

7471014-tail-aware-arbitrage-of-energy-storage-systems-in-volatile-electricity-markets-w

AVTORJI: Mingli Chen, Hongxun Hui, Yonghua Song

POVZETEK: Raziskava predstavlja nov model za arbitražo energijskih shranilnih sistemov (ESS) v volatilnih električnih trgovinah s negativnimi cijenama. Model uključuje inverzno hiperboliko sine transformaciju cijena kao općenitom metriku povratka, te asimetričnu volatilnost model u transformiranoj domeni. Pristup uključuje stabilizirani scenarijski generator s Jensenovim tipom odzivom i omeđenim tržišnim cijenama. Ugotovljeni je pristup zaštito protiv negativnih gubitaka, čim se održava pozitivna arbitražna profitabilnost.

METODA: Prva metoda uključuje inverznu hiperboliku sine transformaciju cijena (asinh) kao općenitom metriku povratka, koja je dobro definirana i podnosi negativne cijene i nule. Druga metoda je asimetrični volatilnost model u transformiranoj domeni s individualnim općenitim Pareto distribucijama za gornje i donje krajnje dijelove. Tretja metoda uključuje rizik upravljanje na osnovi uslovnog rizika (CVaR) u stohastičkom planiranju modelu.

KLJUČNE UGOTOVITVE:

- Inverzna hiperbolika sine transformacija pomaže uključiti negativne cijene, kontrolirajući deformacije i komprimiraju ekstremne promjene.
- Asimetrični volatilnost model generira stabilne scenarije koje uključuju asimetrične ekstremne cijene.
- CVaR bazirano planiranje bilansira arbitražni profit i rizik.

POMEN ZA TRGOVANJE: Raziskava ponuja pristup zaštito protiv negativnih gubitaka u volatilnim trgovinama s negativnim cijenama, čime se može učiniti stabilnije planiranje i operiranje energijskih shranilnih sistemata. To je posebno koristan za trgovce, koji moraju upravljati rizicem negativnih cijena u arbitražnim operacijama.

Tail-Aware Arbitrage of Energy Storage Systems in Volatile Electricity Markets with Negative Prices

Mingli Chen^{a,b}, Hongxun Hui^{a,b,*}, Yonghua Song^{a,b}

^a State Key Laboratory of Internet of Things for Smart City and Department of Electrical and Computer Engineering, University of Macau, Macao 999078, China

^b Hengqin University of Macau Advanced Research Institute, Zhuhai 519031, China

* Corresponding author: Hongxun Hui.

Abstract

The growing volatility of electricity markets creates arbitrage opportunities for energy storage systems (ESSs). However, increasingly frequent negative prices and extreme events undermine conventional price models and expose schedules to severe downside risk. To address these issues, this paper proposes a tail-aware arbitrage framework for ESSs. First, an inverse hyperbolic sine (asinh) transformation of price changes is introduced as a generalized return metric that remains well-defined under negative prices and zero-crossings. The transformation preserves ordinary movements while compressing extremes. Second, a mean-reverting asymmetric volatility model is fitted in the transformed domain, and separate generalized Pareto distributions characterize the upper and lower tails. A stabilized scenario generator incorporates a Jensen-type drift correction and market-price bounds. Third, risk control based on conditional value-at-risk (CVaR) is embedded in the stochastic scheduling model to balance expected profit against downside protection. The proposed method has three appealing features: (i) the asinh metric accommodates negative prices while controlling distortion and compressing extreme changes; (ii) separate upper- and lower-tail models generate stable scenarios capturing asymmetric price extremes; and (iii) CVaR-based scheduling balances arbitrage profit and downside risk under

constraints. Case studies on Australian National Electricity Market data show that the proposed framework substantially reduces realized tail losses while retaining positive arbitrage profit. Keywords: Energy storage system; Energy arbitrage; Asinh transformation; Tail risk; Conditional value-at-risk.

1. Introduction

The transition toward low-carbon electricity supply changes both the formation of spot prices and the economic value of flexibility. Price formation, revenue sufficiency, and the remuneration of energy-limited resources have been identified as interdependent challenges in zero-carbon markets [1]. The price effects of renewable generation depend on demand, available generation capacity, and the time of delivery [2]. Unequal price responses to increases and decreases in Nordic wind generation further indicate that renewable-driven price effects need not be symmetric [3]. Negative prices therefore require explicit treatment rather than interpretation as measurement errors: their occurrence reflects the interaction of supply conditions, operating constraints, and market incentives [4]. Energy storage systems (ESSs) can exploit temporal price differences and absorb surplus generation; comparative welfare analysis confirms the benefits of storage arbitrage alongside interregional transmission [5]. Nevertheless, storage market participation also depends on regulatory arrangements, market power, and battery lifetime considerations [6]. For an independent ESS, favorable expected spreads can coexist with substantial losses when realized charging or discharging prices differ sharply from anticipated values.

Recent storage research has developed increasingly detailed representations of the operational and commercial decisions. Battery operating policies have been derived for arbitrage and peak shaving under uncertainty [7], whereas renewable generation and storage sizing models

account for optimized participation in multiple electricity markets [8]. Energy and frequency-regulation participation have been coordinated over different time scales while considering battery health [9]. Strategic bidding by a price-making ESS has also been formulated under limited market information [10], and the amplify model represents the remaining flexibility of pooled batteries subject to multiple commitments [11]. Complementary studies address customer-cost and network-voltage trade-offs in virtual power plants (VPPs) [12] and internal battery physics in stochastic grid-forming operation [13]. These contributions establish that realistic storage decisions require explicit operational constraints and market assumptions. For price-taking arbitrage, the associated methodological requirement is a representation of uncertain price trajectories that allows rare adverse outcomes to influence a feasible charging and discharging schedule.

The economic assessment of arbitrage also depends on how battery utilization and the scheduling horizon are represented. Comparisons of optimization formulations incorporating dynamic efficiency and cycle-counting degradation demonstrate the sensitivity of net profit to these modeling choices [14]. A subsequent wind-storage study jointly considers wholesale arbitrage, curtailment mitigation, and battery ageing in the sizing and scheduling problem [15]. Artificial end-of-horizon depletion has been addressed by learning the terminal value of stored energy within two-stage stochastic programming [16]. At the investment and market-design level, block orders have been introduced for day-ahead and real-time arbitrage under distributionally robust chance constraints on wind uncertainty [17]. These studies clarify that operational profit reflects both the treatment of uncertainty and the physical, terminal-state, and market-participation assumptions imposed on storage.

Electricity price forecasting provides the statistical basis for such decisions. Interpretable econometric specifications have been developed across day-ahead to annual horizons, emphasizing fundamental regressors and the treatment of nonstationarity [18]. GARCH and stochastic-volatility models with Student-t innovations remain competitive with machine-learning methods when appropriate explanatory variables and feature selection are used [19]. Extreme learning machines have been combined with bootstrap intervals for German and Finnish day-ahead prices [20]. Event-driven learning further supports frequently updated forecasts of energy and regulation prices [21], while grid-physics-informed stacked learning incorporates capacity constraints and temporal adaptation into real-time price prediction [22]. Incorporating seasonal attention into bidirectional long short-term memory networks improves the prediction of extreme imbalance prices in Great Britain [23]. The result emphasizes the relevance of temporal structure when forecasting exceptional price movements. Despite these advances, satisfactory average forecast accuracy does not itself establish an adequate representation of both extreme tails for storage-risk assessment.

Negative-price observations have also received direct methodological attention. Probabilistic neural-network classification has been integrated with extreme learning machine (ELM) forecasts and bootstrap prediction intervals to distinguish negative- and positive-price conditions [24]. Indicator saturation has been used to represent spikes, structural shifts, seasonality, and autocorrelation in Australian electricity prices [25]. These studies demonstrate that negative values and market discontinuities can be accommodated without imposing a strictly positive price process. The remaining representation issue considered here concerns price changes used in conditional volatility estimation. An unshifted logarithmic return is undefined at zero or across a sign change, while a shifted logarithm depends on an externally selected offset. An invertible

transformation must therefore remain valid across zero while balancing the preservation of ordinary fluctuations against the compression of exceptionally large changes.

Probabilistic forecasting extends the research beyond individual price predictions. Distributional neural networks estimate predictive price distributions and have been assessed using both probabilistic scores and trading outcomes [26]. Normalizing flows have been employed to represent multivariate intraday price differences [27] and to generate conditional full-day price scenarios with periodic retraining to accommodate changing market regimes [28]. Bayesian hierarchical forecasting quantifies parameter uncertainty [29]. Online conformal recalibration improves the coverage reliability of neural-network ensembles [30], while multi-distribution Bayesian learning addresses uncertainty through distributional fusion [31]. For the Australian NEM, point and quantile forecasts have been combined through quantile regression averaging, with spike filtering, residual post-processing, and alternative training windows addressing highly volatile prices [32]. These approaches provide important alternatives for representing forecast uncertainty and dependence. For a multistep storage schedule, however, marginal distributional accuracy or interval coverage must be complemented by a representation of sequential price excursions and their asymmetric effects on the lower tail of realized profit.

Extreme-event research provides a more targeted basis for this requirement. The effects of heavy tails and dependence on the deliverability of electricity hedges have been examined [33], and asymmetric positive- and negative-tail connectedness has been documented across Australian regional markets [34]. Spike probabilities in day-ahead-minus-real-time (DART) spreads have been forecast and assessed for their value as trading signals [35]. More directly, DART spreads have been modeled using a covariate-dependent mixture with separate positive and negative generalized Pareto spike regimes [36]. In the broader econometric literature, conditional extreme-

value models allow time-varying tail shape and scale [37], while a subsequent formulation jointly models extreme value-at-risk and expected shortfall using an integrated tail-shape parameter [38]. These studies establish a methodological basis for distinguishing the two price extremes and assessing their severity. The present study extends this research through an invertible price-change representation and a bounded, multistep scenario generator for physically constrained ESS arbitrage.

The relationship between forecasting and operational value is increasingly explicit in recent research. Data-centric task-oriented forecasting has been developed for photovoltaic-battery operation [39], while decision-focused learning has been integrated with data-driven robust optimization to improve downstream energy-management performance [40]. Both studies challenge the assumption that smaller forecast errors necessarily produce better decisions. At the scheduling level, correlated price uncertainty has been represented through data-driven intervals within a two-stage robust VPP bidding formulation [41]. These approaches address the consequences of uncertainty through forecast training or uncertainty sets. The present study instead retains an explicit conditional price-scenario distribution and examines how its adverse profit realizations affect the scheduling objective.

Conditional value-at-risk (CVaR) is already widely used for this type of operational risk control. CVaR has been combined with distributionally robust scheduling of VPPs aggregating distributed storage and photovoltaics [42]. Storage power has been allocated between arbitrage and renewable-output balancing, with spot-market performance assessed under different risk levels [43]. Related applications include risk-aware sizing and control of building portfolios with thermal storage [44], stochastic scheduling of hydrogen-based flexible loads [45], and risk-based planning of seasonal hydrogen storage and batteries [46]. Expected profit, CVaR, and upside preference

have been incorporated into bilevel virtual bidding [47]. Collectively, these studies establish the relevance of risk preferences across storage technologies and market arrangements. For ESS arbitrage, their operational implications depend on the scenario distribution, the source of uncertainty, and the storage exposure being controlled.

Evidence from direct battery-arbitrage comparisons also cautions against assuming that a risk term necessarily improves realized performance. An assessment of risk measures and rainflow-based degradation costs in the Iberian market finds that the examined stochastic strategies with risk measures underperform a deterministic point-forecast strategy [48]. These findings make the distinction between model-implied risk control and realized economic outcomes particularly relevant. Accordingly, an empirical assessment should report the profitability cost of downside protection and its dependence on risk preference, in addition to the optimization objective.

Distributional reinforcement learning provides another route to uncertainty-aware storage decisions. Return distributions have been learned for battery bidding in joint energy and frequency-control markets using an LLM-coordinated framework [49]. Distributional reinforcement learning has also been formulated for battery arbitrage in the Belgian imbalance settlement mechanism, allowing risk preferences to modify the learned policy while limiting daily battery cycling [50]. This work demonstrates that distributional risk sensitivity can be embedded directly in sequential control. Such policy-learning methods and the scenario-based optimization studies above address complementary aspects of market participation. The research need motivating this paper is the joint treatment of three requirements within a transparent price-taking ESS formulation: a price-change representation with an interpretable scale that remains valid across zero; independently characterized upper and lower extremes that can be reconstructed into stable multistep price scenarios; and a quantified relationship between scenario-based downside protection and realized

arbitrage income. In particular, nonlinear inverse transformation can amplify shocks and alter conditional drift, so adequate fitting in transformed space does not by itself ensure suitable trajectories in the original price units.

To address these gaps, this paper develops a tail-aware framework for ESS arbitrage in electricity markets with negative prices. It combines an asinh-based generalized return metric, a mean-reverting asymmetric GARCH model with separate upper and lower extreme-value tails, and CVaR-based risk control. Cases studies on Australian National Electricity Market (NEM) data show that the proposed framework reduces realized downside risk while retaining arbitrage profit. The main contributions are as follows:

1) An inverse hyperbolic sine (asinh) transformation of electricity price changes is proposed as a generalized return metric for markets with negative prices. Unlike the log-return, the metric is well-defined for all price changes including zero-crossings. The metric is exactly invertible in closed form for scenario generation and adaptively compresses extreme movements while preserving typical ones.

2) A tail-aware stochastic price model is developed in the transformed domain. A mean-reverting Glosten–Jagannathan–Runkle GARCH (GJR-GARCH) component with a diurnal volatility profile captures the central dynamics, while separate GPD models characterize the upper and lower tails.

3) CVaR-based risk control is embedded in the stochastic ESS scheduling model. A risk weight balances expected scenario profit against average profit in the worst scenario outcomes. Evaluating the complete weight range reveals the realized tradeoff between profitability and downside protection.

The remainder of this paper is organized as follows. Section 2 introduces the asinh transformation and derives its scale-parameter selection rule. Section 3 develops the tail-aware price model and the scenario-generation procedure. Section 4 formulates the CVaR-based ESS arbitrage optimization. Section 5 presents numerical results for the Australian NEM, and Section 6 concludes the paper. The research framework of this paper is illustrated in Fig. 1.




Fig. 1. Overall framework of the proposed tail-aware ESS arbitrage method.

2. Electricity Price Change Modeling via Asinh Transformation

2.1 Failure of the Log-return under Negative Prices

Let p_t denote the spot price at interval t . The classical log-return is defined as:

$$r_t^{\log} = \ln\left(\frac{p_t}{p_{t-1}}\right) \quad (1)$$

The log-return is defined only when the ratio p_t/p_{t-1} is positive, which requires p_t and p_{t-1} to have the same sign and to be nonzero. Recent electricity-price studies explicitly accommodate negative-price observations [24] [25]. Retaining these observations exposes three limitations of the log-return.

First, the ratio in (1) is undefined whenever $p_{t-1} = 0$.

Second, the log-return is undefined when the price changes sign. When p_t and p_{t-1} have opposite signs, the ratio is negative and its logarithm does not exist. Such sign changes occur whenever the price crosses zero between two intervals, which is common when negative prices are frequent.

Third, when p_t and p_{t-1} are both negative, the ratio is positive and the logarithm is computable. However, the log-return indicates the wrong direction of the price change. Consider a price rising from $p_{t-1} = -100$ to $p_t = -50$. The price increases by 50, whereas the log-return is $\ln(0.5) = -0.6931$, whose negative sign indicates a decrease. As shown in Fig. 2, the log-return reverses the true direction of the movement.




Fig. 2. Failure of the log-return under negative prices.

Therefore, in any electricity market where negative prices occur, the log-return metric is either undefined or sign-reversed, and any model built on the log-return becomes unreliable. A price-change metric that remains valid for modern electricity markets is required.

2.2 Asinh Transformation of Price Changes

Seasonality is removed before the return metric is applied. Let $\tau(t) \in \{0, \dots, S-1\}$ denote the interval t on a grid of S intervals per day, and let $m_{\tau(t)}$ denote the median price at slot $\tau(t)$ over a trailing window of W days. The deseasonalized price deviation is defined as:

$$x_t = p_t - m_{\tau(t)} \quad (2)$$

The trailing median price is robust against spikes inside the window and uses only information available before interval t , which keeps the construction causal for out-of-sample application.

The generalized return metric is defined as the inverse hyperbolic sine of the scaled deviation change:

$$y_t = \text{asinh}\left(\frac{\Delta x_t}{\theta}\right) = \ln\left(\frac{\Delta x_t}{\theta} + \sqrt{\left(\frac{\Delta x_t}{\theta}\right)^2 + 1}\right) \quad (3)$$

where $\Delta x_t = x_t - x_{t-1}$ denotes the deviation change and $\theta > 0$ is a scale parameter with the units of price.

Four properties make (3) a suitable return metric for markets with negative prices.

Property 1 (Full domain): Let $v = \Delta x_t / \theta$ denote the scaled deviation change. The map $v \mapsto \text{asinh}(v)$ is defined, smooth, and strictly increasing on the whole real line. Therefore, the metric remains well-defined under negative prices, zero prices, and sign changes.

Property 2 (Exact invertibility): The price path is recovered in closed form through:

$$p_t = m_{\tau(t)} + x_{t-1} + \theta \sinh(y_t) \quad (4)$$

Eq. (4) allows simulated paths of the transformed metric to be mapped back to price paths exactly.

Property 3 (Adaptive compression): The derivative of the transformation satisfies:

$$\frac{d}{dv} \operatorname{asinh}(v) = \frac{1}{\sqrt{1+v^2}} \quad (5)$$

Eq. (5) equals one at the origin and decays as $1/|v|$ for large $|v|$. Two asymptotic regimes follow. Inside the quasi-linear zone, the metric reduces to a linear scaling:

$$\operatorname{asinh}\left(\frac{\Delta x_t}{\theta}\right) \approx \frac{\Delta x_t}{\theta}, \quad |\Delta x_t| \ll \theta \quad (6)$$

Inside the compression zone, the metric follows a logarithmic law:

$$\operatorname{asinh}\left(\frac{\Delta x_t}{\theta}\right) \approx \operatorname{sgn}(\Delta x_t) \ln\left(\frac{2|\Delta x_t|}{\theta}\right), \quad |\Delta x_t| \gg \theta \quad (7)$$

where $\operatorname{sgn}(\cdot)$ denotes the sign function. Small price changes are preserved by (6), and extreme price changes are compressed by (7), as shown in Fig. 3.




Fig. 3. Geometry of the asinh return metric.

Property 4 (Symmetry and monotonicity): The transformation is odd and strictly increasing. Upward and downward movements of equal magnitude receive equal weight, and the ordering of price changes is preserved.

Table 1 contrasts the log-return with the proposed metric.

Table 1. Comparison Between the Log-Return and the Asinh Metric

Property	Log-return	Asinh metric
Domain	$p_t p_{t-1} > 0$	$\Delta x_t \in \mathbb{R}$
Negative prices	Undefined or sign-reversed	Well-defined
Zero-crossing	Undefined	Well-defined
Invertibility	Closed form	Closed form
Tail compression	Logarithmic	Logarithmic

2.3 Closed-Form Selection of the Scale Parameter

The scale parameter θ locates the boundary between the regimes (6) and (7). A closed-form selection rule is developed from two interpretable inputs: a distortion tolerance and a coverage requirement.

The degree of linearity retained at a scaled change $v = |\Delta x_t|/\theta$ is quantified through the linearity ratio:

$$R(v) = \frac{\text{asinh}(v)}{v}, \quad v > 0 \quad (8)$$

The ratio measures the fraction of the linear reference preserved by the transformation. Values near one indicate negligible distortion, and values near zero indicate strong compression. Differentiation of (8) is calculated as follows:

$$R'(v) = \frac{h(v)}{v^2} \quad (9)$$

$$h(v) = \frac{v}{\sqrt{1+v^2}} - \text{asinh}(v) \quad (10)$$

Differentiation of (10) is calculated as follows:

$$h'(v) = (1+v^2)^{-3/2} - (1+v^2)^{-1/2} = -v^2(1+v^2)^{-3/2} < 0 \quad (11)$$

The function $h(v)$ starts at $h(0) = 0$ and decreases strictly according to (11). Therefore, the ratio $R(v)$ decreases strictly and monotonically, with $\lim_{v \rightarrow 0^+} R(v) = 1$ and $\lim_{v \rightarrow \infty} R(v) = 0$. The monotonic decrease guarantees a unique boundary for any distortion tolerance $\delta_0 \in (0,1)$:

$$v_{\delta_0} = R^{-1}(1 - \delta_0) \quad (12)$$

Representative boundary values are listed in Table 2. A tolerance of $\delta_0 = 0.12$ yields $v_{\delta_0} = 1.008$, and the quasi-linear zone reduces to the convenient form $|\Delta x_t| \lesssim \theta$. The scale θ is then determined through a coverage requirement on the quasi-linear zone. Let $F_{|\Delta x|}$ denote the empirical distribution of absolute price changes. The linear-zone coverage is defined as:

$$L(\theta) = F_{|\Delta x|}(\theta v_{\delta_0}) \quad (13)$$

Table 2. Linearity Thresholds for Representative Tolerances

Tolerance δ_0	Boundary v_{δ_0}	Quasi-linear zone
0.01	0.248	$ \Delta x_t \leq 0.248\theta$
0.05	0.587	$ \Delta x_t \leq 0.587\theta$
0.10	0.893	$ \Delta x_t \leq 0.893\theta$
0.12	1.008	$ \Delta x_t \leq \theta$
0.20	1.478	$ \Delta x_t \leq 1.478\theta$

The coverage is non-decreasing in θ because distribution functions are non-decreasing. The requirement $L(\theta) \geq 1 - q$ states that at least a fraction $1 - q$ of all price changes fall inside the quasi-linear zone, while the remaining fraction q is reserved for compression. The smallest scale parameter θ^* meeting the requirement follows in closed form:

$$\theta^* = \frac{Q_{1-q}(|\Delta x|)}{v_{\delta_0}} \quad (14)$$

where $Q_{1-q}(|\Delta x|)$ denotes the empirical $(1 - q)$ -quantile of the absolute price changes. Here, θ^* is the minimum scale parameter satisfying the coverage requirement $L(\theta) \geq 1 - q$.

3. Stochastic Price Model with Extreme-Value Tails

3.1 Mean Equation and Volatility Structure

The working series is scaled as $\tilde{y}_t = 100 y_t$ for numerical stability of the likelihood optimization, a standard practice in GARCH estimation. The mean equation includes a reversion regressor built from the lagged deviation:

$$\tilde{y}_t = a_0 + \varphi w_t + \varepsilon_t \quad (15)$$

where a_0 denotes the intercept; φ denotes the reversion coefficient; $w_t = \text{asinh}(x_{t-1}/\theta)$ denotes the reversion regressor; and ε_t denotes the innovation. A negative coefficient φ pulls the deviation back toward zero and captures the rapid decay of price excursions after spikes and negative price events. The regressor is expressed in the same asinh geometry as the response, which keeps the equation well-scaled for arbitrarily large deviations.

The innovation is decomposed into a deterministic diurnal factor, a conditional volatility, and a standardized shock:

$$\varepsilon_t = s_{\tau(t)} \sigma_t z_t \quad (16)$$

where $s_{\tau(t)}$ denotes the diurnal factor at slot $\tau(t)$; σ_t denotes the conditional volatility; and z_t denotes the standardized shock. The shocks z_t are independent and identically distributed, and follow the standardized Student- t distribution with ν degrees of freedom. The diurnal factor s_τ is estimated as the root mean square of $\tilde{y}_t$ per time-of-day slot, smoothed with a circular moving average and normalized to a unit geometric mean. The factor captures the pronounced intraday volatility pattern of electricity markets, in which scarcity windows are several times as volatile as overnight troughs. The factor concentrates simulated price spikes in the hours of elevated volatility.

The conditional variance follows a GJR-GARCH (1,1) recursion on the diurnally standardized innovation $\tilde{\varepsilon}_t = \varepsilon_t/s_{\tau(t)}$:

$$\sigma_t^2 = \omega + (\alpha + \gamma \mathbb{1}\{\tilde{\varepsilon}_{t-1} < 0\})\tilde{\varepsilon}_{t-1}^2 + \beta\sigma_{t-1}^2 \quad (17)$$

where ω denotes the baseline variance; α denotes the shock coefficient; β denotes the variance persistence coefficient; γ denotes the asymmetry coefficient; and $\mathbb{1}\{\cdot\}$ denotes the indicator function, equal to one when the condition holds and zero otherwise. The indicator term allows shocks of different signs to move the variance differently. In equity markets the leverage effect makes $\gamma > 0$. Electricity prices exhibit the opposite pattern: upward shocks accompany scarcity and raise subsequent volatility more than downward shocks, and a negative γ is expected. The persistence of (17) equals $\alpha + \gamma/2 + \beta$ under a symmetric innovation. All parameters in (15)–(17) are estimated jointly by maximum likelihood.

3.2 Extreme Value Tails of the Standardized Residuals

The Student- t density captures the central mass of the standardized residuals $z_t = \tilde{\varepsilon}_t/\sigma_t$, while the far tails may require separate calibration. Let z denote a generic realization of the standardized shock z_t . The POT method is applied to each tail separately. For a threshold u and exceedance $e = z - u > 0$, the Pickands–Balkema–de Haan theorem motivates a high-threshold GPD approximation. Recent conditional POT formulations are discussed in [37] [38]. The GPD is expressed as follows:

$$G_{\xi, \beta_u}(e) = \begin{cases} 1 - \left(1 + \frac{\xi e}{\beta_u}\right)^{-1/\xi}, & \xi \neq 0 \\ 1 - \exp\left(-\frac{e}{\beta_u}\right), & \xi = 0 \end{cases} \quad (18)$$

where the shape parameter ξ controls tail heaviness and the scale parameter β_u depends on the threshold. The threshold is selected using two standard diagnostics. First, the GPD mean-excess function $M(u)$ is affine in the threshold:

$$M(u) = \mathbb{E}[z - u | z > u] = \frac{\beta_{u_0} + \xi(u - u_0)}{1 - \xi} \quad (19)$$

where u_0 denotes a reference threshold, and the empirical plot of $M(u)$ is required to be approximately linear above the selected threshold. Second, the estimated shape parameter is required to be stable across neighboring threshold choices.

Given n_u exceedances $e_1, \dots, e_{n_u}$, the parameters are estimated by maximizing the GPD log-likelihood:

$$\ell(\xi, \beta_u) = -n_u \ln \beta_u - \left(1 + \frac{1}{\xi}\right) \sum_{i=1}^{n_u} \ln \left(1 + \frac{\xi e_i}{\beta_u}\right) \quad (20)$$

The upper tail and the lower tail are estimated independently, because scarcity spikes and negative price events arise from different physical mechanisms and need not share tail indices.

The complete innovation distribution splices the two GPD tails onto the truncated Student- t body. Let u^+ and u^- denote the upper and lower thresholds, and let π denote the tail probability on each side, equal to one minus the threshold quantile. Sampling proceeds by inverse transform. A uniform draw ρ lands in the lower tail ($\rho < \pi$), the central body ($\pi \leq \rho \leq 1 - \pi$), or the upper tail ($\rho > 1 - \pi$), and the corresponding inverse distribution is applied:

$$z = \begin{cases} -\left[u^- + \frac{\beta^-}{\xi^-} \left((1 - \rho')^{-\xi^-} - 1\right)\right], & \rho < \pi \\ F_{t_v}^{-1} \left(F^- + \frac{\rho - \pi}{1 - 2\pi} (F^+ - F^-)\right), & \pi \leq \rho \leq 1 - \pi \\ u^+ + \frac{\beta^+}{\xi^+} \left((1 - \rho')^{-\xi^+} - 1\right), & \rho > 1 - \pi \end{cases} \quad (21)$$

where ρ' denotes an independent uniform draw; $\xi^\pm$ and $\beta^\pm$ denote the tail parameters; $F_{t_v}^{-1}$ denotes the inverse Student- t distribution; and $F^\pm$ denote the Student- t distribution values at the thresholds. The spliced draws are truncated to the range of the in-sample residuals, $z \in [z^-, z^+]$, and re-centered to zero mean. The residual bounds act as an empirical safeguard against extrapolation beyond the estimation sample; the market floor and cap are imposed separately in price space.

3.3 Stabilized Scenario Generation with Jensen-Type Drift Correction

Scenario generation inverts the transformation through (4). The inversion introduces a subtle but consequential bias. Conditional on the information set $\mathcal{F}_{t-1}$, the expected deviation change implied by (3), (15), and (16) follows from the hyperbolic addition identity:

$$\mathbb{E}[\Delta x_t | \mathcal{F}_{t-1}] = \theta [\sinh(\mu_t) K_c(\zeta_t) + \cosh(\mu_t) K_s(\zeta_t)] \quad (22)$$

where μ_t denotes the conditional mean of y_t and $\zeta_t = s_{\tau(t)} \sigma_t$ denotes the effective volatility. The two correction moments are defined as:

$$K_c(\zeta) = \mathbb{E} \left[\cosh \left(\frac{\zeta Z}{100} \right) \right] \quad (23)$$

$$K_s(\zeta) = \mathbb{E} \left[\sinh \left(\frac{\zeta Z}{100} \right) \right] \quad (24)$$

The factor K_c exceeds one and grows rapidly with ζ , in direct analogy with the lognormal correction factor $e^{\sigma^2/2}$. Therefore, a naive generator that sets μ_t from (15) alone can exhibit systematic drift distortion because the nonlinearity of $\sinh$ alters the conditional mean when innovations in y -space are reconstructed in price space. Under heavy-tailed innovations the moments in (23) and (24) are finite only on a truncated support, which provides a second, independent justification for the truncation introduced above. Both moments are precomputed by Monte Carlo integration on a grid of ζ values and evaluated by interpolation during simulation.

The drift correction targets a mean-reverting conditional price change estimated directly in price space:

$$\mathbb{E}[\Delta x_t | \mathcal{F}_{t-1}] = -\kappa x_{t-1} \quad (25)$$

where $\kappa \in (0,1)$ denotes the per-step reversion rate obtained by least squares from Δx_t on x_{t-1} .

Substituting (25) into (22) and linearizing $\sinh(\mu_t) \approx \mu_t$, $\cosh(\mu_t) \approx 1$ for the small conditional means encountered in practice yields the drift correction in closed form:

$$\mu_t = \frac{-\kappa x_{t-1} / \theta - K_s(\zeta_t)}{K_c(\zeta_t)} \quad (26)$$

Eq. (26) plays the same role as the $-\sigma^2/2$ correction in lognormal asset models: the term K_s/K_c offsets the skew of the truncated innovation and the division by K_c offsets the convexity inflation, so that (25) is targeted under the linearization used in (26). Subsequent clipping may further alter the simulated conditional mean.

Three physically motivated bounds complete the generator. First, simulated prices are clipped to $[p^{\text{floor}}, p^{\text{cap}}]$, where p^{floor} and p^{cap} denote the market price floor and cap, respectively; the deviation state is reset consistently after clipping. Second, the one-step transformed move is bounded by the largest move the market can produce:

$$|y_t| \leq y^{\text{max}} = \text{asinh}\left(\frac{p^{\text{cap}} - p^{\text{floor}}}{\theta}\right) \quad (27)$$

The variance recursion is fed with the correspondingly bounded shock. Third, the conditional variance is capped at a high quantile $\tilde{\sigma}^2$ of the in-sample filtered variance, which prevents self-reinforcing volatility growth during long simulations. The variance cap is a numerical stabilization safeguard, whose influence depends on the selected in-sample quantile. The complete procedure is summarized in Algorithm 1.

Algorithm 1: EVT-Based Scenario Generation

- 1: **Input:** deviation state x_0 , filtered variance σ_0^2 , seasonal profile m_τ , diurnal profile S_τ , horizon T , scenario count N .
 - 2: Draw spliced innovations $z_{i,t}$ from (21) for $i = 1, \dots, N$, $t = 1, \dots, T$; truncate to $[z^-, z^+]$ and re-center.
 - 3: **For** $t = 1, \dots, T$ (vectorized over scenarios i):
 - 4: $\varsigma \leftarrow S_{\tau(t)} \sqrt{\sigma_{t-1}^2}$; look up $K_c(\varsigma)$, $K_s(\varsigma)$.
 - 5: Compute the corrected drift μ_t from (26).
 - 6: $y_t \leftarrow \mu_t + \varsigma z_{i,t}/100$; clip to $[-y^{\text{max}}, y^{\text{max}}]$ per (27).
 - 7: $x_t \leftarrow x_{t-1} + \theta \sinh(y_t)$; $p_t \leftarrow \min\{\max\{m_{\tau(t)} + x_t, p^{\text{floor}}\}, p^{\text{cap}}\}$; reset $x_t \leftarrow p_t - m_{\tau(t)}$.
 - 8: Update σ_t^2 by (17) with the bounded shock; cap at $\tilde{\sigma}^2$.
 - 9: **Output:** price scenarios $\{p^{(i)}\}_{i=1}^N$.
-

3.4 Model-Implied Mean Price Path

Under the corrected dynamics (25), the conditional mean of the deviation decays geometrically, and the model-implied mean price path over a horizon of k steps is obtained in closed form:

$$\hat{p}_{t+k} = m_{\tau(t+k)} + x_t (1 - \kappa)^k \quad (28)$$

where $\hat{p}_{t+k}$ denotes the mean price forecast k steps ahead. This closed-form path provides a diagnostic reference for the scenario generator; the scheduling objective in Section 4 uses the sample mean of scenario profits.

4. CVaR-Based ESS Arbitrage Optimization

4.1 Storage Model and Scenario Profits

A price-taking ESS schedules charging and discharging over a horizon of T intervals of length ΔT . Let $P_t^{\text{ch}} \geq 0$ and $P_t^{\text{dis}} \geq 0$ denote the charging and discharging power, and let E_t denote the state of charge (SOC) in energy units. The storage dynamics are given as follows:

$$E_t = E_{t-1} + \left(\eta^{\text{ch}} P_t^{\text{ch}} - \frac{P_t^{\text{dis}}}{\eta^{\text{dis}}} \right) \Delta T, \quad t = 1, \dots, T \quad (29)$$

where η^{ch} and η^{dis} denote the charging and discharging efficiencies, respectively. The ESS operation is subject to the SOC limits, the charging and discharging power limits, and the constraint that prevents simultaneous charging and discharging, expressed as follows:

$$\text{SOC}^{\min} E^{\max} \leq E_t \leq \text{SOC}^{\max} E^{\max} \quad (30)$$

$$0 \leq P_t^{\text{ch}} \leq b_t P^{\max} \quad (31)$$

$$0 \leq P_t^{\text{dis}} \leq (1 - b_t) P^{\max} \quad (32)$$

where $\text{SOC}^{\min}$ and $\text{SOC}^{\max}$ denote the lower and upper SOC limits, respectively; $E^{\max}$ denotes the energy capacity; $P^{\max}$ denotes the power rating; and $b_t \in \{0,1\}$ is a binary variable that prevents

simultaneous charging and discharging. To ensure daily energy neutrality, the terminal energy level of the ESS is constrained to equal its initial value:

$$E_T = E_0 \quad (33)$$

which prevents the optimization from liquidating stored energy for artificial end-of-day profit. Degradation is priced per unit of energy throughput:

$$c^{\text{deg}} = \frac{C^{\text{cap}}}{2 N^{\text{cyc}} E^{\text{max}}} \quad (34)$$

where C^{cap} denotes the capital cost of the storage and N^{cyc} denotes the cycle life. For the price vector $\mathbf{p} = (p_1, \dots, p_T)$, the realized profit is the linear function:

$$\Pi(\mathbf{p}) = \sum_{t=1}^T [p_t (P_t^{\text{dis}} - P_t^{\text{ch}}) - c^{\text{deg}} (P_t^{\text{ch}} + P_t^{\text{dis}})] \Delta T \quad (35)$$

Negative prices enter (35) with the correct economics: charging at a negative price is paid, and discharging at a negative price is penalized.

4.2 CVaR of the Scenario Profit Distribution

Let $\mathbf{p}^{(i)}$ denote the i -th equiprobable price scenario generated by Algorithm 1, and let $\Pi^{(i)}$ denote the corresponding scenario profit defined in (35), where $i = 1, \dots, N$. For a CVaR confidence level $\alpha_c \in (0, 1)$, define $\beta_c = 1 - \alpha_c$ as the lower-tail probability. Accordingly, CVaR represents the expected profit in the lower β_c -tail of the scenario-profit distribution. Using the standard auxiliary-threshold representation employed in recent CVaR-based scheduling and bidding formulations [42] [47], it is calculated as follows:

$$\text{CVaR}_{\beta_c}(\Pi) = \max_{\vartheta \in \mathbb{R}} \left\{ \vartheta - \frac{1}{\beta_c} \mathbb{E}[(\vartheta - \Pi)^+] \right\} \quad (36)$$

where ϑ is an auxiliary profit threshold and $\mathbb{E}[\cdot]$ denotes expectation with respect to the scenario-profit distribution. The positive-part operator is defined as:

$$(\vartheta - \Pi)^+ = \max\{\vartheta - \Pi, 0\} \quad (37)$$

At the optimum, ϑ equals the β_c -quantile of the profit distribution, commonly referred to as the profit VaR. Profits above ϑ make no contribution to the shortfall term, whereas profits below ϑ are penalized according to their distance from the threshold. Therefore, (36) gives the average profit in the lower β_c -tail. Because CVaR is expressed in units of profit, a larger value indicates better downside protection.

For N equiprobable scenarios, the expectation in (36) can be replaced by the sample average. The resulting linear representation is calculated as follows:

$$\text{CVaR}_{\beta_c} = \max_{\vartheta, \{\zeta_i\}} \left\{ \vartheta - \frac{1}{N\beta_c} \sum_{i=1}^N \zeta_i \right\} \quad (38)$$

$$\zeta_i \geq \vartheta - \Pi^{(i)}, \quad \zeta_i \geq 0, \quad i = 1, \dots, N \quad (39)$$

where ζ_i is the shortfall of scenario profit $\Pi^{(i)}$ below ϑ . If $\Pi^{(i)} \geq \vartheta$, then $\zeta_i = 0$; otherwise, $\zeta_i = \vartheta - \Pi^{(i)}$. Because each $\Pi^{(i)}$ is linear in the ESS schedule, (38)–(39) preserve the mixed-integer linear structure and require no additional binary variables.

4.3 Complete Day-Ahead Formulation

The day-ahead ESS schedule is obtained by maximizing a weighted sum of the expected profit and the lower-tail profit:

$$\max_{\mathbf{d}, \vartheta, \{\zeta_i\}} \left\{ (1 - \lambda) \frac{1}{N} \sum_{i=1}^N \Pi^{(i)} + \lambda \left(\vartheta - \frac{1}{N\beta_c} \sum_{i=1}^N \zeta_i \right) \right\} \quad (40)$$

$$\text{s.t. (29)-(33), (39)}$$


where $\mathbf{d} = \{P_t^{\text{ch}}, P_t^{\text{dis}}, E_t, b_t\}_{t=1}^T$ is the day-ahead scheduling vector and $\lambda \in [0, 1]$ is the risk weight.

5. Numerical Results

5.1 Data and Experimental Setup

All experiments use five-minute spot prices from the New South Wales (NSW) region of the Australian National Electricity Market. The full-year 2025 sample is used for model estimation.

446 Data from the first half of 2026 are reserved for out-of-sample evaluation. Fig. 4 and Table 3
 447 summarize the empirical characteristics of the 2025 sample. The ESS parameters follow Table 4.



448 (a) (b)

449 (c)

450 Fig. 4. Empirical characteristics of NSW five-minute spot prices in 2025: (a) full-sample price
 451 distribution; (b) monthly median intraday price; and (c) monthly share of negative-price intervals.

Table 3. Descriptive Statistics of NSW Five-Minute Spot Prices, 2025

Statistic	Value	Statistic	Value
Observations	105,120	Std. dev. (AUD/MWh)	483.4
Mean (AUD/MWh)	103.3	Skewness	28.1
Minimum (AUD/MWh)	−999.99	Kurtosis	872
Maximum (AUD/MWh)	20,300	Negative-price share	14.3%
Zero-crossings	2,993	Price floor / cap	− 1,000 / 20,300

Table 4. ESS and Optimization Parameters

Parameter	Value	Parameter	Value
$E^{\max}$ (MWh)	1	$\eta^{\text{ch}}, \eta^{\text{dis}}$	0.90
$P^{\max}$ (MW)	0.5	$\text{SOC}^{\min}$	0.1
$E_0/E^{\max}$	0.5	$\text{SOC}^{\max}$	0.9
C^{cap} (AUD)	100,000	N^{cyc}	5,000
α_c	0.95	c^{deg} (AUD/MWh)	10
β_c	0.05	N	200

Fig. 4(a) shows that the prices cover both negative and positive values and exhibit pronounced tail behavior. Fig. 4(b) presents the monthly median intraday profiles, revealing a typical midday depression and evening elevation without being dominated by isolated price spikes. Fig. 4(c) shows clear monthly variation in the occurrence of negative prices.

5.2 Effectiveness of the Generalized Return Metric

Table 5 quantifies the effect on the full 2025 sample. The residual excess kurtosis is exactly the volatility clustering that the GARCH component is designed to absorb.

Table 5. Distributional Effect of the Transformation (2025 Sample)

Statistic	Raw Δp_t	Transformed y_t
Range	$[-19,939, 19,879]$	$[-8.29, 8.29]$
Standard deviation	395.5	1.13
Skewness	1.46	-0.02
Kurtosis	1,375	6.2

The scale is selected by the rule of Section 2.3 with the tolerance $\delta_0 = 0.12$ and the coverage fraction $q = 0.35$, as illustrated in Fig. 5. The linearity ratio (8) crosses $1 - \delta_0 = 0.88$ at $v_{\delta_0} = 1.008$, and the coverage curve (13) reaches the target $1 - q = 0.65$ at $\theta^* = 10.0$; the scale $\theta = 10$ AUD/MWh is adopted.




Fig. 5. Data-driven selection of the asinh scale parameter: (a) linearity ratio and distortion-tolerance boundary; and (b) linear-zone coverage and minimum feasible scale.

The two panels in Fig. 5 implement the rule in complementary directions. Fig. 5(a) maps the allowed 12% distortion to the unique scaled-change boundary $v_{\delta_0} = 1.008$. Fig. 5(b) then evaluates how much of the empirical change distribution lies inside that boundary as θ varies. The first scale satisfying the 65% coverage requirement is approximately 10 AUD/MWh; choosing the

smallest feasible value preserves the specified fraction of ordinary changes in the quasi-linear region while compressing larger movements.

5.3 Volatility Model Estimation

Table 6 reports the joint maximum-likelihood estimates of (15)–(17) together with the shifted-log benchmark, which applies the identical pipeline (deseasonalization, diurnal profile, reversion regressor, GJR-GARCH- t) to the shifted-log metric with a shift constant of 1,200 AUD/MWh. Three contrasts emerge.

Table 6. Volatility-Model Estimates and Residual Diagnostics

Parameter	Proposed method	Shifted-log benchmark
a_0	0.509 (0.267)	− 0.007 (0.001)
φ	− 5.633 (0.146)	− 1.480 (0.058)
ω	475.0 (23.2)	0.051 (0.001)
α	0.270 (0.007)	0.410 (0.028)
γ	− 0.133 (0.006)	− 0.142 (0.011)
β	0.759 (0.007)	0.661 (0.009)
ν	24.7 (1.85)	3.51 (0.033)
Persistence $\alpha + \gamma/2 + \beta$	0.962	1.000
Residual kurtosis	3.4	2,232
Residual skewness	0.12	27.7

Note: Standard errors are reported in parentheses for the estimated coefficients.

First, the proposed model is stationary with a persistence of 0.962, well below unity, whereas the shifted-log model remains at the integrated-GARCH boundary with a persistence of 1.000, a symptom of unabsorbed extreme observations.

Second, the standardized residuals of the proposed model have a well-behaved central distribution, with a kurtosis of 3.4 and a skewness of 0.12. The shifted-log residuals retain a

kurtosis of 2,232 and a skewness of 27.7: the volatility filter has failed because the near-linear transformation transmits raw spikes directly into the innovation series.

Third, the reversion coefficient is strongly significant with a half-life of roughly two intervals, and the asymmetry coefficient confirms the inverse leverage effect of electricity prices anticipated in Section 3.1. The estimated tail-thickness parameter indicates that the central body is adequately captured after the diurnal profile and GJR recursion have been absorbed; the remaining extremes are delegated to the EVT component.

Fig. 6 collects the volatility-model diagnostics. The lag one autocorrelation of the innovation magnitude falls by about 80% after filtering, and the remaining autocorrelations lie close to the 95% confidence band. The diurnal factor reaches an evening peak of 1.70 and an overnight trough of 0.60, giving a peak-to-trough ratio of approximately 2.8; the numerical contrast is reported here rather than inside the panel to keep the graphic uncluttered. The Student- t quantile plot shows close central agreement but systematic upper- and lower-tail departures outside the central-fit boundaries, which motivates separate EVT tails. Finally, the news-impact curve and its dashed comparison guides show that a positive standardized shock produces the larger conditional-variance increment, confirming the inverse leverage effect expected for electricity prices.




Fig. 6. Volatility-model diagnostics: (a) autocorrelation of innovation magnitudes before and after filtering; (b) estimated diurnal volatility factor; (c) Student- t quantile plot with central-fit boundaries; and (d) asymmetric news-impact curve.

5.4 Tail Estimation

Fig. 7 presents matched upper- and lower-tail diagnostics for the standardized residuals. In Figs. 7(a) and 7(b), the empirical mean-excess functions $M(u)$ become approximately affine beyond the selected 95th-percentile thresholds. In Figs. 7(c) and 7(d), most empirical exceedance quantiles lie close to the 45° reference line throughout the fitted ranges. The agreement provides a direct distributional check that the two fitted GPDs reproduce the corresponding observed tails.




Fig. 7. Two-sided EVT diagnostics for the standardized residuals: (a) upper-tail mean-excess plot; (b) lower-tail mean-excess plot; (c) upper-tail empirical-versus-fitted quantiles; and (d) lower-tail empirical-versus-fitted quantiles.

Table 7 reports the upper- and lower-tail GPD parameter estimates for the standardized residuals, with the standard errors of the shape and scale estimates reported in parentheses. The larger lower-tail shape estimate supports the use of separate rather than symmetric GPD components.

Table 7. Generalized Pareto Tail Estimates (Standardized Residuals)

Tail	Threshold u	Exceedances	$\hat{\xi}$ (s.e.)	$\hat{\beta}$ (s.e.)
Upper	1.668	5,256	0.042 (0.014)	0.452 (0.009)
Lower	1.535	5,256	0.121 (0.016)	0.374 (0.008)

5.5 Out-of-Sample Validation Design

The out-of-sample evaluation covers 180 days in the first half of 2026 using a rolling day-ahead design. For each day, the seasonal profile and conditional variance are updated using the preceding 28 and 7 days, respectively, and price scenarios are generated from the latest observed state. The schedule obtained from (40) is evaluated against the subsequently realized five-minute prices using (35), without using information from the dispatch day. All risk weights are evaluated with identical inputs, ESS constraints, scenario counts, and daily random seeds. The case $\lambda = 0$ serves as the risk-neutral benchmark, while larger values of λ place greater emphasis on downside protection.

5.6 Out-of-Sample Results

Figs. 8 and 9 summarize realized performance over all 180 evaluation days. The risk-neutral endpoint earns 4,858.6 AUD in total, or 27.0 AUD per day, but its worst day is -132.7 AUD and its realized 5% CVaR is -39.3 AUD. At $\lambda = 0.1$, dispersion decreases but downside protection does not yet improve: the 5% CVaR slightly deteriorates to -41.0 AUD. At $\lambda = 0.2$, the worst day improves to -46.0 AUD and the 5% CVaR to -29.9 AUD, while 45% of the risk-neutral mean profit is retained and 151 of 180 days produce nonzero profit. This is the first stable balance point in the realized risk-return sequence.




Fig. 8. Tail-day profit mitigation at $\lambda = 0.2$.




Fig. 9. Out-of-sample performance across risk weights

Three findings follow from the out-of-sample outcomes.

1) The risk weight sequence contains a genuine threshold rather than a mechanically monotone improvement. At $\lambda = 0.1$, mean daily profit falls by 44% while the 5% CVaR worsens by 4%, so this point does not provide effective risk control. At $\lambda = 0.2$, the worst-day loss is reduced by 65% and the 5% CVaR improves by 24% relative to risk neutrality, with 151 of 180 days still producing nonzero profit. Therefore, the improvement begins only when the CVaR term is strong enough to alter the exposures that dominate realized losses.

2) Fig. 8 resolves the day-level source of the tail improvement at $\lambda = 0.2$. Among the nine days in the risk-neutral worst 5% tail, eight realize a higher profit under the proposed schedule; these eight improved days are displayed, while the single non-improving day is excluded from the graphic but retained in all aggregate statistics. The largest gain occurs on 18 May, when the worst daily loss contracts from -132.7 to -46.0 AUD. On 7 March, 25 June, and 27 June, the proposed schedule raises realized profit to zero and removes the loss completely; on the remaining displayed days, it reduces rather than eliminates the loss. The heterogeneity is expected because the schedule controls exposure ex ante rather than enforcing an ex-post profit target. If the protected schedule avoids the intervals in which adverse prices materialize, the remaining spread can cover round-trip and degradation costs and the realized profit reaches zero. If adverse prices persist across several charging or discharging intervals, the 0.5 MW power limit, the 10–90% SOC window, round-trip losses, the terminal-SOC constraint, and degradation cost prevent full recovery, although the loss depth is still reduced.

3) Fig. 9 reveals the limit of additional risk aversion. Beyond $\lambda = 0.2$, the worst-day profit and 5% CVaR generally continue to improve, but total and mean profit contract rapidly. At $\lambda = 1$, total profit falls to 522.2 AUD and mean daily profit to 2.9 AUD, even though the worst day and 5% CVaR improve to -32.6 and -14.6 AUD. Therefore, the sequence supports $\lambda = 0.2$ as a practical knee: it secures a large initial reduction in realized downside risk before the profitability cost becomes dominant.

The two visual summaries have complementary roles. Fig. 8 fixes $\lambda = 0.2$ and shows the paired realized-profit movement on the eight improved days within the risk-neutral worst 5% tail. Arrow length reports the improvement magnitude, and crossing the zero-profit line distinguishes loss elimination from partial loss mitigation. Fig. 9 displays total profit, worst-day profit, mean

daily profit, and 5% CVaR over the full weight interval using all 180 evaluation days, including the one tail day not displayed in Fig. 8. The color scale is rescaled separately within each row, while the cell labels report the realized values. Together, the figures show both the aggregate effectiveness and the heterogeneous day-level response of CVaR-based risk control without changing the complete-sample assessment.

6. Conclusion

This paper develops a tail-aware framework for price-taking ESS arbitrage under negative electricity prices and asymmetric extremes. Three contributions arise from the formulation. First, an invertible asinh transformation of deseasonalized price changes is calibrated using explicit distortion and coverage criteria, preserving approximate linearity for ordinary fluctuations while compressing extreme changes. Second, a conditional scenario generator combines mean reversion, diurnal asymmetric volatility, and independently estimated upper and lower generalized Pareto tails; a Jensen-type drift correction and simulation bounds address nonlinear reconstruction effects. Third, the resulting scenarios support a physically constrained CVaR scheduling formulation whose realized risk-return trade-off is quantified through rolling out-of-sample comparisons. These contributions provide an explicit methodological basis for translating asymmetric price uncertainty into storage scheduling decisions.

The empirical analysis uses NSW five-minute electricity prices from 2025 for model estimation and data from the first half of 2026 for out-of-sample evaluation. The proposed transformation substantially reduces the influence of extreme price changes and supports a stationary volatility specification. When $\lambda = 0.2$, the CVaR-based schedule reduces the worst-day loss by approximately 65% and the magnitude of the realized 5% CVaR loss by approximately 24%, while retaining 45% of the risk-neutral mean profit. These findings confirm that the proposed

framework can effectively mitigate extreme-loss exposure while maintaining economically meaningful arbitrage returns.

Future work will extend the conditional price model through online regime detection and adaptive parameter updating. The scheduling framework can also be expanded to coordinate energy arbitrage with ancillary service participation under joint price and operational risks.

Acknowledgements

This work was supported by the Science and Technology Development Fund, Macau SAR (0050/2025/AIJ, 001/2024/SKL, and 0002/2025/EQP).

References

- [1] Zhou Z, Botterud A, Levin T. Price formation in zero-carbon electricity markets – Fundamentals, challenges, and research needs. *Renew Sustain Energy Rev* 2025;211:115250. <https://doi.org/10.1016/j.rser.2024.115250>.
- [2] Shimomura M, Keeley AR, Matsumoto K, Tanaka K, Managi S. Beyond the merit order effect: Impact of the rapid expansion of renewable energy on electricity market price. *Renew Sustain Energy Rev* 2024;189:114037. <https://doi.org/10.1016/j.rser.2023.114037>.
- [3] Tselika K, Tselika M, Demetriades E. Quantifying the short-term asymmetric effects of renewable energy on the electricity merit-order curve. *Energy Econ* 2024;132:107471. <https://doi.org/10.1016/j.eneco.2024.107471>.
- [4] Rafizadeh N. The economics of negative price phenomenon in renewable-integrated electricity markets. *Energy Econ* 2026;153:109086. <https://doi.org/10.1016/j.eneco.2025.109086>.
- [5] Yagi C, Takeuchi K. Electricity storage or transmission? Comparing social welfare between electricity arbitrages. *Energy Econ* 2024;140:107969. <https://doi.org/10.1016/j.eneco.2024.107969>.

- [6] Bulut M, Özcan E. How to build a state-of-the-art battery energy storage market? Challenges, opportunities, and future directions. *J Energy Storage* 2024;86:111174. <https://doi.org/10.1016/j.est.2024.111174>.
- [7] Su H, Feng D, Zhou Y, Hao X, Yi Y, Li K. Impact of uncertainty on optimal battery operation for price arbitrage and peak shaving: From perspectives of analytical solutions and examples. *J Energy Storage* 2023;62:106909. <https://doi.org/10.1016/j.est.2023.106909>.
- [8] Bechlenberg A, Luning EA, Saltık MB, Szirbik NB, Jayawardhana B, Vakis AI. Renewable energy system sizing with power generation and storage functions accounting for its optimized activity on multiple electricity markets. *Appl Energy* 2024;360:122742. <https://doi.org/10.1016/j.apenergy.2024.122742>.
- [9] Ma Q, Wei W, Mei S. Health-aware coordinate long-term and short-term operation for BESS in energy and frequency regulation markets. *Appl Energy* 2024;356:122363. <https://doi.org/10.1016/j.apenergy.2023.122363>.
- [10] Xie D, Xu Y. Strategic bidding of price-maker energy storage systems in electricity markets with limited information. *Appl Energy* 2025;390:125824. <https://doi.org/10.1016/j.apenergy.2025.125824>.
- [11] Tiemann PH, Nebel-Wenner M, Holly S, Frost E, Nieße A. Amplify: Multi-purpose flexibility model to pool battery energy storage systems. *Appl Energy* 2025;381:125063. <https://doi.org/10.1016/j.apenergy.2024.125063>.
- [12] Song H, Gu M, Liu C, Amani AM, Jalili M, Meegahapola L, et al. Multi-objective battery energy storage optimization for virtual power plant applications. *Appl Energy* 2023;352:121860. <https://doi.org/10.1016/j.apenergy.2023.121860>.

- [13] Chen Y, Zheng K, Feng C, Huang J, Guo H, Zhong H. Optimal grid-forming BESS management incorporating internal battery physics. *Appl Energy* 2025;385:125448. <https://doi.org/10.1016/j.apenergy.2025.125448>.
- [14] Grimaldi A, Minuto FD, Brouwer J, Lanzini A. Profitability of energy arbitrage net profit for grid-scale battery energy storage considering dynamic efficiency and degradation using a linear, mixed-integer linear, and mixed-integer non-linear optimization approach. *J Energy Storage* 2024;95:112380. <https://doi.org/10.1016/j.est.2024.112380>.
- [15] Grimaldi A, Minuto FD, Perol A, Casagrande S, Lanzini A. Techno-economic optimization of utility-scale battery storage integration with a wind farm for wholesale energy arbitrage considering wind curtailment and battery degradation. *J Energy Storage* 2025;112:115500. <https://doi.org/10.1016/j.est.2025.115500>.
- [16] Han D, Heo S. End-effect mitigation in renewable energy systems with energy storage using value function approximation of terminal energy level. *Appl Energy* 2025;401:126785. <https://doi.org/10.1016/j.apenergy.2025.126785>.
- [17] Belmondo Bianchi A, Rijnaarts HHM, Shariat Torbaghan S. Uncertainty-aware energy storage investment planning through arbitrage in DA and RT markets using novel block orders. *J Energy Storage* 2025;110:115197. <https://doi.org/10.1016/j.est.2024.115197>.
- [18] Ghelasi P, Ziel F. From day-ahead to mid and long-term horizons with econometric electricity price forecasting models. *Renew Sustain Energy Rev* 2025;217:115684. <https://doi.org/10.1016/j.rser.2025.115684>.
- [19] Kapoor G, Wichitaksorn N. Electricity price forecasting in New Zealand: A comparative analysis of statistical and machine learning models with feature selection. *Appl Energy* 2023;347:121446. <https://doi.org/10.1016/j.apenergy.2023.121446>.

- [20] Loizidis S, Kyprianou A, Georghiou GE. Electricity market price forecasting using ELM and Bootstrap analysis: A case study of the German and Finnish Day-Ahead markets. Appl Energy 2024;363:123058. <https://doi.org/10.1016/j.apenergy.2024.123058>.
- [21] Sai W, Pan Z, Liu S, Jiao Z, Zhong Z, Miao B, et al. Event-driven forecasting of wholesale electricity price and frequency regulation price using machine learning algorithms. Appl Energy 2023;352:121989. <https://doi.org/10.1016/j.apenergy.2023.121989>.
- [22] Yi Y, Chen X, Tian Z, Zhang Y. Grid physics-informed and time-adaptive stacked learning for real-time electricity price forecasting: a short-term to mid-term approach. Appl Energy 2026;410:127502. <https://doi.org/10.1016/j.apenergy.2026.127502>.
- [23] Deng S, Inekwe J, Smirnov V, Wait A, Wang C. Seasonality in deep learning forecasts of electricity imbalance prices. Energy Econ 2024;137:107770. <https://doi.org/10.1016/j.eneco.2024.107770>.
- [24] Loizidis S, Venizelou V, Kyprianou A, Georghiou GE. Integrating PNN classification and ELM-Bootstrap for enhanced Day-Ahead negative price forecasting. Appl Energy 2025;392:126013. <https://doi.org/10.1016/j.apenergy.2025.126013>.
- [25] Apergis N, Pan WF, Reade J, Wang S. Modelling Australian electricity prices using indicator saturation. Energy Econ 2023;120:106616. <https://doi.org/10.1016/j.eneco.2023.106616>.
- [26] Marcjasz G, Narajewski M, Weron R, Ziel F. Distributional neural networks for electricity price forecasting. Energy Econ 2023;125:106843. <https://doi.org/10.1016/j.eneco.2023.106843>.
- [27] Cramer E, Witthaut D, Mitsos A, Dahmen M. Multivariate probabilistic forecasting of intraday electricity prices using normalizing flows. Appl Energy 2023;346:121370. <https://doi.org/10.1016/j.apenergy.2023.121370>.

- [28] Hilger H, Witthaut D, Dahmen M, Rydin Gorjão L, Trebbien J, Cramer E. Multivariate scenario generation of day-ahead electricity prices using normalizing flows. *Appl Energy* 2024;367:123241. <https://doi.org/10.1016/j.apenergy.2024.123241>.
- [29] Nickelsen D, Müller G. Bayesian hierarchical probabilistic forecasting of intraday electricity prices. *Appl Energy* 2025;380:124975. <https://doi.org/10.1016/j.apenergy.2024.124975>.
- [30] Brusafferri A, Ballarino A, Grossi L, Laurini F. On-line conformalized neural networks ensembles for probabilistic forecasting of day-ahead electricity prices. *Appl Energy* 2025;398:126412. <https://doi.org/10.1016/j.apenergy.2025.126412>.
- [31] Shao Z, Zhu G, Han Y, Zha J, Yang C, Li F. Multi-distribution fusion based Bayesian deep neural network for short-term probabilistic electricity price forecasting. *Appl Energy* 2025;401:126712. <https://doi.org/10.1016/j.apenergy.2025.126712>.
- [32] Cornell C, Dinh NT, Pourmousavi SA. A probabilistic forecast methodology for volatile electricity prices in the Australian National Electricity Market. *Int J Forecast* 2024;40:1421–1437. <https://doi.org/10.1016/j.ijforecast.2023.12.003>.
- [33] Billimoria F, Mays J, Poudineh R. Hedging and tail risk in electricity markets. *Energy Econ* 2025;141:108132. <https://doi.org/10.1016/j.eneco.2024.108132>.
- [34] Pham SD, Do HX, Nepal R, Jamasb T. Tail risk connectedness in the Australian National Electricity Markets: The impact of rare events. *Energy Econ* 2025;141:108123. <https://doi.org/10.1016/j.eneco.2024.108123>.
- [35] Galarneau-Vincent R, Gauthier G, Godin F. Foreseeing the worst: Forecasting electricity DART spikes. *Energy Econ* 2023;119:106521. <https://doi.org/10.1016/j.eneco.2023.106521>.

- [36] Forgetta A, Godin F, Augustyniak M. Distributional forecasting of electricity DART spreads with a covariate-dependent mixture model. *Energy Econ* 2025;144:108332. <https://doi.org/10.1016/j.eneco.2025.108332>.
- [37] D’Innocenzo E, Lucas A, Schwaab B, Zhang X. Modeling extreme events: time-varying extreme tail shape. *J Bus Econ Stat* 2024;42:903–917. <https://doi.org/10.1080/07350015.2023.2260439>.
- [38] D’Innocenzo E, Lucas A, Schwaab B, Zhang X. Joint extreme value-at-risk and expected shortfall dynamics with a single integrated tail shape parameter. *J Bus Econ Stat* 2026;advance online publication. <https://doi.org/10.1080/07350015.2026.2619541>.
- [39] Huang X, Zhao T, Huang B, Zhang Z, Yue M. Advancing energy system optimization via data-centric task-oriented forecasting: An application in PV-battery operation. *Appl Energy* 2025;378:124753. <https://doi.org/10.1016/j.apenergy.2024.124753>.
- [40] Wang G, Lin Z, Zhong X, Ji H, Li P, Chen Y, et al. Decision-focused learning integrated with data-driven robust optimization for energy management. *Appl Energy* 2026;408:127343. <https://doi.org/10.1016/j.apenergy.2025.127343>.
- [41] Ma Y, Li Z, Liu R, Liu B, Yu SS, Liao X, et al. Data-driven interval robust optimization method of VPP bidding strategy in spot market under multiple uncertainties. *Appl Energy* 2025;384:125366. <https://doi.org/10.1016/j.apenergy.2025.125366>.
- [42] Wang Y, Huang W, Chen H, Yu Z, Hu L, Huang Y. Optimal scheduling strategy for virtual power plants with aggregated user-side distributed energy storage and photovoltaics based on CVaR-distributionally robust optimization. *J Energy Storage* 2024;86:110770. <https://doi.org/10.1016/j.est.2024.110770>.

- [43] Mei S, Tan Q, Trivedi A, Srinivasan D. A two-step optimization model for virtual power plant participating in spot market based on energy storage power distribution considering comprehensive forecasting error of renewable energy output. *Appl Energy* 2024;376:124234. <https://doi.org/10.1016/j.apenergy.2024.124234>.
- [44] Yu MG, Pavlak GS. Risk-aware sizing and transactive control of building portfolios with thermal energy storage. *Appl Energy* 2023;332:120474. <https://doi.org/10.1016/j.apenergy.2022.120474>.
- [45] Chen M, Cao X, Zhang Z, Yang L, Ma D, Li M. Risk-averse stochastic scheduling of hydrogen-based flexible loads under 100% renewable energy scenario. *Appl Energy* 2024;370:123569. <https://doi.org/10.1016/j.apenergy.2024.123569>.
- [46] Zhou S, Han Y, Zalhaf AS, Lehtonen M, Darwish MMF, Mahmoud K. Risk-averse bi-level planning model for maximizing renewable energy hosting capacity via empowering seasonal hydrogen storage. *Appl Energy* 2024;361:122853. <https://doi.org/10.1016/j.apenergy.2024.122853>.
- [47] Xiao D, Peng Z, Lin Z, Zhong X, Wei C, Dong Z, et al. Incorporating financial entities into spot electricity market with renewable energy via holistic risk-aware bilevel optimization. *Appl Energy* 2025;398:126449. <https://doi.org/10.1016/j.apenergy.2025.126449>.
- [48] Camuñas García-Miguel PL, Alonso-Martinez J, Arnaltes Gómez S, Rodríguez-Amenedo JL. Impact of risk measures and degradation cost on the optimal arbitrage schedule for battery energy storage systems. *Int J Electr Power Energy Syst* 2024;157:109883. <https://doi.org/10.1016/j.ijepes.2024.109883>.

- 753 [49] Zhang B, Li C, Chen G, Xu Z, Dong Z. LLM-coordination in auto-bidding of frequency
754 regulation: Cross-attention distributional reinforcement agentic learning. Appl Energy
755 2025;401:126702. <https://doi.org/10.1016/j.apenergy.2025.126702>.
- 756 [50] Karimi Madahi SS, Claessens B, Develder C. Distributional reinforcement learning-based
757 energy arbitrage strategies in imbalance settlement mechanism. J Energy Storage
758 2024;104:114377. <https://doi.org/10.1016/j.est.2024.114377>.
- 759