MIT CEEPR
Center for Energy and Environmental Policy Research

Working Paper Series

How Uncertainty Shapes Electricity Storage Decisions: Dispatch Policies and Capacity Portfolios under Renewable Drought Risk

Cem Keske, Bjarne Steffen, and John E. Parsons

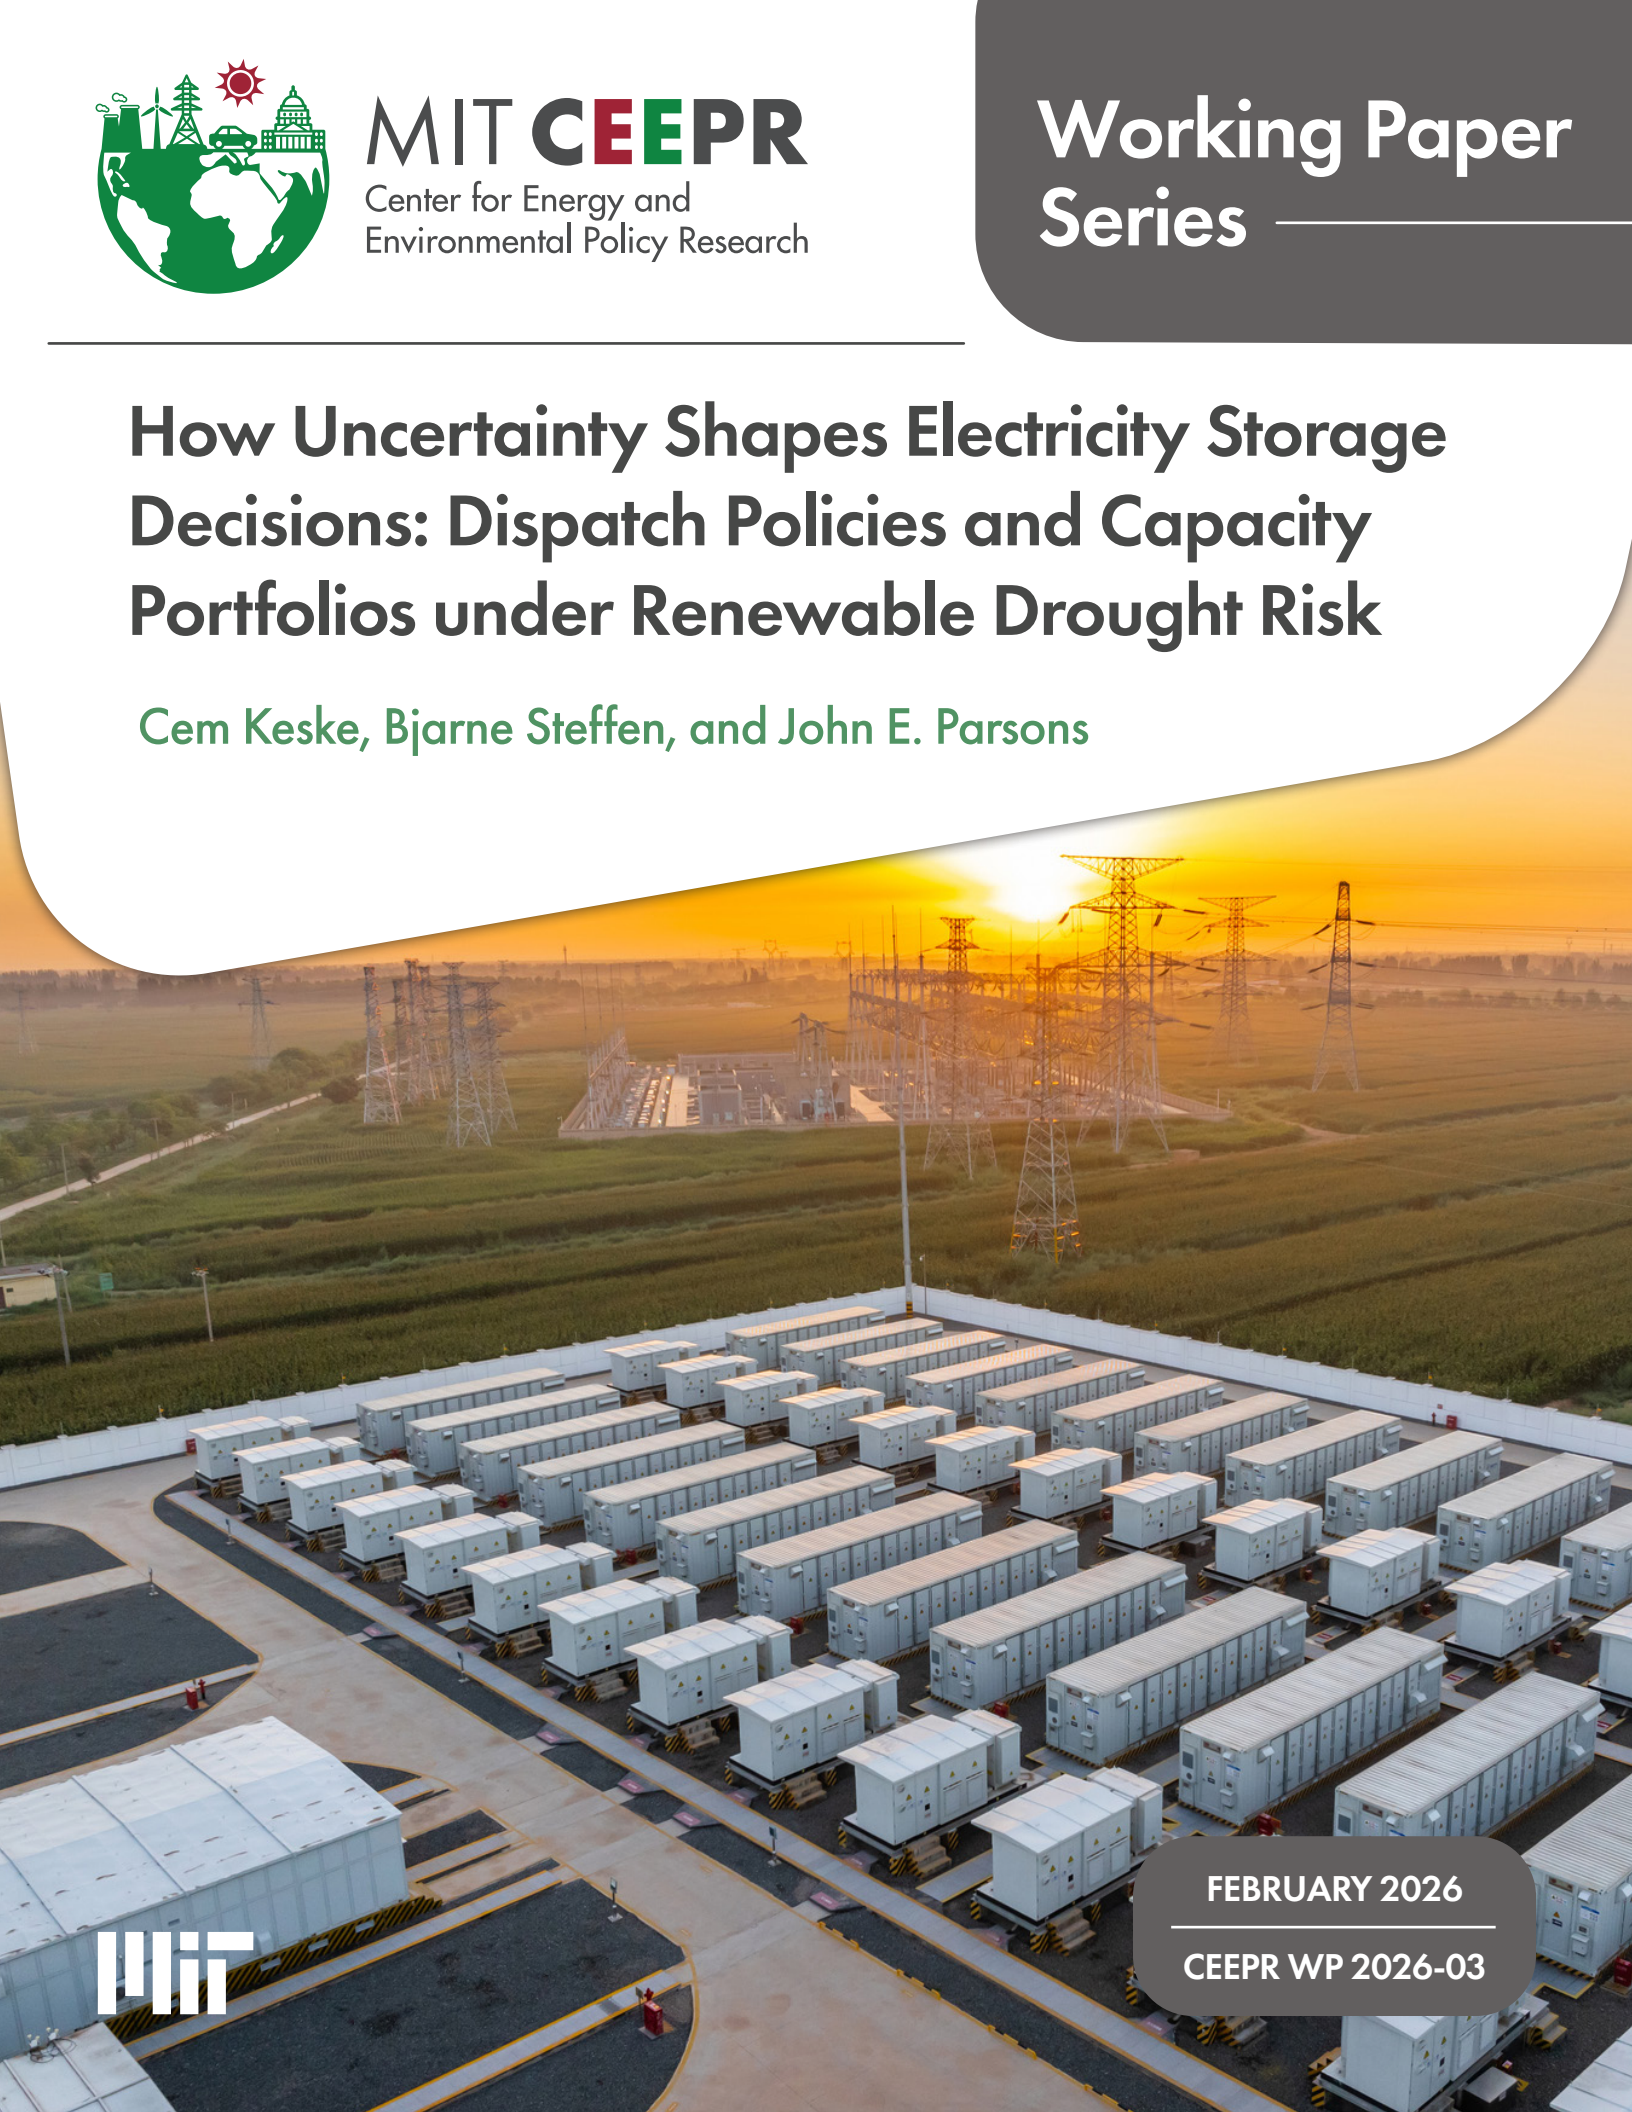

FEBRUARY 2026

CEEPR WP 2026-03

Working Paper Series.

Since 1977, the Center for Energy and Environmental Policy Research (CEEPR) has been a focal point for research on energy and environmental policy at MIT. CEEPR promotes rigorous, objective research for improved decision making in government and the private sector, and secures the relevance of its work through close cooperation with industry partners from around the globe. Drawing on the unparalleled resources available at MIT, affiliated faculty and research staff as well as international research associates contribute to the empirical study of a wide range of policy issues related to energy supply, energy demand, and the environment.

An important dissemination channel for these research efforts is the MIT CEEPR Working Paper series. CEEPR releases Working Papers written by researchers from MIT and other academic institutions in order to enable timely consideration and reaction to energy and environmental policy research, but does not conduct a selection process or peer review prior to posting. CEEPR's posting of a Working Paper, therefore, does not constitute an endorsement of the accuracy or merit of the Working Paper. If you have questions about a particular Working Paper, please contact the authors or their home institutions.

How Uncertainty Shapes Electricity Storage Decisions: Dispatch Policies and Capacity Portfolios under Renewable Drought Risk

Cem Keske,* Bjarne Steffen,† and John Parsons‡

February 12, 2026

Abstract

With increasing shares of intermitting renewables, electricity storage assets are increasingly considered as a way to ensure resource adequacy. Using storage assets to that end, however, involves trade-offs with time-spread arbitrage. To study these trade-offs, we model the optimal dispatch of an electricity storage asset in the face of uncertain renewable generation. We then embed this within a classic capacity expansion model. We use a parameter driven stochastic process to model the uncertainty, which enables us to analyze how the inter-temporal structure of the uncertainty shapes the optimal dispatch and the optimal portfolio of capacities. We show how the optimal dispatch policy changes with the decay time of renewable drought episodes, and how it varies with the storage asset's duration and round-trip efficiency. We also illustrate how the dispatch of one storage asset depends on the features and state of charge of other storage assets in the system, and demonstrate conditions under which one asset can charge while another is discharging. Finally, we show how these drivers affect the optimal portfolio of capacities, system cost, and loss-of-load events.

*ETH Zürich and MIT Center for Energy and Environmental Policy Research, MIT Building E19-411 400 Main Street, 4th Floor, Cambridge, MA 02142-1017, ckeske@mit.edu

†ETH Zurich, Albert Einstein School of Public Policy and Energy and Technology Policy Group, Clausiusstrasse 37, 8092 Zurich, Switzerland, bjarne.steffen@gess.ethz.ch

‡MIT Sloan School of Management and MIT Center for Energy and Environmental Policy Research, MIT Building E19-411 400 Main Street, 4th Floor, Cambridge, MA 02142-1017, jparsons@mit.edu

This paper is based upon Cem Keske's Masters Thesis, Optimal Operation of Grid-scale Electricity Storage for Resource Adequacy, March 1, 2025.

The authors acknowledge research funding from the MIT Energy Initiative's Future Energy Systems Center under the project "Stored Energy Reserve Market for Grid Resource Adequacy". We also appreciate the expert research assistance of Natalie Hatzigeorgiou.

1 Introduction

Electricity storage can serve the electric grid in a variety of ways. Historically, it has enabled utilities to exploit of economies of scale in generation capacity in the face of predictable diurnal variations in hourly load. Different storage technologies have been matched with different generation technologies. The large-scale expansion of pumped hydro storage in many countries between 1960 and 1990 complemented investments in baseload nuclear and coal plants. Thermal storage has been integrated into concentrated solar power plants to economize on the size of the power island in the face of dramatic, predictable diurnal variations in insolation. Grid-scale battery storage is proving a valuable complement to the large scale penetration of solar PV, especially with regard to managing the regular evening ramp. Storage is also useful to manage the stochastic variation in load and renewable generation at short time scales. In many electric grids, battery storage is becoming an important source of frequency regulation and fast response operating reserves.

Attention has also begun to be paid to how storage assets can help manage stochastic variation at longer time scales. Hydro reservoir dominated grids have long managed the inventory of water in the face of seasonal and interannual variation in water inflows. Now other grid operators managing the increasing penetration of wind and solar PV generation are learning how to assess the danger of renewable resource droughts and attend to the stock of stored energy in their systems. Many grid operators have been working to update their metrics for resource adequacy. Capacity markets in many systems are being adapted to incorporate more sophisticated measures of how each resource contribute to system capacity, including storage assets (a.k.a. energy duration limited resources).

Our model is tailored to help illuminate how storage serves resource adequacy, especially in light of the danger of renewable resource droughts. We eliminate any deterministic calendar variations (e.g., daily cycles), and focus exclusively on optimization against a stationary uncertain inflow of renewable energy. In spirit then, this is a model more inspired by the volatility characterizing wind dominated systems than solar PV dominated systems. Our model is constructed with hourly granularity, focusing on the provision of energy, so we do not consider the provision of services at a shorter time scale such as frequency regulation, operating reserves, or other ancillary services. Subject to that proviso, the storage asset is optimized in the face of variability at all other time scales, whether hourly variations or the long-run danger of droughts.

Using a parametric model of uncertainty allows us to efficiently characterize key economic tradeoffs shaping the optimal dispatch strategy. For example, we can explicitly calculate the conditional probability of negative net load and therefore describe the tradeoff involved in discharging when net load is low, as well as the tradeoff in setting the optimal state-of-charge when net load is high and loss of load events threaten. Our solution of the optimal dispatch policy makes explicit how storage duration creates value - by enabling energy

to be carried through sustained deficits - and it implies diminishing marginal value of state-of-charge, since additional stored energy is most valuable when the battery is scarce and less valuable as it approaches full.

Owing to our parameterization of the renewable stochastic process, we decouple the expected decay time of renewable inflow shocks from the long-run distribution of renewable production. We show that:

- As the decay time of the renewable droughts increases - and thus their expected duration increases - the incentive to maintain a high state-of-charge increases.
- A higher round-trip efficiency of the storage asset increases its activity.
- With two storage assets that differ in duration and round-trip efficiency, efficiency governs activity near potential curtailment, whereas duration governs activity near scarcity events.
- Under certain system conditions, it is optimal for one storage asset to charge while the other discharges; more generally, each asset's dispatch is conditional on the other's state-of-charge.
- In a multi-storage system, the marginal value of stored energy in any one storage unit depends on the state of charge of the others. Therefore, the state of charge of one storage unit affects the optimal valuation of energy in the other storage units.
- For multiple storage assets with similar duration and round-trip efficiency serving the same system, system-optimal dispatch (guided by the optimal valuation of stored energy) tends to rebalance state-of-charge across units, equalizing their SoCs and returning to a symmetric (lock-step) operating point.
- Complementary investments in thermal capacity adjust endogenously to the availability of storage.
- In our framework, perfect foresight overstates the value of storage relative to stochastic optimization by much more than prior estimates, likely because we exclude predictable diurnal structure in net load. With diurnal patterns present, perfect foresight and stochastic optimization respond similarly, diluting the differences driven by unpredictable variability that our model highlights.

The next section provides a brief literature review. Section 3 presents our modeling framework, while Section 4 describes the methodology for solving the optimal dispatch problem under our framework. In Section 5, we extend the dispatch optimization to effective-capacity optimization. Section 6 reports the numerical results and discusses their implications for power systems.

2 Literature Review

Our model is closest to the work of Geske and Green (2020), who optimize storage operation and thermal plant capacity in the face of an uncertain, but diurnal net load estimated from historical data. We differ from Geske and Green (2020) in two respects. First, we abstract completely from predictable diurnal patterns and focus exclusively on the uncertainty of episodic droughts of renewable inflows. Second, we use a parametric representation of uncertainty which enables us to identify how both storage dispatch and thermal capacity investments vary as either the likelihood or the expected duration of droughts increases.

There are many strands of literature on the place of storage assets in grid optimization. Steffen and Weber (2013), Schmitz et al. (2013), Korpås and Botterud (2020), and Antweiler and Müsgens (2024) analyze the role of storage when the size of the energy reservoir is unlimited.

Another strand, including McConnell et al. (2015), Junge et al. (2022) and Ruhnau and Qvist (2022) model the limited energy capacity, but employ a perfect foresight assumption to model optimization of storage dispatch. There are many variations on this strand. For example, Moret (2017) employs multiple scenarios and robust optimization. Other work on stochastic optimization include, for example, Gorenstin et al. (1993), Evans et al. (2013), Grillo et al. (2016), Williams et al. (2019), and Xu et al. (2020).

Alongside these different optimization models, is a line of research attempting to capture the complex intertemporal dynamics of renewable inflows. For example, Glasbey (2001), Chen et al. (2010), and Pfenninger and Staffell (2016).

Especially relevant to our research is the literature on the challenge of defining and quantifying renewable energy "droughts". This includes Bhatnagar et al. (2022), Ruhnau and Qvist (2022), Somani et al. (2024), and Bracken et al. (2024). Ruhnau and Qvist's key insight is that multiple scarcity periods can occur in close succession, preventing storage systems from fully recovering between events. To quantify this, they introduce the concept of maximum energy deficit, which accounts for both the duration and cumulative impact (volume) of scarcity events. They document the impact on capacity investments recommended from the widely used perfect foresight models. Our model implicitly does the same, however, the storage dispatch is stochastically optimized.

3 A Model of a Grid with Stochastic Renewable Energy Inflows

We model an electric grid with a variable renewable energy technology subject to a stochastic availability factor. The grid also has a set of thermal generation technologies that can be represented in a simple supply stack, ordered by the fixed marginal cost of generation. Finally, the grid also has a set of energy

storage technologies. The grid is modeled as a copper plate, so the effects of transmission constraints are neglected. The load is assumed to be constant over time and exogenously fixed, so all variability and uncertainty arises from the availability of the renewable energy availability.

We start with an exogenously specified set of capacities. The renewable generation is exogenous, and the optimal dispatch of the thermal generation can be easily specified, so the remaining challenge is to solve for the optimal dispatch strategy for the storage technologies.[1] We take a system perspective, so the optimization is designed to minimize system cost.[2] We analyze how the strategy varies with the parameters describing the stochastic process governing the uncertain inflows of variable renewable energy and the technical parameters of the battery.

After solving the optimal dispatch problem, we move on to consider the optimal portfolio of generation and storage capacities. We keep the capacity of VREs fixed, and solve for the complementary portfolio in other technologies. We analyze how the complementary portfolio varies with the parameters describing the stochastic process governing the uncertain inflows of variable renewable energy and the technical parameters of the battery, and we document the impact on system cost and loss of load events, among other metrics.

3.1 Stochastic Energy Inflows

We model an uncertain hourly resource availability factor $X_t^{\text{avl}} \in [0, 1]$. Given the renewable generation capacity, $\mathcal{R}$ (MW), this translates into uncertainty in the hourly renewable generation, R_t (MWh), with

$$R_t = \mathcal{R} \cdot X_t^{\text{avl}}. \tag{1}$$

The constant hourly load is D' (MWh), and we define net load, D_t as the difference between load and renewable generation,

$$D_t = D' - R_t. \tag{2}$$

3.1.1 Availability Factor

We model the renewable 'availability factor' X_t^{avl} as a bounded mean reverting process, more specifically, as an Ornstein-Uhlenbeck (OU) process truncated to the interval [0,1]. We construct X_t^{avl} from the unbounded OU process,

$$dX_t = \theta_X(\mu_X - X_t)dt + \sigma_X dW_t \tag{3}$$

where W_t is a Wiener process. This process has three parameters:

[1] Given the resulting net load that should be met by the thermal generators, thermal generation dispatch is straightforward to determine as it follows the merit order and is therefore not an endogenous decision variable in the model.

[2] Under certain simplifying assumptions, the system optimal dispatch strategy is also the profit maximizing strategy for a private owner of the battery. In future work we explore the profit maximizing strategy under alternative assumptions about the wholesale market design and how it deviates from the system optimum.

- μ_X, the mean to which the process reverts,
- σ_X, the instantaneous volatility, and,
- θ_X, the speed of mean reversion.

Several considerations motivate our use of the OU process. First, we are interested in exploring how changes in the structure of the uncertainty shape the outcome of the optimization, so we choose to represent the uncertainty using a parametrically defined process.

Second, the OU process has an exact discretization over time obtained by sampling the continuous-time solution using the following recursive formula:

$$X(t+\Delta t) = \mu_X + e^{-\theta_X \Delta t}\left(X(t) - \mu_X\right) + \sigma_X \sqrt{\frac{1-e^{-2\theta_X \Delta t}}{2\theta_X}} Z_t, \tag{4}$$

where Z_t are i.i.d. standard normal random variables: $Z_t \sim \mathcal{N}(0,1)$.

Third, the OU process has convenient mathematical properties, such as a Gaussian long-term distribution. This not only simplifies calculations but also facilitates finding analytical expressions for key values, such as the expected number of hours in a year with lost load. The OU process' stationary distribution (also called the 'long-term' distribution, or distribution after mixing) is: $X_\infty \sim \mathcal{N}(\mu_X, \Sigma_X^2)$, where $\Sigma_X^2 = \sigma_X^2/2\theta_X$ is the stationary variance.

We focus on parameterizations which fix the long-term distribution while altering the intertemporal dynamics that influence the optimal strategy of a storage asset with limited energy capacity. To help define our set of parameterizations, we first translate the parameter for the speed of mean reversion to a parameter measuring time decay, τ_{95_X},

$$\tau_{95_X} \equiv -\ln(0.05)/\theta_X, \tag{5}$$

which is the time it takes for the expected value of the process to decay by 95% towards its mean, conditioned on the current value. It can be interpreted as an expected duration of a period of shortage (a drought) or a period of surplus. We now shift from describing the OU process by the parameters $[\mu_X, \sigma_X, \theta_X]$ to describing it by the three parameters $[\mu_X, \Sigma_X, t_{95_X}]$, with μ_X and Σ_X specifying only the long-term probability distribution of the process and τ_{95_X} specifying only the intertemporal dynamics. In our parameterizations we hold μ_X and Σ_X constant and focus on varying τ_{95X}. Since varying τ_{95X} is equivalent to varying θ_X, in order to simultaneously hold Σ_X constant, we have to simultaneously vary σ_X in lockstep,

$${\sigma_X}^2 = -2\ln(0.05) \cdot \Sigma_X \cdot \tau_{95_X}. \tag{6}$$

In this case, shortening the decay time (increasing the speed of mean reversion) requires simultaneously increasing the instantaneous volatility in order to keep the long-term distribution unchanged. As the duration of droughts shortens, their frequency increases.

Fourth, and especially important to our methodology for identifying the optimal strategy, the OU process is amenable to approximation with a bounded state space while simultaneously retaining the extreme events of interest to this study, droughts. Droughts involve *persistently* low availability across multiple hours rather than isolated extreme values. Truncating the process between its six-sigma values slightly reduces the likelihood of extreme moves in any single hour without significantly reducing the likelihood of drought events. With a bounded state space, we can employ numerical methods that approximate it as a finite state space.

The renewable resource availability factor, X_t^{avl}, evolves as a truncated OU process with support $[\mu_X - 6\Sigma_X, \mu_X + 6\Sigma_X]$. Specifically, we fix $\mu_X = 0.5$ and $6\Sigma_X = 0.5$ (i.e., $\Sigma_X = 1/12$), so that the state space for the availability factor is the unit interval, $X_t \in [0, 1]$. Observing a renewable availability of either 0 or 1 in a given hour are each six-sigma events. During simulations, the values that occur outside of these bounds are truncated to these bounds, that is, the probability of jumping out of the state space is embedded in staying at the boundary value of six-sigma. The effect of this rounding is negligible as the probability of observing a value outside of these bounds is less than 1 hour per 100 years. The long-term average of the availability factor is 50%. Observing an availability factor below 25% and above 75% are three-sigma events. This means that, on average, the availability factor will be larger than 75% for only 12 hours in a year.

Figure 1 illustrates three processes with different decay times (τ_{95_X}), each returning to its stationary state after an unlikely event ($x_0 = \mu_X - 3\Sigma_X$). All processes share the same mean ($\mu_X = 0.5$) and stationary standard deviation ($\Sigma_X = 1/12$), ensuring identical long-term expected values and confidence intervals in the stationary state. When a 3-sigma event occurs, there is always a chance that it could escalate into an even more extreme deviation, exceeding $3\Sigma_X$. In the fastest-decaying process ($\tau_{95_X} = 1$ day), extreme deviations typically disappear within a day. Conversely, in the slowest process ($\tau_{95_X} = 1$ week), there is a probability that deviations will persist and even intensify over several days before stabilizing. Since all processes share the same μ_X and Σ_X, they exhibit the same long-term distribution, but different intertemporal dynamics.

To further illustrate how decay time influences intertemporal dynamics, Figure 2 compares two time series simulations with same long-term distribution ($\mathcal{N}(0.5, (1/12)^2)$), but different decay times (one day and two weeks). The process with a shorter decay time exhibits frequent fluctuations, making it a *fast-moving process*, whereas the process with a longer decay time fluctuates less often, resulting in a *slow-moving process*. The fast-moving process enters scarcity states far more frequently than the slow-moving one. However, while it experiences more scarcity events, these events tend to be shorter in duration. Conversely, the slow-moving process encounters scarcity events less often but remains in those states for extended periods. Since both processes share the same stationary distribution, their total cumulative scarcity volume over the

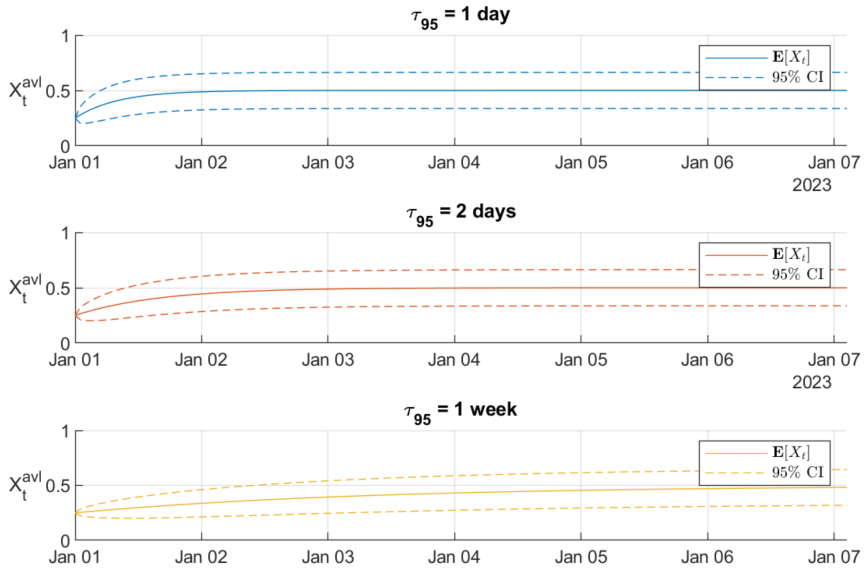

Figure 1: Decay and Mixing Time of Renewable Availability Process. The probabilistic evolution of three processes is illustrated, each simulated with decay times $\tau_{95_X} = 1$ day, 2 days, and 1 week. All processes share the same stationary distribution $\mathcal{N}(\mu_X, \Sigma_X^2) = \mathcal{N}(0.5, (1/12)^2)$ and start from a drought state of $3\Sigma_X$ (availability factor of 0.25). The mixing time, defined as the time required for each process to approach its stationary state, is approximately equal to its respective decay time: 1 day, 2 days, and 1 week.

long term remains the same. However, the distribution of scarcity events differs:

- The *fast-moving process* experiences *frequent but smaller* scarcity events.
- The *slow-moving process* experiences *fewer but larger* scarcity events.

Figure 3 presents the measured frequency distributions of 50-year simulations of processes with different decay times but the same μ_X and Σ_X. Despite their differing intertemporal behaviors, all processes converge to the same stationary distribution, $\mathcal{N}(0.5, (1/12)^2)$. While decay time affects short-term volatility and the persistence of scarcity periods, it does not alter the long-term statistical properties of renewable availability.

3.1.2 Decay rate and energy capacity

One strand of the storage literature has shown how load-duration curve (LDC) methods can be used to determine the efficient storage power capacity of the system - see Steffen and Weber (2013), Schmitz et al. (2013), Korpås and Botterud (2020), and Antweiler and Müsgens (2024). These methods rely on the long-term distributions of stochastic parameters rather than intertemporal dynamics, providing a solution for sizing the power capacity of the storage under the assumption of infinite energy capacity. In our modeling framework, while

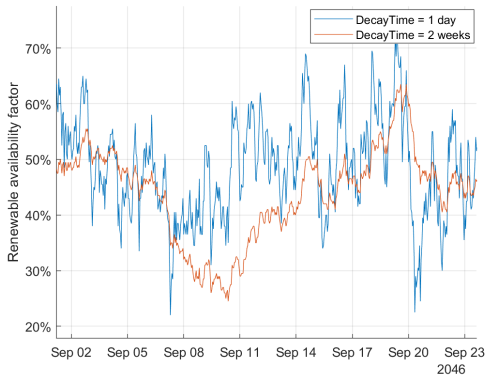

Figure 2: Time series of two simulated renewable availability processes with different decay times. The blue process has a decay time of 1 day, while the orange process has a decay time of 2 weeks. Both processes share the same stationary (long-term) distribution, $\mathcal{N}(0.5, (1/12)^2)$

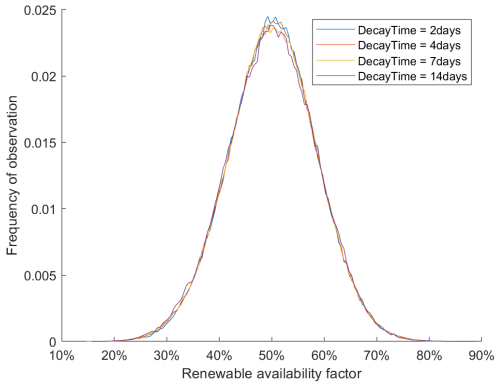

Figure 3: Long-term distribution of the renewable availability factor for different decay times (τ_{95_X} = 2 days, 4 days, 1 week and 2 weeks). The distribution is derived from a simulation spanning 50 years. All processes have the same stationary distribution, $\mathcal{N}(\mu_X, \Sigma_X^2) = \mathcal{N}(0.5, (1/12)^2)$.

the process parameters μ_X and Σ_X influence the long term distribution, the decay rate does not. Therefore, the decay rate (τ_{95_X}) is a parameter that primarily affects and interacts with limited energy capacity in storage. While μ_X and Σ_X are primarily linked to the optimal power capacity of the storage, the effect of τ_{95_X} is only observed when energy capacity is constrained. This parameter will be later used to analyze the operational behaviors caused by the limited capacity of the storage.

3.1.3 Renewable Generation

Since renewable generation is the product of the fixed quantity of renewable capacity and the availability factor, it also follows a bounded OU process with $R_t \in [0, \mathcal{R}]$ and process parameters, $[\mu_R, \Sigma_R, t_{95_R}]$,

- $\mu_R = \mathcal{R} \cdot \mu_X = \mathcal{R}/2$, is the mean renewable generation,
- $\Sigma_R = \mathcal{R} \cdot \Sigma_X = \mathcal{R}/12$ is the long-term volatility of renewable generation, and
- $t_{95_R} = t_{95_X}$ is the 95% decay time (therefore θ_X remains unchanged).

It can be simulated using the formula:

$$R(t + \Delta t) = \mathcal{R}\mu_X + e^{-\theta_X \Delta t}\left(R(t) - \mathcal{R}\mu_X\right) + \mathcal{R}\sigma_X \sqrt{\frac{1 - e^{-2\theta_X \Delta t}}{2\theta_X}} Z_t, \qquad (7)$$

where $Z_t \sim \mathcal{N}(0, 1)$ are standard normal random variables.

3.1.4 Net Load

Hourly load, D', is constant, so net load, $D_t = D' - R_t$, is a simple translation of the renewable generation, and therefore is also a bounded OU process on the state space $D_t \in \mathcal{D} \equiv [D' - \mathcal{R}, D']$, with parameters, $[\mu_D, \Sigma_D, t_{95_D}]$,:

- $\mu_D = D' - \mathcal{R}/2$, the mean of the net load,
- $\Sigma_D = \Sigma_R = \mathcal{R}/12$, the long-term standard deviation of the net load process, and
- $t_{95_D} = t_{95_R} = t_{95_X}$, is the 95% decay time (therefore θ_X remains unchanged).

Thus, the net load D_t can also be simulated using the recursive OU equation, now adjusted for the constant demand D':

$$D(t + \Delta t) = \mu_D + e^{-\theta_X \Delta t} \left(D(t) - \mu_D \right) + \mathcal{R}\sigma_X \sqrt{\frac{1 - e^{-2\theta_X \Delta t}}{2\theta_X}} Z_t, \tag{8}$$

where $Z_t \sim \mathcal{N}(0, 1)$.

3.2 Thermal Generator Supply Stack

Alongside the variable renewable generation and the storage, there are N thermal generators (indexed n) that can meet the demand with their hourly power production $(g_t^1, ..., g_t^N)$, with marginal generation costs c_v^1 to c_v^N, $c_v^n > 0$. Each generator's production is limited by its power capacity $(g_{\max}^n)$ and is non-negative. These assumptions reduce the thermal generation problem to a simple supply stack in which thermal generator capacity is ordered by marginal cost and load net of renewable generation and storage action is supplied by the lowest cost generators in the stack. Any unmet demand is denoted as l_t^{LL} and is valued at c_{LL}, the Value of Lost Load (VoLL).

3.3 Storage

There are M storage units (indexed m), each with an energy capacity of $e_{\max}^m$, power capacity of $p_{\max}^m$ (same both for charging and discharging), round-trip efficiency, η_m, and initial state of charge, $e_0^m \in \mathcal{E}^m \equiv [0, e_{max}^m]$. Given the state of charge, e_t^m, at time t, the feasible range of charge (when positive) or discharge (when negative) is $p_t^m \in [-e_t^m, e_{max}^m - e_t^m]$. The intertemporal storage constraint relates the energy in the storage at the end of the hour to the energy at the beginning of the hour:

$$e_{t+1}^m = e_t^m + p_t^m \qquad \forall m, \forall t, \tag{9}$$

We use the 'discharge-equivalent' definition of reservoir capacity, and account for the round-trip efficiency through the energy balance constraint where charging the storage with one unit of discharge-equivalent power requires taking a larger quantity of power from the system as Equation 11 shows.

4 Dispatch Optimization

The objective function is the minimization of expected discounted cost over time:

$$\min \mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^t \left(\sum_{n\in[N]} g_t^n c_v^n + l_t^{\mathrm{LL}} c_{\mathrm{LL}}\right)\right], \tag{10}$$

where $\gamma < 1$ is the discount factor.[3]

The energy balance constraint, which must be satisfied in all hours, requires that net load equal the sum of thermal generation, the sum of storage net charge, and the amount of lost load:

$$D' \leq R_t + \sum_{n\in[N]} g_t^n + \sum_{i\in[M]} \left([p_t^m]^- - [p_t^m]^+/\eta_m\right) + l_t^{\mathrm{LL}} \qquad \forall t \tag{11}$$

where $[p_t^m]^+$ is the charging power of the storage($[p_t^m]^-$) is the discharging power of the storage and η_m is the round-trip efficiency.

Since hourly renewable generation is given by the stochastic availability factor, and since the cost minimizing dispatch of thermal generators can be trivially specified using the supply stack, the optimization problem is reduced to choosing the storage actions, $p_t^1, ...p_t^M, \forall t$. An operating policy, $\pi_{\mathcal{P}}$, specifies the storage actions, $(p_t^1, ..., p_t^M)$ conditional on the system state, where the system state is defined by the current state of charge vector, $(e_t^1, ..., e_t^M) \in \mathcal{E} \equiv \mathcal{E}^1 \times ... \times \mathcal{E}^M$ and the current net load, $D_t \in \mathcal{D}$. Therefore the state space is $\mathcal{S} = \mathcal{E} \times \mathcal{D}$, and an operating policy maps the current system state onto a set of actions (power setpoints), $\pi_{\mathcal{P}} : \mathcal{S} \to \mathcal{P}$. We now turn to the methodology for determining the optimal operating policy.

4.0.1 A Markov Decision Process

We apply a Markov Decision Process-based approach, and for ease of exposition we first use the single-storage problem, $M = 1$, and later extend it to multiple storage units. We discretize the storage's state of charge into the set $\mathcal{E} = \{\hat{e}_1, \hat{e}_2, \ldots, \hat{e}_{|\mathcal{E}|}\}$, with equally spaced energy levels. We also define a discrete, equally spaced, and symmetric *feasible action set* $\mathcal{P} = \{-\hat{p}_K \ldots, -\hat{p}_1, 0, \hat{p}_1 \ldots, \hat{p}_K\}$ representing possible hourly storage power setpoints that the storage can have. To ensure a well-behaving model, we use the same spacing for these two sets, limiting the new state of charge after an action ($e' = e + p$) to remain within $\mathcal{E}$.

Additionally, we discretize possible net load states with $\mathcal{D} = \{\hat{d}_1, \hat{d}_2, \ldots, \hat{d}_{|\mathcal{D}|}\}$, equally spaced, and denote space between two net load states as δ.

[3]While the discount factor assures that this infinite sum converges, in fact, for our parameterizations, the short mixing times of the stochastic process, 1 day to 1 week, are the dominant factor. Mixing happens before the discount rate significantly reduces future values. However, for the sake of completeness, the discount rate is kept in the formulation.

We define the overall system "state" using the *state-space*: $\mathcal{S} = \mathcal{E} \times \mathcal{D}$. All pairs $s = (e, d)$, with $e \in \mathcal{E}$ and $d \in \mathcal{D}$ are in the state space. The state space has $|\mathcal{S}| = |\mathcal{E}| \times |\mathcal{D}|$ elements[4].

4.0.2 Evolution of Net Load

Since net load transitions are driven by exogenous stochastic processes, they are independent of storage actions (i.e., e_t does not influence d_{t+1}):

$$\mathbb{P}(d_{t+1} \mid d_t, e_t) = \mathbb{P}(d_{t+1} \mid d_t) \tag{12}$$

Let $\omega(\cdot \mid d_t)$ denote the conditional distribution of the next net load given the current net load d_t. Since net load transitions follow a stochastic process, the next net load state d_{t+1} is not deterministic but instead follows a probability distribution:

$$d_{t+1} \sim \omega(\cdot \mid d_t). \tag{13}$$

The stochastic evolution of net load can be represented by a discrete transition matrix $\mathbf{\Lambda}$, which defines the probability of moving from one net load state to another. The entries of $\mathbf{\Lambda}$ are given by:

$$\mathbf{\Lambda}_{ij} = \mathbb{P}\left(d_{t+1} = \hat{d}_j \mid d_t = \hat{d}_i\right), \tag{14}$$

where:

- $\hat{d}_i$ and $\hat{d}_j$ are discrete net load states.
- $\mathbf{\Lambda}_{ij}$ represents the probability of transitioning from $\hat{d}_i$ to $\hat{d}_j$.

We can construct this transition matrix based on the Ornstein-Uhlenbeck process, using Equation 8:

$$\mathbf{A}_{ij} = \mathbb{P}(\hat{d}_j - \frac{\delta}{2} < D_t < \hat{d}_j + \frac{\delta}{2} \mid D_{t-1} = \hat{d}_i) \tag{15}$$

where δ is the (uniform) distance between discrete net load states.

Instead of modeling a single realized trajectory of net load d_t, we model the time evolution of its probability distribution over the discretized state set$\mathcal{D} = \{\hat{d}_1, \hat{d}_2, \ldots, \hat{d}_{|\mathcal{D}|}\}$. Let $\mathbf{d}_t = [d_t^{(1)}, \ldots, d_t^{(|\mathcal{D}|)}]$ be a row vector, where

$$d_t^{(j)} := \mathbb{P}(d_t = \hat{d}_j), \qquad j = 1, \ldots, |\mathcal{D}|. \tag{16}$$

where $d_t^{(j)} \geq 0$ for all j, and $\sum_{j=1}^{|\mathcal{D}|} d_t^{(j)} = 1$.

Given an initial probability distribution for the net load, $\mathbf{d}_0$, the future probability distribution of net load evolves according to the transition matrix $\mathbf{\Lambda}$:

$$\mathbf{d}_t = \mathbf{d}_0 \mathbf{\Lambda}^t. \tag{17}$$

[4]The notation $|\mathcal{S}|$ represents the cardinality of the set $\mathcal{S}$, i.e., the number of elements in $\mathcal{S}$.

In the long run, the system reaches a stationary distribution $\mathbf{d}_\infty$, which satisfies:

$$\mathbf{d}_\infty = \mathbf{d}_\infty \mathbf{\Lambda}. \tag{18}$$

This equation states that $\mathbf{d}_\infty$ is a left eigenvector of $\mathbf{\Lambda}$ corresponding to eigenvalue 1. Intuitively, this means that as $t \to \infty$, the probability distribution stabilizes, regardless of the initial condition $\mathbf{d}_0$. The existence of a stationary distribution is guaranteed as $\mathbf{\Lambda}$ is derived from an Ornstein-Uhlenbeck process, and the speed of convergence to the stationary distribution is governed by the decay rate of the net load process, explained in Section 3.

4.0.3 Evolution of the Complete System

A global transition matrix ($\mathbf{T} \in \mathbb{R}^{|\mathcal{S}|\times|\mathcal{S}|}$) can be formulated to describe the transition probability from the state $s_t = (d_t, e_t)$ to $s_{t+1} = (d_{t+1}, e_{t+1})$ under a specified policy π and transition function ω. For each state transition:

- e_{t+1} is given by the deterministic operation policy: $e_{t+1} = e_t + \pi_{\mathcal{P}}(d_t, e_t)$
- d_{t+1} is sampled from the distribution defined by the stochastic net load transition model: $d_{t+1} \sim \omega(\cdot \mid d_t)$, as specified by the transition matrix $\mathbf{\Lambda}$.

As with the net load transition matrix, the stationary distribution of the system state is the (left) eigenvector of the global transition matrix associated with eigenvalue 1:

$$\mathbf{s}_\infty = \mathbf{s}_\infty \mathbf{T}, \qquad \sum_{s\in\mathcal{S}} (\mathbf{s}_\infty)_s = 1, \quad (\mathbf{s}_\infty)_s \geq 0.$$

One straightforward way to compute this eigenvector is to iterate

$$\mathbf{s}_{i+1} = \mathbf{s}_i \mathbf{T},$$

starting from any initial distribution $\mathbf{s}_0$. This yields the long-run distribution over net load and state-of-charge under the policy $\pi_{\mathcal{P}}$, quantifying how often the storage is typically full, empty, or operating in intermediate regions.

4.0.4 Solution Algorithm

In any hour, given any net load and any storage action, (d_t, p_t), it is straightforward to calculate the lowest cost mix of thermal generation and to determine the minimum system cost using a supply stack curve. Denote this value by $f_v(d_t, p_t)$.

Since the state space $\mathcal{S} = \mathcal{D} \times \mathcal{E}$ and the action space $\mathcal{P}$ are finite, and the hourly cost function $f_v(d_t, p_t)$ is bounded for all valid state-action pairs, Markov Decision Process theory ensures the existence of at least one policy that minimizes the total expected cost over an infinite horizon.

To find the optimal policy, π^*, and therefore the optimal values of system states, $V^*(e,d)$, we use value iteration, starting with an all-zero value function:

$$V_0(e,d) = 0 \quad \forall (e,d) \in \mathcal{S}$$

We iteratively update the value function for each state $(e,d) \in S$ with the following rule:

$$V_{t+1}(e,d) = \max_p \left[-f_v(d,p) + \gamma \sum_{d' \in \mathcal{D}} \mathbb{P}(d'|d) V_t(e+p,d') \right]$$

After each iteration, the value function $V_t(e,d)$ approximates the optimal expected future cost-to-go from state (e,d) over an increasingly long horizon. Define the per-iteration value increment

$$\Delta_t(e,d) \equiv V_{t+1}(e,d) - V_t(e,d).$$

As t increases, these increments become approximately state-independent, because the net load process approaches its stationary distribution. We declare convergence when the span of the increments is below a prescribed tolerance $\varepsilon > 0$, i.e.,

$$\max_{(e,d)\in\mathcal{S}} \Delta_t(e,d) \; - \; \min_{(e,d)\in\mathcal{S}} \Delta_t(e,d) \; \leq \; \varepsilon.$$

At convergence, we obtain the optimal value function $V^*(e,d)$, from which we derive the optimal storage policy, by choosing, in each state (e,d), the action that attains the Bellman maximum:

$$\pi^*(e,d) \in \arg\max_{p\in\mathcal{P}(e)} \left\{ -f_v(d,p) + \gamma \sum_{d' \in \mathcal{D}} \mathbb{P}(d'|d)\, V^*(e+p,d') \right\},$$

where $\mathcal{P}(e) \subseteq \mathcal{P}$ denotes the feasible actions given state of charge e (i.e., those with $0 \leq e + p \leq e^{\max}$).

4.0.5 Extension of the model for two Storage Assets

In the case of two storage assets, one dimension is added to the system state for the second state-of-charge:

$$\mathcal{S} = \mathcal{E}_1 \times \mathcal{E}_2 \times \mathcal{D} \tag{19}$$

The action space is also expanded to include the simultaneous operation of both storages:

$$\mathcal{P} = \{(p_1, p_2) | p_1 \in \mathcal{P}_1, p_2 \in \mathcal{P}_2\}, \tag{20}$$

where p_1 and p_2 are the charging/discharging actions of storage units 1 and 2, respectively.

We expand the *action policy* function to encompass the optimal actions for both storage units *jointly*, mapping the system states to a combined action pair ($\mathcal{E}_1 \times \mathcal{E}_2 \times \mathcal{D} \to \mathcal{P}_1 \times \mathcal{P}_2$):

$$\pi_{\mathcal{P}} : (e_{1,t}, e_{2,t}, d_t) \to (p_{1,t}, p_{2,t}) \tag{21}$$

The action policy can be decomposed into two *individual action policy* functions, one for each storage unit:

$$\pi_{\mathcal{P}}(e_{1,t}, e_{2,t}, d_t) = \left(\pi_{\mathcal{P}}^{(1)}(e_{1,t}, e_{2,t}, d_t), \pi_{\mathcal{P}}^{(2)}(e_{1,t}, e_{2,t}, d_t)\right) \tag{22}$$

where:

$$p_{1,t} = \pi_{\mathcal{P}}^{(1)}(e_{1,t}, e_{2,t}, d_t), \quad p_{2,t} = \pi_{\mathcal{P}}^{(2)}(e_{1,t}, e_{2,t}, d_t) \tag{23}$$

Thus, each function $\pi_{\mathcal{P}}^{(i)}$ (for $i = 1, 2$) maps the system state to the action of the respective storage unit, assuming the other storage unit will be operated under the system's power policy $\pi_{\mathcal{P}}$. However, it's important to note that individual storage units **do not independently optimize** their own operations. Rather, the system planner optimizes a single, unified action policy that coordinates both storage units together. The individual action policies are just the components of the global policy function.

We break down this global policy into individual storage policy components to make it easier to analyze how each battery's operation is influenced by the other. For example, the optimal charging or discharging strategy for one storage asset will depend on whether the other asset is nearly full or empty. When determining the best action for one storage unit, it is always assumed that the other unit will also follow the system's optimal policy. Therefore under an optimal policy, an action is chosen considering the optimal action will be performed by the other storage too.

The Bellman Equation for the two-storage case is as follows:

$$V^*(e_1, e_2, d) = \max_{p_1, p_2} \left[-f_v(d, p_1, p_2) + \sum_{d' \in \mathcal{D}} \mathbb{P}(d'|d) V^*(e_1 + p_1, e_2 + p_2, d') \right] \tag{24}$$

where the hourly thermal generation cost function, f_v now considers the actions of both storage units.

The value iteration algorithm extends similarly to the two-storage case. We initialize the value function:

$$V_0(e_1, e_2, d) = 0, \quad \forall (e_1, e_2, d) \in \mathcal{S} \tag{25}$$

Then, iteratively update:

$$V_{t+1}(e_1, e_2, d) = \max_{p_1, p_2} \left[-f_v(d, p_1, p_2) + \sum_{d' \in \mathcal{D}} \mathbb{P}(d'|d) V_t(e_1 + p_1, e_2 + p_2, d') \right] \tag{26}$$

This process continues until convergence. The same stopping condition as the one-storage case applies: iteration stops when the increase in state values stabilizes across all states, ensuring that extending the lookahead horizon does not alter the optimal policy.

The optimal policy is then obtained by selecting the optimal storage actions that maximize the cumulative future reward function (action-value):

$$\pi^*_{\mathcal{P}}(e_1, e_2, d) = \arg \max_{(p_1,p_2)\in\mathcal{P}_1\times\mathcal{P}_2} Q^*(e_1, e_2, d, p_1, p_2) \tag{27}$$

where the action-value function Q^* extends as:

$$Q^*(e_1, e_2, d, p_1, p_2) = -f_v(d, p_1, p_2) + \sum_{d'\in\mathcal{D}} \mathbb{P}(d'|d) V^*(e_1 + p_1, e_2 + p_2, d') \tag{28}$$

5 Capacity Optimization

In optimizing capacity investments, we first focus on how the availability of storage changes the capacity investments in thermal generation. We fix the load and installed renewable capacity. We solve for the optimal portfolio of thermal capacity with and without a storage asset.

5.1 Investment Costs

Denote by c^n_f the unit investment cost for capacity of thermal generator n (\$/MW). This cost includes the present value of any fixed operating and maintenance cost over the life of the generator.

The overall problem extends to finding the capacities of generators that would enable the minimization of the total investment and operation costs together:

$$\min_{g^1_{\max},\cdots,g^N_{\max}} \left(\sum_{n\in[N]} c^n_f g^n_{\max} + \min_{p^1_t,\cdots,p^m_T} \mathbb{E}\left[\sum_{t=0}^{T} \gamma^t f_v(d_t, p^1_t, ..., p^M_t) \right] \right) \tag{29}$$

5.2 Optimization

Following Equation 29, the total cost of a generation portfolio $\mathbf{g}_{\max}$ and a policy $\pi_{\mathcal{P}}$ is given by:

$$C_T(\mathbf{g}_{\max}, \pi_{\mathcal{P}}) = \sum_{n\in[N]} c^n_f g^n_{\max} + \mathbb{E}\left[\sum_{t=0}^{T} \gamma^t f_v(d_t, p^1_t, ..., p^M_t) \right] \tag{30}$$

The expected value notation is kept in the equation, because d_t evolves stochastically.

The solution presented in Algorithm 1 solves the problem by iteration, following Geske and Green (2020). First, generator investment decisions are optimized using a screening curve methodology.[5] Given the selected capacities, the optimal battery dispatch strategy is found via dynamic programming. This updates the post-storage net load-duration curve, enabling re-optimization of investment decisions. The process iterates, alternating between investment optimization and strategy optimization, until convergence.

Algorithm 1 Iterative Investment and Expected Operational Cost Minimization, adapted from Geske and Green (2020)

1: $\mathbf{g}_{\max}$ is investment decision variable (to optimize)
2: $\pi_{\mathcal{P}}$ is an operational policy (to optimize)
3: $C_T(\mathbf{g}_{\max}, \pi_{\mathcal{P}})$ is the expected total cost (of investment and operation for a duration T) under policy $\pi_{\mathcal{P}}$, if investment decisions $\mathbf{g}_{\max}$ are made, conditional on the stochastic process of net load.
4: **Initialize:** Set initial investment decision $\mathbf{g}^*_{\max,0}$.
5: $i \leftarrow 1$
6: **repeat**
7: Solve operational problem under uncertainty, with a fixed investment capacity (i.e. find the operating policy that minimizes expected operation cost):

$$\pi^*_{\mathcal{P},i} = \arg\min_{\pi_{\mathcal{P}}} C_T(\mathbf{g}^*_{\max,i-1}, \pi_{\mathcal{P}})$$

8: Solve investment problem using the net load-duration curve with fixed operational policy:

$$\mathbf{g}^*_{\max,i+1} = \arg\min_{g} C_T(g, \pi^*_{\mathcal{P},i})$$

9: Update $i \leftarrow i + 1$
10: **until** Convergence: $\mathbf{g}^*_{\max,i+1} \approx \mathbf{g}^*_{\max,i}$ and $\pi^*_{\mathcal{P},i+1} \approx \pi^*_{\mathcal{P},i}$
11: **Return:** Optimal investment $\mathbf{g}^*_{\max}$ and operational policy $\pi^*_{\mathcal{P}}$

6 Results for Numerical Examples

6.1 Dispatch Optimization of a Single Storage Asset

We use a numerical example to illustrate the result of the model in a simple case. The results in this section are all for a single storage asset and for a thermal stack with only two generators. We use a base case for a comprehensive review of the outputs. We then vary three individual parameters to show how

[5]The screening curve method is a common approach used to determine the least-cost generation mix by comparing the cost structures of different generation technologies given an expected load-duration curve. Refer to the *base case without storage* in Steffen and Weber (2013)'s work for more details.

Table 1: Parameter inputs for base case single storage asset example, and inputs for variation cases.

Parameter		Base Case	Variations 1	2	3
OU process					
Mean	μ_X	0.5			
Std dev, long-term dist	Σ_X	1/12			
Decay time	τ_{95X}	2 days	2 days	2 wks	2 days
Load (MWh)	D'	150			
Marginal Cost ($/MWh)					
Renewable		0			
Baseload	c^1_v	40			
Peaker	c^2_v	80			
VoLL	c_{LL}	18,000			
Gen Capacity (MW)					
Renewable	R	200			
Baseload	g^1_{max}	60			
Peaker	g^2_{max}	40			
Storage Capacity					
Power (MW)	p^1_{max}	8			
Energy (MWh)	e^1_{max}	64	400	400	64
Duration (h)	e^1_{max}/p^1_{max}	8	50	50	8
RT efficiency	η_1	90%	90%	90%	99%

the solution changes. Table 1 contains the full set of input parameters for the base case and shows these variations, too.

In the base case the constant load is 150 MW and the system has 200 MW of renewable capacity. When the renewable availability factor is greater than 75%, there is excess renewable generation and therefore curtailment. This is a three-sigma event, with a frequency of 0.135%. When the factor is below 75%, meeting load requires thermal generation. Given the parameters for the OU process, net load is centered on μ_D=50 MW and a long-term standard deviation of Σ_D=16.7 MW, with a $\pm 6\Sigma_D$ range of [-50,+150]. The thermal supply stack has two generators: 60 MW of a baseload generator with a marginal cost of $40/MWh, and 40 MW of a peaker with a marginal cost of $80/MWh.[6] Whenever net demand exceeds 100 MW, the thermal generation is insufficient to serve all of the load. The value of lost load is $18,000/MWh. The storage unit has 8 MW of power capacity, an 8-hour duration (64 MWh of energy capacity), and a round-trip efficiency of 90%.

Since the net load follows an Ornstein–Uhlenbeck process with a Gaussian stationary distribution, exceeding 100 MW corresponds to a three-sigma event under our parametrization, which occurs approximately 11.8 hours per year.

[6] We chose the generator parameters to roughly match the supply stack in Geske and Green (2020)

Consequently, in approximately 11.8 hours per year net load exceeds the capacity of thermal generation, requiring either power from the storage or load shedding. In another approximately 11.8 hours per year net load would be negative, requiring either charging by storage or renewable curtailment. By adjusting the decay rate parameter of the net load process, it is possible to control whether these 11.8 hours are more or less likely to occur as isolated one-hour events throughout the year or as occasional prolonged sequence of events. A higher decay rate in the renewable availability process results in less frequent but more severe renewable inflow shortages; for instance, the system can be tuned to experience one-hour shortages each month on average (*fast-moving* process) or a single 36-hour shortage occurring approximately once every three years. In our base case, we set the decay time equal to 2 days, meaning that after observing an individual three-sigma extreme value, its effect is expected to decay by 95% within two days. In other words, drought events are generally expected to subside within two days.

6.1.1 Value function

We solved the base case problem using a discrete step size of ΔD=1 MW. It will be convenient to discuss the results by dividing the state space into four regions of net demand: the *curtailment region* as states with non-positive net load ($D_t \leq 0$), *baseload region* where the baseload generator would be marginal absent stored energy ($0 < D_t \leq g^1_{max}$), *peaker region* where the peaker generator would be marginal absent stored energy ($g^1_{max} < D_t \leq g^1_{max} + g^2_{max}$), and *lost load region* where net load exceeds available thermal capacity ($g^1_{max} + g^2_{max} < D_t$).

Figure 4 shows how the value function varies across the state space defined over net demand on the horizontal axis and state-of-charge on the vertical axis. the leftmost panel shows the state space centered on the net demand of 0 and spanning the transition from the curtailment region to the baseload region; the middle panel shows the state space centered on the net demand of 60 and spanning the transition from the baseload region to the peaker region; and, the rightmost panel shows the state space centered on the net demand of 100 and spanning the transition from the peaker region to the lost load region. The heat map within each panel shows the value at each state relative to the lowest valued state within the panel. Within each panel we see that the value is everywhere an increasing function of stored energy and decreasing function of net load.

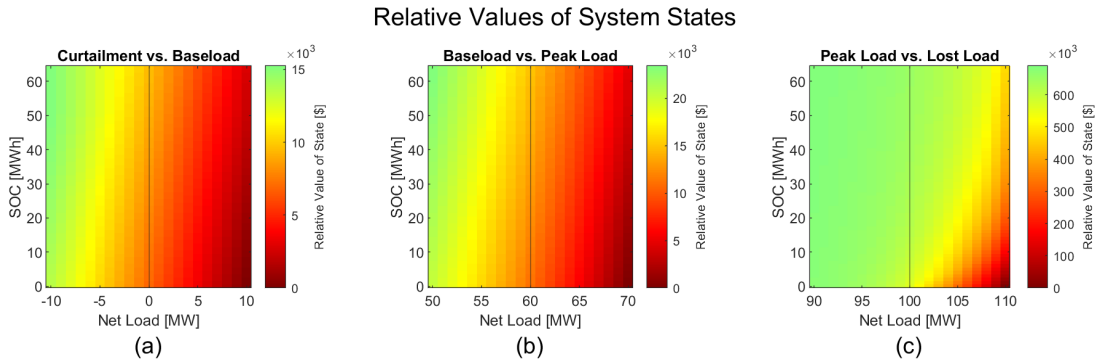

Figure 4: $V(e, d)$: The relative level of the value function is shown around three net load regions centered on the transition boundaries shown as the black vertical lines: 'Curtailment vs Baseload', 'Baseload vs Peaker', and 'Peaker vs Lost Load'. The x-axis shows the net load in MW, while the y-axis shows the state of charge (SoC) of the storage in MWh. The color gradient represents the changing relative level across states, with higher levels shown in green and lower levels in red. These are relative levels because each panel's plot is normalized so that the minimum value in each plot is set to zero. Note that the range of relative levels captured by each panel's color gradient differ across panels.

Looking in the leftmost panel of Figure 4, contrast the \$0 value for the state in the bottom right corner, in which the system has a net load of 10 MW with no stored energy, against the \$15,000 value for the state in the top left corner, in which the system has a net load of -10 MW and 62 MWh of stored energy. Economically, this means that if the optimizer could hypothetically choose its starting point, it would be willing to pay up to \$15,000 to begin operation from the latter state rather than the former. The middle panel of Figure 4 shows a similar pattern, but with a higher gradient.

The rightmost panel of Figure 4 shows the value function around extremely high net load states. The system has 100 MW of generation capacity, so observing a net load of more than 100 MW means that the system is relying on the storage unit to avoid lost load. The green region in the figure represents states with similar, relatively high value, indicating that, when in these states, the storage does not anticipate a significant probability of an imminent loss of load event. The net load is either far enough below the boundary, or the existing state-of-charge is high enough to make lost load an unlikely outcome. Moving towards the yellow and orange regions, the optimizer assigns decreasing values to these states, suggesting an increased likelihood of potential loss of load. If we decrease the state-of-charge below 10 MWh and increase the net load even further, the chances of experiencing a blackout becomes almost certain, so these states' values are very low, and they are shown in red.

6.1.2 Marginal value function

Figure 5 illustrates the marginal value of stored energy, which is defined as the marginal increase/decrease in the *state value* function corresponding to a marginal increase/decrease of the stored energy. Within the Markov Decision Process framework, this is computed by evaluating the change in the state value when the storage contains one additional unit of energy. Consequently, the marginal value of stored energy reflects the optimizer's assessment of the "price" of the stored energy at a specific state, effectively indicating how much the system optimizer would be willing to pay for an additional unit of stored energy in that state. The way in which stored energy is valued provides insights into the storage operation strategy.

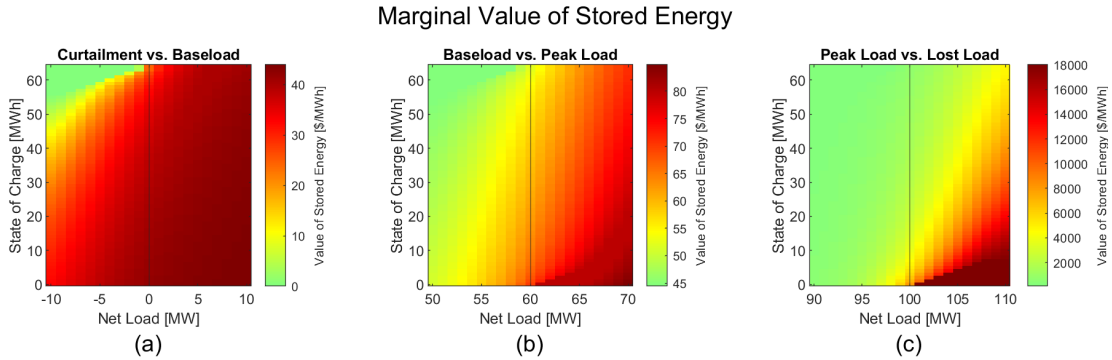

Figure 5: The marginal value of stored energy $(\partial V(e,d)/\partial e)$, around three net load regions centered on the transition boundaries shown as the black vertical lines: 'Curtailment vs Baseload', 'Baseload vs Peaker', and 'Peaker vs Lost Load'. The x-axis shows the net load in MW, while the y-axis shows the state of charge (SoC) of the storage in MWh. The color gradient represents the changing marginal value across states, with higher marginal values shown in red and lower marginal values in green. Note that the range of marginal values captured by each panel's color gradient differ across panels.

In the leftmost panel, which is the region around the boundary between the curtailment and baseload regions, the marginal value of stored energy varies between 0 and 44 \$/MWh. The value of \$44/MWh reflects the fact that the marginal cost of generation in the baseload region is \$40/MWh, and the round-trip efficiency of storage is 90%, so an extra unit of stored energy avoids purchasing 1.1 MW at \$40/MWh. This value arises deep within the baseload region, where the most likely outcome if one unit of stored energy is lost is that the lost energy will be replaced by baseload generation, making the effective value nearly equivalent to the cost of charging the storage with baseload power.

In the middle panel, which is the region around the boundary between the baseload and peaker regions, the marginal value of stored energy has a smooth gradient between the marginal costs of the baseload and peaker generators, adjusted for the round-trip efficiency. It is interesting to examine the marginal

value when net load is either slightly below the transition boundary 60 MW, or slightly above. When slightly below, the marginal value is always higher than \$44/MWh which is the efficiency-adjusted cost of purchasing baseload generation, so the storage should be charging. When slightly above the boundary, the marginal value of stored energy is less than \$80/MWh, so the storage should be discharging to displace generation from the peaker. In this region, the marginal value incorporates the fact that due to the mean-reverting characteristics of the OU process, it is more likely that the process will revert back towards 50 MW than that it will increase further. As a result, the stored energy is valued at less than the full avoided cost of peaker generation, because there is a good chance it can be recharged later at the cheaper baseload price.

In the rightmost panel, the marginal value of stored energy ranges up to the value of lost load. This happens in the bottom right corner, where either the state-of-charge is very low or the net demand is very high, so that the probability of lost load is approaching 1. At any given net load, such as 105 MW, increasing the state-of-charge decreases the marginal value of stored energy. Going from empty to half-full sees the marginal value drop from \$18,000/MWh to \$5,000/MWh, reflecting the decline in the probability of lost load from 100% to 27%, by a ratio that approximately matches 5/18.

We can also see in the rightmost panel that it is possible for the marginal value of stored energy to be zero even when load shedding is 100% certain. To see this, consider a net load of 110 MW. In this case, load shedding is inevitable because the storage's power limit restricts output to a maximum of 8 MW - meaning load shedding will occur regardless of the energy capacity. However, looking at the state where the storage is fully charged, we can see that it still values the stored energy at \$5000/MWh and not at \$18000/MWh. This is because the system perceives that it likely has enough *energy* to manage the scarcity event, even though it lacks sufficient *power* to prevent a blackout. This highlights an important key result: even during a loss-of-load event, it may be optimal to assign a lower value to stored energy than the direct cost of lost load.

6.1.3 Storage dispatch policy

Next, we analyze the optimal dispatch of the storage. Figure 6 shows the dispatch policy across the entire state-space.

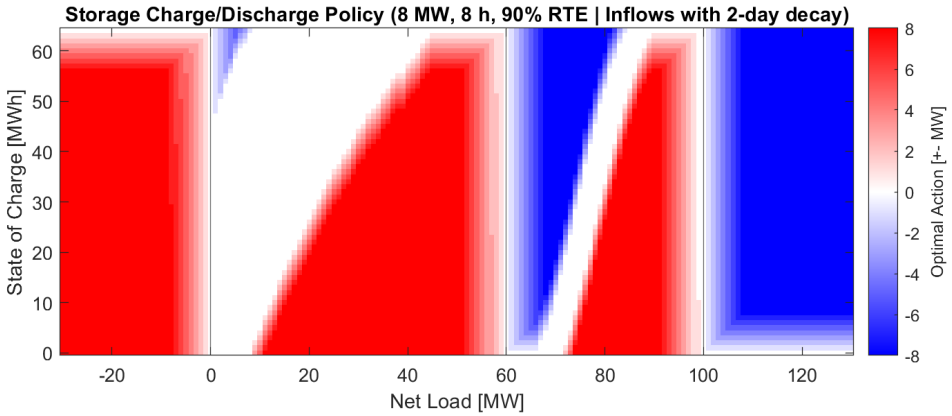

Figure 6: The optimal dispatch policy of a 8 MW, 8h storage with 90% round-trip efficiency in a system with a "fast-moving" renewable inflow process with 2-day decay time. The blue regions show the system states where the preferred action is to discharge, while the red regions show the charging-optimal states.

The red regions in Figure 6 indicate the system states where the storage charges, while the blue regions correspond to states in which the storage discharges. The white region represents states where charge is maintained, i.e., with neither charge nor discharge.

Focusing on the net load region around 100 MW, we see that as net load exceeds the capacity of thermal generation, it is optimal to discharge to prevent load shedding. Initially, when the net load is only slightly greater than thermal capacity, the storage only discharges enough to avoid the lost load.

For net load states where the peak-load generator is active and the system is approaching a loss-of-load event (90–100 MW), the storage charges at full power. This serves as a hedge against the risk of a potential supply scarcity that could push the net load above 100 MW. However, if the net load decreases, the optimal action shifts to discharging, with the storage wanting to reach an empty state before the net load drops to 60 MW. This is because, once the net load falls below 60 MW, the storage can store energy from the cheaper baseload generator. Therefore, it is preferable for the storage to be empty before this transition occurs, ensuring it can take advantage of lower-cost charging opportunities.

Consider the case where the storage is fully charged, and the net load slowly declines from 40 MW down to 0 MW. Until the net load approaches approximately 5 MW, the storage remains idle and does not discharge. However, once the net load reaches 0-5 MW, the storage begins discharging, potentially releasing up to 14 MWh of energy, bringing its energy state down to 48 MWh. This is done in order to capture free energy from a possible curtailment event. Notably, the storage does not fully deplete its stored energy to maximize the chance of capturing free energy later. The rationale for this decision is straightforward: discharging at that moment incurs a cost of approximately \$4.00/MWh, because if curtailment does not occur, the storage will need to recharge from baseload generation, incurring round-trip efficiency losses and

increasing the effective cost by approximately 10%. In contrast, if the storage successfully recharges later with curtailed energy, it captures energy at zero cost and later shifts it to the baseload region, yielding a saving of approximately \$40/MWh. Given these trade-offs, the expected volume of curtailment does not justify discharging more than 14 MWh in this case.

6.1.4 Contrast with perfect foresight dispatch

We present a time series simulation in Figure 7. Specifically, a 50-year sample of the Ornstein-Uhlenbeck process is simulated for the renewable inflows, and the storage is operated on this sample both with perfect foresight optimizer, and with the stochastic dynamic optimizer. The reader is referred to the figure caption for a detailed description of the figure components. The simulation starts on July 23. The perfect foresight optimizer has complete knowledge of the incoming net load. Consequently, at the start of the simulation, it maintains its state of charge and only dispatches energy when the net load exceeds the total generation capacity. Additionally, the optimizer is aware of the magnitude of the impending drought and recognizes that the total energy shortfall will exceed the available storage capacity, making a blackout unavoidable. Given this foresight, an optimal perfect-foresight system planner would ensure that the storage is fully charged before the scarcity period begins. If one marginal unit of energy is lost prior to this period, the optimizer would immediately compensate by recharging, incurring a cost of \$88 /MWh. Therefore, until the scarcity period commences, the stored energy is valued at \$88 /MWh by the perfect foresight optimizer. Once the scarcity period begins at 6 AM on June 24, the optimizer immediately assigns a value of \$18,000/MWh to the stored energy. This is because, with perfect foresight, it recognizes that the available storage will be fully depleted before the scarcity period ends, leading to unavoidable lost load. Consequently, each additional unit of stored energy is valued at \$18,000 /MWh, reflecting the value of lost load.

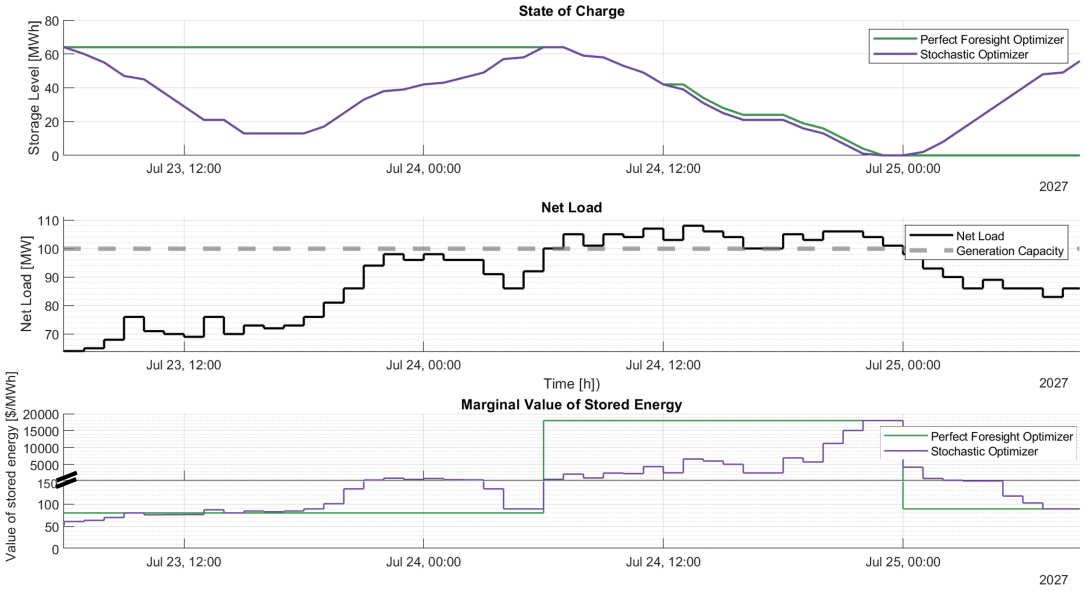

Figure 7: The operation simulation of an 8 MW, 8-hour storage system under both perfect foresight and stochastic optimization. The top panel of the figure displays the state of charge of the storage over time. The middle panel represents the net load trajectory and the total generation capacity with a gray dashed line. The bottom panel shows the valuation of stored energy under the two optimization paradigms. This valuation plot is structured across two Y-axes: one spanning values from 0 to 150, and another capturing a broader range from 150 to 20,000, enabling a comprehensive view of the price fluctuations throughout the simulation period.

Analyzing the stochastic optimizer in Figure 7, at the beginning of the simulated period (around 12:00 on July 23), the stochastic optimizer begins discharging the storage. At this stage, the net load is not yet high enough to pose an imminent risk of scarcity. However, since the system is utilizing the peakload generator, the stochastic optimizer seeks to minimize its use to reduce the operation cost. As the net load starts rising around 6 PM on July 23, the optimizer starts to re-evaluate the value of stored energy. At this point, it perceives the stored energy to be more valuable than the cost of running the peak generator, and starts charging in anticipation of a potential renewable energy shortfall. By 6 AM on July 24, the net load surpasses the total available generation capacity. The perfect foresight optimizer, which has full knowledge of future conditions, recognizes that this situation will inevitably lead to a loss of load due to insufficient stored energy. In contrast, the stochastic optimizer, operating without perfect foresight, does not immediately anticipate this outcome. Consequently, it does not assign the value of lost load to the stored energy until the end of the day, July 24, when it is finally certain that the storage reserves will be depleted before the scarcity ends.

In Figure 7, after the scarcity period ends at midnight on July 25, the net

load drops back below the total generation capacity, though it remains relatively high. The stochastic optimizer, anticipating the possibility of another supply shortage, begins charging the storage as a precautionary measure. In contrast, the perfect foresight optimizer, with perfect knowledge of future inflows, recognizes that no further drought will occur and, therefore, does not initiate charging.

In real-world operations, system operators neither possess perfect foresight nor make decisions dynamically based on a complete stochastic representation of inflows. The appropriate valuation of stored energy likely falls between the two optimization approaches. A more detailed discussion on this topic will be provided in the discussion section.

Valuation of stored energy To gain deeper insights into the valuation of stored energy, we utilize value-duration plots. These plots are analogous to price-duration curves and illustrate the value of stored energy as a function of the fraction of time that value is exceeded. Although the price of electricity does not always correspond to the value of stored energy in a system, we adopt a system planner's perspective, considering these plots as indicative of how storage influences the system's approach to energy valuation. We simulate a 50-year time series sample of the Ornstein-Uhlenbeck process for the renewable inflow, and operate the storage with perfect foresight optimizer, and the stochastic dynamic optimizer to obtain Figure 8. The gray-shaded band in Figure 8 represents the charge-discharge cost of storage. These values form a range due to the asymmetry in the marginal value of stored energy between charging and discharging. Notably, the cost of stored energy is distinct from generation costs. When storage is charging, the dispatchable equivalent energy required for storage necessitates additional energy dispatch from generators to compensate for round-trip efficiency losses. Consequently, there is an inherent efficiency gap between charge and discharge costs. From a system operator's perspective, when the value of stored energy exceeds this cost range, charging is the optimal action, as stored energy is more valuable than immediate consumption. Conversely, when the value of stored energy falls below the shaded area, the storage will discharge.

Figure 8 shows that the perfect foresight optimizer never overvalues or undervalues stored energy (it can price it perfectly). For the perfect foresight optimizer, five distinct price levels emerge corresponding to the two generators' dispatchable costs, these costs adjusted by the round-trip efficiency, and the value of lost load.[7] The perfect foresight optimizer consistently values stored energy at these boundary levels. The optimizer possesses complete knowledge of future energy use: if it receives an additional marginal unit of free energy or loses one, it precisely determines the implications. Specifically,

[7]In presence of n generators, $2n + 1$ price levels arise ($2n + 2$ if the value of lost load is included). This is a well-established result in power market models; see Antweiler and Müsgens (2024) for a formal proof.

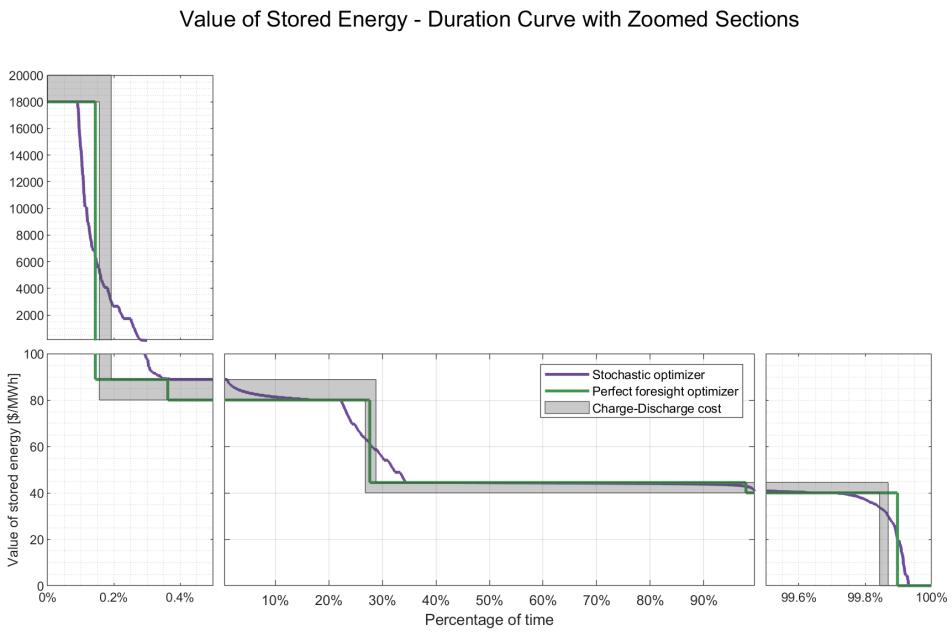

Figure 8: Value-duration curves of stochastic and perfect foresight optimizers. The plots are analogous to price-duration curves and illustrate the value of stored energy as a function of the percentage of total time. The purple line corresponds to the stochastic optimizer while the green line corresponds to the perfect foresight optimizer. The edges of the X axis are zoomed-in, to show the variability in those regions clearer. The gray area shows the cost of charging-discharging the storage. There is an inherent efficiency gap between charge and discharge costs in the system.

it can foresee whether the marginal unit will replace generator's production or, in the case of a loss, the exact additional generation required to compensate for it. In contrast, with the stochastic optimizer, value levels between the discrete price points also emerge. This deviation results from the probabilistic nature of storage utilization. The expected value of stored energy is influenced by the likelihood of different outcomes - gaining or losing a marginal unit - pulling its valuation away from the discrete price levels observed under perfect foresight.

Nevertheless, some constant horizontal segments appear in the stochastic optimizer's value-duration curve in Figure 8. These flat sections correspond to periods when the optimizer is completely certain about the marginal energy's use or replacement. For instance, consider the value of stored energy between 40% and 80% of the time axis. Both the perfect foresight and stochastic optimizers tend to value it at approximately 45 MWh. This suggests that, for the majority of cases, losing a marginal unit of stored energy here would necessitate recharging it from the baseload generator.

The edges of the time percentage axis (specifically the left side of Figure 8), shows substantial valuation differences between the perfect foresight and stochastic models. The system encountered a generation capacity shortage approximately 0.2% of the time. During these periods, the charge cost—the upper bound of the gray area—reached \$20,000/MWh for 0.2% of the time, indicating that the net load surpassed the total available thermal generation capacity and no additional generation capacity remained to charge the storage, so charging one unit would incur a cost of \$20,000/MWh. Additionally, for 0.175% of the time, the stochastic optimizer anticipated an inevitable loss of load due to insufficient stored energy, valuing additional energy at \$18,000/MWh. However, in only 0.1% of cases did the stochastic operator recognize in advance that it would run out of energy. Assuming the perfect foresight optimizer correctly values stored energy, the stochastic optimizer fails to assign the ultimate value of lost load to the stored energy precisely when it should in almost half of the cases (0.175% vs 0.1%).

6.1.5 Comparative Statics

We now illustrate the impact of changing certain parameters of the model: (i) increasing the duration of the storage, i.e., the energy capacity, from 8 hours to 50 hours, (ii) slowing the decay time of the stochastic process (duration of droughts) from 2 days to 2 weeks, and (iii) increasing the round-trip efficiency of the storage from 90% to 99%.

Duration We repeat the simulation keeping all of the base case parameters, except for increasing energy capacity of the storage to 50 hours (400 MWh). The resulting action policy is shown in Figure 9. We observe that, despite the

increased storage capacity, the amount of energy discharged to create space for capturing curtailment events remains 14 MWh. This confirms the assumption that 14 MWh corresponds to the expected volume of curtailment that justifies the efficiency loss.

A similar pattern emerges for the scarcity periods (region around 100 MWh) in Figure 9. The storage does not seek to be fully charged even immediately before a scarcity event (net load at 99 MW). As the expected maximum energy shortfall that justifies charging at peakload does not exceed 120 MWh, it only charges up to that point. To understand this behavior, we can perform a simple back-of-the-envelope calculation. Suppose the system is operating at a net load of 99 MW. In this state, the storage aims to maintain at least 120 MWh of stored energy. The cost of unserved load is \$18,000/MWh, while the cost of charging from the peak-load generator is approximately \$88/MWh. Given these values, the probability threshold at which charging is economically justified can be estimated as $\frac{\$88}{\$18{,}000}$, which corresponds to a probability of approximately 0.48%. The storage expects the likelihood of a scarcity event exceeding 120 MWh to be below 0.48%; if the probability were higher, it would charge more to minimize the expected future costs. Therefore, figure 9 suggests that the stored energy reservoir is oversized for this system. Even in the most extreme scarcity events, it does not utilize its full capacity, indicating that a smaller storage unit could achieve similar risk-mitigation performance.

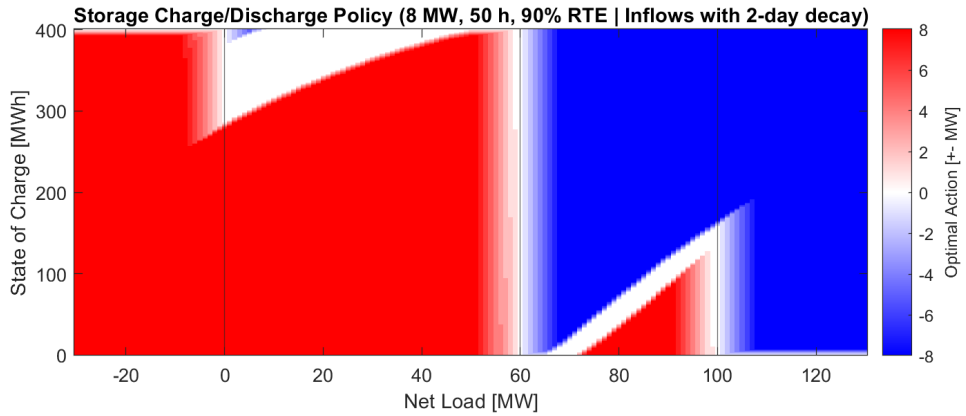

Figure 9: The optimal dispatch policy of a 8 MW, 50h storage with 90% round-trip efficiency, and a "fast-moving" renewable inflow process with 2-day decay time. The blue regions show the system states where the preferred action is to discharge, while the red regions show the charging-optimal states.

Decay time To analyze the impact of decay time, we maintain the 50-h storage while altering the net load process to exhibit slower variations in renewable energy inflows, extending the process decay time from 2 days to 14 days. This adjustment reduces the frequency of extreme events, however, when they do occur, they persist for longer durations and involve larger energy volumes.

The effects of this shift on operational decisions are illustrated in Figure 10. The first key observation is that the storage system is now willing to discharge up to 100 MWh to capture curtailment. This is because, under slow-moving renewable inflows, curtailment events tend to involve larger volumes, incentivizing the storage system to remain emptier to accommodate surplus energy efficiently.[8] Another key observation in Figure 10 is the substantial increase in the expected volume of single renewable energy droughts compared to the previous case. This is evident from the storage system's behavior - prioritizing a fully charged state ahead of scarcity periods. In contrast, under fast-moving renewable inflows, as seen in Figure 9, the storage maintained a more moderate charge level at around 95 MW of net load. This shift highlights the increased precautionary need for stored energy when renewable inflows exhibit prolonged low-output periods.

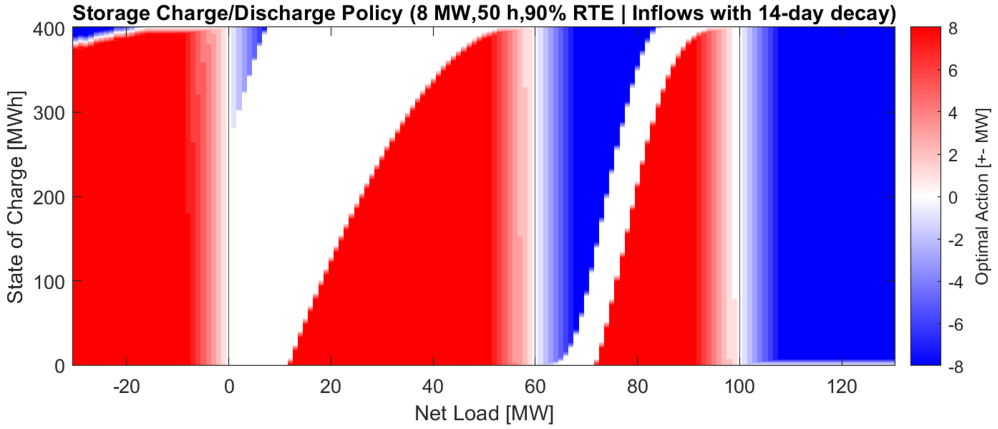

Figure 10: The optimal action policy of a 8 MW, 50h storage with 90% round-trip efficiency. The optimal actions are determined under a "slow-moving" renewable inflow process with 14-days decay time. The blue regions show the system states where the preferred action is to discharge, while the red regions show the charging-optimal states.

Round-trip efficiency Another key factor influencing the decision-making of the storage is the round-trip efficiency of the storage system. To assess its impact, we revisit the system with an 8 MW, 8-hour storage configuration, like the one producing Figure 6, exept now we modify the round-trip efficiency from 90% to an aspirational 99%. The resulting optimal dispatch policy is shown in Figure 11.

[8] At extremely low net loads, around -25 MW, when the storage is fully charged, the figure shows the policy involves discharging, which might seem counterintuitive. However, this behavior is a direct consequence of the initial parametrization of the optimization program that is implemented. In reality, the specific action taken in these states is inconsequential because the slow-moving nature of the process ensures that the storage will have enough time to fully recharge before the net load turns positive. A different initialization would remove this.

The first notable effect is that the white regions in the figure - representing states where the storage maintains its state of charge at a constant level - become even narrower.[9] This is because with higher efficiency, the storage can switch between charging and discharging more without significant energy "regret" associated with the round-trip energy losses.

To understand this in more detail, consider the region corresponding to peakload generation (60 to 100 MW). In Figure 6, the undecided region was noticeably wider (corresponding to a valuation between 80 and 88 $/MWh), implying that the storage hesitated between charging and discharging in many states. However, with higher efficiency, in Figure 11, the financial penalty associated with cycling energy through the storage is minimal. Since the system can recover nearly all of the discharged energy at almost the same price, it follows a more immediate and responsive strategy - charging when the load increases to avoid blackouts and discharging when the load decreases to make space for capturing the baseload generation (the undecided value band shrinks to between 80 and 81 $/MW). As a result, the model exhibits a much narrower undecided region. This shows the direct impact of the round-trip efficiency on storage operation flexibility.

A second key observation in Figure 11 is the storage system's increased willingness to discharge its energy in anticipation of potential curtailments. In Figures 6 and 9, maintaining 14 MWh of available capacity was considered sufficient to capture curtailment events. However, with higher efficiency, in Figure 11, the opportunity cost of prematurely releasing stored energy is significantly reduced, because the system can recover back the energy almost without any losses, if needed. As a result, the storage system can preemptively discharge to near-empty levels, ensuring maximum capacity for potential curtailment events. If the anticipated curtailment does not occur or is smaller than expected, the storage can be charged back from almost the same cost.

[9] At 100% efficiency, these white regions would reduce to a single line.

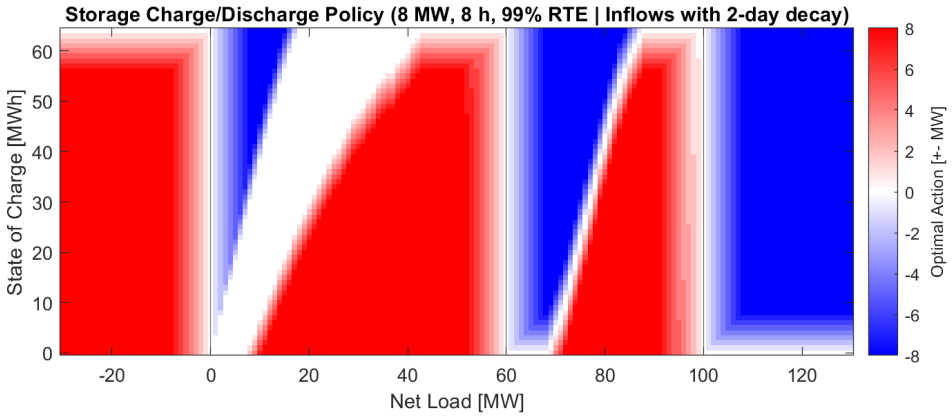

Figure 11: The optimal action policy of a 8 MW, 8h storage with 99% round-trip efficiency. The optimal actions are determined under a "fast-moving" renewable inflow process with 2-day decay time. The blue regions show the system states where the preferred action is to discharge, while the red regions show the charging-optimal states.

6.2 Dispatch Optimization of Two Storage Assets

6.2.1 Inputs and Outputs

The input parameters for the numerical examples of the two storage optimization are shown in Table 2. In all cases, the constant load is 75 MW and the system has 100 MW of renewable capacity. When the renewable availability factor is greater than 75%, there is excess renewable generation. When the factor is below 75%, meeting load requires thermal generation. Given the parameters for the OU process, net load is centered on μ_D=25 MW and a long-term standard deviation of Σ_D=8.33 MW, with a $\pm 6\Sigma_D$ range of [-25,+125]. The thermal supply stack has two generators: 30 MW of a baseload generator with a marginal cost of \$40/MWh, and 20 MW of a peaker with a marginal cost of \$80/MWh. Whenever net demand exceeds 50 MW, the thermal generation is insufficient to serve all of the load. The value of lost load is \$18,000/MWh. We set the decay time equal to 4 days.

We maintain those system input parameters and then consider various combinations of two storage assets with different capabilities. For the base case, storage asset #1 has 4 MW of power capacity, a 25-hour duration (100 MWh of energy capacity), and a round-trip efficiency of 90%. Storage asset #2 also has 4 MW of power capacity, but only a 2-hour duration (8 MWh of energy capacity), and a higher round-trip efficiency of 95%. We then examine how the interaction changes for three alternative parameterizations.

Figure 12 illustrates the mean operating state of charge for the two base case storage assets, when the short duration asset is more efficient. The immediate observation is that the longer duration storage is less responsive at low levels of net load, holding its state of charge for the higher value arbitrage at medium and high levels of net load. In contrast, the shorter duration storage is

Table 2: Parameter inputs for base case two storage assets example, and inputs for variation cases.

Parameter		Base Case	Variations 1	2	3
OU process					
Mean	μ_X	1/2			
Std dev, long-term dist	Σ_X	1/12			
Decay time	τ_{95X}	4 days			
Load (MWh)	D'	75			
Marginal Cost ($/MWh)					
Renewable		0			
Baseload	c^1_v	40			
Peaker	c^2_v	80			
VoLL	c_{LL}	18,000			
Gen Capacity (MW)					
Renewable	R	100			
Baseload	g^1_{max}	30			
Peaker	g^2_{max}	20			
Storage #1 Capacity					
Power (MW)	p^1_{max}	4	4	8	4
Energy (MWh)	e^1_{max}	100	100	400	16
Duration (h)	e^1_{max}/p^1_{max}	25	25	50	4
RT efficiency	η_1	90%	95%	90%	100%
Storage #2 Capacity					
Power (MW)	p^2_{max}	4	4	8	4
Energy (MWh)	e^2_{max}	8	8	16	16
Duration (h)	e^2_{max}/p^2_{max}	2	2	2	4
RT efficiency	η_2	95%	90%	95%	100%

very active immediately around the net load boundaries. For example, during curtailment events and baseload generation periods, the shorter duration storage was utilized to capture a portion of the curtailed energy, while the longer duration storage remained inactive. However, when the system approached scarcity conditions - around a net load of 50 MW - the shorter duration storage was charged to full capacity as a hedge against an imminent energy shortage, whereas the longer duration storage maintained a state of charge of approximately 40%. This suggests that the longer duration storage deemed its current charge level sufficient to bridge potential scarcity gaps.

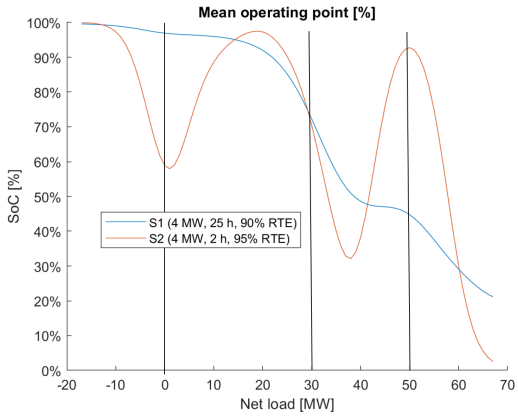

Figure 12: Mean SoC levels of two storages across different net load states, showing each storage's preferred SoC at each net load level. The blue line shows a 4 MW, 25 h storage with 90% round-trip efficiency, while the orange line shows a 4 MW, 2 h storage with 95% round-trip efficiency. The x-axis shows the net load of the system, while y-axis shows the state of charge as a percentage of the reservoir. Base case parameters from Table 2 are used.

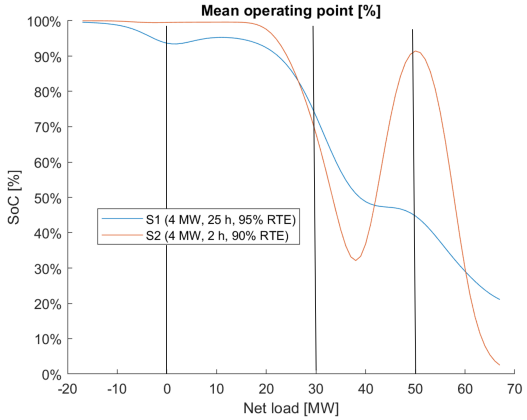

Figure 13: Mean SoC levels of two storages across different net load states, showing each storage's preferred SoC at each net load level. The blue line shows a 4 MW, 25 h storage with 95% round-trip efficiency, while the orange line shows a 4 MW, 2 h storage with 90% round-trip efficiency. The x-axis shows the net load of the system, while y-axis shows the state of charge as a percentage of the reservoir. Case 1 parameters from Table 2 are used.

The operational differences between the two storage types arise from two primary factors: (i) differences in storage duration and (ii) differences in round-trip efficiency. To further disentangle these effects, a secondary simulation was conducted where the storage durations were held constant while swapping their round-trip efficiencies. The results are presented in Figure 13. Examining the system dynamics near curtailment in Figure 13, we observe that the role of capturing curtailed energy shifts, demonstrating that higher-efficiency storage technology is preferentially utilized for curtailment mitigation. This suggests that round-trip efficiency, rather than storage duration, is the dominant factor determining whether it is beneficial to reserve storage ca-

pacity for curtailment capture. However, in contrast to curtailment mitigation, the storage operation under scarcity conditions (net load nearing 50 MW) remains largely unchanged. This indicates that round-trip efficiency does not substantially alter the dynamics of risk mitigation. Instead, storage duration appeared to be the primary determinant of how storage assets are utilized during scarcity periods, whereas round-trip efficiency mainly influenced their role in curtailment avoidance.

Figure 14 presents a time series excerpt from a two-storage optimal operation simulation, performed on a sampled Ornstein-Uhlenbeck renewable inflow process. The figure shows that initially, at midnight on September 6, the net load is met by baseload generation, and both storage units are fully charged. During this time, the stored energy in the long-term storage holds a higher value than that in the short-term storage, primarily due to the lower round-trip efficiency of the former. At 4 AM on September 7, the net load transitions into the peak generation range, increasing generation costs to \$80/MWh. As a result, the value of stored energy in both storage units begins to rise. However, since this value remains below the marginal generation cost of \$80/MWh, both storages commence discharging. By 4 PM on September 7, as the net load approaches the system's peak generation capacity, the stored energy value in the short-term storage exceeds \$84/MWh (the cost of charging the small storage using peak generation[10]). For that reason, the short-term storage starts charging, as it values the stored energy higher than the charging costs from the generator. Meanwhile, the long-term storage still attributes a lower value to its stored energy and continues discharging. Around 6 AM on September 8, the net load decreases and begins fluctuating between the baseload and the peakload generator. This continues until noon, September 9. During this time, the system's generation cost oscillates between \$40 and \$80 /MWh, as the peaker is turned on and off. Throughout this period, both storage units value energy between these price levels, leading them to charge whenever the net load decreases under the baseload generation capacity and discharge whenever it moves back toward peak generation. Consequently, both storages engage in arbitrage.

[10] As the round-trip efficiency of the small storage is 95%, it requires approximately \$84/MWh for charging

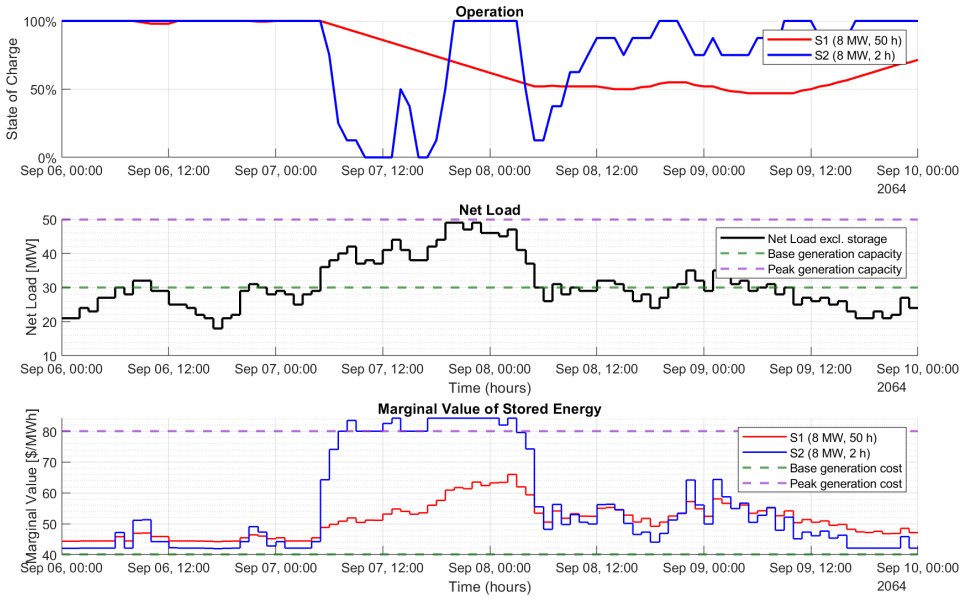

Figure 14: The operation simulation of a two-storage grid. A medium-term, 8 MW, 50-hour storage and a short-term, 8 MW, 2-hour storage is used. The top panel of the figure displays the state of charge of the two storages over time. The middle panel represents the net load trajectory and the base and peak generation capacity with a horizontal dashed lines. The bottom panel shows the valuation of stored energy of two storages, along with the horizontal dashed lines showing the marginal costs of the generators. Both storage units start the simulation fully charged. Case 2 parameters from Table 2 are used.

The difference in the valuation of stored energy between the two storage units is not solely attributable to their round-trip efficiencies. The primary driver of this divergence is the disparity in energy capacities. While the long-term storage can accommodate a larger amount of energy and is therefore less sensitive to rapid fluctuations in generation costs, the short-term storage lacks this buffering capacity. A sudden spike in generation costs can deplete its reserves within two hours, causing the valuation of its stored energy to be adjusted more frequently in response to load conditions.

This dynamic raises a critical question: To what extent does the valuation of one storage unit depend on the state of charge of the other? How does the state of charge of one storage unit influences the other? To explore these interactions, we simulate the same period in Figure 14, this time starting with two empty storage units, as shown in Figure 15.

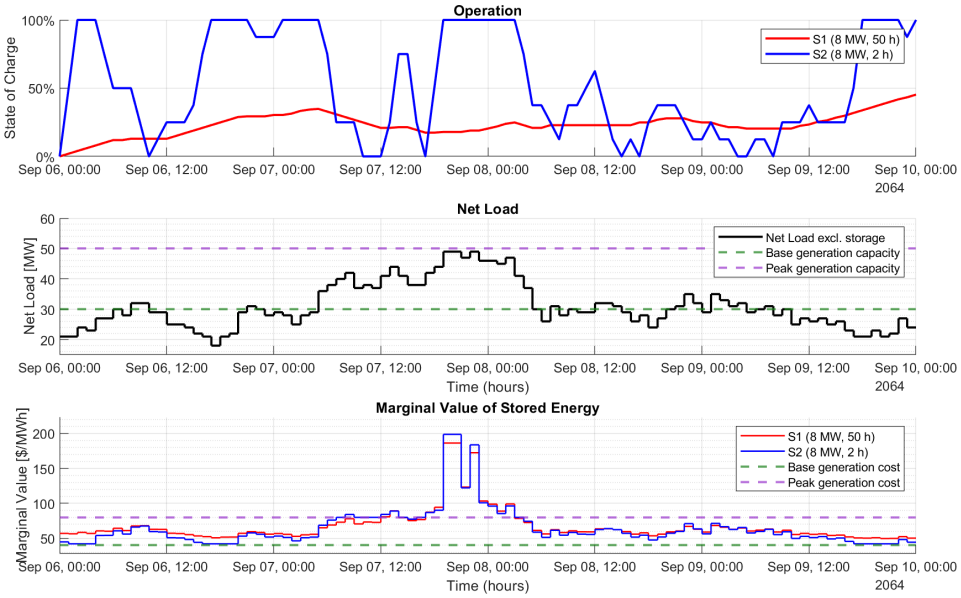

Figure 15: The operation simulation of a two-storage grid. A medium-term, 8 MW, 50-hour storage and a short-term, 8 MW, 2-hour storage is used. The top panel of the figure displays the state of charge of the two storages over time. The middle panel represents the net load trajectory and the base and peak generation capacity with a horizontal dashed lines. The bottom panel shows the valuation of stored energy of two storages, along with the horizontal dashed lines showing the marginal costs of the generators. Both storage units start the simulation with an empty reservoir. Case 2 parameters from Table 2 are used.

The first key observation in Figure 15 is that the operational behavior of the short-term storage from 4 AM on September 7 until the end of the day remains largely unchanged from Figure 14, despite significant fluctuations in the value of its stored energy. Furthermore, the longer-duration storage now holding a minimal amount of energy, the value of its energy rises significantly, together with the short-term storage.

Another notable effect in Figure 15 is the disappearance of the divergent operation observed in the previous case in Figure 14 - where one storage charged while the other discharged. This shift occurs because the overall system now contains less stored energy. With lower reserves, the long-term storage also assigns a higher value to its energy, preventing further discharging even when the system is in the peakload region.

As a result, when the total energy stored in the system is insufficient to bridge scarcity periods, both storage units value their own energy reserves higher. Consequently, the reduced energy availability in the long-term storage does not only impact its own valuation but also elevates the perceived value of energy in the short-term storage.

6.2.2 Energy sharing

Assume a system with two identical storage units, as in case 3 in Table 2. Under optimal operation, both storage units would follow the same control strategy, maintaining identical state-of-charge profiles throughout operation. Consequently, their behavior is equivalent to that of a single, aggregated storage unit with the combined energy and power capacity of both.

The Markov decision process framework developed in this thesis allows to analyze scenarios in which the two storage units deviate from this synchronized, or "lockstep," operation mode[11].

A notable behavior observed in our framework during two-storage simulations is *energy sharing*: the optimal policy tends to equalize the state of charge across storage assets, effectively reallocating energy between them as a hedge against extreme events.

To isolate the effects of energy sharing from other influencing factors, we simulate a system with two identical storage units, without round-trip energy losses. Each unit has 4 MW power capacity, 4-hour energy capacity, and 100% round-trip efficiency (case 3 in Table 2).

Figures 16 and 17 illustrate the optimal action policy of one storage unit under two distinct scenarios: when the other storage is empty and when it is

[11]The "lockstep" operation is optimal because, given identical power and energy constraints, this strategy prevents premature saturation of one storage unit while the other still has capacity, maximizing the system's total available energy and power headroom at all times. Since divergence offers no advantage but introduces the risk of constraint saturation, keeping both SoC's equal is always at least as good as any alternative strategy, and is typically the best choice. Therefore, optimal operation has to be in "lockstep".

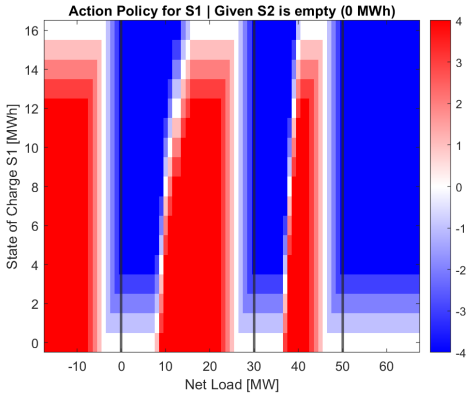

Figure 16: The optimal action policy of a 4 MW, 4h storage with 100% round-trip efficiency, when another identical, but empty storage is present in the system. The blue regions show the system states where the preferred action is to discharge, while the red regions show the charging-optimal states.

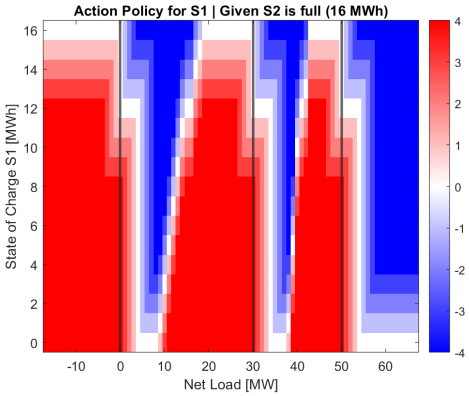

Figure 17: The optimal action policy of a 4 MW, 4h storage with 100% round-trip efficiency, when another identical, and fully charged storage is present in the system. The blue regions show the system states where the preferred action is to discharge, while the red regions show the charging-optimal states.

fully charged.

Figure 16 depicts the optimal action of one storage unit (S1) when the other storage (S2) is empty. Focusing on net load levels around 50 MW - the total generation capacity of the system - we observe that if S1 has stored energy (specifically, if its state of charge exceeds approximately 4-5 MWh), it begins discharging before the net load reaches 50 MW. This behavior emerges because S1 strategically wants to discharge energy to facilitate the charging of S2. For instance, if the net load is 48 MW and S2 is empty, it seeks to charge in preparation for a potential load increase to hedge against a blackout. However, S2 can charge only up to 2 MW; exceeding this would push the system into a loss-of-load state. In this scenario, if S1 starts discharging, it effectively enables S2 to charge at a faster rate. Consequently, when a scarcity period is anticipated and S2 is empty, S1 preemptively discharges energy to help S2 charge faster. Figure 17 shifts the perspective to the case where S2 is full. Now around 50 MW, we see that even if the net load is marginally higher than 50 MW, which exceeds the generator capacity of the system - if S1 has an energy level lower than 11 MWh - it should charge. At a net load of 51 MW, despite the net load being higher than the peak generator capacity, S1 is not helping the system by discharging, instead, it is charging. This happens because under the system-optimal policy, S1 knows that energy to charge S1 would come from S2, which is full, and there would be no loss of load. For example, imagine S1 has 7 MWh of stored energy, and the system is at net load of 51 MW. Knowing that S2 is full, Figure 17 suggests that S1 should charge 3 MW under the optimal policy. This might be counter-intuitive at the

beginning, however, knowing that S2 will discharge 4 MW, this operation will: (i) avoid losing load, as the combined action of the storages is -1 MW (ii) put two storages state of charge closer to each other (S1 will have 10 MWh and S2 will have 12 MWh). Two storages try to distribute their state of charge evenly, so they can go back to the symmetric operating point, where two storage assets have the same state of charge, and operate together in "lockstep".

This behavior reflects the system's natural tendency to balance state of charge between storage units of similar (or identical) storage duration and round-trip efficiency. By aligning their energy levels, the storages return to a symmetric operating point, where they can function in "lockstep".

6.3 Capacity Optimization of Thermal Generation

We present results for a case when there is no storage, a case when there is a short duration storage asset, and a case when there is a long duration storage asset. For all cases, we solve for the optimal portfolio of thermal capacity for two alternative decay times of the stochastic renewable process.

Table 3 shows the parameters that are maintained across all scenarios.

Table 4 shows the parameters of the four scenarios - two duration storage assets and two decay times.[12]

Table 5, adapted from Geske and Green, 2020 shows the parameters for the unit investment costs as well as the variable costs.

Table 6 shows the optimization results for the case of no storage and across the four scenarios for storage and decay time. For each scenario we show results for the stochastic optimization of storage dispatch and the result under the assumption of perfect foresight. Here is a guide to reading the table of results:

- Columns;
  - column [A] shows the results for no storage;
  - columns [B]-[E] show the results when the storage asset is short duration, 4 hours;
    - columns [B]-[C] show the results for a fast moving process; and
    - columns [D]-[E] show the results for a slow moving process;
    - each pair of columns shows the results for stochastic optimization of the storage asset's dispatch as contrasted with the results for perfect foresight optimization;
  - Columns [F]-[I] show the sets of results when the storage asset is long duration, 50 hours.
- Rows;

[12] When there is no storage, the decay time does not change the optimal portfolio of capacity.

Table 3: Maintained Parameter Values

Parameter		Value
OU process		
Mean	μ_X	0.5
Std dev, long-term dist	Σ_X	1/12
Load (MWh)	D'	150
Renewable Capacity (MW)	R	200
Storage Capacity		
Power (MW)	p^1_{max}	8
RT efficiency	η_l	90%

Table 4: Four Scenarios for Storage Duration and Decay Time

	Fast 2-days	Slow 2-weeks
Short duration 4 hour, 32 MWh	1	2
Long duration 50 hour, 400 MWh	3	4

Table 5: Thermal Generation Cost Inputs

Technology	Fixed Cost $/kW-year	Variable Cost $/MWh
Nuclear	387	14
Lignite	189	38
Coal	163	42
Combined Cycle	103	59
Combustion Turbine	55	94
Value of Lost Load		18,000

Table 6: Results for Optimization of Thermal Generation Across Alternative Scenarios

			Short duration, 4 Hr.				Long duration, 50 Hr.			
			Fast, 2-days		Slow, 14-days		Fast, 2-days		Slow, 14-days	
		No Storage [A]	Stoch. Optim. [B]	Perfect Foresight [C]	Stoch. Optim. [D]	Perfect Foresight [E]	Stoch. Optim. [F]	Perfect Foresight [G]	Stoch. Optim. [H]	Perfect Foresight [I]
	System Cost ($M/year)									
[1]	Total	31.5	31.2	30.8	31.4	31.2	30.6	30.1	31.0	30.6
[2]	Fixed	18.1	17.6	18.2	17.9	17.8	17.3	19.0	17.5	18.1
[3]	Variable	13.4	13.5	12.6	13.5	13.4	13.3	11.1	13.5	12.5
	Storage Savings									
[4]	($M/year)		0.3	0.7	0.1	0.3	0.9	1.4	0.5	0.9
[5]	($/kWh/year)		9.38	21.88	3.13	9.38	2.25	3.50	1.25	2.25
[6]	(as % of No Storage)		1.0%	2.2%	0.3%	1.0%	2.9%	4.4%	1.6%	2.9%
[7]	Perf. Fore. Overest			133%		200%		56%		80%
	Capacity (MW)									
[8]	Nuclear	24.00	23.67	26.89	23.56	23.00	24.00	32.89	23.56	26.27
[9]	Lignite	15.00	15.33	13.11	15.44	16.00	18.00	10.11	15.44	13.73
[10]	Coal	15.00	15.00	13.93	14.78	15.00	10.00	9.89	15.00	13.72
[11]	Combined Cycle	13.00	9.78	11.07	12.22	12.48	7.00	7.48	7.89	10.89
[12]	Combustion Turbine	40.00	38.22	34.70	39.89	39.61	41.00	36.63	40.11	35.39
[13]	Total	107.00	102.00	100.00	106.00	106.00	100.00	97.00	102.00	100.00
	Lost Load									
[14]	# Hours	1.7	1.4	1.7	5.8	5.8	1.4	1.4	1.4	1.3
[15]	MWh	7.0	6.7	7.5	9.6	9.7	5.4	5.4	4.6	4.5
	Curtailment									
[16]	# Hours	12.2	9.9	4.5	8.6	7.1	7.2	3.1	8.6	2.2
[17]	MWh	76.7	67.6	31.2	50.6	42.1	47.8	17.2	50.6	8.5

- row [1] shows the total system cost (per annum), which is the metric that is optimized; rows [2]-[3] break that down into fixed and variable cost;
- rows [4]-[6] show the total system cost savings from adding the storage asset, reported in total dollars, in dollars per unit of energy capacity of the storage asset, and as a percent of total system cost when there is no storage;
- row [7] shows the system cost savings under stochastic optimization as a percent of the system cost savings under perfect foresight;
- rows [8]-[13] show the optimized portfolios of generation capacity;
- rows [14]-[15] and [16]-[17] show metrics for hours of lost load and for curtailment, respectively.

In all cases, including the no storage case, total thermal capacity is below the maximum net load, so that it is cost efficient to allow some hours of lost load. Also, in all cases there are many hours of curtailment. Generally, storage reduces both curtailment and lost load, although given a short duration storage asset with a slow moving process, the chosen portfolio of thermal capacity allows an absolute increase in lost load hours. This occurs because a slow-moving renewable process generates longer scarcity episodes that short-duration storage cannot effectively bridge; eliminating those lost-load hours would require substantial additional thermal investment, so the cost-minimizing solution tolerates more lost-load hours while keeping unserved energy (MWh), and thus the cost of lost load, relatively contained.

We can contrast the savings from long versus short duration storage by comparing the values in columns [B] and [F] or columns [D] and [H]. The long duration storage has 12.5 times as much energy capacity and yields 3 times the system cost savings - row [4] - as the short duration storage, for the fast moving process, or 5 times the system cost savings, for the slow moving process. Accordingly, the average unit cost savings - row [5] - for the long duration storage asset are one-fifth the average unit cost savings the short duration asset, for the fast moving process, or two-fifths, for the slow moving process.

We can contrast the savings under a fast moving process versus a slow moving process by comparing the values in columns [B] and [D] or columns [F] and [H]. When the process is slow moving, the short duration storage asset yields one-third the system cost savings as when the process is fast moving. When the process is slow moving, the long duration storage asset yields a little more than one-half the system cost savings as when the process is fast moving. Looking down at rows [14]-[17], we can see that with the slow moving process and the short duration storage asset, there is a significant reduction in curtailment, as compared against the no storage case, but an absolute increase in lost load.

We can contrast the stochastic optimization results against the perfect foresight optimization results by comparing column [B] against [C], [D] against

[E], [F] against [G], and [H] against [I]. Looking at rows [14]-[17], we can see that perfect foresight optimization achieves a much larger reduction in curtailment than the stochastic optimization, whereas the differences in lost load are much smaller and not consistent. Looking at row [13], we can see that perfect foresight enables a reduction in total capacity relative to stochastic optimization. Looking at row [2] we can see that sometimes this reduction in capacity nevertheless involves a greater investment in the costly capital intensive capacity, while looking at row [3], we can see that perfect foresight always enables lower variable costs. Looking at row [7], we see that the cost-saving potential of a storage asset calculated assuming perfect foresight is overestimated by 133%, 200%, 56% or 80% depending upon the scenario. These are all much greater overestimates than the 27% estimated by Geske and Green (2020). This is likely due to the diurnal structure that dominates the Geske and Green model, whereas our model has no predictable diurnal component.

References

Antweiler, W., & Müsgens, F. (2024). The New Merit Order: The Viability of Energy-Only Electricity Markets With Only Intermittent Renewable Energy Sources and Grid-Scale Storage. *Energy Economics*, *145*. https://doi.org/10.2139/ssrn.4702939

Bhatnagar, D., Boff, D. S., Twitchell, J. B., Bedoya Ceballos, J. C., & Somani, A. (2022). *Compensation for long-duration energy storage* (tech. rep.). Pacific Northwest National Lab.

Bracken, C., Voisin, N., Burleyson, C. D., Campbell, A. M., Hou, Z. J., & Broman, D. (2024). Standardized benchmark of historical compound wind and solar energy droughts across the Continental United States. *Renewable Energy*, *220*. https://doi.org/10.1016/j.renene.2023.119550

Chen, P., Pedersen, T., Bak-Jensen, B., & Chen, Z. (2010). ARIMA-based time series model of stochastic wind power generation. *IEEE Transactions on Power Systems*, *25*(2). https://doi.org/10.1109/TPWRS.2009.2033277

Evans, L., Guthrie, G., & Lu, A. (2013). The role of storage in a competitive electricity market and the effects of climate change. *Energy Economics*, *36*. https://doi.org/10.1016/j.eneco.2012.09.016

Geske, J., & Green, R. (2020). Optimal Storage, Investment and Management under Uncertainty: It is Costly to Avoid Outages! *Energy Journal*, *41*(2).

Glasbey, C. A. (2001). Non-linear autoregressive time series with multivariate Gaussian mixtures as marginal distributions. *Journal of the Royal Statistical Society. Series C: Applied Statistics*, *50*(2). https://doi.org/10.1111/1467-9876.00225

Gorenstin, B. G., Campodonico, N. M., Costa, J. P., & Pereira, M. V. (1993). Power System Expansion Planning Under Uncertainty. *IEEE Transactions on Power Systems*, *8*(1). https://doi.org/10.1109/59.221258

Grillo, S., Pievatolo, A., & Tironi, E. (2016). Optimal Storage Scheduling Using Markov Decision Processes. *IEEE Transactions on Sustainable Energy*, *7*(2). https://doi.org/10.1109/TSTE.2015.2497718

Junge, C., Mallapragada, D., & Schmalensee, R. (2022). Energy storage investment and operation in efficient electric power systems. *Energy Journal*, *43*(6), 1–24.

Korpås, M., & Botterud, A. (2020). *Optimality conditions and cost recovery in electricity markets with variable renewable energy and energy storage*, MIT Center for Energy and Environmental Policy Research.

McConnell, D., Forcey, T., & Sandiford, M. (2015). Estimating the value of electricity storage in an energy-only wholesale market. *Applied Energy*, *159*. https://doi.org/10.1016/j.apenergy.2015.09.006

Moret, S. (2017). *Strategic energy planning under uncertainty.* [Masters]. École Polytechnique Fédérale de Lausanne (EPFL).

Pfenninger, S., & Staffell, I. (2016). Long-term patterns of European PV output using 30 years of validated hourly reanalysis and satellite data. *Energy*, *114*. https://doi.org/10.1016/j.energy.2016.08.060

Ruhnau, O., & Qvist, S. (2022). Storage requirements in a 100% renewable electricity system: Extreme events and inter-annual variability. *Environmental Research Letters*, *17*(4). https://doi.org/10.1088/1748-9326/ac4dc8

Schmitz, K., Steffen, B., & Weber, C. (2013). *Incentive or impediment? The impact of capacity mechanisms on storage plants*, Florence School of Regulation.

Somani, A., Barrett, E., Zhou, Z., Chan, G., Middleton, L., Tarel, G., Campbell, A., Botterud, A., Bhatnagar, D., Beckitt, A., & O'Reilley, C. (2024). *An Assessment of Resource Drought Events as Indicators for Long-Duration Energy Storage Needs.* (tech. rep.). Pacific Northwest National Laboratory (PNNL).

Steffen, B., & Weber, C. (2013). Efficient storage capacity in power systems with thermal and renewable generation. *Energy Economics*, *36*. https://doi.org/10.1016/j.eneco.2012.11.007

Williams, P., Potter, C., & Allie, S. (2019). *Battery of the Nation: Operation of storages without perfect foresight* (tech. rep.). Hydro Tasmania.

Xu, B., Korpas, M., & Botterud, A. (2020). Operational Valuation of Energy Storage under Multi-stage Price Uncertainties. *Proceedings of the IEEE Conference on Decision and Control*, *2020-Decem*. https://doi.org/10.1109/CDC42340.2020.9304081

Contact.

MIT CEEPR Working Paper Series
is published by the MIT Center for Energy and Environmental Policy Research from submissions by affiliated researchers.
For inquiries and/or for permission to reproduce material in this working paper, please contact:

General inquiries: ceepr@mit.edu
Media inquiries: ceepr-media@mit.edu

MIT CEEPR

Center for Energy and
Environmental Policy Research

MIT Center for Energy and
Environmental Policy Research
Massachusetts Institute of Technology
77 Massachusetts Avenue, E19-411
Cambridge, MA 02139-4307
USA

ceepr.mit.edu

MASSACHUSETTS INSTITUTE OF TECHNOLOGY