Lawrence Berkeley National Laboratory

Title
The California Demand Response Potential Study, Phase 3: Final Report on the Shift Resource through 2030

Permalink
https://escholarship.org/uc/item/7bx121k6

Authors
Gerke, Brian
Gallo, Giulia
Smith, Sarah
et al.

Publication Date
2020-07-15

DOI
10.20357/B7MS40

Peer reviewed

Lawrence Berkeley National Laboratory

The California Demand Response Potential Study, Phase 3:
Final Report on the Shift Resource through 2030

Brian F. Gerke,[1]* Giulia Gallo,[1] Sarah J. Smith,[1] Jingjing Liu,[1] Peter Alstone,[1,2,3] Shuba Raghavan,[1,4] Peter Schwartz,[1] Mary Ann Piette,[1] Rongxin Yin,[1] and Sofia Stensson[1,5]

[1] Lawrence Berkeley National Laboratory
[2] Humboldt State University
[3] Schatz Energy Research Center
[4] University of California, Berkeley
[5] RISE Research Institutes of Sweden
* Email: bfgerke@lbl.gov

Energy Technologies Area
July 2020

Lawrence Berkeley National Laboratory

The California Demand Response Potential Study, Phase 3:

Final Report on the Shift Resource through 2030

Brian F. Gerke,[1]* Giulia Gallo,[1] Sarah J. Smith,[1] Jingjing Liu,[1] Peter Alstone,[1,2,3] Shuba Raghavan,[1,4] Peter Schwartz,[1] Mary Ann Piette,[1] Rongxin Yin,[1] and Sofia Stensson[1,5]

[1] Lawrence Berkeley National Laboratory
[2] Humboldt State University
[3] Schatz Energy Research Center
[4] University of California, Berkeley
[5] RISE Research Institutes of Sweden
* Email: bfgerke@lbl.gov

July 14, 2020

BERKELEY LAB

<div align="center">**Disclaimer**</div>

This document may contain research results which are experimental in nature. Neither the United States Government, nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not constitute or imply an endorsement or recommendation by the United States Government or any agency thereof, or by The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, or of The Regents of the University of California and shall not be used for advertising or product endorsement purposes.

<div align="center">**LEGAL NOTICE**</div>

This report was prepared as a result of work sponsored by the California Public Utilities Commission. It does not necessarily represent the views of the Commission or any of its employees except to the extent, if any that is has been formally been approved by the Commission at a public meeting. For information regarding any such action, communicate directly with the Commission at 505 Van Ness Avenue, San Francisco, California 94102. Neither the Commission nor the State of California, nor any officer, employer, or any of its contractors or subcontractors makes any warranty, express or implied, or assumes any legal liability whatsoever for the contents of this document.

<div align="center">**Acknowledgments**</div>

Table of Contents

List of Figures

List of Tables

List of Acronyms

AAEE	additional achievable energy efficiency
AC	air conditioner
ACEEE	American Council for an Energy-Efficient Economy
ADM	ADM Energy Evaluation and Research (consulting firm)
AS	ancillary services
ASHP	air-source heat pump
BAU	business-as-usual
BPA	Bonneville Power Authority
BTM	behind-the-meter
CAISO	California Independent System Operator
CARB	California Air Resources Board
CCA	community choice aggregator
CEC	California Energy Commission
CO_2	carbon dioxide
CPUC	California Public Utilities Commission
DER	distributed energy resource
DLC	direct load control
DR	demand response
E3	Energy + Environmental Economics (consulting firm)
EE	energy efficiency
EMS	energy management system
ERWH	electric resistance water heater
ESDER	energy storage and distributed energy resources
EV	electric vehicle
GHG	greenhouse gas
GW	gigawatt
GWh	gigawatt-hour
HPWH	heat pump water heater
HVAC	heating, ventilation and air conditioning
IDSM	Integrated Demand-side Management
IEPR	Integrated Energy Policy Report
IOU	investor-owned utility
IRP	integrated resource planning
kW	kilowatt
kWh	kilowatt-hour
LADWP	Los Angeles Department of Water and Power
LBNL	Lawrence Berkeley National Laboratory
LMDR	load modifying DR
LSE	load serving entity
LSWG	Load Shift Working Group
MW	megawatt
MWh	megawatt-hour
NOAA	National Oceanic and Atmospheric Administration
NREL	National Renewable Energy Laboratory
PCT	programmable communicating thermostat

PDR-LSR	Proxy Demand Resource – load shift resource
PG&E	Pacific Gas and Electric Company
POU	publicly owned utility
RPS	renewable portfolio standard
SCE	Southern California Edison Company
SDG&E	San Diego Gas & Electric Company
SGIP	Self-Generation Incentive Program
SMUD	Sacramento Municipal Utility District
subLAP	sub-load-aggregation-point
T&D	transmission and distribution
TES	thermal energy storage
TOU	time-of-use
V2G	vehicle to grid
VRE	variable renewable energy
WH	water heater
ZEV	zero emission vehicle

Executive Summary

In response to the imperative of the climate crisis, California is in the midst of deployment of low-carbon energy generation and electrified heating and transportation at a rapid pace. In 2018 the grid managed by the California Independent System Operator (CAISO) was already generating 26 percent of its power from non-hydro renewable sources, with an additional ten percent hydroelectric power; and with SB 100, California statute now sets a target of 100 percent carbon free electricity by 2045. Maintaining this fast pace of decarbonization requires action across multiple domains on the grid from generators to loads. This study describes new findings on **the scale of the opportunity for enabling load shifting as a form of demand response (DR)** in California, as part of this renewable energy transition. We use the shorthand term *Shift* to refer to this potential load-shifting DR resource.

The operating principles for the electric power system mean that achieving California's renewable energy goals requires careful planning to build and operate the system in a way that maintains reliability at low cost. The most fundamental of these principles is that generation and demand need to be balanced at all times. Historically this has meant building adequate flexible and dispatchable generators that can be operated to match inflexible and uncontrolled loads. Meeting large fractions of demand with variable renewable energy (VRE) generation introduces a new paradigm, exemplified by the "duck curve" in California. In the times of year it occurs, the duck curve illustrates the net electricity demand that must be met by non-VRE generation; it has sharp peaks in morning and evening (the "tail" and "head" of the duck, respectively) and steep transitions to a deep midday trough (the duck's "belly") resulting from solar generation. Balancing the duck curve is a key priority for integrating increasing levels of VRE on the California grid.

Shift DR is one of several renewables integration planning approaches that can work together to balance a high-VRE system. Others include "overbuilding" renewable energy with the intent to curtail during times of surplus generation, building and operating transmission lines to integrate regional systems, and building and operating energy storage to match generation with loads. Based on the 2019-2020 Integrated Resource Planning (IRP) model for California[1]—which did not consider Shift among its resource options—the least-cost pathway for achieving California GHG goals includes all of these other approaches. The Reference System Portfolio that is used to inform utility planning includes development of energy storage that is operated to serve about 10 GWh of consumption on the average day in 2020, 70 GWh in 2030, and 400 GWh by 2045. There is also significant curtailment of VRE in this least-cost portfolio, with 4 GWh on the average day in 2020, rising to 15 GWh by 2030 and 100 GWh by 2045.

Throughout this report we compare the Shift resource to curtailment and storage as benchmarks, providing context for the results we report related to the quantity, cost, and performance of Shift. The timing of evening peak and curtailment defines times when load shifting is most valuable and would likely be used. The quantity of curtailment and use of storage to balance the system represent how much Shift could be valuable at the upper bound if Shift is inexpensive. The cost of energy storage is also a

---

[1] The IRP results referenced in this section were finalized in CPUC Decision 20-03-028 in Rulemaking16-02-007 (CPUC 2020a), including detailed outcomes from the publicly available version of the RESOLVE model used to support the rulemaking (CPUC 2020b).

useful benchmark for the cost of Shift because it is the most comparable "competing" technology option; we therefore use the cost of storage to set a reference price level below which Shift will be most broadly cost competitive.

Operational data from today's CAISO power system illustrates how the new dynamics of the California grid are already occurring. Curtailment in the spring of 2019 rose to an average value of more than five gigawatt-hours (GWh) per day, as shown in Figure ES-1; this is consistent with curtailment forecasts from the IRP.

Figure ES-1 Average daily curtailment of renewable energy in CAISO, by season, since 2014. Source: CAISO Managing Oversupply web page (CAISO 2019b).

The California DR Potential Study

This report describes the results from Phase 3 of a long-term research effort, known as the California DR Potential Study, by the California Public Utilities Commission (CPUC). Phases 1 and 2 were completed in support of rulemaking R.13-09-011, and identified the range of DR services that could "enhance the role of demand response in resource planning and operational needs." These reports introduced a simplified framework for describing the DR resource in four categories of service: Shape, Shift, Shed, and Shimmy. A key outcome from that work was that Shift could play an important and substantial role in the future of DR. Partly in response to this finding, CPUC decision D.17-10-017 initiated a stakeholder-driven Load Shift Working Group that considered near-term pathways to enabling load shift in California. The present report further supports the CPUC's efforts to enable this category of DR by expanding and improving upon the previously reported findings for Shift in the context of the DR Potential Study.

Scope and focus of the Phase 3 study

The California DR Potential study aims to assess the future potential DR resources that could support the bulk power system within the territories of the three main California electric investor-owned utilities (IOUs): Pacific Gas and Electric Company, Southern California Edison, and San Diego Gas & Electric. Combined, these territories are approximately the same as the CAISO footprint within California. The DR Potential study considers the potential for DR from all customer classes, including customers in the residential, commercial, and industrial sectors. For each sector it focuses on a subset of electrical end uses that are understood to have the potential to provide DR, with some of the most significant being industrial process loads; agricultural/industrial pumping loads; and heating, ventilation, and air conditioning (HVAC) loads in the commercial and residential sectors. Phases 1 and 2 projected California's potential DR resources in years 2020 and 2025; the Phase 3 study extends this forecast horizon to 2030. The Phase 3 study also includes new approaches to estimating the Shift potential from new loads expected over this time horizon from growth in electric vehicle charging and from the electrification of residential space and water heating, and consideration of thermal energy storage in commercial buildings.

The goal of the Phase 3 study is to provide the CPUC and other stakeholders with data-driven insights evaluating how California might use Shift DR in meeting its resource planning needs and operational requirements. We address the following primary questions:

- What is the size and cost of the expected resource base of Shift?
- Where is the Shift resource located; e.g., by end use or geographical region?
- When is Shift needed and when is it available; e.g., how does the shiftable load align with the need to mitigate ramping; how does the resource vary seasonally?
- How will the Shift resource evolve with California's changing grid; e.g., in response to growing electrification or increased VRE penetration?

What is Shift demand response?

DR supports the economic operation of the grid today via services that fall into the Shed category, with participants enabling peak load management by shedding non-critical loads at peak times. From the grid operator's perspective, Shed resources resemble virtual generation capacity to serve peak loads. Over the last two decades, California has developed a valuable Shed DR resource base that can reduce demand at critical times, reducing the need for peak generation capacity by roughly 2 gigawatts (GW), or approximately 4 percent of the annual peak.

Shift is a new DR service concept established in Phase 2 of this study to address the emerging needs from widespread solar and wind power deployment described above. Unlike Shed, which instructs participants to reduce demand at a given time, Shift would instruct participants to change the timing of their demands, shifting energy consumption from one time of day to another. Shift has some similarity to Shed, in that it involves a reduction in demand at a particular time, but in the case of Shift this load reduction is explicitly offset by a load increase at a different time, with the goal of providing an equivalent energy service to the customer with a different timing of the associated electrical load.

From the grid operator's perspective, if Shed resembles virtual peaking generation capacity, Shift resembles virtual behind-the-meter (BTM) storage that can be utilized to smooth variation in the net demand on the grid over hourly to daily timescales in order to more effectively utilize VRE generation. In the Phase 1 and 2 studies, we often used the cost of a natural gas turbine peaker plant as a point of comparison for Shed DR, identifying Shed resources cheaper than that benchmark as being particularly cost-effective. Analogously, in this study we will often compare the cost of Shift to the cost of BTM battery systems and focus much of our attention on the size of the Shift resource that can be obtained more cheaply than such batteries.

In all cases, there is an important difference that should be recognized between the virtual resources provided by DR and the physical resources that would be provided by power plants or batteries. The size of a DR resource varies throughout the day and throughout the year, in parallel with variation in the loads providing DR. Thus, when we report the size of the Shift resource in GWh, we are expressing the quantity of energy that is available to be shifted *on average over the course of the year* (or the time-period being considered, if different from a year) at times when the grid is likely to need Shift. Unlike a physical battery, the Shift resource will not have exactly the same capacity each time it is called. Space cooling loads, for example, will typically be smaller in the winter than in the summer, meaning that the DR resource available from HVAC will vary significantly by season. In the case of Shed, which is typically called only on peak summer days, such seasonal variation is not a significant concern in practice. Because Shift would be used to smooth the daily variation in net loads, not to manage infrequent peak loads—and because it can provide equivalent energy services to customers—it is possible that Shift would be utilized much more frequently than Shed. Investigating the importance of seasonal variation in this context is one of the goals of this study. In that case, we will report the available Shift resource on average over the course of each season. Throughout this report, to help clarify the meaning of our Shift resource estimates, we report the resource as the "average available resource in a Shift event."

Analytical approach

The analytical framework we have developed for the DR Potential Study forecasts DR supply curves for the years 2020, 2025, and 2030 using a bottom-up modeling framework called DR-Futures. The framework consists of two modules:

- **LBNL-Load** is a bottom-up load-forecasting module that capitalizes on a large set of customer smart meter data to project future end-use load shapes for a diverse set of customer clusters.
- **DR-Path** starts from the forecasted load shapes to assess a large number of future pathways to acquiring DR resources, resulting in granular load flexibility potential estimates that can be aggregated into the final supply curve.

The Phase 3 study incorporates updates to LBNL-Load that extend the forecasting horizon to 2030 and add new modeling for growth in light-duty electric vehicle (EV) charging and in the electrification of residential space heating and water heating consistent with the state's climate goals. Updates to DR-Path enable a much more detailed and flexible assessment of the Shift resource, including an improved model for the probability of Shift dispatch, more realistic modeling of operational constraints on Shift, new capability to assess the resource on a seasonal basis, and modeling of new Shift-enabling technologies

such as thermal energy storage. Box 2 in the main report gives a brief summary of the most important updates to the modeling framework.

The final model output is a supply curve for Shift that expresses the quantity of Shift, in GWh, that can be procured at or below a given annualized cost, in $/year/kilowatt-hour. This Shift quantity represents the average amount of energy that could be shifted *each time the resource is dispatched*, within reasonable constraints on the frequency of utilization. In particular, some or all of this resource may be accessible more than once per day (e.g., shifting energy consumption out of the morning demand peak into the late morning, and from the evening peak into the afternoon). This paradigm for reporting the size of the Shift resource *per event* differs from the approach taken in the Phase 2 report, which reported the quantity of Shiftable energy *per day*, under the assumption that two Shift events would occur each day. Importantly, this difference in accounting makes it difficult to compare the results of this study in detail to the Phase 2 results (see section 2.2.2.2 of this report for further discussion of this issue).

The cost accounting is performed from a utility (or DR program administrator) perspective: all costs borne by a utility or a DR aggregator are included in the resource cost, including capital and operating costs for installed technologies, program administration and marketing costs, and customer participation incentives. Any portion of the cost borne directly by the customer[2] is excluded from the resource cost.. Because the supply curve is produced through a bottom-up modeling approach, it can be disaggregated to reveal the various different flexible loads that make up the Shift resource at any given cost level.

Main Findings

Figure ES-2 shows the supply curve for Shift disaggregated by sector (shaded regions) and by end use (colored bars), under our default modeling assumptions, in each of the forecast years considered in this study. Dotted horizontal lines show the cost, in each year, of procuring BTM battery storage as a Shift resource, expressed in the same cost units as the supply curve,[3] as a benchmark for comparison. In light of these primary results, the main findings of the Phase 3 study can be summarized as follows.

1. Today's potential Shift resource could utilize much of the daily VRE curtailment encountered recently, and substantially reduce flexible generation needs, at a lower cost than BTM battery storage. In 2020, the resource that is available at or below the battery benchmark (see Figure ES-2) amounts to 5.3 GWh of Shift resource, primarily provided by commercial HVAC, industrial process, and agricultural pumping loads. A single dispatch of this entire resource would be sufficient, in principle, to utilize much or all of the otherwise-curtailed energy on an average day in spring 2019[4] (see Figure ES-1). Recall also that the supply curve represents the quantity of energy that can be shifted *each time* the

---

[2] This is estimated at the value of any technology co-benefits that accrue to the customer, such as energy savings or non-energy benefits.

[3] The BTM battery cost benchmark also includes all the same cost components as used to compute the Shift costs, including technology costs, operating costs, and program administration costs.

[4] A major caveat to this conclusion is that roughly half of recent curtailment in CAISO has been driven by local transmission constraints, not system-level overgeneration. Mitigating local curtailment with Shift would require utilizing resources that are located in the constrained areas, which may not be large enough to completely offset this category of curtailment.

resource is dispatched. If some significant proportion of this resource could be utilized twice per day (shifting loads from both the morning and evening peaks into midday hours), then this resource could potentially offset more curtailment on above-average curtailment days. The available Shift resource could also shrink the typical evening generation ramp by as much as 50%, reducing the need for costly flexible generation resources.

Figure ES-2. Shift supply curves for the default scenario (Medium technology, 1-in-2 weather), disaggregated by sector (light shaded background areas) and by end use within each sector (shaded bars). BTM battery price thresholds in each year are shown for reference (gray dotted lines). Note: Residential cooling loads are referenced as "space cooling," while commercial loads are labeled as "HVAC," to distinguish between the two resources. EV charging loads are divided into level 1 (wall outlet) and level 2 (dedicated higher-amperage charging station). Agricultural pumping loads are classified under the industrial sector for the purposes of this study and make up the vast majority of the Shift resource from pumping.

2. Today's readily accessible Shift resource will need to grow to support future grid needs. We found above that the Shift resource that is less costly than BTM batteries can potentially address *average* recent curtailment levels, but there is a large day-to-day variation in renewable curtailment levels. The most extreme curtailment day in spring 2019 saw 39 GWh of curtailed renewable generation on the CAISO grid, which is more than could be absorbed by a 5.3 GWh Shift resource, even if all of it could be dispatched twice daily. Moreover, our analysis shows a substantial seasonal variation in both the size of the Shift resource and in the need for Shift, with the largest resource available in the summer, when the need is smallest. Finally, as shown in Figure ES-2, the size of the Shift resource at the battery cost threshold grows only slowly, increasing by less than 20 percent between 2020 and 2030, to 6.3 GWh. Over the same period, renewable generation is expected to grow by some 60 percent, suggesting that the potential need for Shift will rapidly outpace the supply available with the technology options and costs we modeled. As more renewable energy is added to the grid, additional Shift can be deployed alongside additional storage and curtailment to balance the power system, so it will be worthwhile to explore ways to increase the supply of Shift and bring down its costs in the future.

3. New loads from electrification are important potential sources of Shift, but their current enablement costs are prohibitive. Our analysis shows that introducing electrified residential space and water heating can both grow the overall Shift resource and also significantly reduce the seasonal variation in availability by boosting the wintertime resource. However, the estimated costs to enable these sites with current technology deployment through retrofits may currently be more expensive than investing in similarly sized battery storage. Figure ES-2 includes Shift resources arising from projected future EV charging loads. The top panel of Figure ES-3 additionally shows the Shift potential available from future electrification of space and water heating in the residential sector. In each case, the potential is small compared to the most significant other end uses shown, and the costs of enabling these Shift resources are higher than the cost of procuring BTM batteries to provide the same service. This result occurs, in large part, because of high site-level technology costs to enable Shift for these technologies, coupled with relatively small loads at individual sites. As discussed next, however, there may be significant opportunities to capture a larger Shift resource from electrification via a market-transformation approach.
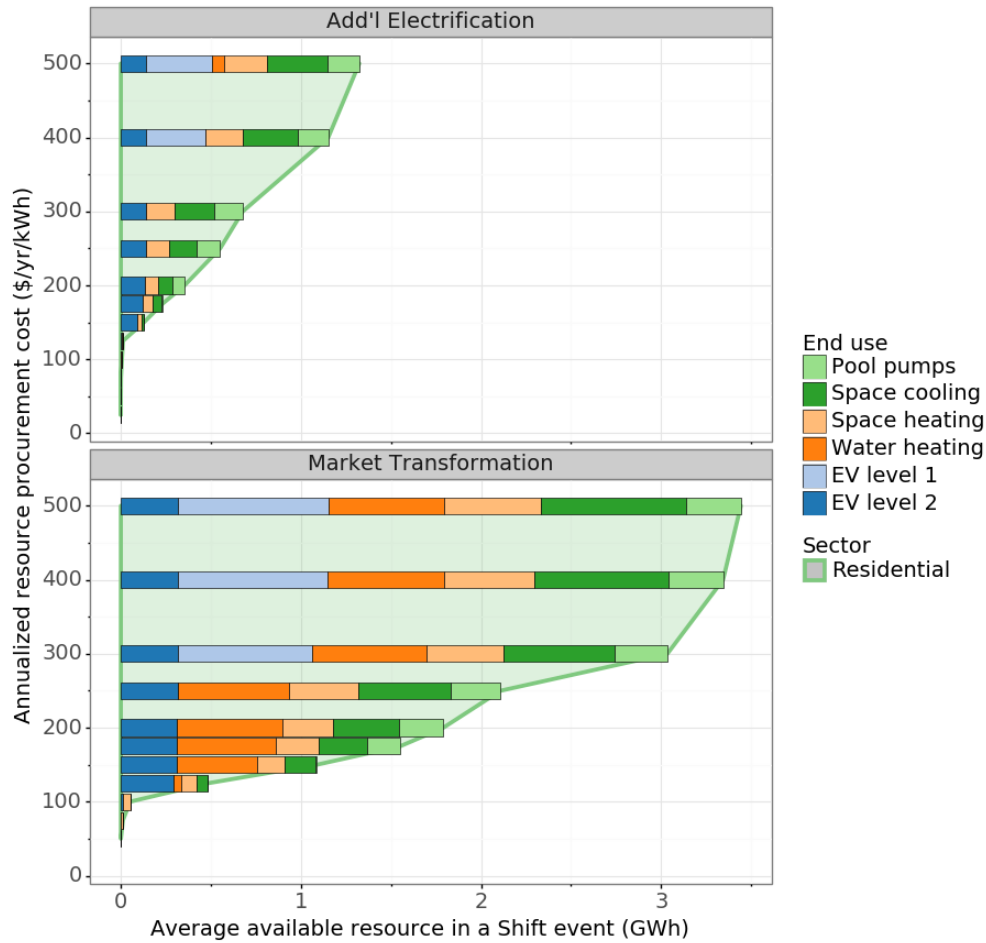
Figure ES-3. Residential supply curves for Shift in 2030, by end use, for model runs incorporating additional electrification loads, under standard modeling assumptions and in a scenario envisioning a major market transformation for Shift-enabling technologies.

4. There are numerous pathways to expand the size of the Shift resource and lower its cost. This study finds that the primary barriers to increasing the cost-competitive Shift resource are high technology costs and low expected customer participation rates used in the default modeling assumptions, which are based on present-day cost data and historical information on price trends and participation in DR programs. However, it is important to recognize that Shift represents a fundamentally new approach to DR; the novel program designs and market structures that accompany its deployment (see Gridworks 2019 for examples) may substantially transform the surrounding market in a way that can dramatically reduce costs and boost customer participation. To give a sense of the opportunity, we included an illustrative scenario envisioning a "market transformation" with significantly reduced costs to enable residential electrification loads, and increased customer participation, which greatly increases the size of the resource as shown in Figure ES-3. As shown, this scenario can triple the aggregate resource—and grow the resource from certain end uses, like water heating, by an even larger factor—while bringing down the overall cost of Shift. These results suggest that it will be important to consider regulatory, R&D,

and program-design pathways to driving down the costs of enabling technologies and increasing customer engagement in newly designed Shift programs.

An additional low-cost pathway to capturing a portion of the Shift resource is improved customer response to time-of-use rates (which are now the default for California customers) or other time-varying retail electricity rates. In the parlance of the DR Potential Study, such a long-term change in customer behavior would be categorized as Shape (whereas Shift is envisioned as a fully, dispatchable resource), but such load reshaping could achieve much of the same benefit at a lower cost. In this study we computed the load shifting impacts from a forecasted customer response to currently planned TOU rates,[5] and we found that the resource thus captured was quite small. These forecasts rely on historical rates of response, however, and a number of the Shift-enabling technologies considered in this study could significantly enhance customers' capacity to shift load in response to simple price signals, potentially lowering the cost of capturing the resource. Rates that have significantly lower prices during times of high renewable energy availability, and higher prices for the opposite conditions, could also provide stronger financial incentives to shift electricity loads than today's rates.

Recommendations and Next Steps

The most important finding in this study is that a sufficient potential Shift resource exists today to address a significant portion of the present-day challenges facing California's power system. Because Shift DR represents a fundamentally new category of DR, with no existing programs that implement it on a large scale, realizing this potential will require development of a technological and policy landscape that supports and values load flexibility. A critical immediate next step will be to initiate pilot programs to test different practical pathways to enabling Shift.[6] Additional research into the real-world capacity for load shifting by different customers and end uses could provide important information, especially in the context of battery storage and improved energy efficiency. Meanwhile, there is an immediate opportunity for new policies to drive wider adoption of flexible loads. Examples include codes and standards that encourage flexible load control at the device level; policies focused on providing real-time electricity price information to consumers; communication and control standards for flexible devices; carbon benchmarking requirements for buildings; and an improved policy analysis paradigm that integrate the co-benefits of energy efficiency, DR, and other distributed energy resources.

As the capabilities of Shift resources are better understood, it is important to consider the resource in the context of other renewable integration strategies. Evaluating the potential for Shift should be considered within the IRP and other grid planning efforts. These integrative studies would help provide context for the value of Shift compared to other approaches, including overbuilding and curtailing renewables, transmission expansion, and energy storage.

---

[5] Moreover, throughout the report, the Shift supply curve is computed *after* applying the forecasted impacts of TOU rates, since these load shifts have already been captured.

[6] The final report of the CPUC's Load Shift Working Group lays out a number of specific pilot proposals. The report can be found at https://gridworks.org/wp-content/uploads/2019/02/LoadShiftWorkingGroup_report.pdf.

To better chart a course for Shift DR in California, there are also a number of modeling improvements that will need to be considered in the next phase of the DR Potential Study. The first is the analysis of additional flexible end uses, such as residential appliances and commercial water heating, which may increase the estimated Shift potential. It will also be important to explore new strategies for enabling Shift, such as bidirectional EV charging. Given that Shift may have customer impacts that are less disruptive than traditional Shed DR, an improved model for customer participation in Shift programs will be a crucial update. Finally, as Shift begins to emerge as a real option in the resource planning context, it will be essential to have a more detailed model of how customers participate in Shift programs, how Shift would be dispatched in real time on the grid, and which other system-level investments it might replace, in order to more accurately determine its value to the grid. At present, the absence of real-world operational data on Shift presents a significant challenge to such modeling. This emphasizes the importance of research, demonstrations, and pilot programs that can generate a better practical understanding of Shift as a resource.

1. Background and Introduction

California's renewable energy goals are leading to rapid changes for electric power grid. In 2018 the grid managed by the California Independent System Operator (CAISO) was already receiving 26 percent of its power from non-hydro renewable sources, with an additional ten percent hydroelectric power (CAISO 2019a); and with SB 100, California statute now sets a target of 100 percent carbon free electricity by 2045. Increased use of variable generation resources affects the grid's operation over multiple timescales. As this transition continues, California will need to plan carefully to ensure resources with the right characteristics are available to meet changing grid management needs. Demand response (DR) has the potential to provide important resources for keeping the electricity grid stable and efficient; defer upgrades to generation, transmission, and distribution systems; and provide customer economic benefits.

This report describes the results from the third phase of research to evaluate the potential size and cost of future DR resources for the three primary investor-owned utilities (IOUs) that provide electricity service in California: Pacific Gas and Electric Company (PG&E), Southern California Edison Company (SCE), and San Diego Gas & Electric Company (SDG&E). Previous research on this topic was presented in an interim report on Phase 1 (Alstone et al. 2016) and a final report on Phase 2 (Alstone et al. 2017). Hereafter, we will refer to these previous two efforts as the Phase 1 report and the Phase 2 report (or study), respectively, and we will refer to the current work as Phase 3.

Unlike the previous two reports, this report focuses exclusively on the *Shift* concept. The Phase 2 report developed an analytic framework that grouped DR services into four core categories: **Shape**, **Shift**, **Shed** and **Shimmy.**[7]

> **Shape** captures DR that reshapes customer load profiles through price response or behavioral campaigns—"load-modifying DR"—with advance notice of months to days.

> **Shift** represents DR that encourages the movement of energy consumption from times of high demand to times of day when there is a surplus of variable renewable energy (VRE) generation. Shift could smooth net load ramps associated with daily patterns of solar energy generation.

> **Shed** describes loads that can be curtailed to provide peak capacity reduction and support the system in emergency or contingency events with a range in dispatch advance notice times.

> **Shimmy** involves using loads to dynamically adjust demand on the system to alleviate short-run ramps and disturbances at timescales ranging from seconds up to an hour.

A significant conclusion from the Phase 2 report was that the Shift resource is of great value to the current and emerging electric system. The load flexibility from Shift represents an additional and significant new potential area of VRE integration value, in addition to traditional load-shedding DR for peak load management that can help reduce the need for costly investments in generation and transmission and distribution (T&D) capacity. The goal of the Phase 3 study is to provide the California

---

[7] Throughout this report, we will capitalize each of these terms when referring to the DR category, to distinguish this usage from common usage of these terms as nouns or verbs. For example, "The Shift service consists of asking customers to shift the timing of their energy consumption."

Public Utilities Commission (CPUC) with data-driven insights evaluating how California might use Shift DR in meeting its resource planning needs and operational requirements. We address the following primary questions:

- What is the size and cost of the expected resource base of Shift?
- Where is the Shift resource located; e.g., by end use or geographical region?
- When is Shift needed and when is it available; e.g., how does the shiftable load align with the need to mitigate ramping; how does the resource vary seasonally?
- How will the Shift resource evolve with California's changing grid; e.g., in response to growing electrification or increased VRE penetration?

The analytical framework we have developed for the DR Potential Study forecasts DR supply curves for the years 2020, 2025, and 2030 using a bottom-up modeling framework called DR-Futures. The framework consists of two modules:

- **LBNL-Load** is a bottom-up load-forecasting module that capitalizes on a large set of customer smart meter data to project future end-use load shapes for a diverse set of customer clusters.
- **DR-Path** starts from the forecasted load shapes to assess a large number of future pathways to acquiring DR resources, resulting in granular load flexibility potential estimates that can be aggregated into the final supply curve.

Portions of the description of our modeling work, the characteristics of Shift DR, and results presented in this report have been presented previously in various forums. These include Phase 1 and Phase 2 reports, two conference papers at the American Council for an Energy-Efficient Economy (ACEEE) Summer Study (Alstone, Piette, and Schwartz 2018; Gallo et al. 2018), and a review article on electric load flexibility analysis (Alstone and Piette 2019). This report significantly extends on those through a broad synthesis of the need for Shift and updated results from our DR potential modeling framework for Shift.

1.1. Background on the DR Potential Study

We developed this work in the context of supporting DR research at the CPUC. This was initiated with Phases 1 and 2 of the 2025 California DR Potential Study (Alstone et al. 2017; 2016), supporting CPUC rulemaking R.13-09-011 (CPUC 2013), whose stated purpose was "Enhancing the Role of Demand Response in Meeting the State's Resource Planning Needs and Operational Requirements." The California DR Potential Study was designed with two goals: (1) to bridge the analysis of distributed energy resources (DERs) with grid investment and operations and (2) to communicate the results of the study clearly to power system policymakers and stakeholders who need to synthesize across those domains. Box 1 shows a summary of findings from the Phase 2 study.

Box 1. DR Potential Study Phase 2 Results Summary

Shape. Our Phase 2 estimates of load reshaping were based on expected response to future time-of-use (TOU) rate structures. We estimated that these prospective rates will result in reductions in the peak load equivalent to approximately 1 gigawatt (GW) (~2 percent of the overall peak), and result in ~2 gigawatt-hours (GWh) of shifted load per day through changes in behavior and schedule (~0.5 percent of volumetric energy demand). These Shed and Shift outcomes that can be achieved with a Shape pathway represent an important and foundational element of the state's DR future. The responses included in the model are only based on historic response, mostly from schedule-based and behavioral changes. As more automated, price responsive systems come online, the expectation is that the scale of impact from load shaping could increase.

Shift. We found a significant emerging opportunity to support the grid with load shifting to capture renewable generation that would otherwise be curtailed, reducing the renewable build-out (and cost) required for renewable portfolio standard (RPS) compliance. Shift can also reduce ramps between low and high demand periods, resulting in lower system operating costs. Our Phase 2 model estimated a daily quantity of 10–20 GWh of cost-effective shiftable energy (equivalent to 2 to 5 percent of daily consumption). Based on the estimated savings from avoided investment and operations costs, we expect the value of this prospective Shift resource to the power system would be $200–$500 million annually by 2030. These estimates for Shift were promising and are being updated, improved, and expanded on in this report.

Shed. Conventional DR programs have focused on Shed to manage demand peaks. In Phase 2 we found a significant role for these resources in the future, but capturing the full value requires a new focus on local capacity and distribution. Existing DR programs circa 2015 were structured primarily to meet *systemwide* peak capacity needs. Because of the significant additions of renewables to the California grid, the outlook at the time suggested little need for new thermal capacity until well after 2025. With no apparent capacity-expansion needs for DR to offset, this suggested a low systemwide value for Shed. However, recent forecasts now show capacity constraints emerging between 2019–2024 (CPUC 2019). This new reality indicates that ***there is indeed near-term value for Shed DR at the system scale,*** updating the original interpretation of the Phase 2 results.

At a conventional value ranging from $50–$100/kilowatt (kW)-year, the available resource in 2025 is 2–5 GW with an annual value between $100–$700 million. Shed can also provide significant additional value at the local level in areas with constrained transmission systems. We described the local resource in a separate technical addendum to the Phase 2 study (Alstone et al. 2017b). There also may be significant ability for Shed to serve distribution system needs and to support reliability in emergencies and contingency events. Overall, the Phase 2 analysis and recent capacity forecasts indicate a steady and growing need for the Shed resource, with an expanded focus on local needs.

Shimmy. In the Phase 2 study we estimated approximately 300 megawatts (MW) of potential for load to stabilize the grid with bidirectional, fast DR providing ancillary services (AS). The estimated value of these grid services is $25 million per year. The specific pathway to creating systemwide value for Shimmy was interestingly related to freeing batteries from the need to provide AS, enabling them to provide Shift instead. This dynamic is similar to the conventional description of price formation for frequency regulation, where the price for the ancillary service is directly related to the opportunity cost of lost revenue in the energy market. Our result reinforces the concept that fast-response Shimmy is a secondary service whose value will be related to opportunity costs for serving load or shifting energy.

Following the completion of the Phase 2 study, with its finding of significant value for Shift, CPUC issued decision D.17-10-017 in the rulemaking (CPUC 2017a), providing for the formation of the Load Shift Working Group (LSWG), to be convened by stakeholders for one year to provide "… a final report on its proposals, which will inform a future rulemaking to consider new models of demand response." Through our engagement with that working group as technical experts, we supported the stakeholder discussions with analysis of the proposed pilots and possible frameworks for implementing load shifting, including an analysis of the possible greenhouse gas (GHG) impacts of shifting load. The analysis presented in this report is responsive to the need for a better and deeper understanding of the Shift resource.

1.2. A brief introduction to Shift

As discussed above, this report focuses on developing a better understanding of the potential for Shift as a resource on the California grid. Because Shift represents a substantially new approach to DR compared to historical practice, this section provides a brief primer on the Shift concept. We start by describing the challenges facing California's grid that Shift can address, then describe what Shift is and discuss the times of day in which grid operators may need to rely on its service. We then examine how the Shift potential can vary across different end uses, and how this potential varies throughout the year. We finish by summarizing the emerging and fast-moving technological trends that may affect California's Shift resource in the coming years.

1.2.1. Renewable grid integration and the need for Shift

As California adds VRE resources, there are new challenges introduced to grid balancing authorities like CAISO arising from the intermittency of the resources. Since there is a rigid need to balance supply and demand at all times, the dispatch of flexible generators and other resources must match not only the demand but also the net effective load that includes the natural cycles and random variability of solar and wind power.

The new challenges of operating the grid with renewables have been typically symbolized in California policy discourse by the "duck curve," initially predicted by CAISO (CAISO 2016). Figure 1-1 (based on actual CAISO operational data from recent years) shows how these predicted duck curves are already common in operational situations (solid black lines). The basic concept behind the duck curve is that one can take the total electricity demand in each hour and subtract the VRE generation that is online in that hour, to find the net load that needs to be served by other generators. This net load has sharp peaks in the morning and evening (the "tail" and "head" of the duck, respectively), a steep downward and upward ramping characteristic from solar generation coming online in the morning and ceasing with sunset in the evening, and a deep trough (the duck's "belly") during the midday hours. The overall effect is that renewables have significantly reduced and delayed the daily peak load into the evening hours evening and added steep new ramps in the net load in the morning and (especially) evening. Throughout the day, additional solar and wind power adds to the short-run variability on the grid as well. The steep ramps and sharp peaks in the net load present particular challenges to a dispatchable generation fleet whose lowest-

4

cost resources (typically natural gas combined cycle plant) are often unable to vary their output sufficiently quickly to accommodate these features.
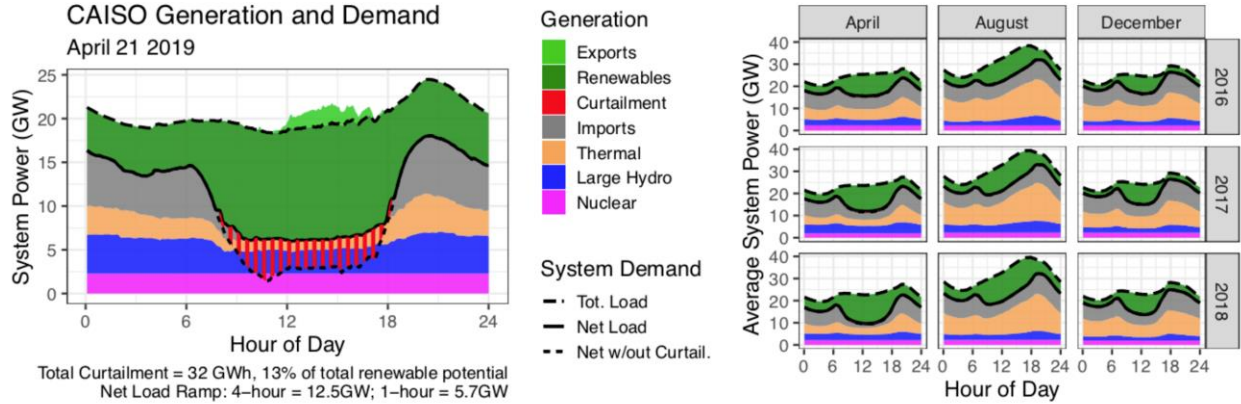
Figure 1-1. CAISO operational profile from April 21, 2019, a representative day from recent operations with high curtailment (left), and average operations for three months in 2016, 2017, and 2018 showing the daily profile for the average of all days in the month (right). Both plots show generation by source, and indicate demand based on different contributions from renewables with and without curtailment. Source: CAISO Managing Oversupply web page (CAISO 2019b).

The large difference in net load between the "belly" and "head" of the leads to renewable generation curtailment in the middle of the day due to operational constraints on the grid. The morning and evening net load peaks must be met by a combination of inflexible base generation (such as nuclear generation) and dispatchable generation (such as natural gas or hydroelectric generation). Most dispatchable generation units have non-negligible startup times and must generate some minimum amount of power in order to remain in operation once started and synced to the grid.[8] Thus there is a minimum amount of non-renewable generation that must be online during the midday hours in order to meet the morning and evening peaks. When this minimum quantity, plus the total VRE generation, exceeds the total load on the grid and potential energy storage charging, the best available balancing option is to curtail renewables by disconnecting them from the grid.[9] Thus the level of curtailment is indicative of the timing and scale of the opportunity for Shift to make use of zero marginal cost surplus generation that is increasingly available.

Figure 1-1 (left panel) shows the hourly generation, by source, on a day with a large amount of curtailment in the spring of 2019 (red hatched region). Between 10 am and 5 pm, each of the non-renewable generation resources reaches its minimum generation level, but the gross demand on the grid is not sufficient to utilize all of the VRE generation, resulting in substantial curtailment. The left panel of Figure 1-1 shows a particularly severe example compared to typical present-day curtailment, but average curtailment has increased rapidly within CAISO in recent years, as shown in Figure 1-2. This chart shows

---

[8] Some hydroelectric powerhouses can sync to the grid without generating, but these may have additional constraints such as minimum environmental flow requirements, or (in wet years like 2019) the need to move water through the system to accommodate future flows.

[9] Figure 1-1 shows an additional option, which is to export power outside of the balancing authority (light green region), but in practice this only broadens the scope of the balancing problem and, as shown in the figure, often will not eliminate it.

that there is significant seasonal variation in curtailment rates, with the lowest amount in summer when gross system demand is highest. But there has been an accelerating rise in curtailment during the winter and spring months, spiking most recently to an average daily curtailment of more than 5 gigawatt-hours (GWh) in the spring of 2019. We can estimate the value of this curtailed generation at $40 per megawatt-hour (MWh), which is the estimated levelized cost of energy for new solar projects reported by the U.S. Energy Information Administration (EIA) (EIA 2019). With that assumption, the curtailment that occurred on the CAISO grid in spring 2019 represents roughly $19 million in potential value from Shift to utilize the available renewable energy.

Figure 1-2. Average daily curtailment of renewable energy in CAISO, by season, from 2014 through 2019. Source: CAISO Managing Oversupply web page (CAISO 2019b).

As even more VRE generation is brought online, the operational constraints of the conventional power system will result in growing opportunities to capture surplus renewable energy. These dynamics present an opportunity for flexible loads to shift the timing of demand to better align with VRE generation, relaxing the constraints on the power system and enabling reliable operation at higher renewables deployment levels. This kind of load shifting would have two primary benefits in the context of the California duck curve:

1. Avoid VRE curtailment, in essence by "filling in" large troughs in the net load, when it is cost-effective to do so, thus increasing the capacity factor of existing renewables.

2. Ease ramping rates and flatten extremes of demand by reducing load during peak periods (similar to Shed DR) and instead using the same energy during periods of low net demand (unlike Shed), reducing the need for conventional power plants and energy storage to meet these needs.

Annual avoided curtailment with Shift of 'x' GWh per day

Shift Resource Size
- x = 100 GWh per day Shifted
- x = 20 GWh per day Shifted
- x = 10 GWh per day Shifted
- x = 5 GWh per day Shifted
- x = 1 GWh per day Shifted

Curtailment Data from CAISO
Assumes all Shifted Energy avoids curtailment
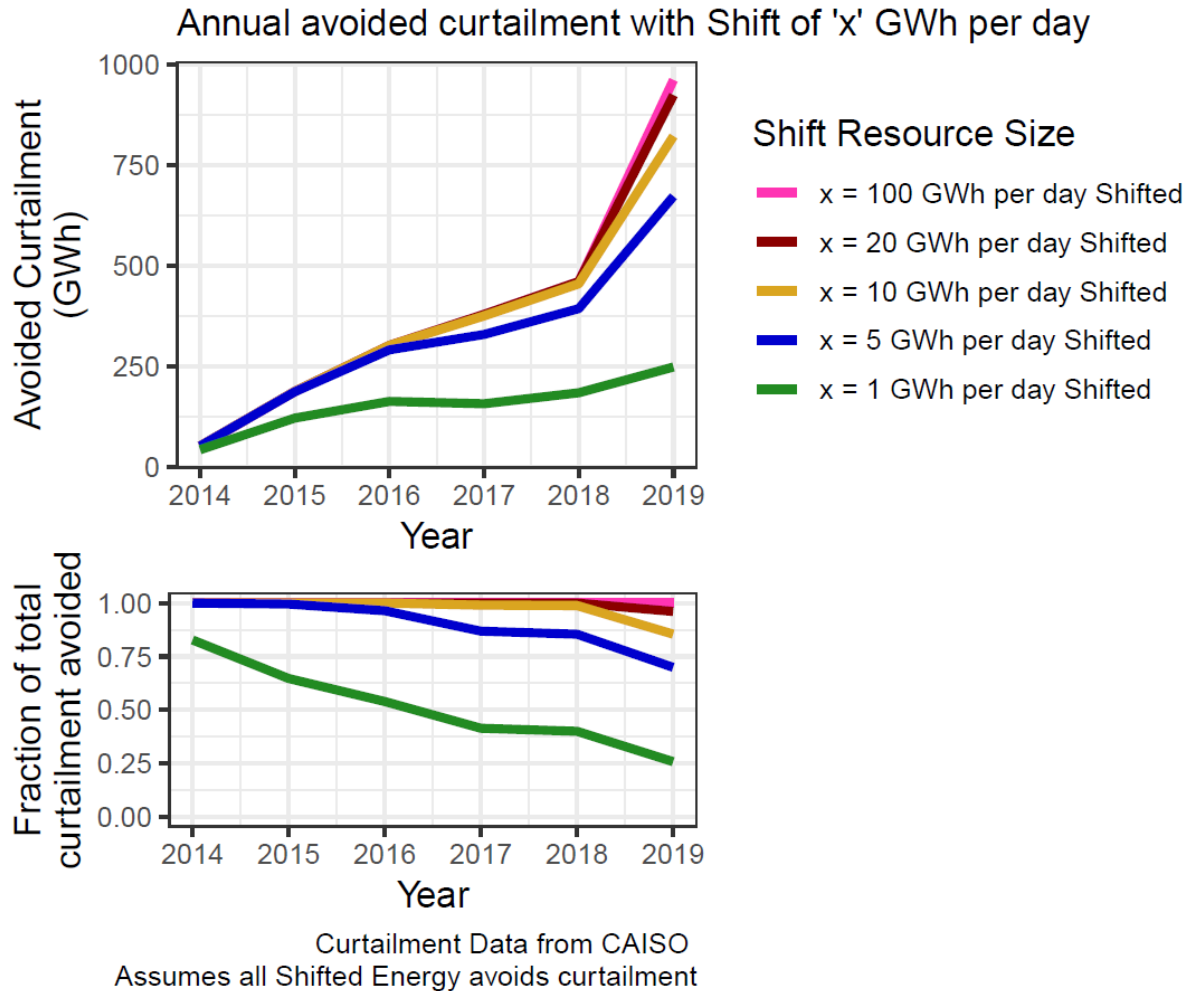
Figure 1-3. The amount of annual renewable curtailment in CAISO, both absolute and as a fraction of the total, that could have been avoided with a Shift resource that can shift a given quantity (denoted "x") of energy per day in a manner that could utilize the curtailed energy.

Figure 1-3 shows an approximation of the impact that Shift could have in the first case: alleviating curtailment of renewables. The figure considers the history of hourly curtailment in CAISO from 2014 through 2019 and estimates the amount of curtailment that could have been avoided in the presence of Shift resources of different sizes, assuming that all shifted energy could be used to offset curtailment.[10] In this picture, a Shift resource that amounts to 10 GWh/day of shifted energy, represented by the yellow

---

[10] There are important technical challenges that this assumption would face in practice. In particular, a substantial fraction of curtailment in CAISO is driven by transmission constraints, not systemwide overgeneration. In these situations, Shift resources would have to be geographically situated in the constrained areas to be of use, but these areas also tend to be sparsely populated, suggesting that sufficient Shift resources may be difficult to obtain.

line, would be sufficient to offset nearly all curtailment that has occurred historically in CAISO. However, the size of the needed resource is growing each year, and therefore in 2019 the same resource only offsets around 80% of the curtailment. The underlying dynamics result from the fact that daily curtailment is a "positive skewed" process, with many days having modest curtailment but some with much higher levels (i.e., long-tail "outliers"). Similar to the way that it is uneconomic to build enough battery storage to capture every possible kWh of renewable generation, it is also the case that installing shift technology will be economic for capturing the typical available surplus energy but is likely to fall short of utilizing all of the surplus on extreme days. As the VRE resource in CAISO grows substantially in the near and medium-term future, there will be further opportunities to grow the capacity of a Shift resource to aid in balancing the grid.

In addition, rapid ramping of flexible resources occurs frequently throughout the year—even on days without significant curtailment—in response to changes in VRE generation at sunrise and sunset (as shown in the right panel of Figure 1-1). By helping to flatten these ramps, Shift could provide value to the grid on nearly every day of the year, since it can eliminate the need for some flexible generation capacity. The value may be greatest, however, on days with substantial renewable curtailment, as shown in the left panel of Figure 1-1, or with especially large evening peaks in the net load, since this affects the total amount of generation infrastructure that is needed to provide adequate capacity.

1.2.2. What is Shift?

Shift refers to an approach to DR in which electricity consumption is moved from one time to another to meet specific grid needs (e.g., to align with available VRE generation and mitigate ramping). At the site level, Shift is manifested as a load reduction (a "shed") over a particular period of time, coupled with a load increase (a "take") at a different time beforehand or afterward.[11] From the system operator's perspective, Shift roughly resembles a virtual fleet of behind-the-meter (BTM) batteries that "charge" during the take part of the cycle and "discharge" during the shed (although the detailed operational constraints on Shift, such as the allowable dispatch frequency and storage period, will not be identical to those for batteries). Depending on the end uses involved and power system conditions, the shed/take cycle of Shift can occur over a period of a few hours, over a diurnal (AM/PM) cycle, or even across several days (e.g., shifting peak weekday loads to the weekend). In this study we focus only on load shifting that occurs over a period of 12 hours or less, which we expect to be both the most straightforwardly achievable and the most frequently utilized form of Shift. Moreover, for many shiftable end uses, a multi-hour shift can often be executed with a minimal impact on the customer's perceived level of energy service: for instance, delaying the operation of an appliance for a few hours, or pre-cooling a building to enable an evening load reduction, may not have noticeable impacts on customer convenience or comfort.[12] This is in contrast to traditional Shed DR, which reduces loads without specifying an offsetting

---

[11] It is worth noting that load sheds in traditional Shed DR may also be partially offset by a take at a different time to make up the lost energy service (so-called "snapback"). One of the primary aims of Shift is to change the timing of loads in a controlled manner to ensure that the shed/take cycle serves to mitigate, not exacerbate, the duck curve.

[12] For example, a 2013 study (Herter and Okuneva 2013) on residential pre-cooling in the context of Shed DR found that for various pre-cooling periods (two or six hours before the peak) there were significantly increased loads in

recovery time; this distinction may make Shift less burdensome to customers, potentially boosting customer participation.

Broadly speaking, the kinds of load shifting that can mitigate the California duck curve can be achieved in two ways: *permanent* load reshaping over the long term and *dynamic* load shifting in real time. The permanent approach includes the kind of long-term changes in customer behavior that might accompany a switch to a TOU tariff; in the context of DR bifurcation, it is load-modifying DR. The dynamic approach involves loads that can be shifted in response to system needs in real time or near-real time (e.g., day-ahead dispatch), either by direct dispatch by the system operator or in response to real-time fluctuations in energy prices. In the context of bifurcation, supply-side resources would fall into this category (as would load-modifying resources that involve exposure to wholesale electricity prices).

In the context of the DR categories used in the DR Potential Study, only the dynamic approaches would be categorized as Shift, whereas the permanent load reshaping approaches would be categorized as Shape. Specifically, in the modeling for this study, we conceptualized Shift as a *supply-side* resource, i.e., a dispatchable resource that is fully integrated into the CAISO wholesale market to support the bulk power system.[13] This approach would require control mechanisms with a high degree of reliability, as well as metering and telemetry of sufficient accuracy to support market settlement, both of which introduce additional costs beyond the basic technological requirements to achieve real-time load shifting. Such a resource would typically also participate in the market as part of an aggregation, with the aggregator controlling the loads for multiple sites in response to ISO dispatches and paying a participation incentive to each customer in the aggregation group. The incentive payments in this arrangement also increase the total cost of the Shift resource.

By including all of these costs, our modeling of the Shift resource represents a conservative accounting of the cost to acquire the resource. It is important to emphasize that much of the Shift potential may be available for lower costs than presented here, either by directly exposing loads to wholesale electricity prices or via Shape-based approaches such as TOU tariffs.[14] Additionally, our modeling considers Shift as a resource for the bulk power system; there may be additional cost-competitive approaches to enabling load-shifting to support the distribution system which may also have positive impacts for the transmission system (see Gridworks 2019 for examples); these are outside of the scope of this study.

---

mid-day and reduced load at peak times, resulting in a total shifted energy of ~3 kilowatt-hours (kWh) per site. At the same time, metrics tracking customer comfort were significantly higher for pre-cooled buildings than for buildings that did not pre-cool.

[13] Although there is currently no CAISO market product meant to accommodate load shifting, the final report of the Load Shift Working Group (Gridworks 2019) identifies several viable pathways to market-integrated load-shifting DR. Importantly, the report also identifies several pathways to Shift that are not market integrated and may be achievable at lower cost than the modeling here indicates.

[14] More broadly speaking, our cost accounting applies all DR procurement costs to procurement of the Shift resource (with the exception of customer-borne costs that achieve co-benefits); total procurement costs for Shift could be lower in the context of multi-use DR applications (e.g., DR technologies that enable both Shift and Shed resources).

1.2.3. When is Shift needed?

To conceptualize Shift as a dispatchable resource, we can assume that Shift resources would be responding to Shed and Take dispatch signals that are the inverse of Generation-Up and Generation-Down dispatch signals communicated by the grid operator.[15] Figure 1-4 shows an idealized Shift dispatch curve throughout the course of a day. The curve partitions the day into segments that we refer to as *Shed periods* and *Take periods:* when the dispatch curve is positive Shift loads are instructed to take additional load, and when it is negative they are instructed to shed load. Each zero crossing of the dispatch curve (red points) represents the central hour of a potential multi-hour load-shifting event.

Figure 1-4. An idealized daily dispatch signal to which Shift resources would respond. Positive values of the dispatch indicate the need for resources to take load; negative values request a shed. To estimate the Shift resource, the DR-PATH model identifies loads that could be shifted within multi-hour windows centered on the zero points (red dots).

Throughout this report, we will use the terms *event* or *dispatch* (used as a noun) to refer to a multi-hour period centered on one of the zero crossings of the dispatch curve (red points), in which customers would be instructed to reduce load during the Shed period and take a corresponding amount of load during the adjacent Take period (which may occur either before or after the Shed period, depending on system conditions).[16] As shown in Figure 1-4, Shed periods typically occur during periods of peak net demand in the morning and evening, when demand is significant but solar generation is not available. In contrast,

---

[15] These signals instruct generation resources to either increase or decrease generation. From a DR perspective they would instruct loads to reduce or increase load, respectively.

[16] In practice, the dispatch signal used to achieve this response could take many forms in addition to generation up or down signals, including real-time price signals, or a pure Take signal with the assumption that offsetting load reductions would occur in response. Understanding the response to such signals is complicated by questions of customer behavior, so for the sake of concreteness we model Shift as a traditionally dispatchable resource here, while acknowledging that some portion of the resource identified could be obtained via alternate program designs.

Take periods occur in the middle of the day, when solar generation is strong, and overnight, when loads are minimal but there is often significant wind generation. (Because of wind's variability, however, an overnight Take period does not necessarily occur every day.)

1.2.4. Which end uses can provide Shift?

To serve as a Shift resource, an electrical end use must satisfy two criteria: it must have an inherent flexibility in timing, and it must be aligned with the system's need for load shifting. A wide variety of electrical end use categories meet the flexibility criterion, including space conditioning and ventilation, water heating, pumping, battery charging, industrial processes, and end uses with fixed-period cycles (e.g., dishwashers). This study considers a subset of these, limited by the ability to disaggregate those loads from the historical meter data and by the available information available to model the end-use technology. The alignment criterion requires that potential Shift resources have a significant amount of load that would typically occur during times when there is need to reduce the system load, along with an ability to increase load at times when there are excess or low-cost generation resources available. We discuss this latter criterion in the next section.

Figure 1-5. Annual average daily load shape in 2025, as forecast by LBNL-Load, with analyzed Shiftable end uses disaggregated (colored bands) and set in the context of the daily Take/Shed cycle from Figure 1-4.

1.2.4.1. The hourly availability of Shift

Figure 1-5 shows the 2025 forecasted average daily statewide load shape from the Phase 2 study, disaggregated into estimated end-use contributions for the various shiftable end uses considered in that study, overlaid on the Shed and Take periods identified in Figure 1-4. The forecasted annual average system level net load curve is also shown for reference, displaying the steep morning and evening ramps, and low mid-day demand, that are rapidly coming to characterize the CAISO system. As shown in Figure 2-3, the heating, ventilation and air conditioning (HVAC) and electric vehicle (EV) charging loads

have natural peaks near the Shed/Take transition hours, which we identified above as the likely central hours of multi-hour Shift dispatches. Hence these loads may have a high Shift potential. By contrast, pool pump loads have their peak during the middle of the day, where it is already largely aligned with system needs. Industrial process and pumping loads have flatter profiles, suggesting that they may have potential to shift a portion of their consumption during times of system need. In essence, the opportunity to participate in Shift comes from a misalignment between the "natural" timing of loads that is incumbent in the power system and the emerging dynamics of VRE generation.
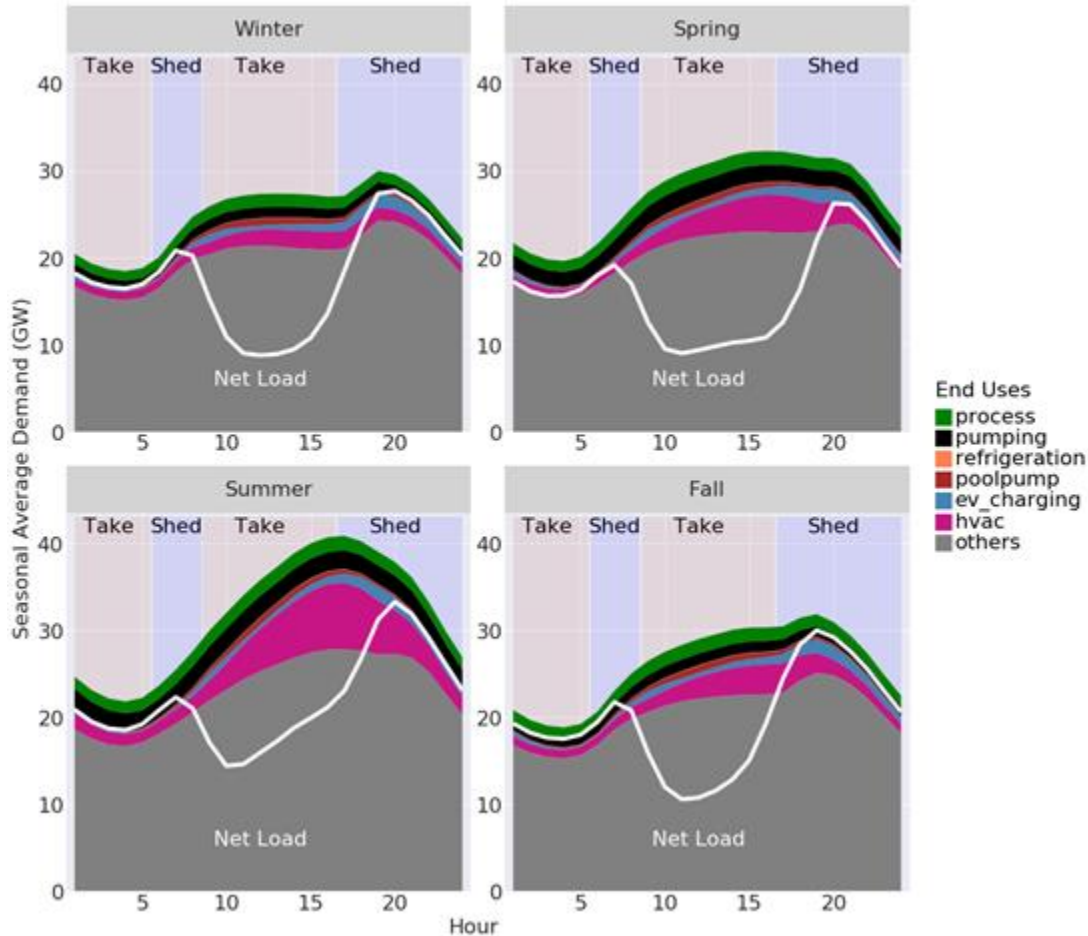
Figure 1-6. Seasonal average daily load shapes for 2025, as forecast by the LBNL-Load model in the Phase 2 study. The loads are broken out by end use and set in the context of the daily net load (white curve) and Shed/Take periods.

1.2.4.2. Seasonal variation in the Shift resource

Just as the potential for a given end use to provide Shift depends on its daily load profile, there is potentially also a strong seasonal component to each end use's value as a Shift resource. Figure 1-6 shows the average system load-shape disaggregation from Figure 1-5, broken out by season. A strong seasonal variation is apparent in the amount of shiftable load, driven primarily by seasonal variation in the HVAC load. The Phase 2 study reported annual averages for Shift DR resources, but when planning for the deployment of Shift DR resources, it will be essential to also account for seasonal variation, since the largest available end uses may not align with times of critical system need. For example, solar curtailment

and ramping may be highest in the winter and spring, when there is relatively high daily insolation but low temperatures limit the total mid-day HVAC load (as indicated by the steeper average net-load ramps in those seasons shown in Figure 1-6). However, the reduced HVAC load in this time of year also means there is less Shift resource available than is available in summer. In this study we developed new modeling capabilities to quantify this seasonal variation more fully.

1.3. Shift and greenhouse gas emissions

Today and in the future, the need to reduce GHG emissions like carbon dioxide (CO_2) means that more zero-emissions resources are needed, changing the dynamics of net load on the grid. At least conceptually, shifting the timing of loads should enable more efficient use of existing variable solar and wind generation; however, there is also a risk in principle of shifting loads in a manner that increases emissions. Understanding the specifics of how much and under what circumstances Shift can help with reducing GHG is important for charting a pathway forward to implementation.

When loads are shifted to better match VRE generation, what would we expect?

- **Reduced cost** to serve load from arbitrage in prices and increased load in low price periods.
- **Reduced peak** loads since "take" tends to happen mid-day and "shed" in peak times.
- **Reduced curtailment** from more demand during curtailment hours to utilize surplus energy.
- **Reduced emissions** from avoided or reduced demand at high emissions times.

We completed a first-order analysis of the plausible links between Shift and GHG from electricity generation as part of analytic support provided to the CPUC Load Shift Working Group as part of the Phase 3 research effort. The group discussed a range of pathways, with various program designs for implementing Shift, and signaling needs to shiftable loads through incentives and dispatch. Our analysis modeled the possible effects of load shifting on the grid using operational data from 2017. It also looked at imposing relatively small (~1 percent of load shifted) and plausibly defined load impacts from different Shift programs designed to address metrics related to the list of possible value streams listed above (e.g., marginal cost of energy, estimated marginal emissions).

Among these pathway options were two distinct approaches: (1) a market-integrated approach to bidding dispatchable loads into the energy market along with other resources, and (2) a load-modifying approach where other incentives or out-of-market instructions lead to shifted loads, like prices, estimates of environmental impact, and target load shapes from distribution system operators. The details of the approach and assumptions are described in Appendix A of this report, along with more synthesis of the results.

In summary, our analysis indicated that for the range of Shift pathways under consideration by the LSWG, the expected impacts of Shift on GHG emissions were nearly always towards lower emissions. With ~1 percent of load shifted or available to shift in 2017, about 120–150 GWh of total curtailment may have been avoided (approximately 50 percent of overall curtailment), thus replacing that quantity of non-renewable generation with zero-carbon VRE. This outcome is similar across the range of load-modifying and market-integrated options we considered. The typical CO_2 savings per shifted megawatt-

hour were 0.10–0.20 tons/MWh. As a point of reference, the typical CO_2 emissions from a natural gas power plant are 0.3 tons/MWh (see Appendix A of this report). This implies that each quantity of shifted energy consumption tended to reduce the emissions from that particular energy service by about half. In addition to these GHG benefits, we identified value from reducing the cost to serve loads and reduced annual peak demand (for most pathways considered).

1.4. Modeling load flexibility in a fast-moving technology landscape

When it comes to energy technologies for decarbonization, California is moving fast on many fronts, with rapidly growing penetration of renewable generation, storage technologies, and electrification of new end uses, all driven by a robust policy portfolio at the state, regional, and local levels, focused on shrinking the state's GHG footprint. At the same time, the landscape of technologies that can enable demand flexibility also is evolving rapidly, with a profusion of new control technologies entering the market, including networked "smart" products such as thermostats, appliances, and EV chargers; as well as technologies supporting new business models for BTM storage, including both chemical (battery) and thermal storage. Because of the central importance of these new technologies to the future of California's demand flexibility, we took care to consider their impacts in this study, analyzing the impacts of building and vehicle electrification, chemical and thermal storage, and networked controls on the potential for Shift.

Because many of the relevant technologies are still in the early-adoption phase, however, there is significant ongoing evolution in the basic design and operation of the devices. Moreover, Shift DR itself represents a fundamentally new approach to implementing and accessing flexible loads on the grid, which may create new market opportunities for many of these technologies. In such a dynamic state of affairs, it can be extremely challenging to accurately forecast future trends in technology price, performance, and adoption. Typical forecasting approaches consider incremental improvements based on historical business-as-usual (BAU) trends, but the market for Shift-related technologies in California is primed for repeated disruption in the coming years, which may result in dramatic and qualitative changes in price, performance, and adoption.

In this study, we base our primary results on official forecasts published by state agencies such as the California Energy Commission (CEC), coupled with relatively conservative approaches to forecasting future trends in technology price and performance, grounded in historical trends for mature technologies. This approach has the advantages of being both measured and well aligned with other forecasting efforts in the state. In the current California market environment, however, this approach also risks significantly underestimating future opportunities for Shift. Therefore, to assess the scale of what the conservative approach may be missing, we also model alternative scenarios involving increased electrification of space and water heating loads and a dramatic transformation of the market for Shift, with markedly lower technology costs and increased customer participation in load shifting programs. This scenario-based approach allows us, in what follows, to paint a fuller picture of the possible future pathways to Shift in California.

2. Study Scope and Methodology

This section gives an overview of the fundamental approach underlying the DR Potential study, including guiding principles, conceptual boundaries, analytical methodologies, and supporting data. We first lay out the geographical and technological scope that bound the overall study. We then summarize the DR-Futures modeling framework that provides the core analytical capacity to support the study, and we describe the primary analytical outputs that we will use to assess the potential for Shift. We finish with a high-level overview of the supporting data that feeds into the modeling effort.

2.1. Scope of the DR potential study

2.1.1. Region and population covered by the study

A primary goal of the California DR Potential Study is to assess the potential future DR resource from sources that fall within the CPUC's purview. To this end, the study covers the geographical region and customer population falling within the areas with electrical service provided by one of the three primary California investor-owned electric utilities: PG&E, SCE, and SDG&E. This includes the territory served by these utilities as of the beginning of calendar year 2014, which is the baseline year for which input data were collected for this study. The region covered by the study includes any areas where wholesale electricity supply is procured by a community choice aggregation (CCA) utility, so long as the actual electrical connection to buildings is maintained by an IOU. This region is approximately the same as the footprint covered by the CAISO wholesale market within the state of California. It does not include the territories of publicly owned utilities (POUs) such as the Sacramento Municipal Utility District (SMUD) or the Los Angeles Department of Water and Power (LADWP). Within the bounds of this territory, the study considers loads arising from utility customers within all customer classes and sectors, including large, medium, and small customers in each of the residential, commercial, and industrial sectors. Throughout this report, agricultural customers are classed within the industrial sector.

2.1.2. Guiding principles for forecast modeling

This study depends on forecasting the growth of different loads within the regional boundaries defined above, to estimate the demand that is available to participate in DR, as a means to inform future policy development. To maximize consistency with other energy policy efforts in California as much as is possible, we base our primary results on official statewide forecasting documents produced by the CEC, typically undertaken under the broad umbrella of the Integrated Energy Policy Report (IEPR), using the versions of each document that were current at the time the analysis was undertaken. Where it is necessary to make a choice between different official forecast scenarios, we tend to use central estimates (e.g., the Mid scenario from a set of High, Mid, and Low options).

Where no appropriate official forecasts are available, we rely on published literature and modeling approaches that attempt to model future scenarios that are consistent with current California policy mandates on emissions and renewable energy generation. We take this approach when modeling

alternative scenarios involving rapid electrification of certain building end uses, which are not included in the most recent official forecasts.

The study also depends on projections into the future of price and performance for DR-enabling technologies. In general, we take a conservative outlook when developing these forecasts, assuming that price and performance will improve slowly and incrementally from the status quo at a rate of a few percent per year, based on historically observed trends for mature technologies. For currently emerging technologies we allow for an accelerated rate of price decline out to 2025, but we largely constrain these drops to be a few tens of percentage points at most, to ensure that our forecasts remain conservative. In an exploratory alternative scenario, we explored the impacts that more dramatic price declines might have on our results, to understand how a fundamentally transformed market for DR in California might look.

2.2. The DR-Futures modeling framework

The California DR Potential Study models future potential for deployment of Shift DR and other California grid DR resources using a bottom-up modeling framework called DR-Futures. This framework leverages a large and rich dataset provided by the three California IOUs, including detailed customer-level data on hourly electricity consumption for hundreds of thousands of customers in calendar year 2014, as well as anonymized demographic information on the full set of roughly 11 million IOU customers.

The DR-Futures framework is subdivided into two core analytical modules: LBNL-Load and DR-Path. We briefly describe these two model components in the next section, and outline the modeling updates that were implemented for Phase 3. LBNL-Load and DR-Path are described in more detail in Appendix B and Appendix C of this report, respectively, as well as in the Phase 2 report. The modeling approach has undergone a number of changes and updates since the Phase 2 report; these are summarized in Box 2, with references to the appropriate sections for more detail.

2.2.1. LBNL-Load

The LBNL-Load module is a "bottom-up" approach for forecasting baseline end-use electricity loads. It capitalizes on IOU-provided demographic data for the full set of more than 11 million utility customers and hourly load data for 220,000 customers across the three California IOUs. Using these data, we develop approximately 3,500 representative customer clusters, each of which is characterized by a typical demographic profile, location, peak demand, and total annual energy consumption. Each cluster's 2014 total hourly consumption is estimated from the available load data, then disaggregated into its constituent end uses based on temperature sensitivity analysis and engineering estimates. These end-use baseline load shapes are then forecasted for several future years in various weather and electrification scenarios. The resulting baseline load shapes are the key input to DR-PATH, our techno-economic model for demand response. Figure 2-1 shows an overview flowchart of this methodology. A full description of the data sources, assumptions, and methodologies used in LBNL-Load in Phase 2 of this study are given in Appendix C of the Phase 2 report.

Box 2. Summary of analytical updates in Phase 3

We made numerous updates to the DR-Futures modeling framework. These changes are summarized below, with references to the relevant sections describing them in more detail.

- The LBNL-Load load forecasting inputs were updated to use more recent demand forecasts from the CEC Integrated Energy Policy Report(Appendix section B.1).
- LBNL-Load was updated to include new forecasts for EV charging loads (Appendix section B.2.1), as well as forecasts for loads from electrification of residential space and water heating consistent with state climate goals (Appendix section B.2.2).
- The DR-Path inputs were updated to use renewable generation assumptions based on the latest CPUC forecasting tool, as well as to modify the market evolution parameters used in the various technology scenarios (Appendix section D.2.2).
- The DR-Path model was substantially rewritten with a modern database back-end, to facilitate faster and more nimble calculation (Appendix section C.2.1).
- DR-Path was upgraded to utilize a more refined conceptualization of Shift as a dispatchable resource (Section 1.2.3), with a new framework for modeling Shift dispatch probability (Section 2.2.2.1.1 and Appendix section C.2.5).
- The outputs of DR-Path were modified to report the Shift resource as a shiftable energy budget available in a single Shift event, not over the course of a whole day as in Phase 2 (Section 2.2.2.1.1 and Appendix section C.2.6; see also Section 2.2.2.2.2 for a discussion of the implications of this change for interpreting results).
- DR-Path was upgraded to enable the calculation of DR resources on seasonal time frames, as well as annual (Section 2.2.2.1.3 and Appendix section C.2.3).
- DR-Path was upgraded to allow limitations on the temporal direction of load shifting by certain end uses (Section 2.2.2.1.2 and Appendix section C.2.4.2).
- The DR-Path technology inputs were updated to include new assumptions for several DR-enabling technologies, including EV chargers, residential water heaters, and the addition of thermal energy storage as a new enabling technology for commercial space conditioning loads (Section 2.2.2.1.4 and Appendix section D.2.3).
- The approach to modeling BTM battery storage was revised to allow for the possibility of procuring arbitrarily large batteries for load shifting (Section 2.2.2.1.5 and Appendix section C.2.9).
- The assumed impacts of time-varying tariffs in DR-Path were updated to utilize CEC forecasts of the load impacts of residential TOU rates (Section 2.2.2.1.6 and Appendix section C.2.8). Since TOU rates are now the default in California, these impacts were also taken to be part of the baseline load, rather than a potential source of Shift (see Section 2.2.2.2.2 for the implications on the results presented here and Section 3.1.3 for a discussion of the implicit Shift resource from TOU rates).
- Because of the changes in the way the Shift energy budget was calculated, it was not possible to follow the Phase 2 approach to computing the system value of Shift. This study focuses on assessing the size of the Shift resource and leaves valuation to future work.
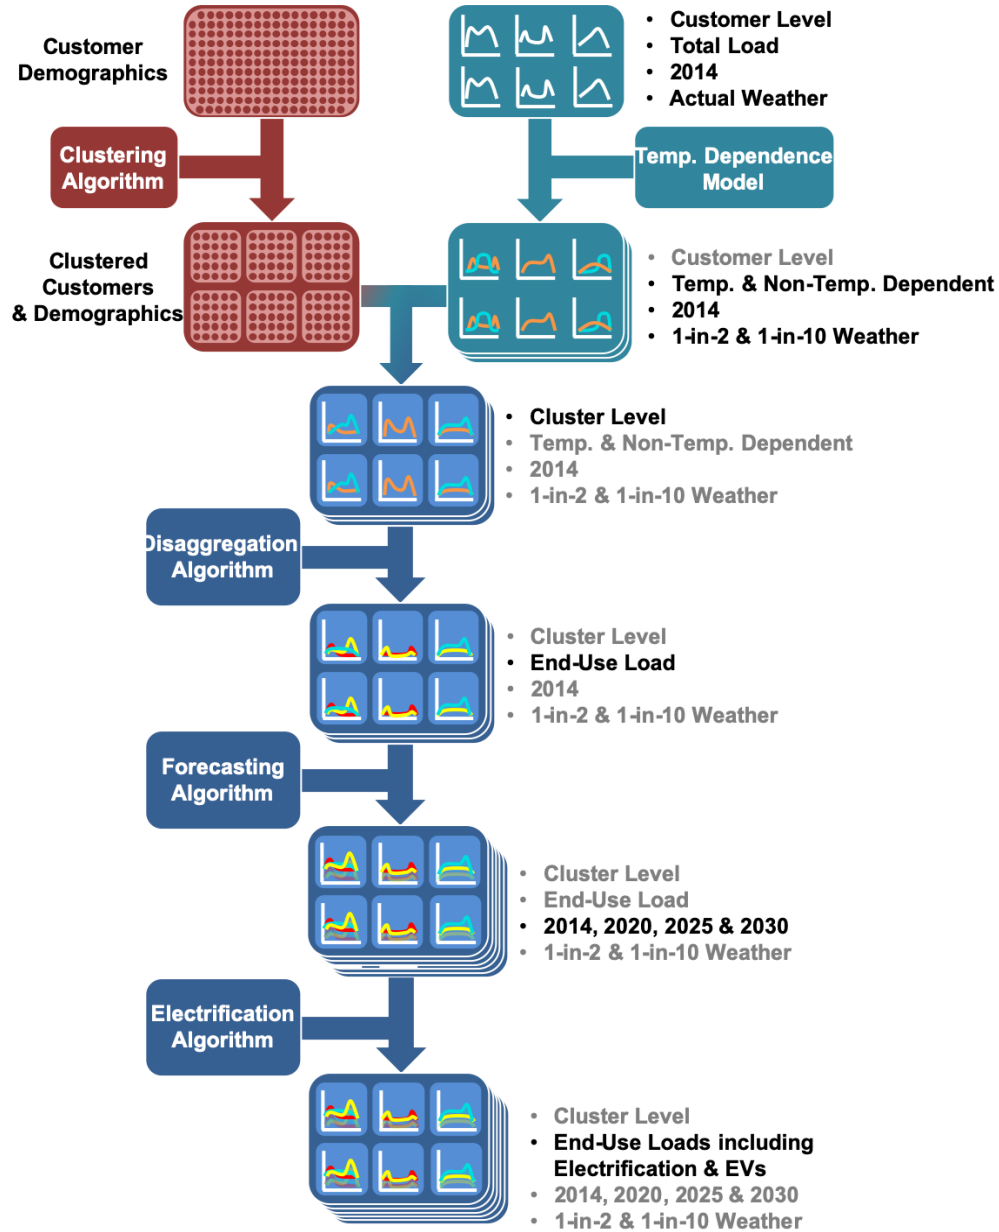
Figure 2-1. Overview of the LBNL-Load methodology for developing customer clusters and forecasting disaggregated end-use load shapes for each cluster.

2.2.1.1. Updates for Phase 3

The current study (Phase 3) involved numerous updates to the LBNL-Load model. These updates are described in detail in Appendix B of this report. Major changes include use of the electricity demand forecasts (Kavalec et al. 2018) from the 2017 IEPR (CEC 2017b) to update the load forecasts and extend them out to 2030, updated assumptions and input data for modeling the growth of EV charging loads, and the addition in an alternative modeling scenario of additional new loads from the electrification of certain building end uses (residential space heating and residential water heating).

2.2.1.2. Overview of modeling outputs

The LBNL-Load module yields a large set of representative customer clusters based on demographic, geographical, and energy consumption information. Table 2-1 summarizes the distribution of the resulting clusters by sector and building type. For each cluster, the module generates a set of forecasted hourly load shapes for forecast years 2020, 2025, and 2030. Load shapes are generated for two different weather scenarios: a 1-in-2 (typical) weather year and a 1-in-10 (extreme) weather year. Load shapes are also generated for two different electrification scenarios: a base scenario, including only additional load from EV charging, and an additional electrification scenario, which also includes forecasted electrification of residential space and water heating. Because we lack sufficient data to forecast the additional electrification scenario in different weather years, this scenario yields forecasts for only the 1-in-2 weather year. In each case, the output load shape outputs are disaggregated by end use for a subset of the end uses in each sector. The load shape disaggregations are summarized in Table 2-2. Not all disaggregated end uses are considered for Shift participation in this study; see section 2.3.1.2 for details.

Table 2-1. Summary of the customer clusters generated by LBNL-Load for this study.

Sector	Building Type	Cluster Count
Residential	**Total**	**493**
Commercial	Office	547
	Retail	509
	Refrigerated warehouse	54
	Other	292
	Total	**1402**
Industrial[a]	Agriculture	286
	Chem/petrol/gas manufacturing	123
	Data center	35
	Food/beverage manufacturing	183
	Water and wastewater	469
	Other manufacturing	393
	Other industrial	125
	Total	**1614**
Grand total		**3509**

[a] Includes agricultural customers.

Table 2-2. List of the end uses for which LBNL-Load explicitly disaggregates load shapes in each sector. Note that not all disaggregated end-uses are considered as targets for Shift participation in this study; see section 2.3.1.2 for details.

Sector	End use
Residential	Plug loads
	Pool pumps
	Space cooling
	Space heating[b]
	Water heating[b]
	EV charging[b]
	Other
Commercial	HVAC
	Lighting
	Refrigeration[c]
	EV charging[b]
	Other
Industrial[a]	Process[d]
	Pumping[e]
	Other

[a] Includes agricultural customers.

[b] Added using a separate forecast model, not disaggregated from customer data.

[c] Estimated for refrigerated warehouses only.

[d] Estimated for all industrial sites except agricultural.

[e] Estimated for agricultural and water/wastewater sites only.

2.2.2. DR-Path

The DR-Path module utilizes the load shape forecasts from LBNL-Load as a basis for estimating the future potential DR resource in California that is available at a given annualized procurement cost. The calculation framework is described conceptually in Appendix C, section C.1 of this report and is fully detailed in Appendix G of the Phase 2 report.

In brief, DR-Path begins by computing, for each of the end-use load shapes for every cluster in the LBNL-Load output, the average quantity of load that would be available, in principle, to provide DR at times of system need.[17] The model then envisions coupling each of these potential DR resources with an array of different technologies that can enable the load to participate in DR. An array of different customer incentive payments is then applied to each end-use/technology combination, and the model estimates the resulting customer program participation levels using a customer decision modeling framework.[18] The result is a tree-like structure encompassing hundreds of thousands of possible future pathways to capturing the potential DR resource in California. Each pathway yields a certain DR resource, based on the associated technology's performance capability, and an associated cost, composed of initial and operating costs for the technology, program administration costs, and customer incentives. At a given procurement cost, some fraction of the available pathways are affordable; the model then chooses the pathways that maximize the resource contribution from each cluster and end use to estimate the available DR potential at that cost. Figure 2-2 shows a schematic diagram of this pathway modeling procedure.
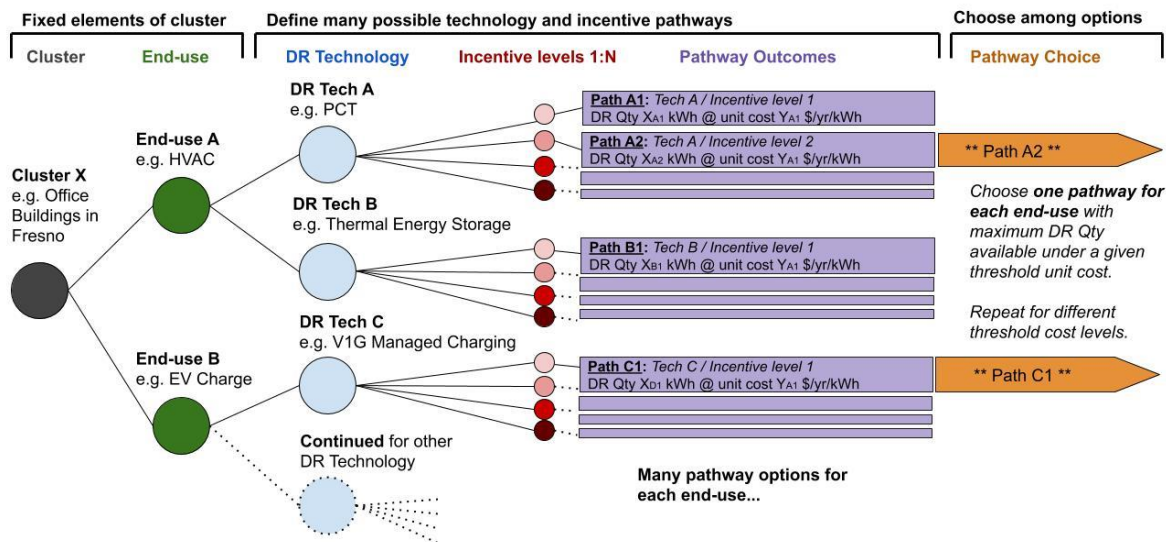
Figure 2-2. A schematic diagram of the DR-Path modeling framework.

2.2.2.1. Updates for Phase 3

This phase of the DR Potential Study included a major software upgrade to the DR-Path module, as well as conceptual updates to the approach to modeling Shift, which yielded version 2.0 of DR-Path. The

---

[17] These times are defined based on a probabilistic model of the timing of DR dispatch for each different category of DR (e.g., Shed or Shift). See Appendix section C.2.5.2 for details.

[18] Importantly, the customer participation model considers only the customer decision to enroll in a DR program, not the decision to participate in any specific DR event. Thus, the model is estimating the *potential* resource for a typical event, not the resource that would actually be realized.

updates are described in detail in Appendix C of this report. Here we summarize the most important conceptual updates incorporated into DR-Path 2.0 for the case of Shift.

2.2.2.1.1. Shift dispatch probability and resource timescale

A critical update to the DR-Path modeling framework concerns assumptions about the relative dispatch probability of Shift events throughout the year. Shift DR is generally assumed to be utilized in such a way as to shift loads within time windows that are centered on the transitions between Take and Shed periods. In Phase 2, all such transitions were assumed equally likely to trigger a dispatch of Shift; however, in reality the need for Shift is likely to be more acute at some Shed/Take transitions than at others, either because the likelihood of VRE curtailment (and associated surplus energy) is especially high in the Take period, or because the generation ramp is especially steep. In Phase 3, we estimated the impact of this varying dispatch probability using three different scenarios: (1) the *ramping* scenario assumes that the probability of Shift dispatch depends primarily on the steepness of the generation ramp that needs to be mitigated; (2) the *curtailment* scenario assumes that the dispatch probability depends on the likelihood of VRE curtailment in the absence of Shift; and (3) the *flat* scenario assumes that all potential Shift dispatches are equally likely (as was done in Phase 2), as a baseline for comparison. As discussed in Appendix C, section C.2.5.2, we use primarily the *ramping* scenario when presenting results in this report.

This probabilistic approach to estimating Shift dispatch also necessitated a change in the way that the Shift resource is reported. In Phase 2, the output of DR-Path for Shift represented an average daily budget of shiftable energy, assuming that Shift would be dispatched twice per day (once at each of the morning and evening ramps). In Phase 3, because the calculation of the Shift resource weights each potential Shift dispatch with a unique probability, it is more natural to report the weighted average Shift resource per *dispatch event*, rather than per *day*. Thus, the Shift resources reported here represent the average quantity of Shift that can be dispatched at a particular time and better captures the shiftable energy available at the most important times of the year in terms of grid needs. This resource could in principle be called upon more than once per day, provided that sufficient time elapses between dispatches;[19] in this sense, it is readily comparable to battery storage. That is, a GWh of Shift has a similar impact, when dispatched, to having a GWh of battery storage available in the appropriate state of charge (either fully charged or fully discharged) to execute the same load shift.[20] When comparing this report to the Phase 2 report, it is important to recognize that the Phase 2 results, which assumed two Shift events per day, will be roughly twice as large and half as expensive as the per-event shiftable energy budget presented in this report, all else being equal (see section 2.2.2.2.2).

---

[19] Shift resources will have a characteristic recovery time, analogous to a recharge time for storage resources, which allows the end use to return to its nominal state of operation under which it is again able to execute a shift. For instance, a pre-cooled building cannot pre-cool again until the temperature has returned to the preferred set point.

[20] By contrast, the reporting paradigm used in Phase 2 would yield a Shift resource that was twice as large as the equivalent battery resource needed to effect the same load shift.

2.2.2.1.2. Shift directionality

The Phase 3 updates to DR-Path also introduced more realistic operational limits on the amount of energy that can be shifted for specific end uses. In Phase 2, all end uses were assumed to be shiftable in either temporal direction (earlier or later in time), and all Shift events were assumed to have a perfectly neutral event on total energy consumption (i.e., the shifted consumption was exactly equal to the baseline consumption.) In Phase 3, certain end uses were constrained to shift load only forward or backward in time: for instance, EV charging load can only shift forward in time, because the baseline behavior assumes that EVs will begin charging as soon as they are connected to the charger, so it is impossible to shift the load to an earlier time because the vehicle is not connected. Certain end uses were also assigned loss factors that applied to each shift: for instance, shifted space cooling loads via pre-cooling was assumed to require more energy than cooling in real time, owing to thermal losses through the building envelope.

2.2.2.1.3. The Shift resource in different time periods

DR-Path 2.0 introduces the capability to calculate the average DR resources that are available in different time periods. In Phase 2, all results reported the DR resource that was available on average over the course of the year, at times of likely dispatch for each DR type (Shed, Shift, or Shimmy). The updated model now includes the ability to compute the average DR resource over time periods other than a year, such as season, quarter, or weekday versus weekend time periods. In this report we use this new capability to examine seasonal variability in Shift.

2.2.2.1.4. Modeling thermal storage technologies

The Phase 3 updates to DR-Path also added two new technologies that utilize thermal storage to enable Shift, in addition to the pre-cooling and pre-heating strategies for shifting space conditioning loads that were already included in the Phase 2 study. The first is thermal energy storage (TES) systems for commercial HVAC, which make ice or chilled water at one time of day to use as a reservoir to provide cooling service without the need to run the building's chiller at another time of day. The second is a strategy to shift water heating load by overheating the water in the storage tank, using a thermostatic mixing valve at the outlet to provide water at a safe temperature, and allowing the heating mechanism to be turned off for some period of time without loss of hot water service (BPA 2018; Delforge and Vokovich 2018). For details of the modeling inputs and assumptions for these technologies, see Appendix D of this report.

2.2.2.1.5. Modeling BTM batteries

From the perspective of the grid, Shift strongly resembles a form of energy storage, which can be "charged" at one time of day (the take) and discharged at a later time to reduce peak loads (the shed). Like batteries, Shift also typically requires a "recharge" time period between dispatches, during which the state of the associated energy service (the "state of charge") needs to be allowed to return to nominal operation before another shift can be dispatched. Thus, it is interesting to consider BTM battery storage as a potential source of Shift to be compared against other Shift-enabling technologies.

DR-Path 2.0 incorporates a conceptually new approach to modeling the potential for Shift using BTM batteries. In the Phase 2 study, DR-Path envisioned utilizing BTM batteries in an opportunistic way to load shift, with the batteries being installed primarily for the purposes for which they are installed today, such as to manage commercial demand charges. In this case, typical batteries have relatively small charge/discharge and storage capacities, relative to the peak demand of the site, and the Shift resource they can provide is limited by their size. In Phase 3, we considered that BTM batteries can be installed for the primary purpose of enabling Shift. In that scenario, arbitrarily large batteries can in principle be installed, with costs that simply scale with the storage and charge/discharge capacities of the battery system.

Appendix section C.2.9 of this report describes our approach to BTM battery modeling in detail. In brief, DR-Path 2.0 assumes that sites in each cluster can be enabled as Shift resources by installing BTM battery systems whose inverter charge/discharge capacity will be sized to handle some fraction of the site's annual peak load—in principle up to 100 percent of this peak value[21]—with associated battery storage that is sufficient to provide 1, 2, or 4 hours of power at full capacity. This makes BTM batteries potentially an extremely large source of Shift. Foreshadowing our findings somewhat, this implies that, once the supply curve price has reached a level where BTM batteries are generally affordable, the battery resource will dwarf all other potential sources of Shift. This makes the procurement price of BTM batteries[22] a natural point of reference for other Shift resources: those that are less expensive than batteries can provide a battery-like grid service at a lower cost than actual BTM storage, while those that are more expensive are likely not cost-effective compared to simply installing a battery at the site.

2.2.2.1.6. Modeling TOU rate impacts

DR-Path modifies the forecasted load shape outputs of LBNL-Load to account for expected consumer responses to time-varying rates. The Phase 2 study applied several different potential scenarios for TOU rates. In Phase 3, the load reshaping for the residential sector was updated to be consistent with the TOU impacts that were forecast as part of the CEC's 2017 IEPR (Kavalec et al. 2018). Since TOU rates will be the default for residential customers following CPUC decision D.15-07-001 (CPUC 2017b), the reshaped load shapes are taken to be the baseline for the Phase 3 analysis, unlike in the Phase 2 analysis, where the available Shape resource was taken to be an effective source of Shift and Shed. (See section 2.2.2.2.2 for of a discussion of how this impacts the supply curves presented in this study.)

---

[21] To be sure, for many sites a battery system of this size is unlikely to be installed. Installing such batteries at all sites would be tantamount to enabling the entirety of California electricity demand as a Shift resource, which is far more Shift than is realistically necessary. This assumption, then, is simply a modeling device to account for the fact that a battery system can, in principle, be sized to handle an arbitrarily large fraction of a building's load, with a cost that simply scales with the quantity of storage and charge/discharge capacity. By making such whole-site batteries available in the model, we can achieve a full accounting of the potential Shift resource from BTM batteries, even if much of it is unlikely to be procured in practice.

[22] Specifically, we compute the cost of procuring BTM batteries explicitly to be used as a Shift DR resource, including all technology costs, operating costs including software controls, and DR program costs for administration, marketing, and outreach.

2.2.2.2. Overview of modeling outputs

The output of DR-Path 2.0 is a large database of future DR pathways representing, for each cluster and end use, the quantity of each type of DR (Shed, Shift, or Shimmy) that can be enabled with a given technology and incentive level, as well as the total cost associated with enabling that resource. Then, by summing together the resources that can be procured for a given price, we can assemble a supply curve for each type of DR, representing the potential quantity of DR that can be procured at a given marginal procurement price (i.e., the aggregate DR resource that can be procured for that price or less). Because this supply curve is built in a bottom-up way from the individual clusters and end uses, it is straightforward to disaggregate the supply curve along various dimensions, to better understand the sources from which the DR resource is being drawn. For instance, one might examine the DR supply curve disaggregated by end use, to understand which end uses are the most inexpensive to enable; or by building type, to understand which types of customers may be most fruitfully targeted for DR program participation.

Figure 2-3. Schematic diagram of a DR supply curve plot for Shift, showing the total DR supply and a disaggregation by source.

2.2.2.2.1. Overview of the Shift supply curve

To help orient the reader to a commonly used supply curve plot in this study, Figure 2-3 shows a schematic diagram of a disaggregated supply curve for Shift. The vertical axis shows the marginal procurement price for Shift (total annualized cost required to install technology, administer the DR program, and pay customer incentives). The horizontal axis shows the quantity of shiftable energy that would be available, on average, during a Shift event and that can be enabled at or below a given annualized marginal cost. The black curve in Figure 2-3 represents the total Shift supply curve—i.e., the

total quantity of Shift that can be enabled at a given marginal cost of procurement. As the price increases, the supply curve steepens, representing diminishing returns to additional investment in more expensive sources of DR (this occurs, for instance, because the sites that are expensive to enable, per unit of DR, also have small resources to offer). The horizontal bars break out the Shift resource by source type (e.g., end use or building type) at a discrete set of prices. It is worth dwelling briefly on the units of measure used to express the Shift quantity, as well as the cost units, since these differ from the units used for traditional Shed DR; Box 3 gives a detailed overview of this topic.

Box 3. Units of measure for Shift and cost modeling components

Because Shift is an energy resource, the units of measure used to express the Shift supply curve differ from the units typically used for Shed DR and other capacity resources. The quantity of Shift is expressed in units of GWh; as discussed in section 2.2.2.1.1, this represents the amount of energy that is available to be shifted at times when Shift dispatches are likely, over the course of the time period being considered (typically a year or a season). The cost is expressed in an annualized form in units of $/yr/kWh; it can be thought of as the annual expenditure that would be needed to procure and maintain an additional kWh of Shift, if the up-front costs were financed at a fixed rate of return over the full lifetime of the equipment needed to enable the resource.

The cost metric includes all costs associated with bringing a DR resource online that are borne by the utility or aggregator. These include:

- Amortized up-front equipment and installation costs, less any costs borne by the customer or site operator (which are estimated as the value of any co-benefits the site receives from installing the enabling technology, such as reduced energy costs).;
- Annual costs for operations and maintenance at the site;
- Annual costs for program administration, marketing and customer outreach; and
- Annual incentives to ensure customer participation in DR programs.

For further discussion of the cost accounting model in DR-Path, see appendix section C.1.3 of this report.

Importantly, the units used here differ from the units that were used to express Shift resources in the Phase 2 report. The cost unit, $/yr/kWh, can also be written in the mathematically equivalent form $/kWh-yr, which is commonly used to report levelized costs for battery storage. The Phase 2 report generalized the denominator of that unit to represent the quantity of Shift procured for a given marginal expenditure in units of GWh-yr. This unit was taken to represent a quantity of shiftable energy that has been secured as a DR resource for a period of a year. This set of units has a significant drawback, however, in that it contains two different units of time (hours and years) placed adjacent to one another, which was a persistent source of confusion for stakeholders. In this report, then, we use a more straightforward unit, GWh, for the Shift quantity, rather than GWh-yr. Crucially, these two units are describing exactly the same quantities; they simply differ in whether they explicitly include the procurement period when expressing the quantity. We also choose to write our cost units as $/yr/kWh, instead of $/kWh-yr, to distinguish the cost of Shift from the levelized cost of battery storage and emphasize the inclusion of additional annual program-related costs ($/yr) that are needed to procure a quantity of shiftable energy (kWh). For more discussion on units of measure for DR resources and costs, see appendix section C.1.2 of this report.

2.2.2.2.2. Comparing the Phase 2 and Phase 3 supply curves for Shift

As discussed in section 2.2.2.1, in this study we made several updates to the modeling approach for Shift in DR-Path. Importantly for the quantities presented in Shift supply curves, when presenting the Shift resource this study presents a quantity of shiftable energy that is expected to be available in each individual Shift event (of which there may be several per day), whereas the Phase 2 study presented its Shift resource in terms of a daily quantity of shiftable energy, under the assumption that Shift would be utilized exactly twice per day. As a result, for exactly the same resource, the paradigm in the Phase 2 study would report a Shift resource, in GWh, that is twice as large, and correspondingly half as expensive in $/yr/kWh, than what would be reported in this study. Thus, the supply curves in the Phase 2 study will appear to show a Shift resource that is substantially larger and cheaper than the one shown in this study, but this is largely due to the way the resource is counted: this study estimates the energy available for a single Shift dispatch, while the Phase 2 study combined the energy available for two Shift dispatches per day.

In addition, we assumed in this study that BTM batteries could be built to arbitrarily large specifications if needed to provide load shifting service. To foreshadow one of our results (see section 3.1.1), a major implication of this assumption is that BTM batteries represent—at least in principle—an extremely large source of Shift. Once the price in the supply curve reaches the price of BTM batteries, the available resource from battery storage quickly dwarfs the resource from flexible loads, since an arbitrarily large amount of battery storage can be purchased at this price. For this reason, we primarily focus on the supply curve for non-battery Shift, using the BTM battery price threshold as a benchmark for comparison. In Phase 2, by contrast, the potential Shift resource was estimated for BTM battery systems that were then typically being installed for non-DR purposes. This represented a much smaller quantity of Shift potential, which was typically included in the supply curve.

Figure 2-4 presents a schematic diagram summarizing the implications of these two changes for the supply curves presented in this study, compared to the Phase 2 study. The solid lines illustrate the supply curves that are most typically presented in the two studies, and the shaded regions represent the qualitative sizes of the estimated BTM battery resource in each study. The Phase 2 non-battery supply curve (dashed grey line) is twice as large and half as expensive as the corresponding Phase 3 curve (solid black line). Importantly, these two curves would represent exactly the same resource; they differ only in the accounting method used to count the size of the resource. The inclusion of BTM battery in the Phase 2 curve then boosts this curve even further. In addition to these changes, the Phase 3 study made changes to the load forecasting methods, the technology inputs, and the assumptions about Shift dispatch probability, all of which would further alter the supply curves. Because of the differences summarized here, however, it is not possible to straightforwardly compare the Shift supply curves in this study to the ones calculated in the Phase 2 study, and direct comparison is not recommended.
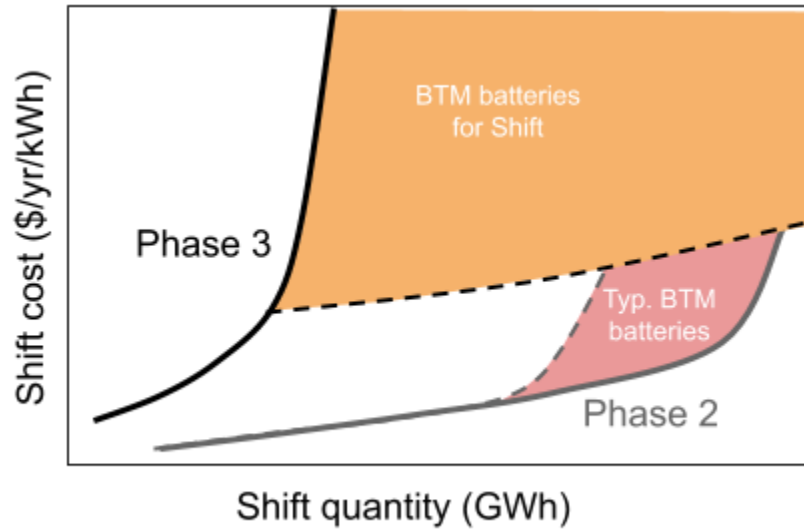
Figure 2-4. A schematic diagram comparing the paradigms for plotting Shift supply curves as presented in this report and in the Phase 2 report. The non-battery supply curves (solid line for Phase 3 and dashed line for Phase 2) represent the same Shift resource, presented using two different approaches for counting the quantity of shiftable energy.

In the Phase 2 study, the supply curve diagrams for Shift and Shed also displayed horizontal bars at zero cost, representing the quantity of effective Shift that could be obtained via Shape (via TOU pricing), which displaced the entire supply curve to the right. Since then, TOU rates have been established as the default for most customers via CPUC decision D.15-07-001 (CPUC 2017b); thus, in this study we consider the resulting Shape impact to be a resource that has already been captured and thus is "baked in" to the underlying load shape. Our Shift supply curves then represent the quantity of load shifting resource that would be available in addition to the Shape that has already been captured. Therefore, we do not show a bar representing new Shape in the supply curves, although we acknowledge that some fraction of the identified Shift resource could potentially be obtained more cheaply as Shape via more stringent TOU rate structures. We discuss the effective Shift impacts of existing default TOU rates in section 3.1.3.

2.3. Modeling Inputs, Assumptions, and Scenarios

The modeling performed for this study in the DR-Futures framework rests on certain critical input data, as well as on assumptions about future trends in the California energy sector that may involve a high degree of uncertainty. In this section we summarize the data used to support the modeling and the most important assumptions, and we describe several alternative model scenarios that we will use in this study to assess a range of different possible futures.

2.3.1. Summary of modeling inputs and assumptions

The modeling inputs used for the DR-Futures framework in this study are described in detail in appendices C and G of the Phase 2 report, with updates summarized in Appendix D of this report. Here, we summarize the primary input data sources, and we highlight the most essential assumptions.

2.3.1.1. LBNL-Load

The primary data inputs for the LBNL-Load module are as follows.

- Demographic information for all customers in the three California IOU service territories, as of 2014, including building type, ZIP code, rate class, peak power demand, and total annual energy consumption, provided by the California IOUs in support of the previous phases of the DR Potential study.
- Hourly load data for approximately 200,000 California IOU customers, covering all 8,760 hours in calendar year 2014; also provided by the California IOUs in previous phases.
- Historical weather data from California National Oceanic and Atmospheric Administration (NOAA) weather stations, for calendar year 2014, and selected from various years to represent a typical (1-in-2) and an extreme (1-in-10) weather year.
- Forecast information for California energy demand growth from the CEC's revised 2018–2030 demand forecast (Kavalec et al. 2018), released as a supplement to the 2017 IEPR (CEC 2017b).
- Forecasted load shapes for EV charging developed for the CEC (Bedir et al. 2018) using the EVI-Pro model (NREL 2019), and CEC forecasts of future demand for EV charging (Bahrenian et al. 2017).
- Typical load shapes for residential space heating (Tarroja et al. 2018) and water heating (Carew et al. 2018).

Throughout this study, where it is necessary to make a choice among different forecasting scenarios, we have selected the CEC's Mid scenarios—for total demand, additional achievable energy efficiency (AAEE), and EV charging demand. The LBNL-Load load shape forecasting rests on an assumption that overall load growth in each sector scales with certain sector-specific indicators tabulated in the IEPR demand forecast.

2.3.1.2. DR-Path

The DR-Path model relies on the following sources of input data.

- Load shape forecasts, by customer cluster and end use, generated by LBNL-Load.
- A database of information on DR-enabling technologies,[23] synthesized by LBNL from a wide range of sources, and including:

  - Cost information, including both fixed and variable initial and operating costs;
  - Site-level co-benefits, expressed as a fraction of the technology cost;
  - Performance information, indicating the fraction of load that can be shed or shifted over a given time period, round-trip loss factors for Shift, and any temporal directionality constraints on load shifting; and

---

[23] Throughout this report, the phrase *enabling technologies* refers to equipment or suites of equipment that could be installed at a customer site to enable particular loads to serve as DR resources providing particular grid services. This phrase encompasses equipment required both for load control and for communications and telemetry to support market settlement. For the purposes of this report, the usage and meaning of *enabling technology* is similar to the term *measure,* which is often used in studies of energy efficiency potential.

○ Expected rates of accelerated price decline or performance improvement for certain emerging technologies.

● VRE electricity generation forecast based on historical solar and wind generation in CAISO, scaled to future years using the CPUC's renewable portfolio standard (RPS) calculator, version 6.2 (CPUC 2016).

● Projected TOU rate impacts, provided with the CEC's IEPR demand forecast (Kavalec et al. 2018).

Embedded in the enabling technology data are implicit assumptions about which end uses are capable of load shifting. Table 2-3 provides a comprehensive list of the end uses that can participate in Shift, in principle, and indicates the specific end uses and enabling technologies that are included as sources of Shift in this study. With the exception of new electrification loads, the enabling technologies considered in this study are all proven technologies that have been successfully used for DR in real-world applications. Thus, the DR potential estimated here represents the resource that is available using current technology.

Throughout the Phase 3 study, we have assumed administration and marketing costs for DR programs that are the same as those assumed in the Phase 2 report. We assumed a discount rate of 7 percent for the purposes of computing a capital recovery factor when annualizing up-front expenditures. All technology costs are reported as net costs made up of the total initial and operating costs, less any site-level co-benefits (i.e., costs borne by the utility or program administrator). In Phase 2, we also considered cost accounting frameworks that included potential market revenues, as well as benefits to the electricity grid; we neglect these two frameworks here, owing to high levels of uncertainty for such revenue streams in the case of Shift, for which no historical data are available.

Table 2-3. Summary of end uses that are capable of participating in Shift, in principle. Shaded cells indicate end uses that are not modeled as sources of Shift in this study; all unshaded cells are included. Red text indicates end uses that have been newly added in this study, while blue text shows end uses whose modeling approach has been significantly updated in this work.

End Use	Sector	Approach/Technology
Space cooling	Res	Pre-cooling with a programmable communicating thermostat (PCT)
Space heating	Res	Pre-heating with a PCT
HVAC	Com	Pre-cooling with a PCT or energy management system (EMS); Thermal storage
Ventilation	Res/Com	Advanced controls
Water heating	Res	Pre-heating and scheduling
	Com	Pre-heating and scheduling
Pool pumps	Res	Dynamic scheduling
Irrigation pumping	Ind (Ag)	Dynamic scheduling
Wastewater pumping	Ind	Dynamic scheduling
Water supply pumping	Ind	Dynamic scheduling/pumped storage
Industrial process	Ind	Dynamic scheduling
Refrigeration	Com	Warehouse pre-cooling
	Res/Com	On-board thermal storage, smart cycling
EV charging	Res/Com	Dynamic scheduling
	Res/Com	Two-way charging
Battery (whole building)	All	Storage, two-way charging
Battery (distributed/point of use)	All	Storage, two-way charging
Plug Loads/Appliances	Res/Com	Dynamic scheduling (e.g., dishwashing)
Lighting	Ind (Ag)	Indoor photoperiod shifting (grow lights)

2.3.2. Description of model scenarios

Several aspects of the DR-Path modeling in this work involve uncertainty about future trends or qualitative outcomes that is most easily handled by developing multiple different scenarios, each of which makes different assumptions about the future.[24] The areas in which we have developed different future scenarios are as follows:

---

[24] One of the outcomes from the Phase 2 study was that the differences in potential related to qualitative uncertainty were much larger than the uncertainty we identified related to trajectories of technology performance and cost within the scenarios. The

- Electrification. We included a Base electrification scenario, which considers only growth in private EV ownership, as well as an Additional electrification scenario, which assumed significant growth in electric space and water heating in the residential sector.

- Weather. We forecasted end use load shapes in both a 1-in-2 (typical) and a 1-in-10 (extreme) weather year. (In the Additional electrification scenario, only a 1-in-2 weather year was modeled, because the input load shape data for the additional electrified loads was based on a typical weather year.)

- Technology improvement. We considered five different scenarios for long term price declines and performance improvement trends for DR-enabling technologies. These trends were assumed to be the baseline rate of improvement for "mature" technologies; faster trends were applied to certain emerging technologies.

    - The Base technology scenario assumed prices and performance frozen at present-day values.
    - The BAU technology scenario assumed price and performance trends consistent with conservatively estimated historical improvement rates.
    - The Mid technology scenario assumed modestly optimistic rates of improvement in price and performance for DR-enabling technologies.
    - The High technology scenario assumed highly optimistic rates of improvement in technology price and performance.
    - The Market Transformation scenario assumed dramatic declines in the costs of DR-enabling technology, as might occur in the presence of a new, disruptive technology.

- Dispatch probability. We considered three different scenarios, representing different approaches by the system operator to dispatching a putative Shift DR resource. These scenarios—called *flat, curtailment,* and *ramping*—are described in detail in appendix section C.2.5.2.

For this study we selected certain permutations of the various scenarios above to construct several different runs of the DR-Path model, representing qualitatively different potential futures for Shift DR in California.

Table 2-4 presents these different model runs. For certain runs, we did not consider all the different forecast years: in the case of the runs with Additional Electrification, year 2020 was not appreciably different from the Base Electrification scenario; and the Curtailment and Flat dispatch probability scenarios were run only for comparison to the Ramping scenario, for which a single forecast year was sufficient.

---

within-scenario uncertainty was identified in Phase 2 with a Monte Carlo analysis. For Phase 3 we are removing this within-scenario uncertainty analysis since it did not result in significantly "better" understanding of the overall uncertainty and comes at a cost of computational complexity and additional challenge to interpret and understand the study results.

Table 2-4. Summary of the various different model runs considered in this study.

Run Name	Dispatch Probability	Technology Scenarios	Weather Scenarios	Electrification	Forecast Years
Reference	Ramping	Base, BAU, Mid, High	1-in-2, 1-in-10	Base	2020, 2025, 2030
Curtailment Mitigation	Curtailment	Mid	1-in-2	Base	2030
Flat Dispatch Probability	Flat	Mid	1-in-2	Base	2030
Additional Electrification	Ramping	Base, BAU, Mid, High	1-in-2	Additional	2025, 2030
Market Transformation	Ramping	Market Transformation	1-in-2	Additional	2025, 2030

Note: All scenarios started from the Mid-demand/Mid-AAEE load forecasts from LBNL-Load and computed technology costs net of site-level co-benefits.

3. Findings

3.1. The size of the Shift resource in California

This section presents the primary results of the Phase 3 study in the Reference model run.
We first present the full supply curve for Shift DR, including all enabling technologies, which shows that BTM batteries come to dominate the supply curve above a certain price. We identify this price threshold as a natural point of comparison for non-battery load-shifting technologies. We then undertake a detailed examination of the Shift resources that can be enabled by non-battery technologies at a cost that is equal to or less than the cost of BTM battery storage. The central finding is that *sufficient load flexibility exists in California today to utilize much of the surplus renewable energy that would otherwise be curtailed, and substantially reduce flexible generation needs, for a lower cost than installing BTM battery storage*.
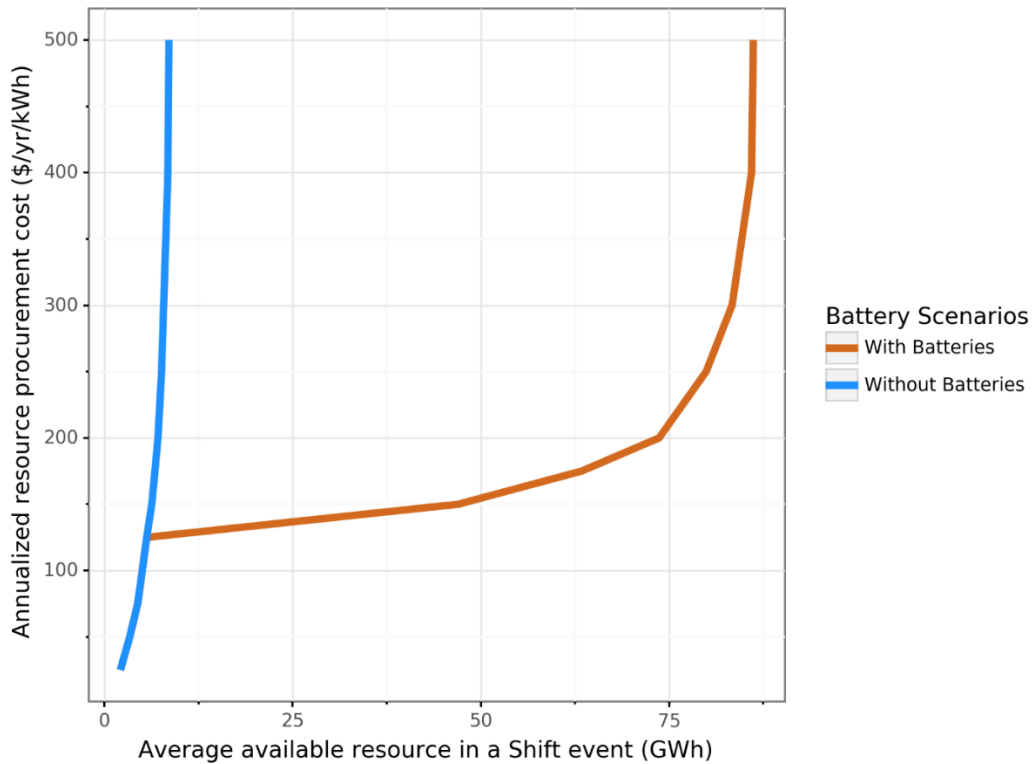
Figure 3-1. The Shift supply curve, in 2030, for the Reference model run, with 1-in-2 weather and Medium technology-improvement assumptions. Different curves indicate the resource that can be enabled with technologies other than BTM batteries (blue) and the additional resource that can be enabled with BTM batteries (orange). (At lower price levels, the orange curve is obscured by the blue.)

3.1.1. The Shift supply curve including BTM battery technology

Figure 3-1 shows the Shift supply curve for the Reference Scenario (for a 1-in-2 weather year with Medium technology improvement trends) in forecast year 2030. Two versions of the supply curve are

shown: the orange curve shows the total available Shift resource estimated by the DR-Path model from any enabling DR technology considered in the model, including BTM batteries. The blue curve shows the resource that is available from enabling technologies other than BTM batteries.

For procurement prices below roughly $150/yr/kWh, the two curves are identical, indicating that BTM batteries are not available at these prices. Above this threshold, however, the curves differ dramatically, with the battery-inclusive resource rapidly growing to exceed the non-battery resource by an order of magnitude, providing tens of gigawatt-hours of load-shifting capability on an as-needed basis. This result is to be expected, given the approach to modeling BTM batteries in DR-Path 2.0 (see section 2.2.2.1.5 and appendix section C.2.9), which assumes that much larger battery systems can be installed than are currently typical, in principle with energy and power capacity to handle as much as the annual peak demand of a building for several hours. These systems would be available to shift load using their full power and storage capacity whenever needed.[25] Since the load-shifting capacity from such arbitrarily large batteries will (by definition) always exceed the individual end-use loads that are available for shifting at a given time—typically by a large margin—once the supply-curve price has reached the price at which battery-enabled Shift can be procured, BTM batteries become the dominant source of Shift in the supply curve. To be sure, battery systems that are large enough to handle entire large-building loads may not typically be built in practice:[26] the Shift supply curve computed under our assumptions indicates the *maximum* Shift resource that could be procured for a given marginal cost, and it would only be necessary to construct the portion of this resource that is sufficient to meet the relevant system need.

Figure 3-1 reflects the basic reality that, if the system need for load shifting is acute enough to make BTM storage a cost-effective option, then batteries will represent both the largest and the most straightforwardly attainable source of Shift. Put differently, if purchasing BTM batteries purely to enable load shifting is a cost-effective option to enable load shifting from the utility's or system operator's perspective, then purchasing batteries will likely make more sense than attempting to enable more load flexibility through controls or other end-use-level technology. The marginal cost of enabling new DR needs to beat the costs of BTM batteries (in dollars per kilowatt-hour-year) to be competitive.[27]

The BTM battery procurement price threshold thus represents a natural point of comparison for non-battery Shift-enabling resources. As shown in Figure 3-1, there are several GWh of non-battery Shift potential available at or below the $150/yr/kWh price threshold at which batteries become the dominant source of Shift. As discussed previously, from the system operator's perspective, Shift resembles BTM storage; thus, the Shift resource that is available below the battery price threshold represents a virtual storage resource that is less expensive than batteries. In the remainder of this Findings section, we will

---

[25] Provided that they are not dispatched so frequently that they do not have time to recharge, which should not be the case if Shift is typically dispatched to mitigate the duck curve.

[26] Although typical residential battery systems are already sized to this standard for smaller households.

[27] There are potential external costs of both DR and storage that may currently not be reflected in the prices of equipment. Both require electronics and electrical systems, and batteries require mining of active materials like lithium and cobalt for the cells. A full comparison of the social costs of these two resources would also account for such externalized costs.

use the battery price threshold[28] as the standard reference price at which we will examine the available non-battery resource in detail.

Before moving on, it is important to emphasize that, in the context of Shift resource cost accounting, a "battery threshold cost" of $150/yr/kWh is *not* the same as a cost of $150/kWh for BTM battery storage. Instead, this threshold represents the annualized cost of the full battery system, plus operating costs, DR program costs, and customer incentive costs. One cannot naively compare present-day or forecasted costs for chemical storage, quoted in $/kWh, against the battery cost threshold used here. Costs for BTM batteries were selected using the same input sources as used for the CPUC's integrated resource planning (IRP) process, for consistency between modeling efforts. For more details on the methodology for computing the cost of BTM batteries as a Shift resource, see appendix section C.2.9.

3.1.2. The non-battery Shift DR resource

Figure 3-2 presents the supply curve of non-battery-enabled Shift DR for the Reference model run, by year, technology scenario, and weather scenario. In each case, the supply of Shift grows rapidly with price below roughly $200/yr/kWh, up to a quantity of several GWh of shiftable energy available per dispatch. At higher prices, there are rapidly diminishing returns to additional investment in Shift-enabling technology, with little additional resource available at marginal prices above $400/yr/kWh.

There is considerable impact from varying assumptions about rates of technological improvement: by 2030, the aggressive assumptions in the High technology scenario can deliver roughly double the Shift resource as the frozen technology assumptions of the Base scenario. By contrast, the different weather scenarios have only a small impact on the estimated Shift resource. The limited weather variability for Shift occurs because Shift has been modeled as a resource that is dispatched on a nearly daily basis, so that the weighted average resource calculated by DR-Path will be dominated by the loads on days with typical weather, not the extreme weather days that will be most sensitive to the different weather assumptions. By contrast, traditional Shed DR would be expected to show more variation with weather assumptions, since it is more likely to be dispatched on days with extreme temperatures.

---

[28] Note that there is not a single "price" for batteries shown in Figure 3-1; rather, the battery resource increases rapidly from zero around $150/yr/kWh, then increases more slowly at higher prices until saturating at around $400/yr/kWh. This increase occurs because batteries have slightly different prices in the residential, commercial, and industrial sector; because fixed costs for DR program administration make smaller sites more costly to enable; and because the modeling also accounts for customer participation rates that increase with higher incentive payments. For the purposes of determining a reference price, though, it is sufficient to make note of the price at which the battery-enabled Shift resource comes to dominate all other forms of Shift, which is $150/yr/kWh in this case.

Figure 3-2. Supply curves of Shift DR enabled by non-battery technologies, by year, technology scenario, and weather scenario, for the Reference model run. Horizontal dotted lines represent the price threshold above which BTM batteries are available to enable load shifting in each year.

In the remainder of this report, we use the Medium technology scenario, with 1-in-2 weather assumptions as our default scenario for reporting detailed results. In each panel of Figure 3-2, the horizontal gray dotted line shows the price threshold for BTM batteries, discussed in the previous section, in that default scenario. Because BTM battery costs are assumed to decline with time, the battery threshold is also progressively lower in each year. Table 3-1 presents the BTM battery price threshold in each year and the quantity of Shift DR available at that price in each weather and technology scenario.[29] Because the battery

---

[29] The battery price threshold used as a reference price in this table uses the Medium technology improvement assumptions in a 1-in-2 weather year. In principle it would be more consistent to use each different scenario's battery price threshold to determine the available resource in that scenario, but we have chosen to use a single threshold in presenting these results for the sake of simplicity.

threshold declines with time, the DR quantity available at this reference price grows somewhat more slowly than would be inferred by inspecting the evolution of the supply curves. In our default scenario (Medium technology, 1-in-2 weather), the available Shift DR resource grows by roughly 20 percent, from 5.3 GWh in 2020 to 6.3 GWh in 2030.

Table 3-1. Available non-battery Shift resource (in GWh) available at the BTM battery price referent, by technology scenario, weather scenario, and year, for the Reference model run.

Year	BTM Battery Price Referent	Base		BAU		Med		High	
		1in2	1in10	1in2	1in10	1in2	1in10	1in2	1in10
2020	$200/yr/kWh	4.6	4.6	5.0	5.1	5.3	5.4	5.8	5.9
2025	$175/yr/kWh	4.4	4.4	5.3	5.4	5.8	5.9	6.8	7.0
2030	$150/yr/kWh	4.0	4.0	5.6	5.6	6.3	6.3	7.9	8.0

It is helpful to consider the 2020 results in the context of present-day CAISO operations. As discussed in section 2.2.2.1.1, Shift resource reported in this study represents an average quantity of energy that can be shifted *per Shift dispatch event* (rather than per day in the Phase 2 report)—and that it may be possible to call on this resource more than once per day, provided there is sufficient time between dispatches, similar to the utilization of battery storage for load shifting. Typical Shift dispatches would tend to occur during the solar-driven generation ramps in the morning and evening,[30] so it may be feasible to dispatch the available Shift resource twice per day,[31] yielding roughly 10 GWh of total shifted energy consumption over the course of the day. This would represent 1 to 2 percent of typical daily CAISO consumption, which ranged roughly between 500–800 GWh in 2018–2019.[32] If each shift takes place within a four-hour window (with two hours of shedding and two hours of taking load), the roughly 2.5 GW demand reduction would represent a ten to twenty percent reduction in net demand during peak hours on typical days, making DR from Shift potentially quite a significant resource.

It is also interesting to compare this estimated Shift resource to the curtailment of VRE and generation ramping that has lately been necessary on the CAISO grid. Figure 1-2 shows that average daily curtailment in the spring of 2019 amounted to more than 5 GWh per day of surplus renewable generation, which is similar to the size of the 2020 Shift resource, suggesting that a single Shift dispatch may be able, in principle, to utilize a substantial portion of the current typical curtailment observed in CAISO. Further supporting this broad conclusion, Figure 1-3 shows that a 10 GWh *daily* budget of shiftable energy—which could be achieved if the 2020 Shift resource were dispatched twice per day—would have been

---

[30] That is, moving load out of the morning and evening peaks into the late morning and afternoon, respectively, to mitigate load peaks and ramping rates.

[31] To be sure, the ability to dispatch twice per day will depend strongly on customers' willingness to respond more than once per day, which would require strategies and enabling technologies that minimize impacts on perceived energy service levels.

[32] See the U.S. Energy Information Administration's Hourly Electric Grid Monitor, at eia.gov/beta/electricity/gridmonitor.

sufficient to eliminate nearly all of the curtailment observed in CAISO in 2019, if it could be deployed optimally. Additionally, the average spring day in 2019 saw a rapid evening ramp in net load, amounting to a need for approximately 10 GW of flexible generation. If Shift occurs in a four-hour window, yielding a ~2.5 GW peak load reduction and an offsetting ~2.5 GW load increase prior to the ramp, then the average ramping need would be cut roughly in half, shrinking by a total of ~5 GW. (An eight-hour shift would cut this reduction in the ramp by half, which would still be a significant impact.) Thus, we find that today's California energy system has sufficient load flexibility available at a cost below BTM battery storage, in principle, to contribute significantly to daily balancing of demand with generation. Developing and operating these Shift resources would have the effect of reducing VRE curtailment and flexible generation requirements.

There is an important caveat to this conclusion, however, in that roughly half of the present-day VRE curtailment in CAISO is driven by local transmission constraints, rather than system-level overgeneration. This means that the surplus energy that drives the value of Shift is not uniformly available across the system. For Shift to offset local curtailment, it must be situated within the constrained regions, so it is important to target Shift resources co-located with renewable energy or in areas with local surplus, if possible. Moreover, as shown in Figure 1-3, the required Shift resource needed to—even nominally— utilize the observed curtailment is growing rapidly, year-on-year. Both of these observations point to a need to bring down the cost of enabling Shift in order to grow the size of the resource. We discuss approaches to increasing the affordable Shift supply in section 3.5.

3.1.3. The effective Shift from time-of-use rates

TOU electricity rates present customers with rates that vary among different time windows throughout the day, to encourage improved alignment between the timing of electrical demand and the needs of the electrical grid. TOU rates have been common for many years for commercial customers in California, and CPUC decision D.15-07-001 made provision for a transition to default TOU rates for residential customers starting in 2019 (CPUC 2017b). In both cases, the time windows for peak pricing are expected to shift toward later hours as increased penetration of solar generation shifts the peak system net load away from the afternoon, where it has historically occurred, and into the evening hours. As discussed in appendix section C.2.8 of this report, the DR-Path model reshapes the forecasted load shapes from LBNL-Load to account for the expected inputs of new and changing TOU rate structures. In particular, this study uses the TOU load impacts forecasted as part of the IEPR demand forecast (Kavalec et al. 2018) to reshape the residential load shapes, in an update from the Phase 2 study, which used earlier estimates of customer price responsiveness and a wider array of potential TOU rate scenarios.

Customer response to TOU rates is categorized as a Shape resource in the nomenclature of DR types that we use in this study. However, TOU rates are intended, at least in part, to encourage customers to shift their electricity consumption away from Shed periods and into Take periods, at least insofar as these periods occur predictably at the same times on each day. The energy consumption thus shifted by customer response to TOU rates can be considered as an effective Shift resource that has been captured within the Shape resource. We can estimate the size of this Shape-as-Shift resource in DR-Path by comparing the reshaped load shapes to the input load shapes from LBNL-Load and computing the

average amount of consumption that has moved from Shed to Take periods, weighted by the relative need for Shift at each Shed/Take transition.[33]

Given the input assumptions for this study, we find that the reshaped loads in DR-Path represent a Shape-as-Shift resource ranging from 100 to 150 MWh in 2030. This resource is quite small compared to the four to eight GWh of dispatchable Shift potential shown in Table 3-1. It is possible that some fraction of this larger resource could be captured as Shape, for a lower cost than enabling a fully dispatchable resource, if new technologies or program designs could be used to boost customer responsiveness to TOU rates from currently assumed levels. (Indeed, the Phase 2 report, which considered a wider range of prospective TOU rates and customer response levels, identified a significantly larger Shape-as-Shift resource.) We discuss this possibility further in section 4.2.3.

3.1.4. The Shift resource by sector and end use

If the Shift resource is to be enabled, it will be important to understand exactly which loads are able to provide flexibility to the grid. Figure 3-3 presents the default-scenario supply curve, disaggregated by sector and end use, with BTM battery price referents shown in each year for comparison.

At the lowest procurement prices, the bulk of the available resource comes from the industrial sector by shifting loads for processing and pumping[34], with a small contribution from commercial HVAC loads. As the procurement price increases moderately, the resource available from each of these end uses grows, with the most rapid growth in the commercial HVAC resource, enabled in part by the availability of thermal energy storage (TES) technology at higher costs, as discussed in the next section. A small resource also becomes available from commercial refrigeration loads (which are restricted in our modeling to refrigerated warehouse loads). Only at prices near or above the battery threshold do other end uses become available as Shift resources.

In particular, in our Reference model run, all residential Shift resources have procurement prices that exceed the BTM battery threshold. On its face, this suggests that, under our reference modeling assumptions, residential loads may not be a cost-effective source of Shift DR (at least in the context of an incentive-based DR program, which is the program structure implicitly assumed in DR-Path). The Reference model run, however, contains a number of assumptions regarding residential technology costs and participation rates, which, while based on the best available data, may turn out to be conservative in the context of new program designs that may emerge to support Shift. See section 3.5 for discussion of a model run that uses more relaxed input assumptions that yield a more substantial residential Shift resource.

---

[33] This can be done using the same weighting function that is used to compute Shift features in DR-Path. See appendix section C.2.5.2.

[34] Pumping loads shown are primarily agricultural irrigation, with a small contribution (approximately 6% of the resource) from wastewater pumping. Importantly, water supply pumping was not included in the estimation of Shiftable end uses in this study (see Table 2-3) owing to limited information on enabling technologies. This end use could contribute additional Shift resource from pumping in practice.

Figure 3-3. Shift supply curves in the default scenario (Medium technology, 1-in-2 weather), disaggregated by sector (light shaded areas) and by end use (shaded bars). BTM battery price thresholds in each year are shown for reference (gray dotted lines). Residential cooling loads are referenced as "space cooling," while commercial loads are labeled as "HVAC" to distinguish between the two resources. Pumping loads are primarily for agricultural irrigation with a small contribution from wastewater pumping.

3.1.5. Enabling Technology Focus: Thermal Energy Storage

To better understand how the Shift supply curve is built up from the diverse shiftable loads available at individual customer sites—and how these can be enabled for Shift participation by different enabling technologies—it is helpful to focus briefly on a single end use and the impacts of a particular enabling technology. Here we explore the impacts of TES technology as a means for shifting commercial cooling loads. This enabling technology was newly included in the DR-Path model for Phase 3. (For details of the modeling inputs and assumptions, see appendix section D.2.2). TES systems enable the building's cooling system, or an auxiliary system, to create large quantities of ice or chilled water in storage tanks, which can be utilized at a later time to provide cooling service to the building without utilizing the electrical cooling system. As we shall see, this technology has the potential to unlock a significant Shift resource.

The typically envisioned approach to shifting HVAC load is to use programmable thermostats (labeled Tstat in Figure 4-4) in smaller buildings, or an energy management system (EMS) in larger buildings, to pre-cool the building below its normal thermal set point during a Take period, allowing cooling loads to be reduced by some fraction at times when load shedding is needed (Ali, Lenzen, and Huang 2018; Herter and Okuneva 2013; Keeney and Braun 1997), since the building can warm by several degrees while maintaining occupant comfort. With TES, the approach would instead be to maintain the building set point during the Take period, but to charge the storage tank simultaneously, and then discharge it during the Shed period. As discussed in more detail in section 3.5, for typical Shed/Take periods expected in California, this can potentially allow buildings to shed all of their cooling load for several hours during times of system need.

Figure 3-4 compares the load shifting capabilities and costs of two different categories of Shift-enabling technology for the commercial HVAC end use. Each point in the figure represents the cost or performance of a particular technology as installed at a typical site in one of the commercial customer clusters developed in LBNL-Load and input into DR-Path. Points are color-coded by building size, and overlaid box plots show the shape of the overall distribution of costs and shiftable loads. As shown, TES can enable a significantly larger amount of load to be shifted as compared to EMSs and thermostats;[35] however, TES is also significantly more expensive per unit of shiftable energy enabled, only becoming available for large buildings at approximately $100/yr/kWh, and at substantially higher prices for medium and small buildings.

---

[35] Note that the shiftable loads are shown on a logarithmic scale; though the increase appears small on the plot, it is in fact nearly a factor of two.
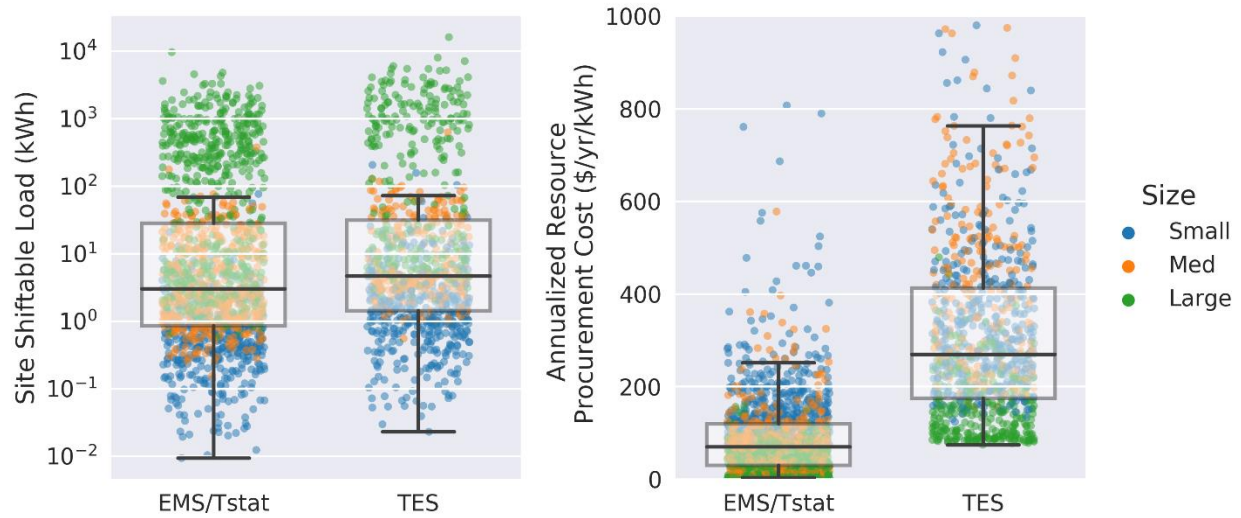
Figure 3-4. Shiftable loads and technology costs (circa 2015) for two different Shift-enabling technologies for commercial HVAC. Left: the sizes of the loads that can be enabled as Shift resources, for representative sites in each of the commercial customer clusters included in the modeling. Right: the associated cost to install each technology at each site. Note that the plot at left is shown on a logarithmic scale covering many powers of ten: although the increase in shiftable load for TES appears small, it is in fact nearly a factor of two.

Figure 3-5 shows how this dynamic translates to the supply curve for Shift. The figure shows the Shift supply curve in 2030, for the commercial sector only, in the Reference model run. In this case, the supply curve is disaggregated by end use (shaded regions) and by the enabling technology that provides the Shift resource (colored bars). To construct this supply curve, DR Path has considered each procurement price level on the vertical axis in turn and selected the DR-enabling technology that can produce the largest Shift resource in each cluster, for a price at or below the price under consideration. Thus, at the lowest procurement prices, the Shift resource comes entirely from EMSs and thermostats, because TES systems are more expensive than these prices can support. Once the procurement price is large enough, TES systems begin to become affordable,[36] and since TES can enable a larger Shift resource at each site, DR-Path selects this technology.

---

[36] TES systems become affordable in Figure 3-5 at a slightly lower price than the minimum TES price shown in Figure 3-4 because the latter figure shows circa-2015 prices, which are assumed to decline slightly by 2030.
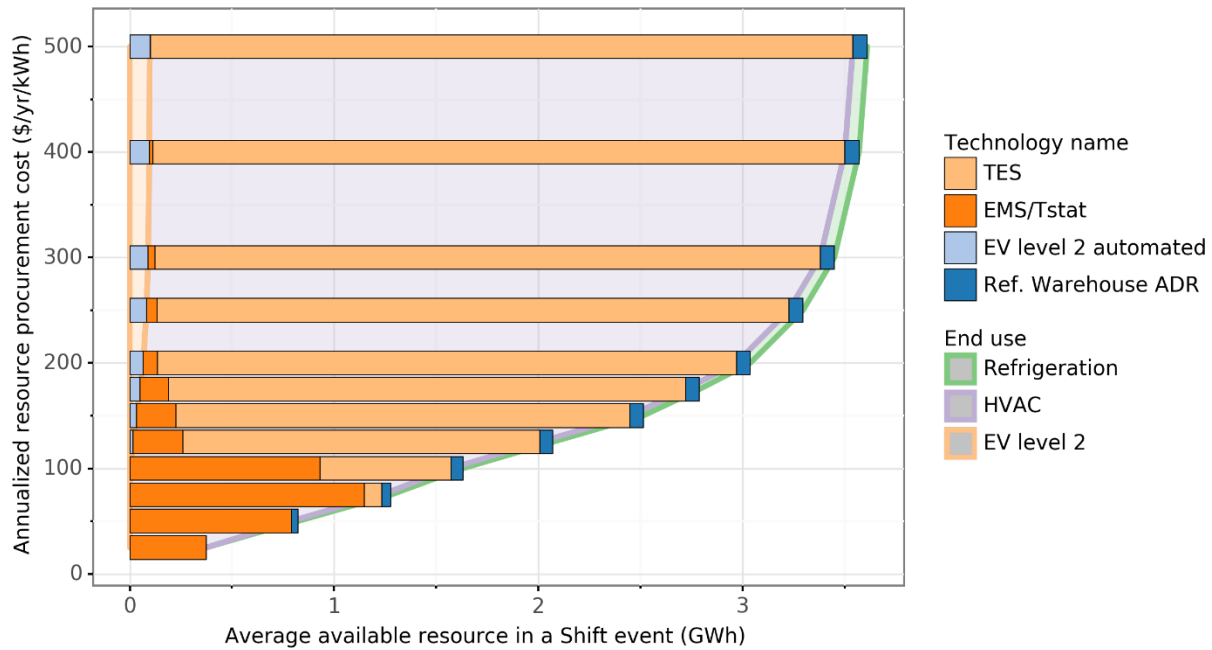
Figure 3-5. The 2030 Shift supply curve in the commercial sector only, disaggregated by end use (shaded regions) and by the enabling technology that is used to provide Shift (colored bars). A small amount of HVAC load can be enabled as a Shift resource by EMSs and programmable thermostats at low costs. At higher costs, TES becomes affordable, which unlocks a substantially larger resource.

The TES-enabled shiftable resource then grows rapidly with increasing price, eventually becoming the dominant technology and enabling more than 3 GWh of Shift, which is a substantial fraction of the resource shown in Figure 3-3. Importantly, note that, at the 2030 BTM battery price threshold of $150/yr/kWh, TES has enabled more than 2 GWh of Shift, indicating that thermal storage technology has a substantial capacity to support load shifting at a lower cost than BTM battery storage.

3.2. The location of the Shift resource in California

In order to enable the Shift resource presented in the previous section, it would be helpful to know which specific customer groups to target for program participation. In this section we explore the physical location of the potential California Shift resource, both in terms of geographical location and in terms of the specific types of buildings that have flexible loads that can be tapped for Shift DR.

3.2.1. The geographical distribution of Shift

Figure 3-6 focuses on the Shift resource that is available at the BTM battery price referent in each year (i.e., the bars underlying the gray dotted lines in Figure 3-3), and we consider how this resource breaks down by end use across the three IOU service territories. Interesting variations are apparent in the resource distributions by end use for the different service territories. In particular, although the overall

Shift resource is similar in size for the PG&E and SCE service territories, PG&E has a much larger availability from pumping loads, reflecting that utility's larger agricultural customer base, whereas SCE has a larger EV charging resource at this price level, reflecting its more urbanized service territory.

Figure 3-6. The Shift resource available at the BTM battery price referent (gray dotted lines in Figure 3-3), disaggregated by utility and end use.[37]

Figure 3-7 presents a more granular geographical breakdown of the Shift resource in forecast year 2030,[38] showing the average GWh of Shift available at the BTM battery price threshold, by end use, in each of the sub-load-aggregation-point (subLAP) regions of the CAISO grid.[39] There is substantial variation in the overall size of the Shift resource within each subLAP, as well as in the distribution by end use. Significant geographical localization is apparent for Shift arising from certain key end uses. For instance, the bulk of the pumping resource is seen to reside in the PGF1 and PGLP subLAPs, which cover the major agricultural regions near Fresno and Bakersfield, respectively, while the largest EV resource is to be found in the SCEW subLAP, which includes Orange County with its sizeable population of automobile commuters.

---

[37] Because of the large number of end uses shown, and the small quantities shown for some end uses, some colors may be difficult to distinguish. To aid in readability, in this chart and all similar fixed-price bar charts, the bars are shaded in the same order (from left to right) as the colors shown in the legend (from top to bottom).

[38] We limit this figure to forecast year 2030 for simplicity, since, as shown in Figure 3-6, the end-use breakdown of the resources tends to maintain similar proportions from year to year.

[39] (See appendix section D.1.1 for maps showing the geographical footprints of the subLAPs. Note that the subLAPs used in this study are as defined in 2014, the year for which the initial customer load shape data were collected.)

Figure 3-7. The Shift resource available in forecast year 2030 at the BTM battery price referent ($150/yr/kWh), disaggregated by subLAP region and end use.

3.2.2. The building types where Shift resides

In Figure 3-8, we present the supply curve for Shift in model year 2030,[40] for our default scenario assumptions in the Reference model run, disaggregated according to the building type hosting the shiftable load, rather than the end use, as shown in previous figures. At all prices, the Shift resource in the industrial sector is fairly evenly distributed across a wide range of building types, whereas in the commercial sector, roughly two thirds of the Shift resource arises in office buildings (the model does not distinguish among residential building types).
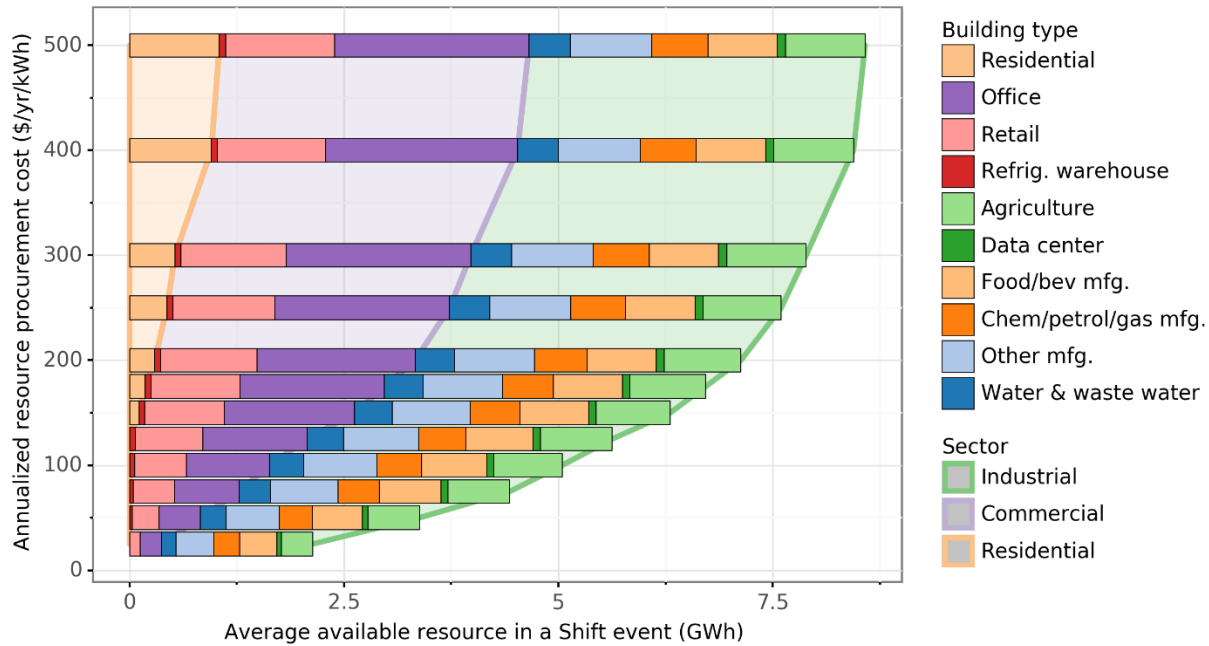
Figure 3-8. The 2030 Shift supply curve, disaggregated by sector (light shaded areas) and building type (colored bars).

It is also interesting to investigate the specific end uses that contribute Shift resources within each building type. Figure 3-9 shows this disaggregation of the Shift supply in 2030, at the BTM battery price threshold of $150/yr/kWh, also broken out by utility service territory. Key conclusions from this figure are that the pumping resource is primarily to be found at agricultural sites[41] and primarily in PG&E's service territory; that the relative importance of residential and office commercial buildings varies by utility, and that level 2 (fast) EV charging is the predominant residential end use that can be enabled for Shift at an expected cost below the battery price referent, given the Reference model assumptions.

---

[40] The proportional breakdown by building type is similar in other years.
[41] Shiftable loads from the Water and waste water sector include both pumping and process loads, and the resource shown is primarily from the latter end use.

Figure 3-9. The Shift resource available in forecast year 2030 at the BTM battery price referent ($150/yr/kWh), disaggregated by utility service territory, building type, and end use.

3.3. Seasonal variation in the Shift resource

As discussed in section 1.2.4.2, there is expected to be significant seasonal variation in both the need for Shift DR and in the available resource. In this section we examine this variability in more detail within the results of the DR-Path model. When examining seasonal Shift resources in this study, for the sake of simplicity we considered four "seasons," each covering three consecutive months, with "spring" beginning on March 1.

First, we looked at seasonal variability in the *need for* Shift by examining the average expected number of daily Shift dispatches in each season. This value was computed from the DR filter function that is used to represent the relative probability that different potential Shift events will be dispatched (see appendix section C.2.5 for details). Although the Shift filter represents a *relative* (not absolute) dispatch probability, by assuming that the most probable dispatch is 100 percent likely to occur, we can convert the Shift filter into time series of absolute dispatch probabilities. From those, the daily average number of expected Shift dispatches can then be straightforwardly computed. Figure 3-10 presents the resulting dispatch frequency of Shift DR, by season and for the entire year, under the ramping-mitigation dispatch probability model used in the Reference modeling run (see appendix section C.2.5.2 for details). As shown, if Shift is used to mitigate steep ramps, it is most likely to be dispatched in the winter months and least likely in the summer months, with Shift dispatches occurring roughly two thirds as frequently in summer winter than in winter. Spring and fall have roughly equal dispatch probability for Shift, slightly below the winter dispatch frequency and slightly above the annual average. Overall, Shift is likely to be dispatched slightly more than once per day under the assumptions used here.

Figure 3-10. The relative frequency of Shift dispatch, by season, in the Reference model run for forecast year 2030. The plot shows the expected average number of daily dispatches of Shift DR in each season, and over the course of the full year, under the assumption that the most likely dispatch in the year is 100 percent likely to occur, and all other potential shifts have dispatch probability that scales according to the "ramping" assumption used in the Reference model run.

Perhaps not surprisingly, the seasonal variation in Shift dispatch frequency depends strongly on the assumed probability model used for Shift dispatch. In our reference scenario, we assume that the probability of Shift dispatch in any given hour is proportional to the size of the generation ramp required in the surrounding four-hour period. Figure 3-11 shows the total net demand in the IOU service territories, as forecasted in 2030 by LBNL-Load, for each day of the year and averaged by season. Curves are plotted relative to the daily peak to facilitate comparison of the morning and evening ramps in generation that would be required on each day. The absolute size of the morning and evening ramps is smaller in the summer than in other seasons (despite the fact that the absolute peak is higher), and the day-to-day variability is also somewhat enhanced compared to the fairly self-similar spring days. Both of these features will lead our dispatch probability model to assume a lower likelihood of Shift dispatch on summer days than in the other three seasons, and this is reflected in Figure 3-10. For a discussion of the differences in Shift dispatch seasonality in different dispatch probability scenarios, see appendix section C.2.5.2.

Figure 3-11. Forecasted system-level net load (gross demand less solar and wind generation) for 2030, normalized to each curve's peak value, in each day of the year (solid gray lines) and averaged by season (dashed black lines). There is evident variability in the size of the morning and evening ramps, both seasonally and from day to day within each season.

We next consider seasonal variation in the *availability of* Shift. Figure 3-12 shows the Shift supply curve by season, for the Reference model run with our default technology and weather assumptions. (For details on the calculation of seasonal supply curves, see appendix section C.2.3.) Substantial variation by season is evident, with roughly 30 percent more shiftable energy available in summer than in winter. Spring and fall resources are intermediate between the summer and winter values and are relatively close to the annual average. This result is to be expected since, as shown in Figure 3-3, a large fraction of the Shift resource arises from pumping and HVAC loads, which will generally be larger in the summer.

Thus, we find a noteworthy tension between the seasonal need for Shift and availability of Shift: the need for Shift is largest in the winter, when its availability is lowest, and the need is smallest in summer when the availability is highest. It will be important to consider this dynamic in detail when planning for real-world implementation of Shift DR. One significant potential mitigating factor will be electrification of space and water heating loads, which will increase the shiftable wintertime load. We explore this factor in more detail in the next section.
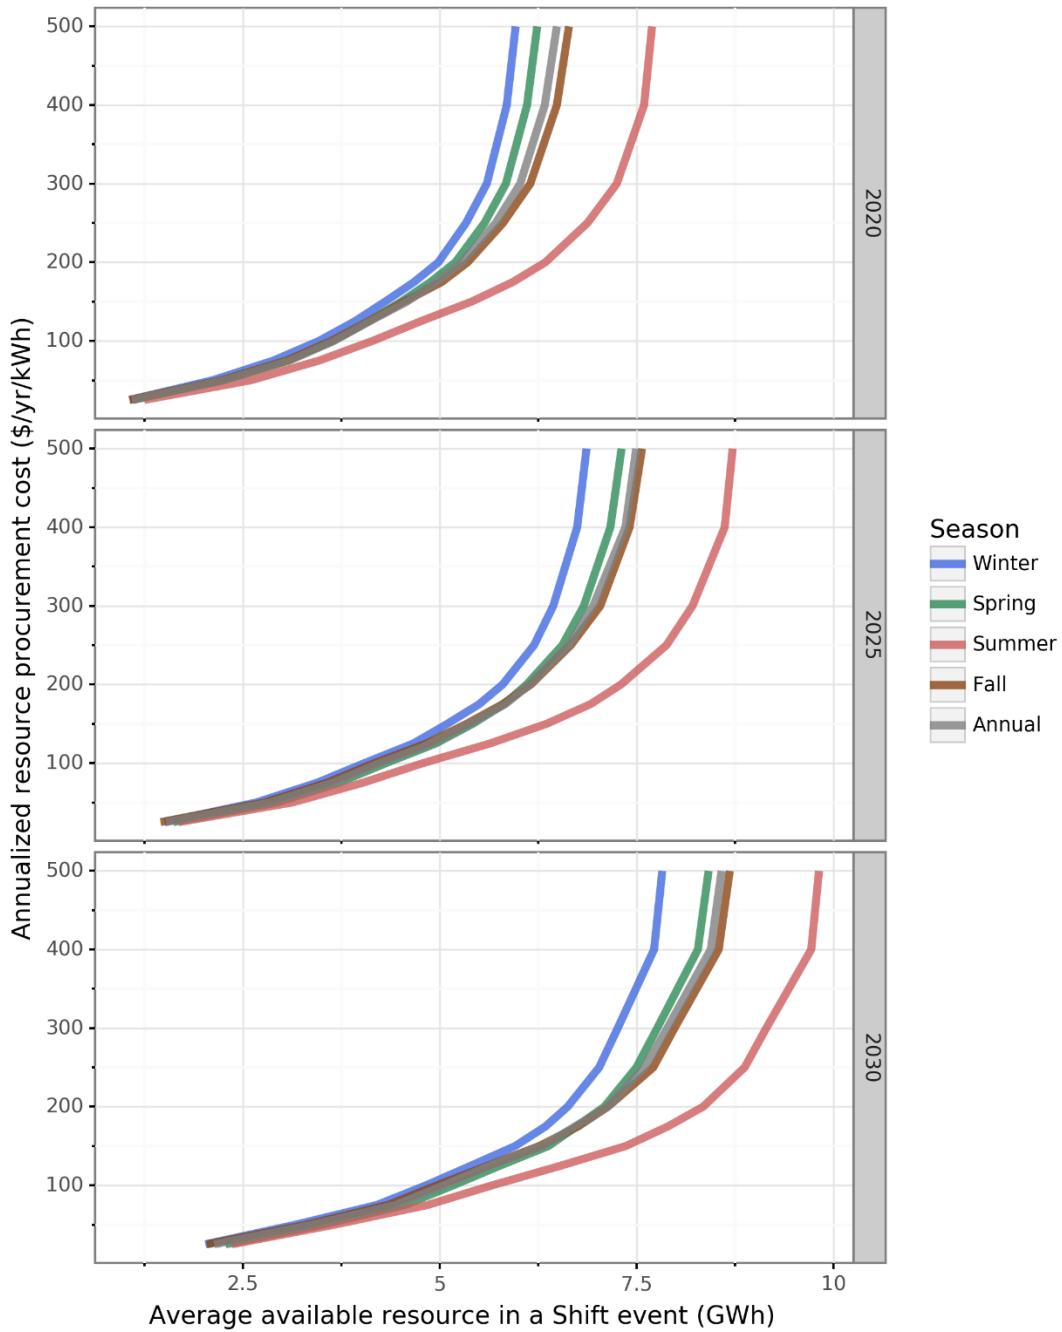
Figure 3-12. Seasonal supply curves for Shift, for the Reference model run, with default technology and weather assumptions (Medium technology, 1-in-2 weather). Seasonal supply curves are computed by averaging the available resource over a single season, rather than the full year.

3.4. Emerging Building Electrification Loads as a Shift resource

As discussed in section 2.3 and appendix section B.2.2, this study includes an alternative forecast scenario that assumes substantial electrification of space heating and water heating in the residential sector, in addition to the electrification of private vehicles that is included in the Base electrification scenario. In this section we examine the impacts that these new electrical loads can have on the Shift resource in California. Specifically, we compare results from the Additional electrification model run to the results of the Reference model run, which uses the Base electrification inputs (see

Table 2-4). Because the adoption of electrification technologies proceeds slowly in the early years of the Additional Electrification forecast, there is negligible difference from the Base Electrification scenario in forecast year 2020; therefore, we focus here only on 2025 and 2030.

3.4.1. Impacts of electrification on the annual average Shift resource

Figure 3-13 shows the supply curves for Shift in 2025 and 2030, by technology scenario and with 1-in-2 weather assumptions, for the Reference model run and the Additional electrification run. There is a small but non-negligible increase in the available Shift resource under the Additional electrification scenario at procurement prices above roughly $200/yr/kWh. There is little difference between the two scenarios at lower prices.

To understand the source of the increased Shift resource from electrification, we can look at the Shift supply curve disaggregated by end use, focusing on the residential sector, where the additional forecasted electrification loads reside. Figure 3-14 shows this supply curve disaggregation in 2025 and 2030. In both years, no residential Shift resource is available below $100/yr/kWh. Very small Shift resources from space heating loads become available at the lowest prices shown, growing to a resource that is slightly smaller than the space cooling resource at the highest prices shown. Water heating loads, by contrast, are only available at prices above $400/yr/kWh, and the resource remains relatively small in both years.

Figure 3-13. Shift supply curves in 2025 and 2030, by technology scenario and with 1-in-2 weather assumptions, comparing the Additional Electrification and Base Electrification scenarios.

Figure 3-14. The Shift supply curve for the residential sector, disaggregated by end use, in 2025 and 2030, under the Additional electrification scenario.

Careful examination of the water heating end use in Figure 3-14 reveals a potentially surprising outcome: from 2025 to 2030, the Shift resource from electrified water heating becomes both smaller and more expensive, despite the fact that adoption of electric water heating grows over this period in the forecast model. This result arises because, at the same time the electrification of water heating is growing in the residential sector, the stock of electric water heating equipment is becoming more energy efficient, as the existing stock of gas and electric-resistance water heaters are replaced with electric heat pump water heaters. Heat pump water heaters use approximately threefold less electricity than electric resistance

water heaters to perform the same service, but the cost of technically enabling an electric water heater to perform DR is dominated by fixed costs (see appendix section D.2.2). As the load per site shrinks, the cost to enable the load for DR increases per unit of enabled load. As a result, although the total *available* water heating load increases substantially between 2025 and 2030, the total *affordable* load at a given procurement price for Shift becomes smaller. This represents an unusually severe case of competition between energy efficiency (EE) and DR, in which the adoption of efficient technology yields a reduction in the available DR resource for the affected end use, even in the context of a rapidly increasing contribution to load from that end use at the system level.

3.4.2. Impacts of electrification on the seasonality of Shift

Because space heating, a highly seasonal load, represents the bulk of the additional Shift resource in the Additional electrification scenario, it is reasonable to ask what impact electrification has on the seasonal variation observed in the previous section. Figure 3-15 compares the seasonal supply curves for 2030 in the Base and Additional electrification scenarios with the default technology and weather assumptions. The additional Shift resource from electrification loads significantly reduces the amount of seasonal variation in the supply curve, in particular increasing the winter supply curve to be nearly identical to the spring, fall, and annual average supply curves.

At the same time, although these residential electrification loads can reduce seasonal variation in the *supply* of Shift, they have a negligible impact on seasonal variation in the *need* for Shift because they represent a relatively small fraction of the overall CAISO system load. The seasonal dispatch frequencies shown for the Reference model run in Figure 3-10 do not change significantly when plotted for the Additional electrification run. Thus, we find that the growth in residential electrification considered here can substantially mitigate the tension between the need for Shift and availability of Shift that we uncovered in the previous section.

Figure 3-15. Seasonal supply curves for Shift in 2030, with default weather and technology assumptions, for the Base and Additional electrification scenarios.

3.5. Imagining a transformed market for DR technology

The basic assumptions for the DR-Path model were based on the costs and load-shifting capabilities of present-day DR-enabling technologies, with incremental technological advances assumed in the future, resulting in relatively slow and smooth trends in technology cost and performance. These choices reflect the priorities of the regulatory forum in which the model assumptions were originally developed and a goal of grounding the research in our understanding of technology as it is deployed in the field, mostly providing Shed service to manage peak loads. In the case of Shift, however, the resource has never been implemented at scale, and the future potential strongly depends on a number of emerging technologies (e.g., EV charging and smart thermostats). In this situation, qualitative changes in the market are likely, but difficult to predict, meaning that the incremental approach taken in DR-Path may be quite

conservative. To develop a sense for the potential impacts of disruptive market changes, in this section we examine the primary barriers that constrain the Shift resource estimated in DR-Path, and we construct a scenario that substantially reduces certain barriers to gain a sense of the potential for Shift in a substantially transformed DR market.

Three primary factors limit the availability of Shift (or any DR) as a grid resource within the DR-Path modeling framework:

- *Technology performance*, as represented by the fraction of the affected load that can be shifted, and the duration over which the shift can occur, for a particular piece of enabling technology (e.g., a PCT).
- *Technology cost*, which, for Shift, is expressed in annualized equipment and operating costs per unit of consumption (i.e., $/yr/kWh for Shift) that the technology would enable for DR participation if installed at a particular customer site.
- *Customer participation rates*, expressed as the fraction of customers who choose to participate in a DR program for a given annual incentive payment (also expressed per unit of participating DR).

Figure 3-16 shows the role these factors play in constraining the available Shift resource in the Additional Electrification model run. A green bar indicates the technical shift potential for each shiftable end use— that is, the total quantity of energy consumption that is available to be shifted on average during times of system need, if every customer site were enabled to shift 100 percent of the relevant end-use load. The orange bars show the quantity of Shift that could be enabled by installing specific Shift-enabling technologies that can be obtained for a total technology cost of $500/yr/kWh[42] or less. Even at that fairly high price, there are large quantities of technical Shift potential that cannot be enabled, e.g., for sites that have relatively small loads. The blue bars then show the additional impact of customer participation rates on the available Shift resource. These bars indicate the amount of customer load that would actually be expected to participate in DR programs, at a total cost level of $500/yr/kWh including customer incentives, based on DR-Path's customer participation model. As shown, limited customer participation can result in a significant further reduction in the available Shift resource.

---

[42] This is the maximum value that we have shown in the supply curve plots. Above this price there are minimal returns to additional investment.

Figure 3-16. The impact of cost and customer participation on the potential for Shift. Each green bar shows the technical potential (total amount of energy that is shiftable in principle) for a particular end use in 2030. The orange bars show the fraction of the technical potential that can be enabled for a technology cost of $500/yr/kWh. The blue bars show the quantity that is available from enabled and participating customers at a total cost of $500/yr/kWh. (As elsewhere in this report, HVAC refers to commercial space conditioning loads, while space cooling and heating refer to residential end uses.)

Figure 3-17 illustrates the impact of equipment cost on the overall cost-competitiveness of enabling Shift at different sites. Shown are the annualized costs per unit of shiftable energy ($/yr/kWh, circa 2015) to install and operate Shift-enabling technology, for each shiftable end use, at a typical site in each one of the customer clusters modeled in DR-Path. The variation in cost can be large from site to site for a given end use, because each site has a different amount of shiftable load that is available to be captured. For instance, the (residential) space cooling and space heating end uses will typically be enabled for Shift by

installing a PCT, which has a fixed cost per household. Since some households have relatively small space conditioning loads (e.g., homes in mild climate zones or with tight construction), the cost per unit of shiftable energy will be higher for these sites compared to sites with large cooling loads. As shown in Figure 3-17, the variation in site-level loads drives a variation in technology cost, per unit of DR enabled, that exceeds an order of magnitude. In addition to the cost variation at the customer site level, there is also substantial cost variability among the different end uses because a different enabling technology is being considered for each end use, with a different underlying equipment cost. This cost variation by site and end use means that certain loads that are technically shiftable will not be available at a given procurement price, because the cost to install the technology is too high. For instance, the minimum BTM battery cost is shown as a vertical dashed line in Figure 3-17. Any point in the figure that sits at a higher cost will not be cost-competitive with a battery. In the context of Figure 3-16, any point that sits above $500/yr/kWh will not be counted in the corresponding orange bar.[43] Technology costs can thus explain much of the reduction between the technical potential and available resource for end uses such as water heating or space conditioning loads.
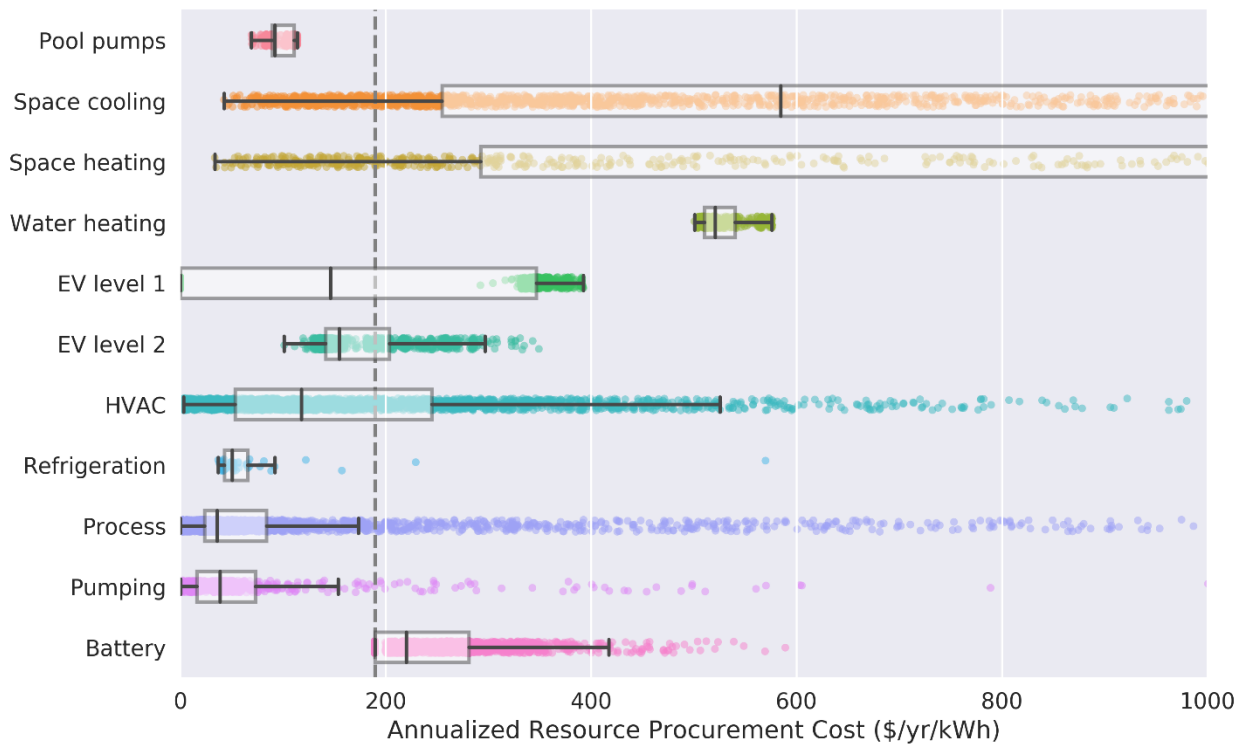
Figure 3-17. Installed costs, circa 2015, for Shift-enabling technologies, by end use, for a typical site in each customer cluster modeled in DR-Path. The minimum battery cost is shown as a vertical dashed line for reference. Box plots show the interquartile range (IQR, middle 50 percent of cluster costs), and whiskers show points beyond the IQR that are within 1.5 times the magnitude of the IQR.

Technology performance levels for Shift also explain a portion of the decline in resource between the green and orange bars in Figure 3-16, since most Shift-enabling technologies are not capable of shifting

---

[43] Note, however, that Figure 3-17 presents costs circa 2015. These costs will be forecasted to decline somewhat by 2030 in DR-Path, so some additional points (e.g., for water heating) may fall below the $500 price level.

100 percent of the energy consumed by an end use during a particular Shed period. For instance, Figure 3-17 shows that EV charging technologies are all less costly than $500/yr/kWh, but the available resource for this end use in Figure 3-16 is nevertheless smaller than the technical potential. This is because the technical potential shown for EV charging represents the energy consumption over one half of a twelve-hour window. We assume that customers will be unwilling to forgo all charging for a full six-hour period, but will instead demand some minimal amount of charging in case of emergency, thus allowing only a portion of their consumption to be shifted.

The final reduction to the participating resource shown in Figure 3-16 (blue bars) is driven by the customer participation rates predicted at a given customer incentive level by DR-Path's propensity-score model. This model, described in detail in Appendix F of the Phase 2 report, uses historical data on participation in DR programs by different customer classes to predict the fraction of customers within each cluster who will participate in DR at a given incentive level. The blue bars then represent the quantity of load that is technologically enabled and participating in Shift DR, if the $500/yr/kWh procurement price represents a budget that must cover both the technology costs and the customer incentive payments, as well as any program administration or marketing costs. Because DR participation by residential customers has been quite limited historically, the resulting estimates of available load to participate are substantially diminished for residential end uses (pool pumps, space heating/cooling, water heating, and EV charging). If customers could participate in larger numbers with lower barriers to entry, the available resource could be larger, perhaps substantially.

3.5.1. Shift in a scenario with accelerated technology progress

There are numerous reasons to believe that at least some of the cost and participation assumptions used in our primary results are fairly conservative. For instance, our technology cost inputs for water heating assume that communication technology and a thermostatic mixing valve would need to be installed as an add-on to the existing water heater, incurring significant on-site labor costs. If instead, the relevant technology were incorporated directly into the appliance, the overall enablement cost could fall dramatically, owing to efficiencies in the manufacturing process. Additionally, the DR-Path propensity score model is based on data from DR programs using traditional enabling technologies for Shed, such as direct load control (DLC) switches, that may reduce customers' comfort or convenience and suppress participation rates. In the case of Shift DR, there are often strategies, such as pre-cooling or thermal storage, under which customers' perceived level of service is unchanged. Thus, Shift may be able to achieve significantly higher levels of participation than traditional Shed DR programs.

Both of these lower-cost and higher-participation pathways could significantly increase the fraction of the technical Shift potential that can be brought online as a participating resource, compared to what is shown in Figure 3-16. Such changes also represent qualitative transformations in the DR market that are difficult to forecast reliably, however, so in our primary results we have considered more incremental evolution in costs and participation rates. Nevertheless, it is interesting to explore what the impacts might be of a more dramatic shift in the landscape for DR technologies and participation. Therefore, we performed a separate Market Transformation model run, in which we modified the inputs of the Additional Electrification run to assume that costs would decline, and participation rates would increase, both by a factor of three to ten,

for certain key end uses and enabling technologies in the residential sector. These changes represent optimistic best guesses about the scale of improvement that might be possible with some qualitative shifts in manufacturing or DR program design; thus, the results should not be taken as realistic forecasts, but rather as an illustrative example of the increased Shift resource that may be available.

Figure 3-18 compares in the 2030 residential-sector Shift supply curve in the Additional Electrification model run, which we use as a baseline, and the Market Transformation model runs. The increase in the Shift resource is dramatic. In the baseline scenario, the residential sector contributes only a small fraction of a GWh of Shift at the $150/yr/kWh battery threshold, and the water heating end use only begins to become available at $500/yr/kWh. In the Market Transformation scenario, which dramatically decreases water heating technology costs and boosts residential participation, there is more than a GWh of Shift available at the battery threshold, and water heating is the dominant contributing end use at this price. Other end uses, such as pool pumps and space cooling, also see significant increases owing to the assumed higher customer participation rates and improved performance of the Shift-enabling technology in the Market Transformation scenario. These results serve to demonstrate the scale of the untapped Shift resource that could be made available in a transformed market for Shift DR.

Figure 3-18. Residential supply curves for Shift in 2030, by end use, for the Additional Electrification and Market Transformation model runs.

4. Discussion, Recommendations, and Next Steps

Over the last five years, the influence of California's aggressive action to add renewable energy generation to the grid have begun to reshape the dynamics of the power system, resulting in negative prices, frequent curtailment of VRE generation as a balancing strategy, and ever-steeper ramping of the net load. Our study is designed to help reveal the opportunities to use flexible loads that can shift the timing of demand to help manage a transition to a low-carbon energy system. It reveals that there are both significant opportunities and significant challenges ahead. The scale and impact of Shift in the future depends strongly on the technology and market pathway California goes down, with qualitatively different outcomes depending on the choices that will be made in the coming years related to research and development, pilot programs, market design, and regulatory actions. In the sections below, we synthesize the scale of the potential Shift resource, identify pathways consistent with deployment to balance VRE at large scale, discuss potential barriers, and describe the necessary effort to better understand and capture the opportunity. We finally discuss possible modeling updates that could allow future iterations of this work to more thoroughly account for the Shift potential in California and more precisely match it to system needs.

4.1. The Shift resource and the need for Shift

4.1.1. The size of the Shift resource

Even based on conservative assumptions, Shift could be a substantial resource and contribute significantly to the integration of renewable energy in the California grid. In the default scenario, Table 3-1 shows a potential 5.3 GWh of non-battery Shift in 2020, at a cost equivalent to or less than that of BTM batteries, growing to 6.3 GWh in 2030. Recall that these estimates represent the annual average quantity of energy consumption that is available to be shifted over a period of several hours, at times of system need. Although there are some realistic limitations on how frequently this resource could be called,[44] it is plausible that it could be utilized twice per day, to mitigate both the morning and the evening ramps. As a point of comparison for scale, the average day in the spring of 2019 saw approximately 5 GWh of curtailed VRE generation and an evening net load ramp of approximately 10 GW (as discussed in section 3.1.2); the potential 2020 Shift resource in California may be large enough, in principle, to absorb all of that average VRE curtailment, and meet up to half of the ramping need,[45] at a cost equal to or below that

---

[44] For most shiftable loads, there will be a recovery period following the shift, during which the building must return to normal operation before shifting again. For instance, a building that is pre-cooled and then allowed to warm to slightly above its usual thermostat set point must undergo cooling to regain its set point before participating in another Shift event. This recovery period is analogous to the period required for a BTM battery to recharge after a discharge event before discharging a second time.

[45] As discussed in section 3.1.2, if the Shift takes place in a four-hour window, with two hours of Shed and two hours of Take, the average load reduction during the Shed period would be ~2.5 GW, with a corresponding load increase of the same amount during the Take period, shrinking the overall ramp size by ~5 GW. This estimate likely represents the maximum ramping reduction that could be achieved.

of BTM batteries (with the caveat that curtailment driven by local transmission constraints may be more difficult to offset with Shift, as also discussed in section 3.1.2).

Furthermore, our preliminary analysis (presented in section 3.5) of the potential for market transformation in Shift-enabling technologies indicates that the resource base could be much larger. The technical potential for Shift in the residential sector is approximately two to ten times what is considered economically available and market accessible based on the default model assumptions. With qualitative changes in the DR market, and technology advances that are plausible given emerging technology, the cost-competitive Shift resource in the residential sector could be triple what we find using default scenario assumptions. At present, though, no market structure exists for dispatchable Shift DR products, leaving this potential resource largely inaccessible. To ensure a reliable and low-cost future power system in California, it will be valuable to explore the most straightforward ways of tapping into this novel form of DR.

It is also worth briefly noting that, although this study focuses on Shift DR as a means to facilitate VRE integration by mitigating steep ramps and curtailment, the Shift resource can also have ordinary generation-capacity value to the extent that it provides peak load reduction. We briefly discuss the peak-management implications of Shift in Box 4.

4.1.2. Seasonal variation in the Shift resource and the need for Shift

A key upgrade to the modeling approach in this study compared to previous work is the ability to assess seasonal variation in the capabilities and need for Shift. The results reveal significant seasonal variation in both the size of the Shift resource and in the frequency with which it is likely to be dispatched. These effects are in tension with one another: the resource is smallest in winter, when it is most likely to be called upon, and largest in summer when the dispatch frequency is lowest. In spring and fall, however, when the dispatch frequency is also high, the available resource is close to the annual average value, so the annual value remains a reasonable representation of the overall resource for much of the year (see section 3.3 for details). Nevertheless, the reduced resource availability in winter may reduce the value of Shift compared to the potentially more stable load flexibility that would be provided by BTM batteries. This assessment will ultimately also depend on other seasonally dependent uses of storage, which is beyond the scope of this study (e.g., to increase resilience to blackouts, manage demand charges, or provide other grid services). It will be important to consider approaches to enabling a portfolio of Shift that matches the seasonal variation in the resource with needs on the grid and capabilities of complementary resource like BTM batteries.

4.1.3. Batteries and Shift

By comparing the cost of Shift to the cost of BTM batteries, we are able to put the resources in a priority order: whatever the value of load flexibility to the grid, Shift resources that are less costly than BTM battery storage should be procured first. This price referent approach is a useful benchmark for considering the scale of the DR resource, and is similar to accepted practice for valuing peak load reductions by comparing the cost of Shed to the avoided cost of new generation and T&D infrastructure.

Box 4. The Shed implications of Shift

This study focuses on Shift as a means to mitigate steep ramps and curtailment in the context of high renewables penetration. Because shed periods will tend to coincide with net load peaks, Shift will also naturally provide similar benefits to Shed DR by reducing peak demand. One can roughly estimate the implied Shed resource associated with a given Shift resource as follows.

The response to a particular Shift dispatch will typically involve reducing load for two to four hours during the defined shed period (with an offsetting load increase of similar duration during the corresponding take period). In that sense, a 1 GWh Shift resource would correspond, very roughly, to 0.25 to 0.5 GW of Shed.

However, it is also important to consider the timing of the need for Shed, which occurs primarily at the highest peaks in system net load, as compared to Shift, which could be dispatched daily. Thus, while the probability of dispatching Shift to mitigate steep ramps is significant throughout the year, with a gentle seasonal variation (see Figure 3-10), the Shed probability is concentrated in the summer and early fall. As discussed in section 3.3, the summertime Shift resource is significantly larger than the annual average. This suggests that the implied Shed impact of a given Shift resource may be larger than the simple estimate above would imply, although a more detailed calculation would be required to fully quantify this effect.

Thus, we can conclude that the 6.3 GWh of Shift resource available in 2030 would correspond, at a conservative estimate, to 1.6 GW of implicit Shed, which assumes a four-hour load reduction and no gain from seasonal variation. Depending on the details of the loads involved and the timing of the system peak, the implicit Shed could be considerably larger.

Moreover, it is important to highlight that we have not undertaken a detailed study of the costs for storage, and rely on other work to define the expected costs. In this study we also have ignored certain practical constraints (e.g., building codes or space constraints) that may limit the ability to install BTM batteries in some buildings. Additionally, we do not consider potential future constraints on storage related to scarcity in rare minerals or unpriced externalities related to mining and production of batteries. In some cases, enabling Shift may be the easiest path to load flexibility, even if it is nominally more expensive than BTM storage based on our assumptions. It is worthwhile, then, to consider pathways to increasing the overall quantity of Shift and the quantity that is available at low cost.

4.2. Pathways to enabling and expanding California's Shift resource

There are a range of possible future pathways to enabling Shift on the California grid, both in the near term and in the longer-term context of an evolving technological landscape. In this section we describe strategies that may help to boost the size of the available resource at a given cost. We follow this with a discussion of the next steps and improvements in modeling Shift that will be important to consider in future phases of the DR Potential Study, and the needed data to inform such efforts, to refine our ability to chart a course for load flexibility in California.

4.2.1. The low-hanging fruit

As with any new resource, it will be advisable to begin implementing Shift DR via the sources that provide the largest resource for the lowest cost—the low-hanging fruit, so to speak. Sections 3.1 and 3.2 can serve as a roadmap to effective targeting of the most readily available Shift resources. As shown in Figure 3-3, the three end uses that provide the largest and the least costly sources of Shift are industrial process loads, industrial pumping loads, and commercial HVAC loads. There is, however, considerable geographic variation in the relative importance of these end uses for providing Shift (Figure 3-7), with much of the pumping resource in particular being concentrated in agricultural regions. Additionally, the relative importance of different building types to providing Shift also shows considerable variation by region (Figure 3-9). For a given end use, the most effective technological pathway to enabling Shift may vary depending on programmatic budgets and resource acquisition targets. For instance, Figure 3-5 shows that, at low per-unit procurement costs, a moderate amount of commercial HVAC can be enabled to provide Shift via pre-cooling strategies using connected thermostats and energy management systems. At higher costs, however, installing TES systems can enable a significantly larger quantity of Shift. It will be important for the design of future DR programs focused on Shift to take into account these and other factors relevant to the specific sectors, regions, and end uses they are targeting.

4.2.2. Developing next-generation Shift

Work could commence today to design and deliver retrofit programs and technology deployment for Shift in the sectors with low-hanging fruit, but how could new resources emerge by "growing the orchard?" Part of the reason for the relatively slow growth in the Shift resource that is available at a lower cost than BTM batteries (Table 3-1) is that the cost of batteries also declines over the same period, so the battery cost threshold moves down the supply curve, even as the supply curve grows (see Figure 3-2). One might conclude from this that battery storage will become a growing part of the solution to California's load flexibility needs. However, we note that much of the potential load flexibility identified in Table 3-1 would be enabled by emerging technologies, and it may be possible to deploy market transformation strategies (e.g., polices or program designs) that bring down the prices of Shift-enabling technologies more rapidly than battery costs. This would expand the pool of load flexibility that is available more cheaply than BTM batteries, especially as new electrification loads come onto the grid.

Beyond capturing the relatively straightforward, near-term sources of Shift, then, this study also explored a number of pathways to expanding the size of the potential Shift resource and to enabling a larger fraction of that resource at a lower cost. Broadly speaking, these pathways can be divided into three categories: new end uses and enabling technologies, improved customer engagement, and strategies for cost reduction. In the area of new end uses and technologies, we investigated future scenarios for the electrification of private transportation and of residential space and water heating. We found that these new electrification end uses represent (Figure 3-14) a potential source of Shift that is growing over the forecast period, but the overall resource is small and expensive relative to the primary end uses discussed above. Nevertheless, because of the introduction of a significant load in space heating, the electrification technologies have the potential to significantly reduce seasonal variation in the available Shift resource (Figure 3-15).

New end uses like electrified heating and transportation represent important opportunities to "bake in" flexibility from the start and deploy devices with capabilities built-in to communicate and respond to the future grid. Given the importance of electrification for reducing emissions, these sectors are poised for significant growth. Instead of requiring relatively costly retrofits, integrated systems from the factory or at initial installation represent a key market transformation opportunity for Shift that should be carefully considered.

Furthermore, as Figure 3-16 shows, the Shift supply curve estimated by DR-Path represents only a small fraction of the technical Shift potential embedded in customer loads, even at the highest procurement costs considered, especially for loads in the residential sector. The primary factors limiting the quantity of Shift captured in the supply curve are technology costs and customer participation rates. Finding ways to reduce these barriers may greatly increase the attainable Shift resource.

To understand the cost barriers and how they may be overcome, it is instructive to consider the case of water-heating electrification. In this study we have envisioned technologically enabling this load to participate in Shift by using a communicating control system to heat the water in a storage-type water heater to a higher temperature than desired during Take periods, then using a thermostatic mixing valve at the tank outlet to provide water at the desired temperature, allowing the system to cease heating altogether during Shed periods, and relying on the stored excess heat to provide adequate hot water service. This technological configuration would require additional installation work by a trained plumber, as well as an uncommon and relatively expensive add-on technology in the mixing valve. Because individual heat-pump water heaters have relatively small average loads during typical Shift windows, and because each water heater would need to incur the same technology cost, the average cost per unit of Shift enabled is quite high (see Figure 3-17). This scenario represents the approach that would most likely be taken given present-day technology options; however, in a future in which load shifting can provide significant benefit to the customer (either through TOU rate savings or direct program incentives), water heater manufacturers would be incentivized to incorporate the communications and mixing-valve technologies into the appliance itself, which would substantially reduce both the technology and installation costs. In this way, the opening of new market opportunities, such as a wholesale market product for Shift, can have a dramatic impact on reducing the costs of related technologies. Recognizing the opportunity, the Bonneville Power Authority (BPA) has recently developed a market transformation plan for this technology (BPA 2018).

The customer-participation limitations on the modeled Shift resource arise because the DR-Path customer-participation model is based on historical rates of participation in Shed-type DR programs, which have tended to require customers to forgo a measure of comfort or convenience in order to participate in DR (e.g., by turning off their air conditioners during the hottest peak hours), and customers must weigh the participation incentive against the loss in energy service. By contrast, Shift participation may be achievable with little or no impact on customer comfort or convenience. Again considering the example of water-heating electrification, the strategy of using the storage tank as a thermal storage reservoir means that the customer's hot water demands can be fully met at the same time as a load shift is being executed. In a real-world example, pre-cooling strategies have been shown to improve comfort ratings reported by building occupants when participating in Shed programs (Herter and Okuneva 2013).

From the customer's perspective, then, there may be little downside to weigh against the participation incentive, and developing customer-engagement strategies that emphasize this may substantially boost participation rates above what has been observed historically for Shed programs. As with the cost example, the creation of new program designs for Shift may, in and of itself, substantially transform the associated market. Figure 3-18 shows that such a market transformation can yield dramatic increases in the available Shift resource.

4.2.3. Low-cost pathways to Shift via Shape

Finally, although this study focuses on Shift as a dispatchable, market-integrated resource, there may be a significant opportunity to capture a portion of the resource as a Shape resource at low costs, via an improved response to varying electricity rates, either with fixed TOU rates, or via real-time pricing (RTP) that exposes customers to time-varying prices in the wholesale electricity market.[46] As discussed in section 3.1.3, the baseline load shapes for this study included load shifting impacts from the forecasted customer response to TOU rates from the IEPR, and we found that the resource thus captured was quite small. These forecasts rely on historical rates of response, however. A number of the Shift-enabling technologies considered in this study include new networked communicating controls that could respond just as effectively to a price signal as to a dispatch signal from the grid operator, potentially boosting customer response. Empowering customers to minimize their electricity costs using these technologies in the context of TOU or RTP programs could result in significant load shifting, without the need to pay explicit incentives for participating as a dispatchable supply-side resource, thus reducing the cost of capturing the resource.[47] There is also a significant opportunity in this area for third-party firms to provide technologies or services that facilitate customer load shifting in response to time-varying rates, delivering value both to customers (via reduced electricity costs) and to the grid.

A potential down side to TOU rates is that they can change only slowly (over a period of years) as part of the rate-setting process, and thus they are unable to respond to day-to-day variation in grid needs, or to Take and Shed periods that evolve rapidly with time. But to the extent that the grid has some relatively stable and predictable long-term need for daily load shifting (e.g., due to a large solar generation fleet), enabling customers to respond more effectively to TOU rates could represent a least-cost pathway to capturing a significant portion of the Shift resource. In the case where the daily need for Shift is variable (e.g., in a future with high wind penetration), the RTP approach coupled with communicating devices that can read and respond to digital price signals (so-called *prices-to-devices*) could capture a significant component of the available Shift at a lower cost than a market-integrated supply-side approach.

---

[46] Hybrid approaches are also possible: for instance, the Oklahoma Gas and Electric utility offers a variable TOU rate whose peak retail price varies according to the day-ahead wholesale price, while the peak-period definition remains fixed (Oklahoma Gas & Electric 2018).

[47] In addition, program administration costs would likely be lower in price-based programs, because there is no need to compute baseline load shapes for market settlement. In addition, utility program marketing costs to attract customers may be lower if customers are facing default TOU rates and actively seeking means to minimize their electricity bills.

4.3. Recommendations for research and policy

In conducting our research to explore the Shift potential of customer loads and technologies we have identified several core findings related to the DR potential. Here we offer a few recommendations, informed by those findings, for technology pilots and policy concepts that may be able to move California forward. Each of these areas is an important topic for further research. A future with flexible demand that can shift to use more renewable energy when it is available will require a combination of advances in technology, policy, pricing, business, markets, and organizational models for serving electricity users.

Our recommendations are based on considering a vision for future DR, and working to identify near-term actions that are consistent with the vision. What would it look like if a range of customers have loads that can be controlled to make the best use of available VRE generation and support other grid operations?

- Flexible loads, schedulable processes, and behind-the-meter storage assets all act in synchrony to connect with clean energy as it is available.
- The communication and telemetry links that enable this coordination are resilient, affordable, and adaptable, and do not introduce significant costs or discomforts to the customer.
- Customers have an appropriate incentive to shift their net demand based on the available generation and local grid state at their point of connection.
- California's utilities, CCA's enterprises, and communities are organized in structures that provide safe, reliable, and clean electricity.
- New models for organizing and regulating businesses are in place to unlock DR potential.

Achieving this vision will require a substantial and sustained effort, and the concepts we offer are a non-exhaustive set of potential avenues.

4.3.1. The importance of technology demonstrations and pilots

A central theme of the discussion in this section is that creating a market for Shift may, in and of itself, serve to boost the available resource by driving a market transformation. An essential next step to charting a path for Shift in California, then, is to explore various pathways for enabling Shift and begin to estimate the potential market response. The LSWG final report (Gridworks 2019) develops a conceptual framework and a set of criteria for describing and assessing potential load-shifting programs, and it lays out six well-defined proposals for pilot programs. We refer the reader to that report for further discussion of potential path-finding real-world pilots for load shifting in California that could be explored soon.

More broadly speaking, research demonstrations and pilots are needed to prove out the currently available technology for Shift and identify a deployment and R&D pathway that takes advantage of today's opportunities while building new capabilities. These pilots can be designed to help learn about the performance of control technology, the practical challenges around integration and deployment, cost-effectiveness, organizational capabilities, strategies for increased customer participation, and more. Importantly, pilots also have the function of teaching us what we did not know to ask but need to know in order to get the technology to work.

4.3.2. New electric loads and storage

Electric vehicles and electrified space and water heating could represent a major growth in the loads served by the grid in California if trends hold. Energy storage is also growing, which has significant Shift implications since chemical and thermal storage have inherent flexibility. Research is needed to understand the opportunities to make these loads controllable in the right way for Shift.

There is an opportunity, as new DER devices like BTM batteries and EV chargers are deployed, to make sure they include built-in interoperable controls and communications. Otherwise, costly retrofits may be needed. This suggests that we need statewide technology policy that encourages investment in demand-side energy technology that is flexible. As this wave of assets is deployed along with increasing levels of solar and wind power, there will be a need to both better integrate with VRE and support a grid that will be under stress.

4.3.3. Dynamic tariffs

Dynamic prices that are in sync with the cost of serving loads could dramatically simplify Shift, allowing customers to simply pay for what they use and avoiding the need to set an assumed baseline load in order to compensate customers. Incentives will then implicitly in place to use more renewable energy during low-price periods, while reducing demand at peak times. TOU pricing is a start, but, as mentioned, has limited flexibility; therefore, there is growing interest in exploring how we can better communicate TOU or other dynamic prices to devices with faster-changing automated signals using machine-readable digital tariffs.

If dynamic prices are available, the same control technology that could be used to develop market-integrated Shift or Shed may also be able to respond to dynamic prices. Given the reduced need for telemetry (since bill-reduction incentives are simply settled at the existing meter), simpler technology systems should also be able to get prices to devices. Pilots could explore new schemes for price-responsive devices and how existing controls could respond to prices. One fundamental open question regards the price elasticity of electricity use, especially in the presence of automation technology. We will need to understand when, and under what conditions, the California energy system can harness this elasticity for Shift, and we do not know which specific technologies will best support customers' ability to take load during mid-day low prices and reduce load during high price periods.

4.3.4. Communication and control standards

Getting signals to and from thousands or millions of customer loads and devices is a challenge and requires careful planning to ensure interoperability between the cacophony of technology platforms that control these loads. It will be important to explore both simple one-way and more secure two-way communication technologies to support various needs for loads to connect and respond to signals from utilities, CCAs, aggregators, and CAISO.

Defining clear communication and load control standards can help to reduce the cost to integrate new technology and accelerate deployment. It takes collective action to develop standards, however, with

coordination between technology integrators on both the load and grid operations side. Support from public sector organizations can be important to convene private sector businesses that could all use the resulting standards, and it would be appropriate to investigate what is needed in this area to support Shift.

California is actively engaging in public discussions and supporting policies to develop relevant standards with the CEC Load Management Standards effort, plus related activities such as in Title 24 and Senate Bill 49 that are developing requirements for load-flexible appliances and equipment. National smart grid standards efforts have required that standards come from formal standards development organizations, and that these organizations develop certification and compliance testing. Standards such as IEEE 2030.5, OpenADR, and CTA-2045 are all at various stages in development and implementation with different communication architectures and business models. Similar efforts are underway to standardize communication with EVs, which have some differences from the aforementioned standards, but some overlap as well.

Another important innovation could be requiring machine-readable digital tariffs. If dynamic prices are part of a pathway to increase Shift it will be vital that they are available on an up-to-date basis by control systems that could manage loads in response. Even without dynamic prices, it is important for DR loads to be controlled in a way that is bill-friendly and optimizes between grid-level opportunities and direct site benefits.

4.3.5. Field research on customer load flexibility

Our study suggests that there are opportunities for DR across many types of customers and categories of loads. From industrial customers and large water pumping operations to the households and businesses in our neighborhoods, there are both existing and emerging opportunities for loads to respond. Field research projects testing and developing these emerging technologies (prior to larger scale testing in pilot programs) will be important to show how different kinds of DR affect the service provided by loads to customers, helping to calibrate the response to various needs and opportunities on the grid. Since there are opportunities to shift loads and capture VRE on the margin frequently throughout the year, it will be important that the fundamental quality of energy service is good for customers and that the controls are imperceptible in the background (as opposed to infrequent peak demand reduction, which may be more noticeable).

4.3.6. Carbon benchmarking

Connecting demand with clean energy is a fundamental element of reducing the carbon intensity of serving loads in the electricity sector. As Shift is deployed, one important benchmark of success is the GHG performance of the technology or the overall customer load. An approach of carbon benchmarking could involve tracking the total GHG associated with serving loads with various strategies (e.g., not responsive versus using technology to Shift) and compare them. This involves more than just the arbitrage of reducing loads sometimes and increasing in other times, but also may involve changes in the total energy consumed (i.e., the joint carbon impacts of EE and DR strategies). Methods for benchmarking are being established for energy storage through the Self-Generation Incentive Program (SGIP); it may be possible to modify these for use with flexible loads as well. It will also be useful to

explore the value of providing carbon savings data to customers, enabling them to reduce their own carbon footprints; for instance, by using the real-time emissions signal provided by the nonprofit WattTime or providing building-level assessments of carbon use per square foot, similar to present-day energy intensity benchmarking and disclosure.

4.3.7. Integrating EE, DR, and storage

Across our analysis work a common theme has been that DR technology is most valuable and cost-effective when it is part of a multi-use, integrated demand-side management (IDSM) strategy. If DR controls can be deployed along with EE upgrades they can deepen the pool of opportunities and value for both. Once installed, these controls should be able to both participate in Shift and also provide additional peak load reduction (i.e., Shed DR) at critical times. To achieve the accelerated deployment of controls and amplified benefits to customers that this would provide, it will be important to create opportunities for technology firms to operate across these domains. It would be useful if pilots were to focus on trying to deploy multi-use EE and DR to understand customer experiences and organizational and regulatory needs.

Similarly, it will be important to consider how to integrate EE and DR with behind-the-meter storage valuation and control. Many customers are considering investing in storage to increase resiliency and minimize disruptions from power outages (especially in the context of recent public safety power shutoffs), as well as to increase self-utilization of rooftop PV generation. Customers that do install energy storage for such purposes might be able to use their storage to shift loads during optimal times, and co-located PV and battery systems could help mitigate some overgeneration and ramping challenges. There is a substantial opportunity to begin capturing this value during a period of accelerated growth for BTM storage.

4.3.8. Incorporating Shift in Resource Planning

A key next step is integrating load shifting into resource planning models that underlie the IRP process in California. This work is underway, and can leverage the existing knowledge of the capabilities of Shift to estimate the scale of the economically favorable resource by including it as an option along with other renewable integration strategies like energy storage and overbuilding with curtailment.

As a point of reference, in the Reference System Portfolio for the 2019-2020 IRP (CPUC 2020a), the estimated curtailment is ~2% of RPS-qualified renewables in 2020, rising to ~8% by 2030 and ~20% by 2045 (CPUC 2020b). If one values avoiding this curtailment at the $40 per megawatt-hour (MWh), which is the levelized cost of energy for new solar projects circa 2020[48], the curtailment forecasted to occur on the CAISO grid in 2020 represents roughly $50 million in potential value, rising to $200 million annually

---

[48] As estimated by the U.S. Energy Information Administration (EIA)

by 2030 and $1.6 billion by 2045. Similarly, the IRP Reference Portfolio projects the annual cost of installing new battery storage at $100 million in 2020, $1 billion in 2030, and $8 billion in 2045.

These estimates for the scale of curtailment and storage show that new approaches to load flexibility, if they are cost competitive, could result in significant value. Simply put, utilizing renewable energy when there is surplus available can reduce the total cost of achieving clean energy goals, if the cost of aligning demand with generation is lower than overbuilding generation or installing energy storage. Integrating Shift into IRP modeling can help to reveal the scale of this opportunity and lead to incorporation of flexible demand and demand response at appropriate levels in grid planning.

4.4. Future directions for modeling Shift in the DR Potential Study

Despite the extensive set of modeling scenarios considered in this work, it is impossible for any study to provide a fully comprehensive accounting of all possible future scenarios for load shifting in California. Here we summarize several important topical areas that were not addressed in this study, which may be important targets for future research. Each of these future modeling directions would be aided by data flowing from pilots, demonstration projects, or larger scale deployment efforts related to the pathways and policy recommendations discussed above.

4.4.1. Modeling additional Shift sources and pathways

Table 2-3 indicates several different end uses and shift-enabling strategies that were not included in the modeling for this study, some of which are at lower technical readiness levels than the proven applications studied here. Additional shiftable end uses that were not modeled, and potential strategies for enabling them, are as follows:

- Refrigeration loads (including both residential refrigeration and commercial refrigeration loads beyond the refrigerated warehouses currently included in the modeling). These loads can be shifted either using pre-cooling strategies such as super-cooling freezer compartments, or by using phase-change materials to provide thermal storage.
- Residential appliance loads that run in fixed cycles (e.g., clothes washing, dish washing). These could be controlled by communicating technologies that execute the cycles in response to a price or dispatch signal, within constraints on cycle completion time set by the customer.
- Agricultural lighting for indoor application (i.e., grow lights). These are often used to simulate the diurnal photoperiod corresponding to a particular growing season, which could be shifted to correspond with periods of high VRE generation.
- Commercial ventilation loads. Although it is unlikely that these could achieve a 100 percent load shed for any significant period of time and still comply with health and safety requirements, modern variable air volume (VAV) systems operating under intelligent controls (e.g., demand-controlled ventilation systems) may be able to achieve adequate ventilation by ventilating aggressively during Take periods and more conservatively during Shed periods, using a similar strategy to pre-cooling for space cooling loads.

- New loads from the electrification of commercial space heating[49] and water heating. Similar strategies to those considered here for residential loads would also apply in the commercial sector.
- Charging loads for fleets of medium- and heavy-duty EVs. Since these vehicles tend to exist in centrally managed commercial fleets, coordinated changes to overall operations may be possible to enable load shifting.
- Additional strategies for light-duty EV charging loads. There are two major additional strategies that could significantly increase the available Shift resource from this end use.
  - Enabling Shift by incentivizing commuters to charge their EVs at work, rather than at home, thus shifting their demand to the midday Take period. This approach would require a reconsideration of the basic design of DR-Futures load clusters, which currently assumes that electricity demand in a cluster is inseparable from the buildings in which it is forecast to occur.
  - Two way (vehicle-to-grid, or V2G) charging of EVs, which allows EV batteries to discharge to the grid during Shed periods, which may dramatically increase the value of EVs as a Shift resource.

This work also performs only a cursory assessment of the potential for major qualitative changes in the customer adoption and operations landscape for DR in the context of Shift. As discussed in the previous section, in many cases load shifting can be achieved with much less disruption to customer comfort or convenience than traditional load-shedding DR. This could substantially boost customer willingness to participate in DR programs, as could new money-saving opportunities for customers facing TOU tariffs. Moreover, many potential Shift opportunities rely on emerging or uncommon technologies whose future costs and performance may change dramatically in a market landscape that includes significant utilization of Shift. The Market Transformation technology scenario considers one potential future with dramatically lower technology costs and higher customer participation. However, DR-Path generally assumes technology improvements that are smoothly varying and incremental on the current status quo. It also rests on a customer participation model, developed for Phase 2, that was based on historical participation rates in Shed DR. As new data on load-shifting approaches to DR become available in future pilot programs, it will be important to revisit and refine these components of the modeling.

In addition, in this study we calculated the load shifting from customer TOU response in a separate analysis of the Shape resource, while envisioning Shift as a purely dispatchable, supply-side resource. In fact, as discussed in the previous section, many of the technologies that can enable Shift would also empower customers to respond more effectively to TOU or RTP rates. This would capture some portion of the potential Shift resource without the need to pay any participation incentives, thus reducing its cost. Because the results of this study assume that all Shift resources, beyond the existing forecasts of TOU response, will participate as incentivized supply-side resources, they should be taken as a relatively conservative assessment of the potential for Shift. In future work, it will be interesting to allow for a more dynamic interplay between the Shape and Shift resources, in which the installation of a Shift-

---

[49] The current modeling includes existing electrified heating loads in the commercial sector, which are explicitly disaggregated in LBNL-Load, but it does not model growth in the fraction of commercial heating loads that are electrified.

enabling technology can yield some fraction of the shiftable load as a Shape resource, at no additional cost. Real-world data on TOU responses in the presence of different enabling technologies will be important for informing this future model component.

4.4.2. Modeling the value of Shift to the California grid

Our focus in this work was to provide a detailed accounting of the cost of Shift, but we have not firmly quantified the *value* of Shift to the grid in this study. Put differently, we have presented a supply curve for Shift but not a corresponding demand curve. In large part this is because it is highly uncertain how Shift, presently a speculative resource, would be utilized in practice in the context of real-time dispatch.[50] This uncertainty makes it difficult to assess the value that Shift would bring to the grid, in terms of reducing either the total cost of production or long-term capacity and infrastructure investments. As real-world approaches to enabling Shift take shape, it will be important to develop detailed models for determining their value in the context of expected future generation and storage portfolios. Instead we rely on the price of BTM batteries as a reference against which to benchmark the scale of the resource (see section 4.1.2). This "battery price referent" approach is a useful comparison since Shift is conceptually similar to load management with battery storage. As storage is deployed on the grid, Shift may be thought of as a way to avoid costs or augment the performance of these valuable assets. It also recognizes that the cost of additional energy storage sets a kind of ceiling on the costs that could be associated with DR that achieves similar goals

A full valuation framework for Shift would require that the resource be included in a detailed dispatch and production cost model of the projected future CAISO grid, with consistent assumptions between such a model and DR-Futures, to determine the specific generation resources that Shift would offset under different grid conditions and the reduction in total production cost that this would yield for a given generation portfolio. The operational characteristics of Shift inferred from the dispatch modeling would then need to be used to inform a capacity expansion model, to determine the quantity of Shift that is cost-effective to procure when compared to all other competing resource types. Such a modeling effort was well beyond the scope of the current study. Moreover, for this approach to be accurate it would require a detailed understanding of the operational characteristics of Shift, such as the achievable temporal window and response depth for different end uses, and the recovery time required between Shift events. Since dispatchable Shift products do not currently exist in CAISO or any market, there is a lack of relevant data to support such an effort. In future work, as pilot programs begin to fill in data gaps, it will be important to develop an increasingly fully featured approach to modeling Shift's value to the grid.

4.4.3. Data as a key limitation

The previous sections describe several key data gaps that will limit progress in the accurate modeling of California's Shift resource. Although it is reasonable to imagine that customer participation rates in Shift

---

[50] The Phase 2 study used the RESOLVE model from Energy + Environmental Economics (E3) to estimate a demand curve for Shift by inputting the total estimated Shift resource as a daily budget of shiftable energy. Since our modeling approach in Phase 3 yielded Shift resources on a per-event, rather than per-day basis, this approach to using RESOLVE was not viable for this study. We did not re-run RESOLVE or set up a new production cost model to estimate updated demand curves for Phase 3 of the study.

programs will be higher than historically observed for Shed DR, there are currently no such programs whose participation rates can be observed, limiting the ability to improve the modeling of customer participation. In addition, without well-designed studies to measure customer response to TOU rates with and without key Shift-enabling technology, it is difficult to accurately account for the additional Shape resource that such technology will deliver. Finally, without real-world operational data on Shift as a dispatchable, market-integrated product, it will be difficult to build the appropriate power system models to accurately assess the value of Shift to the grid. This paucity of data to guide future modeling efforts only sharpens the importance of implementing pilot programs such as proposed in the LSWG final report (Gridworks 2019), in order to bring Shift into better focus.

5. Conclusion

This report describes the results from Phase 3 of the modeling and analysis effort known as the California Demand Response Potential study. Whereas the first two study phases covered a variety of grid services referred to as Shape, Shift, Shed, and Shimmy, this report focuses on Shift, in an effort to better understand the resource and explore potential pathways to enabling it. The study took this focus because the Phase 2 study indicated that Shift may have the greatest potential future economic value, but there is no history in California of using DR to intentionally move energy consumption from one time to another, rather than simply shedding demand during critical peaks. A stakeholder-involved effort at the CPUC in 2018, the Load Shift Working Group, explored a number of near-term pathways to piloting the Shift resource in California (Gridworks 2019); this study attempts to widen the scope to estimate the full future potential for Shift in California and long-term pathways to achieving it.

Throughout this Phase 3 study, we have used the cost of a BTM electric battery as a benchmark against which to compare the cost of procuring Shift. This is similar to the price-referent approach taken in the Phase 1 study, in which we compared the cost of Shed DR to the cost of a natural gas combustion turbine peaker plant. This approach is conceptually helpful, in that it allows us to identify the virtual storage resource that can be obtained at a cost below that of physical battery storage. However, this benchmarking approach does not allow us to place a firm value on the potential Shift resource, in terms of the benefit it would provide to the grid. Computing such a valuation would require a good understanding of how Shift would be utilized operationally. Shift is a purely conjectural resource at the current time, so, absent real-world data on utilization, it is difficult to accurately characterize how it would impact the cost of operating the grid. Hence, for this study, we used the BTM battery benchmark as a useful point of comparison.

Our analysis found that, under the default set of assumptions, a 5.3 GWh of Shift is currently available at a cost below that of BTM batteries, on average over the course of the year (with the value ranging from 4.0 to 8.0 GWh, depending on the specific assumptions chosen). A central conclusion of this study is that, if this resource could be brought online and dispatched twice daily, it would be sufficient to utilize *all* of the more than 5 GWh of average daily VRE curtailment that occurred on average in the spring of 2019, while also substantially easing the need for flexible generation capacity to meet ramping needs. This result emphasizes the importance of finding pathways to rapidly enabling Shift. The primary end uses that can provide this resource, as identified under our default assumptions, are industrial process and pumping loads, as well as commercial HVAC, with significant geographical variation in the relative importance of these three end uses. New pilot programs targeting these key end uses in the appropriate regions may represent the most effective near-term approach to tapping into California's existing Shift resource.

We have also found, however, that the most readily accessible resources will not grow fast enough to keep pace with the expected expansion of renewable generation in California. Under our default assumptions, the Shift resource that is less costly than BTM batteries is expected to increase by less than 20 percent by 2030, as set against a 60 percent increase in renewable generation. Additionally, we observe substantial seasonal variation in the size of the Shift resource, which is in some degree of conflict with expected seasonal variation in the need for Shift on the grid. Thus, in addition to pursuing the most

straightforwardly available Shift, it will be important to explore pathways to increasing the size and diversity of the available resource.

Electric vehicle charging is an emerging end use, poised for significant growth, which possesses some natural flexibility; however, we find in this study that the timing and location of EV charging limits its potential as a Shift resource if the strategy used is limited to shifting energy consumption at the site of the baseline charging load. More novel strategies for load flexibility, such as two-way (V2G) charging, or incentivizing commuters to charge at their workplace, rather than at home, could significantly increase the EV resource; this will be an important topic for future work.

Residential loads represent another plausible target for expanding the Shift resource. Under our default assumptions, the residential sector yields a relatively small Shift resource at all price levels, compared to the commercial and industrial sectors; and nearly all of this resource is costlier than BTM batteries. We also explored the potential for residential-sector Shift from the future electrification of space and water heating. As winter-peaking end uses, these loads have the capacity to significantly reduce the seasonal variation in Shift; however, under our default assumptions, the available Shift resources from these loads are small relative to non-residential sources, and enabling them is significantly more expensive than the cost of BTM battery storage.

Under our default modeling assumptions, the primary factors limiting the size and affordability of potential Shift resources are customer willingness to participate and the costs of enabling technologies. Shift represents a fundamentally new approach to DR, and new program designs are possible with Shift that could have significantly improved customer acceptance. Moreover, Shift is often enabled by emerging technologies whose costs may fall dramatically if Shift creates a new opportunity to drive adoption. To investigate the potential for Shift in a market landscape transformed by the presence of a new DR paradigm, we considered an alternative scenario with dramatically improved customer participation and drastically reduced technology costs in the residential sector. We found that in this scenario the residential Shift resource had the potential to increase by several times and become cost-competitive with BTM batteries.

This study demonstrates that there is an urgent need to continue to evaluate the opportunities to procure low-cost load flexibility in support of California's future electric grid. We have demonstrated that there is a sizeable Shift resource available today, which has the capacity to significantly ease the present-day challenges of renewable integration on the California grid; however, the best pathways to enabling this resource remain unclear. In this study we have mainly envisioned Shift as a fully market-integrated, dispatchable resource, but there are also numerous other programmatic approaches that could capture some fraction of the available load flexibility. These include load-modifying (Shape) approaches such as TOU and RTP tariff structures, as well as "market-informed" approaches that allow customers to respond to wholesale market conditions without being directly dispatched. Specific, detailed proposals for pilot programs in each of these dispatchable, market-informed, and load-modifying categories are laid out in the LSWG final report (Gridworks 2019). An important next step will be to use these or other pilot programs to gain a better understanding of the market and operational dynamics of Shift, to provide critical path-finding for California's future renewable and flexible electricity grid.

References

Ali, Syed Muhammad Hassan, Manfred Lenzen, and Jing Huang. 2018. "Shifting Air-Conditioner Load in Residential Buildings: Benefits for Low-Carbon Integrated Power Grids." *IET Renewable Power Generation* 12 (11): 1314–23. https://doi.org/10.1049/iet-rpg.2017.0859.

Alstone, Peter, and Mary Ann Piette. 2019. "Electric Load Flexibility Analysis for the Clean Energy Transition." *Notre Dame Journal on Emerging Technologies* in review.

Alstone, Peter, Mary Ann Piette, and Peter Schwartz. 2018. "Integrating Demand Response and Distributed Resources in Planning for Large-Scale Renewable Energy Integration." In *2018 Summer Study on Energy Efficiency in Buildings*. Pacific Grove, CA: ACEEE.

Alstone, Peter, Jennifer Potter, Mary Ann Piette, Peter Schwartz, Michael A. Berger, Laurel N. Dunn, Sarah J. Smith, et al. 2016. "Interim Report on Phase 1 Results. 2015 California Demand Response Potential Study: Charting California's Demand Response Future." California Public Utilities Commission. https://www.cpuc.ca.gov/WorkArea/DownloadAsset.aspx?id=10632.

———. 2017. "Final Report on Phase 2 Results. 2025 California Demand Response Potential Study: Charting California's Demand Response Future." California Public Utilities Commission. https://www.cpuc.ca.gov/WorkArea/DownloadAsset.aspx?id=6442452698.

Bahrenian, Aniss, Jesse Gage, Sudhakar Konala, Bob McBride, Mark Palmere, Charles Smith, and Ysbrand van der Werf. 2017. "Transportation Energy Demand Forecast, 2018-2030." CEC-200-2017-010. California Energy Commission.

Baroiant, Sasha, John Barnes, Daniel Chapman, Steven Keates, and Jeffrey Phung. 2019. "California Investor-Owned Utility Electricity Load Shapes." CEC-500-2019-046. California Energy Comission. https://ww2.energy.ca.gov/2019publications/CEC-500-2019-046/CEC-500-2019-046.pdf.

Bedir, Abdulkadir, Noel Crisostomo, Jennifer Allen, Eric Wood, and Clement Rames. 2018. "California Plug-In Electric Vehicle Infrastructure Projections: 2017- 2025." CEC-600-2018-001. California Energy Commission. https://www.nrel.gov/docs/fy18osti/70893.pdf.

BPA. 2018. "CTA-2045 Water Heater Demonstration Report." Bonneville Power Administration. https://www.bpa.gov/EE/Technology/demand-response/Documents/Demand%20Response%20-%20FINAL%20REPORT%20110918.pdf.

CAISO. 2016. "Fast Facts: What the Duck Curve Tells Us about Managing a Green Grid." https://www.caiso.com/Documents/FlexibleResourcesHelpRenewables_FastFacts.pdf.

———. 2019a. "2018 Annual Report on Market Issues and Performance." California Independent System Operator. http://www.caiso.com/Documents/2018AnnualReportonMarketIssuesandPerformance.pdf.

———. 2019b. "Managing Oversupply." 2019. http://www.caiso.com/informed/Pages/ManagingOversupply.aspx.

———. 2019c. "Open Acces Same-Time Information System (OASIS)." 2019. http://oasis.caiso.com/mrioasis/logon.do.

Callaway, Duncan S., Meredith Fowlie, and Gavin McCormick. 2017. "Location, Location, Location: The Variable Value of Renewable Energy and Demand-Side Efficiency Resources." *Journal of the Association of Environmental and Resource Economists* 5 (1): 39–75. https://doi.org/10.1086/694179.

Carew, Nick, Ben Larson, Logan Piepmeier, and Michael Logsdon. 2018. "Heat Pump Water Heater Electric Load Shifting: A Modeling Study." Ecotope. https://ecotope-publications-database.ecotope.com/2018_001_HPWHLoadShiftingModelingStudy.pdf.

CEC. 2017a. "Building Climate Zones California, 2017." 2017. https://www.energy.ca.gov/maps/renewable/BuildingClimateZonesMap.pdf.

———. 2017b. "Final 2017 Integrated Energy Policy Report." CEC-100-2017-001-CMF. California Energy Commission. https://efiling.energy.ca.gov/getdocument.aspx?tn=223205.

CPUC. 2013. "Order Instituting Rulemaking to Enhance the Role of Demand Response in Meeting the State's Resource Planning Needs." R.13.09.011. California Public Utilities Commission. http://docs.cpuc.ca.gov/PublishedDocs/Published/G000/M077/K151/77151993.PDF.

———. 2016. "RPS Calculator v6.2 Spreadsheet Model." https://www.cpuc.ca.gov/WorkArea/DownloadAsset.aspx?id=11513.

———. 2017a. "Decision Adopting Steps for Implementing the Competitive Neutrality Cost Causation Principle, Requiring an Auction in 2018 for the Demand Response Auction Mechanism, and Establishing a Working Group for the Creation of New Models of Demand Response." D.17-10-017. California Public Utilities Commission. http://docs.cpuc.ca.gov/PublishedDocs/Published/G000/M198/K319/198319901.PDF.

———. 2017b. "Decision on Residential Rate Reform for PG&E, SCE, and SG&E and Transition to Time-of-Use Rates." D.15-07-001. California Public Utilities Commission. http://docs.cpuc.ca.gov/PublishedDocs/Published/G000/M153/K110/153110321.PDF.

———. 2020a. "2019-2020 Electric Resource Portfolios to Inform Integrated Resource Plans and Transmission Planning." D. 20-03-028. California Public Utilities Commission. https://docs.cpuc.ca.gov/PublishedDocs/Published/G000/M331/K772/331772681.PDF.

———. 2020b. "RESOLVE Model Inputs and Results Used for 2019 IRP Reference System Plan Decision." California Public Utilities Commission. https://www.cpuc.ca.gov/General.aspx?id=6442464143.

Delforge, Pierre, and Joseph Vokovich. 2018. "Can Heat Pump Water Heaters Teach the California Duck to Fly?" In . Pacific Grove, CA: ACEEE.

EIA. 2019. "Levelized Cost and Levelized Avoided Cost of New Generation Resources in the Annual Energy Outlook 2019." U.S. Energy Information Administration. https://www.eia.gov/outlooks/aeo/pdf/electricity_generation.pdf.

Gallo, Giulia, Brian Gerke, Jingjing Liu, Mary Ann Piette, Peter Schwartz, and Peter Alstone. 2018. "Mobilizing the Anti-Duck Brigade: Technology and Market Pathways for Load-Shifting Demand Response in California." In *2018 Summer Study on Energy Efficiency in Buildings*. Pacific Grove, CA: ACEEE.

Garcia, Cary, and Chris Kavalec. 2016. "California Energy Demand Updated Forecast 2015 – 2025." CEC-200-2016-016-SD. California Energy Comission. https://ww2.energy.ca.gov/2014publications/CEC-200-2014-009/CEC-200-2014-009-SD.pdf.

Gridworks. 2019. "Final Report of the California Public Utility Commission's Working Group on Load Shift." https://gridworks.org/wp-content/uploads/2019/02/LoadShiftWorkingGroup_report.pdf.

Haeri, Hossein, Lakin Garth, Joan Wang, Jim Stewart, Mark Osborn, Shawn Shaw, Mary Knipe, et al. 2018. "Demand Response Potential in Bonneville Power Adminsitration's Public Utility Service Area Final Report." Bonneville Power Administration. https://www.bpa.gov/EE/Technology/demand-response/Documents/180319_BPA_DR_Potential_Assessment.pdf.

Herter, Karen, and Yevgeniya Okuneva. 2013. "SMUD's 2012 Residential Precooling Study – Load Impact Evaluation." Sacramento Municipal Utility District. http://www.herterenergy.com/pdfs/Publications/2013_Herter_SMUD_ResPrecooling.pdf.

Hledik, Ryan, Judy Chang, and Roger Lueken. 2016. "The Hidden Battery: Opportunities in Electric Water Heating." The Brattle Group. https://brattlefiles.blob.core.windows.net/files/7167_the_hidden_battery_-_opportunities_in_electric_water_heating.pdf.

Itron. 2018. "2016-2017 Self-Generation Incentive Program Impact Evaluation." CPUC SGIP Working Group. https://www.cpuc.ca.gov/General.aspx?id=7890.

Kavalec, Chris. 2014. "California Energy Demand Updated Forecast, 2015-2025." CEC-200-2014-009-CMF. California Energy Comission. https://ww2.energy.ca.gov/2014publications/CEC-200-2014-009/CEC-200-2014-009-CMF.pdf.

Kavalec, Chris, Ashish Gautam, Mike Jaske, Lynn Marshall, Nahid Movassagh, and Ravinderpal Vaid. 2018. "California Energy Demand 2018-2030 Revised Forecast." CEC-200-2018-002-CMF. California Energy Commission. https://efiling.energy.ca.gov/getdocument.aspx?tn=223244.

Keeney, Kevin, and James Braun. 1997. "Application of Building Precooling to Reduce Peak Cooling Requirements." *ASHRAE Transactions* 103 (1): 463–69.

Lazard. 2018. "Lazard's Levelized Cost of Storage Analysis—Version 4.0." https://www.lazard.com/media/450774/lazards-levelized-cost-of-storage-version-40-vfinal.pdf.

NREL. 2019. "EV Insfrastructure Projection Tool (EVI-Pro)." 2019. maps.nrel.gov/cec.

Oklahoma Gas & Electric. 2018. "Standard Pricing Schedule: R-VPP." https://www.oge.com/wps/wcm/connect/c41a1720-bb78-4316-b829-a348a29fd1b5/3.50+-+R-VPP+Stamped+Approved.pdf?MOD=AJPERES&CACHEID=ROOTWORKSPACE-c41a1720-bb78-4316-b829-a348a29fd1b5-mhatJaA.

Pindyck, Robert S. 2019. "The Social Cost of Carbon Revisited." *Journal of Environmental Economics and Management* 94 (March): 140–60.

Raghavan, Shuba V., Max Wei, Dan Johnson, Florin-James Langer, and Daniel M. Kammen. 2020. "Decarbonizing Residential Space Heating and Cooling in a Changing Climate." *Energy Policy* (in review).

Raghavan, Shuba V., Max Wei, and Daniel M. Kammen. 2017. "Scenarios to Decarbonize Residential Water Heating in California." *Energy Policy* 109 (October): 441–51. https://doi.org/10.1016/j.enpol.2017.07.002.

Ricke, Katharine, Laurent Drouet, Ken Caldeira, and Massimo Tavoni. 2018. "Country-Level Social Cost of Carbon." *Nature Climate Change* 8 (10): 895–900. https://doi.org/10.1038/s41558-018-0282-y.

Siler-Evans, Kyle, Inês Lima Azevedo, and M. Granger Morgan. 2012. "Marginal Emissions Factors for the U.S. Electricity System." *Environmental Science & Technology* 46 (9): 4742–48. https://doi.org/10.1021/es300145v.

Smith, Margaret, and Jonathan Castellano. 2015. "Costs Associated With Non-Residential Electric Vehicle Supply Equipment." US Department of Energy. https://afdc.energy.gov/files/u/publication/evse_cost_report_2015.pdf.

Tarroja, Brian, Felicia Chiang, Amir AghaKouchak, Scott Samuelsen, Shuba V. Raghavan, Max Wei, Kaiyu Sun, and Tianzhen Hong. 2018. "Translating Climate Change and Heating System Electrification Impacts on Building Energy Use to Future Greenhouse Gas Emissions and Electric Grid Capacity Requirements in California." *Applied Energy* 225 (September): 522–34. https://doi.org/10.1016/j.apenergy.2018.05.003.

U.S. EPA. 2016. "The Social Cost of Carbon Fact Sheet." https://www.epa.gov/sites/production/files/2016-12/documents/social_cost_of_carbon_fact_sheet.pdf.

WattTime. 2019. "Working With Regulators." 2019. https://www.watttime.org/solutions/regulators/.

Yin, Rongxin, Doug Black, Mary Piette, and Klaus Schiess. 2015. "Control of Thermal Energy Storage in Commercial Buildings for California Utility Tariffs and Demand Response:" LBNL-1003740. Lawrence Berkeley National Laboratory. https://escholarship.org/uc/item/8f47t077.

Appendix A. Methods and Assumptions for Greenhouse Gas Analysis of Shift

This emerging opportunity to add Shift DR is directly responsive to the emerging priorities and constraints placed on the grid by climate change. There has always been a priority for grid planners to find the least cost pathway to deliver reliable and safe electric power; using DR to Shed load at peak times contributes by reducing the need for additional transmission and generation to carry loads at peak times. Today and in the future, incorporating the harms and costs from globally significant emissions like carbon dioxide (CO_2) and local air pollution means that more zero-emissions resources are needed, changing the dynamics of net load on the grid. Shifting the timing of loads should enable more efficient use of variable solar and wind generation conceptually, and the specifics of how much flexibility and load shifting "helps" with reducing greenhouse gas (GHG) emissions is important for charting a pathway forward to implementation. This section describes a first order analysis of possible GHG impacts if Shift were available.

The analysis described below was completed in the process of supporting the California Public Utilities Commission (CPUC) Load Shift Working Group, which met 11 times through 2018, with contributions from 85 stakeholders representing 63 organizations (notably including all three major investor-owned utilities (IOUs) in California, the California Independent System Operator (CAISO), a number of demand response (DR) industry representatives, and public interest intervenors). The final report from that group describes a summary of the outcomes (Gridworks 2019). A range of pathways that were discussed in the report and refined in the group were all designed as different options to incentivize and dispatch load shifting. These possible options form the basis for our analysis below, which assesses the potential load impacts and resulting changes in grid performance for each.

A.1. Conceptual framework for GHG and Shift

When loads are shifted to better match renewable generation, what would we expect?
- **Reduced cost** to serve load from arbitrage in prices and increased load in low price periods.
- **Reduced peak** loads since "take" tends to happen mid-day and "shed" in peak times.
- **Reduced curtailment** from more demand during curtailment hours.
- **Reduced emissions** from more demand during zero or low emissions times, and avoided demand at high emissions time.

To drive our analysis we use two approaches to represent how load shifting can support low-GHG grid operation: (1) a compliance framework where load shifting is a cost-reducing measure to meet binding renewable portfolio standard (RPS) targets by utilizing otherwise-lost clean energy, or (2) an arbitrage framework based on timing demand to reduce the marginal emissions of the grid.

In the **compliance framework** for value, loads that are able to **utilize available zero-carbon power** by increasing demand during a curtailment event can avoid the need to turn down renewables. Otherwise the renewable generators would be curtailed. Instead of a lost opportunity, the output of the renewable generators "counts" towards achieving the RPS standards set by California. By avoiding a lost opportunity to generate, this also reduces the need to build additional renewables that would be required to hit the same RPS targets. The most straightforward method for valuing this Shift is through assessing the avoided cost of additional new-build wind or solar in the future. Based on the expected cost for building new utility-scale solar PV in 2023, the implied societal opportunity cost is $40/megawatt-hour (MWh) (EIA 2019). If load shifting is less expensive, it can be an economically optimal choice as part of a portfolio of investments in the context of the binding RPS cap. We do not do additional analysis with this framework, but describe it here to clarify that this is one way to conceptually describe the role of load shifting. In this framework, the assertion is that RPS caps are binding and will drive the overall level of emissions on the grid. Various investments are made to meet this cap, and load shifting can be part of those investments to reduce the overall cost.

The **arbitrage framework** for estimating the value of Shift presents another way to frame the opportunity. When loads are reduced at peak times and increased in times of curtailment, the effects on the grid are that the marginal generators serving peak loads (which tend to be fossil fuel units) are turned down. The marginal emissions during curtailment times are zero if there are no constraints on the otherwise-curtailed solar or wind generators serving the load. The typical marginal emissions from a natural gas power plant is about 0.3 tonnes CO_2 per MWh. If every shifted megawatt-hour resulted in the reduction of emission by 0.3 tonnes, the implied value based on the social cost of carbon—a highly uncertain parameter (U.S. EPA 2016; Ricke et al. 2018; Pindyck 2019)—ranges from $10–$200 per shifted MWh. The analysis described in brief below assesses the potential gains from load shifting using an arbitrage framework.

A.2. Analysis methods for estimating GHG impacts

In the analysis presented below, we assess both the scale of avoided curtailment (the first framework) and the implied emissions savings related to marginal GHG (the second framework) for a range of pathways to achieve Shift. The basis for the analysis is retrospective, using characteristics of the grid during the 2017 operations year and assessing the implied outcomes with various load-shifting concepts. Based on the actual profiles for demand, solar and wind generation, market prices, and estimated marginal emissions, we estimate what value shifting loads would have provided to the grid. *Value* here is defined in terms of reduced curtailment, reduced cost to serve load, and reduced emissions due to arbitrage.

Since there are no active, large-scale Shift programs in California, the approach we take is to impose speculative load shifts on the demand profile using a set of scenarios that represent a range of possible dispatch profiles; these are designed to cover the range of concepts from the LSWG and support understanding of the sensitivity of grid value.

The datasets from 2017 that we combined as a basis for this analysis are illustrated in Figure A-1. The sources of load and price data for the analysis were the CAISO Managing Oversupply web page (CAISO

2019b) and CAISO OASIS (CAISO 2019c), an online portal to access CAISO prices. In order to estimate the impacts of changes in load on operational emissions, we use estimated marginal emissions. These estimates were provided to us by WattTime (WattTime 2019) and were originally developed and used in support of a California energy storage incentive program (Self-generation Incentive Program, or SGIP) impact evaluation (Itron 2018). These hourly estimates are based primarily on the marginal prices in the real time energy market. These prices reveal information about the efficiency and emissions of the marginal generator because the bids are typically based on the marginal operating costs of power plants, and it is possible to infer the cost of fuel based on known information about the power sector. In these emissions estimates, periods with typical prices between $30–$50/MWh can translate to emissions based on the cost of natural gas (enabling a conversion from dollars to energy quantity) and estimates of the emissions intensity of natural gas burned in power plants.

When the prices are low or negative, this indicates renewable generators are likely on the margin, since there is no additional cost to operate them once built. Prices can be negative due to congestion and the opportunity value of production-based incentives and tax credits that would be gained by producers when they operate. In practice, negative pricing typically indicates that somewhere on the system a renewable generator is being curtailed. Increased load at these times, if it results in a reduction in curtailment and is served with renewables, would not result in additional emissions. Conversely, when loads are reduced during times with higher emissions, this would reduce the load on the marginal generator, reducing fuel consumption and emissions. This is the basic framework that underlies our analysis to estimate the overall emissions impacts from changes in load associated with Shifting and flexible demand—reducing the load in some hours and increasing in others.

We established the expected load impacts (reduction and increase) due to shifting based on a range of concepts that were under discussion as part of the Load Shift Working Group in 2018. These were distilled to six concepts for distinct deployment pathways with different combinations of market frameworks and dispatch methods. These are summarized in Table A-1. The concepts are broadly defined in two categories (Gridworks 2019):

1. **Load modifying:** Flexible loads modify the timing of demand to optimize based on prices, emissions, or other priorities. The concepts described by the LSWG included a range of priority targets and various approaches to communicate these and support customer response.
2. **Market integrated:** Dispatchable loads are controlled through integration with the energy market, coordinating bids and dispatch through the CAISO. This proposed concept is more narrowly defined, based on the Proxy Demand Resource–Load Modifying Resource market integration model that was launched in 2018 by CAISO for behind-the-meter storage, simulating flexible loads that could respond to the same signals and result in the same net load effects as battery storage.

[A]: Net Load
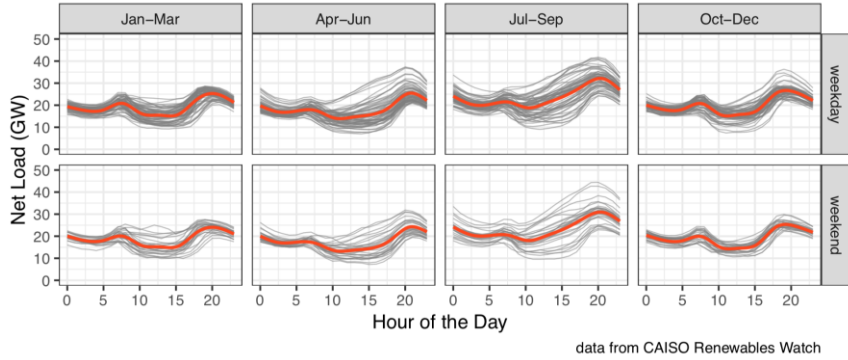
[B]: Energy Market Prices
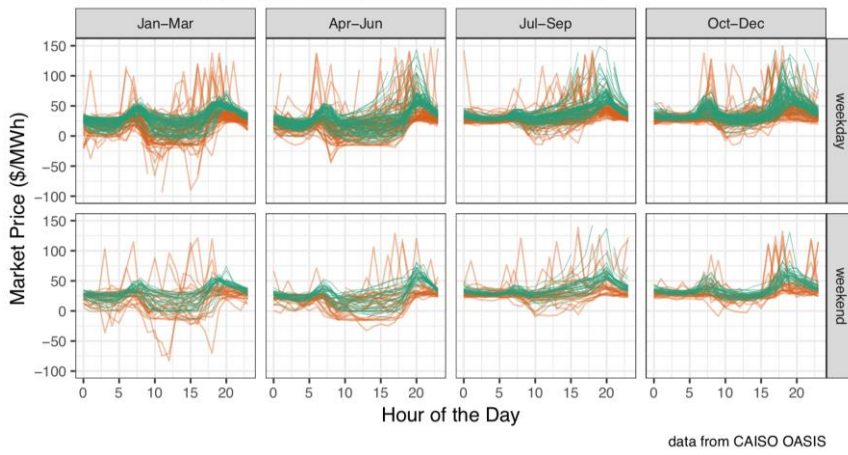
[C]: Estimated Marginal Emissions
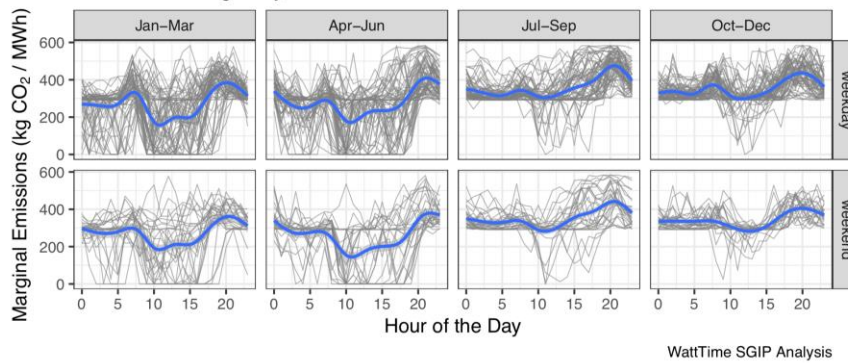
Figure A-1. Operations data and marginal CO₂ estimates from 2017. Data sources indicated.

Table A-1. Summary of the six concepts for Shift presented in the LSWG final report.

Name *Type*	Description
Load Shift Resource 2.0 (LSR2.0) *Market Integrated*	An expanded scope version of the CAISO Proxy Demand Resource–Load Shift Resource (PDR-LSR) that includes flexible loads along with behind-the-meter storage. Loads are dispatched by the system operator to increase during times with negative pricing on the CAISO market.
Critical Consumption Period (CCP) *Load Modifying*	A retail load increase framework where incremental load increases are requested by load-serving entities (utilities) during times of curtailment, and customers choose to respond or not. The concept is designed for large customers who can significantly increase load when instructed.
Market-informed demand automation service (MIDAS) *Load Modifying*	Establish a price or emissions signal that is published through an application programming interface (API), enabling connected devices and systems to respond.
Pay for a load shape (P4LS) *Load Modifying*	Target load shapes are defined, and end-use control is oriented towards achieving those targets. The target could be based on expected prices, emissions, or curtailment at either the system or local scale and can be updated daily, weekly, monthly, etc.
Market integrated distribution service (MINTDS) *Load Modifying*	This framework is similar to the LSR 2.0 concept, but would include a distribution system layer as a primary focus. Loads are dispatched to minimize the impacts on the distribution system first, and any remaining Shift capabilities are dispatched through LSR 2.0.
Distribution load shape *Load Modifying*	This framework is similar to the P4LS, with a focus on defining load shapes that are primarily responsive to distribution system constraints and also to system-level needs.

Source: LSWG Final report (Gridworks 2019).

For both categories, the specifics of how load impacts were defined and applied is described in more detail below. For any particular pathway, the load impact time series includes both load increase (Take) and decrease (Shed) periods. These load impact time series are then multiplied by the hourly price and emissions data in each hour and summed to develop an estimate of the operational changes for each:

- The change in overall emissions based on emissions arbitrage / net effect of load changes on the operation of the marginal units.
- The change in costs to procure energy on the real time market, a price arbitrage analysis.

We also estimate two additional features from the load impacts that are helpful for describing the value to the power system:

- The quantity of curtailment avoided (based on the sum total increased load during times with negative prices, and capped at the reported total curtailed energy in each hour). This is useful for communicating the fraction of curtailment that might be avoided by particular strategies.
- The change in annual peak net load (based simply on the difference in the maximums). This is an additional potential pathway to value for flexible loads, since reduced annual peak demand leads to a reduced cost for peak capacity. Grid operators in California (and many other regions) make payments to generators for these services that represent a significant annual cost.

A.3. Context and Caveats

Because of the context of this study and simplifying assumptions, it is important to keep the following in mind when interpreting these results and considering their applicability for other places and time periods:

1. **The California ISO footprint is a special case** with positive correlation between prices, net load, and emissions. Note in Figure A-1 how the typical daily shape of loads, prices, and emissions estimates are similar. This makes it "easy" to choose a target to prioritize for compared to places with negatively correlated features. For example, at times when coal generators are the marginal unit, there are typically lower prices but higher emissions compared to gas, which may be the marginal unit at other times of day. This dynamic exists in regions with significant remaining coal fleets, like the Midwestern United States (Callaway, Fowlie, and McCormick 2017). Since California does not have significant coal generation, the trade-offs are typically between gas and zero-carbon generators.

2. **As more renewables are added, the opportunity for avoiding curtailment grows.** We described these trends previously. One can take our results using 2017 as a conservative case since the value of and opportunity to Shift continues to grow.

3. **This analysis is designed to be a first order view** into the range of emissions and grid impacts that are plausible from load shifting. The goal is not to estimate the magnitude precisely for each resulting metric, but to understand the likely "direction" of change related to shift and the order of magnitude of the opportunity. Other approaches using more detailed analytic techniques would be appropriate for estimating the magnitude more precisely currently and for future scenarios, e.g., estimating marginal emissions with dispatch models and statistical learning methods (Callaway, Fowlie, and McCormick 2017; Siler-Evans, Azevedo, and Morgan 2012). We are aware of three sources of error related to our approach:

    a. We did not account for elasticity in prices or marginal emissions in this analysis. By omitting this, we expect that these first-order estimates are biased towards higher apparent gains from price and emissions arbitrage.

    b. We did not account for spillover gains or losses to other participants in the market from improved price stability and changes in the operation of the energy and ancillary services markets.

    c. We did not complete a statistical analysis of the annual peak capacity, instead simply reporting the maximum.

These trade-offs are made in exchange for simplicity and understandability of our results by stakeholders and policymakers, and reflect the priorities of the forum we developed them for.

A.4. GHG Impacts from Load Shift

A.4.1. Load Modifying DR

We modeled Load Modifying DR load impacts using an approach that roughly simulates the behavior of flexible loads operating in response to signals like prices and estimated emissions. The general algorithm we defined is described below:

1. Identify a "target" signal that the loads are responding to; this signal provides information about when to consume more or less energy and is an hourly time series. The options we include are day-ahead energy market prices, real-time prices, estimated marginal CO_2 emissions, and the systemwide net load.
2. Find the average value of the target for each hour of the day, including all of the days over a specified "averaging period." The averaging periods we included were daily, weekly, monthly, and over three-month periods. For the daily period, there is no averaging. Just the raw target signal is passed through. For the weekly period, the average of the target signal is calculated for each hour of the day (8 AM, 9 AM, etc.) based on the whole week. A similar approach is used for longer averaging periods.
3. Take the "inverse" of the averaged target load shape (multiply each value by negative one). This is the same as "flipping over" the time series. Finally rescale this inverted time series so that on each day the sum total of the time series is one. This way, the rescaled time series represents the fraction of the flexible demand that should be consumed in each hour of the day if a load is responding to the target signal.
4. Adjust the system demand so that it is reduced by approximately 1 percent in each hour (the portion that is "flexible"), and reallocate the sum total of this flexible demand using the rescaled target load shapes from the previous step.
5. Compare the baseline (unadjusted) system load to the adjusted system load to estimate the impacts from shifting.

The net effect of this algorithm is that approximately 1 percent of the system demand is defined as "flexible," and is dispatched *as if it were optimized* according to a price, emissions, or net load signal that is updated with varied frequency. Figure A-2 shows a set of these target load shapes for the three-month averaging period. Note that because of correlation between loads, prices, and emissions the targets are broadly similar.

Normalized Target Load Shape
Facets by 3-month season and Day Type

data: CAISO Renewables Watch, CAISO OASIS, Watttime SGIP Analysis

Figure A-2. An illustration of the target load shape for flexible loads based on the algorithm we use in this analysis. The targets are differentiated by color, and these represent three-month averaging periods.

The outcome for this analysis of load modifying DR impacts is summarized in Figure A-3. Overall, most of the options that were considered were broadly similar. This makes sense among the target types, given the correlation in California between loads, prices, and emissions. Significant reductions in CO_2 were apparent for all options, both operationally and through reduced curtailment. With 1 percent of load shifted in 2017, about 120–150 GWh of curtailment would be avoided (approximately 50 percent of overall curtailment). The typical CO_2 savings per shifted megawatt-hour were 0.12–0.15 tons/MWh, which is about half of the typical emissions from a natural gas power plant. This implies that these "shifted quantities" of energy tended to reduce the emissions of that energy service by about half.

Because the concept of load modifying DR involves daily activity, there are spillover gains in terms of the peak load reduction on the power system. Based on the way we constructed assumptions in the analysis (including the illustrative assertion that 1 percent of load is flexible), we expect 1.5–2 GW lower annual peak loads, which is approximately 3 to 4 percent of the overall CAISO peak. This represents a significant level of performance and is on the same order of magnitude as current-day existing DR programs. We emphasize that this outcome (and others related to magnitude) is not a predictive estimate, and is illustrative based on the assumptions.

The implied savings in the energy market based on this analysis is relatively modest, about $80 million to $100 million annually in 2017. It is important to keep in mind, however, that the opportunity for savings could grow with the growth in curtailment hours in the future.

Only modest benefit was apparent from frequent versus infrequent updates of the target signal. The three-month average signal achieves a significant fraction of the overall value available from more frequently updated signals. Two outliers that contravene this trend are the daily updated responses to real-time price and CO_2. In these cases, the targeted daily updates provided enough information to improve the performance of energy market savings (for real-time price) by a factor of 2x–3x, and improve the performance with respect to emissions (for the CO_2 target) by a factor of +30 percent.

Figure A-3. Results of load shifting for various load modifying DR options. Each option is coded as "LMDR" for load modifying DR, followed by the target signal type ("rtm" for real time prices, "dam" for day-ahead prices, "netload" for the systemwide net load, and "co2" for estimated marginal CO_2 emissions), and with a label indicating the averaging period used. The four plots show the estimated grid impacts for various metrics.

A.4.2. Market Integrated DR

The approach we took for modeling Market Integrated DR was different from the Load Modifying approach described above. For this category we attempted to mimic the kinds of load impacts one might expect for flexible loads that are dispatched by a central grid operator. The basis for this is the CAISO Proxy Demand Response–Load Shift Resource that was developed and deployed in 2018. For this energy market mechanism, a behind-the-meter storage asset is able to place bids to consume energy (which looks like turning down a generator to the energy market). These bids will always be negative (i.e., the battery will be paid to charge when energy market prices are negative). The concept of our modeling extends this framework to loads that could "take" more (increase demand) in these times when there are negative prices.

We modeled these kinds of flexible loads participating in the energy market as follows:

1. We started by defining the total capacity of flexible loads in terms of power: how many gigawatts of load are available to "take" during times of needs. We modeled three of these "bid quantities:" 500 MW, 1 GW, and 2 GW
2. During hours when there is curtailment, the total "take" quantity was defined as the maximum of the available flexible load and the total curtailment quantity. In this way, we limited the total load increase to be less than or equal to the total curtailment.
3. The total load that was increased during the curtailment hours was added up. We assumed that some fraction of this energy that was served during curtailment times could lead to reduced demand in other hours (a load "shed"). This shed was allocated throughout all of the non-curtailment hours in proportion to the systemwide net load. We modeled three different fractions of shed: 0 percent, 50 percent, and 100 percent of energy neutrality. A 0 percent shed means that

even if there are times when loads are increased to capture curtailment, there is no decrease in load during other times of day. A 100 percent energy neutral shed means that for every kilowatt-hour that is increased, there is a decreased kilowatt-hour at another time of day.

The results of this illustrative modeling exercise are shown in Figure A-4. This plot shows the case of 100 percent energy neutral behavior of loads. It is notable that because most of the curtailment events are in the first six months of the year, those months are the times with the vast majority of the events.

Figure A-4. Illustrations of the load impacts from market integrated DR. Each line on the plot is an individual day of activity, for various bid quantities.

The outcomes of the load impacts for these scenarios are shown in Figure A-5 and summarized here:

- The level of curtailment using market integrated PDR was significantly reduced. Even with a 500 MW flexible load, nearly half of the 2017 curtailment would be reduced. There were diminishing returns to more load bid into the market given 2017 dynamics. About 50 to 70 percent of the curtailment in 2017 could have been avoided with 500 MW to 2 GW of dispatchable PDR.
- The operational CO_2 reductions depend on also having some fraction of load shed in non-take periods, since it is possible to have some non-zero marginal emissions in a period with negative pricing. This meant that in the cases where there was no "shed" on the other side of the load increase, the overall CO_2 emissions were increased. In real operational contexts one would expect this to be a rare and unlikely outcome, since many flexible loads achieve their flexibility by changing the scheduling of processes or demand.

- The energy market cost savings are modest, only $5–$20 million annually in 2017. On the basis of the gains per shifted MWh, the highest unit gains went to lower bid quantities, on the order of $40/MWh. These are modest arbitrage opportunities.

- It is not shown in the plots, but the total quantity of load shifted in market integrated PDR is about one-tenth that of the LMDR options modeled. This is because PDR is a more "targeted instrument" to reduce curtailment only in hours when it is happening versus load modifying DR that is more durable and persistent.

- Based on the assumptions of our modeling, there were no reductions in the peak load on the system from this form of market integrated DR, since there happened to not be any curtailment of renewable energy on the annual peak load day in 2017.

Figure A-5. Results from load impacts of modeled DR for market integrated options. The various scenarios include 500, 1,000, and 2,000 MW bid quantities and 0, 50, and 100 percent "Shed Fractions" that define the level of energy neutrality. The bars for these scenarios are labeled accordingly on the left.

A-11

Appendix B. LBNL-Load Phase 3 updates

This appendix summarizes major changes to the LBNL-Load model that were implemented for the Phase 3 study. Major changes include use of the 2017 IEPR (CEC 2017b) to update the load forecasts and extend them out to 2030, updated assumptions and input data for modeling the growth of electric vehicle (EV) charging loads, and the addition of new loads from the electrification of certain building end uses (residential space heating and residential water heating). In Phase 2 of the study (Alstone et al. 2017), maximum behind-the-meter battery capacity was estimated for each cluster as a part of LBNL-Load (appendix section C-6.4 of the Phase 2 report). This step has been removed from the LBNL-Load model, and batteries are now fully handled in the DR-Path model (see appendix section C.2.9 of this report). Unchanged aspects of the model include (but are not limited to): investor-owned utility (IOU) demographic and hourly load data inputs, clustering methodology and resulting clusters, temperature dependence modeling and disaggregation, and all end-use disaggregation methods except for residential water heating.

B.1. New IEPR Forecast Data

The net load forecast in the *2025 CA DR Potential Study* was based on the California Energy Commission's (CEC's) 2015 Integrated Energy Policy Report (IEPR). This forecast data source was updated using the CEC's 2017 IEPR. We used the "Mid Demand Case" with Middle Savings Scenario "MidAAEE" (adjustments for additional achievable energy efficiency [EE]). This update allowed us to extend our forecast from Year 2025 (in Phase 2) to Year 2030 plus using more recent forecast for year 2020 and 2025. Year 2014 remains the baseline year after the update.

The load forecasting process consists of two large steps: (1) preprocessing raw IEPR data and (2) forecasting load.

   (1) Regarding raw data preprocessing, there are a few changes in the IEPR forecast data content and format that are worth noting.

        Economic and Demographic Assumptions: The predictor of industrial sector load growth by the CEC was "Manufacturing Output in 2009$" in the 2015 IEPR. It was replaced with "Total Non-Agricultural Employment" numbers in the 2017 IEPR. This led to small differences in historical and forecasted loads through Year 2025.

        IOU Demand versus IOU Planning Area Demand: Adjustment factors have been used to address the fact that the customer load data used in our studies represent only the California IOUs and their associated community choice aggregation (CCA) customers, but not the other load serving entities (LSEs) in the IOU planning areas, such as the municipal utilities. In the last few years, CCAs have grown significantly in number and in total load served, reducing the loads served directly by the IOUs. This shift from IOU directly served loads to CCA served loads did little to change the adjustment factors and only led to small differences in historical and forecasted loads through 2025.

Energy Efficiency Impact by End Use: In the 2017 IEPR, the data for additional achievable energy savings (AAEE) (in GWh) uses a different format and contains more end uses for each sector compared to the 2015 IEPR data. It led to minor changes to forecast numbers through 2025. Each sector has the following end uses in the AAEE forecasts:

- Residential: HVAC, whole-building, and other
- Commercial: HVAC, lighting, whole-building, and other
- Industrial: process, agricultural pumping, data center, and other
- Other: no breakdown

(2) Regarding load forecasting, LBNL-Load computes the baseline scenario using 2014 and a number of forecast scenarios, including three forecast years (2020, 2025, and 2030), two weather scenarios (1-in-2 and 1-in-10), and with or without AAEE impact.

The 8,760-hour time-series of disaggregated load profile for each cluster in baseline year 2014 was the basis for computing the hourly load profiles in the forecast years. This was done taking into consideration the **load growth** and **AAEE savings** in each IOU sector.
Load growth ratio for an IOU sector is calculated as the ratio of sector's total MWh multiplied by the IOU's adjustment factor in the forecast year divided by that in 2014. This growth ratio is then applied to all end uses in the same IOU sector's forecasted hourly load profile. The **IOU adjustment factors** are calculated by dividing an IOU's annual total load by the IOU planning area's annual total load in each forecast year.
AAEE impact is calculated as MWh savings and then converted into savings percentages by dividing by the sector's total MWh load for each forecast year, IOU, and end use. The "AAEE savings percentages" are adjusted by the ratio of "non-other" loads in each sector's total loads. The 8,760-hour load profiles of each end use are then multiplied by a factor reflecting the AAEE savings percentage.

Application of the IEPR forecast generates 8,760-hour cluster- and end use-level load profiles for each forecast year. A summary of the forecast results is shown in Table B-1, demonstrating how the total customer count and annual energy use of each utility and sector is forecasted to change from 2020–2030 for a single scenario.

Table B-1. Summary of forecasted customer counts and annual energy consumption for three forecast years by utility and sector. Values shown are for the Mid-AAEE, 1-in-2 weather scenario after forecasting but before additional electrification. Values may not sum to totals shown due to rounding.

Utility	Sector	Customer Count			Annual Energy Use (GWh)		
		2020	**2030**	**% Change**	**2020**	**2030**	**% Change**
PG&E	Commercial	561,700	632,300	12.6	38,400	44,300	15.4
	Industrial	165,300	171,400	3.7	25,200	26,400	4.8
	Other	439,500	439,500	0.0	10,000	10,300	3.0
	Residential	5,123,800	5,659,700	10.5	33,900	42,000	23.9
	Total	6,290,400	6,903,000	9.7	107,500	123,000	14.4%
SCE	Commercial	569,000	650,100	14.3	37,800	42,700	13.0
	Industrial	105,900	112,800	6.5	25,600	27,200	6.3
	Other	257,900	257,900	0.0	8,700	9,000	3.4
	Residential	4,844,400	5,289,600	9.2	34,200	42,900	25.4
	Total	5,777,200	6,310,400	9.2	106,200	121,800	14.7
SDG&E	Commercial	141,100	163,200	15.7	10,000	11,400	14.0
	Industrial	18,500	18,600	0.5	1,700	1,800	5.9
	Residential	1,340,800	1,452,600	8.3	7,700	9,500	23.4
	Total	1,500,400	1,634,500	8.9	19,400	22,700	17.0
Total		13,568,000	14,847,900	9.4	233,100	267,500	14.8

B.2. Modeling Electrification Load Shapes

B.2.1. Electric Vehicles

A 2018 CEC report estimated statewide EV charging station needs for 2017—2025, using the Electric Vehicle Infrastructure Projection (EVI-Pro) tool developed by National Renewable Energy Lab (NREL) (Bedir et al. 2018; NREL 2019). Data were obtained from NREL for the 24-hour charging profiles in 10-minute increments by county (58 in California) and day type (weekday and weekend). Separate profiles are given for three location types: residential, workplace, and public charging. All workplace and public chargers are Level 2 (240 V), while residential chargers can be either Level 1 (120 V) or Level 2. Example load profiles for two counties (Bedir et al. 2018) are shown in Figure B-1.
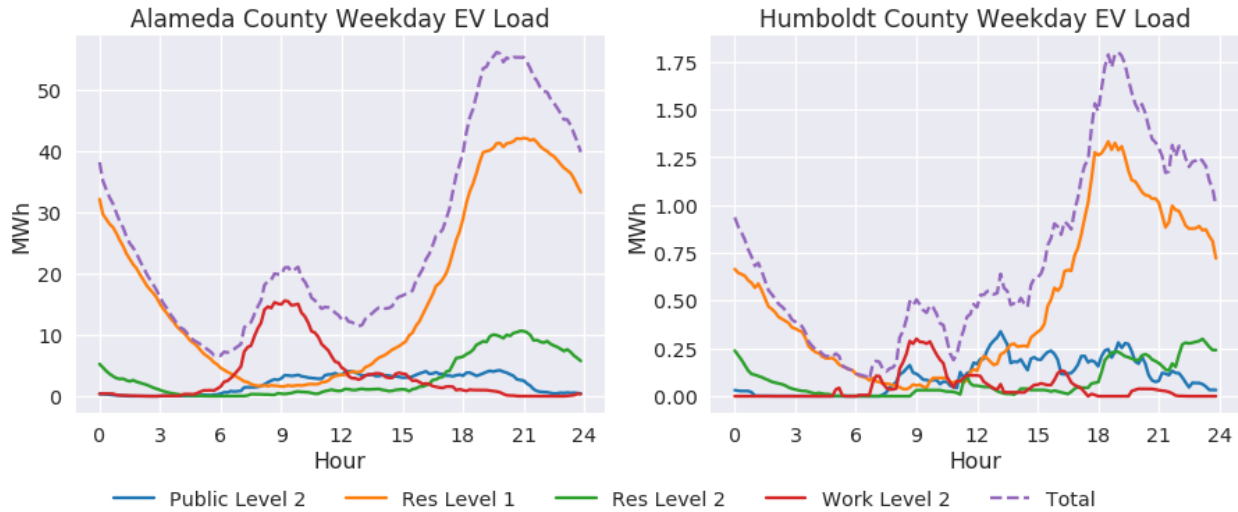
Figure B-1. Weekday EV load profiles for (a) Alameda County and (b) Humboldt County for four charging types.

At the time of modeling EV loads for the current phase of the DR Potential study, the EVI-Pro data were the most recent EV load shape estimates available via CEC analysis. Since then, additional work by ADM Energy Research and Evaluation for the CEC Demand Analysis Office has generated EV load profiles based on metered data as opposed to travel surveys, with special consideration given to the impacts of EV-specific time-of-use (TOU) electricity rates (Baroiant et al. 2019). A comparison of load shapes between these two sources is shown in Figure B-2 for reference; the primary difference for residential EVs is that EVI-Pro assumes owners plug in their vehicles as soon as they arrive home in the evening, whereas in the newer load shapes charging is deferred until lower TOU rates kick in during the overnight hours. For commercial (i.e., workplace) loads, the EVI-Pro load shapes similarly show a peak when individuals arrive to work, while the ADM load shapes show a more distributed demand throughout the morning and early afternoon. While the DR-Path model does apply TOU adjustments to the end uses in this study, these are based on the expected typical TOU rate structures and not EV-specific tariffs. This study would therefore overestimate the amount of load available to be shifted from the evening into the late-night hours, since a portion of the load will already have been shifted by the presence of EV-specific TOU rates. In essence, the Shift potential for EV charging presented in this study can be taken to include the resource that would be captured by EV-specific TOU rates. The DR-Potential team will explore incorporation of EV-specific TOU load impacts and TOU-based EV data in future phases of this work.

Figure B-2. Comparison of EV load profiles between the EVI-Pro model and the recent CEC load shapes report.

To add EV charging loads to the LBNL-Load output load shapes, the load profiles from EVI-Pro are aggregated to an hourly profile, concatenated into an annual profile, and then scaled up based on the total EV count assumed in the state for the given year. The CEC report on charging infrastructure (and therefore, the NREL data) assumed approximately 1.3 million electric vehicles in 2025, based the Zero Emission Vehicle (ZEV) target in California Executive Order B-16-2012 and assumptions regarding what portion of ZEV would be fuel cells (Bedir et al. 2018). For this study, we use the mid-case assumptions from the CEC Transportation Energy Demand Forecast (Bahrenian et al. 2017), which estimate approximately 2.2 million vehicles in 2025, as shown in Figure B-3. This estimate is lower than the CEC mid-case forecast from 2014, which was used in Phases 1 and 2 of this work. Since our study only models California's three IOUs, we then scale this statewide EV adoption down based on the total portion of statewide electricity used in the IOUs, which is approximately 60 percent (Garcia and Kavalec 2016). This was not done in the first two phases of this study, and therefore the current estimate is further lowered relative to those reports.

Figure B-3. On-road zero-emission vehicle stock forecasts. The green line shows the CEC Mid case from the 2014 Demand Forecast (Kavalec 2014), which was used in Phases 1 and 2 of this work. The current work (orange line) uses the CEC Mid case (solid blue line) from the 2017 Demand Forecast (Bahrenian et al. 2017) scaled down to represent only the portion of vehicles in IOU territory.

The result is an annual hourly county-level IOU EV demand profile for the four charging types, for which the aggregated statewide load is shown in Figure B-4. This demand is then disaggregated across relevant clusters: residential clusters for home level 1 and level 2 charging, commercial office clusters for workplace charging, and commercial retail clusters for public charging. There are many clusters in any given county, and as the geography used for clustering is utility sub-load-aggregation-point (subLAP), a given cluster may also have customers across many counties. The following process is therefore used to distribute the demand from counties to clusters. Total demand in each cluster-county combination is calculated by multiplying the cluster's average customer load by the number of customers in that cluster in each county, as determined from ZIP-code data for each customer. Then, county-level EV demand is distributed across the relevant clusters according to the portion of total relevant county demand in each cluster. For example, as shown for the example in Table B-2, cluster Res1 would receive 50 percent of the home level 1 and level 2 charging demand for Alameda County, plus 5 percent of the demand from Contra Costa County, while Res2 would receive 20 percent of the Alameda County demand. The exercise is repeated for all cluster-county combinations for the relevant demand for each charging type, resulting in cluster-specific EV demand profiles which are appended to the LBNL-Load output data files. The model also calculates and stores the number of EVs assumed to be in each residential cluster to enable the DR-PATH model to calculate costs that are based on charger count. Residential cluster EV count is simply the total number of EVs assumed scaled by the fraction of total residential EV load that occurs in the cluster.

Figure B-4. 2025 Statewide EV demand from LBNL-Load.

Table B-2. Example electricity demand (MWh) and assignment of county-level EV loads to LBNL-Load clusters.

Cluster	Demand in Alameda County	Demand in Contra Costa County	...	Total Cluster Demand	EV Demand Assigned
Res1	500	100	...	600	50% of Alameda County, 5% of Contra Costa County
Res2	200	0	...	200	20% of Alameda County
…	...	...	...	...	...
Total County Residential Demand	1000	2000	...		

B.2.2. Electrification of Residential Space and Water Heating

This study analyzes the demand response potential from the electrification of residential water heating and space heating in the context of California's greenhouse gas (GHG) reduction targets. California State Bill 32 (SB 32), passed in 2006, requires the state's annual GHG emissions in 2030 to be 40 percent below the 1990 level. Residential and commercial buildings accounted for 30 percent of the total natural gas demand from on-site natural gas combustion and 11 percent of the overall state's annual GHG emissions in 2016 (CEC 2017b). Leveraging the declining carbon intensity of grid electricity (aided by California SB 100 targeting zero-carbon electricity generation by 2045, and California Executive Order B-55-18 targeting full carbon neutrality by 2045), electrifying building heating end uses can help decarbonize the building sector. With many more commercial building types and energy end uses, and with sparse available data, this study does not attempt to analyze electrifying commercial heating in this phase, leaving this topic to future work.

The additional electricity demand from future installations of electric residential space and water heaters is needed to compute the demand response potential of these end uses. In LBNL-Load, we model this demand using: (1) expected total annual electricity demand by typical electric water heaters (considering two types, electric resistance and heat pump) and space heaters (air source heat pumps) for each of California's 16 building climate zones (CEC 2017a), (2) hypothetical hourly profile of electricity demand for each end use for each climate zone, and (3) a model for the potential penetration of electric heating appliances into the installed stock, gradually replacing the current stock of predominantly gas-fired appliances.

B.2.2.1. Residential Water Heating Assumptions

Residential electric water heaters (WHs) are typically one of two types: electric resistance (ERWH) or heat pump (HPWH). HPWHs are typically three to four times as efficient as ERWH and thus can reduce the anticipated additional electricity demand from electrifying water heating. Year-long hourly load profiles for HPWH were obtained from a simulation model, HPWHSim (Carew et al. 2018). The simulation can be configured to the design parameters of any specific commercial water heating technology. With annual hourly profiles of hot water draws, inlet water temperatures, and surrounding air temperatures as inputs, the simulation model produces hourly load electricity demand profiles for the entire year. The annual energy demand for hot water varies by climate zone. For each of California's 16 building climate zones we chose the hot water demand of a typical three-bedroom house with a 50-gallon storage tank to simulate the hourly load profile of a HPWH and ERWH.

Currently, 90 percent of the residential water heaters in California are natural gas-based, 6 percent are ERWH, and the rest are propane-based. The share of HPWH is negligible today. For the annual change in market share from ERWH and HPWH we used one of the decarbonizing scenarios from a study (Raghavan, Wei, and Kammen 2017) that develops scenarios to decarbonize residential water heating in California based on a stock turnover model. The authors assume water heaters are replaced only on their retirement. We used one of the scenarios where electrification or replacement of retiring natural gas and propane-based water heaters to HPWH is gradually phased in beginning in 2020, with 60 percent, 90 percent, and 100 percent of the retiring NG water heaters being replaced by HPWH in the years 2030, 2040, and 2050, respectively. Additionally, all retiring ERWH are assumed to be replaced by HPWH. In this scenario, the 2050 GHG emissions fall to 80 percent below 1990. Figure B-5 captures the resulting shift in the annual market share between water heating technologies assumed over the next decade for the state.

Figure B-5. Market share of residential water heaters under the chosen decarbonizing scenario.

B.2.2.2. Residential Space Heating Assumptions

For electrification of space heating, high efficiency air-source heat pump (ASHP) space conditioners are assumed to gradually replace natural gas-based space heating. ASHPs can heat and cool buildings, and thus can replace both a natural gas-based furnace and an air conditioner (AC) in a home. Annual demand and hourly load profiles of space heating a single-family home in each of the 16 climate zones was simulated using EnergyPlus modelling software for a CEC-sponsored study (Tarroja et al. 2018).

As for ASHP adoption rates, we use a study (Raghavan et al. 2020) that develops technically feasible decarbonization scenarios to transition away from natural gas-based heating and electric AC for cooling. The scenario utilized in this study is such that, for existing homes: (1) Beginning in 2020, a share of retiring central ACs are replaced with ASHP rising to a 75 percent replacement rate by 2030, with the existing natural gas heating being displaced by ASHP heating in these homes; and (2) beginning in 2020, a smaller share of retiring natural gas furnaces (in homes without AC) are replaced by ASHP, so that by 2030 50 percent of all the retiring natural gas furnaces are replaced by ASHP. For new homes built between 2020 and 2025, the scenario states that only 20 percent use ASHP for heating, while 80 percent will adopt natural gas furnaces. However, all homes built from 2025 onwards are assumed to adopt ASHP. Under this scenario, the reduction in GHG emissions in 2030 and 2050 is estimated to be about 35 percent and more than 90 percent by 2030 and 2050, respectively. The resulting assumed statewide market share of each technology is shown in Figure B-6.

Figure B-6. Share of homes by installed heating technology under the chosen decarbonizing scenario.

B.2.2.3. Modeling of Residential Heating Electrification in LBNL-Load

Water heaters are incorporated into the model by first disaggregating them from the 2014 baseline load shapes according to the penetration and load shape assumptions described above. This is done in the Disaggregation Algorithm in LBNL-Load (see Figure 2-1) alongside other end uses such as pool pumps. Electric space heaters are assumed to have negligible load in 2014, and therefore are not disaggregated from the baseline load shapes.

Space and water heating load shapes are updated in the Forecasting Algorithm (see Figure 2-1) after all other end uses are forecasted to the appropriate years and scenarios. For space heating, the key input is hourly energy use for an average-sized home in each climate zone. The climate zone of each customer in a given cluster (based on ZIP code in the demographic file) is used to create a weighted average load for each cluster. This load is then scaled according to the cooling load of the cluster relative to other clusters in their subLAP; this is done to adjust the load from the "average-sized" home to one that is closer to the assumed size (and therefore cooling energy use) of the cluster. Finally, the load shape is multiplied by the total number of customers in the cluster and the assumed penetration for the given scenario and year. Water heating load is forecasted similarly, using a climate zone weighted average load shape for each cluster. This load shape is then multiplied by annual energy use for the given water heater type, the total number of customers in the cluster, and the assumed penetration for the given scenario, year, and water heater type.

B.2.3. Summary LBNL-Load results for Phase 3

LBNL-Load produces cluster- and end use-level load shapes for all forecast years and scenarios. All figures in this section are for the Mid-AAEE, 1-in-2 weather scenario with residential space and water heating electrification. They show various loads for the average day in January in July, in 2025 and 2030. Statewide residential load is shown in Figure B-7, where in January there are two peaks: a morning one caused by heating electrification and a steep evening ramp from plug loads and other end uses. In July,

residential loads have a single evening peak, which is significantly higher than the January peak due to cooling load. Electrification loads from the residential sector, including space and water heating and EV charging, are highlighted in Figure B-8. Heating is the most prominent of these loads in the winter, with a morning peak followed by decrease during the day and a ramping up throughout the evening and night. Residential EV load peaks in the evening and is consistent year-round due to the model methodology using a singular load shape for both weekday and weekend. Water heating energy use is slightly higher in the winter, particularly in the evening and at night when outside temperatures drop, while in the summer the peak is in the late morning due to morning water consumption.

Figure B-9 shows the statewide commercial load, with a mid-afternoon peak that is similarly larger in July due to increased HVAC demand, while Figure B-10 shows statewide industrial load. The total statewide demand forecast is shown in Figure B-11, disaggregated by sector. In developing these load shapes, sector assignments were based on customer tariffs. A small number of customers' tariff information did not clearly specify a sector; these were assigned to the category "other." In July, the peak demand occurs in the late afternoon and is approximately 40 GWh and 45 GWh on average in 2025 and 2030, respectively. The average day in January has a more moderate peak around 30 GWh and 34 GWh in 2025 and 2030, respectively.

Figure B-7. 2025 and 2030 average statewide daily residential load, for two representative months, as forecasted by LBNL-Load, disaggregated by end use and including electrification of space and water heating.

Statewide Residential Electrification Load - Average Day in a Month
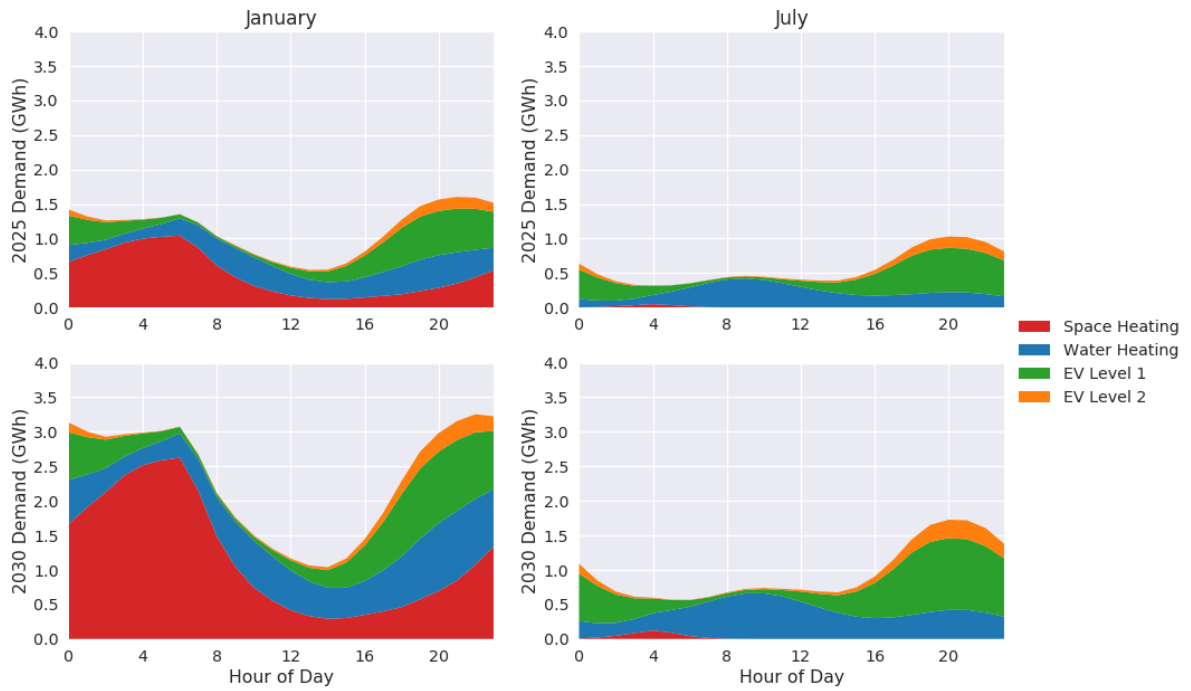
Figure B-8. 2025 and 2030 average daily loads from newly electrified end uses in the residential sector, for two representative months, as forecasted by LBNL-Load.

Statewide Commercial Load - Average Day in a Month
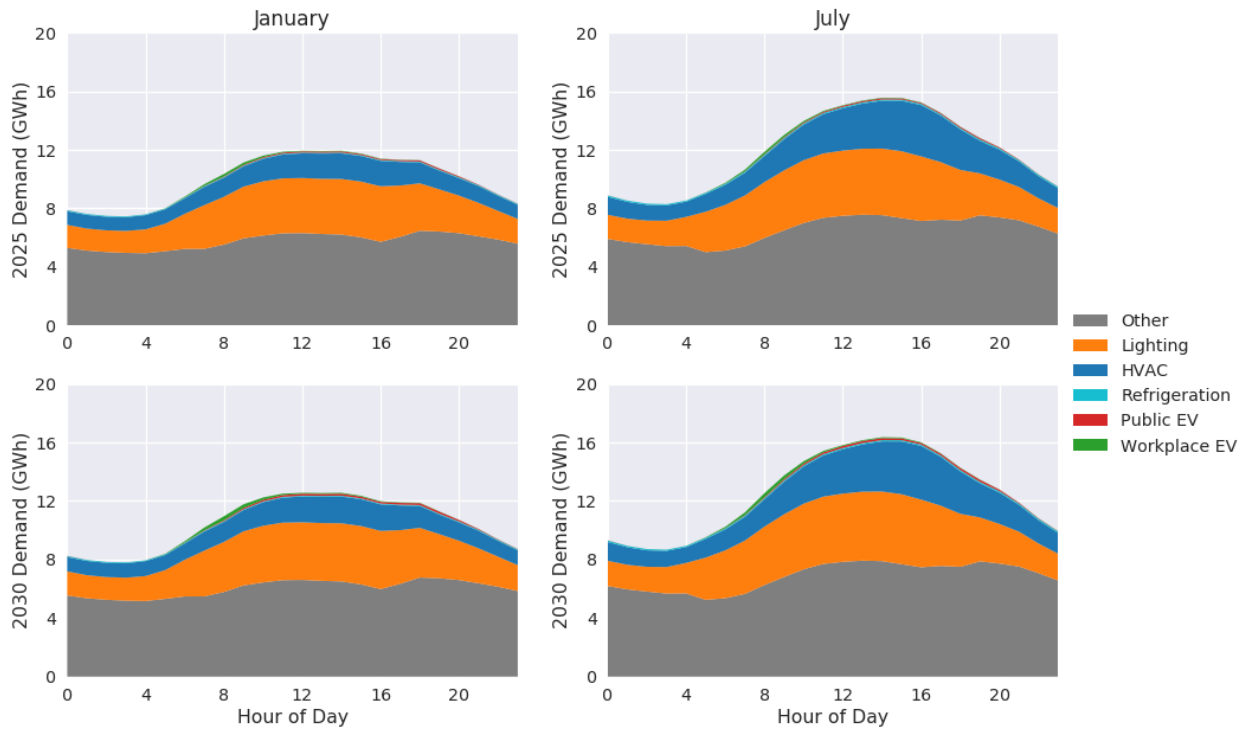
Figure B-9. 2025 and 2030 average statewide daily commercial load, for two representative months, as forecasted by LBNL-Load, disaggregated by end use.

Figure B-10. 2025 and 2030 statewide average daily industrial load, for two representative months, as forecasted by LBNL-Load, disaggregated by end use.

Statewide Total Load - Average Day in a Month
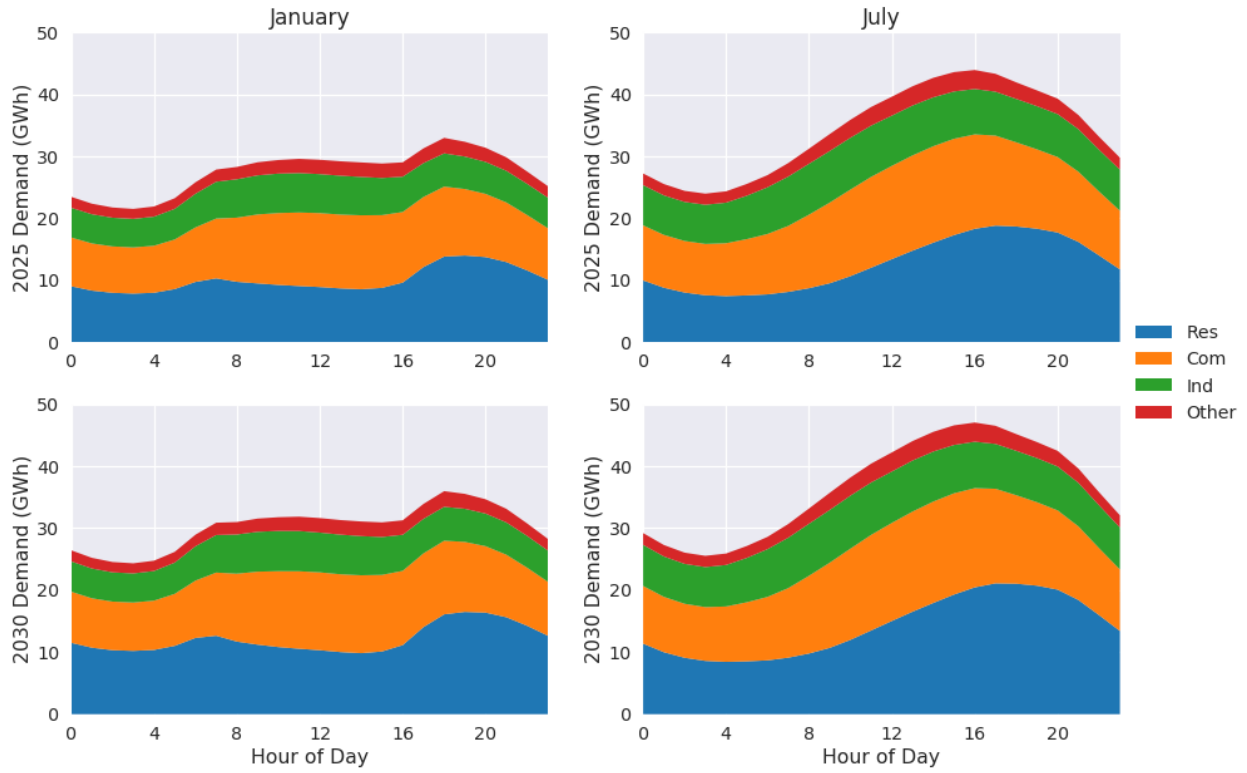
Figure B-11. 2025 and 2030 average total daily statewide load, for two representative months, as forecasted by LBNL-Load, disaggregated by sector. A small number of customers whose sector could not be clearly assigned were assigned to the category "other."

Appendix C. DR-Path: Detailed Methodology

This section describes LBNL's DR-Path model for estimating the quantity of demand response (DR) that is available within the California Independent System Operator (CAISO) territory for a given cost of procurement. In the following, we first provide a conceptual overview of the model, its inputs and functioning, making reference to detailed explanations provided in the previous phases of the DR Potential Study. We then provide detailed descriptions of key updates to the model that were implemented in the course of the Phase 3 work. Because Phase 3 focuses on Shift demand response, throughout this section we primarily describe the components of DR-Path that are relevant to calculating the Shift resource; there is additional modeling functionality for computing other DR resources that are not described in detail here.

C.1. Overview of the DR-Path model

The DR-Path model is described in detail in Appendix G of the Phase 2 report. Here we describe the functioning of the model at a conceptual level, and we detail updates made to the model for Phase 3.

In brief, DR-Path combines the cluster end-use load shape forecasts from LBNL-Load with a database of DR measures, and their associated costs, to estimate the quantity of DR that can be delivered for a particular cost, yielding a supply curve for the DR resource that can be disaggregated in various ways down to the level of the individual cluster.

C.1.1. Calculation layers

Figure C-1 presents a schematic representation of the various steps in the DR-Path model. Starting with the load shapes from LBNL-Load, the DR-Path calculation proceeds through a series of calculation layers that correspond roughly to the columns shown in the figure.
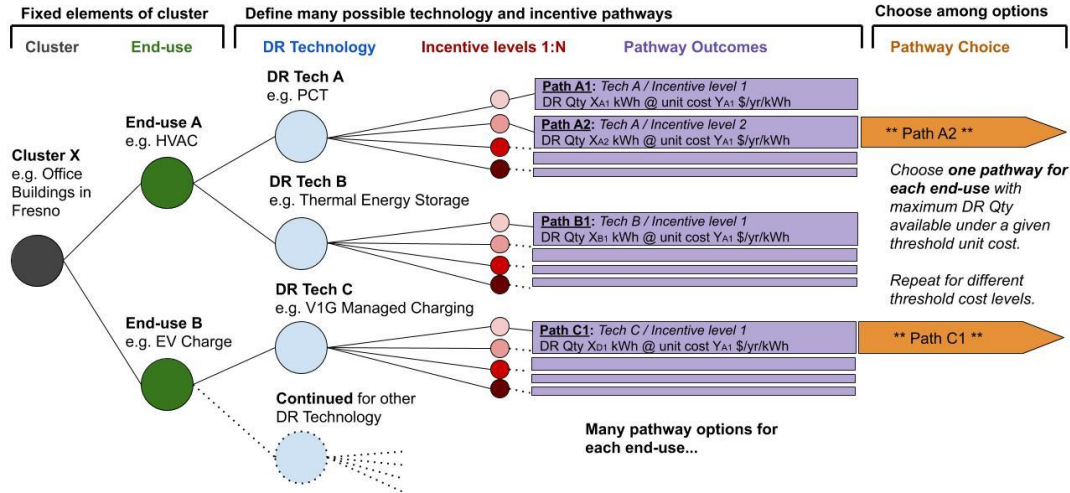
Figure C-1. A schematic diagram of the DR-Path modeling framework.

The load-reshaping layer. First, it is important to capture the load-reshaping effects of any changes in time-of-use (TOU) electricity tariffs that have occurred since the LBNL-Load baseline year, or that are expected to occur in the future.[51] In the current work, the TOU impacts for residential customers were derived from the California Energy Commission (CEC) Integrated Energy Policy Report (IEPR) forecasts, as discussed in section C.2.8 below. For the commercial and industrial sectors, load shapes were modified to account for expected shifting of the existing TOU rate windows, as described in Appendix E of the Phase 2 report.

The feature-calculation layer. Next, for each cluster end-use load shape, DR-Path computes a set of *DR features,* which are single-value representations of the maximum potential DR resource that can be provided by the load shape, for a given DR type, on average over a particular period of time (e.g., a year or season). The updated methodology for computing Shift DR features is described in section C.2.2 of this appendix; the methodologies for Shed and Shimmy feature calculation are unchanged and can be found in Appendix G of the Phase 2 report, sections G-4.5 and G12.2, respectively. This layer corresponds to the leftmost two columns of Figure C-1: for each cluster and end use, we compute a fixed value corresponding to the maximum DR resource it can deliver in principle.

The site-installation layer. The model then considers pairing each DR feature with different technologies that can capture a portion of the maximum potential resource represented by the feature. Corresponding to the DR Technology column in Figure C-1, this layer envisions installing different DR-enabling technologies at a typical site within each cluster, computing the DR resource that could be captured with each technology (i.e., the average kWh of shiftable energy available at times of system need), as well as the total annualized cost of installation, financing, and operation for the technology as installed at the site ($/yr). These can be combined to yield the cost per unit resource (e.g., $/yr/kWh of Shift) to enable DR

---

[51] Although accounting for TOU impacts is nominally a part of the loadshape forecasting task, this step was historically included in the DR-Path module to allow multiple different TOU rate scenarios to be explored, as in the Phase 2 study. In this phase, we consider only a single scenario for TOU rates.

from a particular end use within a particular cluster using a particular technology. The result of this layer is a large set of potential site installations that could be used to enable loads to participate in DR. The detailed site-level cost calculations are described in appendix section G-5.1, of the Phase 2 report. The technology cost and performance inputs are presented in appendix section G-9, of the Phase 2 report, with updates and new inputs in Appendix D of this report.

The incentive/adoption layer. The next step is to calculate, for each potential site installation, the fraction of sites that would choose to adopt the technology and participate in a DR program for a given financial incentive. This step, which corresponds to the Incentive Levels column of Figure C-1, pairs each site installation with an array of different incentive levels, expressed as an annual payment per unit of DR resource provided (e.g., $/kW-yr of Shed). The incentive level, along with certain attributes of the cluster members and DR technology, are used as inputs to a propensity score model that computes the fraction of customers in the cluster that would be expected to enroll in a DR program for a given site installation and incentive level. Importantly, the propensity score model estimates only the customer decision to enroll in a DR program, not the decision to participate in any specific DR event. Thus, the model is estimating the potential resource for a typical event, not the resource that would be realized in practice. The propensity score model is described in detail in Appendix F of the Phase 2 report. The site-level resource is combined with the participation rate to yield the total potential DR resource available from the cluster at each incentive level. The incentive payments, along with marketing and administration costs for the DR program, are then added to the site enablement costs from the previous layer to generate the total cost per unit of DR enabled along each particular pathway from cluster to incentive level (the Pathway Outcomes column of Figure C-1).

The pathway selection layer. The pathway calculations in the previous layers correspond to a large number of potential future pathways to capturing DR resources from different sites and end uses. To estimate the total quantity of DR that is available for a given procurement price, the model selects, for each site and end use, the pathway[52] that provides the maximum quantity of DR, from among all pathways whose cost is at or below the price being considered (the Pathway Choice column in Figure C-1). The model outputs can then be used to build a DR supply curve by repeating this selection procedure for an array of different procurement prices. Because each quantity on the supply curve is an aggregation across many different cluster and end-use pathways, it is straightforward to disaggregate the supply curve by end use, building type, region, or other cluster attributes.

C.1.2. Measuring the DR quantity and cost

The feature calculation layer of DR-Path calculates the maximum DR resource that is available from each cluster and end use for each different DR type (Shed, Shift, or Shimmy). The site-installation layer then computes the actual resource that can be accessed by installing a given enabling technology. The resulting quantity of DR represents the response (kW of Shed or Shimmy, or kWh of Shift) that the system could be expected to provide, on average over the course of a year, weighted by the likelihood of dispatch for

---

[52] In the case of ties, all pathways that maximize the DR resource are selected and given equal weight. For instance, if there are three pathways that provide the same quantity of DR at or below a given procurement price, each pathway is selected with a weight of one-third.

each type of DR in each hour of the year. In the case of Shed, this quantity is analogous to a *capacity* resource: it represents a quantity of proxy generation capacity. Thus, we measure the Shed resource using the same units as used to measure capacity in power systems, namely GW, with the dual caveats that (1) like GW of generation capacity this unit (as with the nameplate capacity of a generation unit) represents the maximum power-reduction capability of the DR resource, not the instantaneous load reduction delivered at any given time; and (2) the Shed quantity is reported as an *average* value available at times that the power system is likely to need Shed, and it will vary with time depending on customers' particular energy usage patterns. Similarly, the Shimmy resource represents capacity for the ancillary services market, so it can be measured in the same units and comes with the same caveats.

The situation is slightly different for Shift, which represents a quantity of *energy* that can potentially be called upon to shift by a few hours temporally, whenever there is a need (provided that it is not called on repeatedly in quick succession). In this sense, whereas Shed and Shimmy represent virtual *generation capacity* resources, Shift represents a virtual form of *energy storage capacity*. The amount of battery or other storage on the grid is typically measured in GWh, so we use the same unit to quantify the Shift resource in this report.

To compute the costs of procuring DR as a grid resource, we must include one-time up-front costs for enabling technology and installation, as well as ongoing costs for operations and maintenance, program administration and marketing, and customer participation incentives. In this situation, the most straightforward approach is to amortize the up-front costs over the lifetime of the equipment being installed, and report the total annualized cost ($/yr) of procuring a unit of DR resource (kWh of Shift or kW of Shed or Shimmy). Then the total cost metric becomes $/yr/kWh for Shift, or $/yr/kW for Shed or Shimmy. These units can also equivalently be written as $/kWh-yr or $/kW-yr, and the latter unit is commonly used to report capacity and Shed DR procurement costs in the utility industry.

Importantly, the units used here differ somewhat from the units that were used to express DR resources in the Phase 2 report. The Phase 2 report expressed costs as $/kW-yr and $/kWh-yr and generalized the denominators of those units to represent the DR quantity procured for a given marginal expenditure in units of "GW-yr" or "GWh-yr." These units were taken to represent the quantity of demand or energy consumption that could be secured as a DR resource for a period of a year—using this set of units, one could make the statement that, at a marginal cost of $100/kWh-yr, an expenditure of $100,000 in a given year would secure 0.01 "GWh-yr" of Shift. This set of units has a significant drawback, however, in that it contains two different units of time (hours and years) adjacent to each other, which was a persistent source of confusion for stakeholders. It was also inconsistent with typical industry practice for Shed DR, which reports annualized costs in $/GW-yr but resources in GW (not GW-yr). In this report, then, we use a more straightforward set of units for DR quantity—namely GW and GWh—which modifies the example statement above to say that a Shift resource of 0.01 GWh can be maintained at an annual expenditure of $100,000, for as long as that expenditure continues. We also choose to write our cost units for Shift as $/yr/kWh, instead of $/kWh-yr, to avoid confusion about time units and emphasize that the reported costs represent an annualized cost ($/yr) to procure a quantity of shiftable energy (kWh). Crucially, these two different sets of units are describing exactly the same quantities; they simply differ in whether they explicitly include the procurement period when expressing the quantity.

C.1.3. Summary of cost calculations

The site-installation and incentive/adoption layers of DR-Path generate estimates of the total cost to enable each site and end use to participate in DR, in units of annualized cost per unit of DR resource. Appendix section G-5 of the Phase 2 report presents the DR-Path cost model in full detail. The approach to computing DR resource costs was unchanged in Phase 3. Here we briefly summarize the components of the DR-cost model and walk through a set of simple examples to orient the reader.

The DR cost calculations in DR-Path include the total, "all-in" cost that is required to enable a certain amount of load to participate in DR, regardless of who pays (be it the customer, the aggregator, or the utility), under the assumption that all costs to enable DR for grid services will ultimately be borne by society, either through direct customer costs or through long-term changes in electricity rates. Certain benefits that accrue to different market participants can then be subtracted from the total resource cost to estimate the costs from the perspective of different market actors. In the current study, we present costs from the perspective of the utility or program administrator, subtracting off the value of any co-benefits from installing specific DR-enabling technologies at a given customer site (which are assumed to be borne by the customer).

The DR-Path cost calculations are based on a number of different input costs that fall into the following categories.

- **Technology costs** represent the costs of installing and operating DR-enabling technology at a customer site. These costs are made up of several different subcomponents:

  - **Fixed initial costs** represent equipment and installation costs that have a fixed value per customer site (e.g., the cost of installing a communicating thermostat in a residence).
  - **Variable initial costs** represent equipment and installation costs that vary with the size[53] of the load being enabled for DR (e.g., the cost of installing a number of smart EV chargers at an office building to meet demand).
  - **Financing costs** represent the total cost of financing the initial costs over the lifetime of the investment (this can, alternatively, be thought of as the capital recovery factor to convert the initial costs into an annualized value).
  - **Fixed operating costs** represent annual costs of operating the technology that have a fixed value per customer site (e.g., annual subscription costs for residential smart EV charging software).
  - **Variable operating costs** represent the annual costs of operating the technology that vary with the size of the load being enabled for DR (e.g., software subscription costs for commercial smart EV charging, levied on a per-charger basis).

---

[53] The "size" of the load is measured by the peak annual demand from a given end use. This is the maximum amount of demand response that the end use could deliver, in principle, if called during the peak hour; it can be thought of as analogous to the nameplate capacity of a generation unit. The quantity of DR actually procured for a given cost is weighted to account for variation in the load at times of system need, as discussed in section C.2.2; this is analogous to applying a capacity factor to derate the nameplate capacity.

- **Program costs** represent the costs to the utility or aggregator of running a DR program. These fall into two subcategories:

  - **Marketing costs** represent the program outlays required to attract and enroll customers, on a per-enrolled-customer basis.
  - **Administration costs** represent the costs of program administration per enrolled customer.

- **Incentive costs** represent the financial incentives that must be paid to participating customers, on a per-unit-resource basis (i.e., $/kW-yr or $/yr/kWh).
- **Site co-benefits** represent the financial value of non-DR co-benefits that accrue to the site owner or occupant. Examples include electricity bill savings from smart thermostat operation and improved space utilization or occupant satisfaction from networked lighting controls. When computing costs from the utility perspective, a technology's co-benefits (if any) are assumed to reduce the technology costs by some fraction, representing the fraction of the technology costs that can be assigned to site-level improvements instead of provision of DR, which the customer would be expected to bear.

C.1.3.1. Example cost calculations

To illustrate how the cost model produces resource costs for different sites and technologies, we consider calculating the total resource costs associated with enabling Shift DR for HVAC loads at four example sites.

In the commercial sector, we consider installing an energy management system (EMS) whose cost scales with the size of the peak load being handled, at $250/kW, with negligible fixed initial costs and negligible operating costs. We compute the cost of installing this technology at two example sites:
- A large office building with a peak HVAC load of 500 kW, but whose Shift potential is comparatively small at 300 kWh, since the building is sparsely occupied during the morning and evening Shift windows.
- A large retail building with a peak HVAC load of 600 kW, whose Shift potential is also large at 600 kWh since the building is open for business during the morning and evening Shift windows.

In the residential sector, we consider installing a programmable communicating thermostat (PCT) having a fixed initial cost of $300 per site, a fixed operating cost of $15/yr for a software subscription, and no significant variable costs. We compute the cost of installing this technology in two example homes.

- A small home with a peak cooling load of 1 kW and a Shift potential of 0.15 kWh.
- A large home with a peak cooling load of 4 kW and a Shift potential of 0.4 kWh.

In both of these example homes, the relatively small Shift potential reflects the fact that residential HVAC loads will be well below their peak values during the winter and spring months when Shift is most likely to be needed, as well as the assumption that a PCT can only shift a fraction of the residential space cooling load.

Figure C-2 shows how the different components of the cost model come into play in calculating a total Shift resource cost for each of these example sites. The top two panels show the total annualized cost breakdown for enabling each site to perform Shift and enrolling it into a DR program. The bottom panel shows how these site-level costs translate into a total Shift resource cost, including customer incentive payments. The site-level costs are several orders of magnitude higher for the commercial sites than for the residential sites, reflecting the much higher costs associated with large-scale EMS installation over a simple PCT installation. (In addition, given the much lower residential technology costs, the costs of program administration and marketing make up a much more significant fraction of the per-site costs.) When translated into a cost per kWh of Shift, however, the residential-sector costs dwarf the commercial-sector costs, owing to the much larger shiftable loads that have been enabled in the latter case. In this example, then, despite the much higher absolute costs at the commercial sites, enabling these sites represents a much more cost-effective pathway to procuring Shift.

Figure C-2. Example site-level cost calculations for enabling different sites to perform Shift (top panels) and total resource costs for the Shift obtained at each site (bottom panel).

C.2. DR-Path modeling updates for Phase 3

As part of the Phase 3 effort, the DR-Path model underwent a significant overhaul to facilitate addressing the specific research topics that were the focus of Phase 3, as well as to increase the modularity and flexibility of the model to support future research directions. Broadly speaking, the updates consisted of a major upgrade and restructuring of the DR-Path software code; a substantial redesign of the analytical approach to computing DR features for Shift; and several smaller updates related to TOU rate impacts, market revenue calculation, and the approach to modeling behind-the-meter (BTM) batteries as a DR

resource. The following sections describe the updates in detail. The updated DR-Path model will be released as open-source software alongside the publication of the present report.

C.2.1. DR-Path software upgrades

The DR-Path modeling software utilized for research in phases 1 and 2 was made up of a series of scripts in the R statistical package, which were designed specifically to address the research questions addressed in the earlier phases, and which handled data input and output via text files in the comma-separated values (CSV) format. We will refer to this R-based model as DR-Path version 1. To support Phase 3 and future work, Lawrence Berkeley National Laboratory (LBNL) developed DR-Path 2.0, a modular and thoroughly tested software package, written in the Python language, with data management and processing handled in a SQLite relational database.

C.2.1.1. Description of the modules making up DR-Path 2.0

The DR-Path 2.0 software package is written in the Python language and is structured into a hierarchical set of modules that implement different key pieces of functionality and different layers of the modeling framework. By separating the functionality into separate logical units, the modular design allows the individual technical and analytical components of the package to be independently tested, as well as enabling future upgrades and modifications to different layers to be implemented straightforwardly without the need to modify unrelated parts of the framework. The underlying modules making up DR-Path 2.0 are as follows.

- The **db** module defines the structure of the DR-Path database and includes tools for interacting with the database from within the DR-Path software, including loading inputs into the database, performing queries and aggregations, and extracting results.
- The **loadshape** module handles processing and calculations performed on the input cluster end-use load shape forecasts from LBNL-Load, culminating in the calculation of DR features for each load shape. It consists of three main submodules:
  - The **reshape** module implements the **load reshaping** layer of the DR-Path framework. It applies reshaping algorithms to the input load shapes to account for the effects of TOU rates for different TOU impact scenarios that can be defined in the input data files.
  - The **aggregate** module computes the aggregated system-level load, both *gross* demand as constructed from the reshaped cluster load shapes, and *net* demand constructed from the gross demand and an input forecast of intermittent renewable generation. (See Appendix D of this report for a description of the input renewable forecast.)
  - The **precalculated_features** module implements the **feature calculation** layer of the DR-Path framework. It contains algorithms to process the reshaped load shapes from the reshape module into single-valued features representing each load shape's technical potential to provide different types of DR. For each DR type, the feature is based on consideration of the system-level need for DR and the amount of load the cluster can deliver coincident with that need, on average over a given period.

- The **futures** module handles calculations downstream of the DR feature calculation. It builds potential future pathways to capturing the technical DR potential in each cluster and end use and selects the pathways that maximize the resource at a given price.

  - The **installations** module implements the **site-installation** layer of the DR-Path framework. It considers coupling each cluster and end use with a variety of different enabling technologies, whose cost and performance details are specified in an input file and loaded into the database. The result is a large table of potential site installations of DR-enabling technology, including site-level cost and performance information for each, as computed by the cost_perf_calculations module.

  - The **cost_perf_calculations** module handles the calculations of DR performance for each site installation. It also computes the total cost of the DR resource from each cluster, including initial, financing, and operating costs at the site level, as well as program costs for incentives, administration, and marketing.

  - The **adoption** module implements the **incentive/adoption** layer of the DR-Path framework. It envisions encouraging participation for each site installation using a range of different monetary incentive levels, and it computes the resulting participation rates based on a propensity score model. The result is a table containing a large number of incentive/adoption scenarios, including total per-unit resource costs, as computed by the cost_perf_calculations module, and cluster participation rates, as computed by the propensity score model.

  - The **pathways** module implements the **pathway selection** layer of the DR-Path framework. It selects the DR-maximizing pathways at a given procurement price and saves the results in the database. It also constructs a series of user-specifiable database queries that allow the user to aggregate the pathway selections up to yield a total DR resource available for a given price, representing a single point on the DR supply curve. The module also includes a tool to generate complete supply curves at different user-specified levels of aggregation.

- The **settings** module allows the user to specify configuration information, key assumptions, and desired outputs from the model prior to running the model. User-specifiable information includes the location of input data, the particular modeling scenarios to be run, and certain single-valued or low-dimensional inputs, such as the discount rate or the assumed transmission and distribution loss factors for each modeled utility service territory.

C.2.1.2. The DR-Path database: Key tables and calculation pathways

The DR-Path software package relies on an underlying relational database, implemented in the SQLite database management system. The final output of the DR-Path model is a database consisting of a number of interrelated tables that store key input data as well as intermediate and final results from the various calculations within the model. The relationships between the tables help to specify the connections between the various data products and enable fast queries of the output dataset to address a wide range of different research questions.

The DR-Path database tables can be subdivided into five different groups relating to different topics within the model, as follows:

- The **scenarios** group consists of a set of database tables specifying the different model scenario inputs relating to demand forecasting assumptions, TOU impacts, renewable generation, and rates of technological improvement.
- The **cluster load shapes** group consists of a set of tables containing the cluster load shape inputs from LBNL-Load, their TOU-reshaped counterparts, descriptive information regarding the customer clusters and end uses, and the DR Features corresponding to each load shape and each DR type.
- The **aggregate load shapes** group consists of a set of tables containing aggregate system-level information, including gross and net demand, renewable generation profiles, wholesale price information, hourly TOU impact factors, and the modeled likelihood of dispatch for different DR types in each hour of the year.
- The **technologies** group consists of a set of tables describing the DR-enabling technologies that are available to be paired with different clusters and end uses. The tables describe the cost and performance of specific end use technologies, control strategies, telemetry specifications, and information on the reliability and timescale with which different technologies can be dispatched to provide DR.
- The **futures** group consists of a set of tables that contain the main outputs of the model, including a table of site installations, a table of incentive/adoption scenarios, and a final table of selected DR pathways at specified procurement prices.

C.2.2. Calculating Shift DR features

The updates to DR-Path in Phase 3 included a major update to the methodology for calculating Shift DR features (which represent the technical Shift potential for a particular cluster and end use). This section gives an overview of the feature calculation methodology in general, the specific approach used for Shift in Phase 3, and the details of changes from the approach used in Phase 2 and their implications for interpreting model results.

In general, DR features in DR-Path are computed by multiplying the load available for DR, from each cluster and end use, by a system-level *DR filter* for the particular type of DR being considered, where the DR filter represents the relative probability that the given type of DR will be dispatched in any given hour. Summing the result over all hours in the period under consideration, and appropriately normalizing the filter, yields the weighted average amount of load that is available to provide that type of DR at times of system need. As shown in Equation (C-1), the calculation is:

$$F_{c,e} = \frac{\left(\sum_{j \in P} L_{c,e,j} w_j\right)}{\sum_{j \in P} w_j}$$

(C-1)

where $F_{c,e}$ is the DR feature for a particular DR service type (Shed, Shift, or Shimmy) associated with cluster c and end use e; $L_{c,e,j}$ is the load that is available to participate in that DR service type in hour j;

w_j is the DR filter for the relevant service type in hour j; and P is the set of hours making up the time period over which the feature is being calculated (typically the full calendar year).

Shed DR provides a simple example of how this formula is applied. The model assumes that Shed DR is only likely to be dispatched during the very highest peaks in net system demand, so the DR filter w_j is set to zero for all but the top 250 hours of the year in terms of net load. In the top 250 hours, the filter is taken to be $w_j = 1/r_j$, where r_j is the rank of hour j in terms of total net load (so that the annual peak hour has $r_j = 1$, the next-highest peak hour has $r_j = 2$, and so on).[54] The available load in each hour $L_{c,e,j}$, in the case of Shed, is just the value of the end use load shape in that hour. This approach can be thought of as conceptually similar to computing a capacity factor for a supply-side renewable generation resource. In this formulation, the "nameplate" capacity for a given cluster end-use would be represented by the annual peak demand for that load shape. By weighting the resource by the hours in which it is most likely to be valuable, we effectively "de-rate" the resource to better represent its value to the system.[55]

When calculating Shift features, the approach is more complex, to account for the need to shift energy from one period to another, rather than shedding instantaneously. In particular, we consider shifting energy consumption about a particular "pivot" hour that represents the transition between adjacent Shed and Take periods. Figure C-3 is a schematic representation of the feature calculation components, for the HVAC end use in a selected commercial cluster on a single summer day. The top panel shows the HVAC load for the entire cluster on June 17. The middle panel shows an idealized "dispatch curve" for Shift DR, representing the total load decrease or increase that would be needed to smooth the net demand out to a daily average value, effectively eliminating the duck curve (see appendix section C.2.5, below, for the details of this curve).

---

[54] This weighting strategy is based on modeling of DR dispatch described in the Phase 1 report (Alstone et al. 2016).
[55] This approach for estimating Shed capacity is conceptually similar to computing an effective load-carrying capacity (ELCC) for a renewable or distributed generation resource, although our probabilistic model is focused on the value of the resource from an economic-dispatch perspective, rather than a reliability perspective (i.e., we do not consider the value of Shed for mitigating generator outages). In the case of Shift, the comparison to capacity factors is not as salient, because the model assumes that Shift can have significant value, and be dispatched, on nearly every day of the year, as discussed in the next sections.
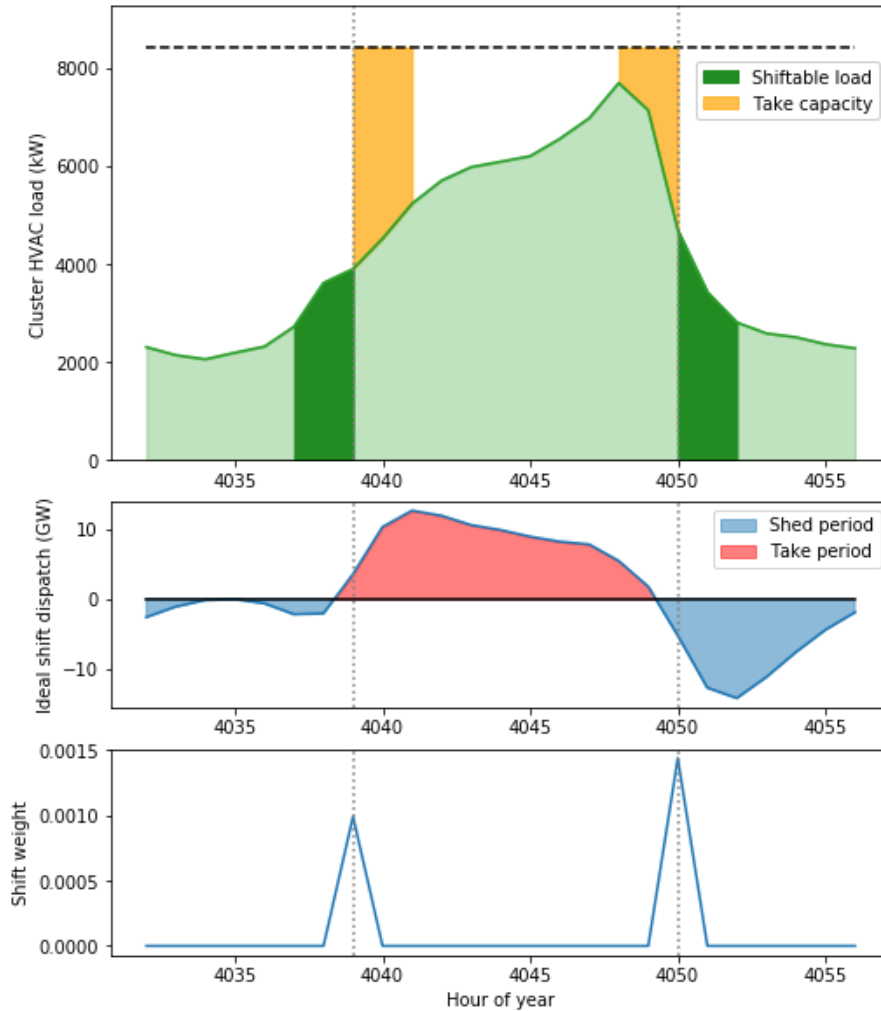
Figure C-3. A schematic diagram showing the various components of the Shift feature calculation, over the course of a single day. Top: The HVAC load for a single cluster (light green) and the annual maximum HVAC load (dashed line), with shading indicating the available shiftable load (dark green) and the corresponding capacity to take load at different times (orange). Middle: An idealized dispatch curve for Shift, representing the change in load that would smooth out the system-level net load to its local average value. The zero-crossings for this curve (dotted lines) represent the pivot hours about which Shift is valuable. Bottom: The weight assigned to each potential Shift hour, representing the relative modeled likelihood that a dispatch would occur around that hour (see appendix section C.2.5 for details).

Although this curve represents a much larger quantity of Shift DR than could realistically be procured, the *sign* of the curve determines whether a particular hour is in a Shed period or a Take period, and thus, the zero-crossings in the dispatch curve represent the pivot hours about which energy consumption should be shifted to meet system needs. In the following two sections, we describe the calculation of the dispatch curve, the available shiftable load $L_{c,e,j}$, and the Shift DR filter w_j in detail.

C.2.3. Calculating the Shift resource in different periods

Equation (C-1) for DR feature calculation involves a summation over a particular time period P. In Phase 1 and Phase 2 of this study, this was taken to be an entire year, but this is not the only period over

which DR features may be calculated. DR-Path 2.0 adds functionality to calculate DR features over an arbitrarily defined time period. This enables calculation of supply curves that represent the average available resource in a particular period of interest. Potential calculation periods could include seasons, quarters, or weekdays versus weekends. In the present study, we have calculated supply curves for the full year and for the four seasons of the year, defined for simplicity as the three-month periods of December–February (winter), March–May (spring), June–August (summer), and September–November (fall).

C.2.4. Calculating the shiftable load

The vertical dotted lines in Figure C-3 mark the location of the hours during which the dispatch curve crosses zero. To calculate Shift features, DR-Path sums the energy consumption within a specified time window during the Shed periods immediately adjacent to these hours and considers shifting it into a time window of the same size immediately opposite the pivot hour.[56] It is important, however, to consider the realistic ability of the end use in question to take more load at any given time: if an air conditioner is running at full capacity on a hot day, for instance, it will not be possible to shift load from the evening into the afternoon. DR-Path accounts for this limit by computing the maximum load for the given load shape, over the course of a full year, and assuming that this represents the maximum possible output for the cluster and end use being analyzed. Then the shiftable energy consumption in each hour can be taken as the smaller of the total consumption during the Shed window (dark green shaded region in Figure C-3) and the total capacity to take additional load during the Take window (orange shaded region in Figure C-3).

C.2.4.1. Modeling energy losses during Shift events

A minor complicating factor in computing the total shiftable load, as described above, is that, for some end uses, shifting consumption from one time to another is likely to lead to an increase in the amount of energy needed to meet the relevant energy service demand. For example, pre-cooling a building may require more energy than cooling in real time, since the internal temperature may need to be brought to a lower point than would otherwise be necessary to ensure occupant comfort throughout the shift. These increases in energy consumption during the Take phase of a given shift can be considered as equivalent to round-trip efficiency losses for behind-the-meter storage. These losses were ignored in Phase 2 of this study, which assumed all shifts to be energy neutral. The Phase 3 update to DR-Path allows users to specify round-trip loss factors for each individual modeled end use.

Including these loss factors, the available shiftable load in Equation (C-1) can be written as:

---

[56] In the Phase 3 work, we computed Shift features using windows of 4, 8, and 12 hours, centered on each pivot hour (i.e., Shed and Take windows of 2, 4, and 6 hours, respectively). Each Shift enabling technology in the site-installation layer was assumed to operate over one of these windows and was coupled with the appropriate Shift feature for each end use and cluster.

$$L_{c,e,j} = \min\left[\sum_{i \in S_j} \ell_{c,e,i}, \sum_{i \in T_j} (P_{c,e} - R_e \ell_{c,e,i})\right]$$

(C-2)

where $\ell_{c,e,i}$ is the load in hour i, $P_{c,e}$ is the annual peak load, R_e is the round-trip loss factor associated with shifting end use e, and S_j and T_j represent the Shed and Take hours within the specified window immediately surrounding the pivot hour j.

C.2.4.2. Constraints on Shift directionality

The Phase 3 update to DR-Path also considered the fact that certain end uses may have constraints on the temporal direction in which they can shift their energy consumption. For example, in the absence of shifting, EV charging will tend to commence at the moment the vehicle is connected to the charger. Such a load cannot be shifted earlier in time, since the vehicle will not have been disconnected from its charger previously. Thus, EV charging load is generally constrained to only be shifted to later hours.

To account for situations like this, DR-Path 2.0 allows users to specify certain end uses as forward-shiftable or backward-shiftable only. In this cases, Equation (C-2) is subject to an additional constraint, such that S_j must precede T_j (or vice-versa for backward-only loads), or else $L_{c,e,j}$ is set to zero.

C.2.5. The Shift DR filter: Estimating the dispatch probability for Shift

In addition to the quantity of load available for DR $L_{c,e,j}$, we must also specify the DR filter w_j in order to calculate DR features for shift using Equation (C-1). The DR filter represents the relative probability that DR of a given service type will be dispatched in any given hour of the year. In the case of Shift, DR-Path assumes that shift will only be dispatched to occur about the pivot hours where the dispatch curve crosses zero; hence, the Shift DR filter is set equal to zero in all hours that are not pivot hours. In the Phase 2 study, the Shift filter was then assumed to be constant in all pivot hours—that is, all shifts were assumed to be equally likely to be dispatched. This simple assumption was used in Phase 2, since there was little information on how Shift DR would be dispatched in real-world conditions. DR-Path 2.0 improves on this situation by adding a framework that allows various dispatch strategies for Shift to be modeled via a varying Shift filter. This facilitates exploration of how the Shift resource could vary depending on the specific grid needs being addressed. In the following, we describe the methodology for computing the idealized Shift dispatch curve and identifying the pivot hours, as well as several different approaches to developing a Shift DR filter.

C.2.5.1. Constructing idealized Shift dispatch curves

In Phase 2, the dispatch curve for Shift DR was derived from external modeling performed by Energy + Environmental Economics (E3), using their capacity expansion model RESOLVE. It was notable that those dispatch curves closely resembled inverted versions of the daily net load curve, with RESOLVE's dispatch model calling for load increase when the net load was low and load decrease when the net load

was high. In Phase 3, we developed an approach to estimating an idealized Shift dispatch curve directly from the modeled net load curve, which has the advantages both of simplicity and of making the dispatch curve calculation endogenous to the model.

To calculate the net load, DR-Path sums up all the input cluster-level load shapes from LBNL-Load to yield a gross demand curve for the CAISO system. It then subtracts off an input forecast for hourly solar and wind generation (see appendix section C.2.7 for details of the renewable forecast methodology) to yield the net load in each hour of the forecast year. To convert this into a dispatch curve for Shift DR, we begin with the assumption that an ideal (if unrealistic) Shift resource would allow the system operator to eliminate all short-timescale variation in the net load, in favor of a smooth curve that varies slowly over a timescale of a day or so, and which can be met by a relatively steady generation by non-intermittent resources, which are never called on to provide significant ramping. Thus, to estimate a Shift dispatch curve, we compare the forecasted net load curve to its 48-hour rolling average value; the difference represents our idealized dispatch curve for Shift. The choice of a 48-hour rolling average was intended to reflect timescales often considered by generators, who may base decisions about whether to be online on the previous day's operations and the demand forecast for the following day. Figure C-4 shows an example of the idealized dispatch curve derivation over the course of a week, with blue and red shading indicating the Shed and Take periods, respectively, that would arise from efforts to use load shifting to bring the net load closer to its local average value.
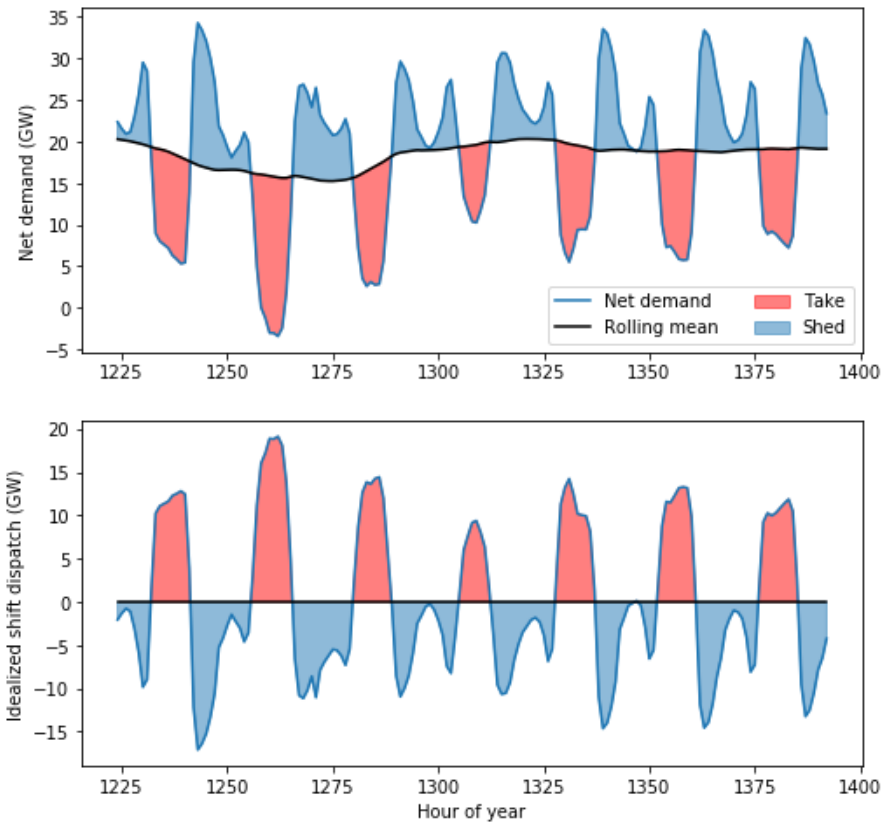
Figure C-4. An example of converting a net load curve (top) into an idealized dispatch curve (bottom). The dispatch curve is simply taken to be the difference between the net demand and its average value over the surrounding 48 hours. Red shading indicates Take periods, when the net load is lower than average, and blue shading indicates Shed periods.

C.2.5.2. Modeling the probability of Shift dispatch

As discussed previously, the hours during which the dispatch curve crosses zero are treated in DR-Path as the "pivot" hours around which load shifting will occur. That is, each zero-crossing in the demand curve represents a distinct potential Shift event within the model, and so the DR filter w_j for Shift will have non-zero weight only in these hours, with $w_j = 0$ in all other hours. As also discussed previously, the Shift filter is meant to represent the relative probability that each potential Shift event will be dispatched in practice. Accurately forecasting this probability represents a significant modeling challenge, since Shift DR remains a largely hypothetical service type, and little to no information exists on its real-world dispatch.

Owing to the lack of relevant data, the Phase 2 study did not attempt to estimate the probabilities of different potential Shift events, treating each zero-crossing in the dispatch curve as equally likely to yield a real-world Shift dispatch. This is unlikely to reflect reality in a practical Shift program, since some pivot hours occur during small fluctuations in the net load (e.g., moderate wind events during the overnight hours) and are not likely to generate a strong need for load shifting. To improve on this modeling, in Phase 3 we developed three different dispatch scenarios intended to qualitatively represent the relative

probability of Shift dispatch when Shift is used to meet different grid needs. The three scenarios were as follows:

- **Flat dispatch scenario.** This scenario is equivalent to the Phase 2 modeling approach: every zero-crossing of the dispatch curve is assumed equally likely to trigger a Shift dispatch. This scenario is used as a basis for comparison, to assess the importance of assumptions about the dispatch probability.

- **Ramping mitigation scenario.** This scenario corresponds to a Shift service that is primarily deployed to flatten steep ramps in the net demand curve. The magnitude of the Shift DR filter w_j in each pivot hour is assumed to be proportional to the size of the corresponding four-hour ramp, defined as the maximum-to-minimum range in the net load during the four-hour window centered on the end of the pivot hour.

- **Curtailment mitigation scenario.** This scenario corresponds to a Shift service that is primarily deployed to avoid curtailment of renewables. Informed by CAISO data and LBNL dispatch modeling of future California grid scenarios, we assumed that the curtailment probability in each hour depends on the load as a percentage of the peak load experienced in the surrounding 48 hours.[57] Above a certain percentage, we assumed zero probability of curtailment. Below a certain percentage (representing the minimum generation level required to have enough online capacity to meet the local peak), we assumed 100 percent probability of curtailment. Between these two thresholds, which we took to vary with time as California's generation becomes more flexible, we assumed a linearly increasing probability of curtailment from 0 percent to 100 percent.[58] In this scenario, the Shift filter w_j for each pivot hour was set to be proportional to the maximum curtailment probability that occurred in the adjacent take period.

Figure C-5 shows the Shift supply curves in the 2030 forecast year, as calculated using the DR filters from each of these three scenarios. The ramping and curtailment scenarios both show a noticeably larger potential Shift resource at each price level than the flat scenario. This occurs because the ramping and curtailment scenarios both tend to give higher weight to the daytime pivot hours (i.e., the shift opportunities around sunrise and sunset), when the shiftable loads tend to be large, whereas the flat model gives equal weight to the pivot hours that occasionally occur during overnight hours, reducing the average size of the resource. Notably, though, there is little difference between the average resource whether utilized for ramping or curtailment mitigation. We find similar results when considering the supply curves computed for each of the four seasons of the year, with the flat scenario yielding a smaller resource than the ramping or curtailment scenarios, but with little difference between the latter two scenarios. Both the ramping and the curtailment scenarios are likely to be more realistic than the flat scenario in representing the average Shift resource as it would be used in practice. Because the ramping model is substantially

---

[57] As discussed below the supply curve that results from this model is similar to the supply curve for the ramping scenario. Since the latter is conceptually simpler, we chose to focus on it for our primary analysis. Thus, we have given only a brief, qualitative description of the curtailment model here, rather than describing it in full detail.

[58] This model was not intended to represent precisely the probability of curtailment, which depends on a complex interplay of factors; rather, it is intended to provide a qualitative representation of curtailment probabilities as part of an exploratory analysis.

simpler than the curtailment model, we used the ramping scenario to calculate our primary results throughout the remainder of this study.

Figure C-5. Shift supply curves in 2030 as computed using the different DR filters described in this section, corresponding to the flat, curtailment, and ramping scenarios. The curtailment and ramping scenarios have a significantly larger resource than the flat scenario but are similar to one another.

It is also interesting to explore the impact of these different scenarios for Shift dispatch on the utilization of the Shift resource throughout the year. To compare the frequency of Shift dispatch throughout the year and from scenario to scenario, we assume that, in each scenario, the most probable Shift dispatch is 100 percent likely to occur, and all other pivot hours have dispatch with likelihood given by the relevant Shift filter, relative to that most likely dispatch. Given this assumption, we can then calculate the weighted average number of shift dispatches per day in a given period:

$$\bar{D} = \sum_{d \in P} \sum_{j \in d} \frac{w_j}{N_d}$$

(C-3)

where d represents a day in period P, which contains N_d days.

Figure C-6 shows this daily average, for each season and for the full year, in each of the dispatch scenarios considered. The flat scenario sees significantly more shifts per day, on average, since all shifts are assumed equally likely to occur, so that even small fluctuations in the net load will trigger a dispatch

if they are associated with a zero crossing. By contrast, in the ramping and curtailment scenarios, only substantial swings in the net load are likely to trigger a Shift dispatch, so fewer dispatches occur on average. The seasonal variation in Shift dispatch is also quite different among the different scenarios. For instance, the spring sees the fewest opportunities for Shift in the flat scenario but the most frequent dispatch when Shift is utilized to mitigate curtailment.

Figure C-6. Charts showing the relative need for Shift DR, by season of the year, in each of the three dispatch probability scenarios considered in this study. The relative need for shift is expressed as the average number of load shift events that would be called upon, per day, under the assumption that the most probable shift across the whole year is 100 percent likely to be dispatched.

C.2.6. The timescale of the Shift resource

A significant difference between the Phase 2 and Phase 3 work relates to the timescale of the Shift resource that is reported in the supply curve. In Phase 2, because all shifts were assumed equally likely, it was possible to compute the shiftable load (Equation (C-2)) over a full day, so the supply curve presented in Phase 2 reported an effectively *daily* quantity of shiftable energy consumption, under the assumption that two shifts would occur each day (one in the morning and one in the evening).

In the current phase, each pivot hour receives a different weight, so it is more natural to report the average quantity of shiftable energy *per dispatch*. Because the Phase 2 work explicitly computed the quantity over two assumed daily shifts, this means that the supply curve presented in this study will report a Shift resource that is smaller than what would have been reported in Phase 2 (given the same inputs) by roughly a factor of two, with a price ($/yr/kWh) that is roughly twice as high. It is important to recognize this difference when comparing the results presented here to the Phase 2 results.

C.2.7. Renewable generation forecasting

To compute a net load curve in DR-Path, it is necessary to have a forecast of renewable generation on the CAISO grid for each forecast year. In Phase 3, this was generated using the same methodology as in Phase 2, in which the actual hourly solar and wind generation in 2015 was scaled up according to the expected growth in each resource as computed using the California Public Utilities Commission's (CPUC's) Renewable Portfolio Standard (RPS) Calculator. For Phase 3, the forecast was updated to use

the most recent RPS calculator, version 6.2,[59] and the forecast years were updated to include 2030 in addition to 2020 and 2025.

C.2.8. Modeling TOU rate impacts

The load-reshaping layer of DR-Path modifies the forecasted load shape outputs of LBNL-Load to account for expected consumer responses to future TOU rates. In Phase 3, the load reshaping for the residential sector was updated to be consistent with the TOU impacts that were forecast as part of the CEC's 2017 IEPR (Kavalec et al. 2018). The IEPR forecasts provide absolute system-level load impacts from TOU response, in the context of the various IEPR demand forecast scenarios. To apply these to the individual LBNL-Load load shapes, we computed the fractional TOU load impact, in each hour of the year, for each of the IEPR high, mid, and low demand scenarios, and we applied these fractional impacts as multipliers on the LBNL-Load load shapes. Figure C-7 shows the fractional impacts for one week in the mid-demand scenario. The impacts are quite small, on the order of a few percent. Though these impacts are somewhat different in the other demand scenarios, in no case are the differences large enough to have a substantive impact on the DR-Path results. Thus, throughout this study, we use the mid-demand impacts only.

Figure C-7. The TOU impact multiplier for residential load shapes, as derived from the 2017 IEPR load forecasts for the mid-demand scenario.

C.2.9. Updated modeling approach for BTM batteries

This study introduces a conceptually new approach to considering BTM batteries as a DR enabling technology. The Phase 2 study considered that BTM batteries might be installed at some sites for purposes other than DR (e.g., demand-charge management in commercial applications) and that the charging and discharging of these batteries could be utilized for DR. This meant that the size of the DR resource from BTM batteries was restricted by the size of typical BTM battery systems that would be installed for non-DR purposes, and the battery charging and discharging was considered as a separate end use load shape that could be utilized for DR.

---

[59]See https://www.cpuc.ca.gov/RPS_Calculator/

In the current study, we instead consider that BTM batteries could be installed explicitly as a resource to provide Shift DR. In this picture, batteries are conceptualized as a DR-enabling technology, rather than a separate end use whose load can be shifted. Batteries then serve to enable some fraction of the total building load to participate in Shift DR, and battery systems can be sized to accommodate arbitrarily large loads as needed, up to the building's total peak demand, which is set as an upper limit on reasonable BTM battery systems (since there would be little reason to place a larger system than that behind the meter, rather than building grid-scale storage).

C.2.9.1. Calculating DR features for batteries

Recall that the DR features computed by DR-Path represent the maximum amount of load that can be shifted in principle for a given end use. These are then coupled with DR enabling technologies that can access some fraction of that maximum available resource. In the case of BTM batteries, we assume that these can, in principle, be installed with an inverter sized to handle the entire building's peak load. However, if the battery is used to provide Shift DR, the amount of shiftable energy will be restricted by the round-trip efficiency of the system.

Based on these assumptions, DR-Path 2.0 computes a set of "battery" features for each cluster, using the same calculation as laid out in Equation (C-1),but using a modified approach to computing the available load $L_{c,e,j}$. For Shed, the maximum available load that can be enabled by a BTM battery is just the peak total site load:

$$L_{c,batt,j} = \max_j\left(\ell_{c,tot,j}\right),$$

$$(C\text{-}4)$$

where $\ell_{c,tot,j}$ is the total average site-level load in cluster c in hour j. For Shift, the maximum shiftable energy is set by the peak site load, the Shift window, and the battery round-trip efficiency:

$$L_{c,batt,j} = \frac{1}{2}TR_{batt}\max_j\left(\ell_{c,tot,j}\right),$$

$$(C\text{-}5)$$

where T is the duration of the shift window being considered (e.g., 4 or 8 hours) and R_{batt} is the assumed round-trip efficiency of a BTM battery system. The factor of 1/2 accounts for the fact that half of the Shift window is in a shed period and half is in a take period. This equation says that a battery can be used, in principle, to take the peak building load in one half of the shift window and shed the same amount, less round-trip losses, in the other half. In this study, we assumed a battery round-trip efficiency of $R_{batt} = 0.8$.

C.2.9.2. Cost and performance modeling

A BTM battery system will have two primary drivers of its cost: the energy storage capacity of the battery, in kWh, and the charge/discharge capacity of the inverter and/or charge controller, in kW. The cost accounting framework for DR-enabling technologies in DR-Path computes initial technology costs

using two different cost inputs: (1) a fixed cost to install the technology at a site, in dollars per site, and (2) a variable cost that scales with the size of the load under control, in dollars per kW of peak load.

To integrate BTM battery systems into the set of DR technology inputs, we specify several prototypical system configurations having different ratios of inverter capacity to battery storage capacity, or *C-rates*. For instance, a battery system with a C-rate of 0.25 has a battery that is large enough, relative to its maximum discharge rate, to allow it to supply power for four hours at full output. Then the variable initial cost of the battery system can be written, in dollars per kW, as

$$V_{batt} = V_{inv} + \frac{1}{C}V_{sto},$$

(C-6)

where V_{inv} is the variable cost, in dollars per kW, of the inverter and charge controller; V_{sto} is the variable cost, in dollars per kWh, of the storage; and C is the C-rate of the prototype system. The cost in dollars per kW is what is needed as an input to DR-Path, for consistency with the modeling of other variable technology costs (see appendix section C.1.3); however, V_{batt} is also easily converted to units of dollars per kWh by multiplying it by the C-rate.

When computing the total amount of DR resource that is enabled by a given technology, DR-Path relies on a set of input parameters for each technology that specify the fraction of the relevant DR feature that the technology will enable to participate in DR for a given duration. To compute the performance parameters for batteries, we first assume that the battery's state of charge will be actively managed to allow it to participate fully in a given DR event. For Shed DR, this means a battery with C-rate C can provide load an amount of load shedding resource, over a duration T, given either by its maximum charge/discharge capacity, or by a fraction $(CT)^{-1}$ of that maximum capacity, whichever is smaller. Thus a battery with $C = 2$ can provide Shed at 100 percent of its maximum capacity for two hours, but only 50 percent of its maximum capacity for four hours.

For Shift DR the performance parameters depend on the duration T of the Shift window (which covers both the Shed and the Take period) and on the C-rate, such that a battery can shift a quantity of energy given either by the entire Shift feature, or by a fraction $2(CT)^{-1}$ of the feature, whichever is smaller. Thus a battery with $C = 2$ can shift 100 percent of its storage capacity over a four-hour window but only 50 percent of its storage capacity over an eight-hour window.

To compute the final cost of batteries as a DR resource, in the case of Shed or Shimmy, we can simply annualize the technology cost per kW, V_{batt}, by multiplying by the appropriate capital recovery factor, and then add program and incentive costs. To compute the annualized cost per kWh of Shift, we must first convert V_{batt} into units of dollars per kWh, and we also must account for round-trip efficiency of the battery system, which reduces the amount of energy that the battery can shift, relative to its nominal capacity. The formula for the procurement price of Shift from BTM batteries is thus given by:

$$P_{batt} = V_{batt} \times \frac{C}{R} \times f_{CR} + \frac{P_{admin} + P_{mkt}}{F} + P_{incent},$$

(C-7)

where P_{admin} and P_{mkt} are the program administration and marketing costs, in units of dollars per enrolled site; F is the Shift feature for the site under consideration; P_{incent} is the annual customer incentive cost, in \$/yr/kWh; R is the round-trip efficiency; and f_{CR} is the capital recovery factor that converts total cost into annualized costs, which has units of yr^{-1}.

It is important to note that the resulting representation of the cost of BTM batteries as a Shift resource, with units of \$/yr/kWh, is *not* the same as the cost of the battery storage itself, with units of \$/kWh. Instead, it represents a full-system cost, plus program costs, converted to the same units as the other Shift resources considered in this study.

To explore a range of reasonable battery systems, intended for different DR durations, in this work we modeled prototype battery systems with C-rates of 1.0, 0.5, and 0.25 (i.e., a "one-hour," "two-hour," and "four-hour" battery), each of which can be installed in any building in each of the residential, commercial, and industrial sectors to handle an arbitrarily large desired fraction of the building's peak load.

Appendix D. Detailed Modeling Inputs

The modeling inputs used in for DR-Futures in the Phase 2 study are detailed in the appendices of the Phase 2 report. There were some updates to these inputs for the Phase 3 study. This appendix provides a brief, self-contained summary of the modeling inputs for the LBNL-Load and DR-Path modules, as well as a detailed description of updates implemented for this work.

D.1. Summary of inputs

D.1.1. LBNL-Load

The input datasets for LBNL-Load are described in detail in appendix section C-1 of the Phase 2 report. In brief, the LBNL-Load input data fall into the following categories.

Customer demographic data. IOU-provided non-identifying demographic information on nearly the full IOU customer cohort circa 2014, including energy consumption information, ZIP codes, and building use categorization data.

Customer load shape data. IOU-provided hourly or 15-minute Smart Meter data for all hours in calendar year 2014, for a sample of approximately 200,000 customers in the IOU service territories.

SubLAP definitions. The customer clustering algorithm groups customers within individual subLAPs, which are the main regional partitions of the CAISO grid. See Figure D-1 for the subLAP definitions circa 2014, which were used in this study, based on IOU-provided mappings between ZIP codes and subLAPs.

Geographic and weather data. The model depends on historical weather station data from the National Oceanic and Atmospheric Administration (NOAA) to perform weather normalization, and it depends on geographical information from the U.S. Census Bureau to map customers to weather stations.

Demand and efficiency forecasts. Forecasting of load shapes to future years is based on the demand and energy efficiency (EE) forecasts associated with CEC's IEPR. See the next section for details of the forecasts used in this study.

Data on end-use saturation and load shapes. The load shape disaggregation depends on a variety of studies to determine the saturation and expected growth of certain end uses within the building stock and in some cases their typical load shapes.

Figure D-1. In 2014, PG&E service territory comprised 16 subLAPs, SCE service territory comprised 6 subLAPs, and SDG&E service territory comprised a single subLAP. Source: CAISO.[60]

D.1.2. DR-Path

The input data sources for the DR-Path module fall into the following categories:

LBNL-Load outputs. The summary of customer clusters and the disaggregated end-use load shapes for each cluster, as output from LBNL-Load, are used as the customer load shape inputs to DR-Path.

TOU impact model assumptions. These input files represent hourly forecasts of the impact of changing TOU rates on customer load shapes, based on a model for customer price response. The approach to developing these in Phase 2 is described in Appendix E of the Phase 2 report; this work uses a different approach for the residential sector as discussed below.

---

[60] California ISO. Proxy Demand Resource (PDR) & Reliability Demand Response Resource (RDRR) Participation Overview. https://www.caiso.com/Documents/PDR_RDRRParticipationOverviewPresentation.pdf

Hourly renewable generation forecasts. Hourly forecasts of renewable generation are based on actual historical generation (see appendix section B-2 of the Phase 2 report), which is then projected into the future using the California Public Utilities Commission's (CPUC's) Renewable Portfolio Standard (RPS) calculator.

DR-enabling technology cost and performance data. The DR-Path calculations of DR resources and associated costs are based on a database of cost and performance data for a wide variety of DR-enabling technologies. The data used for Phase 2 are described in detail in appendix section G-9 of the Phase 2 report. Updates to these data for Phase 3 are described below.

Customer participation propensity model. The DR-Path resource calculation depends on tabulated outputs of a model of different customers' propensity to participate in DR programs. This model is described in Appendix F of the Phase 2 report.

Miscellaneous assumptions. This category includes inputs such as assumed transmission and distribution loss factors, customer discount rates, and rates of technology improvement.

D.2. Updates for Phase 3

D.2.1. LBNL-Load

There were three primary updates to the LBNL-Load inputs for Phase 3. First, the forecasting of load and EE used the forecast data from the 2017 IEPR (Kavalec et al. 2018), updating from the 2015 IEPR data used in Phase 2. The implications of this change are summarized in appendix section B.1 of this report. Second, the model utilized new forecasts of EV charging load shapes, based on a recent CEC study (Bedir et al. 2018), as well as updated assumptions about EV penetration. The implications of these changes are summarized in appendix section B.2.1 of this report. Finally, LBNL-Load has been updated for Phase 3 to include forecasts of electrified residential space and water heating loads, based on stock turnover models and load shapes custom-built for this effort from a variety of different sources detailed in appendix section B.2.2 of this report.

D.2.2. DR-Path

The DR-Path inputs underwent several updates for the Phase 3 study. The most notable update involves significant new technology cost and performance inputs, which are detailed in the next section. In addition to this, we also updated the input data for reshaping residential loads in response to TOU rates, to align with CEC forecasts associated with the 2017 IEPR (Kavalec et al. 2018), as discussed in appendix section C.2.8 of this report. We also updated the forecasts of renewable generation to use the expected growth of solar and wind generation given in v6.2 of the CPUC's RPS calculator (CPUC 2016). Finally, we made some updates to the inputs used for the different technology scenarios (Base, BAU, Med, and High) described in section 2.3.2 of this report to include assumed rates of evolution in price, technology performance (i.e., fraction of load that can be shed or shifted), and customer participation. The assumptions used for these scenarios are summarized in Table D-1.

Table D-1. Technology Scenario parameters used for DR-Path in the Phase 3 study

Scenario	Price Decline Rate*	Performance Improvement Rate	Participation Growth Rate
Base	0	0	0
BAU	0.01	0	0.01
Med	0.02	0.01	0.02
High	0.04	0.02	0.03

*All rates are reported in terms of fractional change per year.

D.2.3. New DR-Path technology inputs

As discussed in appendix section C.2.4 of this report, DR-Path 2.0 introduced new assumptions about the performance of Shift, including both constraints on the temporal direction of load shifting for certain end uses and factors to account for energy losses (i.e., increases in total energy consumption) associated with Shift for certain end uses. In this study, we assumed that EV charging was constrained to shift forward in time only, since the baseline behavior is to begin charging the instant the vehicle is plugged in, meaning that the load cannot be shifted to earlier times since the EV is disconnected from the grid at those times. We also assumed that refrigeration loads could only be shifted to earlier in time—that is, that refrigerated warehouses would need to pre-cool before shedding load (rather than increasing their set point and recovering it later) to comply with food safety standards. We also assumed that space heating, space cooling, refrigeration, and water heating loads all incurred a 5 percent energy penalty when shifting load, owing to thermal losses in the pre-cooled or pre-heated space or water. This factor is based on LBNL's synthesis of a number of studies of precooling and thermal storage strategies (Delforge and Vokovich 2018; Herter and Okuneva 2013) as well as LBNL building modeling performed to inform this study, all of which indicate energy losses on the order of a several percent. We selected a consistent 5 percent value for the sake of simplicity in this study. For BTM batteries, we assumed a 25 percent loss factor, corresponding to an 80 percent round-trip efficiency (using the formula $1/0.8 = 1.25$), or a roughly 90 percent one-way charge/discharge efficiency.

To develop cost and performance inputs for individual DR-enabling technologies in this study, we started from the Phase 2 technology inputs and performed a literature review to seek any relevant new data. The new approaches we took to modeling EVs and BTM batteries in this study necessitated updates to the technologies associated with those end uses, and we also developed new inputs for TES technologies and for technologies associated with residential electric water heating. For all other end uses, we found no new data or studies suggesting that updates to the Phase 2 inputs were needed, so we utilized the same technology inputs as were utilized in Phase 2. Those inputs are summarized in appendix section G-9 of the Phase 2 report. Our new and updated technology inputs are summarized below and the specific input values are presented in Table D-2.

We derived our BTM battery inputs from an analysis of the levelized cost of storage (Lazard 2018) that was used to develop inputs for CPUC's 2019 integrated resource planning (IRP) process, in the interest of consistency across CPUC modeling efforts. The Lazard analysis gives sector-specific cost ranges for BTM battery storage, broken out into costs for charge/discharge capacity, in dollars per kW, and costs for battery storage, in dollars per kWh. For our inputs in this analysis, we selected costs in the middle of the

presented ranges; specifically, we used commercial sector costs of $225/kW and $500/kWh, and for the residential sector we use costs of $150/kW and $700/kWh. Inputting these values into the methodology described in appendix section C.2.9.2 of this report yields overall costs, in dollars per kW, for battery systems having a variety of C-rates, as presented in Table D-2. We also assume that annual operating cost outlays are required to secure licenses to communications and control software for the battery system. We also acknowledge the rapid recent (and expected future) decline in BTM battery costs by assuming a rate of price decline that is accelerated relative to the underlying rates assumed for the overall technology scenario (Table D-1): we assume that residential costs will decline by 20 percent, and commercial costs by ten percent, by 2020, in addition to the underlying decline from the technology scenario. Performance inputs were developed based on the assumption that batteries can shift up to 100 percent of load for as long as they have sufficient storage capacity—so, for instance, a battery with a C-rate of 1.0 can shift 100 percent of load for one hour, but only 50 percent of load for two hours, since the battery storage will be exhausted after one hour. Because we are envisioning potentially large battery systems that can handle up to a site's peak load, we assume co-benefit fractions that are quite small, to represent the battery system sizes that may commonly be purchased for purposes other than enabling Shift.

In conjunction with the new modeling of EV charging loads in Phase 3, we also updated the technology input assumptions for EV charging. These values replace the battery EV and plug-in hybrid EV Level 2 charging inputs that were used in Phase 2. For commercial EV charging, we assumed that Shift will be enabled using a Level 2 smart charger, with variable initial costs based on an assessment of charging infrastructure costs from the U.S. Department of Energy (Smith and Castellano 2015), along with an assumption for operating costs associated with a control software subscription. In the residential sector, we assumed a cost for Level 2 smart charging based on a survey of retail sites, as well as operating costs associated with software. In both cases we assumed Level 2 charging could shift nearly all of the relevant load for up to eight hours, owing to the superior speed of Level 2 charging. In the residential sector we also considered manual and automated Level 1 charging, assuming an initial cost for automated charging based on a survey of retail websites offering smart chargers along with an operating cost for software, whereas manual load shifting was assumed to be cost-free. We assumed that the Shift performance of these strategies was reduced relative to Level 2 charging, owing to the slower charging speeds and the need to complete a full charge, with an even greater reduction in performance in the manual case, owing to the burden on the consumer to participate. Co-benefit fractions for these technologies were carried over from Phase 2.

For TES systems, we derived performance estimates from an LBNL building performance modeling performed for this study to assess the load-shifting capabilities of TES systems. Typical outputs of the building models are shown in Figure D-2, which shows typical operation of a TES system for peak shaving (top panel), and the performance of the same system when used to shift load away from evening hours (bottom panel). The conclusion one can draw from this is that typical present-day TES system configurations (intended for mid-day peak reduction) have the capability to shift 100 percent of load out of the evening hours. Based on this, we assume that TES systems will be able to shift 100 percent of load during typical times of system need. We developed cost estimates based on the building modeling and a study of TES systems utilized in the context of Shed DR study (Yin et al. 2015). The study gives estimated TES system costs in large buildings of $275/ton-hr of storage. The building modeling implies

typical system efficiencies of 1 kW/ton, and up to four hours of potential load shifting capability, implying a variable initial cost of $275/ton-hr × 1 ton/kW × 4 hr = $1100/kW. We assume a 30 percent co-benefit fraction for this technology, which is consistent with assumptions for other HVAC DR technologies. In medium and small buildings, the report gives an estimated cost of $2,170/kW, although this includes the cost of the primary cooling unit, which would be purchased regardless of the TES system. We took the cost of the cooling unit to represent half of the total system cost, and so assumed a co-benefit fraction of 0.5 in this case.

Figure D-2. Modeled TES system operation for a typical system under strategies to reduce mid-day peak load (top) and shift evening load (bottom). A system that is sized for the common peak-shaving strategy shown at top is able to shed 100 percent of evening load in a Shift scenario.

For electrification of residential water heating we assumed that Shift will be achieved by use of a communicating control that overheats the water in the storage tank, coupled with a thermostatic mixing valve that delivers water at the desired temperature (BPA 2018; Delforge and Vokovich 2018). We developed cost estimates for this technology based on synthesis of data in public reports (Haeri et al. 2018; Hledik, Chang, and Lueken 2016) and expert estimates. We assumed that this technology had a significant customer co-benefit because, in addition to enabling Shift, it effectively increases the water heater's storage capacity, allowing a greater quantity of domestic hot water to be drawn before exhausting

the storage tank. Performance assumptions were based on modeling in a study specifically focused on load-shifting strategies for heat pump water heaters (Delforge and Vokovich 2018).

For electrification of residential space heating, we assumed that the technologies and costs were identical to those for space cooling, so we used the Phase 2 space cooling technology inputs.

Table D-2. New and updated cost and performance data for DR-enabling technologies in Phase 3.

End Use	Sector	Enabling Technology	Equipment and Installation Costs ($)	Variable Initial Costs ($/kW)	Fixed Operating Costs ($/yr)	Variable Operating Costs ($/kW/yr)	Co-benefit Fraction*	Average 1-hour Shed** (Fraction)	Average 2-hour Shed** (Fraction)	Average 4-hour Shed** (Fraction)
Battery	Res	Residential battery, C=1.0	0	850	120	0	0.03	1	0.5	0.25
		Residential battery, C=0.5	0	1,550	120	0	0.05	1	1	0.5
		Residential battery, C=0.25	0	2,950	120	0	0.10	1	1	1
	Ind	Industrial battery, C=1.0	0	725	0	33	0.03	1	0.5	0.25
		Industrial battery, C=0.5	0	1,225	0	33	0.05	1	1	0.5
		Industrial battery, C=0.25	0	2,225	0	33	0.10	1	1	1
	Com	Commercial battery, C=1.0	0	725	0	33	0.03	1	0.5	0.25
		Commercial battery, C=0.5	0	1,225	0	33	0.05	1	1	0.5
		Commercial battery, C=0.25	0	2,225	0	33	0.10	1	1	1
EV charging	Res	Level 1 smart charging	50	0	120	0	0.75	0.8	0.45	0.45
		Level 1 manual charging	0	0	0	0	0.00	0.20	0.11	0.11
		Level 2 smart charging	200	0	120	0	0.75	0.9	0.9	0.9
	Com	Level 2 smart charging	0	500	0	33	0.75	0.95	0.95	0.9
HVAC	Com	TES, large buildings	0	1,100	0	0	0.30	1	1	1
		TES, small/med buildings	0	2,170	0	0	0.50	1	1	1
Water Heating	Res	Communicating mixing valve	315	0	50	0	0.50	0.95	0.8	0.4

* The co-benefit fraction is the fraction of the cost that is defrayed by co-benefits that accrue to the customer.

** The average N-hour shed represents the fraction of load that can be shed for N continuous hours. This value is also used to compute the fraction of shiftable load in a Shift window of 2N hours.

Appendix E. Summary of software and data release

This appendix summarizes the release of software and data associated with this report. The release is available at https://buildings.lbl.gov/download-phase-3-california-demand-response. It consists of public downloads available in three different categories, as follows.

Documentation

Users may download a zip file that yields four descriptive documents when extracted:

- A summary of the software and data release (identical to the information in this appendix).
- A data dictionary in Microsoft Excel describing the data fields in the cluster profile data outputs from the LBNL-Load software module.
- A data dictionary in Microsoft Excel describing the tables in the database outputs from the DR-Path software module.
- An entity-relationship diagram (ERD) in PDF format, describing the structure of the DR-Path database.

Software

Users may download zip files containing the open-source Python software for each of the modules in the DR-Futures module, in the form that they existed when developing this report:

- The LBNL-Load module, in the file lbnl-load_phase3.zip
- The DR-Path module, in the file dr-path_v2.0.zip

Each module contains a readme file with basic installation and usage instructions, as well as a license file describing any conditions on use. Both modules consist primarily of Python code. They were primarily developed on Ubuntu Linux and are most likely to run successfully in a similar environment. It should be feasible to install these packages on other systems (e.g., Windows or macOS), but those platforms are not explicitly supported.

Basic operation of the software on small datasets is possible on personal computers, but operation on the full dataset made available here requires a research-quality server with significant memory and multiprocessing capabilities, beyond what is typically available on a consumer-grade computer.

Data

Users my download zip files containing the data outputs of LBNL-Load and DR-Path, for both the Base Electrification scenario and the Additional Electrification scenario from this report. In each case, the data represents customer clusters that differ slightly from the actual clusters used in this report, in order to meet the required anonymization conditions for the public release of customer energy consumption data.

Thus, results obtained from these files may not be identical to the results presented in this report. For each scenario, the following files[61] are available.

- The LBNL-Load cluster load profile outputs, in the file cluster_profiles.tgz
  - This file is a gzipped TAR archive, which can be extracted using common tools such as WinZip or 7zip, as well as the Linux gzip tool.
  - When extracted, this archive yields one directory for each relevant weather scenario and forecast year. Each directory contains two types of file:
    - One cluster summary file, containing descriptive data on each customer cluster.
    - For each cluster, one file containing the disaggregated end use load shapes for that cluster.
    - Data columns in these files are described in a data dictionary that is available as part of the Documentation download.
- The DR-Path database output, in the file dr_path.db.gz
  - This file is a gzipped SQLite database file.
    - It can be viewed in a simple manner using widely available tools such as SQLite DB Browser.
    - The recommended method for interacting with the database is to install the DR-Path module and use the data retrieval tools it contains, as described below.
  - The database tables and structure are described in the data dictionary and ERD available as part of the Documentation download.

To get started using the DR-Path software module for interaction with the DR-Path database, one should first download the DR-Path module and follow the installation instructions in the readme file. After successfully installing the module, place the unzipped dr_path.db file within the dr_path module folder. With these steps complete, a user can extract a basic supply curve by issuing the following commands in a Python session running within the DR-Path `pipenv` virtual environment:

```
from dr_path.db import io

dem = models.DemandScenario.get(id=1)
tou = models.TOUImpactScenario.get(id=1)
ren = models.RenewableScenario.select().where(
      models.RenewableScenario.weather == dem.weather)[0]
ts = models.TechnologyScenario.get(name='Base')

sup = io.get_supply_curve(dem, tou, ren, ts, year=2030,
                          timescale='year', dr_type='shift',
                          disaggregate_by=['sector','end_use'])
```

The variable sup now contains a supply curve for forecast year 2030, disaggregated by sector and end use, for the first demand scenario and TOU Impact scenario in the database, for the Base Technology scenario.

---

[61] Caution: the individual files are quite large, comprising several gigabytes (GB) each in zipped format, and will more than double in size when extracted. The largest file is larger than 40 GB when extracted

More substantial examples of interacting with the database and producing plots of the outputs can be found in the `postprocessing` subfolder within the `dr-path` module directory.