

Risk-aware battery arbitrage optimisation using Quantile-Based Price Forecast Error Adjustment

Parna Imannezhad, Magnus Soederberg, Neda Todorova,
Phillip Wild

Abstract

Battery operators participating in electricity markets face uncertain price forecasts that impact revenue, particularly under volatile conditions. This paper presents a risk-aware optimisation framework that adjusts pre-dispatch price signals using quantile-based historical forecast error bounds. By leveraging the 5th and 95th percentiles of past forecast errors, the model introduces dynamic price thresholds to guide charge and discharge decisions under uncertainty. It is further extended through a lightweight dispatch layer that preserves a small state-of-charge reserve and enables targeted releases under cap-proximate or high-volatility conditions, improving robustness to unforecast extreme events. The approach is evaluated on 12 representative days from the Australian National Electricity Market in 2025 and compared against both a perfect foresight benchmark and a baseline strategy that relies solely on pre-dispatch prices. Results demonstrate that the model improves revenue outcomes in most scenarios while reshaping the distribution of risks across trading intervals. It offers a practical decision-support tool for battery operators navigating uncertainty in short-term price forecasts.

Keywords: *Battery arbitrage optimisation; Risk-aware dispatch; Forecast error quantiles; Conditional Price-at-Risk; Price volatility.*

JEL Codes: *L94, Q40, D81, C61.*

1 Introduction

Battery Energy Storage Systems (BESS) are increasingly employed in energy-only markets, such as the Australian National Electricity Market (NEM), to perform price arbitrage by charging during low-price intervals and discharging when prices are high. A substantial body of research shows, however, that arbitrage performance is highly sensitive to price-forecast accuracy. Weber and Lu [1] demonstrate that even modest deviations between pre-dispatch and real-time prices can erode revenues by up to 50%, with conservative bidding reducing risk exposure but also lowering expected returns. Prakash et al. [2] show that short-term forecast errors can reduce annual arbitrage revenues by 20–64%, with shorter-duration batteries most affected. Their study concludes that improving forecast quality and operator decision frameworks such as through probabilistic information or dynamic re-optimisation, can materially enhance arbitrage performance under uncertainty. Real-world operational evidence aligns with these findings: the University of Queensland's MPC-operated Tesla Powerpack identified inaccuracies in AEMO's pre-dispatch forecasts as a primary barrier to effective arbitrage, motivating the need for hybrid or probabilistic control strategies [3].

Complementing operational studies, recent financial econometric work highlights why uncertainty in the NEM is inherently difficult to manage. Pham et al. [4] show that electricity price returns in the NEM, particularly in South Australia, deviate substantially from normality and exhibit pronounced skewness and excess kurtosis, as confirmed by formal statistical diagnostics. They employ CAViaR (Conditional Autoregressive Value-at-Risk), which directly estimates time-varying VaR^1 at chosen quantiles without assuming a parametric distribution, by modelling VaR as an autoregression with asymmetric responses to past positive and negative returns. Their goal is to characterise how extreme risks propagate across NEM regions via time-varying connectedness, demonstrating that prices are heavy-tailed and highly interconnected. Billimoria and Simshauser [6] similarly emphasise sensitivity to forecast uncertainty, using CVaR^2 and VaR measures to quantify risk exposure for merchant BESS in South Australia. They model uncertainty by applying Gaussian random variations with zero mean and fixed standard deviations (10%, 20%, or 30%) to spot prices and propose a yardstick contract that compensates operators based on their deviation from a hypothetical perfect-foresight benchmark as a potential financial hedge. Practical implementation, however, requires robust ex-post calculation mechanisms and agreed performance thresholds.

Internationally, risk-averse optimisation approaches in electricity markets similarly emphasise the value of incorporating distributional information. Sheikahmadi et al. [7] incorporate CVaR into a two-stage stochastic procurement model to manage wholesale price uncertainty. Janczura and Wójcik [8] demonstrate that probabilistic forecasts combined with risk measures such as VaR, CVaR, standard deviation and the Sharpe ratio improve short-term trading strategies in German and Polish markets. Li et al. [9] introduce a worst-case CVaR framework for risk-averse market clearing in Singapore's demand-response program, reducing exposure to rare but severe outcomes. Barquín et al. [10] propose an innovative approach that iteratively adjusts scenario probabilities to overweight tail-risk events when price dispersion is high, enabling risk-aware scheduling without embedding CVaR constraints directly in the core optimisation problem. More recently, Sajeev et al. [11] developed a stochastic, risk-aware optimisation framework for battery arbitrage in which price uncertainty is represented through scenario-based modelling and managed using risk measures such as CVaR. Chai et al. [12], reviewing over sixty studies, highlight the rapid growth of probabilistic and quantile-based electricity price forecasting, underscoring the increasing importance of distribution-driven decision models.

Together, these studies underline that forecast error is a first-order determinant of arbitrage profitability in the NEM. Proposed responses range from static bid-banding and dynamic control strategies to contractual hedges and probabilistic methods, but they share the recognition that battery operators must explicitly manage the risks created by price uncertainty. Recent contributions highlight non-normal price behaviour, and the use of quantile-based risk measures further support the case for developing optimization frameworks that incorporate empirical distributional information. Despite these advances, most existing studies treat uncertainty as an exogenous stochastic perturbation or apply VaR/CVaR only to aggregate portfolio outcomes, rather than to link the forecast-error distribution directly to interval-level optimization decisions. Similarly, while probabilistic forecasting methods are increasingly recognized, their operational integration into day-ahead battery scheduling frameworks remains limited.

This study addresses these gaps by developing a risk-aware battery arbitrage optimization model that adjusts pre-dispatch prices using empirical forecast-error distributions and quantile-based risk measures. Operational assumptions reflect real-world practice: decisions are made once per day based on AEMO's pre-dispatch price path published around 00:00 AEST; and a single daily cycle is enforced to reflect economic and degradation considerations. Its performance is benchmarked against perfect-foresight and unadjusted pre-dispatch strategies. Building on the logic of distribution-driven approaches and

¹ Value-at-Risk (VaR) at level α (e.g., 5%) estimates the worst expected loss that will not be exceeded with $(1-\alpha)$ confidence [5]. For example, a 5% VaR of -200 \$/MWh implies that in 95% of cases, losses will not exceed 200 \$/MWh.

² Conditional Value-at-Risk (CVaR), also known as Expected Shortfall, extends VaR by measuring the average loss given that the VaR threshold has been breached, thereby focusing on tail risk (Rockafellar & Uryasev [13]; Artzner et al. [14]).

prior evidence that forecast errors lead to mistimed charging and discharging decisions at specific dispatch intervals, the model embeds Conditional Price-at-Risk (CPR) adjustments, defined as quantile-based and empirically derived margins from historical pre-dispatch errors, within the dispatch problem. This formulation enables battery operators to internalise asymmetric price risks at the interval level, yielding a more realistic and data-driven approach to decision-making under uncertainty.

While the CPR-based day-ahead model materially improves performance under most market conditions, it remains insufficient under extreme-volatility conditions characterised by short-lived, poorly forecast scarcity events. To address this limitation, we introduce a hybrid extension that couples the CPR day-ahead plan with a lightweight real-time decision layer. This layer preserves a small state-of-charge reserve and authorises targeted releases during idle, high-price intervals or cap-proximate conditions, allowing the battery to respond causally to realised scarcity without re-optimising the full schedule.

Across the 12-day case study, perfect-foresight revenues total \$1.41 million, while the unadjusted pre-dispatch strategy captures approximately 55% of this value. Incorporating risk awareness through CPR and its hybrid extension improves outcomes, increasing revenues by around 10% relative to pre-dispatch and reducing average CVaR₅ losses by approximately 3%, indicating less severe downside exposure from an operator's perspective. Overall, the resulting hybrid strategy lifts value capture across all 12 days to about 60–61% of the perfect-foresight benchmark, depending on the reserve setting, using only day-ahead pre-dispatch forecasts and historical forecast-error information available ex-ante, without compromising operational constraints or the model's single-cycle, risk-aware structure.

Although the empirical analysis focuses on the NEM, the modelling framework has broader relevance for wholesale electricity markets that rely on day-ahead or short-term forecast-based scheduling. Many high-renewable systems experience pronounced price volatility and divergence between forward signals and realised outcomes, including gross-pool markets with central dispatch, such as Ireland and Singapore, as well as net-pool markets. In particular, markets such as CAISO, ERCOT, PJM, ISO-NE, and NYISO publish day-ahead price or forecast signals that guide storage decisions, yet these can deviate sharply from real-time prices during scarcity events or rapid net-load changes.

The CPR approach developed here is intentionally distribution-agnostic and forecast-agnostic: it requires only historical forecast errors relative to realised prices, making it readily transferable across market designs. As renewable penetration and merchant exposure increase globally, the need for risk-aware dispatch rules that internalise forecast uncertainty, rather than relying on point forecasts alone, is likely to become increasingly important across diverse electricity market architectures.

The remainder of the paper is organised as follows. Section 2 outlines the modelling framework and CPR-based methodology. Section 3 describes the case-study design and benchmark-day selection. Section 4 presents empirical results and aggregate performance. Section 5 introduces the hybrid strategy extending the CPR framework. Section 6 concludes and discusses future research directions.

2 Methodology and Model Framework

2.1 Alternative Methods Considered and Motivation for the CPR Approach

Before adopting the Conditional Price-at-Risk (CPR) adjustment as the core of our risk-aware optimisation framework, we evaluated several alternative methods for incorporating forecast uncertainty into day-ahead BESS scheduling. Documenting these trials is important, as they illustrate both the limitations of more conventional risk treatments and the rationale for selecting a quantile-based approach.

2.1.1 Monte Carlo and Scenario-Based Simulation

A natural starting point was simulation-driven uncertainty modelling, where Monte Carlo sampling of historical forecast errors generates many alternative day-ahead price paths. While appealing in principle by providing a full probabilistic distribution of outcomes, this approach proved impractical at the five-minute resolution. Thousands of scenarios were required for convergence; runtimes became prohibitive; and results were highly sensitive to simulation seeds. Such computational burdens are incompatible with real-world operational timelines, where day-ahead plans must be generated quickly on a daily basis using the available pre-dispatch information.

2.1.2 Parametric Error Fitting

We then explored parametric representations of forecast-error distributions, fitting normal, Student-t, and Generalised Extreme Value (GEV) distributions. None offered a satisfactory characterisation of the empirical data. Normal fits under-represented tails: t-distributions captured heavy tails but not skewness, and GEV fits frequently failed to converge. A related mean-variance adjustment ($\mu \pm \gamma\sigma$) was also trialled, but the presence of extreme outliers, with errors exceeding $-15,000$ \$/MWh in some intervals, distorted both μ and σ and produced implausible adjusted prices, sometimes even highly negative expected values. These outcomes highlight the difficulty of imposing a parametric structure on the fat-tailed, asymmetric, and volatility-clustered errors observed in NEM forecasts.

2.1.3 CVaR-in-Objective Formulations

We also experimented with embedding CVaR directly into the optimisation objective, maximising expected revenue minus a λ -weighted CVaR penalty. Although theoretically sound and widely used in financial portfolio optimisation, this formulation required substantial expansion of the optimisation problem, including auxiliary variables and binary constraints. Solutions were highly conservative, and revenues were materially depressed unless λ was tuned on a case-by-case basis, undermining robustness, interpretability, and operational simplicity. Given the need for a practical day-ahead dispatch tool, this complexity is regarded as unacceptable.

2.1.4 Motivation for Adopting CPR

The above-mentioned limitations collectively motivated the adoption of CPR as a tractable and empirically grounded alternative. CPR bypasses scenario proliferation, avoids parametric mis-specification and circumvents the auxiliary-variable complexity of CVaR-in-objective models. By applying empirical quantiles from recent forecast-error distributions (e.g., 5th and 95th percentiles) directly to adjust pre-dispatch prices, CPR produces asymmetric, distribution-informed thresholds that reflect the fat-tailed and non-normal behaviour documented in NEM prices. The resulting adjustments are transparent, computationally efficient, and well-suited to real-time operational environments.

2.2 Core Model Formulation

The optimisation framework generates a day-ahead charging and discharging schedule for a grid-scale battery in the NEM under imperfect price information. Australian Energy Market Operator pre-dispatch prices provide the only forward-looking market signal available to storage operators, yet these forecasts can diverge markedly from realised 5-minute prices, particularly during volatile conditions. To address this, the model embeds historical forecast-error information directly into dispatch decisions through a risk-aware adjustment mechanism.

At its core, the model is formulated as a mixed-integer linear optimisation problem that selects charging power, discharging power, and state-of-charge trajectories across 288 five-minute intervals. The formulation respects all key technical constraints and a single daily cycle to reflect practical operating behaviour and battery longevity considerations. The battery is modelled as a price taker, consistent with the operational reality of contemporary BESS units in the NEM.

The framework's distinctive feature is the incorporation of CPR adjustments. Rather than using raw pre-dispatch prices, interval-level prices are modified based on the empirical distribution of recent forecast errors, raising charging thresholds and lowering discharging thresholds when systematic bias is detected. This reduces exposure to arbitrage losses from unexpected price spikes or collapses. A volatility-responsive scaling factor further moderates the adjustment strength, allowing strong responses under high uncertainty while preserving stability in normal conditions.

This architecture preserves computational tractability while embedding interval-level price risk into dispatch decisions. The final optimisation maximises expected net revenue using CPR-adjusted price signals, consistent with the heavy-tailed and asymmetric price behaviour observed in the NEM. Resulting schedules are benchmarked against perfect-foresight and unadjusted pre-dispatch strategies to quantify the value added by CPR-based risk adjustment under realistic market uncertainty.

2.3 Forecast Error Modelling and Risk-Aware Dispatch Implementation

To incorporate forecast uncertainty into the dispatch model, we perform a quantile-based analysis of historical price forecast errors for each 30-minute interval of the trading day. In this framework, only the initial pre-dispatch forecast published around 00:00 AEST is used to represent day-ahead market expectations; subsequent hourly updates issued by AEMO are intentionally excluded to preserve a single-decision structure consistent with day-ahead operational planning.

Forecast errors are defined as the difference between AEMO's pre-dispatch price and the corresponding realised spot price, computed over a rolling four-week window of historical data. This window is long enough to capture prevailing market conditions, yet short enough to adapt to changing regimes. From these empirical error distributions, we derive interval-level quantiles and tail expectations, which form the basis of the CPR adjustments applied in the optimisation model.

Step 1: Data Alignment

Pre-dispatch prices are published at 30-minute resolution, whereas spot prices and the battery dispatch model operate at 5-minute resolution. To align these, we aggregate real prices to the 30-minute level. For each 30-minute interval $i \in \{1, \dots, 48\}$, the average realized price is:

$$\pi_i^{real} = \frac{1}{6} \sum_{k=1}^6 \pi_{i,k}^{real}, \quad (1)$$

where $\pi_{i,k}^{real}$ is the 5-minute real price in subinterval k of 30-minute interval i . These 30-minute averages are then paired with the corresponding pre-dispatch prices.

Step 2: Forecast Error Definition

For each day d and interval i , the forecast error is defined as:

$$\varepsilon_i^d = \pi_{i,d}^{pred} - \pi_{i,d}^{real}, \quad (2)$$

where $\pi_{i,d}^{pred}$ is the predispach price and $\pi_{i,d}^{real}$ is the realized average spot price for the same interval. For each interval i , we collect ε_i^d over the most recent 28 days to form the empirical error distribution used in subsequent steps.

Step 3: Interval-Level Risk Metrics (VaR and Tail Expectations)

From the 28-day error sample for each interval i , we compute:

Volatility (overall uncertainty): Volatility is measured as the standard deviation of the forecast errors:

$$\sigma_i = std(\varepsilon_i), \quad (3)$$

where ε_i denotes the sample of errors for interval i . This captures how dispersed the errors are around their mean, providing a measure of overall forecast uncertainty. Higher volatility implies a greater likelihood that the pre-dispatch signal may deviate materially from realized prices.

Tail quantiles (VaR) at level α : VaR_α represents the forecast-error threshold that is exceeded with probability α in each tail of the empirical distribution. In this context, it quantifies the potential worst-case underestimation (downward tail) or overestimation (upward tail) of market prices in each interval. These quantile thresholds serve as statistical boundaries for the subsequent CPR adjustment. We compute:

$$VaR_\alpha^{down}(i) = quantile_\alpha(\varepsilon_i), \quad VaR_\alpha^{up}(i) = quantile_{1-\alpha}(\varepsilon_i). \quad (4)$$

We set $\alpha = 5\%$ by default, noting that smaller values (e.g., 2.5%) emphasize rarer and more severe forecast errors, while larger values (e.g., 10%) yield milder adjustments. The implications of alternative α choices are examined in Appendix B.

Quantile-based tail expectations (CPR adjustments): To summarize not just the threshold but the average behavior in the tails, we compute conditional means as:

$$Adj_i^{down} = \mathbb{E}[\varepsilon_i | \varepsilon_i \leq VaR_i^{down}], \quad Adj_i^{up} = \mathbb{E}[\varepsilon_i | \varepsilon_i \geq VaR_i^{up}]. \quad (5)$$

Typically, $Adj_i^{down} \leq 0$ captures the average downside deviation when real prices exceed forecasts (underestimation risk) and $Adj_i^{up} \geq 0$ captures the average upside deviation when real prices fall short of forecasts (overestimation risk). These conditional means provide magnitude-aware tail measures (more informative than a single quantile) and are used as data-driven offsets to tighten charge and discharge thresholds in the optimization model. As noted earlier, since we apply CVaR-style expectations to forecast price deviations, we term this construct Conditional Price-at-Risk or CPR.

Step 4: Price Adjustment Mechanism

CPR adjustments are used to modify pre-dispatch prices into risk-aware charging and discharging thresholds. For each interval i , the adjusted thresholds are defined as:

$$\text{Charging threshold: } \tilde{\pi}_i^{charge} = \pi_i^{pred} - \gamma_i Adj_i^{down}. \quad (6)$$

Since $Adj_i^{down} \leq 0$, the adjustment raises the effective charging threshold, discouraging charging unless prices remain sufficiently low after accounting for underestimation risk.

$$\text{Discharging threshold: } \tilde{\pi}_i^{discharge} = \pi_i^{pred} - \gamma_i Adj_i^{up}. \quad (7)$$

Because $Adj_i^{up} \geq 0$, the adjustment lowers the discharging threshold, ensuring that discharging occurs only when prices remain high after subtracting a risk premium.

Adjustments are applied conditionally:

$$\text{Discharge adjustment only if } Adj_i^{up} > 0, \quad \text{Charge adjustment only if } Adj_i^{down} < 0. \quad (8)$$

Step 5: Dynamic Scaling and Safeguard Conditions

The factor γ_i modulates the intensity of CPR-based adjustments. Instead of a static or linear mapping, we employ a sigmoid function that maps interval-level volatility into a bounded risk-sensitivity coefficient. Similar sigmoid transformations have been widely used in probabilistic modelling [15] and in energy price forecasting under uncertainty [16]. The sigmoid function defined in Eq. (9) below provides a smooth and bounded transition between regimes, ensuring that adjustments increase gradually with volatility rather than through abrupt shifts.

$$\gamma_i = \gamma_0 \cdot \left(\frac{1}{1 + e^{-(\sigma_i - v_{\text{thresh}})/100}} \right), \quad (9)$$

where σ_i is the historical volatility for interval i , $\gamma_0=0.2$ is the base risk-sensitivity cap, and $v_{\text{thresh}}=100$ \$/MWh is a volatility threshold. This specification keeps γ_i small when volatility is low and allows it to increase smoothly as volatility rises, reflecting the nonlinear escalation of market risk. Parameter values can be re-tuned for different datasets and were selected via sensitivity testing to balance revenue performance against tail-risk mitigation.

To ensure compatibility with 5-minute settlement, the 30-minute CPR metrics (volatility, γ_i^{down} , γ_i^{up} , and γ_i) are replicated across the six 5-minute subintervals within each half-hour. This preserves the empirical structure of pre-dispatch forecasts while allowing risk-adjusted prices to be applied directly in the 5-minute dispatch model.

To further enhance robustness, a conditional safeguard is applied: when the standard deviation of predispach prices falls below 100 \$/MWh while the average historical interval volatility exceeds 200 \$/MWh, the model reverts to the unadjusted predispach trajectory. This rule ensures that unnecessary skewness is not introduced into dispatch decisions under otherwise stable price regimes. These values were tested against historical distributions and may be re-tuned for different datasets or operational contexts. Conceptually, this logic aligns with the iterative CVaR methodology proposed by Barquin et al. [10], in which scenario probabilities are reweighted to emphasise tail risk only when price dispersion is significant. Whereas their approach achieves this adaptively through dual re-optimisation, our framework embeds similar risk-awareness more simply via percentile filtering and threshold-based moderation. This design avoids overreaction in calm periods while enhancing arbitrage performance under price uncertainty.

2.4 Objective Function

The battery operator is assumed to be a price taker and selects charging and discharging schedules to maximise arbitrage revenue over 288 five-minute trading intervals. Revenue is evaluated using the CPR-adjusted price signals developed in Section 2.3, which incorporate interval-level forecast-error information into the decision process. The optimisation problem is formulated in equation (10) as:

$$\text{Max}_{p_t^{\text{dis}}, p_t^{\text{ch}}} \sum_{t=1}^{288} [p_t^{\text{dis}} \cdot \tilde{\pi}_t^{\text{discharge}} - p_t^{\text{ch}} \cdot \tilde{\pi}_t^{\text{charge}}] \Delta t, \quad (10)$$

where p_t^{dis} and p_t^{ch} denote discharging and charging power at interval t , $\tilde{\pi}_t^{\text{discharge}}$ and $\tilde{\pi}_t^{\text{charge}}$ are the CPR-adjusted discharging and charging price thresholds, and Δt is the 5-minute interval expressed in hours. This optimisation is subject to the following storage dynamics and operational constraints, which ensure that the battery operates within its technical limits.

$$\text{Storage dynamics: } S(t) = S(t-1) + (\mu_c p_t^{\text{ch}} - \mu_d^{-1} p_t^{\text{dis}}) \Delta t, \quad t = 1, \dots, 288. \quad (11)$$

where S_t is the state of charge at the end of interval t , and μ_c and μ_d are charging and discharging efficiencies (both set to 0.92). The initial condition is $S_0 = S_{\text{init}}$.

Operational constraints: Charging and discharging power are bounded by the battery's nameplate rating:

$$0 \leq p_t^{\text{ch}} \leq P_{\text{max}}, \quad 0 \leq p_t^{\text{dis}} \leq P_{\text{max}}, \quad t = 1, \dots, 288. \quad (12)$$

State of charge must remain within allowable limits:

$$S_{\text{min}} \leq S_t \leq S_{\text{max}}, \quad t = 1, \dots, 288. \quad (13)$$

Binary mode variables prevent simultaneous charging and discharging:

$$\text{Charge}_t + \text{Discharge}_t \leq 1, \quad \text{Charge}_t, \text{Discharge}_t \in \{0,1\}. \quad (14)$$

To reflect conservative cycling behaviour, at most one charging-to-discharging transition is permitted per day:

$$\sum_{t=1}^{288} \text{Transition } CD_t \leq 1, \quad \text{Transition } CD_t \in \{0,1\}, \quad (15)$$

where $\text{Transition } CD_t$ indicates a transition from charging to discharging between intervals $t-1$ and t .

Parameterisation: Numerical experiments consider a 25 MW / 50 MWh battery with a two-hour duration. The initial state of charge is set at 5 MWh, and the state-of-charge is restricted to 10–90% of capacity (5–45 MWh). Variable operating costs are set to zero, as operational and longevity effects such as efficiency losses, State-of-charge (SoC) bounds, and limits on daily cycling are already embedded through model constraints and reflect the principal economic components of wear and degradation.

2.5 Assumption of Day-Ahead Single-Plan Optimisation

The present model assumes that the battery commits to a single daily operating plan, optimised at the start of the day using only day-ahead predispach signals published around 00:00 AEST. This assumption is central to the analysis and provides a deliberately conservative, worst-case framework, since dispatch decisions cannot be revised as new market information becomes available.

In practice, battery operators in the NEM can update their bids throughout the day as AEMO revises predispatch forecasts, and they also observe realised prices from past intervals (post-ante) before submitting revised offers for subsequent intervals. However, even with this flexibility, operators face a persistent trade-off: securing immediate, certain revenues versus holding capacity for potentially higher but uncertain returns later in the day. Our day-ahead single-plan assumption magnifies this trade-off to a scale directly comparable with the baseline predispatch strategy. This design choice enables clearer analysis of how the proposed risk-aware model reshapes dispatch incentives under uncertainty, without the confounding influence of rolling re-optimisation.

3 Case Study Design and Day Selection Criteria

To evaluate the robustness of the proposed risk-aware arbitrage strategy, we conduct a systematic statistical analysis of real-time electricity prices in the South Australian (SA) region of the NEM. For each trading day, we compute descriptive statistics including the mean, standard deviation, minimum, maximum, and intraday price range of 5-minute spot prices to quantify both the level and structure of price volatility. Based on these statistics, we apply transparent, SA-specific thresholds to classify daily price behaviour into distinct volatility and structural regimes.³ This classification forms the basis for selecting a diverse and representative case-study portfolio.

Daily market behaviour is categorised into the following regimes:

Calm: low-volatility conditions characterised by a standard deviation below \$60/MWh, an intraday range below \$300/MWh, and maximum prices under \$300/MWh.

Moderate volatility: intermediate conditions, defined as a standard deviation between \$60–200/MWh or an intraday range between \$300–2,000/MWh, without scarcity pricing.

Two-regime (midday–evening) structure: days exhibiting at least 15 negative-price intervals during the solar midday period (10:00–15:00), followed by a pronounced evening peak (17:00–20:00) with maximum prices exceeding \$300/MWh.

Extreme spike: days containing scarcity events, defined as at least one price \geq \$10,000/MWh or multiple spikes \geq \$5,000/MWh.

Negative-to-positive whipsaw: days featuring sharp reversals across both ends of the price distribution, with minimum prices \leq -\$100/MWh and maximum prices \geq \$1,000/MWh within the same day.

This framework ensures that the case study portfolio spans the full spectrum of operational conditions observed in South Australia. From the broader 2025 dataset, we select 12 benchmark days that collectively capture calm periods, moderate deviations, structural two-regime behaviour, extreme scarcity pricing, and intraday whipsaw events. This targeted case-study design aligns with the operator viewpoint, where battery dispatch decisions are made on a day-ahead basis under price uncertainty and risk must be managed at the level of individual trading days rather than long-run averages. Accordingly, the analysis prioritises mechanism-level insight and stress testing of dispatch behaviour under uncertainty, rather than estimation of annualised revenues, which can obscure the impact of rare but value-dominant events. Similar focused or scenario-based evaluations are common in the storage arbitrage literature when analysing risk exposure, forecast error, and operational robustness under extreme price dynamics (e.g. Billimoria and Simshauser, 2023; Weber and Lu, 2024).

³ We adopt a pragmatic regime classification informed by NEM scarcity settings and South Australian price stylized facts documented in official market reports [17, 18]. Thresholds such as standard deviation bands and price ranges are researcher-defined and calibrated to the empirical distribution of South Australian prices, rather than formal regulatory classifications. The "extreme spike" regime is consistent with the AER's use of prices above \$5,000/MWh as a scarcity indicator, while the remaining regimes reflect typical price dispersion reported in AEMO's Quarterly Energy Dynamics. Reports available at:

<https://www.aemo.com.au/energy-systems/major-publications/quarterly-energy-dynamics-qed>.

Table 1 reports descriptive statistics for the selected days, and the subsequent classification panel assigns each date to its corresponding regime. The selected days include:

Calm days: 2-Jun-25 and 14-Jun-25.

Moderately volatile days: 26-Jan-25, 7-Mar-25, 10-Apr-25, and 20-May-25.

Two-regime days: 26-Aug-25 and 15-Sep-25.

Extreme spike days: 1-Feb-25, 12-Feb-25, 12-Jun-25, and 2-Jul-25, with price peaks between \$15,000–\$17,500/MWh.

Whipsaw behaviour: 1-Feb-25, where deep negative prices (–\$125/MWh) rapidly transitioned into a record high (\$17,412/MWh).

By spanning these distinct regimes, the selected dataset provides a balanced and challenging testbed for evaluating the CPR-based optimisation model. It enables us to assess not only improvements in expected revenue but also the model's resilience to sharp forecast deviations, tail-risk outcomes, and highly stressed market conditions.

Table 1. Summary statistics of real 5-minute prices for selected case study days in South Australia

| Date | Mean (\$/MWh) | Std Dev (\$/MWh) | Max | Min | Range |
|------------|---------------|------------------|-----------|---------|-----------|
| 26/01/2025 | 55.73 | 114.02 | 952.67 | -265.59 | 1,218.26 |
| 1/02/2025 | 480.78 | 2,534.82 | 17,411.99 | -125.32 | 17,537.31 |
| 12/02/2025 | 467.78 | 2,170.42 | 17,500.00 | 20.27 | 17,479.73 |
| 7/03/2025 | 43.04 | 69.45 | 320.7 | -32.01 | 352.71 |
| 10/04/2025 | 91.37 | 104.79 | 1,000.00 | -30.44 | 1,030.44 |
| 20/05/2025 | 142.01 | 114.01 | 937.37 | -22.50 | 959.87 |
| 2/06/2025 | 148.85 | 54.37 | 349.03 | 50.01 | 299.02 |
| 12/06/2025 | 1,635.93 | 3,644.86 | 16,983.46 | 99.53 | 16,883.93 |
| 14/06/2025 | 86.02 | 56.13 | 231.31 | -17.51 | 248.82 |
| 2/07/2025 | 1,993.66 | 3,273.00 | 15,103.84 | -0.57 | 15,104.41 |
| 26/08/2025 | 19.41 | 88.71 | 1001 | -13.52 | 1014.52 |
| 15/09/2025 | 2.17 | 73.01 | 815.73 | -21 | 836.73 |

4 Analysis of Results

4.1 Overview of Risk-Aware vs Predispatch Performance

Table 2 reports daily revenues and downside risk metrics across the 12 benchmark days for the predispatch, risk-aware, and perfect-foresight strategies. Perfect foresight serves only as an unattainable upper bound; the central comparison is between the predispatch benchmark and the proposed risk-aware strategy, which applies a 28-day empirical error window to adjust predispatch prices. This subsection provides a high-level assessment of overall performance before examining operational dynamics in detail.

Performance is evaluated using both revenue outcomes and interval-level downside risk measures computed across the 288 five-minute trading intervals. Mean P&L provides the average profit per interval and reflects overall profitability. VaR₅ quantifies the worst loss expected under normal operating conditions, whereas CVaR₅ measures the expected loss conditional on this tail event. The Sortino ratio benchmarks mean returns against downside volatility, providing a measure of risk-adjusted profitability. These downside measures are particularly relevant for storage operators, as large interval losses can impose financial pressure even when total daily revenues remain positive.

Across the benchmark set, the risk-aware strategy generally improves mean interval returns and moderates severe downside outcomes relative to the predispatch benchmark. Gains are most consistent on moderate-volatility and two-regime days, where more selective charging and filtering of noisy price signals improve timing and reduce tail exposures. On extreme-spike days, outcomes are more heterogeneous. In two of the four extreme cases, the risk-aware strategy either reduces CVaR₅ losses or achieves higher revenues than pre-dispatch, while in the remaining cases performance deteriorates due to unforecast scarcity timing or premature energy release. Perfect foresight consistently outperforms both strategies, as expected, but is included only as a reference point for the maximum attainable value. Overall, Table 2 indicates that integrating empirical forecast-error information into dispatch decisions yields measurable benefits across a broad set of conditions.

Table 2: Comparison of Daily Revenues and Downside Risk Measures for Predispatch, Risk-Aware, and Perfect Foresight Strategies

| Date | Strategy | Revenue (\$) | Mean_P&L | VaR ₅ | CVaR ₅ | Sortino |
|-----------|-------------------|--------------|----------|------------------|-------------------|---------|
| 26-Jan-25 | Risk-aware | 12,315.80 | 42.763 | -0.000 | -0.000 | ∞* |
| | Perfect Foresight | 16,285.90 | 56.548 | 0 | -0.000 | ∞* |
| | Predispatch | 11,173.70 | 38.798 | 0 | -0.000 | ∞* |
| 1-Feb-25 | Risk-aware | 11,305.60 | 39.256 | 0 | -0.065 | 41.475 |
| | Hybrid | 112,693.44 | 391.3 | 0 | 0 | 47.287 |
| | Perfect Foresight | 264,595.80 | 918.735 | -0.000 | -0.000 | ∞* |
| | Predispatch | 11,809.30 | 41.005 | 0 | -0.000 | ∞* |
| 12-Feb-25 | Risk-aware | 190,322.50 | 660.842 | -197.533 | -290.705 | 9.525 |
| | Perfect Foresight | 196,025.90 | 680.646 | -97.119 | -144.369 | 19.499 |
| | Predispatch | 192,208.00 | 667.389 | -251.738 | -317.185 | 8.771 |
| 7-Mar-25 | Risk-aware | 7,669.60 | 26.631 | -0.000 | -1.582 | 1.953 |
| | Perfect Foresight | 8,841.30 | 30.699 | -0.000 | -0.000 | ∞* |
| | Predispatch | 7,635.50 | 26.512 | -0.000 | -1.616 | 1.938 |
| 10-Apr-25 | Risk-aware | 5,782.00 | 20.076 | 0 | 0 | ∞* |
| | Perfect Foresight | 9,734.50 | 33.8 | 0 | -0.000 | ∞* |
| | Predispatch | 5,698.20 | 19.785 | 0 | -0.000 | ∞* |
| 20-May-25 | Risk-aware | 10,456.00 | 36.308 | -23.096 | -44.402 | 3.489 |
| | Perfect Foresight | 14,328.00 | 49.924 | -0.000 | -9.87 | 16.194 |
| | Predispatch | 10,348.00 | 35.934 | -26.25 | -51.66 | 2.953 |
| 2-Jun-25 | Risk-aware | 2,289.70 | 7.95 | -175.020 | -242.320 | 0.13 |
| | Perfect Foresight | 7,188.40 | 24.96 | -141.15 | -166.76 | 0.6 |
| | Predispatch | 1,635.90 | 5.68 | -175.017 | -298.933 | 0.08 |
| 12-Jun-25 | Risk-aware | 362,671.80 | 1,259.28 | -324.490 | -384.798 | 13.13 |
| | Hybrid | 377,446.48 | 1,310.58 | -324.490 | -384.798 | 14.57 |
| | Perfect Foresight | 445,241.60 | 1,545.98 | -280.04 | -289.80 | 20.53 |
| | Predispatch | 341,983.80 | 1,187.44 | -283.23 | -368.65 | 12.77 |
| 14-Jun-25 | Risk-aware | 4,922.80 | 17.09 | 0 | -0.65 | 4.83 |
| | Perfect Foresight | 7,079.10 | 24.58 | 0 | 0 | ∞* |
| | Predispatch | 4,923.00 | 17.09 | 0 | -0.65 | 4.84 |
| 2-Jul-25 | Risk-aware | 93,296 | 323.94 | -1,157.4 | -2,279.4 | 0.47 |
| | Hybrid | 122,677.92 | 425.97 | -1,157.4 | -2,279.4 | 0.63 |
| | Perfect Foresight | 415,633.70 | 1,443.17 | -372.858 | -499.997 | 11.979 |
| | Predispatch | 178,187.50 | 618.707 | -1,157.396 | -2,307.819 | 0.893 |
| 26-Aug-25 | Risk-aware | 7,255.20 | 25.191 | -0.000 | -0.000 | ∞* |
| | Perfect Foresight | 10,115.90 | 35.124 | -0.000 | 0 | ∞* |
| | Predispatch | 7,199.10 | 24.997 | -0.000 | 0 | ∞* |
| 15-Sep-25 | Risk-aware | 3,400.70 | 11.808 | -0.000 | 0 | ∞* |
| | Perfect Foresight | 10,629.10 | 36.907 | -0.000 | 0 | ∞* |
| | Predispatch | 3,387.70 | 11.763 | -0.000 | 0 | ∞* |

Note: "∞" indicates an undefined or extremely large Sortino ratio resulting from zero downside deviation (i.e., no negative return intervals).

Note: Grey-shaded rows report outcomes from the hybrid extension of the risk-aware strategy, which is formally introduced and discussed in Section 5. These rows are included here solely to enable direct comparison of portfolio performance metrics within a single table.

Table 3 complements this quantitative assessment by summarizing qualitative market characteristics, including daily price structure, forecast accuracy, and strategy behavior, across the same 12 days. It highlights how each strategy responds to specific market regimes such as moderate volatility, extreme scarcity spikes, calm days, and duck-curve patterns. Detailed operational analysis for all benchmark days is presented in Section 4.2 and Appendix A.

Table 3: Summary of benchmark day types, price characteristics, and strategy performance

| Date | Primary category | key stats/shape | Predispatch vs Real Prices (Forecast performance) | Risk-aware vs Predispatch strategy (CPR performance) |
|-----------|------------------------------|---|--|---|
| 26-Jan-25 | Moderate volatility | Calm day, short late spike; benign tails | Real spike earlier than PD | RA \approx +10% > PD via better timing |
| 1-Feb-25 | Extreme spike (also whipsaw) | Min $-\$125$, Max $\$17,412$ | PD missed the spike | Both missed cap; PD > RA slightly on ramp |
| 12-Feb-25 | Extreme spike | Max $\$17,500$; spike forecast | PD signaled spike in right window | PD > RA revenue; RA tighter tails |
| 7-Mar-25 | Moderate volatility | $\sigma = \$69$; range = $\$353$ | PD \approx Real (close track) | RA > PD small edge (rev & CVaR) |
| 10-Apr-25 | Moderate volatility | Single $\$1,000$ blip; broad evening rise | PD \approx Real; midday spike ignored (1-cycle constraint) | RA > PD small; tails ~ 0 |
| 20-May-25 | Moderate volatility | One brief real-price blip; evening ramp | PD over-forecast | RA > PD small; RA tighter tails |
| 2-Jun-25 | Calm | $\sigma = \$54$, range = $\$299$, Max < $\$300$ | Close alignment | RA \gg PD (timing uplift) |
| 12-Jun-25 | Extreme spike | Max $\$16,983$; multiple $>\$10k$ | PD roughly in window but overstated | RA > PD revenue; RA worse tails; Sortino \sim equal |
| 14-Jun-25 | Calm | Low range; minor late variation | PD \approx Real | RA = PD (safeguard condition is triggered and RA reverts to PD) |
| 2-Jul-25 | Extreme spike | Max $\$15,104$; multi-cluster | Only late cluster truly capped | RA < PD (early release depleted SoC) |
| 26-Aug-25 | Two-regime (duck) | Long morning negatives; eve $\$230-\390 | PD missed short noon spike; understated evening plateau | RA > PD small (avoids late top-ups) |
| 15-Sep-25 | Two-regime (duck) | Long negatives; short late spike | PD only partially caught late spike | RA \approx PD (tiny RA edge via charging selectivity) |

Note: PD is AEMO pre-dispatch prices when comparing forecast signals with realized prices, and the Predispatch strategy when referring to dispatch outcomes based on unadjusted AEMO pre-dispatch prices.

Note: RA is Risk-aware strategy using CPR-adjusted thresholds.

Across the benchmark days, several regime-specific patterns emerge. On moderate-volatility days (26 Jan, 7 Mar, 10 Apr, 20 May), the risk-aware strategy consistently delivers small but meaningful gains by refining charge-discharge timing and filtering out noisy forecast signals. On calm days (2 Jun, 14 Jun), forecast and realised prices are closely aligned, leaving little room for adjustment; nevertheless, small timing improvements can still translate into incremental gains, as observed on 2 Jun. On two-regime "duck-curve" days (26 Aug, 15 Sep), both strategies navigate long negative-price periods, but the risk-aware model achieves marginal improvements through more selective midday charging and avoidance of inefficient late-day top-ups. The largest discrepancies arise on extreme-spike days (1 Feb, 12 Feb, 12 Jun, 2 Jul), where pre-dispatch forecasts often mis-time or understate scarcity events. On these days, the risk-aware strategy can either reduce downside tail losses at the cost of slightly lower revenues relative to pre-dispatch or outperform pre-dispatch in terms of revenue while accepting somewhat heavier tails (e.g., 12 Feb and 12 Jun). However, performance is not uniform across all extreme conditions. On 1 Feb, both forecast-based strategies fail to capture the unanticipated super-spike, while on 2 July, early energy release leads the risk-aware strategy to underperform relative to pre-dispatch. These regime-level patterns motivate a closer operational examination in Section 4.2, allowing the model's behaviour to be assessed across different market conditions and to better understand how performance varies from day to day.

4.2 Operational and Dispatch Dynamics Across Representative Days

To illustrate how the proposed strategy behaves across different market regimes, we examine representative benchmark days spanning calm, moderate-volatility, and extreme-spike conditions. For each selected day, the figures compare realised and pre-dispatch prices, CPR-adjusted charge and discharge thresholds, and the resulting charge-discharge schedules and

state-of-charge (SoC) trajectories under the pre-dispatch, risk-aware, and perfect-foresight strategies. These examples highlight the operational mechanisms behind the performance patterns summarised in Tables 2 and 3. Detailed interpretations for the remaining benchmark days are provided in Appendix A.

26 January 2025 (Moderate Volatility):

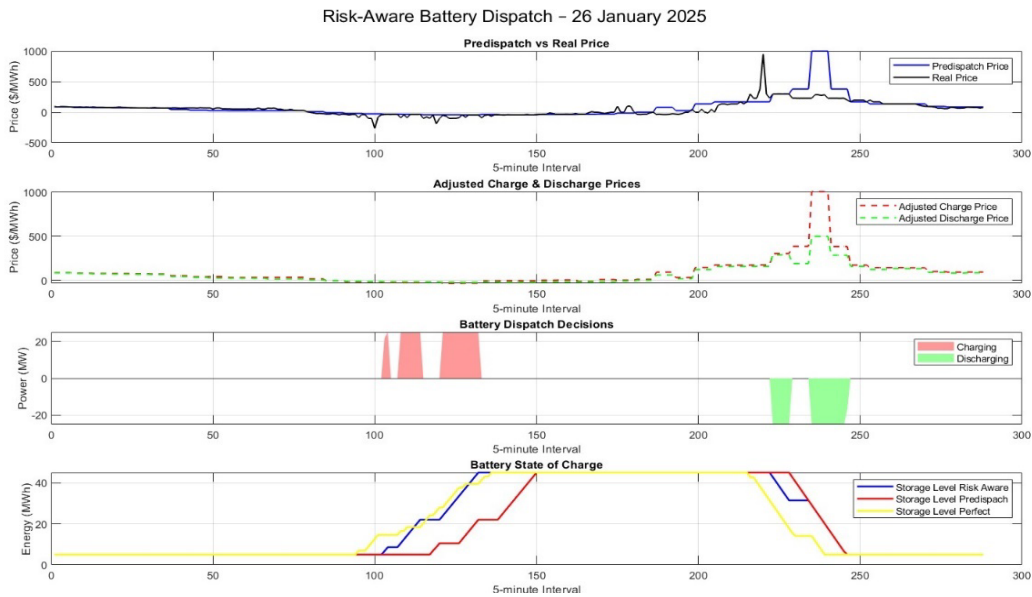


Figure 1: Electricity price trajectories and battery operational responses under three strategies on 26 January 2025

Figure 1 shows dispatch outcomes on 26 January 2025, a moderate-volatility day characterised by generally subdued prices, a brief negative-price pocket ($\approx 100\text{--}130$), and an evening ramp in which realised prices rise slightly earlier than indicated by pre-dispatch forecasts (≈ 220 rather than $230\text{--}240$).

Perfect foresight charges during the low-price window ($\approx 95\text{--}135$) and discharges cleanly across $\approx 220\text{--}250$, capturing the full price ramp and achieving the highest revenue (\$16,285.89).

Risk-aware begins charging around interval 105 and reaches a high state of charge by ≈ 135 . It enters the evening ramp earlier than pre-dispatch, pauses briefly, and then discharges selectively after ≈ 236 , earning \$12,315.80, a 10.2% uplift relative to pre-dispatch.

Pre-dispatch charges later ($\approx 120\text{--}150$) at higher prices and also delays discharge, missing the earliest portion of the realised plateau and yielding \$11,173.70.

Downside risk is negligible: VaR_5 and CVaR_5 are effectively zero for all strategies (Table 2). Overall, on this benign day with a single evening opportunity window, the risk-aware strategy improves timing and outperforms predispatch, while perfect foresight remains the upper bound.

1 February 2025 (Unforecasted Super-Spike):

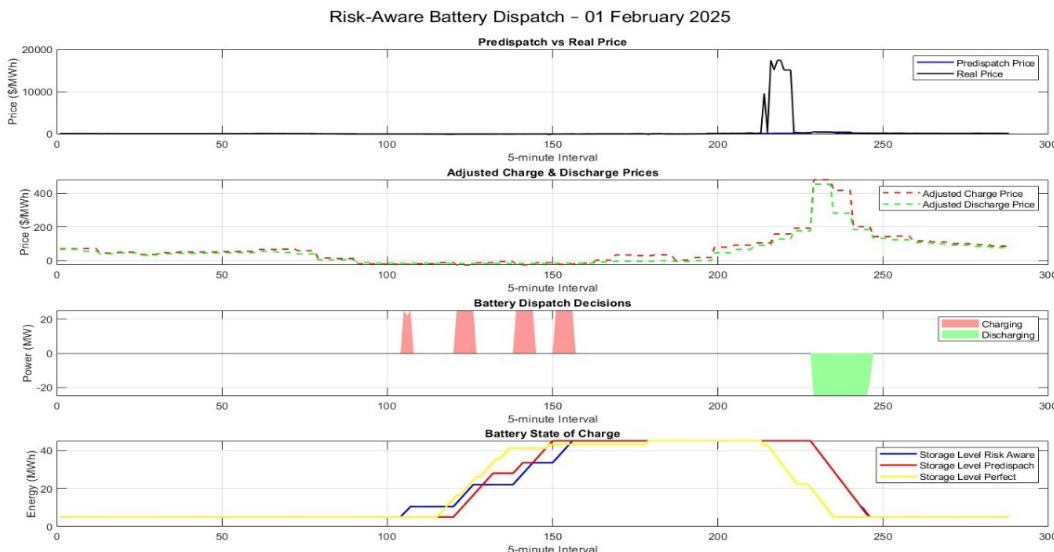


Figure 2: Electricity price trajectories and battery operational responses under three strategies on 1 February 2025

Figure 2 shows dispatch outcomes on 1 February 2025, an exceptionally volatile day ($\sigma \approx \$2,539/\text{MWh}$; range $\approx \$17,537/\text{MWh}$). Prices fall to $-\$125/\text{MWh}$ in the late morning before surging to $\$17,412/\text{MWh}$ in the early evening. The predispatch forecast completely misses this super-spike, indicating only a mild evening plateau.

Perfect foresight charges during the negative-price midday period and discharges entirely through the scarcity window, capturing the spike and achieving the highest revenue ($\$264,596$).

Risk-aware charges slightly earlier than perfect foresight but, with no spike signalled day-ahead, discharges only during the late-day ramp and plateau after the spike has passed, yielding $\$11,305$.

Predispatch begins discharging slightly later than the other strategies, capturing more of the pre-spike price ascent and earning $\$11,809$, marginally outperforming the risk-aware strategy.

Risk metrics show negligible downside exposure: VaR_5 and $\text{CVaR}_5 \approx 0$ for both forecast-based strategies, and Sortino ratios are effectively infinite. Risk-aware exhibits a very small negative tail ($\text{CVaR}_5 = -0.065$). Overall, because the super-spike is not signalled day-ahead, both strategies miss it; predispatch gains a slight advantage through earlier release, while perfect foresight remains the upper bound.

2 June 2025 (Calm Conditions):




Figure 3: Electricity price trajectories and battery operational responses under three strategies on 2 June 2025

Figure 3 shows dispatch outcomes on 2 June 2025, the calmest day in the benchmark set ($\sigma = \$54/\text{MWh}$; range $\approx \$299/\text{MWh}$). Prices remain moderate throughout the day, with two modest morning spikes ($\approx 50-115$) and a late-day rise ($\approx 240-270$), creating low-risk but time-sensitive arbitrage opportunities.

Perfect foresight charges early to prepare for the morning spikes and delays depletion until after interval 240, capturing the highest afternoon prices and achieving the highest revenue ($\$7,188$).

Risk-aware adopts a slightly more conservative morning profile, shifting part of its charging to intervals 145–148 when prices are lower. It also pauses during a mid-afternoon dip (230–235) before discharging into the later ramp, behavior that is absent under Predispatch. These refinements yield $\$2,290$, substantially outperforming Predispatch.

Predispatch charges earlier at higher prices and discharges prematurely around intervals 210–235, missing the strongest late-day prices and earning only $\$1,636$.

Risk metrics reinforce these patterns: the risk-aware strategy achieves a higher mean return (7.95 vs 5.68), lower downside severity ($\text{CVaR}_5 -242$ vs -299), and a slight Sortino improvement (0.13 vs 0.08), indicating a more favorable risk-return profile. Overall, on this low-volatility day, the risk-aware model improves both timing and profitability relative to Predispatch by aligning decisions more closely with realized price dynamics while remaining below the perfect-foresight upper bound.

12 June 2025 (Extreme Spike Day)




Figure 4: Electricity price trajectories and battery operational responses under three strategies on 12 June 2025

Figure 4 illustrates dispatch outcomes on 12 June 2025, the most volatile day in the benchmark set ($\sigma = \$3,645/\text{MWh}$; range = $\$16,884/\text{MWh}$). Real prices reach the Market Price Cap ($\$16,983/\text{MWh}$) and exhibit multiple spikes above $\$10,000/\text{MWh}$, creating an exceptionally challenging environment for forecast-based strategies.

Perfect foresight begins charging earliest, ensuring full capacity before the main spike, and discharges precisely within the peak cluster. It achieves the highest revenue ($\$445,242$).

Risk-aware charges later but still reaches full SoC before the spike window. It begins discharging slightly earlier than Perfect foresight, avoids reacting to transient dips ($\approx 220-225$), and remains selective enough to capture a portion of the spike. This filtering of false signals yields $\$362,672$, outperforming Predispatch.

Predispatch follows the raw forecast trajectory, leading to mis-timed actions, including charging during intervals that appear low in the forecast and discharging fully in response to forecasted cap events that do not fully materialise. Although it captures part of the spike, these inefficiencies limit revenue to $\$341,984$.

Risk metrics reflect this trade-off: Risk-aware improves mean interval returns (1,259 vs. 1,187) and total revenue but experiences deeper tail losses ($\text{CVaR}_5 = -324$ vs. -283), leading to similar risk-adjusted performance (Sortino ≈ 13 for both). Perfect foresight maintains a superior balance, combining the highest mean return (1,546) with the strongest Sortino ratio (20.5).

2 July 2025 (Hyper-Volatile Multi-Spike Day):




Figure 5: Electricity price trajectories and battery operational responses under three strategies on 2 July 2025

Figure 5 illustrates dispatch outcomes on 2 July 2025, a hyper-volatile day characterised by multiple spike clusters ($\sigma \approx \$3,273/\text{MWh}$; range $\approx \$15,104/\text{MWh}$). Real prices remain moderate through the morning before erupting into successive peaks, with the highest prices concentrated in a late-day cap event.

Perfect foresight times charging to retain inventory and concentrates nearly all discharge within the $\sim 240\text{--}280$ interval window, capturing the late-day cap event and achieving the highest revenue (\$415,634) with relatively controlled downside Sortino ≈ 12).

Risk-aware responds more aggressively to elevated volatility. Morning spike signals combined with lower adjusted discharge thresholds around intervals 220–250 trigger early filling followed by premature release. As a result, the battery is largely depleted before the decisive late-day cap event, limiting revenue to \$93,296.

Predispatch does not respond to the morning spike signals and delays charging until around interval 170. Although some energy is released during earlier price increases, a sufficient state of charge is preserved to monetise part of the late spike around intervals 220–245, yielding \$178,187.

Risk metrics highlight this divergence: Tail-risk measures (VaR_s and CVaR_s) are similar for the Risk-aware and Predispatch strategies because losses occur in the same early intervals. However, risk-adjusted performance differs sharply. The risk-aware strategy exhibits a lower Sortino ratio (0.47) than predispatch (0.89), reflecting reduced upside capture rather than increased downside exposure.

4.3 Aggregate Findings and Limitations

Across the twelve benchmark days, the risk-aware CPR strategy demonstrates consistent advantages over the baseline pre-dispatch approach. By elevating charging thresholds and moderating discharging decisions, CPR frequently avoids inefficient trades during noisy intervals and aligns dispatch more closely with realised price dynamics. These improvements are most evident under calm to moderately volatile conditions, where forecast errors are systematic rather than extreme.

Two important limitations nevertheless emerge. First, when scarcity events are largely unforecasted, both pre-dispatch and CPR-based strategies fail to capture the full value of price spikes, and outcomes depend primarily on timing luck rather than optimisation, as observed on 1 February. Second, on highly volatile days with multiple spike clusters, such as 2 July, volatility-scaled CPR adjustments can misfire, triggering premature discharge and leaving insufficient state-of-charge for the terminal scarcity event.

These findings indicate that while CPR materially improves upon raw pre-dispatch forecasts, further refinement is required under the most challenging market regimes. As an initial refinement, we examine sensitivity to the tail parameter α (2.5%, 5% baseline, and 10%) to assess the robustness of the CPR formulation. Most benchmark days exhibit minimal sensitivity: tightening the tail from $\alpha = 5\%$ to $\alpha = 2.5\%$ produces virtually identical dispatch paths and risk outcomes because the corresponding quantiles yield similar CPR adjustments once scaled, and state-of-charge constraints limit behavioural divergence. By contrast, loosening the tail can matter on cap-proximate days. On 2 July, setting $\alpha = 10\%$ tempers premature morning discharge and preserves more energy for the late cap spike, increasing revenue from approximately \$93.3k ($\alpha = 5\%$) to \$101.5k, while VaR and CVaR metrics remain broadly unchanged. Full results are reported in Appendix B. Overall, $\alpha = 5\%$ remains a reasonable default, balancing robustness and responsiveness within a single day-ahead planning framework. While α -sensitivity analysis provides a modest robustness check, it does not resolve the fundamental limitations of the CPR approach under extreme or weakly forecast scarcity conditions. Addressing such cases requires a more substantive extension of the framework.

Accordingly, this study develops a hybrid strategy that combines the static, risk-aware CPR schedule with a limited real-time decision layer. The core idea is to reserve a small portion of the battery's state-of-charge for unanticipated spike events and to permit targeted overrides of the CPR plan only when realized prices approach the market cap. This design introduces flexibility precisely where CPR is most constrained, while preserving the single-plan, risk-aware structure of the base model.

5 Hybrid Risk-Aware Dispatch

5.1 Motivation

Cap-price events in the NEM are rare but often dominate annual arbitrage revenues. While the CPR-based strategy improves performance under moderate and typical volatility conditions, two classes of days remain particularly challenging: (i) unforecast or weakly forecast scarcity events, where predispatch signals fail to indicate the true spike window and the real spike becomes observable only once prices surge in real time (e.g., 1 February), and (ii) hyper-volatile, multi-cluster days, where early price bursts prompt CPR-driven discharging, leaving insufficient state-of-charge for the late-day terminal spike (e.g., 2 July).

To address these limitations without abandoning the single-plan framework, we introduce a lightweight wrapper that adds a limited degree of real-time adaptability as a hybrid strategy. The wrapper allocates a modest state-of-charge reserve ex-ante within the CPR optimisation and governs its release through simple, causal trigger rules. When realised prices provide clear evidence of emerging scarcity, the wrapper authorises targeted discharge of this reserved energy. In extreme cases, particularly when prices approach the market cap (approximately VOLL⁴), the wrapper may override the CPR instruction and trigger immediate discharge, reallocating energy that CPR would otherwise schedule for later intervals. Such reallocation entails limited additional risk, since prices are bounded above by the cap, whereas the opportunity cost of failing to capture a scarcity event can be substantial.

The next subsection formalises the design of this layer, detailing the reserve allocation rule, the trigger conditions for release, and the override mechanism for cap-proximate intervals.

5.2 Design Overview (Operator Logic)

The hybrid framework preserves CPR as the primary day-ahead schedule and introduces two lightweight decision rules that act only when realised prices provide clear causal evidence of emerging scarcity. These rules operate without re-optimising the CPR plan and intervene only during intervals in which CPR remains idle, while preserving the single-cycle structure and avoiding any reliance on future price information.

(i) Reserved capacity

A small fraction of energy is reserved ex-ante by elevating the state-of-charge floor from its operational minimum (either S_{\min} or the longevity threshold $S_{\text{longevity}}$) to:

$$S_{\text{floor}} = S_{\text{longevity}} + S_{\text{reserve}}, \quad (16)$$

where S_{reserve} typically ranges between 5–15% of total capacity. This reserve is excluded from CPR's optimisation, ensuring that a buffer of energy remains available for unanticipated scarcity events.

(ii) Targeted release triggers

Two causal release mechanisms govern when and how the reserve is deployed:

- **Cap-shock rule:** If two of the last three realised 5-minute prices exceed $0.85 \times \text{VOLL}$, the wrapper authorises discharge in the next interval. Energy is released first from the reserved buffer; if the event warrants further action, additional energy may be reallocated from CPR's scheduled future discharges, with the corresponding amount debited from later intervals to maintain energy balance and respect technical limits.
- **Late-window insurance:** From a predefined late-day gate (e.g., interval 220), if two of the last three realised prices exceed a high-price floor defined as:

$$\max \{95\text{th percentile of CPR-adjusted discharge prices in the late window, } \$3000/\text{MWh}\},$$

The wrapper releases energy only during intervals in which CPR remains idle. This mechanism avoids double-selling, focuses reserve usage on genuine scarcity indicators, and prevents premature deployment earlier in the day.

Under this design, CPR continues to determine the core charging and discharging schedule, including all risk-aware threshold logic and the one-cycle-per-day constraint. The hybrid wrapper is deliberately minimal and targeted. It intervenes only when strong real-time evidence indicates that scarcity is unfolding, and every action at interval $t+1$ is based solely on information available up to interval t . When reserve energy is partially or fully released, the state-of-charge floor relaxes dynamically to reflect the reduced reserve (see Appendix D.1 and D.2 for details). Outside these trigger conditions, the CPR schedule remains fully intact.

5.3 Case Studies: 1 February (cap cluster) and 2 July (multi-spike day)

5.3.1 1 February — cap-cluster with strong evidence

On 1 February, realised prices enter a sustained near-cap band during the late-afternoon window (Figure 2). The CPR base plan remains conservative on this day because the magnitude and persistence of the spike are not fully anticipated in the day-ahead pre-dispatch signals. As a result, CPR alone under-commits to the extreme event.

The wrapper detects a valid 2-of-3 cap-proximity trigger and authorises a targeted intervention. It releases the full reserved energy (≈ 2.25 MWh) and initiates an additional discharge burst of ≈ 9.06 MWh by reallocating energy that CPR would

⁴ Value of Lost Load) denotes the market price cap, representing the maximum willingness to pay to avoid involuntary load shedding.

otherwise schedule for later intervals. All reallocated energy is subsequently back-debited from future CPR discharges to maintain energy balance and respect the one-cycle constraint.

This intervention produces a clear upward shift in daily revenue relative to the base CPR outcome (Figure 6). Because overrides occur only under near-cap conditions, where upside potential is bounded by VOLL, the wrapper introduces no additional downside risk. Quantitatively, the cap-proximity trigger reallocates energy into genuine scarcity intervals, lifting revenue by approximately \$102 k relative to the CPR baseline while maintaining feasibility through the dynamic SoC floor and back-debiting mechanism. Full revenue and risk metrics for this event are reported in Table 4.

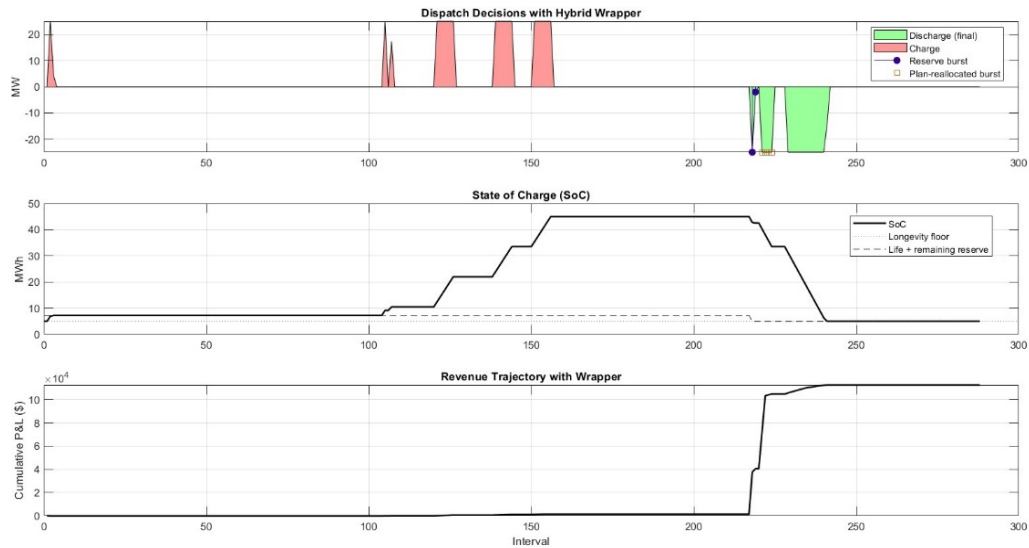


Figure 6: Dispatch, state-of-charge, and revenue trajectories with the hybrid strategy on 1 February

The top panel shows charge and discharge actions under the hybrid strategy, including reserve and plan-reallocated bursts triggered by the cap-proximity rule. The middle panel plots the state-of-charge trajectory with the longevity (life) floor and remaining reserve. The bottom panel shows the cumulative revenue path, highlighting the step increase generated by wrapper interventions during the near-cap event.

5.3.2 2 July — Multi-Spike with Idle Opportunities

Figure 7 illustrates battery behaviour under the hybrid strategy on 2 July 2025, a hyper-volatile day characterised by multiple scarcity clusters during the mid-day and late-evening periods (see Figure 5). The CPR base plan discharged the main early spike cluster but left several late-day intervals idle despite market prices remaining well above the risk-adjusted discharge threshold.

The wrapper intervenes only during CPR-idle intervals, releasing brief reserve bursts (≈ 2.08 MWh) while leaving the main CPR discharge block intact. With a 5% reserve fraction, these targeted releases lift total revenue to \$102,487, an increase of approximately \$20,610 ($\approx 25\%$) relative to the CPR baseline. All interventions remain compliant with the single-cycle constraint and respect the dynamic SoC floor. No plan reallocation under the cap-shock rule occurs because, by the time cap events emerge, no CPR-scheduled energy remains available to borrow.




Figure 7: Dispatch, state-of-charge, and revenue trajectories with the hybrid strategy on 2 July (Multi-Spike)

Because wrapper effectiveness depends on the amount of energy reserved ex-ante, we conduct sensitivity tests for reserve fractions of 10% and 15% (see Appendix D.3). Increasing the reserve from 5% to 10% materially improves performance by allowing larger targeted releases during high-price idle intervals, leading to a substantial revenue uplift relative to the CPR baseline and an improvement in risk-adjusted returns.

By contrast, further increasing the reserve to 15% yields diminishing returns. Although more energy is technically available for release, the larger reserve increasingly constrains the underlying CPR schedule by limiting normal discharge opportunities, thereby reducing total revenue. This pattern highlights a clear trade-off between preserving base-plan arbitrage performance and allocating sufficient reserve capacity for opportunistic intervention. Quantitative results for all reserve settings are summarised in Table 4.

To examine the robustness of the hybrid strategy, we also apply the hybrid strategy to 12 June, another cap-event day. Although the baseline CPR strategy already performs strongly on this day, a meaningful gap remains between the risk-aware and perfect-foresight outcomes, leaving scope for improvement. In this case, the wrapper delivers incremental revenue gains and improves downside risk metrics. Detailed results are reported in Appendix D.4.

Table 4: Quantitative results for hybrid strategy performance under different reserve fractions (1 February and 2 July 2025)

| Day | Reserve fraction | Reserve used (MWh) | Plan reallocated (MWh) | CPR total (\$) | CPR+Wrapper (Hybrid) total (\$) | Δ Revenue (\$) | Mean (\$/int) | Sortino |
|-------|------------------|--------------------|------------------------|----------------|---------------------------------|-----------------------|---------------|---------|
| 1-Feb | 5% | 2.25 | 9.06 | 10,700.19 | 112,693.44 | 101,993.26 | 391.3 | 47.287 |
| 2-Jul | 5% | 2.08 | 0 | 81,877.52 | 102,487.10 | 20,609.58 | 355.86 | 0.526 |
| 2-Jul | 10% | 4.14 | 0 | 83,815.78 | 122,677.92 | 38,862.14 | 425.97 | 0.63 |
| 2-Jul | 15% | 6.21 | 0 | 42,426.74 | 101,766.16 | 59,339.41 | 353.35 | 0.523 |

6 Conclusion

The Conditional Price-at-Risk (CPR) framework developed in this study provides a simple and distribution-informed approach for incorporating forecast-error tail risk into battery arbitrage decisions. By adjusting pre-dispatch price signals using empirical forecast-error distributions, CPR improves dispatch timing and moderates downside exposure relative to raw pre-dispatch strategies. Across the benchmark set, CPR consistently outperforms unadjusted pre-dispatch under calm to moderately volatile conditions, demonstrating the value of embedding interval-level risk awareness into day-ahead battery scheduling.

However, the results also show that CPR alone is insufficient under the most challenging market regimes, particularly during extreme and short-lived price spikes that are poorly signalled in pre-dispatch forecasts, yet contribute a large fraction of total annual arbitrage revenues despite occurring infrequently. To address these limitations, the study introduces a hybrid strategy that augments CPR with a lightweight wrapper incorporating two causal mechanisms: a reserve-release rule for idle high-price intervals and a cap-shock override for sustained cap-proximate events. Case studies demonstrate that maintaining a modest state-of-charge reserve of 5 to 10 percent enables these mechanisms to operate effectively, delivering substantial incremental gains on both cap-cluster and multi-spike days without violating operational constraints.

Aggregating results across the twelve benchmark days, the unadjusted pre-dispatch strategy captures approximately 55 percent of perfect-foresight arbitrage value. Incorporating risk awareness through CPR and its hybrid extension raises this capture rate to around 60 to 61 percent, while also reducing average CVaR_s losses by approximately 3 percent. From an operator perspective, this reduction reflects a meaningful moderation of severe downside outcomes, improving revenue robustness without sacrificing tractability or operational realism. Importantly, these gains are achieved using only day-ahead pre-dispatch forecasts and historical forecast-error information available ex-ante, while preserving a single-cycle, day-ahead optimisation structure.

While the proposed hybrid framework enhances robustness under forecast uncertainty, an important avenue for future research is to integrate CPR-based risk adjustments into rolling or multi-stage optimization settings. Such extensions would allow battery operators to update decisions dynamically as new information becomes available, bringing the framework closer to real-world bidding and dispatch practices while retaining its risk-aware foundations.

7 Acknowledgements

The author used the ChatGPT language model to assist with checking and correcting language during manuscript preparation. All output was reviewed and edited by the authors.

8 References:

- Weber, J., & Lu, B. An Open-Source Energy Arbitrage Model Involving Price Bands for Risk Hedging with Imperfect Price Signals. *Energies*. 2024;17, 13.
- Prakash, M., Bruce, A., MacGill, I. The scheduling role of future pricing information in electricity markets with rising deployments of energy storage: An Australian National Electricity Market case study. *Energy Economics*. 2025; 142, 108191.
- University of Queensland. UQ Tesla Battery Q1 2020 operational report. Brisbane: The University of Queensland. 2020.
- Pham, S. D., Do, H. X., Nepal, R., & Jamasb, T. Tail risk connectedness in the Australian National Electricity Markets: The impact of rare events. *Energy Economics*. 2025; 141, 108123.
- Jorion, P. *Value at Risk: The New Benchmark for Managing Financial Risk*. Vol. 3, McGraw-Hill, New York. 2007.
- Billimoria, F., & Simshauser, P. Contract design for storage in hybrid electricity markets. *Joule* 7. 2023; 1663–1678.
- Sheikhahmadi, F., Kazemi, A., & Ghasemi, H. A risk-based approach for modeling the strategic behavior of a distribution company in wholesale energy market. *Applied Energy*. 2018; 214, 24–38.
- Janczura, J., & Wójcik, R. Dynamic short-term risk management strategies for the choice of electricity market based on probabilistic forecasts of profit and risk measures. The German and the Polish market case study. *Energy Economics*. 2022; 110, 106015.
- Li, Y., Huang, J., Liu, Y., Zhao, T., Zhou, Y., Zhao, Y., & Yuen, C. Day-ahead risk averse market clearing considering demand response with data-driven load uncertainty representation: A Singapore electricity market study. *Energy*. 2022; 254, 123923.
- Barquín, J., Centeno, E., and Nogales, F.J. Short-Term Risk Management for Electricity Markets: A CVaR-Based Approach. *IET Generation, Transmission & Distribution*. 2017; 11(17): 4150–4157.
- Sajeev, S. P., Roth, L., Gangwar, T., Adu, J. A., Sakowski, I. A., and Monti, A. Risk-Aware Stochastic Optimization for Battery Energy Arbitrage Dispatch Planning. *IEEE PowerTech*. 2025; Kiel, Germany.
- Chai, S., Li, Q., Abedin, M. Z., & Lucey, B. M. Forecasting electricity prices from the state-of-the-art modeling technology and the price determinant perspectives. *Research in International Business and Finance*. 2024; 67, 102132.
- Rockafellar, R. T., & Uryasev, S. Optimization of conditional value-at-risk. *Journal of Risk*. 2000; 2(3), 21–42.
- Artzner, P., Delbaen, F., Eber, J.M., and Heath, D. Coherent Measures of Risk. *Mathematical Finance*. 1999; 9(3): 203–228.
- Bishop, C. M. *Pattern recognition and machine learning*. Springer. 2006.
- Weron, R. Electricity price forecasting: A review of the state-of-the-art. *International Journal of Forecasting*. 2014; 30(4), 1030–1081.
- Australian Energy Market Operator. *Quarterly Energy Dynamics: Q1 2024*. Australian Energy Market Operator, Melbourne; 2024.
- Australian Energy Regulator. *Wholesale Electricity Market Performance and Price Events*. Australian Energy Regulator, Melbourne; 2023.

Appendix A

12 February 2025:

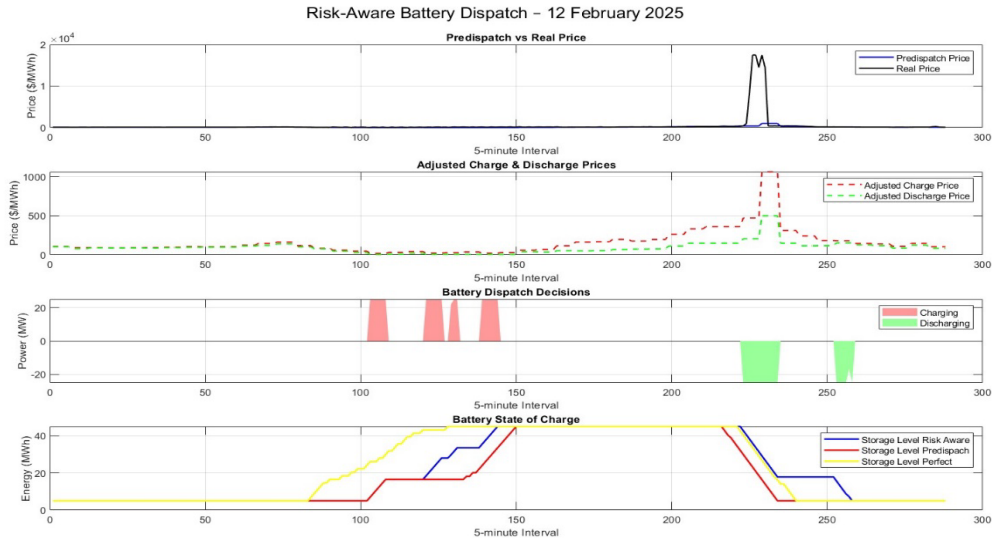


Figure A.1: Electricity price trajectories and battery operational responses under three strategies on 12 February 2025

On this day real prices remain moderate through the morning/early afternoon before ramping into a short cap-level burst (max \$17,500/MWh, min \$20.27/MWh). Dispersion is extreme ($\sigma \approx \$2,170/\text{MWh}$; range $\approx \$17,480/\text{MWh}$), while the day-average is \$467.78/MWh. Predispatch correctly signalled a sharp evening event in roughly the right window, aligning strategies toward the spike.

Perfect foresight charges to full ahead of the ramp and empties directly through the cap window, yielding the highest revenue (\$196,025.94) and the strongest risk profile (VaR₅ -97.12; CVaR₅ -144.37; Sortino 19.50).

Risk-aware charges steadily through ~100–150, reaching full SoC by ~150. Discharge is split across two peaks around the spike, with a small reserve for later intervals. While this slightly reduces peak revenue (\$190,322.47) relative to Predispatch, its tighter downside profile comes from more selective charging, yielding less-negative VaR₅ (-197.53 vs -251.74), less-negative CVaR₅ (-290.71 vs -317.19), and a higher Sortino (9.53 vs 8.77).

Predispatch holds slightly more energy into the steepest part of the first crest and this yields slightly higher revenue (\$192,208.03).

On this type of volatile day with a forecasted evening cap spike that predispatch signalled in the correct window, both strategies captured the event; predispatch achieved slightly higher revenue, while Risk-aware delivered tighter downside tails and better risk-adjusted returns, and Perfect foresight remained the upper bound.

07 March 2025:

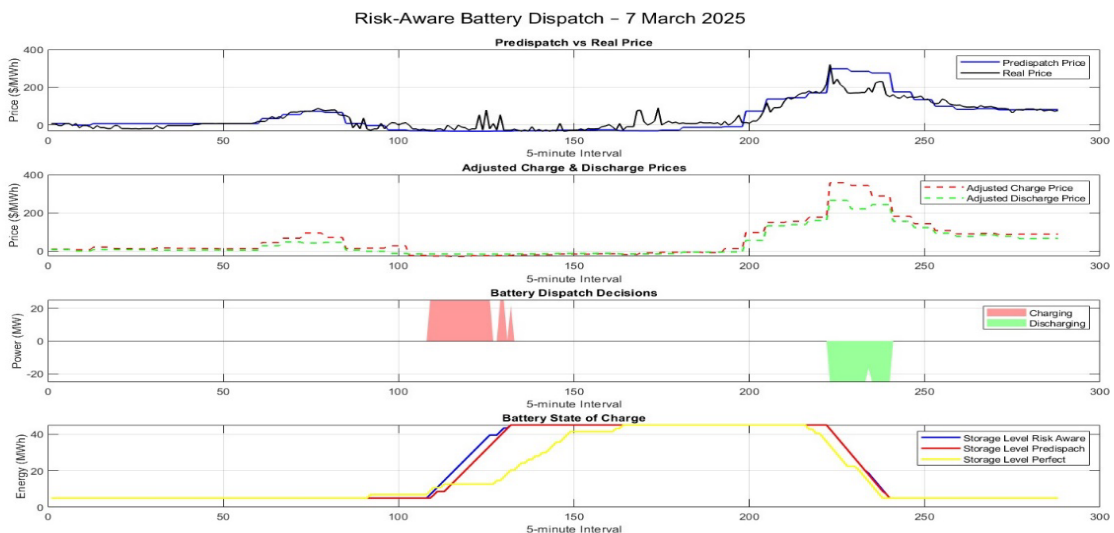


Figure A.2: Electricity price trajectories and battery operational responses under three strategies on 7 March 2025

This is a low-to-moderate volatility day with a single late ramp. Prices remain near zero-\$80/MWh for most of the day, dip briefly into shallow negatives (min -\$32.01/MWh), then climb into an evening plateau (max \$320.70/MWh). Dispersion is

modest compared to spike days ($\sigma \approx \$69.45/\text{MWh}$; range $\$352.71/\text{MWh}$). The figure shows midday charging followed by coordinated evening discharge during the ~220–245 interval ramp/plateau.

Perfect foresight times discharge tightly around the top of the ramp, yielding the benchmark revenue of $\$8,841.27$ and the strongest risk profile (Mean P&L 30.70; $\text{VaR}_5 = 0$; $\text{CVaR}_5 = 0$; Sortino very large).

Risk-aware charges through the midday trough and reaches full SoC slightly earlier than Predispatch. Discharge is paced more evenly across the high-price band, giving the higher revenue than forecast-based strategy ($\$7,669.59$) and a marginally better risk-return mix (Mean P&L 26.63; $\text{VaR}_5 = 0$; $\text{CVaR}_5 = -1.58$; Sortino 1.95).

Predispatch follows a very similar schedule but is marginally less effective on the plateau, with revenue of $\$7,635.45$, slightly weaker downside severity ($\text{CVaR}_5 -1.62$), and a lower Sortino (1.94).

On this type of calm, non-cap day with a clear evening ramp that predispatch and real prices tracked closely, both forecast-driven strategies performed similarly, with CPR adjustments edging ahead by slightly higher revenue and tighter downside risk, while Perfect foresight remained the upper bound.

10 April 2025:

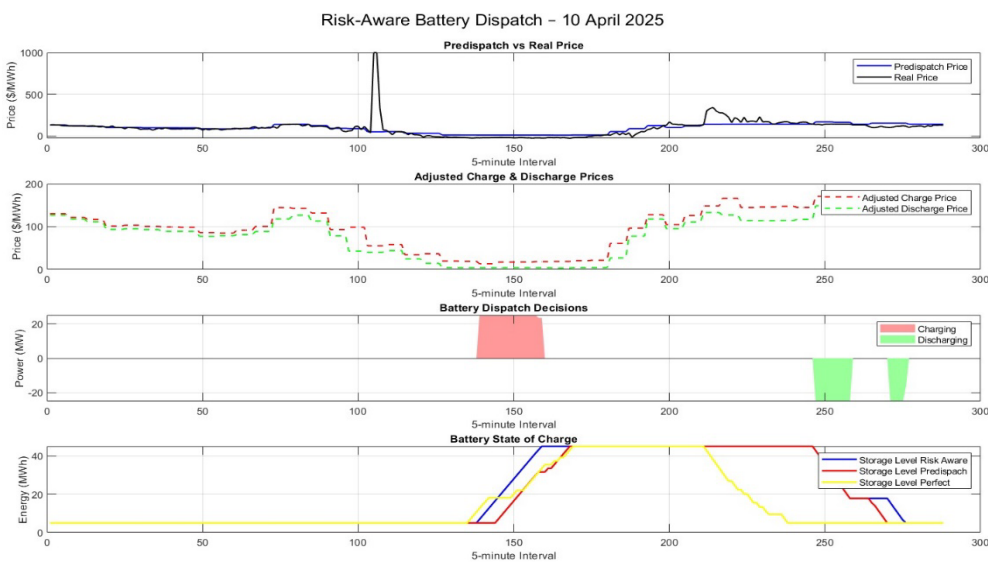


Figure A.3: Electricity price trajectories and battery operational responses under three strategies on 10 April 2025

This is a mostly calm day with one isolated $\$1,000/\text{MWh}$ spike around midday, followed by a broad evening rise. Dispersion is moderate ($\sigma \approx \$104.79/\text{MWh}$; range $\approx \$1,030.44/\text{MWh}$), with max $\$1,000/\text{MWh}$, min $-\$30.44/\text{MWh}$, and a day-average of $\$91.37/\text{MWh}$.

Perfect foresight ignores the midday spike due to the one-cycle constraint, charging afterwards to discharge fully during the evening plateau. This yields the highest revenue ($\$9,734.50$) with essentially no downside risk ($\text{VaR}_5 = 0$; $\text{CVaR}_5 \approx 0$; Sortino $\rightarrow \infty$).

Risk-aware charges through the afternoon trough and discharges smoothly across the evening peak, producing slightly higher revenue than Predispatch ($\$5,782.02$ vs $\$5,698.20$) while keeping tails benign ($\text{VaR}_5 = 0$; $\text{CVaR}_5 = 0$).

Predispatch follows a very similar schedule but is marginally less effective on the plateau, resulting in slightly lower revenue ($\$5,698.20$).

On this type of calm day with a brief unexploited midday spike and a broad evening rise that predispatch and real prices tracked closely, both forecast-driven strategies performed similarly. Risk-aware provided a small but consistent revenue edge with negligible downside, while Perfect foresight remained the upper bound.

20 May 2025:

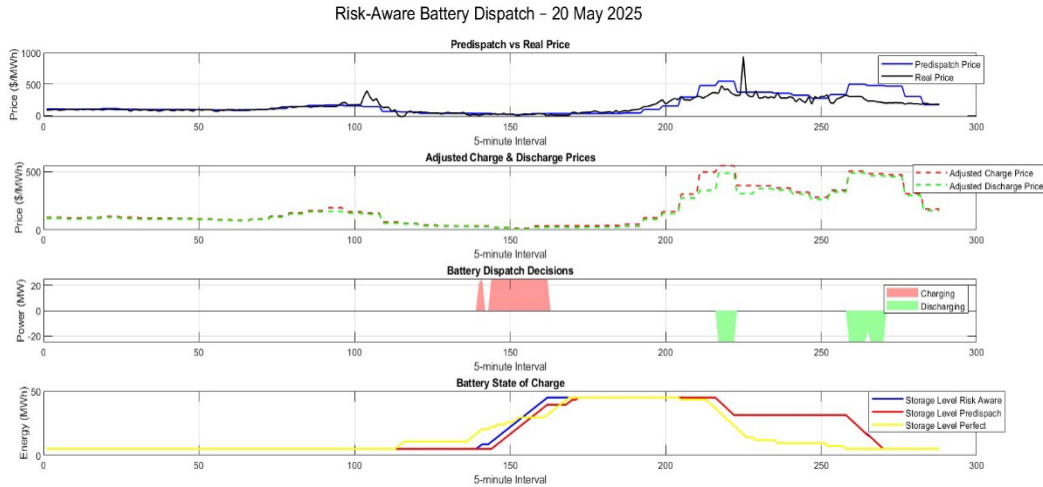


Figure A.4: Electricity price trajectories and battery operational responses under three strategies on 20 May 2025

The day remains relatively quiet through midday, with real prices mostly below approximately \$150/MWh. From around interval 200, a clear evening ramp develops. Real prices exhibit only a brief spike near interval 225, with most value concentrated in the sustained ramp between approximately intervals 200 and 260. In contrast, the Predispach forecast projects two evening plateaus: one broadly aligned with this period at around intervals 210–225, and a second, misplaced plateau at approximately intervals 260–280.

Perfect foresight charges earlier (~120–170) and discharges in one concentrated block across the ramp, capturing the climb. Revenue: ~\$14,378 (highest).

Risk-aware charges slightly later and a touch cheaper (~140–165) than Predispach, then discharges over a window broadly similar to Predispach (two short bursts around ~240–270). The later, lower-cost charging reduces downside exposure more than any difference in discharge timing. Revenue: ~\$10,457.

Predispach charges a bit later (~145–170) at slightly higher prices and begins evening release marginally earlier (~235), consistent with the over-forecast plateau. Revenue: ~\$10,349.

Risk metrics: Relative to Predispach, Risk-aware shows milder downside tails— VaR_5 -23.10 vs -26.25 ; CVaR_5 -44.40 vs -51.66 ; Sortino 3.49 vs 2.95 —driven primarily by the lower average buy cost from the earlier charging window. Perfect foresight shows small tails and the highest Sortino (16.19). This is a moderate volatility day, and CPR make only marginal adjustments. The improvement arises mainly from charging a little earlier/cheaper, not from materially different discharge timing.

14 June 2025:

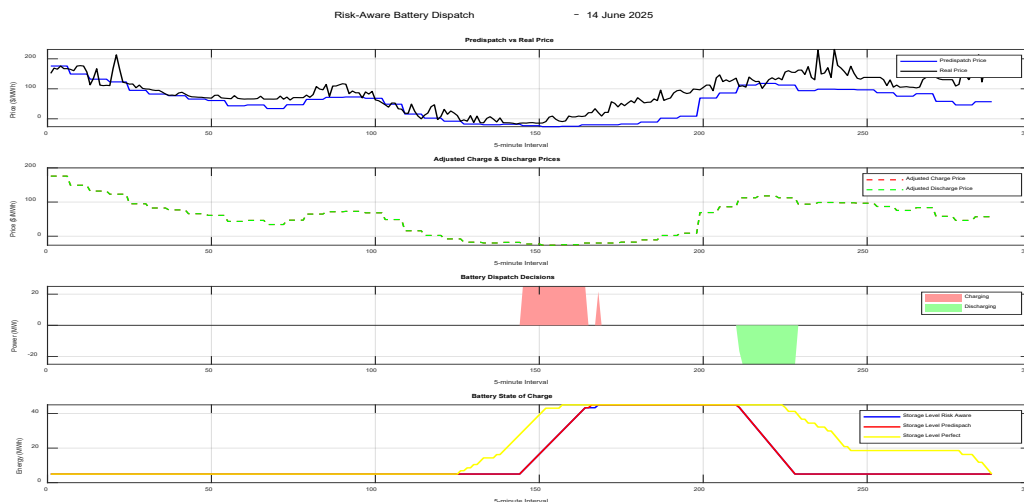


Figure A.5: Electricity price trajectories and battery operational responses under three strategies on 14 June 2025

The price dynamics on 14 June 2025 are calm compared to the other benchmark days. Real and predispach prices track each other closely, with only small deviations. No major spikes or crashes occur, and predispach volatility for the day is relatively low (52) compared with the average historical interval volatility (403).

Because of this, the model's safeguard condition is triggered:

$$\text{Std (Predispatch Prices)} < 100 \text{ and mean (Historical Volatility)} > 200$$

Under this rule, the risk-aware strategy reverts to the unadjusted predispatch trajectory to avoid introducing unnecessary skewness. Consequently, the adjusted and predispatch strategies follow identical price signals. Battery scheduling reflects the limited arbitrage opportunities. The perfect foresight strategy begins charging slightly earlier, taking advantage of lower-than-forecast real prices before interval 150. It then delays discharging to capture a modest late-day lift after interval 225. By contrast, both the risk-aware and predispatch strategies charge and discharge more conservatively, without reacting to this deviation.

Daily revenues amount to \$7,079 for perfect foresight and \$4,923 for both predispatch and risk-aware strategies. Risk metrics confirm the relatively calm nature of this case: downside CVaR is minimal (around -0.65), while Sortino ratios exceed 4.8, indicating very low downside volatility relative to mean returns.

Overall, this case demonstrates the intended robustness of the adjustment logic: when markets are stable, the risk-aware model defaults to the baseline predispatch strategy, avoiding incorrect adjustments that could mislead operational decisions.

26 August 2025:

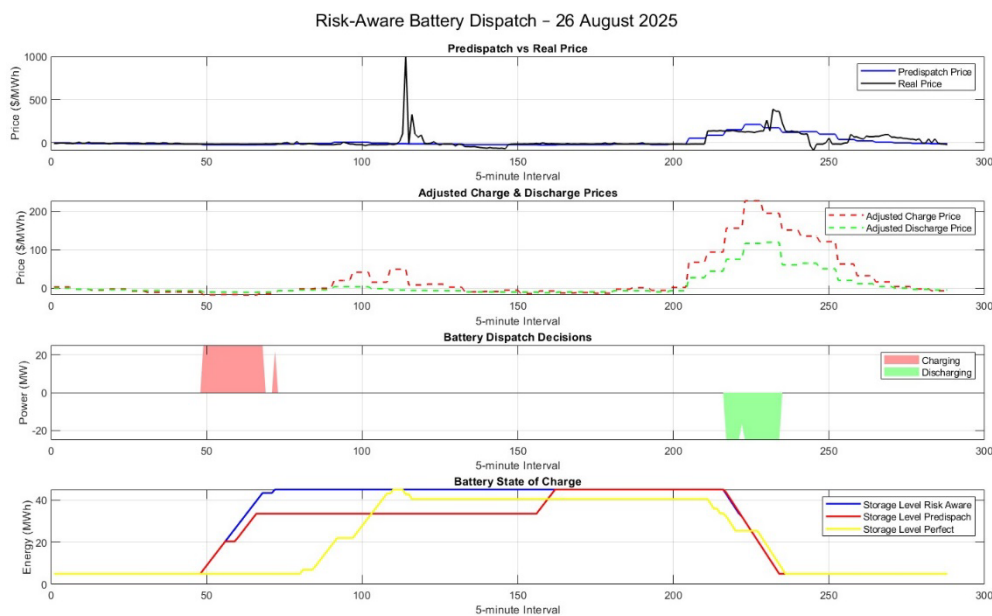


Figure A.6: Electricity price trajectories and battery operational responses under three strategies on 26 August 2025

A long negative trough dominates the morning, followed by a very short real-price spike around midday (~10–15 minutes) that collapses quickly, then a genuine evening ramp into the \$230–390/MWh range. Predispatch does not forecast the midday spike and modestly understates part of the evening lift.

Perfect foresight charges into the deepest negative window (~90–120), but also benefits from the brief midday spike by discharging partially and then emptying through the evening ramp. Highest revenue: \$10,115.90.

Risk-aware charges earlier (~50–70) while prices are already negative, avoids topping up around ~160 when prices are higher, and discharges smoothly across ~220–240. Revenue: \$7,255.20.

Predispatch charges early and then adds an additional charging block around interval ~160 in response to a mild dip in the forecast prices (not in realised prices). It begins discharging in a similar late window, though slightly earlier on average, yielding a total revenue of \$7,199.10.

Risk metrics show negligible downside tails, so the Sortino ratios appear artificially large. Differences are therefore driven by revenue, not risk dispersion.

On this type of “negative-to-evening-ramp” day, where Predispatch missed the short midday spike and understated part of the evening lift, the Risk-aware strategy avoided unnecessary top-ups and held energy for the genuine ramp. The gain over Predispatch is modest since both discharged in nearly the same high-price band, while Perfect foresight gained extra revenue by also exploiting the brief midday spike.

15 September 2025:




Figure A.7: Electricity price trajectories and battery operational responses under three strategies on 15 September 2025

The day begins with a long negative-price trough (many intervals below $-\$20/\text{MWh}$, including several deep dips), followed by a gradual recovery and a short late-afternoon spike band ($\sim 235\text{--}260$) where real prices climb sharply into the hundreds of $\$/\text{MWh}$. Outside this window, prices remain mild.

Perfect foresight charges fully during the deep negative trough in the morning and discharged squarely in the late spike band, producing the highest revenue. Revenue: $\$10,629$.

Risk-aware builds its state of charge in the early afternoon ($\sim 140\text{--}170$) and discharged across $\sim 220\text{--}240$, arriving with a moderate SOC and finishing slightly before the spike ended. Revenue: $\$3,401$.

Predispatch follows a nearly identical trajectory, beginning to charge a bit earlier and also completing discharge slightly earlier in the spike window. Revenue: $\$3,388$.

Downside risk is negligible ($\text{VaR}_5 \approx 0$, $\text{CVaR}_5 \approx 0$), making Sortino ratios extremely large and non-informative. The $\$13$ revenue edge of Risk-aware over Predispatch stemmed mainly from charging behaviour.

On this benign-risk day with a late spike that predispatch forecast only partially, both forecast-based strategies behave very similarly. CPR adjustments give the Risk-aware strategy a slight revenue edge by avoiding higher-priced early charging, while Perfect foresight gains a clear advantage by fully exploiting the deep negatives and discharging directly into the spike band.

Appendix B: Sensitivity of CPR Outcomes to the Tail Parameter α (2.5%, 5%, 10%)

The parameter α determines the tail probability used to construct Conditional Price-at-Risk (CPR) adjustments from the empirical distribution of forecast errors. Smaller values of α (e.g., 2.5%) focus on rarer and more extreme forecast misses, while larger values (e.g., 10%) reflect less severe but more frequent deviations. These quantile-based adjustments are applied subject to the model's conditional rules and are further moderated by the volatility-based scaling factor $\gamma(t)$.

In practice, changes in α do not necessarily translate into different dispatch outcomes. When the empirical forecast-error distribution is relatively flat between the 2.5% and 5% quantiles, or when state-of-charge (SoC) and power constraints are binding, the effective CPR thresholds remain unchanged. In such cases, reducing α from 5% to 2.5% produces identical price adjustments and therefore identical dispatch trajectories.

Table B.1 summarises α -sensitivity for two representative extreme days. On 2 July 2025, characterised by multiple spike clusters and a late cap event, increasing α from 5% to 10% improves revenue by approximately \$8,000. The looser tail results in milder CPR adjustments, reducing premature discharge into the early spike block and preserving additional energy for the late-day cap. By contrast, lowering α from 5% to 2.5% has no effect on dispatch or revenue, as the lower-tail quantiles coincide once scaled and constrained.

Table B.1. Sensitivity of Risk-Aware Revenues and Risk Metrics to α

| Date | α (tail prob.) | Revenue (\$) | Mean P&L | VaR ₅ | CVaR ₅ | Sortino | Notes |
|---------------------------------------|-----------------------|--------------|----------|------------------|-------------------|---------|--|
| 2 Jul 2025 (extreme spikes) | 2.50% | 93,295.86 | 323.944 | -1,157.4 | -2,279.4 | 0.47 | ≈ identical to 5% (flat tail, binding SoC) |
| | 5% (baseline) | 93,295.90 | 323.944 | -1,157.4 | -2,279.4 | 0.47 | Reference run |
| | 10% | 101,499.67 | 352.43 | -1,157.4 | -2,279.4 | 0.51 | Later discharge preserved → +\$8.2k |
| 1 Feb 2025 (mixed, evening spikes) | 2.50% | 11,305.60 | 39.256 | 0 | -0.065 | 41.475 | ≈ identical to 5% (tails negligible) |
| | 5% (baseline) | 11,305.60 | 39.256 | 0 | -0.065 | 41.475 | Reference run |
| | 10% | 11,464.56 | 39.807 | 0 | -0.065 | 41.475 | Slight +\$159 gain (threshold shift) |

Figure B.1 illustrates battery behaviour on 2 July under $\alpha = 10\%$, compared with the baseline case of $\alpha = 5\%$ shown in Figure 5. With $\alpha = 10\%$, the CPR adjustments are less extreme, and the discharge threshold is not driven as far downward during the initial spike sequence. As a result, the strategy discharges less into the early, ultimately sub-dominant spike block and retains more energy for the terminal cap window. Revenue increases to \$101,499.67, although the battery still reaches the final spike with a thinner SoC than under perfect foresight (VaR₅ = -\$1,157.40; CVaR₅ = -\$2,279.40; Sortino = 0.51).

Overall, this sensitivity analysis confirms that α primarily affects outcomes under extreme volatility regimes with multiple competing spike opportunities. Under calm to moderately volatile conditions, and even on some extreme days, dispatch behaviour is dominated by SoC limits and price levels rather than by the precise choice of tail probability. These findings support the use of $\alpha = 5\%$ as a stable default in the main analysis, while motivating targeted extensions such as the hybrid strategy for handling the most challenging spike-driven environments.

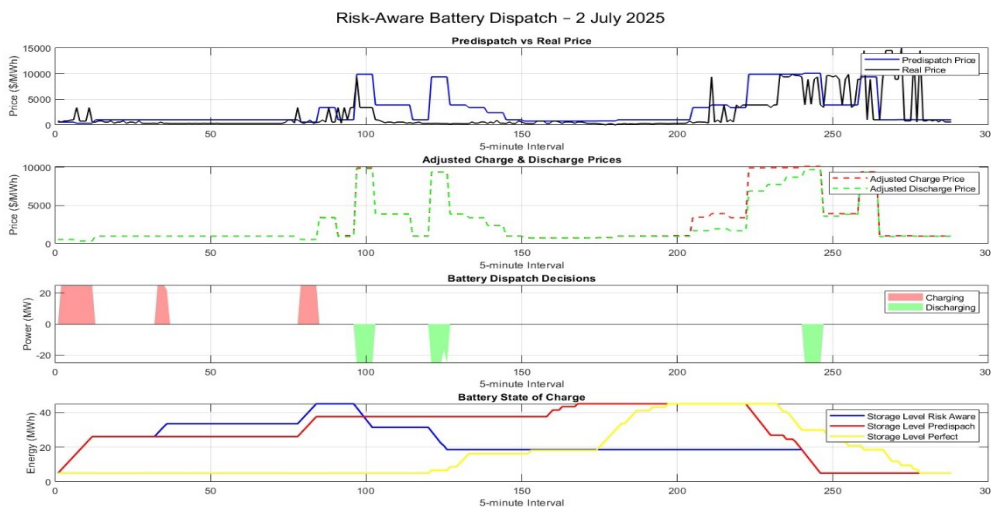


Figure B.1: Electricity price trajectories and battery operational responses under three strategies on 2 July 2025 under $\alpha=10\%$

Appendix C. Calibration of the Volatility Scaling Parameter γ_0

The CPR adjustment incorporates interval-level volatility through a sigmoid scaling factor $\gamma(t)$, which maps historical forecast-error volatility into bounded adjustments of the charge and discharge thresholds. The base parameter γ_0 controls the maximum intensity of this scaling and therefore governs how aggressively volatility expands the CPR adjustment margins. In the main analysis, γ_0 is set to 0.2. This appendix explains the rationale for this choice and demonstrates why larger values do not yield systematic performance improvements.

To assess the behaviour of the scaling mechanism, we re-ran the model under an alternative setting of $\gamma_0 = 0.4$ for a small set of representative days, including two extreme events (1 February and 2 July) and two moderately volatile days (26 January and 20 May). These tests are not intended as a full sensitivity analysis, but rather to evaluate whether a more aggressive volatility response materially improves dispatch outcomes.

On the cap-cluster day (1 February), increasing γ_0 produced almost identical dispatch paths and raised CPR revenue by approximately 1%. This limited response reflects the fact that once prices enter the extreme tail, dispatch decisions are dominated by spike timing rather than by incremental threshold tightening. In such regimes, the sigmoid scaling has little leverage, and higher γ_0 values do not meaningfully alter the effective operating plan. On the multi-spike day (2 July), raising γ_0 to 0.4 reduced total revenue by approximately 6% and did not improve downside risk metrics. The stronger scaling led to over-tightened thresholds during non-critical intervals, prompting earlier energy release and reducing flexibility ahead of the terminal spike.

For moderately volatile days, CPR outperformed raw pre-dispatch under both γ_0 settings. However, $\gamma_0=0.2$ consistently delivered slightly higher revenues and marginally stronger risk-adjusted performance, with Sortino ratio improvements of approximately 1% on 26 January and 0.7% on 20 May relative to $\gamma_0 = 0.4$.

Smaller values of γ_0 were also considered but found to generate only negligible adjustments, effectively collapsing the CPR thresholds back toward the pre-dispatch baseline and offering limited incremental risk mitigation.

Taken together, these results support the choice of $\gamma_0 = 0.2$ as a balanced calibration. This value provides sufficient responsiveness to volatility to correct systematic forecast bias, while avoiding over-tightening of thresholds in periods where price risk does not warrant aggressive adjustment. Higher values offer no systematic robustness benefits and can reduce revenues on complex extreme days. Accordingly, $\gamma_0 = 0.2$ is adopted as the default throughout the paper.

Appendix D: Wrapper Implementation and Supporting Analysis

D.1. Dynamic SoC Floor Formulation

The wrapper dynamically modifies the lower storage bound to preserve both battery longevity (life) and the reserved tail capacity. At each interval t , the active SoC floor is defined as

$$S_{\text{floor}}(t) = S_{\text{longevity}} + S_{\text{reserve}}(t), \quad (\text{D.1})$$

where $S_{\text{reserve}}(t)$ denotes the remaining unutilised reserve capacity. Initially,

$$S_{\text{reserve}}(0) = \rho S_{\text{max}}, \quad (\text{D.2})$$

with ρ representing the reserve fraction (e.g., 5–15%). Following a reserve release of magnitude $\Delta S(t)$, the reserve evolves according to

$$S_{\text{reserve}}(t+1) = \max(0, S_{\text{reserve}}(t) - \Delta S(t)). \quad (\text{D.3})$$

Accordingly, the storage floor relaxes gradually as reserve energy is consumed, while never falling below the fixed longevity limit $S_{\text{longevity}}$.

D.2. Cap-Shock and Late-Window Trigger Logic

Two causal release rules govern the wrapper's interventions: a cap-shock trigger and a late-window trigger. Both rules rely exclusively on realised prices and operate without re-optimising the underlying CPR schedule.

Cap-shock rule: A cap-shock event is detected if at least two of the last three realised 5-minute prices exceed a fixed fraction of the market price cap:

$$\frac{1}{3} \sum_{k=t-2}^t \mathbf{1}(\pi_k \geq 0.85 \pi_{\text{cap}}) > \frac{2}{3}, \quad (\text{D.4})$$

where π_k denotes the realised 5-minute price and π_{cap} is the NEM market price cap.

If the condition is satisfied at interval t , the wrapper authorises a discharge from the remaining reserve in the subsequent interval $t+1$. If the reserve is fully exhausted and cap-proximate conditions persist, additional energy may be borrowed from CPR's scheduled future discharges. Such actions are recorded as plan-reallocated bursts and are back-debited from the corresponding future CPR intervals to preserve energy balance and maintain the single-cycle constraint.

Late-window rule: From a predefined late-day gate t_{gate} (e.g., interval 220), the wrapper may release reserve energy only during intervals in which the CPR plan is idle and realised prices exceed a high-price threshold. Specifically, a release is authorised if at least two of the last three realised prices satisfy:

$$\pi_k > \max\{\text{95th percentile of CPR-adjusted discharge prices in the late window}, \$3000/\text{MWh}\}.$$

This condition focuses interventions on genuine scarcity episodes while avoiding interference with CPR's planned discharge blocks.

In both cases, when a trigger condition is met at interval t , the corresponding discharge action is executed at interval $t+1$, ensuring strict causality and avoiding any use of future price information. If neither trigger fully depletes the reserved energy, any unused reserve at the end of the day is valued at the final 5-minute spot price and added as a terminal revenue term. This ensures that reserved energy is not treated as lost when it is not physically dispatched.

D.3. Sensitivity and Performance Diagnostics

Figures D1 and D2 illustrate how different reserve fractions alter the interaction between the wrapper and the underlying CPR schedule on 2 July. These figures complement the quantitative results reported in Table 4 by highlighting the operational mechanisms through which reserve sizing affects performance.

With a 10% reserve (Figure D1), the elevated SoC floor remains binding for longer into the day, preventing early depletion while still allowing the CPR baseline to execute its main discharge block. The wrapper releases multiple short bursts during CPR-idle intervals late in the day, coinciding with sustained high prices. As a result, additional value is extracted without materially constraining the baseline CPR trajectory. By contrast, under a 15% reserve (Figure D2), although more reserve energy is technically available, the tighter floor reduces flexibility in the CPR baseline, limiting normal discharge opportunities during other intervals.

Taken together, the figures illustrate a clear mechanistic trade-off: increasing the reserve expands the capacity for opportunistic releases, but beyond a moderate level it increasingly interferes with the base CPR plan. Performance therefore peaks at intermediate reserve levels, where idle high-price intervals can be exploited without undermining the baseline discharge structure.

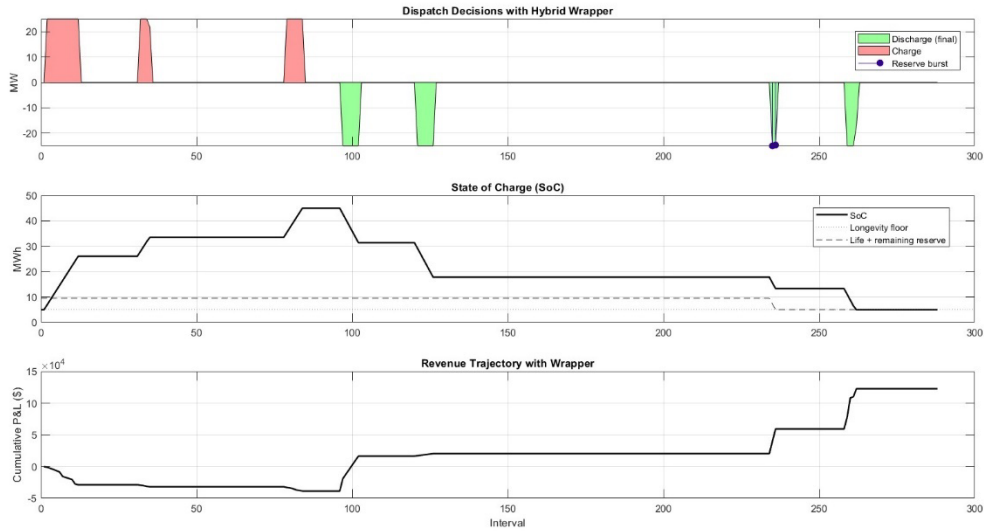


Figure D1. Dispatch outcomes under the hybrid strategy on 2 July 2025 with a 10% reserve fraction

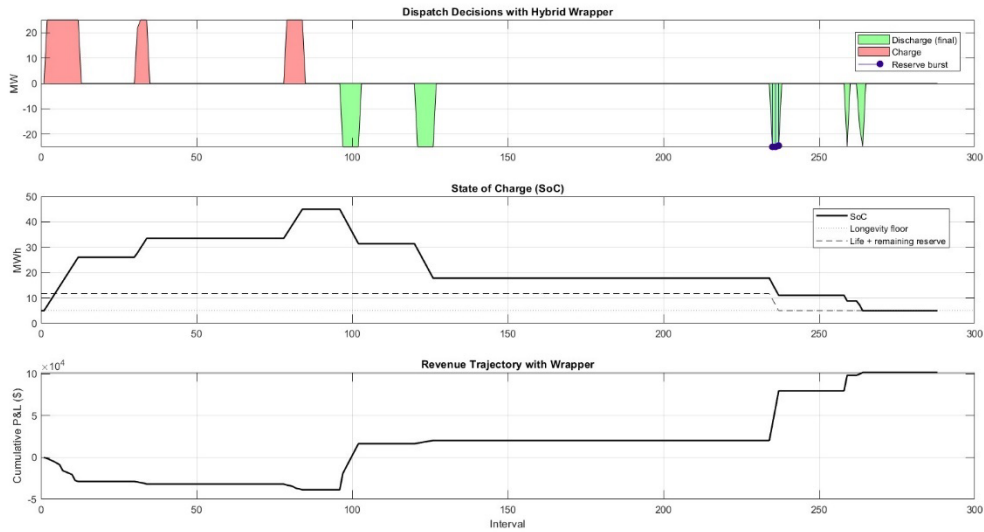


Figure D2. Dispatch outcomes under the hybrid strategy on 2 July 2025 with a 15% reserve fraction.

D.4 – 12 June 2025: Cap-shock burst after full discharge

We include 12 June to examine the wrapper on a day where the baseline risk-aware CPR strategy already performs strongly, earning about 6% more revenue than pre-dispatch (\$362.7k vs \$342.0k) and capturing roughly 81% of perfect-foresight value (\$362.7k vs \$445.2k). Despite this solid baseline, the presence of a short cap-proximate episode makes the day suitable for testing whether the wrapper can add incremental value without degrading downside risk.

CPR with 5% reserve: 12 June is a late cap-band day, with prices tracking a high plateau during intervals 230–245. Under the CPR plan with a 5% reserve (≈ 2.25 MWh), the battery charges to about 90% by $t \approx 167$ and then executes a long, planned discharge through the evening ramp (Figure D3). As reported in Table D1, CPR with reserve delivers \$342,332 in daily revenue, a Sortino ratio of 13.21, and tight tail metrics ($Var_5 = -302$; $CVaR_5 = -374$).

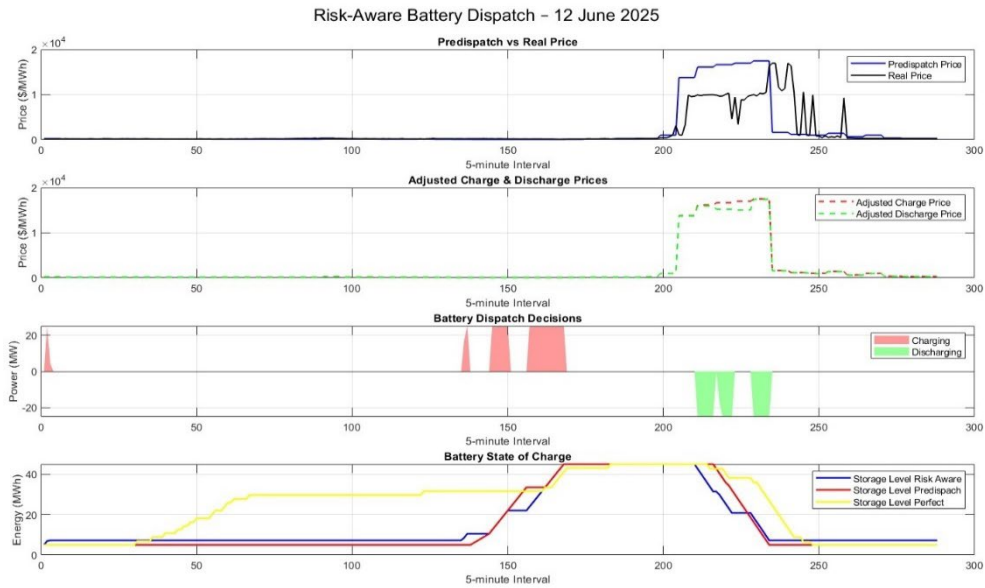


Figure D3: Electricity price trajectories and battery operational responses under three strategies on 12 June 2025 under 5% reserve capacity

Wrapper outcomes: A single, well-timed wrapper intervention and Discharge releases the reserved energy, increasing daily revenue to \$377,446 and improving the Sortino ratio to 14.57. The state-of-charge floor (life plus remaining reserve) relaxes correctly after the burst, and terminal value is zero because the reserve is fully utilised. Relative to CPR without reserve, the hybrid (CPR+Wrapper) strategy lifts revenue by about \$14.8k (+4.1%) while modestly improving downside risk metrics (VaRs: $-324.5 \rightarrow -302.5$; CVaRs: $-384.8 \rightarrow -374.0$).

Wrapper behaviour: Because most energy has already been discharged when the cap band emerges, no idle CPR intervals exist above the late-window threshold (q_{95} , late $\approx \$17,451/\text{MWh}$), so the late-window rule does not activate. Instead, a brief cap-shock episode is detected when two of the last three realised prices exceed $0.85 \times \text{MPC}$ ($\$14,875/\text{MWh}$) around $t = 233-237$. The wrapper releases the remaining reserve at $t = 236$ (25 MW for one 5-minute interval, ≈ 2.08 MWh). No plan reallocation occurs, as no CPR-scheduled energy remains to borrow. The cap condition dissipates immediately thereafter, allowing only a single burst (Figure D4).

Table D1 summarises performance across CPR without reserve, CPR with a 5% reserve, and the hybrid strategy.

Table D1: Performance on 12 June: CPR vs CPR (5% Reserve) vs Hybrid strategy ($\alpha=5\%$)

| Strategy | Day total (\$) | Mean (\$/int) | VaR_5 | CVaR_5 | Sortino |
|------------------|----------------|---------------|----------|----------|---------|
| CPR (no reserve) | 362,671.80 | 1,259.28 | -324.490 | -384.798 | 13.13 |
| CPR (5% reserve) | 342,332.49 | 1,188.65 | -302.49 | -374.01 | 13.21 |
| Hybrid | 377,446.48 | 1,310.58 | -302.50 | -374.02 | 14.57 |




Figure D4. Hybrid strategy dispatch on 12 June 2025 under 5% reserve fraction

Contact

Parna Imannezhad

Centre for Applied Energy Economics & Policy Research (CAEEPR)

Griffith University, Queensland, Australia

Email: parna.imannezhad@griffithuni.edu.au

