Non-Program Residential Battery Storage Identification in High Fire Threat Districts and Disadvantaged Communities

Abraham Torok, Verdant Associates
Ethan Barquest, Verdant Associates
Collin Elliot, Verdant Associates
Brian McAuley, Verdant Associates
Neil Yazma, Pacific Gas and Electric Company

ABSTRACT

Nearly 40 thousand behind-the-meter energy storage systems have been interconnected in California with the help of incentives provided by the Self-Generation Incentive Program (SGIP). In recent years, however, residential customers are increasingly installing systems without any program incentive. Currently, roughly half of systems in California are installed independent of program incentives. While the SGIP helped drive market adoption of battery systems, the program also provided some level of control over how battery systems were used through program rules. This helped ensure systems were used to support grid resiliency and system performance would not lead to increases in greenhouse gas emissions. Further, program rules allowed for monitoring of systems that made it possible to characterize their behavior, validate that they are operating as intended, and determine the systems' grid impacts. With so many systems being installed outside the SGIP, IOUs and regulators are left with blind spots regarding how these systems are used and how they operate without program performance requirements.

This study uses machine learning approaches applied to standard AMI data to identify the operating mode of systems not enrolled in the SGIP program, providing some idea of how various customer segments are operating their storage systems (backup/resilience versus "not resilience" or "active cycling"). Additionally, the study will compare the share of backup/resilience and active storage systems in population segments such as high fire threat district customers, disadvantaged communities, and customers on low-income rates. These can be used to develop strategies to recruit and incentivize these non-program systems to be used in ways that provide more benefits to the electric grid and customers.

Introduction

Nearly 40 thousand behind-the-meter energy storage systems have been interconnected in California with the help of incentives provided by the Self-Generation Incentive Program (SGIP). In recent years, however, residential customers are installing systems without any program incentives. Currently, roughly half of storage systems in California are installed independent of program incentives (Verdant, 2024a). While the SGIP helped drive market adoption of battery systems, the program also provided some level of control over how battery systems were used through program rules. This helped ensure systems were used to support grid resiliency and system performance and would not lead to increases in greenhouse gas emissions. Further, program rules allowed for monitoring of systems that made it possible to characterize the behavior of their storage systems, validate that they are operating as intended, and determine the systems' grid impacts. With so many systems being installed without program incentives, IOUs and regulators are left with blind spots regarding how these systems are used and how they operate without program performance requirements.

Behind-the-meter (BTM) residential energy storage systems incentivized through the SGIP must meet certain eligibility criteria – minimum cycling requirements and roundtrip efficiencies. These criteria differ based on when a customer applied to the program, but overall, they are designed to ensure systems are used for more than personal back-up support and performance does not lead to increases in greenhouse gas (GHG) emissions.

Prior SGIP impact evaluations (Verdant, 2022; Verdant, 2024a) have measured and tracked the performance and operational strategies of these systems using battery telemetry, (charge and discharge data) along with other available data, including PV generation and net metered utility usage. These data provide a more complete picture of an SGIP storage participant's sources of energy for household consumption (PV production, battery discharge, utility delivery) and where the energy produced by the customer is used (customer consumption, battery charging, utility received). This makes it possible to characterize the behavior of their storage systems, validate that they are operating as intended, and determine the system impacts on the grid.

However, these telemetry data are only available for SGIP incentivized systems. One of the objectives of the SGIP is market transformation, where early program intervention helps steer the general storage market and, over time, helps foster a self-sustaining market where incentives are no longer needed to guide the success of BTM installations. Market transformation is a long-term goal, however, but certain customer segments are increasingly installing energy storage systems outside of the program.

This trend is presented below in Figure 1, where cumulative energy storage interconnections collected from DGStats for 2018-2022 are combined with SGIP program data for PG&E, SCE, SDG&E (and statewide). SGIP projects represented almost all interconnected systems in 2018-2019, but more recently, non-SGIP projects have outpaced SGIP rebated interconnections. By the end of 2022, non-SGIP projects represented more than 50 percent of all BTM storage interconnections in IOU service territory (over 44,000 of the 81,000 systems installed statewide).

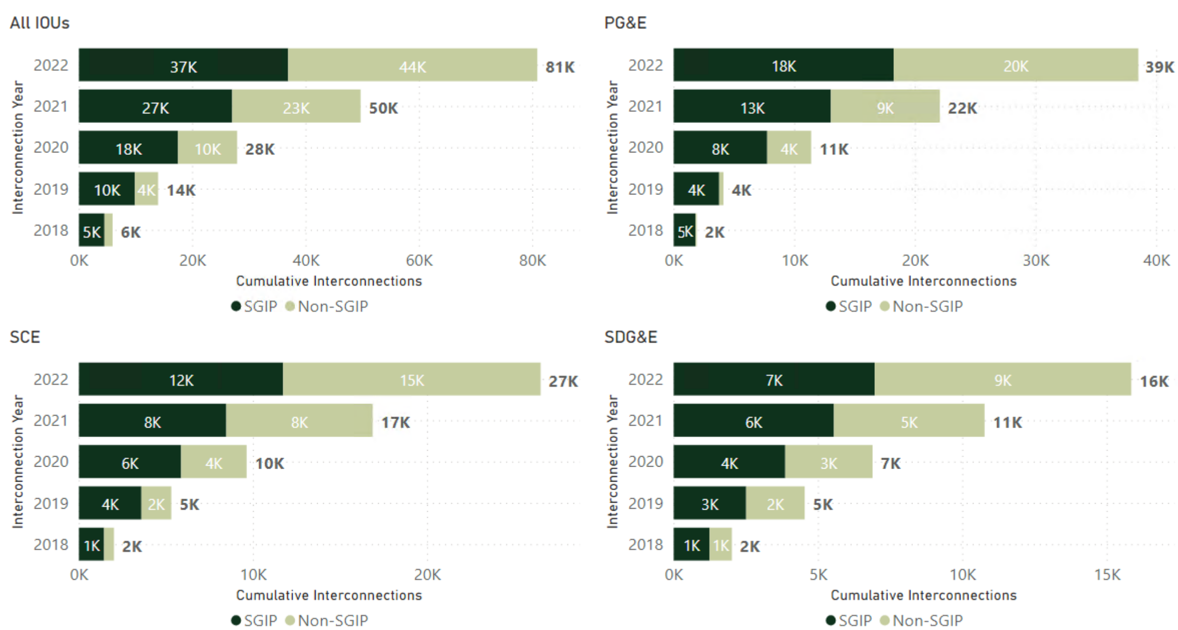

Figure 1: SGIP and Non-SGIP BTM Storage Interconnections by IOU
Source: *DGStats/Verdant Analysis*

SGIP impact evaluations (Verdant, 2022; Verdant, 2024a) have provided insights into the behavior of SGIP incentivized systems, but considerable uncertainty exists around the behavior of non-SGIP energy storage systems. Absent SGIP dispatch and cycling requirements, residential energy storage systems could be sitting idle and used exclusively for backup. Not only is this information useful for California policymakers and grid planners, it presents a potential untapped resource for future California policies seeking to leverage behind-the-meter storage. . However, policymakers and grid planners do not have access to battery telemetry for non-SGIP systems, but utility AMI data are readily available for systems identified in interconnection datasets. These data include channels for delivered load and received load (NEM customer export). Historically, these data sources were insufficient to understand the underlying behavior of the battery storage systems, without access to telemetry data, and likely not without conducting sophisticated analyses.

However, prior evaluations of SGIP battery storage customers have allowed for significant access into how BTM energy storage systems operate and behave under numerous scenarios and across multiple customer segments. While this paper looks to classify non-SGIP storage systems as back-up/resilience and active use, it is valuable to understand the types of storage system operations. Through evaluation efforts, multiple operating modes available to residential customers have been identified. These include:

- **TOU Arbitrage:** Systems are programmed to charge when rates are lowest (and from early on-site PV generation) and discharge when peak rates are in effect.
- **TOU Arbitrage and DR response:** In addition to TOU arbitrage, systems are also programmed to respond to specific grid emergencies through DR program automation. During these periods, the battery is not only limiting grid energy during peak times, but also exporting excess discharge to the grid.
- **Self-Consumption:** Systems are programmed to minimize the consumption of electricity from the utility grid and maximize their use of on-site solar generation. This is also referred to as net zero, where homes are generating as much energy as they are using.
- **Backup/Resilience:** Systems are programmed to maintain a full charge until a grid or local outage. Systems paired with solar PV are capable of riding out longer duration utility power shutoff – sometimes for 3 days – because the system could charge directly from solar, and the solar energy could be used to partially power the home during the day.

Examples of these dispatch strategies from SGIP participants, minus backup/resilience, which is likely very similar to typical residential load[1], are presented in Figure 2, which shows the average daily load profiles for both net load and net load with the battery. Where net load (the dashed line) represents counterfactual net energy consumption without use of the battery and net load with the battery (green line) represent net energy consumption with use of the battery as seen by the utility. Gaps between these lines represent system charging (when solid line is greater than the dashed line) and discharge (when the dashed line is greater than the solid line).

[1] In contrast with other dispatch strategies, the backup/resilience category is based on the absence of any battery usage behaviors, as systems should show minimal usage except during outages, in which case there will not be any AMI data.

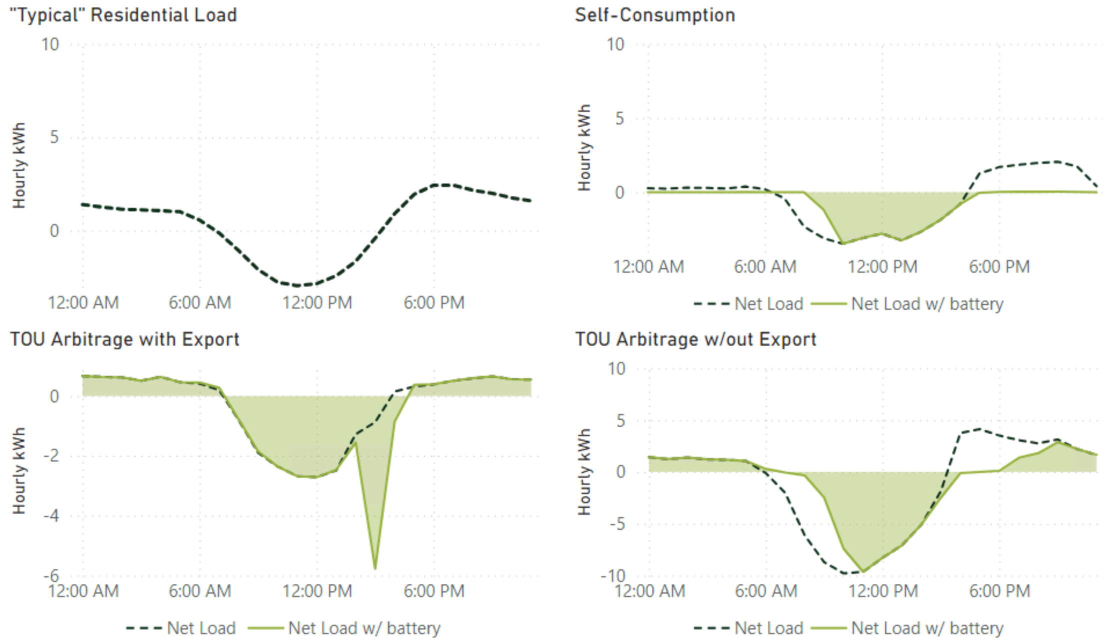

Figure 2: Load Profiles for Battery Storage Behavior
Negative hourly kWh represents export to the grid.

Given this insight and the significant AMI and telemetry data from SGIP participants it is possible to close the gap in understanding how non-program battery storage operates within California. As a result, this paper looks to identify non-program battery storage use cases (backup/resilience versus active use) within PG&E's and SCE's service territories[2] with Machine Learning (ML) based solely on AMI data. Additionally, this paper seeks to identify the share of backup/resilience systems within specific population groups, primarily systems within High-Fire Threat Districts (HFTD), in disadvantaged communities (DAC) and across different rate groups.

Methodology

At a high level there are three main steps for the battery storage identification analysis. These include 1) data collection, 2) training dataset development, and 3) ML model training and validation. For this analysis, data from a variety of sources were used including:

- **Customer information**. These data consist of account IDs and a variety of other attributes that are necessary for segmentation and the identification of control/testing groups. These attributes include climate zones (CZ), SubLAP and LCA, rates, customer fuels (electric only or gas and electric), CARE/FERA status, Disadvantaged Community (DAC) status, weather station(s), NEM status, and any information on generation technologies..
- **AMI data**. The analysis required service point level AMI data, including delivered and received channels, for all sampled households (non-SGIP and control homes).
- **Weather data.** Temperature, irradiance, and precipitation data for all weather stations corresponding to customers included in the analysis.

[2] SDG&E is not a program administrator (PA) to the SGIP like the other electric IOUs, so AMI data collection for this effort was limited to PAs with direct access to AMI data.

- **SGIP program data**: Program data for all SGIP participants.
- **System Interconnection Data**: Data on system characteristics and installation dates for all battery and solar customers.

Training Dataset Development

As described previously, the objective of this analysis is to develop a tool that can classify the battery dispatch strategy (resilience and active use) based on AMI data using machine learning (ML) classification methods.

A fundamental input for employing supervised ML methods is a training dataset that can be used to test and evaluate the performance of different models. For this analysis, we use a set of households with known battery behavior strategies. The initial basis for the households in this training dataset is SGIP customers for which there is sufficient data from their systems to establish clear discharge patterns. Given program rules governing minimum cycling requirements, SGIP systems are expected to be utilized regularly (and not be exclusively in backup mode). To account for this, the training dataset utilizes residential customers with solar interconnections, but no battery storage system, as a proxy for backup/resilience storage systems. As described above, it is expected that customers with storage that only use their system for resilience will look like the "Typical" residential load depicted in Figure 2 above.

Once the customers for the training data are identified, the next step is to leverage their premise-level AMI data to calculate a variety of metrics, or variables, that could be associated with the different dispatch strategies. As an example, customers with systems following time-of-use arbitrage exhibit a marked difference in their loads during various time-of-use windows, and this difference may have high predictive value in classifying these customers. Similarly, metrics that show no meaningful changes to the load profiles are likely predictive of a system being employed for resilience.

Ultimately, the training dataset used for the analysis consists of 28 metrics (or features), 24 of which correspond to the mean hourly load from a given week of the year, and the remaining four were variables indicating season. To build a training set for supervised machine learning, the AMI data from all the no-storage solar systems was transformed to return average hourly load per meter, per week. After many iterations of training and testing, the weekly aggregation was elected as the final approach because it balanced the two limiting factors: data outliers, which are seen more often when data is aggregated over smaller periods of time (each single data point has a larger overall impact on the aggregate) and the lack of modeling data. While working with data at the monthly aggregation level creates much smoother load shapes with less influence from outliers, the reduction in total size of training data was deemed to be too significant.

The first step in the creation of our final machine learning model was to perform algorithm selection. For this process, the training dataset was split into two samples; an in-sample training dataset used for training the model (training dataset) and an out-of-sample dataset used for validation (validation dataset) of the ML model's performance. The in-sample training dataset contained 80% of the overall training dataset's premises. The other 20% make up the out of sample validation dataset.

Residential load shapes can look very different from each other based on the routines and lifestyle. For the training model development, it is important that the data from different weeks of a premise does not get divided into both the training and validation dataset, as this will lead to an overfit model. As a result, a group k-fold validation method was used to control this intra-class variance and prevent any data leakage. The goal of dividing the data in this manner is to maximize the generalizability of the model beyond the training data and avoid overfit of the ML model.

ML Algorithm Methodology

Using the base kit of parameters, the training and validation data were used to compare different supervised machine learning algorithms: a random forest ensemble classifier, linear regression classifier, support vector machine classifier, and finally a 1-dimensional convolutional neural network (CNN). After testing the various approaches, it was determined that the random forest would be best suited for the analysis as it is not restricted by the need to normalize feature ranges. This allows us to compare 2kW, 5kW, and 10kW systems without requiring any prior knowledge of system size, although a CNN was seriously considered as a suitable alternative.

The random forest model is an ensemble learning technique widely used in machine learning and statistical analysis. It operates by constructing a multitude of decision trees during training, each based on a random subset of the features from the training dataset. This randomness injected into the model helps reduce overfitting and increases robustness of the model. Each decision tree serves as a weak learner, providing its prediction based on the subset of features it has been trained on. Through a process called bagging (Bootstrap Aggregating), the model combines the predictions of these individual trees to arrive at a final prediction.

The random forest classifier may have been the best out of the box, but there was still quite a bit of optimization left undiscovered. In the pursuit of the optimal hyperparameters for our model, a random search was carried out to narrow down the areas worth pursuing at a higher granularity. Like the choice of aggregating data at the weekly level, utilizing a random search was the optimal compromise between computational cost and the need for an extensive exploration of hyperparameter space, enabling the identification of effective model configurations without the exhaustive computational requirements of a grid search. After an initial round of random searching through our hyperparameter space, an exhaustive search was done through a much smaller region, leading to the final model specifications.

Model Validation

With the hyperparameters of the model tuned, the final validation dataset is used to validate the model's accuracy by scoring the known out-of-sample sites and comparing the predicted use case against the known use case. Overall, the ML model was able to predict with an accuracy of 90.1% whether a given premise's battery operation was backup/resilience or an active system. The confusion matrix, depicted in Figure 3 (right), details the results of the validation predictions, with the 0 class meaning "Active Use" and the 1 class meaning "Backup/Resilience". In other words, a 0 would mean that the model has classified this load shape as showing signs of repeated battery discharge, while a 1 would indicate that the battery is not discharging to offset load in any way.

While the model is not perfect, it scored the correct system behavior with 90.1% accuracy. Additionally, there appears to be diminishing returns by including more data. Figure 3 (right), below presents the model learning curve which highlights the incremental accuracy of the model by including additional sample data into the training dataset. The point of inflection occurs around ~55,000 data points, where the second derivative of the learning curve changes sign. This indicates a transition in the model's performance improvement rate. Specifically, this point marks where the learning curve changes from accelerating improvement to decelerating improvement. This supports the finding that our model is likely very close to the upper limit of what is possible given our specific set of features and training data.

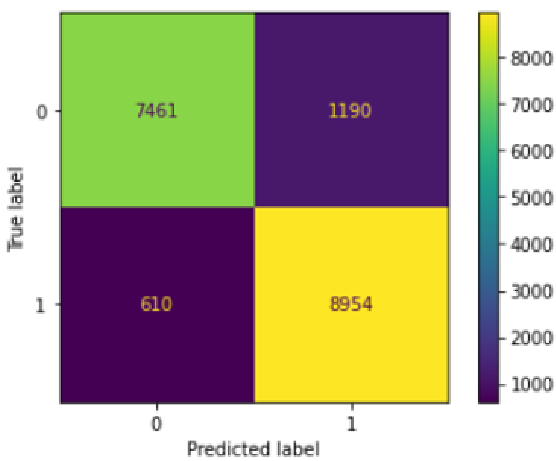

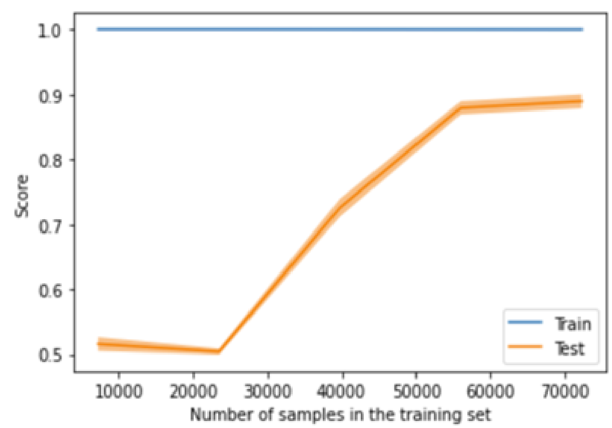

Figure 3: Confusion Matrix of the Random Forest Model Predictions of Storage System Use (Left) and Model Learning Curve (Right)

Figure 4 presents three examples of load shapes classified as Resilience (top row), and although they differ from one another, they all share some key characteristics of a solar residential load shape lacking any sort of load offset from a storage system, despite having one installed. One common characteristic of these backup/resilience examples is evident in all three examples, where the maximum energy usage occurs at or around the peak period. The second-row highlights some of the different shapes seen in the “active use” data (which appear to be correctly identified). The right-most load shape (in the second row) looks to be self-consumption (where the system aims for a net zero load over the course of the day) and the left-two load shape resembles time of use arbitrage without export (where battery capacity is only discharged during the peak period to minimize the overall cost of electricity).

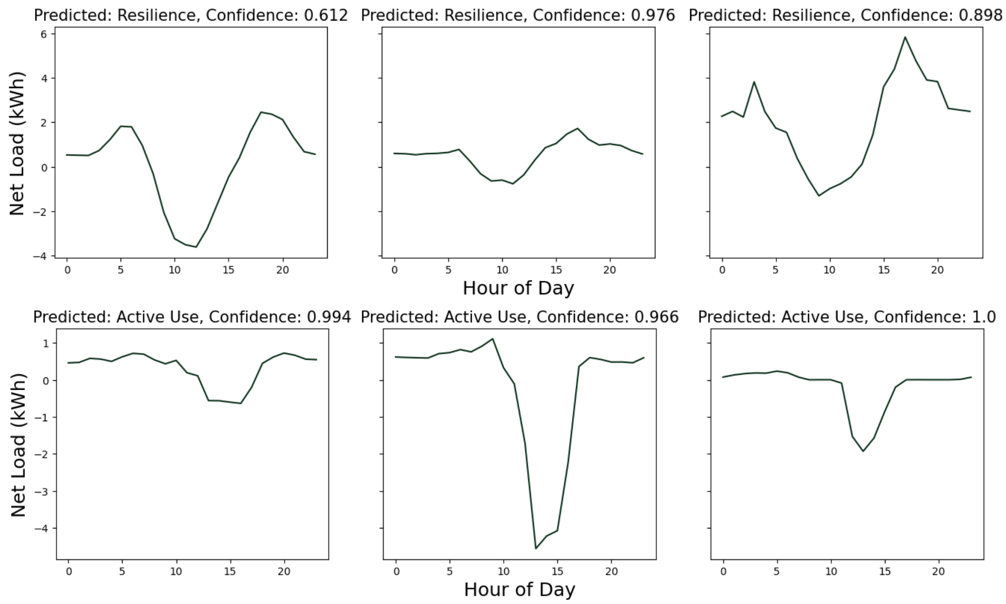

Figure 4: Example of Model Meter Classifications

Results

Overall, it is estimated that 24% (5,532) of PG&E and 16% (2,402) of SCE non-program battery storage systems are used solely for backup/resilience operation and are not cycled on a regular basis. As a result, these systems are not providing any additional benefits to the grid or customer bill reductions, in fact they are slightly increasing their bills due to battery round trip efficiency losses and other system parasitics. Figure 5 presents the count and share of non-program residential battery storage systems by dispatch behavior for both PG&E and SCE.

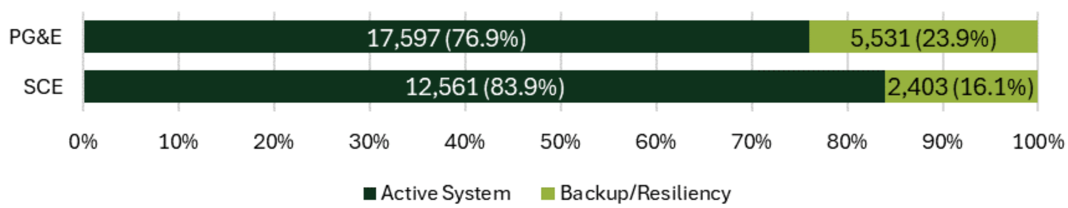

Figure 5: Count and Share of Modeled Non-Program System Operation

Figure 6**Error! Reference source not found.** further breaks down the share of backup/resilience battery storage systems by DAC and CARE/FERA status. DAC customers represent those that are located within a disadvantaged community and CARE/FERA represent residential customers that are on low-income CARE or FERA rates. Customers on the CARE rate receive a 30% to 35% discount on their bills and FERA customers and 18% discount on their bills (if they do not qualify for CARE). As seen, for both PG&E and SCE there are higher shares of backup/resilience systems for customers in DACs and on low-income rates. For PG&E there is a 9% higher share of backup-systems in DACs compared to non-DAC storage systems and a 12% higher share between CARE/FERA compared to non-CARE/FERA storage systems. Similarly, for SCE there is a higher share of backup systems in DAC and low-income groups, however the delta between the comparative shares is smaller; with a 4% higher share for DAC systems over non-DAC and a 9% higher share for CARE/FERA systems over non-CARE/FERA.

However, it should be noted that most systems are in non-DAC and non-CARE/FERA populations and contain the majority of non-active, backup/resilience systems.

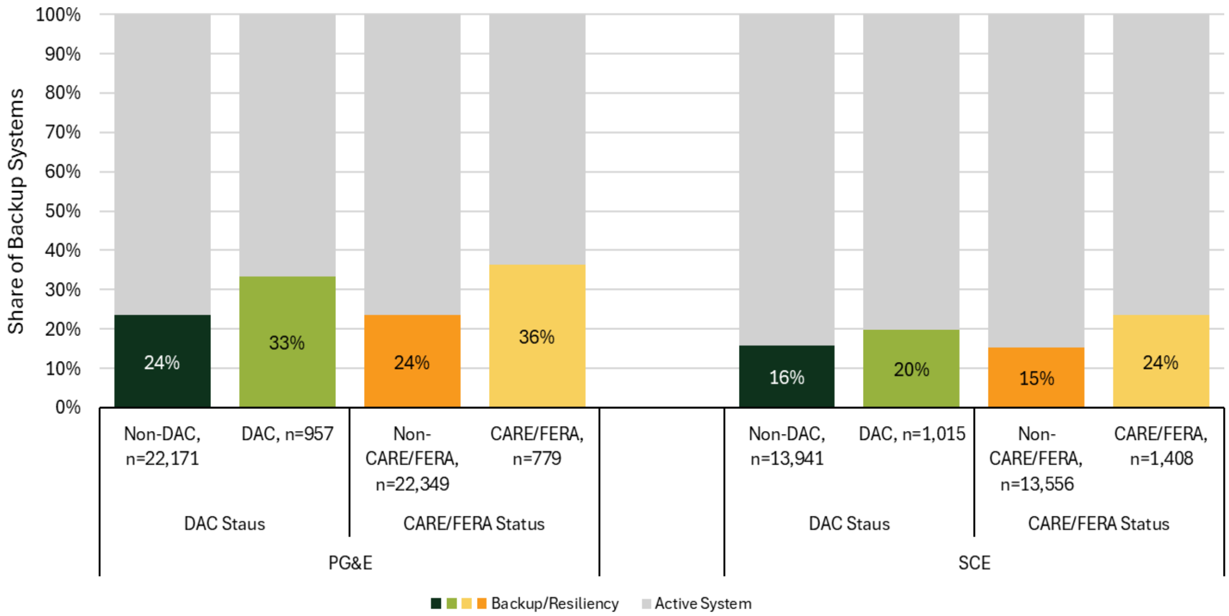

Figure 6: Share of Modeled Non-Program System Operation by DAC and CARE/FERA Status

Anecdotally, the shares of backup systems within DAC and CARE/FERA groups make sense relative to their counterparts. Customers within DACs are more likely to be on CARE and FERA rates. As a result, there may be less financial incentive to actively dispatch batteries to offset bills as these customers already receive a discount on bills through CARE and FERA. However, this is only speculative and additional research is needed to truly understand the underlying causes for these differences.

Figure 7 presents the share of backup/resiliency battery storage systems by HFTD status for PG&E and SCE non-program storage systems. As seen, there is not a substantial difference between SCE customers in and out of HFTDs. For PG&E there is variation between HFTD, with greater shares of customers actively using their battery for storage systems outside of Tier 2 and 3 HFTDs. Notably, there is a higher share of backup/resiliency systems in Tier 2 HFTDs than in the riskier Tier 3 HFTDs. Despite this the PG&E share of Backup/Resilience system only range for 23% of systems in non-HFTDs to 29% in Tier 2 HFTDs. Regardless of whether or not a premise is in an HFTD, there is an overall higher share of resilience systems in PG&E in comparison to SCE.

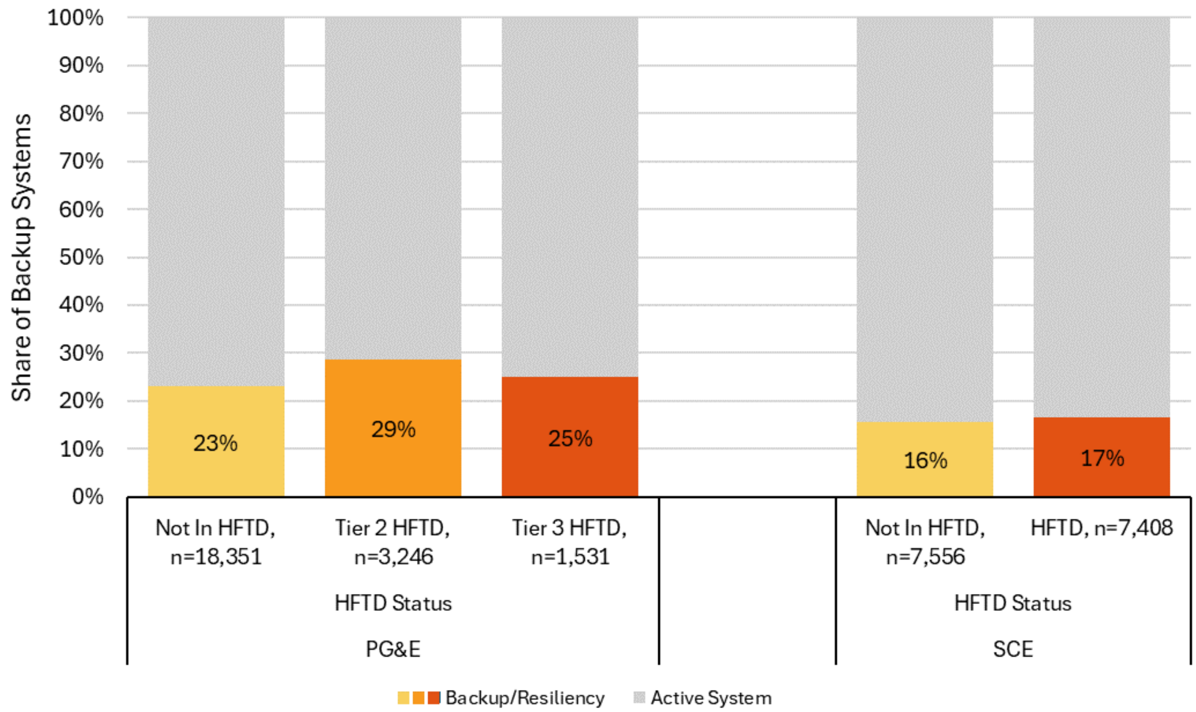

Figure 7: Share of Modeled Non-Program System Operation by HFTD Status
Note: PG&E provided HFTD by Tier while HFTD Tiers were not available for SCE within data provided.

The results between PG&E and SCE correlate with the frequency of Public Safety Power Shutoff (PSPS) events[3] since 2020 which affected more PG&E SGIP customers than SCE SGIP customers (used as a proxy for non-program battery storage customers). In 2020, 21% of PG&E SGIP storage participants were affected by PSPS events, 3% in 2021 and 0% in 2022 (Verdant, 2022, Verdant, 2024a). Whereas 4% of SCE SGIP participants in 2020 were affected by PSPS event, 6% in 2021 and 1% in 2022. While AMI data used in this analysis comes from 2022, it is may be that customers' historic experiences with PSPS events (especially in 2020) influenced the purchasing decision of a storage system and the selected battery operation strategy at a higher incidence for PG&E customers than for SCE customers.

Figure 8 presents the share of backup/resilience systems by different rate groups. The rates have been combined across the IOUs to provide a holistic view of the distribution of system behavior across rates in California and represent the last rate a participant was on in 2022. Given the number of rate tariffs for these customers, they are binned into three groups TOU (time-of-use), non-TOU, and EV (EV TOU and EV non-TOU). Unsurprisingly, non-program systems on non-TOU rates have a higher share of backup/resilience systems (23%) than non-TOU (21%) and EV rates (15%), although this difference is small.

[3] PSPS events are planned outages where powerlines are energized to prevent wildfires and/or provide public safety.

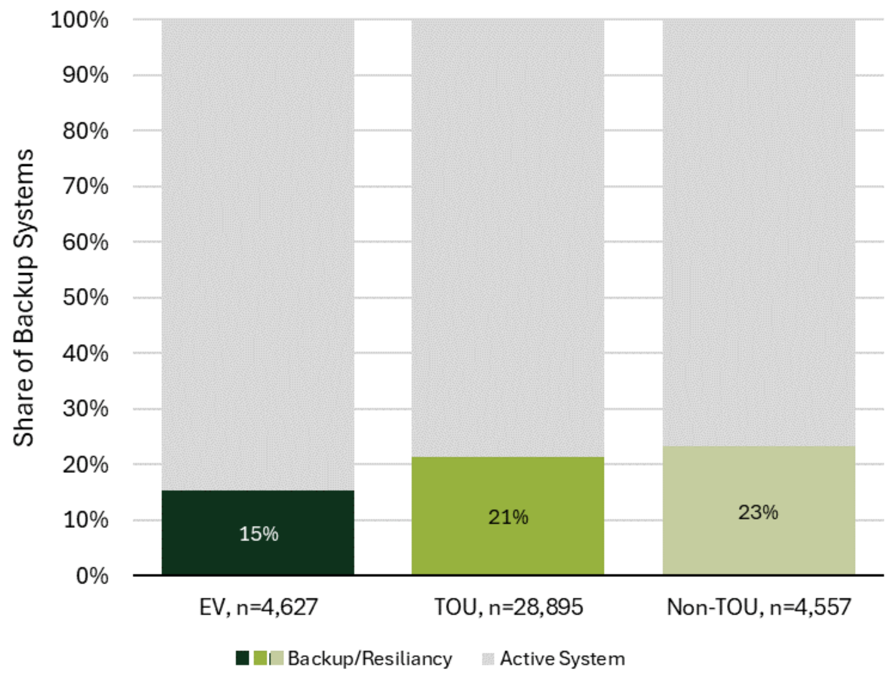

Figure 8: Share of Modeled Non-Program System Operation by Rate Group

TOU rates are expected to encourage daily discharge of battery storage systems during peak periods when retail prices of electricity are higher. In surveys with SGIP customers, TOU rate arbitrage is one of the more commonly cited reasons for installing battery storage (Verdant, 2024b). Although TOU rates have a smaller share of backup/resiliency storage systems than non-TOU rates, the difference in share between them is not substantially large (2%). However, the trend is intuitive. Customers on EV rates had the lowest share of customers with backup/resiliency systems. EV customers may be more aware of the benefits of frequent battery usage, particularly if they charge their vehicles at home during the day. Additionally, EV customers may be more aware of the impact of peak period pricing and therefore use their system often. This analysis frequently saw customers with overnight peak loads for these customers, where EVs appear to be charging during "super" off-peak periods.

Findings

While not perfect, AMI data can be used to create insights about the behavior of residential storage systems to a relatively high degree of accuracy. Leveraging scale-invariant machine learning methods such as the random forest allows for storage system insights without any knowledge of a system's storage capacity or any other scaling factor. (However, prior understanding of system operations in another population is necessary to develop a model of this type). Within our training dataset, we were able to identify backup/resilience and actively used systems with an accuracy of 90%. While a known population of storage installation is a necessary precursor, this information is readily available in utilities' interconnection data. If necessary data conditionals exist, this approach will allow utilities and program managers to identify individual candidates for program offerings or messaging for programs that encourage load shifting.

Additional takeaways and findings include:

- Roughly 24% of PG&E and 16% of SCE non-program battery storage systems are used solely for backup/resilience operation. This represents a significant untapped resource that may provide additional grid benefits if they can be encouraged or incentivized to actively use their battery storage systems.
- Customers in DACs and on CARE/FERA rates have a higher share of systems that are idle compared to their non-DAC and non-CARE/FERA counterparts. Given the customer overlap of these groups and discounted bills, there is likely less financial incentive for these customers to actively use their systems. While these customer segments represent traditionally hard to reach populations, it may be beneficial for utilities to encourage these customers to use their system to participate in load shifting programs rather than encourage active system usage through rate design. However, additional research should be conducted before actions are taken in vulnerable communities.
- Storage system behaviors in HFTDs do not appear to be different from non-HFTDs for SCE (at least in terms of share of backup/resiliency storage systems). However, there is a measurable difference within PG&E's HFTDs. Experience with PSPS events may result in additional uptake of backup/resiliency storage system operation where customers do not have interest in operating their system in another way. Information campaigns informing these customers that their system can provide backup power during an outage even if the battery system is frequently charged and discharged may encourage a switch to a more optimal strategy.
- Customers on non-TOU rates are more likely to leave their battery systems idle. However, the share of non-TOU rates is not substantially different from customers on TOU rates.

References

California Distributed Generation Stats. "Interconnected Project Sites 2023-10-31", https://www.californiadgstats.ca.gov/archives/interconnection_rule21_projects/

Verdant Associates. 2022. *2020 SGIP Energy Storage Impact Evaluation.* San Francisco, CA: California Public Utilities Commision. https://www.cpuc.ca.gov/-/media/cpuc-website/divisions/energy-division/documents/self-generation-incentive-program/sgip-2020-energy-storage-impact-evaluation.pdf

Verdant Associates. 2024a. *2021-2022 SGIP Impact Evaluation.* San Francisco, CA: California Public Utilities Commision. https://www.cpuc.ca.gov/-/media/cpuc-website/divisions/energy-division/documents/self-generation-incentive-program/sgip-2021-2022-impact-evaluation.pdf

Verdant Associates. 2024b. *2023 SGIP Program Performance and Process Evaluation (PPPE).* Unpublished Report.