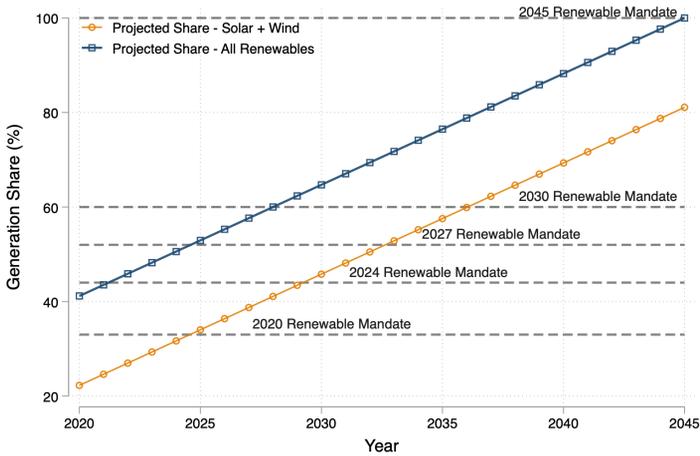
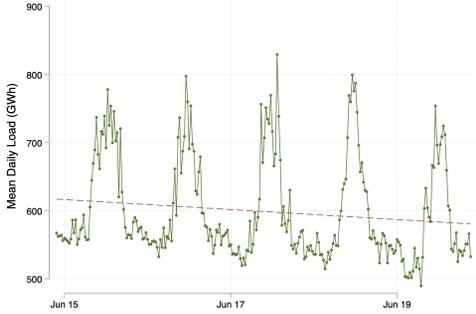


FIGURE A.2.—Renewable Energy Over Time Under the California Renewable Portfolio Standard



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.

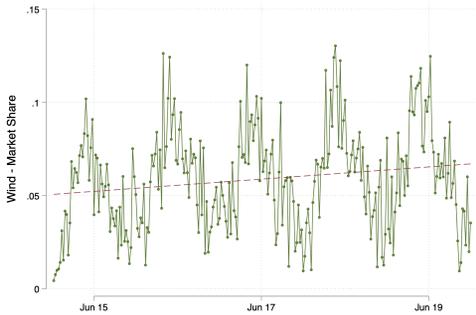
FIGURE A.3.—CAISO Electricity Market Trends



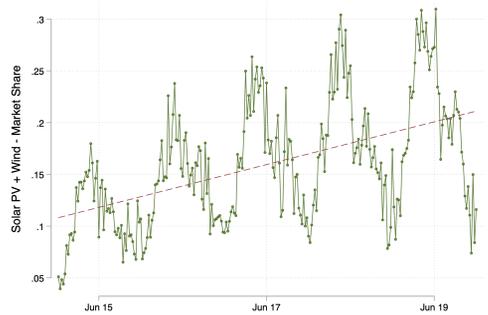
(a) Load



(b) Solar PV Share



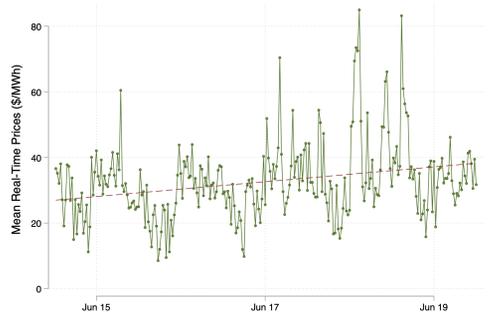
(c) Wind Share



(d) Solar + Wind Share



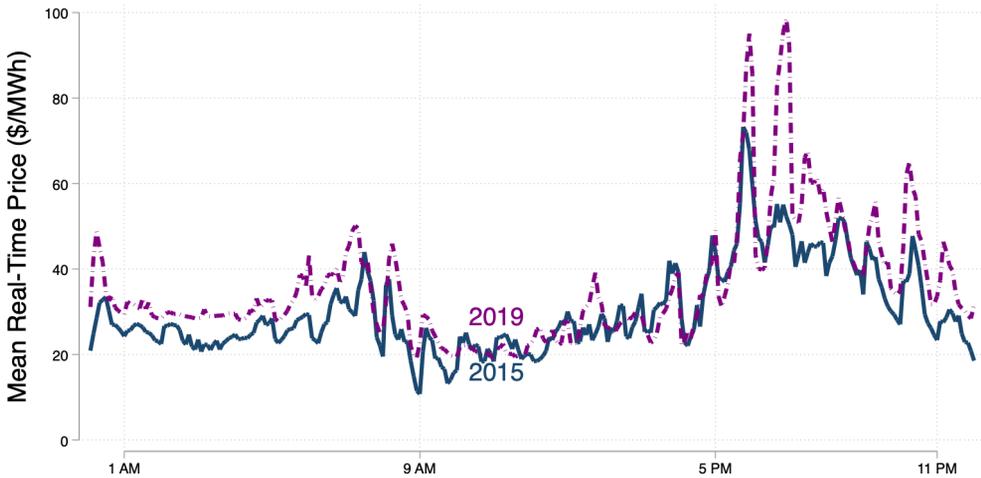
(e) Natural Gas Price (\$/mmbtu)



(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)

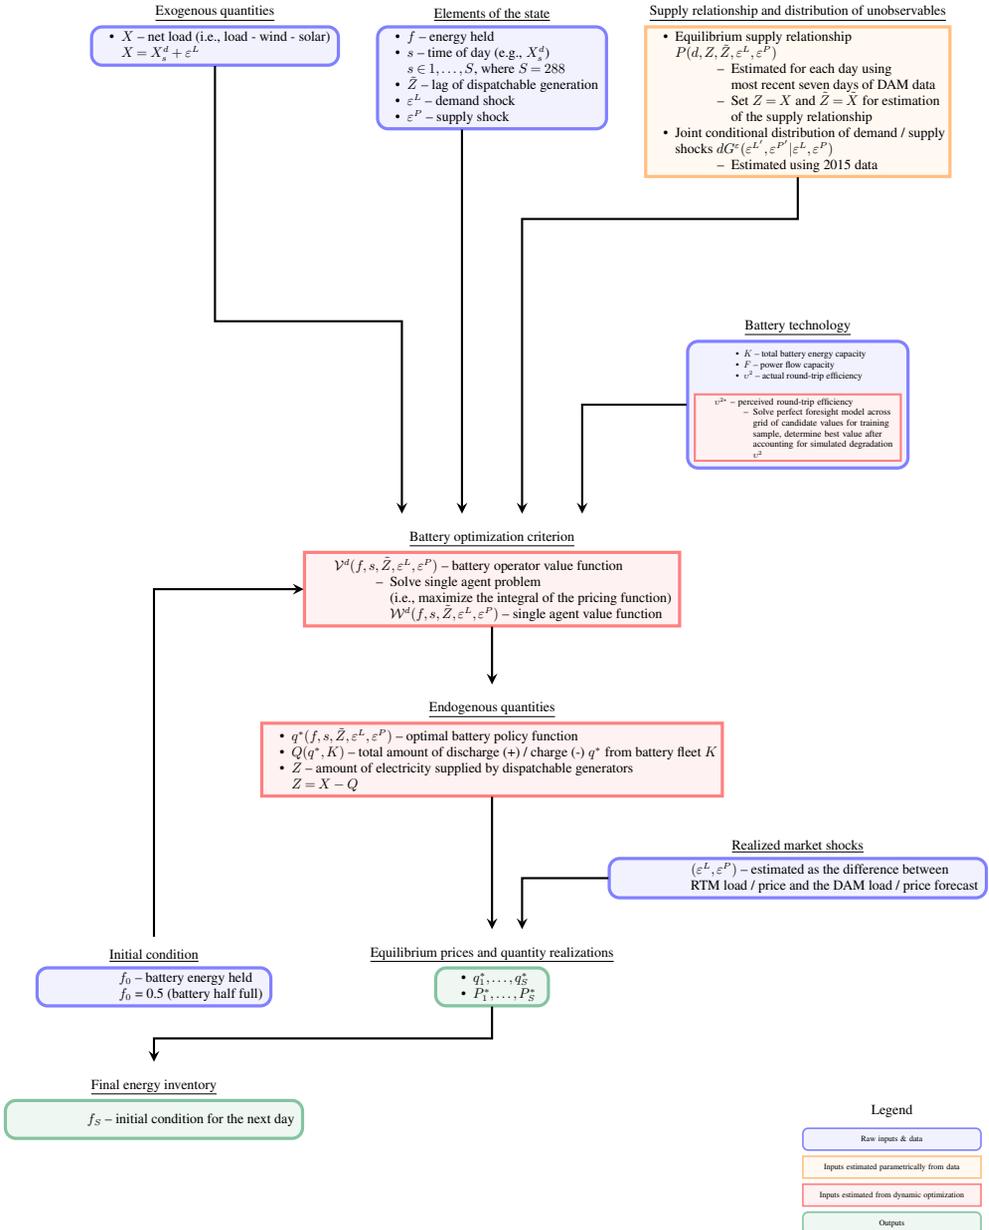


TABLE A.1
SUMMARY STATISTICS FOR ESTIMATED SUPPLY RELATIONSHIP PARAMETERS

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)
$\sigma^{P,Peak}$	0.012	0.010	0.010	0.016	0.016	0.013
$\sigma^{P,Off-peak}$	0.010	0.006	0.008	0.012	0.014	0.010
Observations	17568	105408	105120	105120	105120	420768

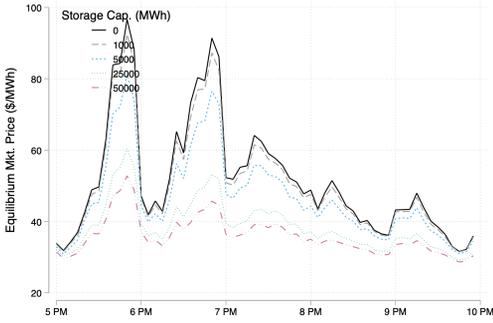
Notes: This table summarizes the estimates of the supply relationship residual (ε_t^P) parameters. The 2015 sample includes only November and December. We report standard errors, clustered by day-of-sample, in parentheses.

TABLE A.3
SUMMARY STATISTICS FOR ESTIMATED NET LOAD MODEL

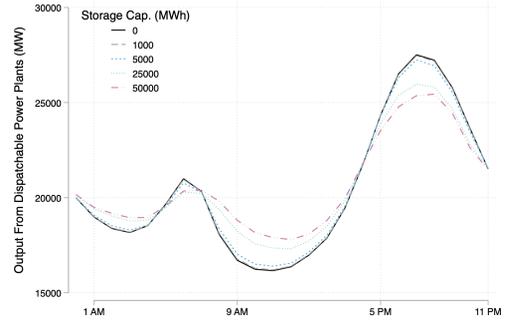
	2015	2016	2017	2018	2019	2016–19
	(a) Dependent Variable: Net Load _t					
Net Load DAM Forecast	0.969 (0.003)	0.950 (0.002)	0.950 (0.001)	0.971 (0.001)	0.955 (0.002)	0.956 (0.001)
Dependent Variable Mean	1794.61	1798.35	1734.13	1687.41	1599.83	1704.99
In-sample RMSE	67.721	83.007	77.494	74.292	80.513	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



(a) Peak Five-Minute Equilibrium Prices



(b) Mean Hourly Output from Dispatchable Generators

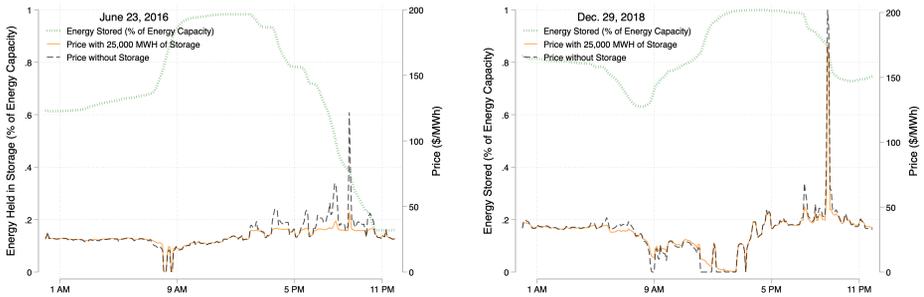
Notes: Each line plots the mean counterfactual outcome for specific storage capacity level across all days during 2016–19.

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

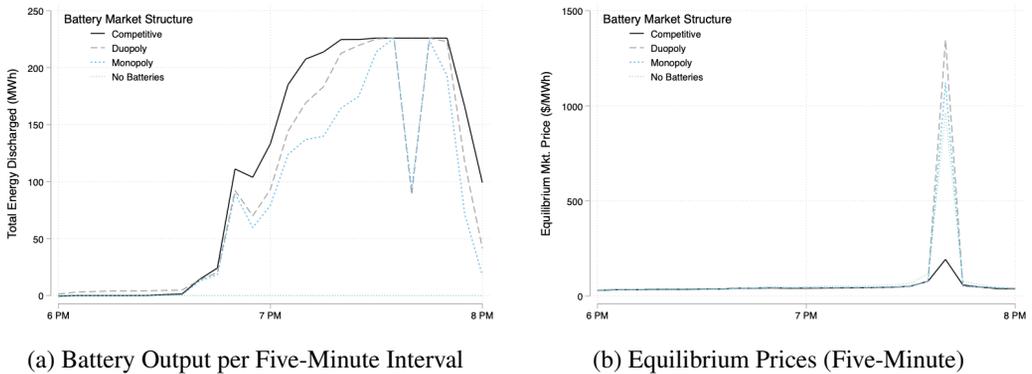
Notes: The first column lists the aggregate battery capacity. The second column indicates the total revenue a battery owner would earn between 2016–19 summed over the least profitable 99 percent of time periods. The third column lists the total revenue a battery owner would earn summed over the most 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)
$\ln(K^*) \times$ 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)	
(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).¹ This changed starting in 2020, with the development of the Gateway Energy Storage System in California, which has 250 MW of capacity battery storage.² 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 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.

¹For an overview of the growth of battery storage projects, see EIA (2022).

²The 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).

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 [Fowlie et al. \(2016\)](#) (henceforth, R/FRR).³ 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).

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 $\nu \in (0, 1)$, θ_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](#).⁴

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

³Another 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. Appendix H provides more details on this point.

⁴Similar 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 $Z/\mathcal{K} = 1$.

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 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)		
	Base Model	R/FRR Functional Form
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 [Fowle 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 $\theta_1 \in [-700, 500]$, $\theta_2 \in [0, 500]$, $\theta_3 \in [1.01, 4]$, $\kappa \in [1, 8]$, $\alpha \in [0, 1]$.⁵ 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) .

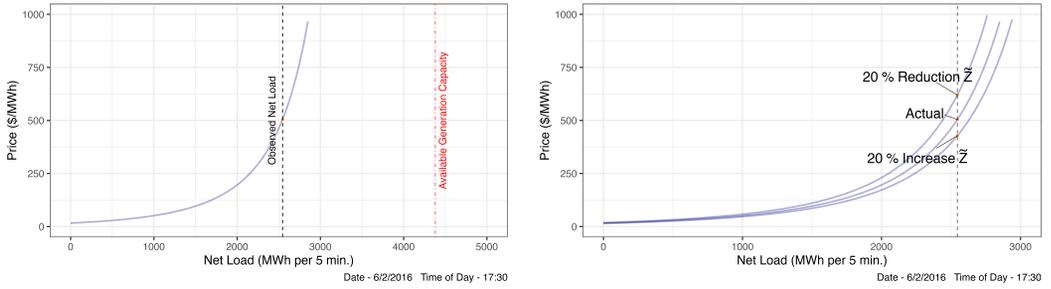
As an example of the features of our approach towards modeling the supply relationship, Figure D.1 provides the supply relationship on June 2, 2016, when net load was approaching the constraint on available generating capacity. From Figure D.1a, at 5:30 PM, the market equilibrium was near an inflection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

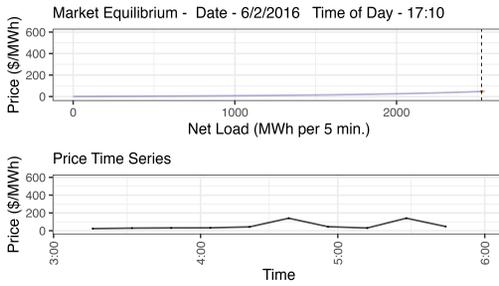
⁵We also compute a perfect foresight model, which uses the same marginal cost curve parameters.

FIGURE D.1.—Time-Varying Supply Relationship Curve

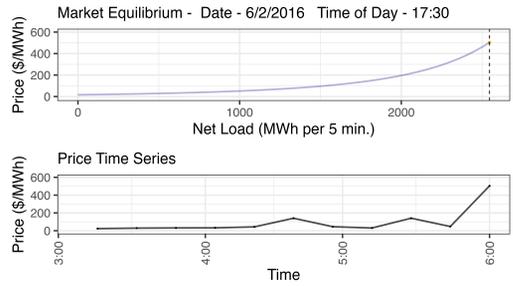


(a) Price Rises at Capacity Constraint

(b) Generation Output at $t - 1$ shifts MC



(c) Equilibrium Before Price Spike Event



(d) Equilibrium During Price Spike Event

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.

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.⁶ 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.

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 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.

⁶We 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.

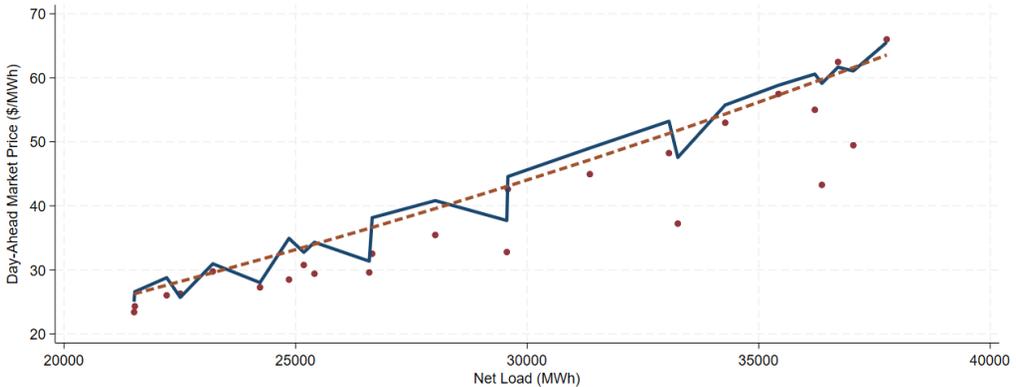
TABLE D.1

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.—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/SMOOTHER

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/smoothing 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/smoothing to estimate the latent process (e.g., [Proietti, 2006](#)).

More specifically, we use the following state space model:

$$A_t = H_t \begin{bmatrix} a_t \\ \phi_t \end{bmatrix},$$

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

where H_t is a deterministically time-varying selection matrix⁷ designed to handle the missing observations of A_t ; M_t ⁸ and U_t ⁹ 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.¹⁰ 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. \(2016\)](#), as we discuss below.¹¹

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.

⁷ H_t iterates between the matrix $[0 \ 1]$ on the last period of each hour (the period we observe A_t , and $[0 \ 0]$ for the first to penultimate period of each hour.

⁸ M_t takes 12 possible values for each period within the hour such that $M_t = [1 \ 0; 1/j(t) \ (j(t) - 1)/j(t)]$, where $j(t)$ is the period within the hour associated with time period t .

⁹ U_t takes 12 possible values for each period within the hour such that $U_t = [1; 1/j(t)]$, where $j(t)$ is the period within the hour associated with time period t .

¹⁰ See [Brave et al. \(2021\)](#) for the explicit recursive formulation of the Kalman filter/smoothing equations for a temporally aggregated series involving an average.

¹¹ With a slight abuse of notation, we express Π and δ as functions of v^p and suppress the argument K in this appendix.

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,¹² 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, $\bar{\sigma}$:

$$g_t = 0.00000000414 \times \tilde{t} \times \exp(1.04(\bar{\sigma} - 0.5)). \quad (\text{F.3})$$

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 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^* .¹³

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.

¹²We use a period length of a week, as we discussed in Section 5.2.

¹³Our 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 .

- Record the batteries' state-of-charge for each five-minute time interval of the simulation.
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:

$$\begin{aligned} \text{Var}[\ln(\tilde{c}_y)] &= \text{Var}\left[y\tau + \sum_1^y \xi_y\right] \\ \Rightarrow \text{Var}[\ln(\tilde{c}_y)] &= \text{Var}[y\tau] + \text{Var}\left[\sum_1^y \xi_y\right] \end{aligned}$$

$$\begin{aligned}
&\Rightarrow \text{Var}[\ln(\tilde{c}_y)|y] = y \times \text{Var}[\xi_y] \\
&\Rightarrow \text{SD}[\ln(\tilde{c}_y)|y] = \sqrt{y} \times \text{SD}[\xi_y] \\
&\Rightarrow \text{SD}[\ln(\tilde{c}_y)|y] = \sqrt{y} \times \sigma.
\end{aligned} \tag{G.3}$$

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.¹⁴

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:

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

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](#)).¹⁵ 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:

¹⁴Figure 2a shows that years that are further in the future tend to have fewer cost projections.

¹⁵Note that the marginal costs in [\(H.1\)](#) do not include ramping costs.

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.

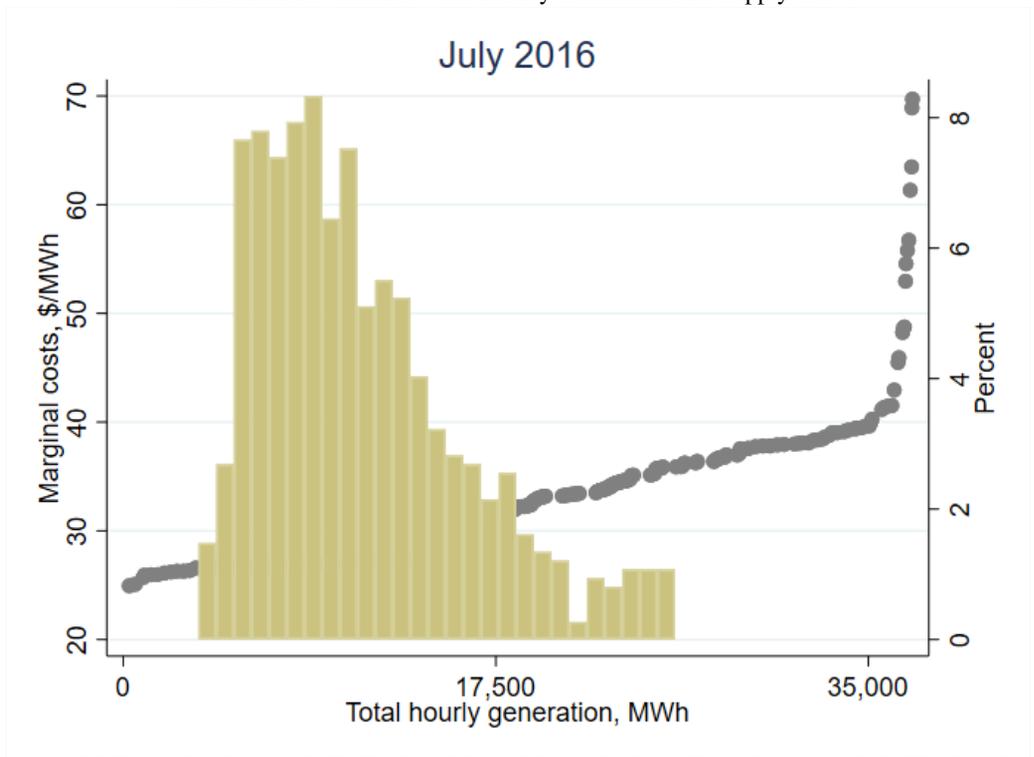
For both approaches, we implemented a robustness check where we restrict the generators in the sample to those in Southern California, which we define as below latitude 36.7378 (essentially south of Fresno).

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 \$40/MWh 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 SP-15 hub. To do this,

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.

we run regressions of the following form:

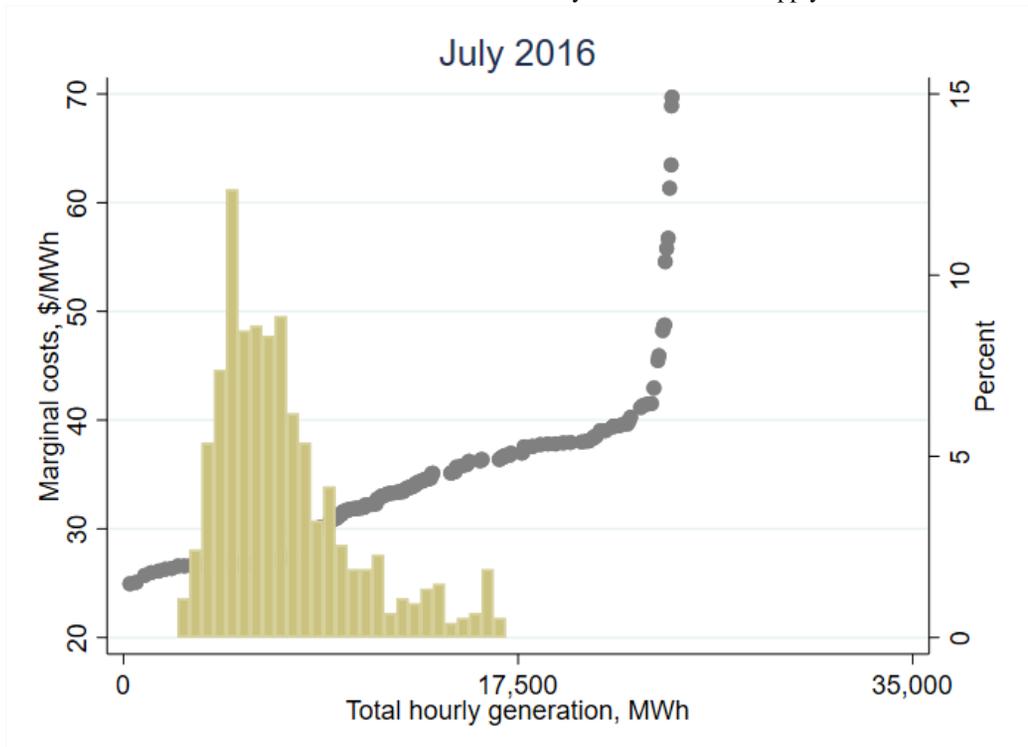
$$P_t^{\text{DAM}} = \beta_0 + \beta_1 MC_t + \varepsilon_t$$

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 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.

REFERENCES

- BRAVE, SCOTT A., R. ANDREW BUTTERS, AND DAVID KELLEY (2021): "A Practitioner's Guide and MATLAB Toolbox for Mixed Frequency State Space Modeling," *Journal of Statistical Software*, Forthcoming. [17]
- 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>. [11]
- 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>. [20]
- COLE, WESLEY J AND ALLISTER FRAZIER (2019): "Cost Projections for Utility-Scale Battery Storage," Tech. rep., National Renewable Energy Lab.(NREL), Golden, CO (United States), accessed on November 26, 2019. [20]
- DURBIN, JAMES AND SIEM JAN KOOPMAN (2012): *Time Series Analysis by State Space Methods: Second Edition*, no. 9780199641178 in OUP Catalogue, Oxford University Press. [17]

- EIA (2022): “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>. [11]
- FOWLIE, MEREDITH, MAR REGUANT, AND STEPHEN P RYAN (2016): “Market-Based Emissions Regulation and Industry Dynamics,” *Journal of Political Economy*, 124 (1), 249–302. [12, 13]
- GOWRISANKARAN, GAUTAM, ASHLEY LANGER, AND WENDAN ZHANG (2022): “Policy Uncertainty in the Market for Coal Electricity: The case of Air Toxics Standards,” Working paper, National Bureau of Economic Research. [20]
- GRAFF ZIVIN, JOSHUA S., MATTHEW J. KOTCHEN, AND ERIN T. MANSUR (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. [20]
- HARVEY, A.C. (1989): *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge University Press. [17]
- PIRRONG, CRAIG (2012): *Commodity Price Dynamics: A Structural Approach*, Cambridge University Press. [12, 13]
- PROIETTI, TOMMASO (2006): “Temporal Disaggregation by State Space Methods: Dynamic Regression Methods Revisited,” *The Econometrics Journal*, 9 (3), 357–372. [17]
- RYAN, STEPHEN P (2012): “The Costs of Environmental Regulation in a Concentrated Industry,” *Econometrica*, 80 (3), 1019–1061. [12, 13]
- SPECTOR, JULIAN (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>. [11]
- XU, BOLUN, ALEXANDRE OUDALOV, ANDREAS ULBIG, GÖRAN ANDERSSON, AND DANIEL S KIRSCHEN (2016): “Modeling of Lithium-Ion Battery Degradation for Cell Life Assessment,” *IEEE Transactions on Smart Grid*, 9 (2), 1131–1140. [17, 18]