Endogenous Use of Electricity Storage in the Long Run[*]

Alex Marsh Hannah Rubenstein Drew Van Kuiken
Jonathan W. Williams Andrew J. Yates

Last Updated: September 24, 2026

Abstract

We study the role of electricity storage in a long-run equilibrium model with multiple storage technologies and region-specific renewable resources. Varying storage costs and efficiency as well as policy across a rich set of simulations, we characterize storage regimes, substitution patterns, and emissions impacts. A central result is that the temporal pattern of storage use is an equilibrium outcome rather than being fixed by technological characteristics; both short- and long-duration storage are commonly used for relatively short-term arbitrage. Nevertheless, technological differences shape the extent to which one type of storage can substitute for another.

Keywords: Electricity Storage, Long Run, Regional Heterogeneity

JEL Codes:Q41, Q42, L94

[*]We thank North Carolina Collaboratory for research funding for this projects.

1 Introduction

Electricity storage is becoming a central component of modern power systems. Rapid declines in the cost of battery technologies, combined with the expansion of intermittent renewable generation, have led to a sharp increase in storage deployment. At the same time, a growing set of technologies—including lithium-ion batteries, long-duration chemical storage, and mechanical systems—offer different capabilities for shifting energy over time. Understanding how these technologies will be deployed, and how they will interact with the evolving mix of generation, is inherently a long-run question. As storage becomes cheaper and more widespread, it will shape not only how electricity is produced, but also when it is consumed and how variability in renewable generation is managed.

In this paper, we study the role of storage in a long-run model of the electricity grid with multiple storage technologies and region-specific renewable resources. We allow for endogenous investment in both short- and long-duration storage. We characterize how storage technologies interact with each other, with renewable generation, and with policy. Varying storage costs and efficiency as well as renewable subsidies across a large set of simulated equilibria, we document patterns of technology adoption, substitution, and emissions outcomes. We find that storage use is determined endogenously in equilibrium rather than being fixed by technological characteristics, although those characteristics still shape outcomes. Both short- and long-duration storage are commonly used for short-run arbitrage. Nevertheless, substitution between technologies is asymmetric.

We develop a long-run equilibrium model of electricity generation and storage that builds on Holland et al. (2022). The model features endogenous investment in generation capacity and hourly operation over a representative year, allowing for realistic variation in renewable output. Our primary extension is to allow for multiple storage technologies that differ both in self-discharge and in round-trip efficiency. In particular, we distinguish between short-duration storage, which has high round-trip efficiency but loses stored energy relatively quickly, and long-duration storage, which has lower efficiency but little or no self-discharge and can retain energy over extended horizons.

We begin by characterizing the range of equilibrium outcomes across model runs. Depending on technology costs and policies, the model generates distinct storage regimes, including equilibria with no storage, only short-duration storage, only long-duration storage, and both technologies. These regimes differ systematically across regions, with storage more prevalent in areas with stronger renewable resources. The results highlight that both technologies can coexist in equilibrium, but that single-technology outcomes remain common, reflecting the discrete nature of storage adoption decisions.

We next examine how the viability of long-duration storage depends on the cost of short-duration storage. We show that the cost threshold at which long-duration storage is deployed rises with the cost of short-duration alternatives, and with improvements in storage efficiency. These relationships are highly nonlinear, with sharp transitions in adoption across parameter values. Statistical analysis confirms that technological and policy variables shift these thresholds in predictable ways, with solar subsidies and carbon pricing increasing the range over which long-duration storage is viable.

We then quantify substitution patterns between storage technologies by estimating own- and cross-price elasticities of storage capacity. We find that both technologies exhibit large own-price elasticities, indicating strong responses to changes in their own costs. However, cross-price effects are asymmetric: short-duration storage responds strongly to changes in long-duration costs, while long-duration storage responds more modestly to changes in short-duration costs.

We also examine how storage is used in equilibrium by analyzing utilization rates. Short-duration storage is used intensively and cycles frequently, consistent with intra-day arbitrage, while long-duration storage is used less intensively but still participates in near-daily arbitrage. These patterns highlight that storage technologies are not tied to fixed operational roles: both short- and long-duration storage are commonly used for short-run arbitrage. At the same time, differences in technological characteristics shape how storage is deployed across conditions and over time.

Finally, we examine emissions outcomes and find that both technologies reduce emissions. A decomposition reveals that this occurs primarily through reductions in fossil capacity, rather than through changes in fossil utilization.

Our paper contributes to a growing literature on the role of electricity storage in decarbonized power systems. This literature largely focuses on environments in which demand for electricity is perfectly inelastic. Within this class of models, the literature can be organized along two dimensions: whether it studies short-run or long-run outcomes, and whether it considers a single storage technology or multiple technologies. In the short run with a single storage technology, Tarel et al. (2024) and Karaduman (2023) analyze how storage interacts with renewable generation and prices when capacity is fixed. In the long run with a single technology, Butters et al. (2025) studies equilibrium investment and operation of storage. A separate set of papers considers multiple storage technologies. In the long run, Junge et al. (2022), Jenkins and Sepúlveda (2021) and Dowling et al. (2020) study investment and operation with multiple storage technologies that differ in their technological characteristics.[1] Recent work has also emphasized that the economic role of storage depends on its interac-

[1]See Selänniemi et al. (2024) for a summary of electricity storage papers.

tion with renewable generation and market equilibrium. For example, Andres-Cerezo and Fabra (2026) show that storage and renewable investments may act as either complements or substitutes depending on the generation mix and the level of renewable penetration

Our paper contributes to this literature by embedding storage investment and operation in a unified long-run economic equilibrium framework that allows for both endogenous demand and endogenous storage utilization. In our model, consumption responds to price signals, dampening extreme scarcity and altering the economic role of storage. Within this setting, we show that storage duration is not an exogenous technological characteristic but instead emerges endogenously from equilibrium utilization patterns. By modeling both self-discharge and round-trip efficiency, and allowing for endogenous investment and operation across a wide range of parameter configurations, we characterize how technology costs and system variability jointly determine not only which storage technologies are deployed, but also how intensively they are used and over what time horizons. This perspective highlights that storage technologies provide the capacity to shift energy across time, while equilibrium conditions determine the effective duration over which that capacity is utilized.

2 Model

We consider a long-run model of investment and generation in the electricity grid that is a modified version of Holland et al. (2022). The main modifications are that we allow for multiple storage technologies rather than one single technology and we allow for losses from charging and discharging and self-discharge over time rather than assuming a "perfect" storage technology.

As with any long-run model, all the variables of interest are endogenous. In our case, this includes electricity consumption, electricity generation and capacity, and storage utilization and capacity. Figure 1 illustrates the basic elements of the model and the flow of electricity from generation to households, as well as into and out of storage.

The model considers a representative year with $T = 8760$ hours. Let $t \in [1, T]$ denote a given hour, $i \in [1, I]$ denote a given generation technology, and $j \in [1, J]$ denote a given storage technology. Electricity consumption is Q_t and the consumer's utility from this consumption is $U_t(Q_t)$.

Electricity generation technology i is defined by capacity K_i, capital costs r_i per unit of capacity, constant operating costs c_i, and capacity factor f_{it}. The capacity factors account for the fact that renewable generation technologies are intermittent. For example, solar technologies have $f_{it} = 0$ at night whereas natural gas technologies have $f_{it} = 1$ for all t. Let q_{it} be the generation from technology i in hour t. Hourly generation must satisfy the

Figure 1: Model Elements and Electricity Flow

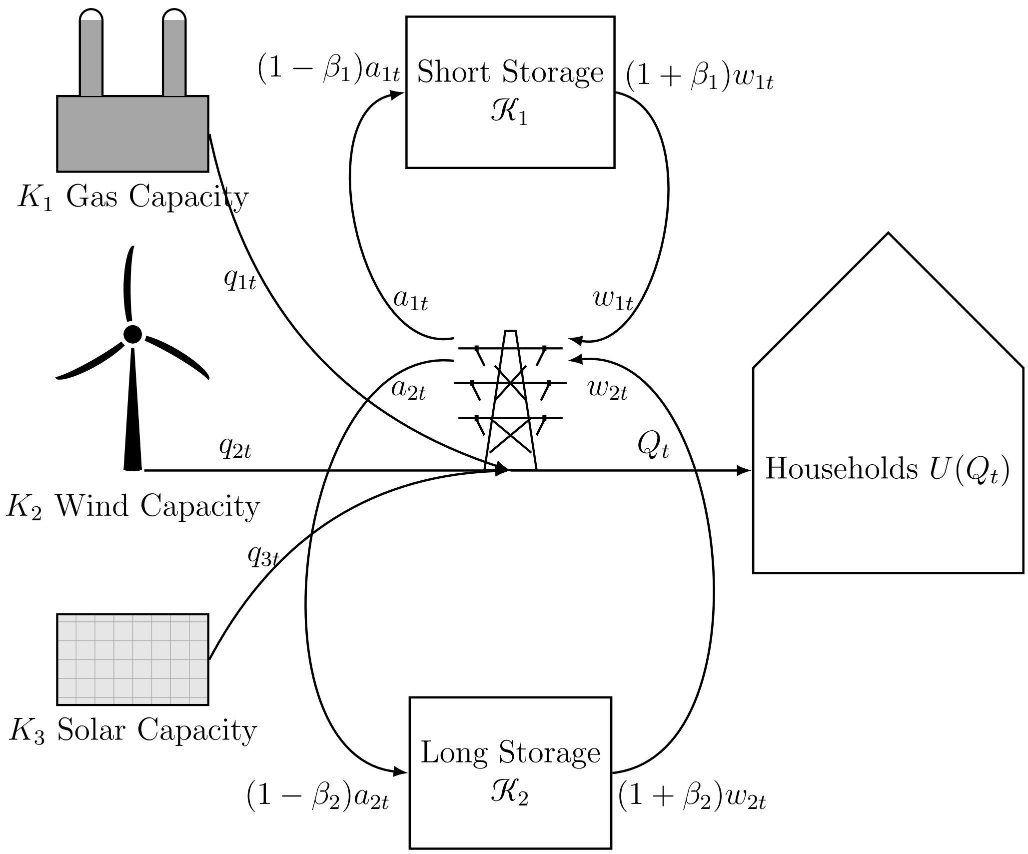

Notes: The arrows show how electricity flows from generation to storage and households. Labels indicate charging losses and delivered energy after discharging.

constraint $q_{it} \leq f_{it}K_i$. Let the external costs from technology i due to emissions of air pollution be given by δ_i per unit of generation.

The capacity of storage technology j is given by $\mathcal{K}_j$. The current state of storage in technology j during hour t is given by S_{jt} where $0 \leq S_{jt} \leq \mathcal{K}_j$. The amount of electricity added to storage technology j in period t (i.e., charging a battery) is given by a_{jt} and the amount of electricity withdrawn from storage (i.e. discharging a battery) is given by w_{jt}. The level of storage satisfies the dynamic equation

$$S_{jt} = \alpha_j S_{jt-1} + (1 - \beta_j)a_{jt} - (1 + \beta_j)w_{jt},$$

where α_j captures self-discharge and β_j captures efficiency losses due to charging and discharging.[2] Long and short duration storage technologies are differentiated by the values for α and β. Long-duration typically has lower self discharge (higher α) and greater efficiency losses (higher β) than short-duration storage. It is useful to characterize efficiency losses with the notion of round trip efficiency, defined as $\frac{1-\beta}{1+\beta}$. Storage capacity has capital costs ρ_j.

In each hour, the demand of electricity must less than or equal to the supply for electricity. We have

$$Q_t + \sum_j a_{jt} \leq \sum_i q_{it} + \sum_j w_{jt}.$$

We assume a long-run competitive equilibrium in which consumers and producers have perfect foresight. Accordingly, we can characterize this equilibrium with the planners problem:

$$\max_{Q_t, q_{it}, a_{jt}, w_{jt}, S_t, K_i, \mathcal{K}_j} \sum_t [U_t(Q_t) - \sum_i c_i q_{it}] - \sum_i r_i K_i - \sum_j \rho_j \mathcal{K}_j, \tag{1}$$

subject to the constraints described above and additional non-negativity constraints on consumption, generation, and storage additions and withdrawals. This formulation allows for elastic demand because consumption Q_t is a choice variable that is selected by the planner rather than being fixed exogenously. The planner trades off the marginal utility of consumption $U_t'(Q_t)$ against the marginal cost of supplying electricity.

[2] We abstract from constraints on charge and discharge rates, which may limit how quickly energy can be moved into and out of storage.

3 Calibration

Our calibration adopts the parameters in Holland et al. (2022) with the exception of storage parameters.

3.1 Holland et al. (2022) Baseline parameters

We begin by reporting the core elements of the Holland et al. (2022) calibration. Table 1 shows the baseline parameters for six generation technologies, including capital costs, operating costs, and emissions intensities for each technology. Combined-cycle plants have relatively low operating costs and are typically used for baseload generation, while peaker plants have higher operating costs but lower capital costs and are used primarily to meet periods of high demand.

Table 1: Baseline Parameters for Generation Technologies

	Operating Cost (\$/MWh)	Annual Capital Cost (\$/MW)	CO_2 Emissions (mt/MWh)
Solar	0	83,274	0
Wind	0	132,602	0
Nuclear	2.38	528,307	0
Coal	22.48	374,608	0.771
Combined Cycle Gas	26.68	79,489	0.338
Peaker Gas	44.13	54,741	0.526

Notes: Source Holland et al. (2022).

While Holland et al. (2022) consider 13 regions, we select four representative cases. These cases create a 2×2 structure defined by relatively high or low wind and solar levelized costs, as reported in Table 2. Levelized costs reflect both capital costs and capacity factors; in our setting, cross-region differences arise from variation in capacity factors, as capital costs are held constant across regions. The four regions—Texas, the Mid West, Florida, and New England—capture the key variation in renewable competitiveness across locations. Figure 2 provides a geographic mapping of the regions.

Table 2: Levelized Costs for Wind and Solar by Region

Region	Solar	Wind
Texas	38.86	37.84
New England	59.22	50.06
Florida	41.97	81.01
Mid West	54.07	43.05

Notes: Calculated from data in Holland et al. (2022). Levelized costs are given by $\rho_i / \sum_t f_{it}$

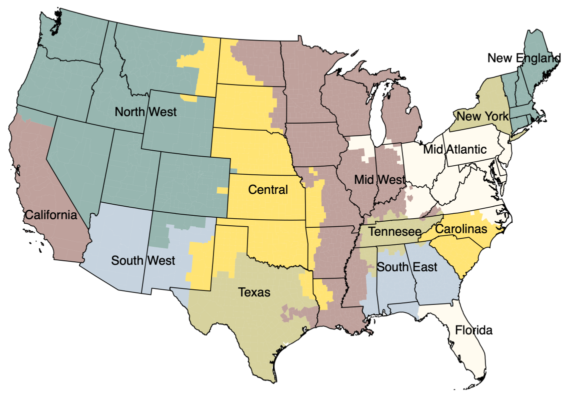

Figure 2: **Map of EIA regions.**

Other elements of the calibration, including the specification of electricity demand and hourly capacity factors for generation technologies, are taken directly from Holland et al. (2022). They assume quadratic utility functions and an hourly elasticity of demand equal to -0.15.

3.2 Storage parameters

Our primary extension relative to Holland et al. (2022) is the introduction of multiple imperfect storage technologies. Storage enables intertemporal arbitrage by shifting electricity across time and thereby interacts closely with intermittent renewable generation. We distinguish between short-duration and long-duration storage technologies, which differ in their ability to retain energy over time and in the efficiency with which stored electricity can be recovered.

We parameterize storage technologies along three key dimensions: capital cost, round-trip efficiency, and self-discharge. In the model, these features are captured by the parameters governing storage capacity costs, efficiency losses during charging and discharging, and decay of stored energy over time. Short-duration storage is characterized by relatively high round-trip efficiency but more rapid loss of stored energy, while long-duration storage exhibits lower efficiency but greater ability to retain energy over extended periods. Table 3 reports the baseline parameter values for each storage technology. These values serve as the starting point for the analysis, while selected parameters—most notably storage costs and efficiency—are varied in the simulations described in the following section.

Table 3: Baseline Parameters for Storage Technologies

	Annual Capital Cost \$/MWh	Round Trip Efficiency $\frac{1-\beta}{1+\beta}$	Self Discharge α
Short-duration	18,935	0.95	0.975
Long-duration	35,962	0.7	1

Notes: Need source of calculations.

3.3 Simulation design

We evaluate the model over a large set of parameter combinations that vary regions, technology costs, policy variables, and storage characteristics (see Table 4).

Table 4: Simulation Design: Summary of Cases

Dimension	**Parameter**	**Values**
Region	Geographic case	Texas, Mid West, Florida, New England
Renewable subsidies	Solar subsidy	{0, 0.25, 0.5, 0.75}
	Wind subsidy	{0, 0.25, 0.5, 0.75}
Carbon policy	Carbon tax (\$/ton CO_2)	{0, 100}
Storage costs	Short-duration cost discount	{0, 0.2, 0.4, 0.6, 0.8, 0.95}
	Long-duration cost discount	{0, 0.2, 0.4, 0.6, 0.8, 0.95}
Storage performance	LD round-trip efficiency	{0.5, 0.7, 0.9}

Renewable subsidies are implemented as proportional reductions in capital costs. A value of 0.5 implies capital costs are half of the baseline value. We consider carbon pricing by allowing for either no tax or a carbon tax of 100 (dollars per ton of CO_2). The tax enters through the marginal cost of fossil fuel generation. Turning to the storage technologies, we consider a wide range of capital cost reductions from baseline to a 95 percent reduction. For long-duration storage we also consider alternative values of round-trip efficiency (0.5, 0.9) that are higher and lower than baseline. All other parameters, including capital costs, generation costs, and short-duration storage performance, are held fixed at their baseline values.

In total, there are 13,824 cases. For each case, we record equilibrium outcomes including generation mix, storage capacity, electricity prices, and emissions. These outcomes form the basis for the empirical analyses in Section 4, including the characterization of storage regimes, threshold effects, and elasticity estimates.

4 Results

This section describes the equilibrium outcomes of the model and how they respond to changes in storage costs and policy parameters. The model features multiple interacting technological and policy parameters, including storage costs, efficiency, carbon pricing, and renewable subsidies. The effects of these variables depend on their joint configuration, so that the impact of any single parameter can vary across different environments. As a result, varying one parameter at a time while holding others fixed at arbitrary values can give a misleading picture of equilibrium behavior. Instead, we solve the model across a large grid of parameter combinations and use regression analysis to summarize how outcomes vary with these inputs. These regressions should be interpreted as reduced-form summaries of equilibrium behavior, capturing average relationships across a rich set of simulated equilibria.

We begin by characterizing the range of storage regimes that arise in equilibrium. We then examine how the system transitions between these regimes as costs vary, focusing on substitution patterns and local elasticities. Next we analyze the temporal utilization of storage technologies to better understand how they are used in equilibrium. Finally we show how storage costs effect emissions.

4.1 Regimes

We begin by characterizing the set of equilibrium outcomes that arise across model runs. To do so, we classify each simulation according to the composition of storage technologies used. Table 5 summarizes equilibrium outcomes across model runs by classifying each simulation into one of four storage regimes: no storage, short-duration storage only, long-duration storage only, and both technologies. The table is descriptive and is intended to illustrate the range of system configurations generated by the model rather than to isolate causal effects.

Several features of the equilibrium outcomes stand out. First, storage is ubiquitous in regions with strong renewable resources (Texas and the Mid West), while equilibria without storage arise only in regions with weaker renewable potential (Florida and New England). Second, equilibria often feature only one storage technology, although both technologies coexist in a non-trivial share of cases across all regions. Finally, these configurations are associated with distinct system outcomes, including differences in generation composition, emissions, and price variability.

Table 5: Equilibrium Outcomes by Storage Regime and Region

Region	Type	Share	$\mathcal{K}_{SD}$	$\mathcal{K}_{lD}$	Solar	Wind	Gas	CO2	PriceSD
Florida	No storage	0.05	0	0	0.11	0.00	0.89	62.51	17.92
Florida	SD only	0.61	126	0	0.50	0.26	0.23	16.19	23.57
Florida	LD only	0.21	0	193	0.59	0.21	0.20	14.02	21.29
Florida	Both	0.13	94	74	0.59	0.26	0.13	9.39	24.19
Mid West	SD only	0.67	185	0	0.19	0.52	0.27	61.41	34.40
Mid West	LD only	0.21	0	396	0.21	0.54	0.23	51.54	31.37
Mid West	Both	0.12	133	209	0.25	0.55	0.18	40.68	34.63
New England	No storage	0.08	0	0	0.02	0.16	0.75	28.85	29.44
New England	SD only	0.60	39	0	0.16	0.43	0.34	12.83	37.67
New England	LD only	0.20	0	85	0.19	0.42	0.32	12.34	35.99
New England	Both	0.11	30	56	0.22	0.48	0.23	8.73	39.25
Texas	SD only	0.65	186	0	0.36	0.48	0.15	19.15	46.64
Texas	LD only	0.21	0	343	0.40	0.48	0.11	14.07	36.46
Texas	Both	0.14	106	161	0.40	0.48	0.11	14.01	42.42

Notes: Each row corresponds to a region–storage regime pair. Storage regimes are defined as: no storage (no short- or long-duration storage), SD only (short-duration storage only), LD only (long-duration storage only), and both. "Share" indicates the fraction of model runs in each regime within a region. Generation shares are expressed as fractions of total generation. Emissions (CO_2) are reported in millions of tons. PriceSD denotes the standard deviation of electricity prices. Values are averages across simulation runs.

4.2 Transitions

Figure 3 illustrates how the viability of long-duration (LD) storage varies with the cost of short-duration (SD) storage. For each value of SD cost, we compute the maximum LD cost at which LD storage is still deployed in equilibrium—our measure of the LD viability threshold. The figure plots this threshold against the underlying SD cost, holding policy fixed and varying only the round-trip efficiency (RTE) of LD storage. Two clear patterns emerge. First, the threshold is increasing in the cost of SD storage, indicating that LD becomes viable over a wider range of costs as short-duration alternatives become more expensive. Second, higher RTE shifts the threshold upward, allowing LD storage to remain viable even at higher cost levels. While the figure provides a transparent visualization of these patterns, it is necessarily limited in how many dimensions it can display. We therefore turn next to an ordered probit specification, which allows us to quantify how policy and technological parameters shift the LD viability threshold across the full set of model outcomes.

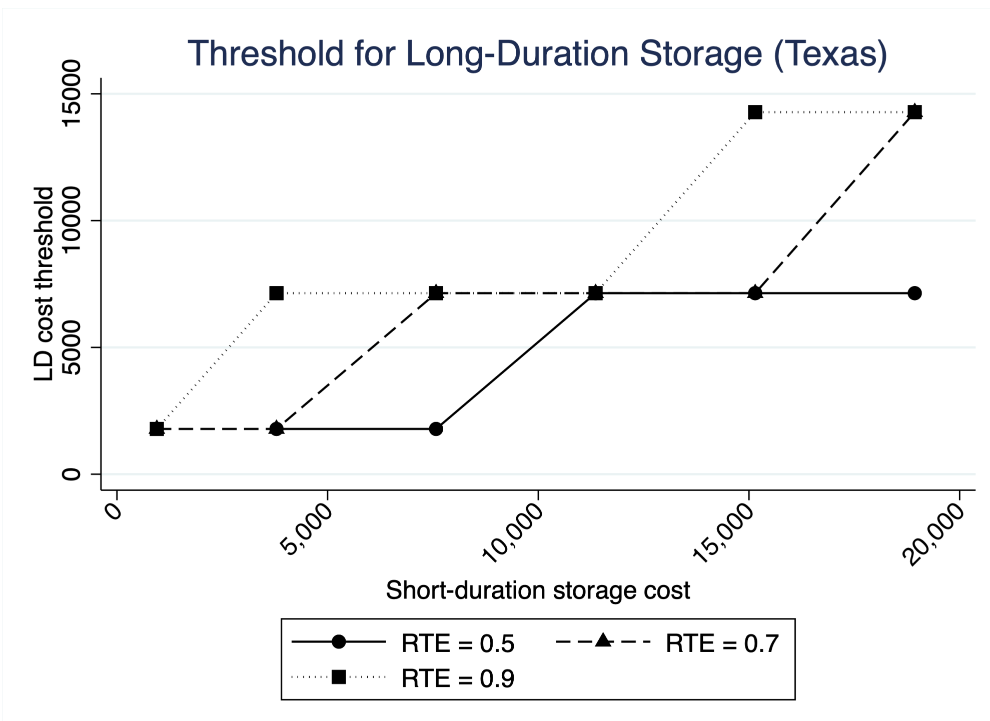

Figure 3: Threshold for Long-Duration Storage

We estimate an ordered probit model in which the dependent variable is an ordered index of the cost threshold at which long-duration storage remains viable. We report average marginal effects on the probability that this threshold lies in the highest category. Explanatory variables are indicators for short-duration cost levels, round-trip efficiency, carbon pricing, and renewable subsidies, along with region fixed effects. The results in Table 6 reveal strong nonlinearities. Increases in short-duration storage costs have negligible effects at low levels but sharply increase the likelihood that long-duration storage remains viable once costs exceed a threshold. Improvements in round-trip efficiency have large and statistically significant effects, substantially increasing the probability that long-duration storage remains viable at high cost levels. Carbon pricing also shifts the viability threshold outward, though to a lesser extent than technological improvements. Solar subsidies exhibit a similar nonlinear pattern, with limited effects at low subsidy levels and larger impacts at higher levels. Wind subsidies have little to no effect.

Because storage profitability depends on the frequency of arbitrage opportunities, technologies that generate regular temporal variation (such as solar) have a larger impact on storage demand than those that generate more diffuse variability (such as wind.)

Table 6: Marginal effects from Ordered probit regression

Variable	Δ Pr(High Threshold)	Std. Error
SD cost = 2	0.000	(0.001)
SD cost = 3	0.000	(0.000)
SD cost = 4	0.004***	(0.001)
SD cost = 5	0.043***	(0.006)
SD cost = 6	0.156***	(0.009)
RTE = 0.7	0.010***	(0.002)
RTE = 0.9	0.091***	(0.005)
Carbon tax	0.055***	(0.003)
Solar subsidy = 0.25	0.001	(0.002)
Solar subsidy = 0.5	0.014***	(0.003)
Solar subsidy = 0.75	0.027***	(0.003)
Wind subsidy = 0.25	-0.000	(0.003)
Wind subsidy = 0.5	0.001	(0.003)
Wind subsidy = 0.75	0.000	(0.003)

Notes: The dependent variable is an ordered index of the cost threshold at which long-duration storage remains viable across simulations. Reported coefficients are average marginal effects on the probability that the threshold lies in the highest category. Robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

To complement the threshold analysis, we next examine how storage technologies adjust along the intensive margin. Rather than focusing on entry and exit, we restrict attention to observations in which both technologies are used and estimate log–log regressions of long-duration (LD) and short-duration (SD) capacity. The dependent variables are the logarithms of installed LD and SD capacity, while the key regressors are the logarithms of LD and SD costs. The resulting coefficients can be interpreted as local own- and cross-price elasticities, conditional on a given technological and policy environment. Additional controls include round-trip efficiency, carbon policy, renewable subsidies, and region fixed effects, mirroring the specification used in the threshold analysis.

The results in Table 7 show that substitution between storage technologies is present, but limited and asymmetric. Both technologies exhibit large own-price elasticities, on the order of -2, indicating that capacity responds strongly to changes in its own cost. However, cross-price elasticities differ substantially: SD capacity responds strongly to increases in LD costs (elasticity = 1.3) whereas LD capacity responds more modestly to increases in SD costs (elasticity = 0.6). Consistent with earlier results, improvements in round-trip efficiency and solar subsidies increase storage deployment, while wind subsidies have little effect on LD and tend to reduce SD demand, reflecting their more diffuse temporal generation patterns.

Table 7: Elasticity Regressions

	Long Duration		Short Duration	
Variable	Coefficient	Std. Error	Coefficient	Std. Error
Log LD cost (own / cross)	-2.145***	(0.055)	1.339***	(0.068)
Log SD cost (cross / own)	0.558***	(0.045)	-2.162***	(0.052)
RTE = 0.7	0.615***	(0.062)	-0.854***	(0.076)
RTE = 0.9	0.731***	(0.090)	-0.775***	(0.076)
Carbon tax	1.405***	(0.062)	0.748***	(0.066)
Solar subsidy = 0.25	0.199**	(0.079)	0.211***	(0.079)
Solar subsidy = 0.5	0.511***	(0.077)	0.581***	(0.081)
Solar subsidy = 0.75	1.267***	(0.082)	1.287***	(0.079)
Wind subsidy = 0.25	0.053	(0.070)	-0.142*	(0.077)
Wind subsidy = 0.5	0.070	(0.071)	-0.339***	(0.074)
Wind subsidy = 0.75	0.145*	(0.080)	-0.484***	(0.087)
MIDW	0.788***	(0.076)	0.170**	(0.080)
NE	-0.963***	(0.084)	-1.639***	(0.081)
TEX	1.128***	(0.076)	0.334***	(0.082)

Notes: The dependent variables are the logarithms of long-duration (LD) and short-duration (SD) storage capacity. Robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Taken together, the threshold and elasticity results describe how storage technologies enter and scale in equilibrium. We next examine how they are used, focusing on the temporal dimension of storage and the effective duration implied by equilibrium patterns of use.

4.3 Storage Cycle Length

We examine how storage technologies are used in equilibrium by constructing a measure of cycle length for both short- and long-duration storage. We first define the average number of charge–discharge cycles per day. Let b_t^c and b_t^d denote charging and discharging flows in hour t, and let K denote installed storage capacity. The number of cycles per day is given by

$$\text{cycles per day} = \frac{1}{365 \cdot K} \cdot \frac{1}{2} \sum_{t=1}^{8760} \left| b_t^c - b_t^d \right| .$$

The factor of one-half ensures that charging and discharging flows are not double-counted, so that a full charge–discharge cycle contributes one unit. We then use the inverse of this measure—days per cycle, or cycle length—as an intuitive way to characterize how short- and long-duration storage technologies are used over time. For example, if the cycle length is equal to 7, the storage technology is, on average, fully charged and discharged once per week.

Table 8: Storage Cycle Length by Region

Region	Storage Type	N	Mean	Min	Max
Florida	LD	1143	1.34	0.56	6.00
	SD	2558	1.08	0.35	2.88
Mid West	LD	1142	2.21	0.66	5.58
	SD	2744	1.39	0.43	4.65
New England	LD	1079	2.61	0.59	7.10
	SD	2456	1.64	0.38	5.58
Texas	LD	1218	2.18	0.69	7.94
	SD	2717	1.47	0.49	6.06

Table 8 reports summary statistics for storage cycle length across regions and storage technologies. Short-duration storage is used very frequently, with typical cycle lengths on the order of one to a few days. Long-duration storage cycles less often, but still at relatively short horizons, with average cycle lengths measured in days to at most a week. Notably, there is substantial overlap across technologies, with long-duration storage often used at time scales comparable to short-duration storage. These results indicate that, in equilibrium, both storage technologies are deployed repeatedly to capture relatively short-run fluctuations rather than being reserved for infrequent, long-horizon arbitrage opportunities.

We estimate log–log regressions of long-duration (LD) and short-duration (SD) storage days per cycle, restricting the sample to observations in which both technologies are used. The dependent variables are the logarithms of cycle length, while the key regressors are the logarithms of LD and SD costs. The resulting coefficients can be interpreted as elasticities of cycle length, conditional on a given technological and policy environment. Additional controls include round-trip efficiency, carbon policy, renewable subsidies, and region fixed effects, mirroring the specification used in the capacity analysis.

The results in Table 9 show that storage cycle length responds systematically to both price and policy parameters, and that these responses differ across technologies. Increases in SD costs lead both technologies reduce cycle length, consistent with a reduction in total storage capacity that increases the value of available arbitrage opportunities. In contrast, increases in LD costs lengthen SD cycle lengths while shortening LD cycle lengths, indicating a reallocation of utilization across technologies rather than a uniform shift in intensity. Technological and policy variables also have systematic effects. Improvements in round-trip efficiency shorten cycle lengths for LD storage, indicating more frequent arbitrage, while effects on SD are smaller and less uniform. Solar subsidies lead to substantially shorter cycle lengths for both technologies, with effects increasing in magnitude at higher subsidy levels, consistent with solar generation creating regular intra-day price variation that increases arbitrage opportunities. In contrast, wind subsidies lengthen cycle times, particularly for LD storage, reflecting the more diffuse and less predictable temporal profile of wind generation. Carbon pricing modestly increases LD cycle lengths but has little effect on SD. Overall, these results reinforce that storage duration is not fixed by technological characteristics: both technologies primarily engage in relatively frequent cycling, with the effective duration of storage determined endogenously by equilibrium conditions and the temporal structure of supply and demand.

These results reinforce a central theme of the paper: storage technologies provide the capacity to shift energy across time, but the effective duration over which that capacity is used is determined endogenously by equilibrium conditions. As costs fall, storage is applied over longer horizons, even though its underlying technological characteristics are unchanged.

Overall, our results show relatively short cycle lengths, indicating that long-duration storage is not used for seasonal balancing. There are several features of the model that provide alternative margins for adjustment across seasons and thereby reduce the need for seasonal storage. First, demand is not perfectly inelastic, so consumption can adjust in response to seasonal variation in renewable availability. Second, most equilibria include a substantial amount of dispatchable generation, which provides an additional margin for responding to seasonal fluctuations of renewables. Finally, wind and solar generation exhibit complementary seasonal patterns, with wind tending to have higher capacity factors in the winter and solar in the summer. Generating seasonal storage would require restricting these alternative adjustment margins.

Table 9: Cycle Length Regressions

	Long Duration		Short Duration	
Variable	Coefficient	Std. Error	Coefficient	Std. Error
Log SD cost	-0.097***	(0.012)	-0.464***	(0.012)
Log LD cost	-0.323***	(0.014)	0.159***	(0.014)
RTE = 0.7	-0.119***	(0.014)	-0.091***	(0.015)
RTE = 0.9	-0.340***	(0.020)	0.069***	(0.021)
Carbon tax	0.114***	(0.013)	-0.002	(0.013)
Solar subsidy = 0.25	-0.118***	(0.017)	-0.087***	(0.018)
Solar subsidy = 0.5	-0.252***	(0.017)	-0.143***	(0.017)
Solar subsidy = 0.75	-0.413***	(0.018)	-0.156***	(0.017)
Wind subsidy = 0.25	0.058***	(0.017)	0.011	(0.016)
Wind subsidy = 0.5	0.148***	(0.016)	0.009	(0.016)
Wind subsidy = 0.75	0.244***	(0.018)	0.063***	(0.019)
Mid West	0.385***	(0.017)	0.155***	(0.017)
New England	0.486***	(0.019)	0.261***	(0.019)
Texas	0.389***	(0.018)	0.172***	(0.017)

Notes: The dependent variables are the logarithms of long-duration (LD) and short-duration (SD) days per cycle. Robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

4.4 Generation and Emissions

We next examine the relationship between storage costs and emissions. Emissions are closely proportional to total fossil generation in our model, though not exactly one-for-one due to differences in emissions intensity between gas combined cycle and gas peaker. The differences lead to trivial differences in elasticities, so for simplicity we focus on total fossil generation as a proxy for emissions.

Table 10: Elasticity Decomposition of Fossil Generation

	Short-duration (SD)	Long-duration (LD)
Fossil generation	0.215	0.215
Fossil capacity	0.194	0.205
Implied utilization	0.021	0.011

Notes: Generation and capacity elasticities are estimated from log–log regressions. All regressions include controls for round-trip efficiency, carbon policy, renewable subsidies, and region fixed effects. The utilization elasticity is computed as the difference between generation and capacity elasticities.

Table 10 shows the results from estimating the elasticity of total fossil generation with respect to short- and long-duration storage costs. We find that fossil generation responds similarly to both technologies: a 1 percent increase in either SD or LD costs raises generation by 0.215 percent.

To better understand the mechanism behind these effects, we decompose the response of fossil generation into capacity and utilization margins (utilization is the ratio of generation to capacity). In the appendix, we show that the elasticity of generation can be expressed as the sum of the elasticities of capacity and utilization. Table 10 reports the results of this decomposition. The reduction in generation is driven almost entirely by changes in fossil capacity, with only a small role for fossil utilization. Despite the similarity in total elasticities, the two technologies operate along slightly different margins. Long-duration storage has a somewhat larger effect on fossil capacity, whereas short-duration storage has a relatively larger effect on utilization.

5 Conclusions

Taken together, our results suggest a reinterpretation of how storage technologies should be understood in equilibrium. Our findings highlight that the effective time horizon over which storage is used is also shaped by economic conditions. In particular, storage duration is not an intrinsic technological attribute, but an endogenous equilibrium outcome determined by the interaction of storage costs, system variability, and arbitrage opportunities. Technologies provide the capacity to store energy over different horizons, but equilibrium determines how that capacity is actually deployed. As storage becomes cheaper and more widespread, this distinction becomes increasingly important: the same technology can be used to balance short-run fluctuations or longer-run imbalances depending on the economic environment. Recognizing this distinction is central to understanding the role of storage in future power systems, where economic conditions—not just technological capabilities—will determine how energy is shifted across time.

References

Andrés-Cerezo, David, and Natalia Fabra (2026). "Storage and Renewable Energies: Friends or Foes?" *The Economic Journal*, published online April 21, 2026. DOI: https://doi.org/10.1093/ej/ueag048.

Butters, R. Andrew, Dorsey, Jackson, and Gowrisankaran, Gautam (2025). "Soaking Up the Sun: Battery Investment, Renewable Energy, and Market Equilibrium." *Econometrica*, 93(3), 891–927.

Dowling, Jacqueline A., Rinaldi, Katherine Z., Ruggles, Tyler H., Davis, Steven J., Yuan, Mengyao, Tong, Fan, Lewis, Nathan S., and Caldeira, Ken. (2020). "Role of Long-Duration Energy Storage in Variable Renewable Electricity Systems." *Joule*, 4(9), 1907–1928.

Holland, Stephen, Mansur, Erin, and Yates, Andrew. (2022) Decarbonization and Electrification in the Long Run, NBER working paper w30082.

Jenkins, Jesse D., and Sepúlveda, Nestor A. (2021). "Long-Duration Energy Storage: A Blueprint for Research and Innovation." *Joule*, 5(9), 2241–2246.

Junge, Cristian, Mallapragada, Dharik, and Schmalensee, Richard (2022). "Energy Storage Investment and Operation in Efficient Electric Power Systems." *The Energy Journal*, 43(6), 1–24.

Karaduman, Ömer. (2023). *Economics of Grid-Scale Energy Storage in Wholesale Electricity Markets*. Stanford Graduate School of Business.

Selänniemi, Andreij, Hellström, Magnus, and Björklund-Sänkiaho, Margareta. (2024). "Long-Duration Energy Storage—A Literature Review on the Link between Variable Renewable Energy Penetration and Market Creation." *Energies*, 17(15).

Tarel, Guillaume, Korpås, Magnus, and Botterud, Audun (2024). "Long-term equilibrium in electricity markets with renewables and energy storage only." *Energy Systems*. https://doi.org/10.1007/s12667-024-00654-y

Appendix

Decomposition of Emissions Responses

We now describe the details of the decomposition used to separate the effects of storage costs on fossil generation into capacity and utilization margins. The approach is based on the accounting identity that relates total fossil generation to installed capacity and its utilization rate.

Let G_i denote total fossil generation and K_i denote total installed fossil capacity in equilibrium i. Total generation is obtained by aggregating over time:

$$G_i = \sum_t g_{it}, \tag{2}$$

where g_{it} denotes fossil generation in period t.

We define the capacity factor of fossil generation in equilibrium i as the ratio of total generation to installed capacity:

$$CF_i = \frac{G_i}{K_i} = \frac{\sum_t g_{it}}{K_i}. \tag{3}$$

This object measures the average utilization of installed fossil capacity.

Using this definition, total generation can be written as

$$G_i = CF_i \cdot K_i, \tag{4}$$

which implies

$$\log G_i = \log K_i + \log CF_i. \tag{5}$$

Estimated elasticities. To quantify the response of the system to changes in storage costs, we estimate the following log–log regressions:

$$\log G_i = \alpha^G + \beta^G_{SD} \log c_{SD,i} + \beta^G_{LD} \log c_{LD,i} + X_i' \gamma^G + \varepsilon^G_i, \tag{6}$$

and

$$\log K_i = \alpha^K + \beta^K_{SD} \log c_{SD,i} + \beta^K_{LD} \log c_{LD,i} + X_i' \gamma^K + \varepsilon^K_i, \tag{7}$$

where $c_{SD,i}$ and $c_{LD,i}$ denote the costs of short- and long-duration storage, and X_i includes controls for policy and technology parameters.

The coefficients β^G_{SD} and β^G_{LD} measure the elasticity of fossil generation (and hence emissions) with respect to storage costs, while β^K_{SD} and β^K_{LD} measure the elasticity of installed fossil capacity.

Implied utilization elasticities. Combining equations (5)–(7) implies that the elasticity of the capacity factor can be expressed as the difference between the generation and capacity elasticities:

$$\beta^{CF}_j = \beta^G_j - \beta^K_j, \quad j \in \{SD, LD\}. \tag{8}$$

That is, the response of utilization to storage costs is not estimated directly, but inferred from the two regressions above.

Interpretation. This decomposition separates the response of fossil generation into two conceptually distinct channels. The capacity elasticity β^K_j captures changes in long-run investment, reflecting how storage affects the need for fossil generation capacity. The utilization elasticity β^{CF}_j captures changes in how intensively existing fossil plants are used. Because emissions are approximately proportional to fossil generation in our model, this decomposition also applies to the emissions results reported in the main text.

In practice, we find that the capacity component accounts for the vast majority of the total response, indicating that reductions in storage costs primarily affect emissions through changes in investment decisions rather than through sustained changes in utilization.