077
dual:Sierra	-0.632	0.084	-0.373	0.066
cdh70:care:dual:Fresno	0.005	0.006	-0.044	0.006
cdh70:care:dual:Kern	0.011	0.005	-0.085	0.004
cdh70:dual:Fresno	-0.017	0.005	0.017	0.005
cdh70:care:dual	0.019	0.004	0.062	0.003
cdh70:dual:hour1	-0.025	0.003	-0.029	0.003
cdh70:care:dual:Sierra	-0.026	0.006	-0.046	0.006
cdh70:care:dual:NorthCoast	0.067	0.007	NA	NA
cdh70:dual:Stockton	0.117	0.029	0.301	0.029
cdh70:dual:Sierra	0.238	0.03	0.168	0.024
cdh70:Sierra:SRonly	NA	NA	0.015	0.002

The terms in the model used for the 2019 analysis is almost identical to that used for 2018; the only differences are the elimination of a coefficient of cdh70 that applied only to SmartRate-only customers in the Sierra LCA, and the addition of a cdh70 coefficient that applies only to dual, care customers in the North Coast / North Bay LCA. The model coefficients are reasonably stable across years in spite of the small number of events, which shows that stochastic variability ('noise') is small compared to the actual relationships being quantified.

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: -0.027

CARE: 0.063

Coefficients that multiply cooling degree-hours:

All: 0.093 (the coefficient denoted 'cdh70' applies to all customer segments)

CARE: -0.071

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 = (-0.027 + 0.063) + (0.093-0.071)/10*(90 F – 70 F) = 0.064 kW per customer. This is exactly the same (to three significant digits) as predicted from last year’s coefficients, although each individual coefficient was slightly different: last year’s calculation was

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

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	2019 Estimate	2019 Standard Error	2018 Estimate	2018 Standard Error
pred_final_hour:after1	0.113	0.025	0.087	0.023
pred_final_hour:after2	-0.146	0.024	-0.169	0.024
pred_final_hour:after3	-0.152	0.025	-0.160	0.022
pred_final_hour:after4	-0.125	0.023	-0.113	0.021
pred_final_hour:after5	-0.086	0.022	-0.080	0.019
pred_final_hour:dual:after1	-0.853	0.032	-0.658	0.030
pred_final_hour:dual:after2	-0.570	0.031	-0.389	0.030
pred_final_hour:dual:after3	-0.357	0.031	-0.203	0.027
pred_final_hour:dual:after4	-0.156	0.029	-0.076	0.026
pred_final_hour:dual:after5	-0.065	0.027	-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 in the third hour after the event = -0.152(final hour load shed) -0.357*(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.

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 36. Ex post load impact from three LCAs

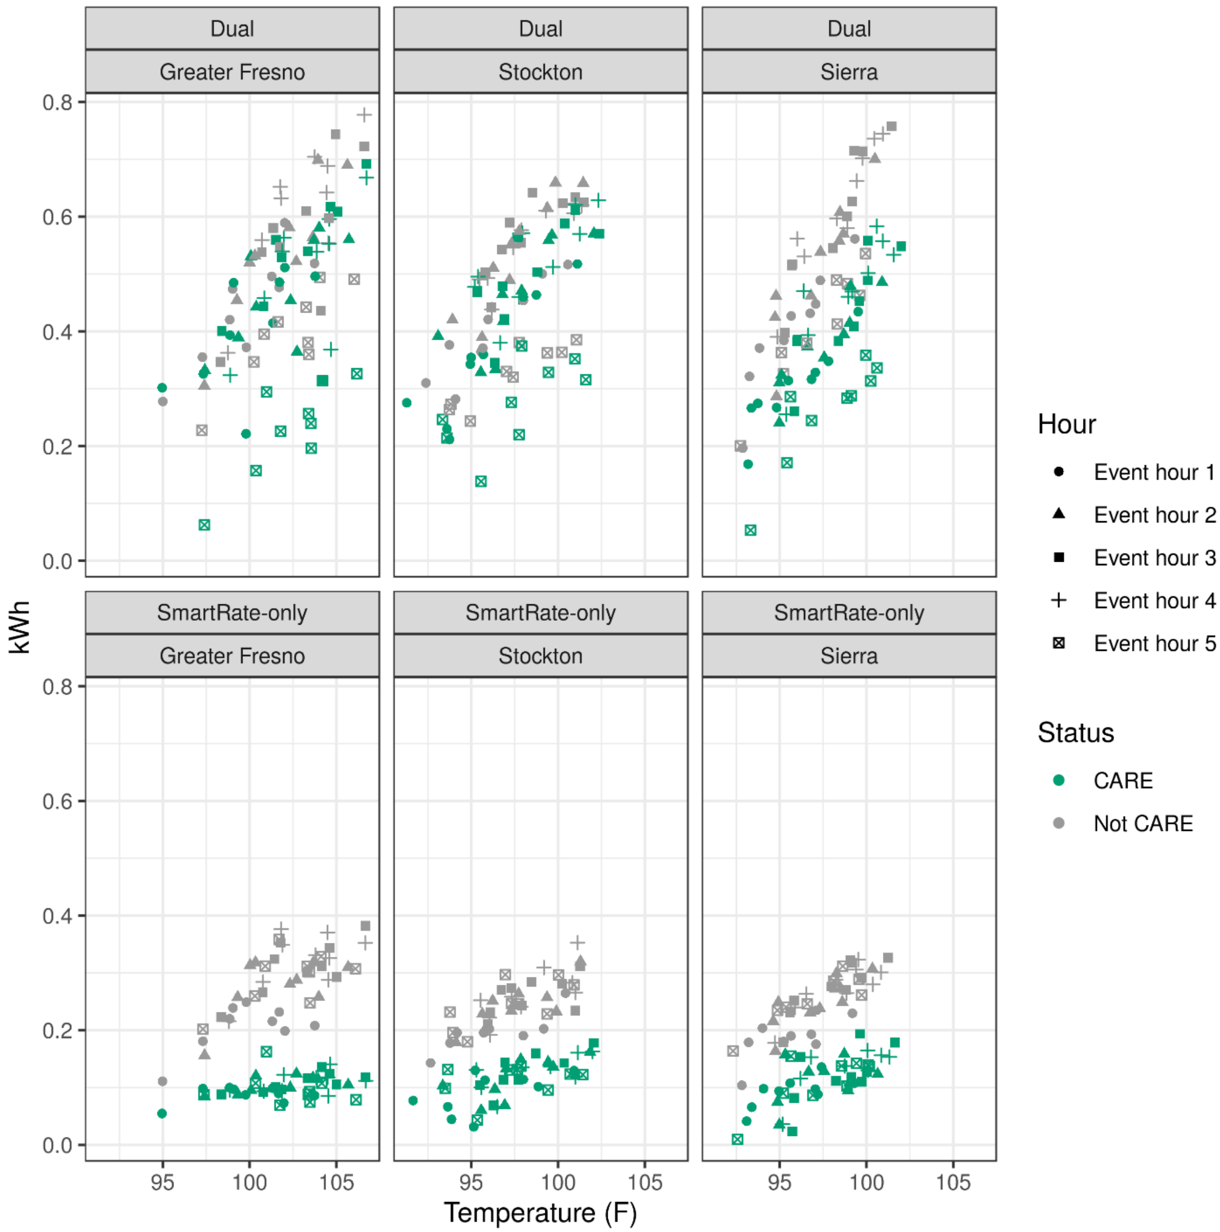

Figure 36 shows ex-post load impact from three LCAs and demonstrates the need for a model that treats customers differently by CARE and dual categories: CARE customers (green) have substantially different load shed at given temperature than non-care customers (gray); Dual-enrolled customers (top row) 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 37: Comparison of ex post load shed to predictions from the ex ante model.

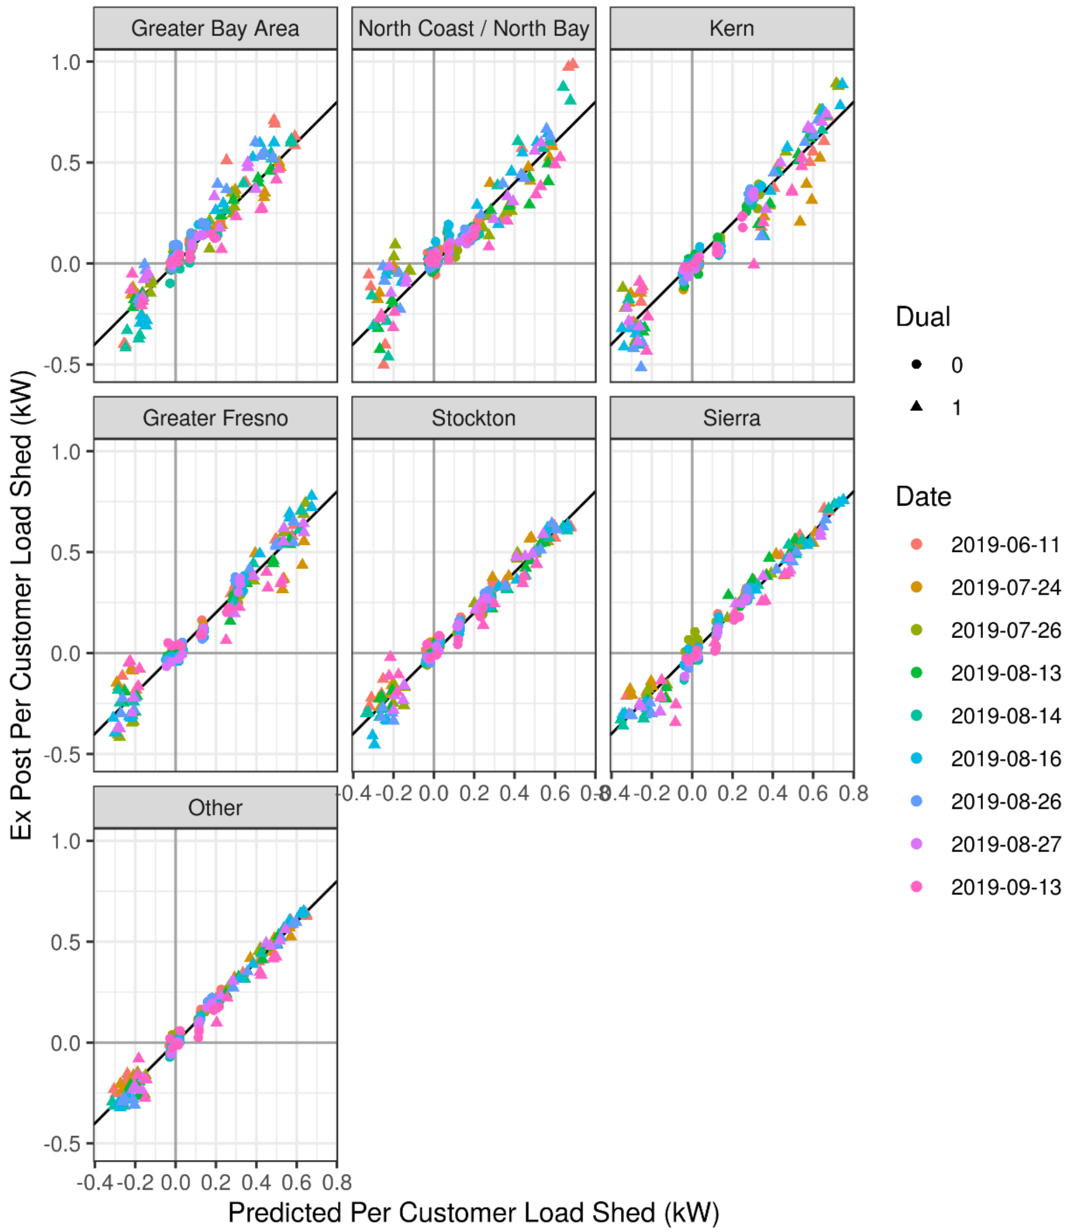

Figure 37 shows the ex post load shed vs the load shed that is predicted by the ex ante model, for the hours of the RA window (4-9 PM), for each event. Values below zero are from the snapback period (7-9 PM).

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 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: Bill Protection Analysis

PG&E provided billing data with program specific credits and charges for all SmartRate™ participants enrolled as of 2019-05-01. This included billing impacts for the period 2019-05-01 through 2019-11-01.

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 2019-05-01 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 66,257 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	74,179
SmartRate™ active during 2019 program year	74,179
Billing customers	71,784
Billing and SmartRate™ data merged	71,784
CARE, LCA, and Dual values present	71,230
3+ months of billing May - September 2019	66,257
Working sample	**66,257**

Average Savings and Refunds Across all Customers

Across all 66,257 customers in our analysis, the average participant saved $30.39. SmartRate™ only participants averaged slightly bigger reductions in their bill than dually enrolled participants. See Table 24.

Table 24. Bill impacts for all customers. Dollar values are the total for the summer of 2019. Numbers discussed in the text are highlighted.

Enrollment Status	Impact	Count of Participants	% of Participants	Avg. Bill Change
SR only	Increased Bill	4,471	8%	$10.75
	Decreased Bill	49,275	92%	-$41.25
	All	53,746	100%	-$30.50
Dual	Increased Bill	1,274	10%	$9.71
	Decreased Bill	11,237	90%	-$39.40

	All	12,511	100%	-$29.69
All	Increased Bill	5,745	9%	$10.52
	Decreased Bill	60,512	91%	-$40.91
	All	66,257	100%	-$30.39

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 2018-05-01, i.e. after the start of the prior year's summer season, as protected.

Overall, 22% of all participants fell into the bill protection status during the 2019 SmartRate™ season. (See Table 25Table 26.) This is slightly lower than the number on bill protection in 2018 (22% compared to 25%, 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	39,117	73%	59%
	Protected	14,629	27%	22%
	All	53,746	100%	81%
Dual	Unprotected	12,343	99%	19%
	Protected	168	1%	0%
	All	12,511	100%	19%
All	Unprotected	51,460	78%	78%
	Protected	14,797	22%	22%
	All	66,257	100%	100%

Of those who were eligible for bill protection (14,797), 93% of bill protected participants experienced bill reductions (shown as a negative bill change in Table 26). The average *bill reduction total across the summer of 2019* was $41.85, but the small number of dually enrolled customers who experienced a bill reduction saw an even larger reduction in their bill ($53.67 for dually enrolled customers compared to $41.71 for SmartRate™ only customers).

Among the 7% of participants that saw their bill increase, it increased by an average of $9.44. 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 average monthly impacts for the summer of 2019. Numbers discussed in the text are highlighted.

Enrollment Status	Impact	Count of Participants	% of Participants	Avg. Bill Change
SR only	Increased Bill	1,037	7%	$9.45
	Decreased Bill	13,592	93%	-$41.71
	All	14,629	100%	-$32.25
Dual	Increased Bill	8	5%	$8.45
	Decreased Bill	160	95%	-$53.67
	All	168	100%	-$45.21
All	Increased Bill	1,045	7%	$9.44
	Decreased Bill	13,752	93%	-$41.85
	All	14,797	100%	-$32.40

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 (51,460). Among this group, 91% experienced a decrease, and the average decrease was roughly the same as those on bill protection ($40.63 compared to $41.85).

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	3,434	9%	$11.14
	Decreased Bill	35,683	91%	-$41.08
	All	39,117	100%	-$29.94
Dual	Increased Bill	1,266	10%	$9.72
	Decreased Bill	11,077	90%	-$39.19
	All	12,343	100%	-$29.48
All	Increased Bill	4,700	9%	$10.76
	Decreased Bill	46,760	91%	-$40.63
	All	51,460	100%	-$29.88

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 71,230 customers with billing data available for 2019. Of the participants who were bill protected and participated in any portion of the summer of 2019 (18,304 SmartRate™ only and 207 dually enrolled participants), we calculate that a total of 1,400 or 8% experienced bill increases and were therefore eligible to receive refunds (Table

28).[14] The average refund amount for SmartRate™ only customers was $9.13. The 15 dually enrolled customers with a refund received somewhat less ($7.11).

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
Dual	Yes	1,390	$9.13
	No	16,914	-
	Either	18,304	-
SR only	Yes	10	$7.11
	No	197	-
	Either	207	-
Table note: Unlike all other tables presented in this section, this table was derived from data for all customers participating in the 2019 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 (and North Coast for dually enrolled and "Other" areas outside LCAs).

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

Enrollment Status	LCA	% of Population*	Avg. Bill Change
SR only	Stockton	16%	-$28.80
	Sierra	20%	-$28.46
	Other	8%	-$20.94
	North Coast and North Bay	5%	-$24.02
	Kern	13%	-$40.13
	Humboldt	-	-
	Greater Fresno Area	23%	-$35.62
	Greater Bay Area	13%	-$22.90
	All	100%	-$30.50
Dual	Stockton	25%	-$27.62
	Sierra	24%	-$27.38
	Other	2%	-$20.54

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

	North Coast and North Bay	4%	-$16.78
	Kern	7%	-$46.66
	Humboldt	-	-
	Greater Fresno Area	20%	-$37.14
	Greater Bay Area	16%	-$22.53
	All	100%	-$29.69

*May not add to 100% due to rounding.

Customers on CARE experience somewhat smaller reductions in their average bill than non-CARE customers. The average bill change for CARE customers ranged from a reduction of $26.93 to a reduction of $28.26 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	32,801	50%	-$32.70
	CARE	20,945	32%	-$26.93
	All	53,746	81%	-$30.50
Dual	Non-CARE	8,186	12%	-$30.39
	CARE	4,325	7%	-$28.26
	All	12,511	19%	-$29.69