

Soaking Up the Sun: Battery Investment, Renewable Energy, and Market Equilibrium

R. Andrew Butters*
Jackson Dorsey†
Gautam Gowrisankaran‡

June 28, 2024

Abstract

Renewable energy and battery storage are seen as complementary technologies that can together facilitate reductions in carbon emissions. We develop and estimate a framework to calculate the equilibrium effects of large-scale battery storage. Using data from California, we find that the first storage unit breaks even by 2024 without subsidies when the renewable energy share reaches 50%. Equilibrium effects are important: the first 5,000 MWh of storage capacity would reduce wholesale electricity prices by 5.6%, but an increase from 25,000 to 50,000 MWh would only reduce these prices by 2.6%. Large-scale batteries will reduce revenues to both dispatchable generators and renewable energy sources. The equilibrium effects lead battery adoption to be virtually non-existent until 2030, without a storage mandate or subsidy. A 30% capital cost subsidy—such as the one in the U.S. Inflation Reduction Act—achieves 5,000 MWh of battery capacity by 2024, similar to the level required under California’s storage mandate.

JEL Codes: L94, Q40, Q48, Q55

Keywords: energy storage, solar PV, renewable energy, dynamic programming, electricity markets

*Kelley School of Business, Indiana University (e-mail: rabutter@indiana.edu).

†University of Texas at Austin (e-mail: jackson.dorsey@austin.utexas.edu).

‡Columbia University, NBER, and CEPR (e-mail: gautamg2@gmail.com).

We thank Jim Bushnell, Ken Gillingham, Kate Ho, Ashley Langer, Derek Lemoine, James Mackinnon, Erin Mansur, Vijay Modi, Mar Reguant, Stan Reynolds, Dan Steingart, Mo Xiao, and numerous conference and seminar participants for helpful comments, discussions, and suggestions. We also thank Filippo Paternollo and Yanhao Wang for excellent research assistance.

1 Introduction

Growth in renewable electricity generation has been dramatic over the past 10 years, in the U.S. and worldwide. By displacing generation from fossil fuels, renewables reduce greenhouse gas emissions. However, almost all recent growth in renewables comes from *intermittent* sources such as solar photovoltaics (PV): a solar farm cannot generate electricity after the sun sets, or when a cloud passes overhead. Absent the ability to store electricity, integrating these intermittent sources into the electricity grid requires the capability both to produce electricity at times with low expected renewable production and to adjust production suddenly when renewable production is unavailable. Intermittency reduces the benefits of renewables through the costs of building, maintaining, and operating additional fossil fuel generators (Bushnell and Novan, 2021; Gowrisankaran et al., 2016; Joskow, 2011). Thus, battery storage is a potentially important complement to intermittent renewable energy: it can lower the social costs of integrating renewables by storing energy when renewable production peaks and releasing it when it plummets.

In tandem with recent growth in renewable energy investment, the capital costs of lithium-ion battery cells fell by 85% from 2010 to 2018 with projections of 50% further cost drops over the next decade (Cole and Frazier, 2019; Goldie-Scot, 2019).¹ Despite these dramatic cost decreases, capital costs are still a central impediment to utility-scale battery storage. In addition, the equilibrium value of large-scale storage investment is limited because each additional storage unit acts as an arbitrageur, smoothing price differentials across time and lowering the value of existing units. Finally, even after capital costs reach a break-even point, companies may defer battery investments to exploit the option value of waiting for additional capital cost declines.

This paper has three main goals related to understanding the economics of battery storage. First, we develop a framework to estimate the equilibrium effects of large-scale battery storage and the complementarities between batteries and renewable energy penetration, in a model that incorporates dispatchable generator market power and the ramping costs that these generators bear when they raise or lower output. Second, we use our methods to calculate which parties would gain and which would

¹Other storage technologies are also expected to have up to 90% lower capital costs within the next decade (U.S. Department of Energy, 2021).

lose from large-scale battery adoption. Third, we evaluate the extent of expected equilibrium battery adoption and how this responds to different policies.

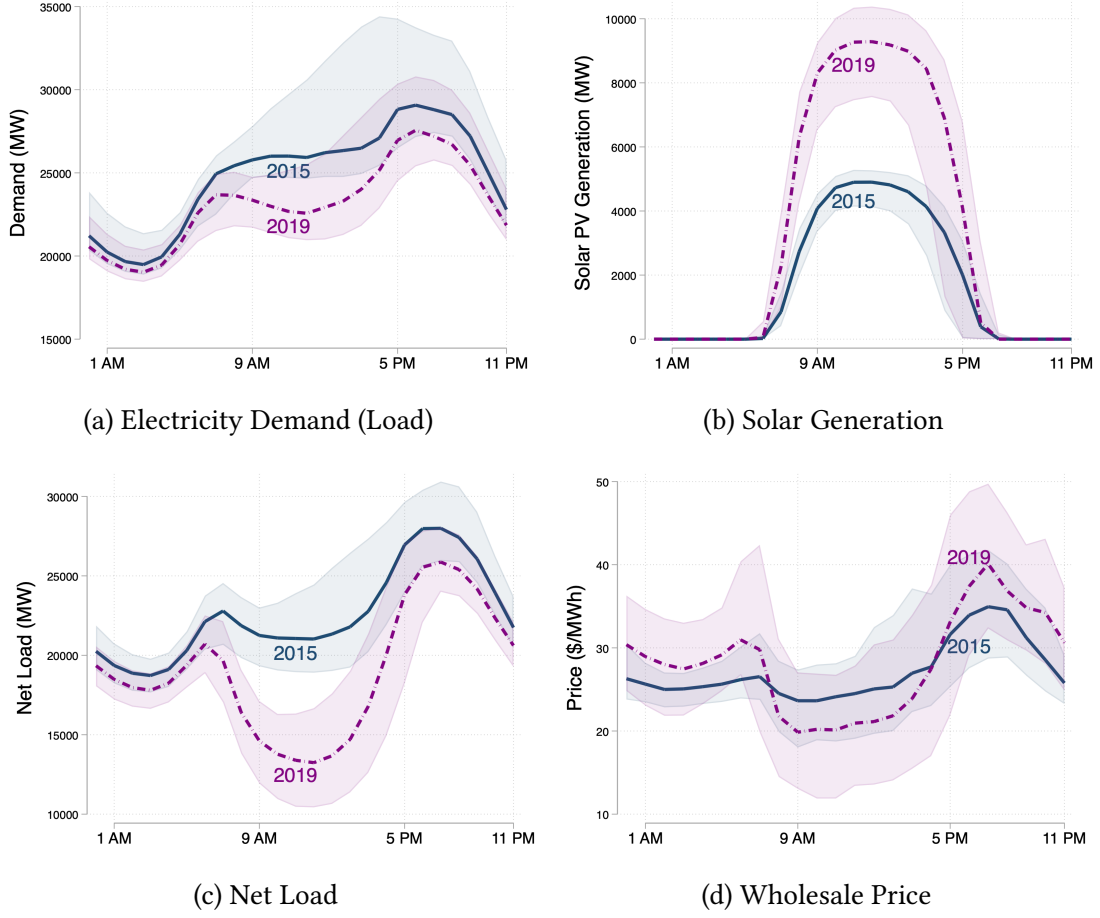
Understanding the complementarities between battery storage and renewable energy is particularly important because many policy proposals have paired renewable energy standards with battery mandates. For example, in conjunction with its aggressive renewable energy standards, California’s 2010 AB 2514 requires utilities to procure 1,300 MW of storage power capacity, operational by 2024.² The state justified the storage mandate on the basis that storage resources can help optimally integrate renewable energy resources and improve grid reliability ([California Secretary of State, 2010](#)). Additionally, implementing a concurrent battery mandate and renewable portfolio standard could be a cost-effective way to achieve renewable energy goals if there is the potential for coordination failures at the investment stage due to these complementarities ([Zhou and Li, 2018](#)). More recently, the 2022 U.S. Inflation Reduction Act (IRA) directly subsidized storage investments by providing federal investment tax credits ([House of Representatives, 2022](#)).

We illustrate the complementarities between renewable energy and storage with California data. Figure 1a displays median electricity demand and Figure 1b displays median solar generation, over the hours of the day and separately for 2015 and 2019. Solar generation in California increased dramatically over this period, but this generation typically occurs in the middle of the day and not in the evening, when demand is highest. Figure 1c displays median *net load*, which is the difference between total demand and intermittent renewable generation, and hence the electricity that is supplied by *dispatchable generators*.³ Net load in 2019 plummets in the middle of the day but rises again in the early evening to a similar level as in 2015, resulting in a curve with two humps. This change in the shape of the net load curve has at least two implications for costs. First, it implies that solar PVs are not producing in the evening when net load, and hence marginal costs, are highest. Second, it implies an increase in dispatchable generators’ ramping costs, as these units now need to change production levels more throughout the day ([Cullen, 2010](#); [Jha and Leslie, 2021](#); [Mansur, 2008](#);

²This size is similar to a large natural gas power station and could serve about 6% of the typical California Independent System Operator (CAISO) load. For the most common 4-hour duration batteries, this corresponds to 5,200 MWh of stored energy capacity.

³Unlike intermittent generators like wind and solar PV power plants, dispatchable generators, which include natural gas and hydroelectric plants, can be started on demand.

Figure 1: Electricity Demand, Solar Generation, and Prices by Year in California



Notes: Each panel shows the hourly median, 25th percentile, and 75th percentile of electricity demand (load), solar generation, net load, and real-time wholesale market price, respectively. Figures calculated by authors from California Independent System Operator data. All prices are for the California South Hub Trading Zone (SP15).

Reguant, 2014). Finally, Figure 1d displays median wholesale electricity prices. Despite the similarity in evening load between 2015 and 2019, median wholesale prices are substantially higher in 2019, suggesting the importance of increased ramping costs and the potential of storage to mitigate these costs.

We address our main goals by developing a new theoretical and estimation framework to understand equilibrium battery operations. A battery operator, acting as a price arbitrageur, must decide at each short time interval whether to pay to add to its energy inventory (i.e., charge) or use some of its inventory to generate revenues (i.e., discharge). This decision is inherently dynamic and stochastic: at each decision point, a battery operator maximizes the expected revenues from selling electricity in

the future net of the costs of purchasing that electricity.

Our model incorporates what we believe are key features of the electricity market: equilibrium effects of utility-scale battery fleets reducing the peaks and valleys of prices, ramping costs—where past generation by dispatchable generators reduces current marginal costs, and dispatchable generator market power. Beyond this, we incorporate predictable within-day fluctuations in net load; a non-linear dispatchable generator supply relationship for the wholesale electricity market that evolves over time; serial correlation of the shocks to net load and the supply relationship; a restriction that charge/discharge policies be based on data that would have been available in real-time to a market participant; a loss in energy from charging and discharging the battery; and degradation of batteries from operations, particularly with deep cycles. We estimate electricity demand and supply relationships using data from the California Independent System Operator (CAISO) from 2015-19.

The battery operations model allows us to address our first two main goals. To address the third goal, we link this model with a dynamic competitive equilibrium battery adoption model by leveraging additional assumptions. Our battery adoption model solves for an equilibrium of investment decisions of potential battery operators. Each year, potential battery operators make an optimal stopping decision, choosing whether to install capacity or wait, given battery installation costs, current and future renewable energy standards, and the mass of existing battery capacity. We use the solutions to the operations model—evaluated at counterfactual battery storage levels—to calculate profits for potential battery operators deciding whether to adopt a new system. To compute the adoption model, we estimate expected future battery capital costs using data compiled by the National Renewable Energy Laboratory.

Our results depend crucially on three main identifying assumptions. First, we assume that the net loads and supply relationships that we identify from the market data are structural and hence, would continue to hold given counterfactual large-scale battery operations. Our rich specification of the supply relationship—with market power, ramping costs, and serial correlation of the residuals—adds to the credibility of this assumption.⁴ Second, we assume that differences between wholesale day-ahead market and real-time market electricity prices reflect changes in dispatchable generation

⁴However, this assumption rules out the possibility that large-scale battery storage would cause fossil fuel generators to retire or price differently conditional on lagged and current generation levels.

capacity unavailability—which storage can help mitigate—instead of common generator cost shocks.⁵ Third, our adoption model—needed only for our third main result—uses weeks in our sample with high renewable generation as a proxy for a future with higher renewable penetration, after controlling for observable attributes of those weeks.

Relation to literature: Our study builds on three main literatures. First, it relates to an engineering and economics literature that investigates the value of storage in wholesale electricity markets. Early engineering papers in this literature modeled the storage decision using a finite-horizon framework and assumed that the storage device operator had perfect foresight about future prices or relied on historical prices when making discharge and charge decisions (e.g. [Sioshansi et al., 2009](#)). Other engineering studies relax the perfect foresight assumption and model storage decisions given uncertainty about future prices (e.g. [Mokrian and Stephen, 2006](#)). Our operations model extends this framework by considering the equilibrium effects of large-scale storage in competitive storage markets. It also relates to several recent economics papers. [Kirkpatrick \(2018\)](#) estimates the effect of recent utility-scale battery installations on electricity market prices and transmission line congestion in California. [Lamp and Samano \(2022\)](#) find that battery operators respond to price incentives at certain hours of the day, which has led to less wholesale electricity price variation. [Holland et al. \(2022\)](#) and [Karaduman \(2021\)](#) also consider the economics of grid-scale energy storage, employing different modeling approaches and data from ours.⁶ [Andrés-Cerezo and Fabra \(2023\)](#) examine theoretically how market structure affects battery investment and usage and through this, consumer surplus.

Second, we contribute to an economics literature that explores the market impacts of new energy technologies. [Wolak \(2018\)](#) and [Bahn et al. \(2021\)](#) measure the environmental and market effects of increases in renewable energy generation. [Feger et al. \(2022\)](#), [Langer and Lemoine \(2022\)](#), and [De Groote and Verboven \(2019\)](#) evaluate the impact of solar subsidies on adoption, while [Gonzales et al. \(2023\)](#) show how investments in transmission infrastructure increase the value of solar energy. Our results on

⁵Online Appendix [D](#) provides evidence supporting this point using auxiliary data on fuel prices, which affect costs similarly across many generators.

⁶Our results that there are large equilibrium effects of storage entry, that storage entry was not profitable during our sample period but would nonetheless raise consumer surplus, and that storage entry lowers solar and wind revenue are all consistent with [Karaduman \(2021\)](#).

the distributional impacts of renewable and battery adoption and market power add to a literature that includes [Bushnell and Novan \(2021\)](#), [Jha and Leslie \(2021\)](#), and [Liski and Vehviläinen \(2020\)](#), with our dynamic equilibrium framework.

Third, our work also relates to the literature on electricity forecasting ([Kanamura and Ōhashi, 2007](#); [Knittel and Roberts, 2005](#); [Weron, 2014](#)) and commodity storage ([Deaton and Laroque, 1992](#); [Pirrong, 2012](#)). Based on this literature, we develop and estimate a model of electricity demand and supply that allows for seasonal patterns, dynamics from ramping costs, and high-frequency cost volatility arising from unanticipated shocks to available generation.

Summary of Results: We find that a very small battery fleet would break even on the wholesale electricity market—i.e., earn enough revenues in the energy market to cover costs—if capital costs were to fall to \$264/kWh and renewable energy share were to increase from 40% (the share in 2019) to 50%, which are both expected to occur by 2024. This break-even figure incorporates both capacity degradation and uncertainty, which significantly limit the expected future profits that batteries can earn as arbitrageurs.

As the battery fleet expands in size, battery operations significantly lower the variation in mean equilibrium prices across hours of the day, in particular lowering prices in the evening peak. However, the marginal effects diminish as the battery fleet increases in size. For instance, the first 5,000 MWh of storage capacity would reduce prices by 5.6% but an increase from 25,000 to 50,000 MWh would only reduce prices by 2.6%. Large battery fleets also allow dispatchable generators to ramp more slowly, and thereby shift the peak production hour from 7 PM to 8 PM. The lower equilibrium prices imply that battery fleets of 10,000 MWh or higher would not be profitable as arbitrageurs by 2024 without subsidies or unless capital costs were to fall far below current expectations. Turning to the distributional consequences of batteries, 1,000 MWh of battery storage would decrease total revenues of dispatchable generators by \$126 million per year. More surprisingly, they would also decrease solar and wind generator revenues by \$14 million annually, as they reduce prices from 3 PM to 5 PM when many solar generators in California are still producing.

Finally, our adoption model—which incorporates the option value of waiting for future cost declines—shows that an ambitious renewable energy standard is not suffi-

cient to encourage large-scale battery adoption on its own. Specifically, battery investment would be negligible until 2030 without storage subsidies or mandates. However, a 30% capital cost subsidy—as specified by the 2022 IRA—yields approximately 5,000 MWh of battery capacity by 2024. This figure is very similar to the capacity required under California’s storage mandate.

2 Data and Institutional Setting

2.1 Storage Resources in the Electricity Market

Recognizing the complementarities with renewable energy, regulators nationally and in California have enacted new policies to increase electricity storage investment. In early 2018, the Federal Energy Regulatory Commission (FERC) issued Order 841, which requires independent system operators (ISO) to remove any existing barriers that would inhibit the participation of storage resources in wholesale markets.

In 2010, the California legislature authorized the California Public Utility Commission (CPUC) to evaluate and determine energy storage targets for the state. Accordingly, the CPUC required the state’s investor-owned utilities to procure 1.3 GW of storage power capacity by 2020,⁷ with installations required to be operational no later than the end of 2024. Since this time, California’s utilities have been adding storage capacity and, by 2019, utilities had at least 126 MW of operational battery power capacity.⁸

Though energy storage technologies such as pumped hydroelectric storage have been established for decades, the majority of recent utility storage installations use battery technologies. Our study focuses on one technology: lithium-ion batteries, which account for over 90% of U.S. battery storage capacity (EIA, 2020). A number of other emerging technologies allow electricity to be stored, including thermal energy storage, mechanical energy storage, and other forms of chemical energy storage, including hydrogen storage. Today, both the high capital cost and low round-trip efficiency of hydrogen storage make this route much less attractive than batteries, except for very

⁷Power capacity is the amount of power that the battery can supply to the grid at any point in time while energy capacity is the maximum amount of energy that the battery can store.

⁸Authors’ calculations based on maximum aggregate output reported by the California Independent System Operators between May 2018 and December 2019.

long duration storage ([Schmidt et al., 2019](#)), though this may change in the future. Importantly, our modeling framework could be used to assess the impact of alternative storage technologies with different physical parameters and cost projections.

Although the stock of utility-scale batteries is growing at a rapid rate, the overall battery fleet remains small. In 2018, there were only 900 MW of aggregate battery power capacity in the U.S., similar to that of two to three combined-cycle natural gas generators ([EIA, 2020](#)).

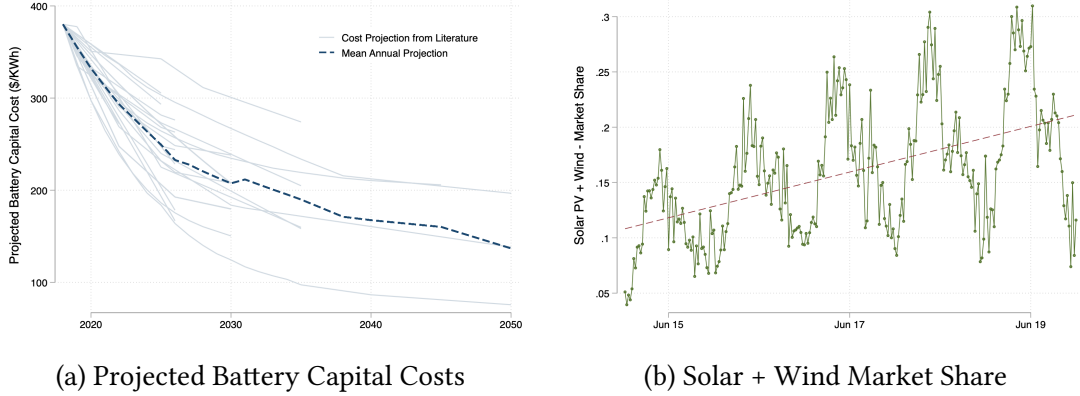
2.2 Battery Storage Costs, Technology, and Market Structure

Our adoption model relies on data on the capital costs of energy storage. Given the large expected declines in utility-scale battery capital costs, we use forward-looking projections from the National Renewable Energy Laboratory ([Cole and Frazier, 2019](#)) to model the evolution of future lithium-ion battery costs. These data compile utility-scale lithium-ion battery cost projections from over 25 publications published between 2016 and 2018.

Figure 2a summarizes the cost projections for battery storage over time in \$/kWh. Each point in the figure represents a normalized cost projection from a single publication for one year (with gray solid lines connecting multi-year projections within a publication), and the dashed line plots the mean projection by year. While most projections anticipate continued declines in capital costs, there remains considerable variation in just how large those declines are anticipated to be.

Batteries vary in their round-trip efficiency and duration. A battery's round-trip efficiency measures the percentage of stored energy that is available for later usage. A battery's duration indicates the amount of time the battery is able to discharge at its rated power capacity. For example, a 2-hour duration battery could discharge at full power capacity for 2 hours. Our study follows [Cole and Frazier \(2019\)](#) and focuses on 4-hour batteries with 85% round-trip efficiency. Four hours is the average duration of batteries operating in California in 2019, though shorter batteries are prominent within other ISO territories ([EIA, 2021](#)). Our round-trip efficiency figure implies that a battery that draws 1 MW of power from the grid can return 0.85 MW of power. Importantly, lithium-ion batteries degrade from repeated use and particularly from deep cycles, a factor that we incorporate in our model.

Figure 2: Battery Capital Cost Projections and Renewable Energy Trends



Notes: The authors constructed Figure 2a using data from [Cole and Frazier \(2019\)](#). Each gray transparent line represents a future cost projection from a single publication. The blue dashed line plots the mean cost projection. The figure reflects all cost projections related to grid battery applications (not electric cars). The authors constructed Figure 2b from CAISO data. It shows the share of electricity generation coming from solar and wind generators for each week between 2015 to 2019.

In general, battery storage is a nascent industry both nationally and in California. In this early stage, the battery market in California has been relatively unconcentrated, with a 2018 Herfindahl-Hirschman Index (HHI) of 1,347.⁹

We model batteries operating as arbitrageurs in wholesale energy markets. However, many of the earliest battery operators earned profits by supplying reserve capacity in ancillary services markets, most commonly regulation up and down. Despite this, three pieces of evidence suggest that future batteries will largely earn revenues from energy arbitrage. First, Figure A.1 in Online Appendix A shows that CAISO procured an average of less than 1,000 MW of hourly regulation reserves in all but three months of our five-year sample, and that the quantity was reasonably flat over time despite the increase in renewable energy over this period (e.g., see Figure 2b). Second, industry experts state that ISOs typically require only limited ancillary services capacity, on the order of 100-400 MW ([Sackler, 2019](#)). Third, more than 80% of the battery capacity added in 2021 in CAISO was used for energy arbitrage ([EIA, 2022a](#)). We view our analysis as pertaining to additional batteries that will earn profits as arbitrageurs rather than existing batteries that operate in ancillary services markets.

⁹Online Appendix B provides more details on battery market structure.

2.3 Operations Model Data

We estimate the main parameters of our operations model with data from CAISO over 2016-19.¹⁰ California restructured its electricity sector in 1998, and consequently designated CAISO the state's new independent system operator. CAISO dispatches over 200 million megawatt-hours of electricity to 30 million consumers each year, accounting for about 80% of electricity demand in California. CAISO runs two distinct wholesale energy markets: a day-ahead market (DAM) and a real-time market (RTM).

On the day before power is delivered, CAISO conducts 24 DAM energy auctions, one for each hour of the day, making available projections of net load prior to the auction. Market participants then submit bids to either buy or sell energy and CAISO computes market-clearing quantities and prices that meet the projected load at the lowest cost.¹¹ On the day of energy delivery, CAISO uses an RTM auction 75 minutes before each delivery hour to adjust generator production in response to unplanned outages or deviations. During the delivery hour, the system operator dispatches the lowest-cost generators every five minutes. The system operator uses reserve operations to meet any unanticipated imbalance within the five-minute interval.

Following FERC Order 841, CAISO has made efforts to integrate new storage technologies into its wholesale markets. CAISO allows batteries to submit either demand bids or supply bids in both day-ahead and real-time energy auctions. We focus on storage operators' final bids in the RTM, where the greatest arbitrage value lies, and which operators make having observed DAM prices. A battery can submit a set of prices and associated quantities at which it is willing to discharge energy, with negative quantities when it would like to charge. We use wholesale electricity prices from CAISO's South-Zone hub (SP-15), because this zone covers the largest share of the California population and currently hosts the most battery storage capacity. We augment the electricity price data with other market data: total load from the CAISO territory, generation by resource type, natural gas prices, and hydroelectric availability.

Notably, California's grid is currently undertaking a dramatic transition away from fossil fuel generation and towards renewable resources that will impact storage invest-

¹⁰We obtained data from the CAISO Open Access Same-time Information System (OASIS) portal. OASIS provides data related to the ISO transmission system and its markets. In some instances, we use CAISO data from 2015 as a training sample.

¹¹CAISO also uses the day-ahead market to secure energy reserves.

ment and operations. As of 2015, California already hosted the largest capacity of solar PV panels in the United States. Figure 2b shows that during the sample period of our study—January 2015 to December 2019—utility-scale solar and wind resources’ market share doubled from 10% to 20%, and exceeded 30% during some weeks. Going forward, state lawmakers have voted to boost renewable energy further under Senate Bill 100, signed in September 2018, which establishes the state’s updated renewable portfolio standard (RPS) (California Secretary of State, 2018). Figure A.2 in Online Appendix A provides details on California’s RPS schedule. The law specifies the share of generation that must come from renewable sources: 44% by 2024, 52% by 2027, 60% by 2030, and 100% by 2045. The figure also projects the share of energy that will come from solar and wind together for each future year that we model—as required by our adoption model—by linearly interpolating the RPS to intermediate years.

Figure A.3 in Online Appendix A provides more details on market trends in CAISO over our sample period. From Figure A.3a, average demand (load) for electricity has remained relatively stable, falling by 7.5%. Figures A.3b, A.3c, and A.3d show the solar, wind, and combined solar plus wind market shares over our sample period, respectively. Average wind power production increased slightly from 5% to 7% of generation, while solar PV’s generation share rose from 6% to 14%. Figure A.3e shows that prices for natural gas, the predominant fossil fuel generation source in CAISO, hovered around \$3/MMBtu for much of the sample period. Figure A.3f shows that mean prices in the real-time market have also trended upwards by nearly 20%. Finally, Figure A.4 in Online Appendix A replicates Figure 1d but with data at the five-minute, rather than hourly, level. It shows that real-time prices have become more volatile within each hour of the day as intermittent renewable generation has expanded.

3 Battery Operations Framework

3.1 Model

In our setting, a fleet of battery operators with total energy capacity K faces a set of dispatchable (typically, fossil fuel) generators. We model and estimate a wholesale electricity pricing function for dispatchable generators where price is a function of both current and lagged production, consistent with existing generator market power

and ramping costs.

Battery operators (or just operators) buy and sell energy in the real-time electricity market in each five-minute time interval, t , with the goal of maximizing their expected discounted profits from arbitrage. Our base model assumes that operators are competitive and take wholesale electricity prices as given, but operators' decisions together affect equilibrium prices.¹²

Batteries are characterized by three technological attributes. First, a battery's power capacity, F , determines what fraction of the battery can be charged or discharged in each five-minute interval, and therefore, how *quickly* the battery can transition from full to empty and vice versa. Second, the round-trip efficiency of the battery, v^2 , is the percentage of energy that is preserved during a full charge/discharge cycle. Finally, a battery's energy capacity degrades at a rate δ that depends on how and how much it is used. We model the capacity degradation rate δ using the [Xu et al. \(2016\)](#) algorithm, which provides an engineering-based formula of the percent of a lithium-ion battery energy capacity that "fades" (or degrades) over any time period as a function of the battery's sequence of charges and discharges.

At every five-minute time interval t , each operator makes a charge/discharge decision in order to maximize the sum of its expected discounted profits over an infinite horizon, using an annual discount factor β . Its decisions are a function of its charge level and the time-varying market state, which characterizes the current and expected future electricity market prices.

We focus on a symmetric equilibrium, where all battery operators start each time interval with the same energy held—which we denote $f \in [0, 1]$ —and then choose the same charge/discharge fraction each time interval—which we denote q .¹³ Each day consists of $S = 288$ five-minute time intervals. Let D denote the number of days within a year, d denote any day in our (multi-year) sample, and $s \in 1, \dots, S$ denote a particular time interval of a day.¹⁴ The five-minute interval discount factor is then

¹²Section 6 examines the impact of battery market power on outcomes.

¹³Because the choice variable is the charge/discharge fraction as a share of capacity and each battery takes the electricity market price as given, the equilibrium can still be symmetric even if batteries have different capacity levels.

¹⁴We use four different indices of time: t denotes a five-minute interval, $s \in \{1, \dots, 288\}$ denotes a five-minute interval within a day, d denotes a sample day (of which we have four years' worth), and, in Section 5, y denotes a calendar year. We need both s and d because our model includes interval-of-day fixed effects and separate parameter estimates by sample day.

$\beta^{\frac{1}{SD}}$.

Let $Q(q, K)$ be the net quantity of electricity supplied to the grid by battery operators at a time interval where this (common) discharge fraction is given by q :

$$Q(q, K) = K \times \left(\mathbb{1}\{q > 0\}qv + \mathbb{1}\{q < 0\}q/v \right).$$

Because $Q(q, K)$ is the net quantity supplied, it will be closer to 0 the lower the efficiency parameter v^2 , for a given charge fraction q and battery capacity K .

We define electricity net load to be the electricity load (or demand) by final users net of the amount produced by intermittent renewable sources (i.e., wind and solar). We assume that net load at time interval t is perfectly inelastic, varies across time, and is partly forecastable. Specifically, we let net load equal $X = X_s^d + \varepsilon_t^L$, where X_s^d is the interval-of-day forecastable mean net load and ε_t^L is unobservable until t . Our inclusion of X_s^d implies that the forecastable portion of net load varies across days in our sample and time intervals within a day.

Let Z be the amount of electricity supplied by dispatchable generators and \tilde{Z} be the amount supplied in the previous time interval. In the absence of storage, $Z = X$, since net load is the amount of electricity that needs to be supplied by dispatchable generators.¹⁵ With battery storage, $Z = X - Q(q, K)$. Next, we define a supply relationship for dispatchable generators (Bresnahan, 1982; Wolfram, 1999), $P^d(s, Z, \tilde{Z}, \varepsilon^L, \varepsilon^P)$. The supply relationship defines the equilibrium price as a function of the time interval of day, Z , \tilde{Z} , and two unobservable terms: the demand unobservable, ε^L , and a supply unobservable, ε^P . As with net load, the d superscript indicates that the supply relationship varies by day, which captures factors such as generator outages, transmission congestion, and changes in fuel prices. We include \tilde{Z} in the supply relationship to allow for ramping costs. All else equal, we would expect prices to be lower with a higher \tilde{Z} because more generators will be available to produce. For simplicity, we assume that operators believe that the forecastable demand and supply conditions of the current day repeat forever.¹⁶

We assume that the residuals ε^L and ε^P have a joint conditional distribution

¹⁵We assume that wind and solar are exogenous and exhausted before dispatchable generation.

¹⁶We focus on batteries that can completely fill or empty within a few hours, so expectations about changes in future days' demand and supply conditions will have relatively little influence on charging decisions.

$dG^{\varepsilon'}(\varepsilon^{L'}, \varepsilon^{P'} | \varepsilon^L, \varepsilon^P)$ that governs how they transition from one interval to the next. At the start of time interval t , operators know ε_t^L and ε_t^P and their joint conditional distribution. This joint distribution allows for serial correlation. This is important because if, for instance, a generator is unavailable during one time interval, it is likely to be unavailable in the subsequent time interval, and this knowledge will then affect the storage operator's charge/discharge decisions and profits.

Given our assumptions,¹⁷

$$\begin{aligned} \mathcal{V}^d(f, s, \tilde{Z}, \varepsilon^L, \varepsilon^P) = & \\ \max_q \left\{ P^d(s, Z, \tilde{Z}, \varepsilon^L, \varepsilon^P) \times (\mathbb{1}\{q > 0\}qv + \mathbb{1}\{q < 0\}q/v) \right. & \\ \left. + \beta^{\frac{1}{SD}} \int \mathcal{V}^d(f - q, s + 1 - \mathbb{1}\{s = S\}S, Z, \varepsilon^{L'}, \varepsilon^{P'}) dG^{\varepsilon'}(\varepsilon^{L'}, \varepsilon^{P'} | \varepsilon^L, \varepsilon^P) \right\}, & (1) \\ \text{s.t. } -Fv \leq q \leq F/v, 0 \leq f - q \leq 1, \text{ and} & \\ Z = X_s^L - Q(q^*(f, s, \tilde{Z}, \varepsilon^L, \varepsilon^P), K) + \varepsilon^L, & \end{aligned}$$

where ε' denotes the value of ε at the next time interval, and $q^*(f, s, \tilde{Z}, \varepsilon^L, \varepsilon^P)$ is the equilibrium quantity discharged at that state and is equal to the value of q that maximizes (1) at every state. Lines 2 and 3 of equation (1) specify the tradeoffs involved in optimization. Each operator faces a cost of P^d/v for energy it charges from the grid. But, charging q increases its energy held at the next time interval from f to $f + q$. If future prices are sufficiently higher than the current price, the operator can earn revenues from charging now and discharging in the future. Lines 4 and 5 detail the constraints that the operator faces: it is limited in the speed of its charge and discharge, its inventory cannot fall below 0 or rise above its capacity, and dispatchable generation, Z , must match net load minus net quantity supplied by operators, Q , given their actions.

To ease computation, we use the fact that battery operators are price takers to recast the competitive battery operations problem as a single-agent decision problem, where the incentives of the single agent correspond to the equilibrium incentives of a competitive battery operator. To understand the appropriate single-agent problem,

¹⁷In the Xu et al. (2016) capacity fading model, battery degradation depends on battery usage, but in a complex and non-linear way and over a long time horizon. For simplicity, we do not model cumulative battery usage that would lead to degradation as a state variable, but rather let battery operators account for degradation in their charging decisions with a heuristic approach; see Section 3.3 for details.

consider a competitive battery market and a state where battery operators are charging, but less than their power or energy capacity. The competitive market will charge energy to the point where the market price equals the expected discounted future revenues from discharging the energy.¹⁸ Now consider a single agent that maximizes a dynamic problem, whose current period criterion function is the negative of the integral of the pricing function—from 0 to the chosen charge/discharge level. The derivative of the integrated pricing function equals price, and therefore, the single agent's first-order condition and corresponding charge/discharge decisions are equivalent to that of a competitive battery operator. We solve:

$$\begin{aligned} \mathcal{W}^d(f, s, \tilde{Z}, \varepsilon^L, \varepsilon^P) = \max_q \Big\{ & - \int_0^Z P^d(s, \zeta, \tilde{Z}, \varepsilon^L, \varepsilon^P) d\zeta \\ & + \beta^{\frac{1}{SD}} \int \mathcal{W}^d(f - q, s + 1 - \mathbb{1}\{s = S\}S, Z, \varepsilon^{L'}, \varepsilon^{P'}) dG^{\varepsilon'}(\varepsilon^{L'}, \varepsilon^{P'} | \varepsilon^L, \varepsilon^P) \Big\}, \quad (2) \\ \text{s.t. } Z = X_s^L - Q(q, K) + \varepsilon^L, & -Fv \leq q \leq F/v, \text{ and } 0 \leq f - q \leq 1. \end{aligned}$$

Equation (2) depends on the pricing function evaluated at different values of Z and \tilde{Z} . Since Z and \tilde{Z} are, respectively, the current and lagged electricity supplied by dispatchable generators, they will adjust based on the charge/discharge decisions of utility-scale batteries. In the special case where the generator market is perfectly competitive, Equation (2) is equivalent to a social planner minimizing expected discounted total costs. For a similar model to ours, [Cullen and Reynolds \(2023\)](#) prove that competitive equilibria and a solution to the planner's problem exist, and that the planner's solution is equivalent to *all* competitive equilibria.

Our modeling approach leverages restrictions on dispatchable generators' supply relationship. We impose:

Assumption 1. *The equilibrium supply relationship for dispatchable generators is a function of only the variables $(d, s, Z, \tilde{Z}, \varepsilon^L, \varepsilon^P)$. In particular, the supply relationship is invariant to installed battery capacity.*

Assumption 1 is an exclusion restriction that allows us to identify the supply relationship and the equilibrium price effects of batteries: it specifies that while battery capacity can affect generators' market-clearing bids, this occurs only through changes

¹⁸Analogous arguments hold for the other cases where battery operators are charging at capacity, not charging or discharging, discharging less than capacity, or discharging at capacity.

in $(d, s, Z, \tilde{Z}, \varepsilon^L, \varepsilon^P)$. It is restrictive for two reasons. First, and most importantly, it imposes that the set of bidding generators be held fixed. If large-scale batteries cause dispatchable generators to exit, this would shift the supply relationship to the left and violate Assumption 1. In addition, it implies that dispatchable generators' dynamic equilibrium bids are only a function of $(d, s, Z, \tilde{Z}, \varepsilon^L, \varepsilon^P)$. This rules out the possibility of changes in Z due to large-scale batteries causing generators to change their equilibrium bids by changing expectations about future values of Z .

While Assumption 1 is restrictive, it does allow generators' bids to respond to battery operations. Specifically, bids are a function of load supplied by dispatchable generators and its lag. For example, when a large-scale battery fleet charges, generators will be called on to supply more electricity in the current period. This in turn will also lower their bidding function in the subsequent five-minute interval as their current higher quantity supplied means that fewer generators will need to ramp up to meet any level of quantity Z next period.

In sum, our generator pricing function is fully consistent with generator optimization only in the absence of battery capacity or if generators act myopically in their bids. Given this, we believe our counterfactual predictions of equilibrium prices are most credible for modest departures from the bidding environments that generate our market-level price and quantity data (i.e., environments with relatively little battery capacity present).

We solve the operations model by discretizing the state elements \tilde{Z} , ε^L , ε^P , and f into 10 dimensions each and solving the single agent problem in (2).¹⁹ We solve the optimization separately for each day in our 4-year main estimation sample and across 9 candidate values of K , resulting in about 13,000 dynamic problems with 2,880,000 states each.²⁰ The infinite horizon solution is very computationally challenging to solve. Instead, we solve for a finite approximation of the infinite horizon model. For

¹⁹We discretize the transitions of ε^L , ε^P by assuming that the innovation to these shocks is independent and normally distributed. We use the Tauchen (1986) procedure to discretize ε^P and use the Rouwenhurst method to discretize ε^L , which avoids the sensitivity of the Tauchen (1986) procedure to very persistent processes (Kopecky and Suen, 2010). We discretize f into 10 equally spaced grid points between zero and one. Similarly, we discretize \tilde{Z} into ten equally spaced grid points, spanning from the minimum to the maximum feasible level of \tilde{Z} . The feasible span of \tilde{Z} at any time is determined jointly by the observed level of net load in the CAISO data and the specified battery capacity, which determines how far net load can deviate from its observed level due to storage operations.

²⁰We also solve the operations model under an (infeasible) assumption of perfect foresight. For this model, we assume that the current and future values of ε^L , ε^P are known to the operator before it makes its charge/discharge decision. The state space for this model is thus much smaller.

each sample day d , we set up a finite horizon model with the base 288 intervals for the day plus 288×3 additional intervals, which repeat the same set of net load and supply parameters. We verified that the policies computed from the finite approximation are virtually identical to the policies from the infinite horizon solution. For robustness to distributional assumptions about $(\varepsilon^L, \varepsilon^P)$, after solving for the optimal policies, we compute counterfactual market outcomes by applying the policies to the *realized* time series of $(\varepsilon^L, \varepsilon^P)$.

3.2 Estimation of Supply Relationship and Net Load

Having described our operations model, we now turn to the estimation of our key structural parameters. We estimate the supply relationship and net load structural parameters with data from the wholesale electricity market and without imposing our structural model of battery optimizing behavior. Our central goal is to develop credible estimates of these parameters using information that operators could themselves observe in real time. This informational component is important because we do not want to inadvertently overstate the value of battery storage as arbitrageurs by providing operators in our model with more information than operators participating in this market would have when forming their operations decisions.

Starting with our supply relationship, $P^d(s, Z, \tilde{Z}, \varepsilon^L, \varepsilon^P)$, we first define *dispatchable generation capacity*, \mathcal{K} , which indicates the maximum quantity that can be supplied by dispatchable generators at any given time interval, and which is a function of \tilde{Z} and ε^P .²¹ This transforms the supply relationship to be a function of capacity utilization, $Z/\mathcal{K} \in [0, 1]$, and \mathcal{K} . We let $\tilde{P}^d(Z/\mathcal{K}, \mathcal{K})$ denote the transformed function.

For our supply relationship to make economic sense, our estimation imposes two monotonicity properties. First, \mathcal{K} should be strictly increasing in \tilde{Z} , because a higher level of generation in the previous time interval will result in more generators available to produce electricity without bearing ramping costs. Second, \tilde{P}^d should be strictly increasing in Z/\mathcal{K} , as greater capacity utilization implies that higher marginal cost generators—such as peakers—must be used, which will tend to drive up market prices.

We choose a simple Cobb-Douglas functional form for dispatchable generation

²¹We use \mathcal{K} for dispatchable generation capacity to distinguish it from battery capacity K .

capacity:

$$\mathcal{K} = \kappa^\alpha \tilde{Z}^{1-\alpha} \exp(\varepsilon^P), \quad (3)$$

where $\kappa > 0$ and $\alpha \in (0, 1)$ are parameters that we estimate.

Our base model imposes a functional form for \tilde{P}^d taken from the commodity storage literature (Pirrong, 2012):

$$\tilde{P}^d(Z/\mathcal{K}, \mathcal{K}) = \theta_1 + \theta_2[\mathcal{K}(1 - Z/\mathcal{K})]^{-\theta_3}, \quad (4)$$

where θ_1 , θ_2 , and θ_3 are parameters to estimate. Equation (4) satisfies our monotonicity conditions: for $\theta_2, \theta_3 > 0$, prices are increasing in capacity utilization and decreasing in capacity. In addition, because this functional form lets price asymptote to infinity as capacity utilization approaches one, it can capture the spikes that occur frequently in wholesale electricity prices (e.g. Borenstein et al., 2002; Knittel and Roberts, 2005).²²

Collecting terms, the structural parameters that we estimate for \tilde{P}^d are $(\alpha^d, \kappa^d, \theta^{1d}, \theta^{2d}, \theta^{3d})$.²³ We estimate these parameters using DAM prices. At the DAM stage, we assume that $\varepsilon^P = 0$. The idea is that when dispatchable generators bid in the DAM market, they do not yet know last-minute changes in capacity, which enter into ε^P . There may still be variation in the observed prices relative to predicted prices, corresponding to measurement or optimization errors. We impose that the deviations are orthogonal to the observable regressors and estimate the parameters using non-linear least squares, choosing:

$$(\hat{\alpha}^d, \hat{\kappa}^d, \hat{\theta}^{1d}, \hat{\theta}^{2d}, \hat{\theta}^{3d}) = \arg \min_{\alpha^d, \kappa^d, \theta^{1d}, \theta^{2d}, \theta^{3d}} \sum_t \left[P_t^{DAM,d} - \tilde{P}^d(Z_t^d/\mathcal{K}_t^d, \mathcal{K}_t^d) \right]^2, \quad (5)$$

where $\mathcal{K}_t^d = \kappa^{d\alpha^d} \tilde{Z}_t^{1-\alpha^d}$; $P_t^{DAM,d}$, Z_t^d , and \tilde{Z}_t^d are data; and t indicates a sample hour.²⁴

We estimate a separate specification for (5) for each day d of our sample, using

²²Online Appendix C provides results from an alternative flexible functional form from the industrial organization literature (Fowle et al., 2016; Ryan, 2012) that imposes similar monotonicity conditions. It would also be possible, though computationally challenging, to estimate non-parametric polynomial specifications for \tilde{P}^d that directly impose these monotonicity properties (Compiani, 2022).

²³We include ‘ d ’ superscripts since we allow these parameters to vary by day of the sample.

²⁴DAM prices and quantities vary at the hourly, not five-minute, level.

data for one week, with all hours in the current and the previous 6 days.²⁵ Because very few batteries (as measured by capacity) engaged in arbitrage during our sample period, we directly substitute net load, X , for the electricity supplied by dispatchable generators, Z (and analogously for its lag, \tilde{Z}). In this step of the estimation, we use the predicted net load—which is load minus solar and wind generation—all as reported by CAISO in its DAM forecasts.

Turning to our estimation of the net load process, given our assumption that net load is perfectly inelastic, we use the predicted net load reported by CAISO in its DAM forecasts as our estimate of the mean net load, $X_{s(t)}^{L,d}$. Since the DAM forecasts are only reported at the hourly frequency, we temporally disaggregate the net load forecasts to the five-minute level using a Kalman filter/smoothing approach; see Online Appendix E for details.²⁶

Turning now to the unobservables ε^P and ε^L , we estimate these values from the RTM. Batteries can bid in the RTM, having observed the sequence of DAM prices and supply relationships, but not future RTM prices. Our idea is that RTM price fluctuations relative to the DAM arise due to unanticipated changes in net load relative to the DAM forecast and unanticipated shocks to availability of generation capacity. Hence, for each time interval, t , we recover the value of ε^P that makes the wholesale electricity price equal to the observed electricity price in the RTM conditional on supply relationship parameters and the realizations of net load.²⁷ Thus, ε_t^P is defined implicitly by:

$$P_t^{RTM,d} = \tilde{P}^d \left(\frac{Z_t}{\kappa^{d\alpha^d} \tilde{Z}_t^{1-\alpha^d} \exp(\varepsilon_t^P)}, \kappa^{d\alpha^d} \tilde{Z}_t^{1-\alpha^d} \exp(\varepsilon_t^P) \right). \quad (6)$$

It is easy to verify that (6) does in fact define a unique ε_t^P for the Pirrong (2012) functional form because prices are monotonically decreasing in \mathcal{K} for a given net load Z , and a higher ε_t^P implies a higher \mathcal{K} and no change in Z .

The assumption that RTM supply relationship fluctuations are due to generator

²⁵Before estimation, we scale both prices and quantities, since they vary considerably both seasonally and across years. Online Appendix D provides further details on our estimation of the supply relationship.

²⁶CAISO market reports indicate that the CAISO day-ahead load forecasts are shaded up to ensure sufficient supply is available. We scale the net load forecasts by 0.95 to reflect this practice. This choice is supported by the empirical relationship between the day-ahead market forecasts and the realized values, see Table A.3, panel (a) in Online Appendix A.

²⁷RTM prices are available at the five-minute level and hence t now indicates a five-minute time interval.

unavailability is important for our analysis. It implies that batteries can mitigate peak prices by supplying energy at times when dispatchable generation capacity is scarce, which will tend to imply important equilibrium effects. Alternatively, if price variations within a day were due to a shock common to all generators (e.g., a common fuel price shock where all generators had the same heat rate), then the difference between RTM and DAM prices would not vary based on the amount of available energy that batteries supplied. In this case, equilibrium effects may be smaller since battery operators supplying energy when prices spike would not affect this fuel premium. Online Appendix D provides evidence of the plausibility of this modeling assumption.

We model the transition of ε^P as an AR(1) process given by:

$$\begin{aligned}\varepsilon_t^P &= \rho^P \varepsilon_{t-1}^P + \sigma_{s(t)}^P \eta_t^P \\ \sigma_{s(t)}^P &= \begin{cases} \sigma^{P,\text{Peak}} & \text{if } s(t) \in 5\text{--}10 \text{ PM} \\ \sigma^{P,\text{Off-peak}} & \text{if } s(t) \notin 5\text{--}10 \text{ PM} \end{cases},\end{aligned}\tag{7}$$

where η^P is a mean zero serially uncorrelated shock with unit variance, ρ^P governs the persistence of changes to available capacity, and $\sigma_{s(t)}$ accommodates any heteroskedasticity that exists across peak (i.e., 5 PM to 10 PM) and off-peak hours of the day.

Finally, we recover the net load unobservable ε^L as the difference between the realized net load, X_t^{RTM} and our forecast of net load from the DAM. We model the transition of ε^L as an AR(1) process given by:

$$\varepsilon_t^L = \rho^L \varepsilon_{t-1}^L + \eta_t^L, \quad \eta_t^L \sim N(0, \sigma^L),\tag{8}$$

where ρ^L and σ^L are parameters to estimate. We estimate each of the AR(1) models using ordinary least squares (OLS) on a training sample in 2015,²⁸ and hold these parameters fixed over the evaluation sample, 2016–19.²⁹ This ensures that the policies would be feasible to estimate and implement given the information set of a market

²⁸Day ahead forecasts for solar and wind are publicly available starting in Nov. 2015. Thus, our training sample includes only data from Nov. and Dec. 2015.

²⁹For our estimates of $\sigma^P(\text{Peak})$, $\sigma^P(\text{Off-Peak})$, we use a robust (and consistent) estimator of the scale for the normal distribution: $1.4826 \times \text{median}_t\{|x_t - \text{median}_j x_j|\}$ (Rousseeuw and Croux, 1993).

participant.

3.3 Calibration of Battery Technology Parameters

Our battery operations model depends on the battery’s storage technology. In many cases, industrial organization economists have structurally estimated technology parameters by imposing the assumption of optimizing behavior (Rust, 1987). However, our sample period includes very little observed battery behavior in the wholesale electricity market. For this reason, we estimate the battery’s technology parameters using engineering estimates, rather than from revealed preferences and structural estimation. Following Section 2.2, we model batteries with a duration of four hours, and thus set $F = \frac{1}{4 \times 12}$. In addition, we model the round-trip efficiency as $v^2 = 0.85$.

The final technology parameter is capacity degradation. Battery capacity degrades over time, a process that is accelerated by use, particularly deep cycles (Xu et al., 2016). Our primary goal here is to specify a model of optimizing behavior by operators in response to degradation. Because degradation is exacerbated by use, operators should defer some marginal arbitrage opportunities to reduce battery degradation. However, modeling optimization accounting for degradation is difficult because degradation is a non-linear function of cycling across multiple subsequent five-minute time intervals. Instead, we endogenize degradation in a heuristic way: we assume that operators account for greater charging intensity leading to more degradation by choosing their charge decisions using a lower perceived efficiency than the true value, v .

Specifically, we solve our model using the 2015 pre-analysis training sample and determine the optimal “perceived” roundtrip efficiency parameter, which we denote v^* .³⁰ This involves finding the value of v^* that would maximize the lifetime value of a battery considering the degradation effects specified by the engineering literature (Xu et al., 2016). For each fleet size K , we solve for a heuristically optimal $v^*(K)$ and then solve the battery operations decision in (2) using $v^*(K)$. An $v^* < v$ leads generators to decline marginal arbitrage opportunities, which affects operations decisions and also equilibrium adoption decisions modeled in Section 5.

Online Appendix F provides implementation details for our capacity degradation

³⁰For computational ease, we use a grid of different candidate perceived round-trip efficiency levels, $v^p \in [.6v, .65v, \dots, v]$ and solve the perfect foresight version of our model.

model. Figure A.5 provides a schema of the different parts of our operations framework for a single day.

4 Main Results

4.1 Supply Relationship and Net Load Parameter Estimation

Table A.1 in Online Appendix A reports sample statistics on the supply relationship parameters. For each supply relationship parameter, we report the mean, standard deviation, and 25th and 75th percentiles of the distribution of all the daily estimates by year. We find a considerable amount of variation in the parameters, even within a year. Some parameters appear to have longer-run trends. E.g., the intercept and slope terms— θ_1 and θ_2 respectively—both trend downward, consistent with the declining natural gas prices. The exceptions to these patterns are the parameters governing the weight on the current dispatchable generation capacity in the Cobb-Douglas capacity function, α , and the parameters governing the curvature of the supply relationship, θ_3 . In the case of α , the estimates center around 0.85 and are fairly stable, indicating the presence of positive and similar ramping costs throughout our sample. In the case of θ_3 , the mean estimates across the year range from 1.07 to 2.07, with a fairly skewed distribution towards 1. Finally, estimates for κ indicate that the scheduled available capacity (relative to the day-ahead forecasted maximum net load) is relatively stable over our analysis sample.

Table A.2 in Online Appendix A reports our parameter estimates for the ε^P AR(1) process. Our estimate of ρ^P used in our simulations—based on the training sample of 2015—is 0.947. We also report (but do not otherwise use) the AR(1) parameters for our evaluation sample. We find that ρ^P falls a little over time—lying within a range of 0.832 to 0.897. Our training sample on- and off-peak estimates of the standard deviations are 0.012 and 0.10, respectively. These estimates are stable over time, with virtually identical overall averages for each evaluation sample year. Across both the training and evaluation samples, comparing estimated $\sigma^{P,\text{Peak}}$ and $\sigma^{P,\text{Off-peak}}$ values, on-peak hours have about 25 percent more volatile changes in ε^P than do off-peak hours.

Table A.3 panel (b) in Online Appendix A summarizes estimation results for the net load model. Our estimate of ρ^L is very close to one—indicating a very high level

of persistence in the day-ahead forecast errors. The parameters governing the AR(1) process (ρ^L, σ^L) are fairly stable across both our training and evaluation samples, with only σ^L exhibiting a modest increase over the evaluation sample.

4.2 Profitability of Small Battery Fleet

We first use our model to estimate the value of a small battery fleet that can charge and discharge energy as arbitrageurs without affecting equilibrium electricity prices. This allows us to evaluate the conditions under which initial battery investments would reach a break-even point and also provides an informative benchmark about battery profitability in the absence of equilibrium effects.

We proceed by evaluating the profits of the small battery fleet for each sample week over the 2016–19 period. In particular, we approximate the profits of a small battery by solving for optimal charge/discharge policies with an aggregate battery capacity of $K = 10$ MWh from (2), and then simulating the weekly returns with these policies.³¹ We then convert each of these weekly observations into a heuristic lifetime value of storage capacity, using a weekly discount factor of $\beta = 0.95^{7/365}$, and a weekly degradation rate from these policies, calculated as in Online Appendix F.

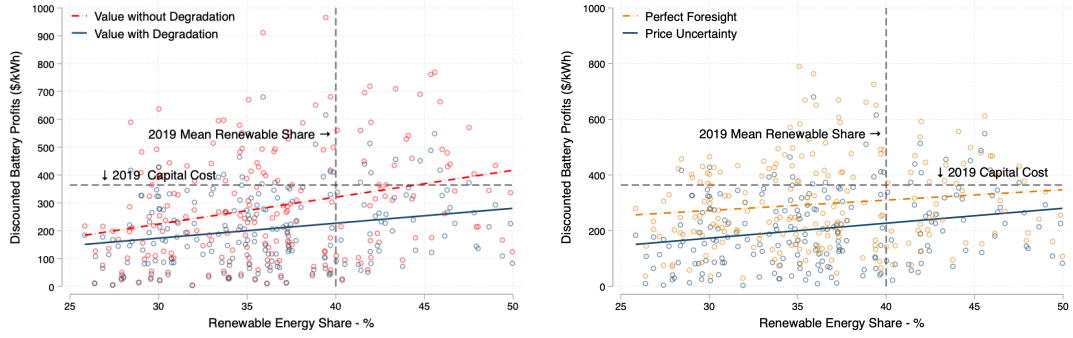
Figure 3 uses these calculations to illustrate these lifetime values relative to capital costs, with and without accounting for degradation. The dashed-red line plots a simple linear fit of the relationship between battery profits and the share of electricity generated by renewable sources, before adjusting for capacity degradation.³² We find a strong positive association between renewable generation and the value of storage.³³ The dashed-grey line shows the expected capital cost per kWh of storage capacity in 2019. Together, these lines show that, absent capacity degradation, lifetime battery profits would exceed the 2019 expected capital cost of storage if the renewable energy

³¹The single-agent Bellman equation policies and returns from (2) for $K = 10$ relative to $K = 0$ (divided by 10) will approximate the small fleet, since price is roughly equal to marginal revenue for a small fleet.

³²We calculate the renewable energy share as the percentage share of solar plus wind generators during the sample week plus 19%. 19% is the mean generation share from non-intermittent renewables, including hydro, geothermal, and biomass generators across the sample period.

³³Our measure of generation from renewables is net of any curtailments, which we do not model explicitly. In the California market over our sample period, curtailment of solar and wind generation was relatively small. However, to the extent that batteries eliminate the need for renewable sources to curtail (by storing their energy when they would have curtailed) that is a channel that may add to the value of a battery fleet with a large presence of renewables, beyond what we find.

Figure 3: Renewable Energy, Degradation, and the Value of Batteries



(a) Battery Degradation and Battery Value (b) Price Uncertainty and Perfect Foresight

Notes: Each point in the scatter plot represents the lifetime profits for a unit of storage capacity based on market conditions during a single week of the sample (assuming conditions during that week repeated in perpetuity). The solid line plots the linear trend for each group. The profits are estimated assuming there are 10 MWh of aggregate operational storage capacity in the market. We rescale the estimated weekly storage value into a perpetuity using a 5% annual discount rate and adjusting for the rate of capacity degradation.

share was above 45%.

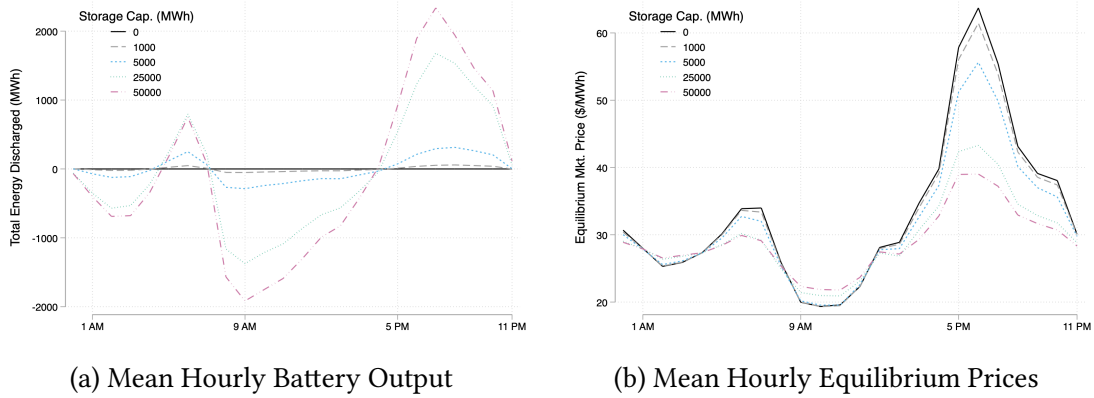
The solid blue line in Figure 3a highlights how capacity degradation (as discussed in Section 3.3) influences the estimated storage values. Degradation from cycling reduces the estimated value of storage investment by 27% on average. Moreover, the impact of degradation is higher with more renewable energy, which is due to batteries cycling more in this case. After accounting for degradation, the first battery unit would earn net profits in the energy market when renewable energy share is above 50.2% and capital costs are below \$264/kWh, as is expected to occur by 2024.³⁴ Alternatively, capital costs would have needed to be 40% lower for the first battery unit to be profitable in 2019. These findings emphasize the significance of accounting for degradation when measuring the value of storage.

Figure 3b compares our baseline storage value estimates—that assume battery operators face uncertainty about future wholesale prices—to the value estimates if battery operators have perfect foresight about future net load and electricity supply curve realizations.³⁵ Our model with uncertainty, which can be feasibly implemented by battery operators, achieves 70% of the theoretical maximum value under perfect foresight. Although our baseline results under uncertainty attain the majority of the perfect-

³⁴Renewable share from California's RPS (Figure A.2) and capital costs from Cole and Frazier (2019).

³⁵In both cases, we adjust the values to account for degradation.

Figure 4: Mean Battery Output and Equilibrium Prices Effects



Notes: Each line plots the mean counterfactual outcome across all days during 2016–19.

foresight value, they should be interpreted as a lower bound for storage value that could be further improved through better forecasting and modeling.

4.3 Equilibrium Effects of Battery Storage

We use our model to estimate the impact of battery operations on equilibrium prices. Figure 4a illustrates the mean simulated battery discharge quantity for each hour of the day for our evaluation sample, 2016–19. Each line in the figure shows battery output for a specific aggregate battery fleet capacity, K .

Across levels of K , batteries discharge the most during the hours where net load is the highest—the evening peak hours of 5–10 PM, but also discharge on average between 5–7 AM. As aggregate battery capacity grows, total discharges increase in the evening and total charges increase during the day.

Figure 4b shows that, as the fleet expands, battery operations exert a strong effect on lowering the variation in hourly mean equilibrium prices. Battery operations have the biggest impact on evening peak prices. Batteries have a relatively small effect on prices during the middle of the day, because the supply relationship is relatively flat during these hours.³⁶ Additionally, Figure 4b shows that the first few units of battery investment would have the largest impact on equilibrium prices, whereas incremental storage investment has a smaller impact on prices. The first batteries will reduce the

³⁶Figure A.6a in Online Appendix A focuses on the evening hours, showing that from 6–7 PM—the hours with the highest average net load—a modest 5000 MWh battery fleet would reduce average prices by over \$10 per MWh.

occurrence of extreme pricing events by discharging during periods when net load approaches the available generation capacity. By doing so, the batteries will reduce prices and also move the equilibrium to flatter regions of the supply relationship, thus reducing the marginal impact of subsequent battery entry on prices.

Table A.4 in Online Appendix A emphasizes this result. It shows that the first 5,000 MWh of storage capacity would reduce evening prices by 10.3% (\$54.25/MWh to \$48.67/MWh) and overall average price by over 5.6% (\$35.92 per MWh to \$33.90 per MWh). In contrast, an increase in capacity from 25,000 to 50,000 would only reduce evening prices by 7.3% (\$39.76/MWh to \$36.84/MWh) and overall mean prices by an additional 2.6% (\$31.02/MWh to \$30.20/MWh).


Figure A.6b in Online Appendix A demonstrates how battery operations would affect the mean generation from dispatchable power generators (e.g., natural gas generators) throughout the day. Unsurprisingly, large-scale storage increases dispatchable generator output during the middle of the day and reduces it in the evening peak hours. Notably though, batteries would also change the times of day with dispatchable generation troughs and peaks. With no battery capacity, the lowest production hour is 11 AM, whereas with a large battery fleet the lowest production period moves an hour later to noon. Similarly, the peak for dispatchable production without battery storage is 7 PM, relative to after 8 PM with a large storage fleet. These patterns demonstrate the importance of ramping costs in modeling storage operations. A competitive battery fleet reduces the *rate* at which dispatchable production increases, spreading the morning ramp down and evening ramp up over more hours.

To further understand how large battery fleets would optimally operate, Figure A.7 in Online Appendix A graphs real-time prices and battery operations for two arbitrarily selected days—June 23rd, 2016 and December 29, 2018—both for a 25,000 MWh capacity. Battery operations change discretely and abruptly during the day. On the left graph, batteries charge substantially in the morning before 9 AM, remain idle throughout the middle of the day, and then discharge at different points in time in the evening. On the right graph, prices are higher in the morning, causing batteries to discharge then. On both days, batteries reach approximately a full state of charge by mid-afternoon, wait several hours, and then discharge in the evening when real-time market prices spike. However, the two days differ in the times at which batteries start

charging and discharging. More generally, and consistent with Figure A.7, we find that (1) battery output at any time period varies considerably across days, but (2), on most days, batteries will fully charge prior to the evening ramp-up period and then wait to discharge until a price spike occurs.

As a result of highly volatile real-time prices, battery operations revenues are highly skewed across five-minute time intervals. From Table A.5 in Online Appendix A, batteries earn the majority of their revenues during the most profitable 1% of time intervals. For a 1000 MWh battery fleet, each 1 MWh of battery capacity would earn \$38,400 during the most profitable 1% of intervals and only \$17,293 across the other 99% of intervals over our sample period. Moreover, battery revenues are very sensitive to equilibrium effects. For instance, battery revenues during the most profitable intervals decline dramatically as aggregate battery capacity rises. For example, an increase in the battery fleet from 100 MWh to 10,000 MWh reduces per-unit revenues by nearly 28% during these intervals.

Figure 5: Battery Value by Aggregate Battery Capacity and Renewable Energy Share



Notes: The sloped lines plot the relationship between the expected lifetime value per kWh of battery investment and the share of renewable energy for selected aggregate battery capacity levels. They represent the best linear fit based off value and renewable energy across each week in our data. The gray horizontal line shows the expected capital cost of battery storage in 2024 (Cole and Frazier, 2019). The vertical line shows the total share of renewable energy (including hydro) based on data, the California RPS, and the authors' calculations.

These findings highlight the considerable decreasing returns to scale in battery storage capacity, which has important implications for the time path of battery invest-

ment. Figure 5 plots linear fit lines of the relationship between a heuristic of the value per unit capacity of a competitive battery market and the share of renewable energy during each week across three aggregate capacity levels—10 MWh, 10,000 MWh, and 50,000 MWh.³⁷

As in Figure 3, the per-unit value of a small battery fleet increases rapidly over time as more renewables enter the market, resulting in a 10 MWh fleet being profitable by 2024. Nevertheless, as more batteries enter the market, each battery’s value shifts downward due to the market equilibrium effects of operations of the preceding battery stock. For example, with a 50% renewable energy share, average battery value falls from \$280/kWh to \$230/kWh when aggregate capacity increases from 10 MWh to 10,000 MWh. These values fall further to \$140/kWh when there is 50,000 MWh of battery storage in the market. Because of equilibrium effects, storage fleets of even 10,000 MWh would not be profitable as arbitrageurs by 2024 without subsidies or unless capital costs were to fall far below current expectations.

4.4 Distributional Effects of Utility-Scale Batteries

Table 1 considers the impact of battery capacity additions on different market participants. Column 1 shows that with 1,000 MWh (1 GWh) of aggregate storage capacity in the market, batteries would have earned an average of \$14 million per year from operations during our sample period. As aggregate capacity increases to 50,000 MWh, the battery fleet flattens the price peaks, resulting in the average operating profits per unit capacity falling to \$4.48 million per GWh-year.

Column 3 indicates that batteries would significantly reduce the total expenditures (price \times load) that load-serving entities need to pay—to generators and storage operators—to meet demand. In particular, a 1,000 MWh battery fleet would reduce mean hourly expenditures for utilities by \$124 million per year.

Column 4 shows the change in dispatchable generators’ revenues with large-scale batteries. These track the expenditures to load-serving entities very closely. For instance, batteries would reduce total revenues of dispatchable generators substantially by \$126 million per year. These results suggest that large-scale battery adoption may

³⁷As in Figure 3, we use the weekly value simulated from the single-agent Bellman equation (2) (which mimics the incentives of the competitive battery market) scaled for degradation and discounting.

accelerate the retirement of dispatchable generators, a feature that we do not model.

Column 5 shows the change in solar and wind revenues with 1,000 MWh of batteries. Surprisingly, the presence of batteries *reduces* solar and wind generators' revenues by \$13 million annually. Although batteries increase prices between 9 AM and 1 PM when solar plants are coming online, they also reduce prices in the mid-afternoon (3 PM-5 PM) when many solar generators are still producing. Summing these impacts, battery operations make intermittent renewable generators slightly worse off.³⁸ These impacts are likely to occur in markets similar to CAISO, though they may not hold universally. The impact of batteries on renewables' profits will depend on the correlation between load and renewable generation across the day, among other factors.

Finally, column 6 investigates the impact of batteries on social surplus (gross of battery capital costs) in the electricity market. This exercise requires that we calculate costs, for which we leverage two additional assumptions: first, that our estimated supply relationships represent marginal costs (i.e., the dispatchable generation is competitively supplied), and second, that the fixed costs of producing 0 in any period are 0. Under these assumptions, a 1,000 MWh storage fleet would have increased gross social surplus by about \$14 million annually during our sample period.³⁹ A larger fleet with 50,000 MWh would have further reduced costs by \$351 million per year.

5 Evaluating Equilibrium Battery Adoption

Our results in Section 4.2 highlighted that a small battery fleet earning profits from electricity arbitrage was not far from breaking even by the end of our sample, while those from Section 4.3 showed that equilibrium effects will dampen the value of large-scale battery fleets. Nonetheless, even the Section 4.3 results do not speak to the equilibrium level of battery adoption, because the break-even constraint does not incorporate the opportunity cost of investment. Specifically, with declining capital costs, by waiting to adopt until *after* the break-even point, a potential operator will lower its expected adoption cost and potentially increase its value. The option value of waiting

³⁸Notably, these results contrast [Gonzales et al. \(2023\)](#), who find that transmission infrastructure investment led to more solar investment in Chile. Transmission investments help integrate renewables by allowing for additional spatial arbitrage whereas storage allows for arbitrage across time.

³⁹Since we assume that demand is perfectly inelastic, a change in gross social surplus is equal to the change in the total cost of electricity generation.

Table 1: Revenue and Costs Across Aggregate Battery Capacity Levels (MWh)

	Battery Profit per GWh Capacity	Total Battery Operation Profits	Δ Load Serving Entities' Expenditures	Δ Dispatchable Generator Revenue	Δ Solar and Wind Revenue	Δ Gross Social Surplus***
10	15.79	0.16	-1.76	-1.74	-0.18	0.15
1000	14.46	14.46	-123.98	-125.55	-12.88	13.96
5000	12.37	61.83	-436.48	-456.38	-41.87	65.52
25000	7.04	175.96	-1,062.34	-1,158.03	-80.13	242.57
50000	4.48	223.99	-1,239.99	-1,379.44	-84.39	351.09

Notes: All variables are annual means in millions of dollars per year over our sample period. Columns 1 and 2 show battery operations profits per unit (GWh) and in aggregate as a function of the total installed battery capacity. “ Δ Load Serving Entities’ Expenditures” is the change in the total price paid by load-serving entities for energy (change in equilibrium price times total load) relative to the $K = 0$ case. “ Δ Dispatchable Generator Revenues”, and “ Δ Solar and Wind Revenue”, are the mean change in annual gross revenues for dispatchable generators and renewable generators respectively. *** “ Δ Gross Social Surplus” is the estimated change in mean total costs of generation relative to the $K = 0$ case under the assumptions that the supply relationship represents marginal cost and that the fixed costs with no net load served are 0, and not accounting for battery capital costs.

will then delay equilibrium battery adoption.

This section develops an equilibrium adoption model that evaluates expected battery adoption rates under different policies, accounting for the opportunity cost of investment. As we detail below, our results here leverage assumptions beyond our operations model. This occurs because potential battery operators need to forecast their option value from waiting instead of adopting, which requires understanding future adoption capital costs and revenues. Additionally, our modeling framework is limited in that it does not consider dispatchable generator retirement, learning-by-doing causing battery capital cost reductions, or energy storage technologies other than lithium-ion batteries. We proceed by developing the modeling framework we use to understand adoption, explaining the calibration and estimation of the underlying parameters, and then presenting our results.

5.1 Model

Our capacity adoption model complements our operations model in Section 3.1 by considering *potential* battery operators at the *annual* level. We assume that there is an infinite mass of ex-ante identical potential battery operators, each of which has the ability to install a fixed-capacity storage system in one year. This capacity, which

we normalize to $k = 1$, is sufficiently small that the potential operator takes future electricity market prices as given in its adoption decision.

Potential operators are forward-looking and solve an optimal stopping problem of when to invest. At each year y , potential operators that have not previously adopted make a binary decision of whether or not to invest in storage capacity. To adopt, they must pay a fixed capital cost, c_y . At year y , agents observe c_y but do not know future adoption costs. We assume that these costs evolve stochastically as a Markov process based on current costs, declining over time in expectation due to technological advances. Agents have rational expectations over future adoption costs and, hence, form accurate distributions over cost trajectories.

Besides costs, a potential operator must also forecast the expected current and future revenues from its system for every future year. We model two important and counterbalancing factors regarding the future path of revenues. First, following our results in Section 4.2, the extra renewable energy capacity in future years will increase revenues. However, from Section 4.3, revenues will decline with greater equilibrium battery capacity, because large-scale storage will flatten equilibrium price peaks.

Thus, the annual per-unit revenues depend on both the year, y , which affects renewable energy generation share, and K , the aggregate capacity of storage present in the market. To simplify the analysis, we assume that potential operators perceive that, apart from these changes, the structural parameters of the operations model—i.e., the distributions of (gross) load by hour and the supply relationship from dispatchable generators—will remain constant in the future and hence are not state variables.

Combining these factors, the potential operator's state is (k, c, y, K) , where $k = 0$ for a potential operator that has not yet adopted and $k > 0$ for existing operators. We

write its Bellman equation as:

$$\begin{aligned}
\mathcal{V}(k, c, y, K) = & \mathbb{1}\{k = 0\} \\
& \left[\max \left\{ \overbrace{\pi(y, K^*) - c + \beta \int \mathcal{V}(\delta(y, K^*), c', y + 1, \delta(y, K^*) K^*) dG^{c'}(c'|c, y)}^{\text{Value from adopting}}, \right. \right. \\
& \left. \overbrace{\beta \int \mathcal{V}(0, c', y + 1, \delta(y, K^*) K^*) dG^{c'}(c'|c, y)}^{\text{Value from waiting}} \right\} \right] \\
& + \mathbb{1}\{k > 0\} \left[\underbrace{\pi(y, K^*)k + \beta \int \mathcal{V}(\delta(y, K^*) k, c', y + 1, \delta(y, K^*) K^*) dG^{c'}(c'|c, y)}_{\text{Value if adoption before } y} \right], \tag{9}
\end{aligned}$$

where $\pi(y, K^*)$ are annual operating profits. The second and third lines in (9) indicate the value to a potential operator that has not previously adopted, and reflect its optimal stopping problem. The last line, which is proportional to k , is the value to a potential operator that has previously adopted. After adoption, the operator makes no further adoption decisions, but the future distribution of battery installation costs will affect future adoption and, hence, future operating profits.

We microfound $\pi(y, K^*)$ using the computed values from the operations model. Specifically, we let:

$$\pi(y, K^*) = \sum_t E[p_t(\mathbb{1}\{q_t^* > 0\}q_t^*v + \mathbb{1}\{q_t^* < 0\}q_t^*/v)], \tag{10}$$

where the generator uses state-contingent optimal charge decisions, q_t^* . These are calculated using (1) at the expected equilibrium capacity K^* and state-contingent equilibrium prices, p_t^* .

We calculate capacity degradation, $\delta(y, K^*)$, from the Xu et al. (2016) engineering model as detailed in Online Appendix F. Degradation is a function of the state, since the state affects battery usage and this usage affects degradation. Because we assume exponential capacity degradation, all batteries at a given state will have the same incentives proportional to their capacity. Thus, we do not need to keep track of battery age as a state variable.

Similar to the operations model, we ease computation by recasting the battery operations problem in (9) as a single-agent decision problem where the incentives of the

single agent correspond to the incentives of the price-taking fleet of battery operators. In an equilibrium with price-taking potential battery operators, the marginal operator sets per-unit adoption cost equal to the marginal operating revenue net of the opportunity cost of adopting. Marginal operating revenue is composed of the weighted sum of prices over the year. The weights are determined by the charge quantity, q^* , which can be positive or negative. Thus, the corresponding single-agent maximization problem is as follows:

$$\begin{aligned} \mathcal{W}(c, y, K) = \max_{K^* \geq K} & \left\{ -E \left[\sum_t \int_0^{Z_t} P^{d(t)}(\zeta, \tilde{Z}_t, \varepsilon_t^P) \zeta \right] \right. \\ & \left. -c(K^* - K) + \beta \int \mathcal{W}(c', y+1, \delta(y, K^*) K^*) dG^{c'}(c'|c, y) \right\} \quad (11) \\ \text{s.t. } Z_t = & X_t^{RTM} + Q(q_t^*, K^*). \end{aligned}$$

In (11), the expectation is taken over the sequences of $(\varepsilon_t^P, \varepsilon_t^L)$ over the time intervals t during the year y and where $d(t)$ indicates the day corresponding to each time interval t , since the supply relationship parameters vary by d . Since the integral is taken up to $X_t^{RTM} - Q(q_t^*, K^*)$ —which is the electricity supplied by dispatchable generators—the first order condition with respect to K^* will weight $P^{d(t)}$ negatively in intervals where there are charges and positively in intervals where there are discharges.

We compute the solution to the adoption model by solving the single-agent adoption Bellman, equation (11). As with the operations model, in the case where the dispatchable generation market prices at marginal cost, the single agent problem is equivalent to the social planner problem, where the social planner minimizes the expected discounted costs of dispatchable generation plus storage capital costs.

5.2 Calibration and Estimation of Parameters

The main computational difficulty in solving the single-agent adoption Bellman equation is to evaluate the integral of expected operations revenues from across aggregate battery capacity states in (11), which we call batteries' *flow return*. The flow return is a function of the optimal charging behavior q^* , which varies based on aggregate battery capacity, K^* . In principle, for each state K^* that we reach in computing the

adoption Bellman equation solution, we could solve for optimizing behavior in the operations model at each time interval, and then plug in the resulting flow return into the adoption Bellman equation. However, this process would be very computationally intensive, especially because we allow the supply relationship parameters to vary across sample days.

In addition, the flow return is also a function of the year y . The calendar year affects the flow return because it affects renewable energy penetration which, per the Section 4.2 results, is complementary to the values batteries can earn.⁴⁰ Unlike with K^* , we do not develop a structural model of how increases in renewable energy penetration would affect the wholesale electricity price and through that, affect operations revenues, but instead, identify this effect from our in-sample variation in renewable energy generation share.

Given these issues, we follow Bod  r   (2022) and Gowrisankaran et al. (2022) and first evaluate the flow return across a fixed grid of states. We then estimate a regression of these flow return values on the state variables and treat the fitted value of this regression structurally. The benefit of this flow return surface approach is that it allows us to predict flow returns without computing the operations model Bellman equation for every state reached in the adoption model solution. The cost is that this approach puts functional form restrictions stemming from our regression on the flow return surface. These restrictions are an approximation to the true and unknown functional form implied by the structural model.

Specifically, we start by evaluating the flow return for eight different counterfactual values of K^* at the weekly level, all relative to $K^* = 0$.⁴¹ We evaluate the flow return over the week by simulating the optimized operations model.⁴² We then regress the flow return per unit capacity (flow return divided by K^*) on battery capacity (K^*), renewable energy generation share, controls for peak electricity demand, natural gas fuel prices, hydroelectricity availability (using the Sacramento Valley water-year index as a proxy), and week-of-year fixed effects. We use the fitted values from the regression—multiplied by K^* and scaled from the weekly to the annual level—as the

⁴⁰As noted above, an important limitation is that we do not allow the supply relationship to change across years y , implying that we are not allowing for dispatchable generator exit in response to greater renewable energy capacity.

⁴¹We use $K^* \in \{10, 100, 1000, 5000, 10000, 15000, 25000, 50000\}$.

⁴²Because we define simulated realized profits at the week level, our sample starts on Friday, Jan. 1, 2016 and ends on Thursday, Dec. 27, 2019. We assume that batteries start each week with 50% charge.

flow return for any state.

To obtain our fitted values, we map each calendar year into a renewable energy generation share that matches California’s legislated RPS schedule, interpolating in years where the RPS is not specified. This is appropriate if our estimated relationship between weekly renewable share on the flow return (conditional on controls) applies at the annual level.⁴³ We believe this assumption is reasonable given our inclusion of week-of-year fixed effects and other controls and that wind and solar production have both been increasing in California.⁴⁴ Our approach will capture the fact that changes in renewable energy production will indirectly affect the within-day net load variation, which will then affect storage systems’ profits. Finally, while our sample includes weeks with up to 50% renewable generation (see Figure 3), our flow return surface extrapolates out of sample above this level.

To solve the adoption model, we also need to estimate the state-contingent battery degradation over a year, $\delta(y, K^*)$, in (11). We estimate this function with the same methods as our estimation of the flow return function, except with the dependent variable being the annualized battery capacity degradation rate, as calculated by the Xu et al. (2016) engineering model.

Last, we calibrate the evolution of battery capital costs over time. We specify the following unit root with drift process for the cost of the storage technology, c_y :

$$c_y = c_{y-1} \exp(\tau) \exp(\xi_y), \quad \xi_y \sim N(0, \sigma_c^2), \quad (12)$$

with c_{2018} as the capital cost of batteries in 2018, the initial year, and τ and σ_c governing the size of the drift and future uncertainty of costs. To the extent that $\tau < 0$, the costs of storage will trend down over time on average. The ξ_y process captures the uncertainty about the size of these future cost declines. We assume that ξ_y are i.i.d. over time.

We estimate two parameters in (12): the magnitude of the downward drift (τ) and the size of the shock process governing the level of cost uncertainty (σ_c). Online Appendix G provides details of this estimation.

⁴³We use the 2019 sample mean values of the above controls for our fitted values.

⁴⁴California does not have separate wind and solar mandates.

Table 2: Battery Flow Return and Degradation by Year and Battery Capacity

	Battery Flow Return Per Unit Capacity (\$/kWh)		Annual Degradation Rate (%)	
	(1)	(2)	(3)	(4)
Constant	134.6*** (43.70)		0.4187*** (0.1529)	
$\ln(K^*)$	-2.832 (2.158)	-2.832 (2.195)	0.0623*** (0.0104)	0.0623*** (0.0106)
Renewable Share (%)	12.47*** (2.609)	10.04** (4.229)	0.0805*** (0.0085)	0.0705*** (0.0092)
$\ln(K^*) \times \text{Renewable Share (\%)}$	-0.6883*** (0.1298)	-0.6883*** (0.1321)	-0.0049*** (0.0006)	-0.0049*** (0.0006)
Observations	1,664	1,664	1,664	1,664
R ²	0.10533	0.41319	0.23211	0.52330
Within R ²		0.09888		0.10697
Controls + week fixed effects		✓		✓

Notes: In columns 1 and 2, the dependent variable is the annual flow return per kWh of storage capacity. Each observation represents a single week of the sample for a single storage capacity. In columns 3 and 4, the dependent variable is the annual capacity degradation due to operations. Columns 2 and 4 include controls for the mean load in the evening peak hours of 5–10 PM over the week, the mean natural gas price over the week, and the Sacramento Valley hydroelectric water year index (WYI) associated with that week. Peak load is the mean load between 5 PM and 9 PM during the week. We cluster standard errors by week of sample.

5.3 Adoption Model Results

Table 2 reports the estimates of our flow return regressions. Column 1 shows results from a specification that regresses the battery flow return on the logarithm of aggregate battery capacity ($\ln(K^*)$), renewable energy share (wind + solar share), and an interaction term. Column 2, our preferred specification, adds week-level controls for mean load in the evening peak hours, mean natural gas price, and the Sacramento Valley hydroelectric water year index (WYI), and week-of-year fixed effects.

The specifications with and without controls yield very similar results, adding to our confidence that the estimates are not being confounded by electricity market changes that are contemporaneous to renewable energy share changes. In our preferred specification, we estimate a negative though not statistically significant base coefficient on $\ln(K^*)$, a positive and significant coefficient on renewable share, and a negative and significant coefficient for the interaction term, consistent with the trends in Figure 3. Overall, our results paint a clear picture of the link between installed

battery capacity, renewable generation, and the value per unit of storage capacity. Per-unit storage value falls quickly as the aggregate storage capacity in the market rises, consistent with the equilibrium pricing impacts of storage we document in Section 4.3.⁴⁵

The third and fourth columns of Table 2 show the regression results with the annual battery degradation rate as the dependent variable.⁴⁶ The estimates in Column 3 indicate that when the solar and wind share equals 30%, and there is a single unit of storage in the market, capacity would degrade at an expected annual rate of 2.8% (2.7% for Column 4) due to cycling. The coefficient on the renewable energy share is positive: as renewable energy increases, the annual degradation rate also rises because batteries engage in more charge-discharge cycles.

Figure 6 provides simulated mean competitive equilibrium adoption paths under a variety of alternative assumptions, all using an annual discount factor of $\beta = 0.95$ and without an explicit battery mandate or subsidy. Throughout each panel of Figure 6, the solid black line shows the expected battery capacity trajectory under our baseline case, in which we assume that: battery capacity degrades as a function of use; potential adopters have rational expectations over future capital costs; renewable energy increases according to the California RPS; and peak load is held fixed at the 2019 mean level. The purple line in Figure 6a plots the expected battery capital cost over time from our estimated capital cost process. The solid black line shows that battery adoption begins very slowly, with the first storage system installed in 2026.⁴⁷ Total capacity reaches 288 MWh by 2030, before increasing sharply and achieving an aggregate capacity of 7,098 MWh by 2035.⁴⁸ A 7,000 MWh storage fleet composed of 4-hour duration batteries can produce 1,750 MW at any instant, similar to the typical output of a large nuclear power plant. This output could serve less than 10% of the typical CAISO load.

The remaining lines in Figure 6 explore several potential factors that may be limit-

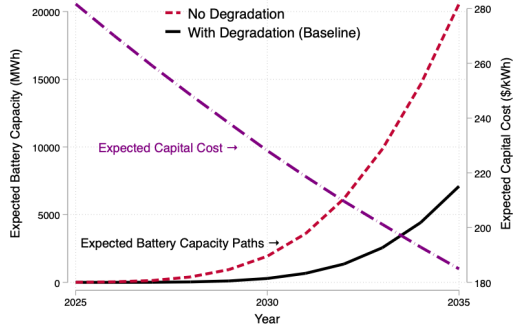
⁴⁵Table A.6 in Online Appendix A shows that the regression estimates are robust to alternative specifications and control variables.

⁴⁶In Table 2, we present results with annual battery degradation rate as the dependent variable for ease of interpretation. However, when implementing the adoption model, we run the same regression with weekly degradation rate as the dependent variable and then rescale the regression predictions to annual degradation rates.

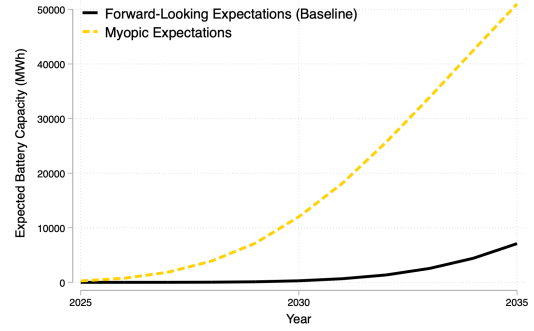
⁴⁷We find that there would be 0.91 MWh of storage in 2026 in expectation.

⁴⁸Because our model does not decompose wholesale electricity prices into costs and markups, we cannot determine whether the competitive battery market would have too much or too little entry.

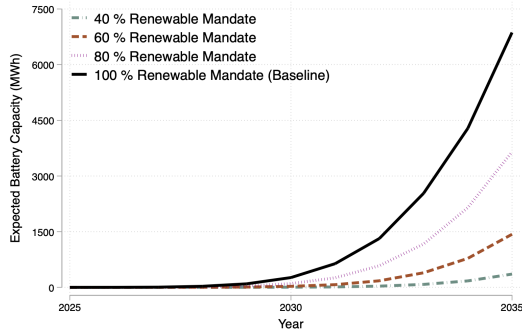
Figure 6: Counterfactual Battery Capacity Adoption Paths



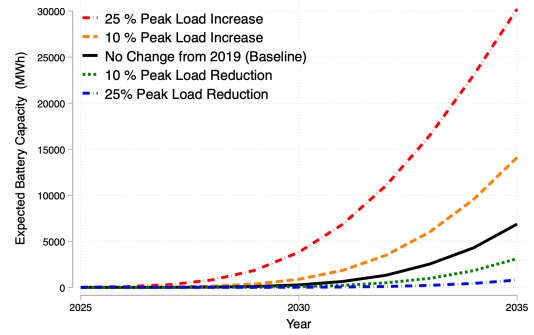
(a) Battery Capacity With vs. Without Degradation



(b) Myopic vs. Forward-Looking Expectations



(c) Renewable Mandates and Battery Capacity



(d) Peak Demand and Battery Capacity

Notes: In Figure 6a, the purple line shows the expected capital cost over time. In all figures, the solid black line plots expected battery capacity under the baseline case with: capacity degradation, forward-looking expectations, 100% RPS, and peak load held constant. The other lines plot expected battery capacity adoption under different counterfactuals. Each figure varies a single parameter, and holds all other assumptions fixed.

ing baseline equilibrium adoption. First, Figure 6a contrasts expected battery capacity over time without capacity degradation to the baseline. When we ignore degradation in calculating the value of storage, adoption starts one year sooner and increases at a much faster pace. In particular, the expected capacity would be roughly three times higher in 2035 (20,560 MWh).

Another factor that encourages potential battery adopters to delay investment is the anticipation of future capital cost reductions. Figure 6b quantifies the influence of future cost expectations on investment by calculating the predicted adoption path for myopic agents. While the forward-looking potential operators in our baseline know the parameters of the stochastic capital cost process in equation (12), myopic potential operators assume that the current capital cost will remain unchanged in future years,

but are otherwise identical to the baseline agents. Under myopic expectations, the first unit of battery investment is expected in 2023, with aggregate battery capacity reaching 12,000 MWh by 2030, and surpassing 50,000 MWh by 2035. These results are striking, as they indicate that expectations of future battery cost declines can play a major role in limiting early adoption.

Another key driver of the battery adoption decisions is the trajectory of future renewable energy generation. Figure 6c measures the effect of changing the renewable portfolio standard on the time path of battery adoption. Specifically, we plot the battery investment path for a 40% RPS by 2045, a 60% RPS by 2045, an 80% RPS by 2045, and a 100% RPS by 2045 (the current policy). With an RPS of 40%—a policy that would hold renewable generation constant at 2019 levels—less than 400 MWh of battery investment would occur by 2035. With the more aggressive renewable energy mandates, storage investment substantially increases: the 2035 expected storage capacity would be 1,430 MWh with the 60% RPS and 3,650 MWh with the 80% RPS.


Figure 6d explores how changes in future electricity load (demand) would change the time path of battery adoption. In our baseline case, Figure 6a, we assumed that peak load would remain constant at 2019 levels in all future years. However, California's peak load may change over time for a multitude of reasons. On the one hand, peak load could decrease over time due to energy efficiency retrofits and adoption of behind-the-meter renewable (e.g., residential solar panels) and storage technologies. On the other hand, rising adoption of electric vehicles could increase peak load if drivers plug in their cars during evening hours. Figure 6d illustrates how different assumptions about future peak load in California would change the trajectory of battery adoption. We evaluate expected battery adoption under five different cases: (1) 25% increase in peak load, (2) 10% increase in peak load, (3) no change in peak load (baseline), (4) 10% decrease in peak load, and (5) 25% decrease in peak load. We find that peak load changes can result in significant changes in expected battery investment. A 25% increase in peak load leads to a massive four-fold increase in capacity by 2035, whereas a 25% decrease in peak load reduces aggregate capacity by more than 80% relative to the baseline case.

These results show that utility-scale battery investment serves as a substitute for other investments that reduce peak load. For instance, energy efficiency retrofits

can reduce electricity demand at times of the day when the grid is most strained (Boomhower and Davis, 2020) while home batteries could also reduce peak household electricity demand. Accordingly, policies that encourage residential storage or energy efficiency investments would reduce the optimal capacity of utility-scale storage investment, while further investments in residential solar might complement them.

Finally, we use our results to evaluate policies. Thus far, our results suggest that a renewable portfolio standard alone is not sufficient to reach the amount of battery adoption stipulated in California’s 2024 battery mandate under AB 2514. Consequently, Figure 7 explores the impact of various government subsidies on battery adoption. Specifically, we compute the expected battery capacity in 2024 for different investment subsidies offered by the government ranging from 0%-40% of the capital cost. For each subsidy level, we assume that the subsidy is available to storage adopters in each year until 2024, and then no subsidy is available thereafter. Notably, we consider a 30% subsidy, similar to the energy storage investment tax credit offered by the 2022 U.S. Inflation Reduction Act (IRA).⁴⁹

Figure 7: Evaluating Battery Adoption Response to Subsidies



Notes: The blue line plots the total installed battery capacity in 2024 for differing levels of up-front subsidies (as a percentage of capital cost). The horizontal pink line indicates the California storage mandate under AB 2514 assuming 4-hour batteries. The vertical green line shows the 30% subsidy offered to storage under the 2022 U.S. Inflation Reduction Act.

Figure 7 shows that very little adoption would occur by 2024 with subsidies below

⁴⁹The IRA includes a 30% energy storage investment tax credit, available through 2025 (House of Representatives, 2022).

25%. However, the 2024 expected battery capacity increases substantially for subsidies ranging from 25-40%. Specifically, the vertical green line shows that the IRA subsidy would increase capacity to over 5,000 MWh. A larger 40% storage subsidy could further boost battery capacity to 25,000 MWh. We estimate that California’s battery mandate, which is equivalent to 5,200 MWh of storage capacity, could be implemented with a 30.4% up-front subsidy. Interestingly, this subsidy is very similar to the more recent 30% federal IRA subsidy that potential battery operators are earning throughout the U.S.

6 Impact of Battery Market Structure

Our results to this point have assumed that batteries are price takers. In this section, we investigate the role of battery market structure in determining battery operations and equilibrium prices. We consider two alternative battery market structures: monopoly and duopoly.

Unlike the competitive battery operator, a monopoly battery operator fully internalizes how an incremental charge (discharge) increases (reduces) equilibrium prices on inframarginal output. A duopoly battery operator partly internalizes this effect. Hence, the monopoly operator corresponds roughly to the worst case from the point of view of the system operator, while the duopoly is an intermediate case.

We let the monopolist make its operations decisions at each five-minute interval at the same point as competitive operations in our baseline model. Since the monopolist optimizes a single-agent dynamic problem, we solve its decisions using a similar approach to our baseline, except that we use operator profits as the flow payoffs.

In contrast, the duopoly model requires solving for the best response functions of a dynamic game. We assume the two battery operators play a Stackelberg-style game: each player chooses its output sequentially within each five-minute time interval and alternates as the first mover across time intervals. We chose a sequential-moves game because, in our finite horizon approximation, this game has a unique equilibrium. However, solving the duopoly model still adds two computational complications. First, it requires adding a state variable to track the second player’s energy held. Second, solving for the first mover’s optimal policy requires a conditional nested search over the second mover’s optimal policy given the first mover’s policy. These complications

substantially increase computational time. Therefore, we only solve the duopoly problem for a single aggregate battery capacity level, $K=10,000$ MWh, composed of two operators each with $K=5,000$ MWh.

Figure 8a shows how battery charge and discharge choices differ across battery market structures, holding aggregate capacity fixed at 10,000 MWh. The grey-dashed line indicates that duopolists tend to operate their batteries more sparingly compared to competitive batteries (solid black line)—charging less during the day and discharging less during evening peak hours. Monopolists, shown with the dotted-blue line, are the most conservative, charging and discharging even less than the duopolists. Intuitively, batteries with market power supply less during the morning and evening peak periods in order to maintain higher prices for sales of their additional stored energy.⁵⁰

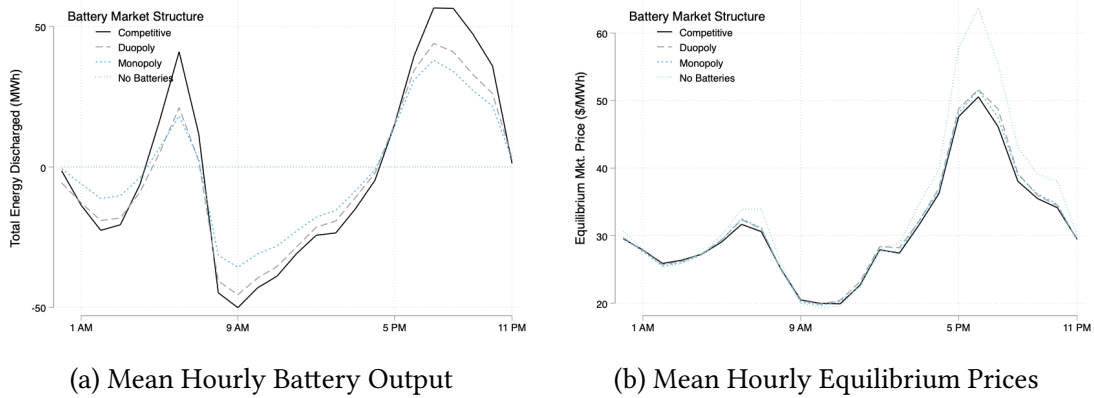
Figure 8b and the top panel of Table 3 illustrate that all three battery market structures lead to substantially lower prices relative to a market without batteries, especially during the evening peak periods. Without batteries, the load-weighted average price during peak hours is \$54.25/MWh. With competitive batteries, this price drops to \$45.04/MWh, highlighting the substantial impact of competitive battery operations in reducing peak hour prices. Both duopoly and monopoly prices during peak hours are slightly higher at \$46.71/MWh and \$46.09/MWh, respectively. Interestingly, monopoly prices during peak hours are slightly lower than duopoly prices. This result is due to the dynamic nature of dispatchable generator supply in our model. Monopolists tend to discharge their batteries less during peak hours compared to duopolists (Figure 8a). This leads to higher lagged net load, \tilde{Z} , during peak hours with monopolists, meaning more generators are running to meet the demand. When a large unexpected price shock occurs during peak hours, both monopolists and duopolists withhold output to maintain higher prices. However, because fewer generators are operating before the shock (higher \tilde{Z}), this can lead to higher relative price spikes under a duopoly than a monopoly.⁵¹ More broadly, we find that with 10,000 MWh of battery capacity, all battery market structures lead to an overall average price reduction between 7 to 9% relative to a market without batteries.

Table 3 Panel B investigates how equilibrium prices differ between a competitive

⁵⁰These findings are consistent with the theoretical results in [Andrés-Cerezo and Fabra \(2023\)](#) that market power in battery operations leads to lower storage utilization.

⁵¹Online Appendix Figure A.8 shows an example of a price spike event on the evening of June 15th, 2019, which leads to higher prices under a duopoly than a monopoly.

Figure 8: Battery Output and Equilibrium Prices by Battery Market Structure



Notes: Each line plots the mean counterfactual outcome across all days during 2016–19. All models assume 10,000 MWh of aggregate battery capacity.

and monopoly battery market across different levels of aggregate battery capacity. For a smaller aggregate battery capacity of 5,000 MWh, the average price for the competitive and monopoly markets are very similar at \$33.90/MWh and \$34.04/MWh, respectively. However, at the higher capacity level of 25,000 MWh, the competitive market achieves an average price of \$31.02/MWh, whereas the monopoly market results in a higher average price of \$33.87/MWh. Intuitively, monopolists behave similarly to price takers when they own small capacities. As capacity expands, monopolists begin to withhold more capacity at peak hours to sustain higher prices.

Altogether, this set of results indicates that battery operations under various market structures can lead to substantial price reductions relative to a market without batteries. In addition, we find that battery market power results in relatively small price distortions until aggregate capacity exceeds 10,000 MWh. This is interesting because it is possible for electricity suppliers to exercise market power at peak times even in relatively unconcentrated markets, due to residual demand being very inelastic at these times (Borenstein, 2002).

Table 3: Equilibrium Prices by Market Structure and Battery Capacity

Panel A: Average Prices (\$/MWh) by Market Structure, $K=10,000$ MWh

	Price (All hours)	Price (5-10 PM)
No Batteries	35.92	54.25
Competitive	32.70	45.04
Duopoly	33.33	46.71
Monopoly	33.06	46.09

Panel B: Average Prices (\$/MWh) by Aggregate Battery Capacity (MWh)

	Competitive	Monopoly
0	35.92	35.92
5000	33.90	34.04
10000	32.70	33.06
25000	31.02	33.87

Notes: This table presents the counterfactual (load-weighted) average real-time market price across 2016-2019. Panel A calculates average prices across all hours and for evening peak hours (5-10 PM) across the competitive, duopoly, and monopoly model. In Panel B, all models assume an aggregate battery capacity of 10,000 MWh. Panel B presents average prices across aggregate battery capacities for a competitive and monopoly battery market, respectively.

7 Conclusion

A significant challenge to meeting the world’s growing demand for energy is that utilities cannot typically store electricity for later use. As the majority of new renewable generation capacity comes from intermittent resources, the interest and potential role for battery storage technology has grown substantially. This paper develops a new framework to understand the equilibrium effects of large-scale battery storage and its complementarities with intermittent renewable energy. We model a number of features that we believe are critical to understanding the incentives to adopt and use storage and the value created by storage: the equilibrium price effects of large-scale battery capacity, dispatchable generator market power and ramping costs, and battery degradation from use. We estimate our model using data from California’s electricity market—which allows us to exploit variation in renewable energy generation over time—but our model can be applied to explore the economic impacts of storage in other markets and contexts.

We find that the equilibrium effects of batteries are large. The first 5,000 to 10,000 MWh of storage capacity will reduce peak hour prices significantly, but further in-

creases will have much smaller marginal impacts. The value that batteries can earn from energy market arbitrage is also significantly increasing in renewable energy penetration. Despite this, utility-scale storage in California is expected to reduce revenues for both dispatchable generators and renewable energy.

Finally, although we are currently not very far from a point where a small battery storage investment could break even in the energy market, utility-scale battery adoption would be limited in the absence of subsidies or mandates, due to the equilibrium effects and because of the option value of waiting for future capital cost declines. We estimate that the 2022 U.S. Inflation Reduction Act storage subsidy of 30% is roughly sufficient to implement California’s 2024 battery mandate of 5,200 MWh (1,300 MW). However, more ambitious policies to encourage large-scale storage will be substantially more costly. Moreover, our results indicate the investment in battery capacity is likely to be sensitive to future capacity costs, which may be affected by trade policies being considered ([The White House, 2024](#)).

While our analysis makes several contributions towards understanding the economics of battery storage investment, our modeling approach has several important limitations. First, we hold fixed the existing dispatchable generation capacity and the associated electricity supply relationship, even though our results imply that utility-scale batteries would lower dispatchable generator revenues and hence would likely lead to retirements. Relatedly, our analysis of the impact of battery market power assumes separate ownership among dispatchable generators and the battery operator(s). We believe that modeling endogenous dispatchable generator retirement and alternative ownership structures—including the potential for load serving entities to own batteries—are useful areas for further research. Second, while we model generator ramping costs, our supply relationship for generators in the presence of ramping costs and large-scale batteries is an approximation to a complex dynamic oligopoly problem. Third, we do not model the impact of storage on grid reliability, an abstraction that interacts with the potential exit of dispatchable generators’ capacity and one worthy of future research. Fourth, we assume that battery costs evolve exogenously, not allowing for battery mandates to lead to declines in production costs through learning-by-doing. Fifth, we use weekly variation in renewable energy over our 4-year sample period and extrapolate to predict the value of storage investment in a world where

more renewable generation exists than we can observe within our sample. Finally, we do not attempt to solve for the optimal storage subsidy to mitigate environmental externalities, given the complex interplay between a combination of mechanisms that incentivize both renewable energy and storage.

References

- Andrés-Cerezo, D. and Fabra, N. (2023). Storing Power: Market Structure Matters. *The RAND Journal of Economics*, 54(1):3–53.
- Bahn, O., Samano, M., and Sarkis, P. (2021). Market Power and Renewables: The Effects of Ownership Transfers. *The Energy Journal*, 42(4).
- Bodéré, P. (2022). Dynamic Spatial Competition in Early Education: An Equilibrium Analysis of the Preschool Market in Pennsylvania. Working paper.
- Boomhower, J. and Davis, L. (2020). Do Energy Efficiency Investments Deliver at the Right Time? *American Economic Journal: Applied Economics*, 12(1):115–39.
- Borenstein, S. (2002). The trouble with electricity markets: understanding california’s restructuring disaster. *Journal of economic perspectives*, 16(1):191–211.
- Borenstein, S., Bushnell, J. B., and Wolak, F. A. (2002). Measuring Market Inefficiencies in California’s Restructured Wholesale Electricity Market. *American Economic Review*, 92(5):1376–1405.
- Brave, S. A., Butters, R. A., and Kelley, D. (2021). A Practioner’s Guide and MATLAB Toolbox for Mixed Frequency State Space Modeling. *Journal of Statistical Software*, Forthcoming.
- Bresnahan, T. F. (1982). The Oligopoly Solution Concept Is Identified. *Economics Letters*, 10(1-2):87–92.
- Bushnell, J. and Novan, K. (2021). Setting With the Sun: The Impacts of Renewable Energy on Conventional Generation. *Journal of the Association of Environmental and Resource Economists*, 8(4):759–796.
- CAISO (2022). A Golden Age of Energy Storage. Available at <http://www.caiso.com/about/Pages/Blog/Posts/A-golden-age-of-energy-storage.aspx>.
- California Energy Commission (2019). Estimated Cost of New Utility-Scale Generation in California: 2018 Update. Available at <https://www.energy.ca.gov/sites/default/files/2021-06/CEC-200-2019-005.pdf>.
- California Secretary of State (2010). Energy Storage Systems. Assembly Bill No. 2514 (2009-2010). Skinner, Chapter 469, Statutes of 2010. Available at https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=200920100AB2514.
- California Secretary of State (2018). California Renewables Portfolio Standard Program: Emissions of Greenhouse Gases. Senate Bill No. 100 (2017-2018). De León, Chapter 312, Statutes of 2018. Available at https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=201720180SB100.
- Cole, W. J. and Frazier, A. (2019). Cost Projections for Utility-Scale Battery Storage. Technical report, National Renewable Energy Lab.(NREL), Golden, CO (United States).
- Compiani, G. (2022). Market Counterfactuals and the Specification of Multiproduct Demand: A Non-parametric Approach. *Quantitative Economics*, 13(2):545–591.

- Cullen, J. A. (2010). Dynamic Response to Environmental Regulation in the Electricity Industry. In *Industrial Organization Seminar*, volume 50.
- Cullen, J. A. and Reynolds, S. S. (2023). Market Dynamics and Investment in the Electricity Sector. *International Journal of Industrial Organization*, 89:102954.
- De Groote, O. and Verboven, F. (2019). Subsidies and Time Discounting in New Technology Adoption: Evidence From Solar Photovoltaic Systems. *American Economic Review*, 109(6):2137–72.
- Deaton, A. and Laroque, G. (1992). On the Behaviour of Commodity Prices. *The Review of Economic Studies*, 59(1):1–23.
- Durbin, J. and Koopman, S. J. (2012). *Time Series Analysis by State Space Methods: Second Edition*. Number 9780199641178 in OUP Catalogue. Oxford University Press.
- EIA (2020). Battery Storage in the United States: An Update on Market Trends. *Washington, DC: US EIA*.
- EIA (2021). Battery Storage in the United States: An Update on Market Trends. Technical report. Available at: https://www.eia.gov/analysis/studies/electricity/batterystorage/pdf/battery_storage_2021.pdf.
- EIA (2022a). Battery Systems on the U.S. Power Grid Are Increasingly Used to Respond to Price. Available at <https://www.eia.gov/todayinenergy/detail.php?id=53199>.
- EIA (2022b). U.S. Battery Storage Capacity Will Increase Significantly by 2025. Available at <https://www.eia.gov/todayinenergy/detail.php?id=54939#:~:text=As%20of%20October%202022%2C%207.8,GW%20of%20battery%20storage%20capacity>.
- Feger, F., Pavanini, N., and Radulescu, D. (2022). Welfare and Redistribution in Residential Electricity Markets With Solar Power. *The Review of Economic Studies*, 89(6):3267–3302.
- Fowlie, M., Reguant, M., and Ryan, S. P. (2016). Market-Based Emissions Regulation and Industry Dynamics. *Journal of Political Economy*, 124(1):249–302.
- Goldie-Scot, L. (2019). A Behind the Scenes Take on Lithium-Ion Battery Prices. *BloombergNEF*. Available at <https://about.bnef.com/blog/behind-scenes-take-lithium-ion-battery-prices/>.
- Gonzales, L. E., Ito, K., and Reguant, M. (2023). The Dynamic Impact of Market Integration: Evidence From Renewable Energy Expansion in Chile. Working paper, National Bureau of Economic Research.
- Gowrisankaran, G., Langer, A., and Zhang, W. (2022). Policy Uncertainty in the Market for Coal Electricity: The case of Air Toxics Standards. Working paper, National Bureau of Economic Research.
- Gowrisankaran, G., Reynolds, S. S., and Samano, M. (2016). Intermittency and the Value of Renewable Energy. *Journal of Political Economy*, 124(4):1187–1234.
- Graff Zivin, J. S., Kotchen, M. J., and Mansur, E. T. (2014). Spatial and Temporal Heterogeneity of Marginal Emissions: Implications for Electric Cars and Other Electricity-Shifting Policies. *Journal of Economic Behavior & Organization*, 107:248–268.
- Harvey, A. (1989). *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press.
- Holland, S. P., Mansur, E. T., and Yates, A. J. (2022). Decarbonization and Electrification in the Long Run. Working Paper 30082, National Bureau of Economic Research.
- House of Representatives (2022). Inflation Reduction Act of 2022. H.R.5376 - 117th Congress (2021-2022). Available at <https://www.congress.gov/bill/117th-congress/house-bill/5376>.


- Jha, A. and Leslie, G. (2021). Start-Up Costs and Market Power: Lessons From the Renewable Energy Transition. Working paper, SSRN 3603627.
- Joskow, P. L. (2011). Comparing the Costs of Intermittent and Dispatchable Electricity Generating Technologies. *The American Economic Review P&P*, 101(3):238–241.
- Kanamura, T. and Ōhashi, K. (2007). A Structural Model for Electricity Prices With Spikes: Measurement of Spike Risk and Optimal Policies for Hydropower Plant Operation. *Energy Economics*, 29(5):1010–1032.
- Karaduman, O. (2021). Economics of Grid-Scale Energy Storage in Wholesale Energy Markets. Working paper. Available at: <https://gsb-faculty.stanford.edu/omer-karaduman/files/2022/09/Economics-of-Grid-Scale-Energy-Storage.pdf>.
- Kirkpatrick, A. J. (2018). Estimating Congestion Benefits of Batteries for Unobserved Networks: A Machine Learning Approach. Working paper, Michigan State University.
- Knittel, C. R. and Roberts, M. R. (2005). An Empirical Examination of Restructured Electricity Prices. *Energy Economics*, 27(5):791–817.
- Kopecky, K. A. and Suen, R. M. (2010). Finite State Markov-Chain Approximations to Highly Persistent Processes. *Review of Economic Dynamics*, 13(3):701–714.
- Lamp, S. and Samano, M. (2022). Large-Scale Battery Storage, Short-Term Market Outcomes, and Arbitrage. *Energy Economics*, 107:105786.
- Langer, A. and Lemoine, D. (2022). Designing Dynamic Subsidies to Spur Adoption of New Technologies. *Journal of the Association of Environmental and Resource Economists*, 9(6):1197–1234.
- Liski, M. and Vehviläinen, I. (2020). Gone With the Wind? An Empirical Analysis of the Equilibrium Impact of Renewable Energy. *Journal of the Association of Environmental and Resource Economists*, 7(5):873–900.
- Mansur, E. T. (2008). Measuring Welfare in Restructured Electricity Markets. *The Review of Economics and Statistics*, 90(2):369–386.
- Mokrian, P. and Stephen, M. (2006). A Stochastic Programming Framework for the Valuation of Electricity Storage. Technical report, 26th USAEE/IAEE North American Conference.
- Pirrong, C. (2012). *Commodity Price Dynamics: A Structural Approach*. Cambridge University Press.
- Proietti, T. (2006). Temporal Disaggregation by State Space Methods: Dynamic Regression Methods Revisited. *The Econometrics Journal*, 9(3):357–372.
- Reguant, M. (2014). Complementary Bidding Mechanisms and Startup Costs in Electricity Markets. *The Review of Economic Studies*, 81(4):1708–1742.
- Rousseeuw, P. J. and Croux, C. (1993). Alternatives to the Median Absolute Deviation. *Journal of the American Statistical Association*, 88(424):1273–1283.
- Rust, J. (1987). Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher. *Econometrica*, pages 999–1033.
- Ryan, S. P. (2012). The Costs of Environmental Regulation in a Concentrated Industry. *Econometrica*, 80(3):1019–1061.
- Sackler, D. (2019). New Battery Storage on Shaky Ground in Ancillary Service Markets. *Utility Dive*. Available at <https://www.utilitydive.com/news/new-battery-storage-on-shaky-ground-in-ancillary-service-markets/567303/>.

- Schmidt, O., Melchior, S., Hawkes, A., and Staffell, I. (2019). Projecting the Future Levelized Cost of Electricity Storage Technologies. *Joule*, 3(1):81–100.
- Sioshansi, R., Denholm, P., Jenkin, T., and Weiss, J. (2009). Estimating the Value of Electricity Storage in PJM: Arbitrage and Some Welfare Effects. *Energy Economics*, 31(2):269–277.
- Spector, J. (2020). LS Power Energizes World’s Biggest Battery, Just in Time for California’s Heat Wave. *Greentech Media*. Available at <https://www.greentechmedia.com/articles/read/ls-power-energizes-worlds-biggest-battery-near-san-diego-just-in-time-for-heatwave>.
- Tauchen, G. (1986). Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions. *Economics Letters*, 20(2):177–181.
- The White House (2024). Biden-Harris administration takes action to strengthen American solar manufacturing and protect manufacturers and workers from China’s unfair trade practices. Fact sheet, The White House.
- U.S. Department of Energy (2021). Secretary Granholm Announces New Goal to Cut Costs of Long Duration Energy Storage by 90 Percent. Available at <https://www.energy.gov/articles/secretary-granholm-announces-new-goal-cut-costs-long-duration-energy-storage-90-percent>.
- Weron, R. (2014). Electricity Price Forecasting: A Review of the State-of-the-Art With a Look Into the Future. *International Journal of Forecasting*, 30(4):1030–1081.
- Wolak, F. A. (2018). The Evidence From California on the Economic Impact of Inefficient Distribution Network Pricing. Working paper, National Bureau of Economic Research.
- Wolfram, C. D. (1999). Measuring Duopoly Power in the British Electricity Spot Market. *American Economic Review*, 89(4):805–826.
- Xu, B., Oudalov, A., Ulbig, A., Andersson, G., and Kirschen, D. S. (2016). Modeling of Lithium-Ion Battery Degradation for Cell Life Assessment. *IEEE Transactions on Smart Grid*, 9(2):1131–1140.
- Zhou, Y. and Li, S. (2018). Technology Adoption and Critical Mass: The Case of the US Electric Vehicle Market. *The Journal of Industrial Economics*, 66(2):423–480.

Online Appendix

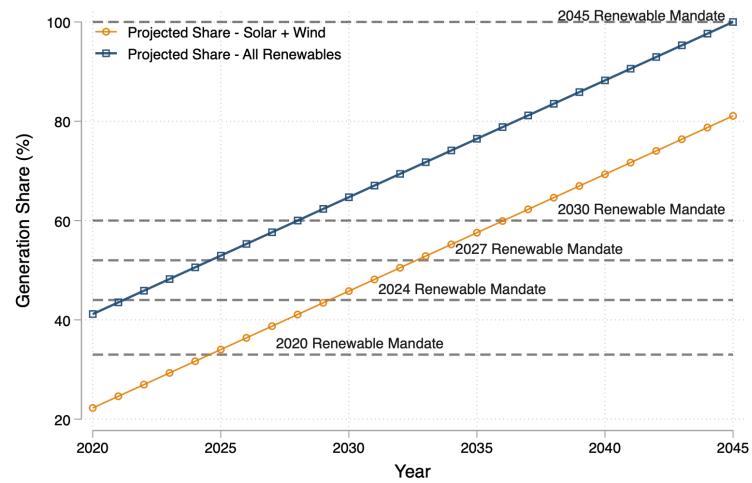
A Additional Tables & Figures Referenced in Main Paper

Figure A.1: Regulation Service Quantity Procured by CAISO




Notes: The figure plots the mean hourly quantity of regulation services procured by CAISO each month. Regulation quantity is calculated as the sum of “regulation up” and “regulation down” quantities in the day-ahead market.

Figure A.2: Renewable Energy Over Time Under the California Renewable Portfolio Standard




Notes: Each horizontal line shows the share of generation that must come from renewable sources in a particular year under the California RPS. The “All Renewables” line shows our linear interpolation of the California RPS. The “Solar + Wind” line shows the assumed path of solar and wind generation in future years.


Figure A.3: CAISO Electricity Market Trends




(a) Load



(b) Solar PV Share




(c) Wind Share



(d) Solar + Wind Share




(e) Natural Gas Price (\$/mmbtu)



(f) RTM Price (\$/MWh)

Notes: Each panel plots the weekly average of a given single variable over the sample period. The solar generation measure does not include distributed generation. The reported market prices are for the CAISO South Zone Trading Hub (SP 15).

Figure A.4: Real-Time Market Prices (Five-Minute Frequency)



Notes: Figure shows the average real-time market price (South Hub - SP-15) for each five-minute interval of the day, separately for 2015 and 2019.

Figure A.5: Operations Model (Single Day)

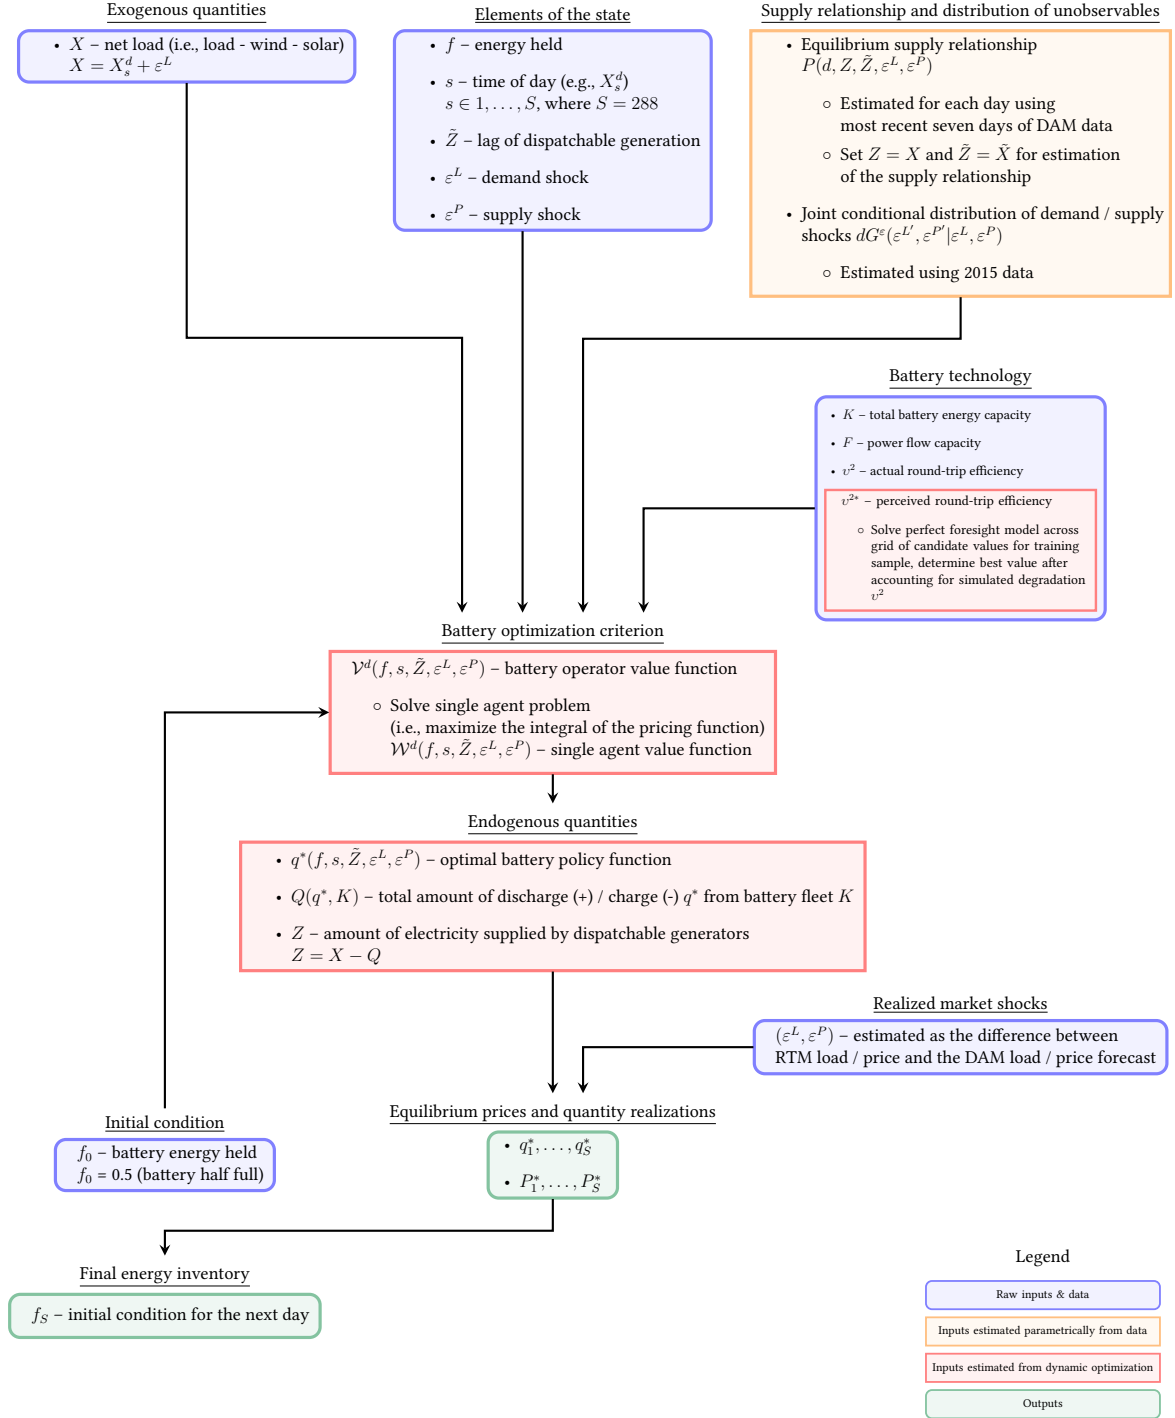


Table A.1: Summary Statistics for Estimated Supply Relationship Parameters

Parameter	2015	2016	2017	2018	2019	2016–19
θ_1						
Mean	-6.45	-27.24	-16.52	-11.71	-10.65	-16.54
Std. Dev.	8.82	25.62	19.50	14.32	14.92	20.21
25th-percentile	-5.20	-52.95	-29.78	-12.16	-9.62	-21.27
75th-percentile	-2.79	-4.80	-1.98	-3.27	-2.77	-2.99
θ_2						
Mean	18.39	161.30	93.27	50.95	45.55	87.82
Std. Dev.	53.17	192.24	144.39	107.31	107.98	149.54
25th-percentile	1.42	3.96	0.81	1.71	1.01	1.47
75th-percentile	15.67	365.00	135.41	23.13	12.96	63.66
θ_3						
Mean	2.07	1.37	1.47	1.16	1.07	1.27
Std. Dev.	1.35	0.89	1.02	0.58	0.39	0.78
25th-percentile	1.01	1.01	1.01	1.01	1.01	1.01
75th-percentile	3.81	1.01	1.01	1.01	1.01	1.01
κ						
Mean	2.18	4.29	3.41	2.67	2.46	3.21
Std. Dev.	1.11	2.82	2.69	2.17	2.07	2.56
25th-percentile	1.42	1.75	1.25	1.38	1.24	1.35
75th-percentile	2.55	8.00	5.46	2.70	2.31	3.87
α						
Mean	0.82	0.90	0.85	0.85	0.83	0.86
Std. Dev.	0.08	0.09	0.15	0.12	0.12	0.12
25th-percentile	0.76	0.83	0.76	0.80	0.73	0.79
75th-percentile	0.88	0.97	0.97	0.95	0.94	0.97
R-squared						
Mean	.	0.85	0.86	0.86	0.85	0.86
Std. Dev.	.	0.05	0.06	0.06	0.10	0.07
25th-percentile	.	0.82	0.84	0.84	0.84	0.84
75th-percentile	.	0.89	0.89	0.89	0.90	0.90

Notes: This table summarizes the means, standard deviations, and 25th and 75th percentiles of the daily estimated supply relationship parameters, and the r-squared for the analysis sample.

Table A.2: Summary Statistics for Estimated Supply Relationship Residuals

	2015	2016	2017	2018	2019	2016–19
Dependent Variable: ε_t^P						
ε_{t-1}^P	0.947*** (0.013)	0.849*** (0.015)	0.897*** (0.017)	0.832*** (0.028)	0.839*** (0.019)	0.861*** (0.013)
Constant	0.005*** (0.001)	0.004*** (0.000)	0.007*** (0.001)	0.010*** (0.002)	0.008*** (0.001)	0.007*** (0.001)
$\sigma^{P, \text{Peak}}$	0.012	0.010	0.010	0.016	0.016	0.013
$\sigma^{P, \text{Off-peak}}$	0.010	0.006	0.008	0.012	0.014	0.010
Observations	17568	105408	105120	105120	105120	420768

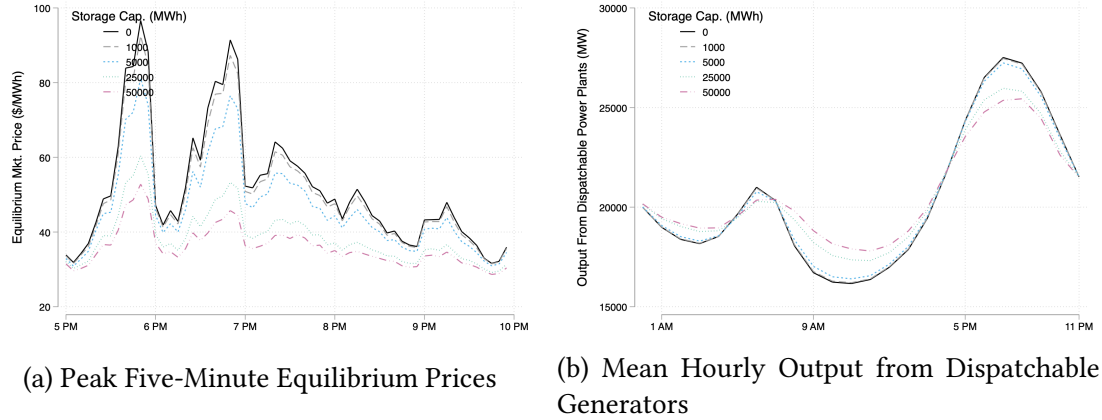
Notes: This table summarizes the estimates of the supply relationship residual (ε_t^P) parameters. The 2015 sample includes only November and December. We report standard errors, clustered by day-of-sample, in parentheses.

Table A.3: Summary Statistics for Estimated Net Load Model

	2015	2016	2017	2018	2019	2016–19
(a) Dependent Variable: Net Load _t						
Net Load DAM Forecast	0.969*** (0.003)	0.950*** (0.002)	0.950*** (0.001)	0.971*** (0.001)	0.955*** (0.002)	0.956*** (0.001)
Dependent Variable Mean	1794.61	1798.35	1734.13	1687.41	1599.83	1704.99
In-sample RMSE	67.721	83.007	77.494	74.292	80.513	80.511
(b) Dependent Variable: ε_t^L						
ε_{t-1}^L	0.996*** (0.001)	0.996*** (0.000)	0.996*** (0.000)	0.995*** (0.000)	0.995*** (0.000)	0.996*** (0.000)
Constant	0.144** (0.043)	-0.014 (0.016)	-0.023 (0.016)	0.178*** (0.023)	0.017 (0.021)	0.032*** (0.009)
σ^L	6.426	7.110	7.245	7.591	8.131	7.530
Observations	17568	105408	105120	105120	105120	420768

Notes: This table summarizes the estimates of the net load model. The 2015 sample, which is used to obtain the parameters of the AR(1) process, includes only November and December. We report standard errors, clustered by day-of-sample, in parentheses.

Figure A.6: Equilibrium Prices Effects and Dispatchable Generator Output




Notes: Each line plots the mean counterfactual outcome for specific storage capacity level across all days during 2016–19.

Table A.4: Equilibrium Prices and Aggregate Battery Capacity

	Price (All hours)	Price (6-9 AM)	Price (10 AM - 3 PM)	Price (5-10 PM)
0	35.92	31.44	25.15	54.25
10	35.91	31.43	25.15	54.23
100	35.84	31.39	25.13	54.04
1000	35.35	31.14	25.01	52.75
5000	33.90	30.33	24.88	48.67
10000	32.70	29.52	24.90	45.04
15000	31.96	29.03	24.95	42.79
25000	31.02	28.62	25.03	39.76
50000	30.20	28.61	25.42	36.84

Notes: Prices reported are in \$/MWh and indicate the load-weighted mean across all five-minute intervals between 2016–19.

Figure A.7: Battery Operations on Selected Days



Notes: The black lines show the observed real-time market price in the absence of battery operations. The orange lines show the equilibrium prices after incorporating storage operations. The green lines in both show the simulated amount of energy held in storage (i.e. the stock) as a percentage of energy capacity on June, 23, 2016 and December, 29, 2018. The simulations assume an aggregate storage capacity of 25,000 MWh.

Table A.5: Skew in Distribution of Battery Operating Profits Across Time Periods

	Time Periods - Other Percentiles	Time Periods - 99th Percentile
Battery Capacity in MWh: 10	17,764.74	40,679.28
Battery Capacity in MWh: 100	18,707.89	41,452.05
Battery Capacity in MWh: 1000	17,292.98	38,399.72
Battery Capacity in MWh: 5000	17,161.93	35,114.04
Battery Capacity in MWh: 10000	16,433.17	32,413.34
Battery Capacity in MWh: 15000	14,961.54	30,118.45
Battery Capacity in MWh: 25000	12,145.14	26,559.53
Battery Capacity in MWh: 50000	7,388.77	20,621.49


Notes: The first column lists the aggregate battery capacity. The second column indicates the total revenue a battery owner would earn between 2016–19 summed over the least profitable 99 percent of time periods. The third column lists the total revenue a battery owner would earn summed over the most profitable 1 percent of time periods. All numbers are in \$/MWh of capacity.

Table A.6: Robustness Checks: Battery Flow Return Regressions

	Battery Flow Return Per Unit Capacity (\$/kWh)			
	(1)	(2)	(3)	(4)
$\ln(K^*)$	-2.832 (2.195)	-2.832 (2.195)	-2.832 (2.196)	-2.832 (2.195)
Renewable Share (%)	10.04** (4.229)	9.857** (4.542)	7.853 (4.779)	17.17 (14.41)
$\ln(K^*) \times \text{Renewable Share (\%)}$	-0.6883*** (0.1321)	-0.6883*** (0.1321)	-0.6883*** (0.1321)	-0.6883*** (0.1321)
Peak Load (Mean)	0.1573* (0.0878)		0.4517** (0.2268)	
Load (Mean)		0.1760 (0.1258)		
Off-Peak Load (Mean)			-0.4900 (0.3610)	
$(\text{Renewable Share})^2$				-0.2735 (0.3318)
Observations	1,664	1,664	1,664	1,664
R ²	0.41319	0.41069	0.41648	0.40806
Within R ²	0.09888	0.09504	0.10393	0.09100
Controls + week of year fixed effects	✓	✓	✓	✓

Notes: The dependent variable is the annual flow return per kWh of storage capacity. Each observation represents a single week of the sample for a single storage capacity. All columns include controls for the mean natural gas price over the week and the Sacramento Valley hydroelectric water year index (WYI) associated with that week. Peak load is the mean load between 5 PM and 9 PM during the week; off-peak load is the mean load at all other times. We cluster standard errors by week of sample.

Figure A.8: Battery Output and Equilibrium Prices on June 15, 2019



Notes: Each line plots a counterfactual outcome on June 15, 2019. All models assume 10,000 MWh of aggregate battery capacity.

B Battery Market Structure

This appendix documents the evolving industry structure of battery storage in the California electricity market over the past several years. Before 2020, most battery storage projects were small in scale (e.g., less than 40 MW).⁵² This changed starting in 2020, with the development of the Gateway Energy Storage System in California, which has 250 MW of capacity battery storage.⁵³ Following this trend, Pacific Gas and Electric (PG&E)—California’s largest investor-owned utility—unveiled the Moss Landing site for battery storage in collaboration with Tesla in June 2022. The Moss Landing facility composed 182.5 MW of the 955.5 MW total storage capacity operated by PG&E. With the Moss Landing Battery Storage Project beginning operations in June of 2022, California’s Independent System Operator (CAISO) had just over 3,160 MW of battery storage capacity, with an additional 700 MW of planned storage capacity scheduled to come online later that month (CAISO, 2022).

Table B.1: Industry Structure of the CAISO Battery Market

	2018	2020	2022
Number of Entities	17	28	70
Total Capacity (MW)	233.3	528.9	4737.8
Top 4 Share	67%	71%	29%
HHI	1347	2522	432
Avg. Capacity (MW)	13.7	18.9	67.7

Notes: Calculations by authors from EIA Form 860. Sample includes all operating battery plants in California in each of the respective years. Market shares based on capacity.

Table B.1 provides some descriptive statistics of the industry structure of California’s battery market between 2018 and 2022. We calculate the statistics using the entity-level capacity information provided by the Energy Information Administration (EIA) in Form 860. This table indicates several patterns. First, the growth in battery capacity from 2018 to 2022 was substantial. For instance, in 2018, the total amount of battery capacity operating in California was negligible, amounting to less than 240 MW. But, battery capacity grew by nearly 2000% over this time frame. Second, there

⁵²For an overview of the growth of battery storage projects, see EIA (2022b).

⁵³The Gateway Energy Storage System is managed and operated by LS Power, which also owns the 40MW Vista Project which came online in 2018 (Spector, 2020).

has been a sizable uptick in the number of firms operating battery storage facilities in California from 17 firms in 2018 up to 70 firms in 2022.

Third, battery market concentration fell markedly from 2020 to 2022 (as measured by capacities) in terms of both the top-four share and the Herfindahl-Hirschman Index (HHI). In the earliest years of battery entry to the California market, total capacity was relatively small and the four largest operating companies owned 71% of battery capacity, implying a battery market HHI of 2,522. In contrast, many new operating companies entered the battery market in 2022, which led to both a large increase in the market's total capacity and a major reduction in the concentration of ownership. Specifically, the combined market share of the four largest operating fell to just 29% and the market HHI dropped to 432 in 2022.

While the statistics presented above are useful in characterizing how the industry structure has evolved, there are a couple of important caveats.

First, the reported market shares and HHIs are based on operating companies' reported names on EIA Form 860. In our calculations, we assume that a different operating company name implies a different firm, but we cannot rule out that unique operating companies may be owned by a common parent company. Second, because the battery market is changing rapidly over time, the current market structure does not necessarily guarantee that such trends will continue in the future.

C Robustness to Supply Relationship Functional Form

We explore the robustness of our results to our chosen functional form for the supply relationship. Our main results are based on the [Pirrong \(2012\)](#) model, which has been used in the commodity storage literature. We re-estimate our model using a functional form for the supply relationship based on the cost function in [Ryan \(2012\)](#) and [Fowle et al. \(2016\)](#) (henceforth, R/FRR).⁵⁴ R/FRR used this functional form to estimate costs for cement plants, noting that this cost function accounts for increasing costs near capacity, which gives the function the “hockey stick” shape common in the electricity generation industry” ([Ryan, 2012](#), p. 1029).

⁵⁴ Another alternative to estimate the supply relationship would be to use generator-level data on heat rates and capacities to infer a market-level dispatch curve using a merit-order approach. We found this approach to be inferior in explaining the behavior of electricity prices in the wholesale markets. Online Appendix [H](#) provides more details on this point.

In our case, we estimate a supply relationship and not a cost function. From Section 3.2, we require that the supply relationship define a unique ε_t^P for any observed price, as in (6). We define a supply relationship based on the R/FRR cost function that is strictly increasing in capacity utilization:

$$\tilde{P}^d(Z|\mathcal{K}, \mathcal{K}) = \theta_4 + \theta_5 Z/\mathcal{K} + \theta_6 \mathbb{1}\{Z/\mathcal{K} > \nu\} (Z/\mathcal{K} - \nu)^2, \quad (\text{C.1})$$

where $\nu \in (0, 1)$, $\theta_4, \theta_5 (> 0)$, and $\theta_6 (> 0)$ are parameters to estimate. The parameter ν represents the point at which the pricing surface starts to bend from linearly increasing in capacity utilization to quadratically increasing. We proceed by estimating our entire model using the supply relationship motivated by R/FRR instead of Pirrong.⁵⁵

Table C.1 compares several of our main results to the results using the R/FRR functional form supply relationship. Each panel summarizes a different aspect of the effects of battery capacity. Broadly speaking, we find that most results implied from the alternative functional form are similar to those from our base model. In Panel A, we see that the R/FRR functional predicts slightly muted equilibrium price effects relative to the base model. Specifically, the base model with 50,000 MWh of battery capacity predicts mean peak prices of \$35.96/MWh versus \$40.30/MWh with the R/FRR functional form. Panel B shows that our estimates of the expected lifetime revenues per unit of battery capacity are relatively similar across the two models. Panel C illustrates that with 50,000 MWh of battery capacity, the R/FRR functional form yields slightly lower estimates for battery profits and also that storage operations would have a smaller impact on both dispatchable generators' revenues and renewable generators' revenues relative to the base model. Finally, Panel D shows that predicted battery adoption between the years 2024 to 2030 is remarkably similar across the two specifications, although we see slightly less adoption under the R/FRR functional form.

Importantly, both functional forms require a similar assumption that the deviations in prices that occur between the real-time market and the day-ahead market reveal changes in available capacity or transmission. However, these results indicate that our

⁵⁵Similar to our approach with the Pirrong (2012) form, we use non-linear least squares to estimate the R/FRR supply relationships, with parameters restrictions as follows: $\theta_4 \in [-700, 500]$, $\theta_5 \in [0, 500]$, $\theta_6 = [0, 100]$, $\nu \in [0, 1]$, $\kappa \in [1.01, 4]$, $\alpha \in [0, 1]$. Unlike with Pirrong, the real-time supply relationship in (C.1) does not asymptote to $P^d = \infty$ at \mathcal{K} . In fewer than 1% of cases with high RTM prices, the observed RTM price implies $Z/\mathcal{K} > 1$. We simply use these prices, rather than restricting (C.1) to $Z/\mathcal{K} = 1$.

baseline results are unlikely to be specific to a functional form choice, a consequence of our flexible approach to estimating supply relationships that vary by sample day.

Table C.1: Robustness of Results to Functional Form

Panel A: Mean Peak Prices by Aggregate Battery Capacity (\$/MWh)

	Base Model	R/FRR Functional Form
1000 MWh	50.44	51.74
10000 MWh	43.57	48.65
50000 MWh	35.96	40.30

Panel B: Present Value of Expected Lifetime Battery Revenues per Unit by Aggregate Battery Capacity (\$/kWh)

1000 MWh	204.83	192.58
10000 MWh	149.21	176.90
50000 MWh	67.80	93.37

Panel C: Annual Operating Profits/Revenues for $K^* = 50,000$ MWh (\$1M)

Battery Profit per GWh Capacity	5.34	4.71
Δ Dispatchable Generator Revenue Relative to $K=0$	-1,392.97	-960.74
Δ Solar and Wind Revenue Relative to $K=0$	-86.69	-33.13

Panel D: Cumulative Battery Adoption by Year (MWh)

2024	0.00	0.00
2026	0.91	0.00
2028	28.99	13.54
2030	263.33	228.89

Notes: Column 1 summarizes key results for our base model that uses the [Pirrong \(2012\)](#) functional form for the supply relationship. Column 2 reports analogous results from an alternative functional form (R/FRR) for the supply relationship based on [Ryan \(2012\)](#) and [Fowlie et al. \(2016\)](#). Panel A compares the counterfactual equilibrium prices across the specifications and calculates peak prices as the mean predicted price between 5 PM and 9 PM across our sample, 2016-2019. Panel B reports the expected lifetime revenues per unit of battery capacity for different aggregate capacity levels. The estimates in Panel B assume that the grid conditions in our sample persist forever (e.g., renewable energy is held constant). Panel C calculates the change in market participants' expected annual profits or revenues during our sample period with 50,000 MWh of battery storage compared to the market with no storage. Panel D uses the adoption model to find the expected cumulative battery capacity over time across the model specifications.

D Details of Supply Relationship Estimation

For each sample day, d , we estimate the supply relationship parameters using net load and price data from the day-ahead market (DAM) over the previous week. Variation in these parameters across sample days may be caused by shifts in natural gas prices, changes in the availability of low-cost generation coming from nuclear power plants and hydroelectric sources, as well as day-to-day changes in generator availability and imports and exports from neighboring states. By using the DAM to estimate the marginal cost curve, our approach allows us to account for market characteristics that vary at a high frequency, while ensuring that our dynamic operations model remains feasible in that it only uses information that would be available to a storage operator in bidding in the real-time market.

Turning to specifics of the estimation of the supply relationship given in (4), we facilitate estimation by standardizing each day's DAM prices and net load forecasts. For the DAM prices, we subtract the median and divide by the interquartile range over the sample window. For net load, we divide by the maximum of that sample window's net load forecast. Finally, we restrict the parameter domain, Θ , to be such that $\theta_1 \in [-700, 500]$, $\theta_2 \in [0, 500]$, $\theta_3 \in [1.01, 4]$, $\kappa \in [1, 8]$, $\alpha \in [0, 1]$.⁵⁶ These restrictions ensure that the supply relationship is monotonically increasing in \tilde{Z} for $Z \leq \mathcal{K}$.

Turning to the structural unobservable (ε^P), conditional on a set of supply relationship parameters for any particular day, we recover a time series of ε_t^P as the shocks required to rationalize the RTM price observed at time t with the realizations of net load and lagged net load. At time t , we obtain:

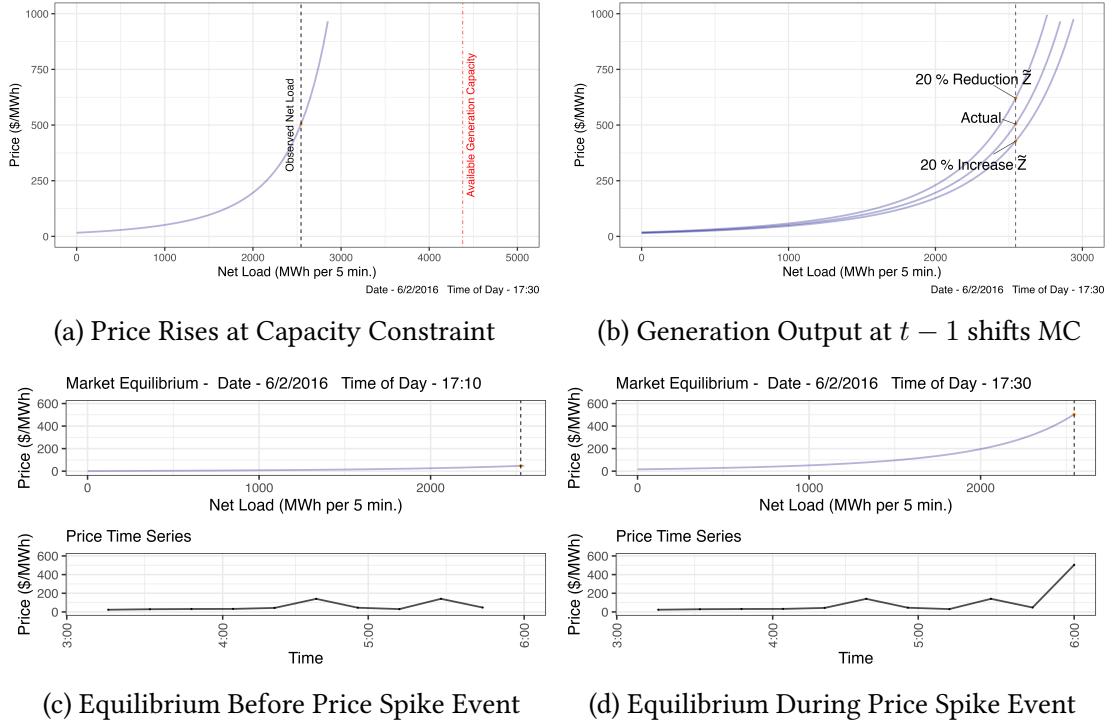
$$\varepsilon_t^P = \ln \left[Z_t + \left(\frac{P_t^{RTM} - \theta_1}{\theta_2} \right)^{-1/\theta_3} \right] - \ln \left[\kappa^\alpha \tilde{Z}_t^{1-\alpha} \right], \quad (\text{D.1})$$

where we use the sample day d estimated values of (θ, κ, α) .

As an example of the features of our approach towards modeling the supply relationship, Figure D.1 provides the supply relationship on June 2, 2016, when net load was approaching the constraint on available generating capacity. From Figure D.1a, at 5:30 PM, the market equilibrium was near an inflection point: an increase in net load

⁵⁶We also compute a perfect foresight model, which uses the same marginal cost curve parameters.

Figure D.1: Time-Varying Supply Relationship Curve



Notes: This figure displays supply relationships for June 2, 2016. Figure D.1a shows the market equilibrium and the implied generation capacity available for a single five-minute interval. Figure D.1b shows how 20% changes in last period's dispatchable generation would shift the supply relationship. Figures D.1c and D.1d show how both the net load and the supply relationship shifts during a period when price increased rapidly over a 20-minute span.

would significantly raise equilibrium price, while a decrease in net load would have a smaller effect in decreasing price. Figure D.1b illustrates the importance of ramping costs in our model. At this same time, a 20% decrease in generation from fossil fuel generators in the previous period (\tilde{Z}) would lead to a substantial price increase, with a smaller price decrease from a 20% increase in \tilde{Z} .

Figures D.1c and D.1d illustrate how our model rationalizes a rapid change in price that occurred in the real-time market. At 5:10 PM on June 2, 2016, the real-time market price was just under \$50/MWh, then at 5:30 PM price rises to above \$500/MWh. As evidenced by the change in the supply relationship curves between 5:10 PM (top sub-panel of c) and 5:30 PM (top sub-panel of d), the model largely rationalizes this price change as being due to a shock in the available generating capacity, ε_t^P —as opposed to an anticipated or unanticipated movement along the curve driven by net load—perhaps due to unplanned generator outages or a transmission congestion event.

The assumption that RTM supply relationship fluctuations are due to generator unavailability—and not a shock common to all generators—is important for our analysis (see Section 3.2). Table D.1 provides evidence regarding the plausibility of our modeling assumptions. It displays the results of several regressions of prices in the day-ahead and real-time markets (and their deviations) on fuel prices as measured by the daily spot price for natural gas. It shows first that daily natural gas prices strongly impact mean P_t^{DAM} . The magnitude is consistent with complete pass-through from natural gas prices to wholesale electricity prices.⁵⁷ A similar pattern holds for P_t^{RTM} . This motivates our estimation of separate supply relationship and demand parameters by sample day. In contrast, when gas prices are high, we find no positive association with P_t^{RTM} being higher than P_t^{DAM} . In other words, gas price variation does not appear to be causing price spikes in RTM prices relative to DAM prices. This lends credence to our assumption that RTM price spikes are due to generator or transmission unavailability, which batteries can then mitigate, rather than common cost shocks.

Table D.1: Regression Results of Day-Ahead (DAM) and Real-time Market (RTM) Prices on Natural Gas Price

	Dependent Variable:						
	P_t^{DAM}	P_t^{RTM}			$P_t^{RTM} - P_t^{DAM}$		
	Mean	Mean	10th	90th	Mean	10th	90th
P^{NG}	10.40*** (0.58)	7.31*** (0.59)	4.62*** (0.54)	11.50*** (1.01)	-3.09*** (0.46)	-5.93*** (0.62)	0.81 (0.63)
R-squared	0.32	0.15	0.09	0.14	0.04	0.06	0.00
Observations	1459	1459	1459	1459	1459	1459	1459


Notes: This table summarizes the coefficient estimates on the natural gas price from several regressions where the dependent variable is a part of the distribution of the daily day-ahead (P_t^{DAM}) or real-time (P_t^{RTM}) market prices or their deviation for that particular day. In all regressions, the unit of observation is a day, and the sample is all days from 2016 to 2019. We calculate the distribution from five-minute or hourly prices over the day. We report heteroskedasticity consistent standard errors in parentheses.

Figure D.2 provides the fit of the supply relationship for June 28, 2016. The maroon dots show the net load forecasts and DAM price realizations. The blue solid line shows the predicted DAM prices as a function of the forecasts of net load from our estimated

⁵⁷We calculate that the median gas generator in California had a heat rate of 8.79, which should be scaled up by approximately 5% to account for losses from gross to net generation. The scaled figure is similar to our estimated coefficient of 10.40.

model. Finally, the orange dashed line shows the predicted DAM prices as a function of the forecasts of net load from a model estimated without ramping costs (i.e., $\alpha = 1$). By allowing for ramping costs, the solid line is able to explain more of the variation in the DAM prices than the dashed line, and hence lies closer to the maroon dots.

Figure D.2: Supply Relationship From Day-Ahead Market



Notes: This figure displays the day-ahead market prices and forecast of net load for each hour for June 28, 2016 (maroon dots). Additionally, the figure displays the estimated supply relationship with ramping costs (blue solid line) and without ramping costs (orange dashed line). The reported market prices are for the CAISO South Zone Trading Hub (SP 15).

E The Kalman Filter/Smoothen

As described in Section 3.2, a complication of our data is that CAISO implements the day-ahead market (DAM) only at the hourly frequency, reporting prices and forecasts for net load that are constant over the 12 five-minute intervals of each hour. Our operations model and the real-time market (RTM) prices use a five-minute frequency. Thus, our estimation procedure needs to accommodate the mixed-frequency nature of the data.

We use the Kalman filter/smoothen to temporally disaggregate (i.e., interpolate) the forecasts of net load to yield a forecast at the five-minute frequency. Generically, assume that a series A_t is observed only every h periods, and what is observed is the average of the interim h periods of the latent process a_t , so $A_t = \frac{1}{h} \sum_{j=0}^{h-1} a_{t-j}$. Our objective is to take the observed series A_t and construct estimates of the latent process a_t such that the implied values of the accumulated version of that series, ϕ_t , match the

observable data (A_t) at the end of the h periods. We cast the problem as a state space model and use the Kalman filter/smoother to estimate the latent process (e.g., [Proietti, 2006](#)).

More specifically, we use the following state space model:

$$\begin{aligned} A_t &= H_t \begin{bmatrix} a_t \\ \phi_t \end{bmatrix}, \\ \begin{bmatrix} a_t \\ \phi_t \end{bmatrix} &= M_t \begin{bmatrix} a_{t-1} \\ \phi_{t-1} \end{bmatrix} + U_t \psi_t, \quad \psi_t \sim N(0, 1), \end{aligned}$$

where H_t is a deterministically time-varying selection matrix⁵⁸ designed to handle the missing observations of A_t ; M_t ⁵⁹ and U_t ⁶⁰ are deterministically time-varying matrices designed to create the accumulated version of the latent process, ϕ_t ; and ψ_t is a serially independent error term that contributes to the time series variation in the latent process of interest a_t . We use the techniques outlined in [Harvey \(1989\)](#) and [Durbin and Koopman \(2012\)](#) to recover an estimate of the latent a_t for each five-minute interval in our sample.⁶¹ We then use these estimates to augment our data on the deterministic portion of net load, X_s^L .

F Modeling Battery Capacity Degradation

We calculate the best perceived efficiency level given capacity as the one that maximizes an approximation of the expected discounted future value accounting for capacity degradation. For any candidate v^p we approximate this value by first calculating the realized profits over 2015 from the solutions to the Bellman equations (2), which we denote $\Pi(v^p)$. We then calculate the capacity degradation, $\delta(v^p)$, using [Xu et al.](#)

⁵⁸ H_t iterates between the matrix $\begin{bmatrix} 0 & 1 \end{bmatrix}$ on the last period of each hour (the period we observe A_t , and $\begin{bmatrix} 0 & 0 \end{bmatrix}$ for the first to penultimate period of each hour.

⁵⁹ M_t takes 12 possible values for each period within the hour such that $M_t = \begin{bmatrix} 1 & 0; 1/j(t) & (j(t) - 1)/j(t) \end{bmatrix}$, where $j(t)$ is the period within the hour associated with time period t .

⁶⁰ U_t takes 12 possible values for each period within the hour such that $U_t = \begin{bmatrix} 1; 1/j(t) \end{bmatrix}$, where $j(t)$ is the period within the hour associated with time period t .

⁶¹ See [Brave et al. \(2021\)](#) for the explicit recursive formulation of the Kalman filter/smoother equations for a temporally aggregated series involving an average.

(2016), as we discuss below.⁶²

This allows us to scale these profits by $1 - \beta - \delta(v^p)$ which provides the expected discounted value, with the approximation that the profits in future years will be similar to profits in the current year. Finally, we solve for v^* as:

$$v^* = \arg \max_{v^p} \frac{\Pi(v^p)}{1 - \beta - \delta(v^p)}. \quad (\text{F.1})$$

In words, v^* is the optimal perceived roundtrip efficiency for a single agent to maximize 2015 expected discounted profits, accounting for the capacity degradation generated by its charge/discharge decisions.

In Xu et al. (2016), capacity degradation depends on: (1) temperature, (2) depth-of-discharge, (3) state-of-charge, (4) calendar time, and (5) number of cycles. For our application, we assume that batteries are operated at 25°C (77°F) throughout the year, which is the Xu et al. base case.

Let K denote the battery's capacity this period, K' denote its capacity next period,⁶³ and g_d be the term that determines degradation between the current period and next period, so that:

$$K' = K \exp(-g_d). \quad (\text{F.2})$$

From Xu et al. (2016), g_d consists of calendar degradation and cycle degradation.

The first component of the degradation function, calendar degradation g_t , is the portion that occurs regardless of how much the battery is charged or discharged. Calendar degradation is a function of elapsed time as well as the battery's mean state-of-charge. Battery capacity will degrade more if the battery is left idle at full state-of-charge relative to if the battery is left idle at 50% state-of-charge. More concretely, at 25°C, calendar degradation is the following function of elapsed time in seconds, \tilde{t} , and the mean state-of-charge during the time elapsed, $\bar{\sigma}$:

$$g_t = 0.000000000414 \times \tilde{t} \times \exp(1.04(\bar{\sigma} - 0.5)). \quad (\text{F.3})$$

The second component of the degradation function, cycle degradation, is degradation attributable to operations. Using the Xu et al. notation, define N to be the total

⁶²With a slight abuse of notation, we express Π and δ as functions of v^p and suppress the argument K in this appendix.

⁶³We use a period length of a week, as we discussed in Section 5.2.

number of cycles that the battery undertakes during a time period, where a full cycle indicates a battery making a roundtrip of charging and discharging; n_i to indicate if cycle i was a full roundtrip cycle ($n_i = 1$) or a half cycle ($n_i = 0.5$) of either charge or discharge; and g_{ci} to be the cycle degradation during cycle i . The cycle degradation g_{ci} depends on the mean state-of-charge during cycle i , σ_i , as well as the depth of discharge of the cycle, $\tilde{\delta}_i$. The depth of discharge indicates what fraction of power was gained or lost during the cycle. Cycle degradation is convexly increasing in the depth of discharge. E.g., cycling from 0% to 100% once is more damaging than cycling from 25–75% twice. Applying [Xu et al. \(2016\)](#) to the case of 25°C,

$$g_{ci} = \exp(1.04(\sigma_i - 0.5)) \times (140000\tilde{\delta}_i^{-0.501} - 123000)^{-1}. \quad (\text{F.4})$$

We combine the different degradation terms to write:

$$g_d = g_t + \sum_i^N n_i g_{ci}. \quad (\text{F.5})$$

From (F.3)–(F.5), capacity degradation g_d is a function of \tilde{t} , $\bar{\sigma}$, n_i , $\tilde{\delta}_i$, and σ_i , $\forall i = 1, \dots, N$.

Following [Xu et al. \(2016\)](#), we perform the following algorithm to simulate capacity degradation on our evaluation sample using v^* :⁶⁴

1. Solve the optimal policy for a given week. Recall that we solve for policies separately for each day within the week and that our policy functions for the evaluation sample incorporate a heuristic approach that limits cycling due to degradation.
2. Use the optimal policy from (1) and the realized stream of net load residuals ε^L , price residuals ε^P , and supply curve parameters across all time periods in the week to simulate charge/discharge actions.
 - Record the batteries' state-of-charge for each five-minute time interval of the simulation.

⁶⁴Our algorithm for the training sample is similar, but we calculate it over the entire 2015 training sample period (rather than separately by each week), use perfect foresight policies, and evaluate it separately across candidate values of v^P .

3. Calculate g_t over the simulation period using (F.3).
 - Use the recorded state-of-charge path to calculate the mean state-of-charge over the simulation period, $\bar{\sigma}$.
 - Over one week, $\tilde{t} = 60 \times 60 \times 24 \times 7 = 604,800$.
4. Feed the recorded state-of-charge path into a rainflow cycle counting algorithm.
 - See <https://www.mathworks.com/matlabcentral/fileexchange/3026-rainflow-counting-algorithm>.
 - The rainflow counting algorithm returns N , n_i , $\tilde{\delta}_i$, and $\sigma_i, \forall i = 1, \dots, N$. In words, it returns the number of cycles and whether each cycle is full or half, and determines the depth-of-discharge and mean state-of-charge for each cycle.
5. Calculate $g_{ci}, \forall i = 1, \dots, N$ using (F.4).
6. Calculate the total degradation rate $\exp(-g_d)$ for each week-long simulation using the above estimates and (F.5) and (F.2).

Finally, we note that this formulation implicitly assumes that both power and energy capacity degrade through cycling. The engineering literature shows that primarily energy capacity should degrade. Therefore, our calculation should provide a lower bound on the social value of storage.

G Details of Battery Capital Costs Estimation

This appendix provides details on our estimation of battery capital costs. The National Renewable Energy Laboratory (NREL) cost projections in Figure 2a motivate the functional form we use in (12). In particular, they demonstrate: (i) a downward trend in costs, (ii) a non-linear trajectory to costs, (iii) an increase in the uncertainty the further we are in the future, and (iv) positive skewness in the distribution of future costs. The downward trend in costs motivates the drift term in our model; the non-linear trajectory motivates the exponential formulation; the increasing level of uncertainty in the forecast uncertainty motivates the unit-root (in logarithms) formulation of the model;

and the positive skewness in the cost assessments justifies the log-normal distribution for the shock process.

Our estimation treats the first year of our sample, 2018, as $y = 0$. We rescale costs in year y to be relative to initial cost c_0 , so that $\tilde{c}_y \equiv c_y/c_0$. Taking logs of both sides of the (rescaled) capital cost evolution equation (12) from Section 5.2, we obtain:

$$\ln(\tilde{c}_y) - \underbrace{\ln(\tilde{c}_0)}_{\ln 1=0} = \tau \times y + \sum_1^y \xi_y. \quad (\text{G.1})$$

We use a method of moments approach to recover the two parameters τ and σ_c . Using (G.1), we derive the following moment conditions. First moment:

$$E[\ln(\tilde{c}_y)] = \tau \times y. \quad (\text{G.2})$$

Second moment:

$$\begin{aligned} \text{Var}[\ln(\tilde{c}_y)] &= \text{Var}\left[y\tau + \sum_1^y \xi_y\right] \\ \Rightarrow \text{Var}[\ln(\tilde{c}_y)] &= \text{Var}[y\tau] + \text{Var}\left[\sum_1^y \xi_y\right] \\ &\Rightarrow \text{Var}[\ln(\tilde{c}_y)|y] = y \times \text{Var}[\xi_y] \\ &\Rightarrow \text{SD}[\ln(\tilde{c}_y)|y] = \sqrt{y} \times \text{SD}[\xi_y] \\ &\Rightarrow \text{SD}[\ln(\tilde{c}_y)|y] = \sqrt{y} \times \sigma. \end{aligned} \quad (\text{G.3})$$

We estimate the parameters τ and σ_c that solve the two moment conditions by estimating two univariate regressions, pooling across the set of cost projections. For the first regression, the dependent variable is $\ln(\tilde{c}_y)$, and the independent variable is y .

For the second regression, the dependent variable is the standard deviation of all the logged cost realizations $\ln(\tilde{c}_y)$ conditional on y and the independent variable is \sqrt{y} . To accommodate the variation in the number of cost assessments over time, the second regression uses weights based on the number of cost projections that were made for that year.⁶⁵

⁶⁵Figure 2a shows that years that are further in the future tend to have fewer cost projections.

Importantly, we do not observe actual realizations of the battery capital cost process, only the set of *projected* cost realizations from [Cole and Frazier \(2019\)](#). Therefore, our estimation treats each cost projection (i.e., each line in Figure 2a) as a realization of the cost process. Our estimates for the cost process are $\hat{\tau} = -0.044$ (with a standard error of 0.001) and $\hat{\sigma}_c = 0.064$ (with a standard error of 0.003). Following [Cole and Frazier \(2019\)](#), our simulations use an initial condition for capital costs in 2018 of $c_{2018} = \$380/\text{kWh}$. Since we use NREL data, our estimates pertain exclusively to lithium-ion battery costs, and do not include alternative storage technologies or account for learning-by-doing.

H Identifying the Counterfactual Supply Relationship with CEMS Cost Data

This appendix considers an alternative way to identify dispatchable generator pricing under counterfactual environments, that we considered but did not use in our main analysis. This method involves calculating the observable components of marginal costs at the generator level. We could then use these calculated costs to recover the sum of markups and ramping costs. Controlling for marginal costs in this way could also potentially allow us to estimate separate ramping costs by generator type.

Towards these ends, we gathered all the generators that report their generation and fuel consumption in the Environmental Protection Agency’s (EPA) Continuous Emissions Monitoring System (CEMS) database in the state of California, calculating their capacity and heat rate following [Gowrisankaran et al. \(2022\)](#). We constructed an observable marginal cost for each generator by using the following formula:

$$MC_{it} = \text{Heat Rate}_{it} \times \text{Fuel Price}_t \times (1.0526) + 2.37, \quad (\text{H.1})$$

where Fuel Price_t is the spot price for fuel (e.g., natural gas) and can vary over time, the scale 1.0526 reflects the adjustment for approximated 5% losses from gross to net generation ([Graff Zivin et al., 2014](#)), and \$2.37 reflects an adjustment for variable operations and maintenance (O&M) costs from the CEC 2019 report ([California Energy](#)

[Commission, 2019](#)).⁶⁶ Next, we constructed an industry observable marginal cost curve by sorting the generators from lowest to highest observable marginal cost and assuming constant observable marginal costs for each generator up to its capacity. When combined with information on total net generation, the industry observable marginal cost curve can be used to predict the market clearing price, absent misspecifications, other costs, and market power.

We defined the set of available generators in the market at each hour with two different approaches:

1. For every hour, we assume that only generators that produced in that hour are available to produce.
2. For every month, we assume that generators that produce at some hour in that month are available to produce at every hour in that month.

For both approaches, we implemented a robustness check where we restrict the generators in the sample to those in Southern California, which we define as below latitude 36.7378 (essentially south of Fresno).

Figure [H.1](#) displays the industry observable marginal cost curve, using July 2016 natural gas prices and method 2 for calculating the set of available generators. We observe the hockey-stick nature of the industry observable marginal cost curve: costs are below \$40MWh for much of the domain of the curve, but tick up sharply after 30,000 MWh. We plot the distribution of the hourly total generation from all the units in the CEMS data during July 2016 on top of the industry observable marginal cost curve. Surprisingly, we do not observe even one hour with net load sufficient to reach the steep part of the cost curve. This figure shows that this cost curve is unlikely to reproduce the observed wholesale electricity price spikes, which are a crucial component of the revenues that batteries earn.

Figure [H.2](#) displays the cost curve and distribution of total hourly generation for July 2016, but now for generators in Southern California. While both the cost curve and distribution of total hourly generation are shifted to the left, we observe the same pattern as in Figure [H.1](#).

Next, we summarize the descriptive evidence of how the two measures of industry observable marginal cost relate to the day-ahead market prices we observe for the

⁶⁶Note that the marginal costs in [\(H.1\)](#) do not include ramping costs.


SP-15 hub. To do this, we run regressions of the following form:

$$P_t^{\text{DAM}} = \beta_0 + \beta_1 MC_t + \varepsilon_t$$

where P_t^{DAM} is the day-ahead market price and MC_t is industry marginal costs (defined using both methods). In some specifications, we include a day-of-sample fixed effect, in which case the coefficient β_1 is identified only from within-day (hourly) variation in market-level marginal costs and day-ahead market prices.

Tables H.1 reports the coefficient estimates from these specifications for all generators in California, while Table H.2 includes only Southern California generators. The tables show that, without fixed effects, industry observable marginal costs explain only a relatively small fraction—21 percent at the highest—of the overall DAM price variation. Method 2 performs better than method 1 in explaining DAM prices. Nonetheless, across specifications and sample, the highest R^2 we observe is 37 percent, implying

Figure H.1: Market-Wide Hourly Generation and Supply Curve




Notes: This figure plots the (sorted) distribution of observable marginal costs for each generator in California (using method 2 to determine available generators) along with the histogram of generation.

that this approach does not predict the majority of the DAM price variation.

We opted not to use this method for our main simulations because of its lack of predictive power and the fact that it cannot predict the price spikes observed in the real-time market. Our supply relationship accounts for four potential forces that we cannot obtain from the CEMS data: market power, ramping costs, imports, and transmission constraints. We believe that these forces may explain some of these discrepancies.

Figure H.2: Southern California Hourly Generation and Supply Curve



Notes: This figure plots the (sorted) distribution of observable marginal costs for each generator in Southern California (using method 2 to determine available generators) along with the histogram of generation.

Table H.1: Results of Day-Ahead Market Prices on Marginal Costs

	Dependent Variable: P_t^{DAM}			
	Industry MC Method 1		Industry MC Method 2	
Marginal Cost	0.73 (0.02)	2.63 (0.34)	1.66 (0.04)	15.24 (1.18)
Constant	2.35 (0.63)	-81.16 (14.87)	-3.16 (0.78)	-313.43 (26.94)
Day FEs		✓		✓
R-squared	0.16	0.07	0.20	0.37
Observations	34988	34988	34988	34988

Notes: This table displays coefficient estimates for regressions of the day-ahead market price on observable marginal costs for all generators in California. We report heteroskedasticity consistent standard errors in parentheses.

Table H.2: Results of Day-Ahead Market Prices on observable Marginal Costs (South CA)

	Dependent Variable: P_t^{DAM}			
	Industry MC Method 1		Industry MC Method 2	
Marginal Cost	0.82 (0.02)	2.28 (0.26)	1.70 (0.04)	11.88 (0.90)
Constant	1.16 (0.67)	-59.22 (10.57)	-4.19 (0.80)	-237.15 (20.71)
Day FEs		✓		✓
R-squared	0.17	0.08	0.21	0.33
Observations	34988	34988	34988	34988

Notes: This table displays coefficient estimates for regressions of the day-ahead market price on observable marginal costs for all generators in Southern California. We report heteroskedasticity consistent standard errors in parentheses.