Bounds on the Arbitrage Value of Faster Battery Dispatch

Conrad Voigt
UC Santa Barbara

November 2nd, 2026

Abstract

How much can a battery gain from resetting output every five minutes instead of every fifteen? We compare optimal long-run arbitrage revenue on nested dispatch calendars with common underlying prices and public information. Under single-rate exponential mean reversion the ratio of revenues lies between a calendar multiplier and its square. Both endpoints are sharp in the Ornstein–Uhlenbeck model. For more general prices, we bound the effect of misspecification by the revenue a feasible battery can earn on two forecast discrepancies. A test based on expected price changes over fixed holding periods establishes a strictly positive premium in 118 of 126 finite-state specifications without solving the inventory control problem. Direct optimization finds a gain in all 126. An application to four months of 2025 data from New South Wales illustrates the cost of assuming a single rate. For one- and four-hour batteries, a rule using archived forecasts earned enough on the synthetic discrepancy to erase the benchmark's implied minimum gain in the sample. The market premium itself remains unidentified.

1 Introduction

A trader who buys the dispatch rights to a battery buys a sequence of future decisions. The power rating limits how many megawatts the trader can move, and the energy rating limits how long the trader can move them. A further contract term specifies how often an instruction can change. To value that term, consider two counterparties entitled to operate identical stores. Both see prices arrive at the same rate. One may reset its output on every five-minute boundary; the other must hold the chosen output for fifteen minutes. A price reversal in the intervening ten minutes can be valuable to the first counterparty. Its value depends on whether the battery has room to charge or energy to sell when the reversal occurs.

Comparing realized hourly and quarter-hourly price spreads does not isolate this right. Different products can have different liquidity, settlement dates and market participants. A historical optimizer can also use prices learned after a dispatch decision. Either feature changes the question. Here the physical price tape is common, revenue is settled on actual delivery, and the trading instruction must be chosen before the slot's average price is observed. We take a long-run objective so that an arbitrary terminal date or a free initial inventory has no recurring value.

Three results connect the price process to this contractual right. First, with single-rate mean reversion, a pair of policy constructions gives sharp multiplicative bounds at finite decision intervals. For five- and fifteen-minute dispatch and a 30-minute shock half-life, the model implies an 11.75–24.88% uplift at any positive capacity. At a two-hour half-life the range falls to 2.90–5.89%. The larger opportunity goes with more rapidly decaying information. The bounds address percentage

revenue. At the same power rating a larger battery can earn a larger dollar increment and still have a smaller percentage uplift.

Second, the policy constructions continue to apply when mean reversion depends on the price state. They identify two particular forecast discrepancies. The value a feasible battery can earn on these discrepancies gives a correction in money per unit time. This matters even for small perturbations. In a controlled three-price experiment, a modest asymmetry raises the short battery's uplift from 51.17% to 54.90%, above the upper bound in the unperturbed model. A bound that accounts for the asymmetry still certifies a gain.

Third, we give conditions for establishing a positive premium from conditional price changes without computing an optimal dispatch rule. When the slow calendar permits the entire store to refill in one slot, cancellation between positive and negative predicted price movements within a block has a direct dollar interpretation. For longer batteries, feasible policies that repeatedly open and close small positions give lower bounds on arbitrage revenue. Together with two upper bounds from conservation of energy, these produce a test across intermediate durations. It can return an inconclusive answer even when the true gain is positive. A controlled experiment measures that distinction. The empirical section asks how far a public decision-time forecast reduces the uncertainty visible in settlement prices alone.

1.1 Relation to previous work

Hodges derives a continuous-time valuation formula for storage under mean reversion and identifies a single price threshold for buying and selling [1]. Carmona and Ludkovski model storage as optimal switching with inventory and operating constraints [2]. Chen and Forsyth develop a semi-Lagrangian approach for a natural-gas storage problem [3]. Faghih, Roozbehani and Dahleh analyze the time-averaged value and price responsiveness of finite-capacity, ramp-constrained storage, with structural results under independent prices [4]. Our contribution concerns a comparison of two attainable stationary values when prices persist across decisions.

Discretizing the control of a diffusion also has an approximation theory. Jakobsen, Picarelli and Reisinger give convergence rates as the piecewise-constant control interval tends to zero [5]. The valuation question here compares two *fixed* intervals. The argument depends on energy conservation, power limits and how a block price is settled. It supplies constants that can be interpreted for a given price half-life, and capacity regimes in which the constants are attained.

Market evidence addresses related but distinct notions of speed. Metz and Saraiva consider arbitrage in German 15- and 60-minute intraday markets [6]. Soini, Parra and Patel disaggregate storage operations by timescale [7]. Schaurecker et al. model high-frequency trading for intraday battery portfolios [8]. Fast trading in a delivery contract can revise a position without changing the physical cadence of dispatch; changing a battery's setpoint can have value even in a price-taking model without an order book. In Australia, the move to five-minute settlement aligned settlement with an existing five-minute dispatch process [13]. This institutional change does not identify our counterfactual premium.

The robustness analysis is related to stability of multistage stochastic optimization under changes in the underlying process [10]. Mukherjee and Panda study decision regret after projection of battery price forecasts and identify the role of energy conservation [11]. Le Roux Tardif relates forecast information to the value of a storage bidding decision under throughput constraints [12]. Our comparison holds the information set fixed and changes the feasible physical decisions. Its correction depends on a forecast discrepancy's direction, timing, and tradability by the asset.

2 Prices, storage, and settlement

A fine slot lasts h hours. Index consecutive fine slots by integers t, and write Y_t for the realized average price during slot t. The trader has information $\mathcal{F}_t$ at the start of that slot and forms the conditional expectation $\mathbb{E}_t$. Prices and information are exogenous and jointly stationary under fine-slot shifts. Assume integrability of all price signals. A coarse slot lasts $H = qh$ for integer $q \geq 2$ and settles at $q^{-1}\sum_{j=0}^{q-1} Y_{t+j}$. The two calendars differ in when a physical action can be taken. At a coarse decision date the trader can use the intervening fine-grid public history. With less frequent observation as well as less frequent control, an additional filtering problem would arise.

A lossless store has usable energy capacity $E > 0$, common charge and discharge power $r > 0$, and duration $D = E/r$ hours. At each decision choose an energy increment b, positive for charging. Starting from energy Q, feasibility on a slot of length δ requires

$$|b| \leq r\delta, \qquad 0 \leq Q + b \leq E. \tag{1}$$

The rate b/δ is held until the next decision. Storage is price-taking. There are no degradation charges, losses, fees or other revenue streams. Splitting the store into virtual accounts later in the proof relies on this linear, symmetric feasible set. Fixed losses or limits shared with other services would require a new comparison.

For an integrable signal z_t known at the decision, define

$$J_\delta(z) = \sup_b -\frac{1}{\delta}\mathbb{E}[b_t z_t].$$

The supremum is over stationary admissible policies on the specified calendar. Policies may use the original public information and independent private randomization. Joint stationarity of inventory and decisions is imposed on their decision grid. Equivalently for comparisons one may draw the calendar phase uniformly before operating: the price model assigns the same ex ante value to each phase. Write

$$f_t = \mathbb{E}_t Y_t, \qquad c_t = \frac{1}{q}\sum_{j=0}^{q-1} \mathbb{E}_t Y_{t+j}, \qquad d_t = \frac{1}{q}\sum_{j=0}^{q-1} \mathbb{E}_t c_{t+j}. \tag{2}$$

The realized slot-average cash flow is $-b$ times the realized settlement. Since b is chosen in advance, its expected rate equals the corresponding expression using the conditional mean. The values of interest are $V_h = J_h(f)$ and $V_H = J_H(c)$. The premium is $V_h - V_H$ in currency per hour or $V_h/V_H - 1$ when the denominator is positive. A constant price shift leaves these values unchanged, since $\mathbb{E}b = 0$. Taking opposite inventories $Q \mapsto E - Q$ reverses every trade, so $J_\delta(z) = J_\delta(-z)$. Both properties concern stationary policies and the symmetric store.

3 Policy transfer and a sharp mean-reversion benchmark

Proposition 3.1 (Calendar transfers)**.** *With the common information and settlement conventions above,*

$$V_h \geq V_H, \qquad V_h \geq J_H(f), \qquad V_H \geq J_h(d).$$

Proof. To reproduce a coarse policy on the fine grid, repeat its power setting for q slots. All intermediate inventories lie on the segment between the feasible endpoints. Randomizing the initial grid phase gives a stationary fine policy.

For the second inequality, take any coarse policy and split the battery into q virtual accounts, each with E/q energy capacity. Operate a scaled copy of the policy in each account on a different coarse phase. On each fine slot execute the entire energy decision of the one active account, scaled by $1/q$. This respects the common physical rating because $rH/q = rh$. Energy feasibility holds account by account. Every fine slot settles its active account's trade against the first-slot price. The stationary expected rate is $-\mathbb{E}[bf]/H$ for the unscaled coarse decision b. The coarse policy used for this construction can be selected to maximize this objective.

For the last inequality, start with a stationary fine policy. At coarse date t choose the sum of the last q fine decisions, including the one just selected at t. This is known then. If Q_s is the virtual fine inventory immediately before decision s, the aggregate moves the physical store from Q_{t-q+1} to Q_{t+1} over the next coarse slot. Both endpoints are feasible and the aggregate size is at most qrh. A two-sided stationary virtual history gives the initial target; bounded initialization effects disappear in long-run implementations. By stationarity and the tower property the resulting rate is

$$-\frac{1}{H}\sum_{j=0}^{q-1}\mathbb{E}[b_{t-j}c_t] = -\frac{1}{h}\mathbb{E}[b_t d_t].$$

Taking the supremum over fine policies proves the result. □

To specialize, suppose the centered spot price is a stationary Ornstein–Uhlenbeck diffusion,

$$dX_s = -\kappa X_s\,ds + \sqrt{2\kappa}\,\sigma\,dW_s, \qquad \kappa, \sigma > 0.$$

Here s denotes continuous time and the average over each physical slot gives Y_t. The trader observes the current spot and its past, and has no advance information about future innovations. Set

$$m(\delta) = \frac{1 - e^{-\kappa\delta}}{\kappa\delta}, \qquad M = \frac{m(h)}{m(H)} = \frac{q}{\sum_{j=0}^{q-1} e^{-\kappa jh}}.$$

The conditional block mean is $m(\delta)X_{th}$, so $f = Mc = M^2 d$. A single exponential conditional mean suffices for these identities. The calendar transfers apply without a distributional assumption on the innovations.

Theorem 3.2 (Sharp finite-calendar comparison)**.** *For every positive capacity and rating in the stationary OU model,*

$$M \le V_h/V_H \le M^2.$$

The upper endpoint is attained for $0 < D \le h$, and the lower endpoint is approached as D tends to infinity at fixed r.

Proof. Substitute the conditional-mean identities into Proposition 3.1 and use positive homogeneity of J. For the small-capacity endpoint, when $D \le \delta$ every end inventory is reachable in one decision. Summation by parts gives

$$V_\delta = \frac{\kappa E\sigma}{\sqrt{2\pi}}\,m(\delta)^2.$$

Indeed the optimal endpoint is E if $X_{th} < 0$ and zero if $X_{th} > 0$. The expected gain per unit of energy and per slot is the expected positive conditional change in the block quote, $m(\delta)(1 - e^{-\kappa\delta})\sigma/\sqrt{2\pi}$. Division by δ gives the expression. For large capacity the attainable limit, justified in Section 5, is $rm(\delta)\mathbb{E}|X|$. Its ratio is M. □

An economic distinction follows from the treatment of the two endpoints. A small store can cross its entire inventory range in either slot, and faster control captures more reversals *and* settles on a shorter block. A large store rarely has to defer a trade because of an energy boundary. The shorter block settlement still matters even there. These endpoint statements do not imply monotonicity in D at every intermediate capacity. The five-minute uplift relative to fifteen minutes for a one-hour battery and a four-hour battery is investigated in the numerical section.

4 Allowing for incorrect conditional prices

Choose a nominal scalar $M \geq 1$. For general stationary prices define the forecast discrepancies

$$\ell = f - Mc, \qquad u = f - M^2 d.$$

The first compares the next fine-slot forecast with the current coarse-block forecast. The second compares it with the average of forecasts for staggered coarse blocks. The latter average must be formed using information actually available at the original decision. It arises because aggregating past fine instructions delays their settlement.

Proposition 4.1 (Additive error allowances)**.** *The original stationary values satisfy*

$$\max\{V_H, MV_H - J_H(-\ell)\} \ \leq V_h \ \leq M^2 V_H + J_h(u).$$

Proof. The functional J_δ is subadditive because it is a supremum of linear expected cash flows. In particular $MV_H = J_H(f - \ell) \leq J_H(f) + J_H(-\ell)$. By Proposition 3.1, the first term on the right is attainable on the fine calendar. Similarly, $J_h(f) \leq M^2 J_h(d) + J_h(u)$ and $J_h(d) \leq V_H$. □

This result turns a forecast discrepancy into the value of trading *that discrepancy* with the original asset. A store that remains full cannot repeatedly buy a persistent negative error. A discrepancy that changes direction around dispatch dates can matter even when it has a smaller mean absolute size. A bound in dollars also requires a scale for the underlying arbitrage value. Multiplying an exact OU price by a constant preserves its zero discrepancies, yet scales the dollar premium.

The additive qualification cannot be removed uniformly. Let three fine-slot prices $(-\varepsilon, 0, \varepsilon)$ repeat with a uniformly random initial phase, observable from past settlements. Every three-slot mean is zero. The coarse store earns zero. The fine store can buy at $-\varepsilon$ and sell at ε once per cycle. Its gain is strictly positive for every $\varepsilon > 0$, although the absolute discrepancies tend to zero with ε. The example concerns control under common public information and a stationary price law. It shows why an out-of-sample forecast error, measured only in currency per MWh, cannot by itself establish a percentage premium.

5 How battery duration changes the comparison

Let z_n be an integrable, stationary quote observed at date $n\delta$. The expected reward from charging b_n units is $-b_n z_n$. Write $J_\delta(z)$ for optimal stationary revenue per unit time, so that $V_h = J_h(f)$ and $V_H = J_H(c)$. On each calendar the quote is evaluated at that calendar's decision dates. With $\mathbb{E}_n$ denoting expectation given the information available at $n\delta$, define

$$\mathcal{D}_{k,\delta}(z) = \mathbb{E}\big|\mathbb{E}_n z_{n+k} - z_n\big|, \qquad A(z) = \inf_{a \in \mathbb{R}} \mathbb{E}|z_n - a|.$$

The first quantity measures predictable changes in the quote over k decisions. The second measures its dispersion around a median.

Proposition 5.1 (Duration-dependent bounds)**.** *For every $E, r, \delta > 0$,*

$$B_\delta(z) := \sup_{k\geq 1,\, k\in\mathbb{N}} \frac{\min(r, E/(k\delta))}{2}\, \mathcal{D}_{k,\delta}(z) \;\leq\; J_\delta(z), \tag{3}$$

$$J_\delta(z) \leq U_\delta(z) := \min\left\{ \min(r, E/\delta) A(z), \frac{E}{2\delta}\mathcal{D}_{1,\delta}(z) \right\}. \tag{4}$$

If $D = E/r \leq \delta$, then $J_\delta(z) = E\mathcal{D}_{1,\delta}(z)/(2\delta)$.

Proof. Fix a holding period k and an energy amount $e = \min(r\delta, E/k)$. Set $s_n = 1$ if $\mathbb{E}_n z_{n+k} > z_n$, and $s_n = 0$ otherwise. At date n, take in es_n units and release the es_{n-k} units reserved k decisions earlier. The net trade is at most e in absolute value. End inventory $e\sum_{j=0}^{k-1} s_{n-j}$ lies in $[0, ke]$. By stationarity the rule earns $e\,\mathbb{E}[s_n(\mathbb{E}_n z_{n+k} - z_n)]/\delta$. The forecast change has mean zero, so its expected positive part is half its expected absolute value. This gives (3).

Every stationary rule satisfies $\mathbb{E}b_n = 0$ and $|b_n| \leq \min(r\delta, E)$. Subtracting a median from z_n gives the first upper bound. For the second, shift the first term in $\mathbb{E}[(Q_n - Q_{n+1})z_n]$ forward by one date to obtain

$$-\mathbb{E}[b_n z_n] = \mathbb{E}[Q_{n+1}(\mathbb{E}_n z_{n+1} - z_n)].$$

End inventory is in $[0, E]$, and the mean forecast change is zero. If $E \leq r\delta$, the $k = 1$ rule chooses either endpoint from any starting inventory and attains this upper bound. □

The tranche rule gives a feasible coarse-calendar revenue even when a full inventory optimization is impractical. For example, if $B_H(c) > 0$, the calendar transfers and Proposition 5.1 give

$$\max\left\{1, \frac{B_H(f)}{U_H(c)}\right\} \;\leq\; \frac{V_h}{V_H} \;\leq\; \frac{U_h(f)}{B_H(c)}.$$

Here $B_H(f)$ applies the coarse-date rule to fine-slot quotes and implements its trades with the staggered accounts. In the residual comparison, the feasible error allowances $J_H(-\ell)$ and $J_h(u)$ can likewise be bounded by $U_H(-\ell)$ and $U_h(u)$. All quantities depend only on quote forecasts and the two physical limits. Intermediate durations can leave a substantial interval because a fixed holding period restricts the dispatch rule.

There is a direct interpretation for any stationary forecast process when $D \leq H$. Index fine slots by t, define $w_t = \mathbb{E}_t f_{t+q} - f_t$, and let J be independent and uniform on $\{0, \ldots, q-1\}$, unobserved at t. Physical block averaging gives $c_t = q^{-1}\sum_{j=0}^{q-1} \mathbb{E}_t f_{t+j}$. The tower property thus gives $\mathbb{E}_t c_{t+q} - c_t = \mathbb{E}_t w_{t+J}$. For a real number x, write $x^+ = \max(x, 0)$ and $x^- = \max(-x, 0)$. Taking the difference between $\mathbb{E}|w_t|$ and $\mathbb{E}|\mathbb{E}_t w_{t+J}|$ yields

$$J_H(f) - J_H(c) = \frac{E}{H}\,\mathbb{E}\left[\min\left\{\mathbb{E}_t w^+_{t+J},\; \mathbb{E}_t w^-_{t+J}\right\}\right]. \tag{5}$$

Indeed, stationarity gives $\mathbb{E}|w_{t+J}| = \mathbb{E}|w_t|$, and for nonnegative a, b one has $a + b - |a - b| = 2\min(a, b)$. Thus the best staggered-account revenue exceeds coarse revenue precisely when the upcoming fine slots have conditional mass on changes of both signs with positive probability. Other fine policies may earn a premium even when the right-hand side vanishes.

Corollary 5.2 (A strict gain under mixing)**.** *Suppose the observed state is a finite stationary Markov chain relative to the trader's full information and the fine quote f is a function of that state. Write P for its one-slot transition matrix and restrict the state space to the support of its*

stationary distribution π. *Set* $K = q^{-1}\sum_{j=0}^{q-1} P^j$, $c = Kf$, *and* $w = (P^q - I)f$. *When* $D \le H$ *and* $V_H > 0$,

$$\frac{V_h}{V_H} \;\ge\; \frac{\|w\|_{1,\pi}}{\|Kw\|_{1,\pi}} \;\ge\; \frac{1}{\tau(K^*)}, \qquad \tau(K^*) = \frac{1}{2}\max_{i,j}\sum_l |K^*_{il} - K^*_{jl}|,$$

where $K^*_{ij} = \pi_j K_{ji}/\pi_i$.

Proof. The staggered-account construction gives $V_h \ge J_H(f)$. Apply the exact short-duration formula to $J_H(f)$ and $J_H(c)$. Here the matrices $P^q - I$ and K commute. The signed measure $\nu_i = \pi_i w_i$ has mass zero, and $(\nu K^*)_j = \pi_j (Kw)_j$. Its positive and negative parts have equal mass. After normalization their images under K^* are mixtures of rows of K^*, whose total variation distance is at most $\tau(K^*)$. This proves the second inequality by the coefficient-of-ergodicity argument of [9]. □

The signal-specific bound uses how much block averaging attenuates the predictable change in the quote. A coefficient strictly below one guarantees a gain across all signals with positive coarse revenue. The coefficient can be close to one when states persist, even when the particular price signal gives a materially larger gain. If $D \le h$, both calendars allow a complete inventory adjustment, and the exact ratio is $q\,\mathcal{D}_{1,h}(f)/\mathcal{D}_{1,H}(c)$ whenever its denominator is positive.

Proposition 5.3 (Large energy capacity)**.** *At fixed* r, δ, *if* z *is stationary, ergodic and integrable on the decision grid, then* $\lim_{E\to\infty} J_\delta(z) = rA(z)$.

Proof. The power bound gives the upper limit. Choose a median a. Consider a desired trade $r\delta\, g(z_n)$, with $g = 1$ below a and $g = -1$ above a. On any atom at a, choose a common value in $[-1, 1]$ so that $\mathbb{E}g(z_n) = 0$. This trade earns $rA(z)$ per unit time at prices centered on a.

Fix $0 < \eta < 1$. Begin a cycle at half capacity and follow the desired trades for N dates. If a proposed displacement from the starting inventory exceeds $r\delta L$, freeze the position for the remainder of this stage, where $L = \lceil \eta N \rceil$. Use the next L trades to reset to half capacity in equal increments. Capacity $E \ge 2r\delta L$ suffices. Ergodicity gives $\max_{j\le N} |\sum_{i=1}^{j} g(z_i)|/N \to 0$ almost surely, so the probability of freezing tends to zero. The expected absolute first-stage loss per unit time on this event vanishes by truncating $|z_i - a|$ at a fixed level, using the event probability for the bounded part and stationarity and integrability for the tail. The reset stage costs at most $Lr\delta A(z)$ in expectation. A uniformly randomized cycle phase makes the construction stationary. Letting N grow and then η decrease to zero proves the claim. □

Thus, if both grids are ergodic and $A(c) > 0$, the large-duration ratio is $A(f)/A(c)$. Short batteries earn from predictable changes between successive decisions. A large store with fixed power can sustain trades near the power limit on either side of the median quote for long periods. These endpoints require different features of the price law, which explains why a single persistence estimate can give misleading calendar comparisons.

6 How much speed premium can be certified?

We first measure the premium in the single-rate benchmark and then hold the stationary price distribution fixed while changing its conditional dynamics. All values are expected profit per unit time at unit power. These calculations use the true transition law so that numerical errors and model errors can be examined separately.

Continuous prices. Table 1 uses a centered Ornstein–Uhlenbeck price with stationary standard deviation one. A block settles against the average price over its entire delivery interval. At a half-hour price half-life, faster dispatch adds more absolute revenue for the four-hour store than for the one-hour store, but the percentage uplift is smaller. Reporting the ratio alone would conceal this distinction. At a two-hour half-life the relative uplift is below four percent for both stores.

Table 1: Unit-variance OU prices. Rates are per hour. The absolute gain is $V_{5\text{min}} - V_{15\text{min}}$.

Half-life (h)	Duration (h)	$V_{15\text{min}}$	$V_{5\text{min}}$	Absolute gain	Uplift
0.5	1	0.297675	0.341679	0.044004	14.78%
0.5	4	0.555185	0.626820	0.071635	12.90%
2.0	1	0.116947	0.121255	0.004309	3.68%
2.0	4	0.347474	0.359116	0.011641	3.35%

Relative value iteration uses the exact Gaussian integral of a linear interpolant on 961 price nodes spanning six stationary standard deviations on either side of zero. Doubling the price resolution from 481 nodes changes each reported rate by less than 0.00026. Extending the boundary to eight standard deviations at the same spacing changes each rate by less than 4×10^{-9}. The finite-grid Bellman residual span, divided by block length, is below 2×10^{-8}. Grid comparisons measure numerical sensitivity and the residual bounds the error for the discretized model.

Changing price persistence without changing marginal prices. Let $g = (-1, 0, 1)$ and $\pi = (1/4, 1/2, 1/4)$. With a the nominal retention and u an asymmetry fraction, set $s_i = a - u(1-a)g_i$. From state i, hold with probability s_i and otherwise refresh from weights proportional to $\pi_j(1 - s_j)$. Detailed balance gives the same π for every parameter pair. At $u = 0$, $Pg = ag$. At $u > 0$, the negative state becomes more persistent. We use $a \in \{0.55, 0.80, 0.92\}$, $u \in \{0, 0.10, 0.25, 0.50, 0.75, 0.90, 0.95\}$, and $E/r \in \{1, 2, 3, 6, 12, 24\}$ fine slots. The trader sees the current state. The fine trade settles on the next realized tick, and the coarse trade settles on the mean of the next three. Thus $h = 1$, $H = 3$, $f = Pg$ and $c = (I + P + P^2)f/3$. The three-state Bellman benchmark enumerates integer inventory knots and all feasible destinations. These knots suffice because the piecewise-linear concave continuation retains integer breakpoints under integer charge limits and linear rewards. The residual span is below 10^{-10} in rate units.

A lower bound without solving the control problem. To evaluate the holding-period lower bound, we use offsets $1 \le k \le \max\{12, 2\lceil E/(rH)\rceil\}$ in B_H. The direct lower ratio is $S = B_H(f)/U_H(c)$. The scaled residual lower ratio is

$$G(m) = m - \frac{U_H(mc - f)}{B_H(c)}.$$

We compare $\max\{1, S, G(M)\}$ with $\max\{1, S, \sup_{m \ge 1} G(m)\}$, where $M = 3/(1 + a + a^2)$. The optimized version is inexpensive on a finite state space. Put $T = P^3 - I$. For a weighted median absolute deviation, breakpoints occur when two components of $mc - f$ cross. For the inventory bound, breakpoints occur when a component of $T(mc - f)$ vanishes. The objective is the larger of the two resulting piecewise-linear functions; its maximum is attained at a breakpoint or at $m = 1$. Neither function grows on its final ray, since both upper bounds on c dominate $B_H(c)$. The choice of scale thus requires only the transition probabilities and three price states.

The actual fine premium is positive in all 126 cases. If the nominal single-rate interval $[M, M^2]$ is applied throughout, the true ratio exceeds its upper end in 19 cases and falls below its lower end

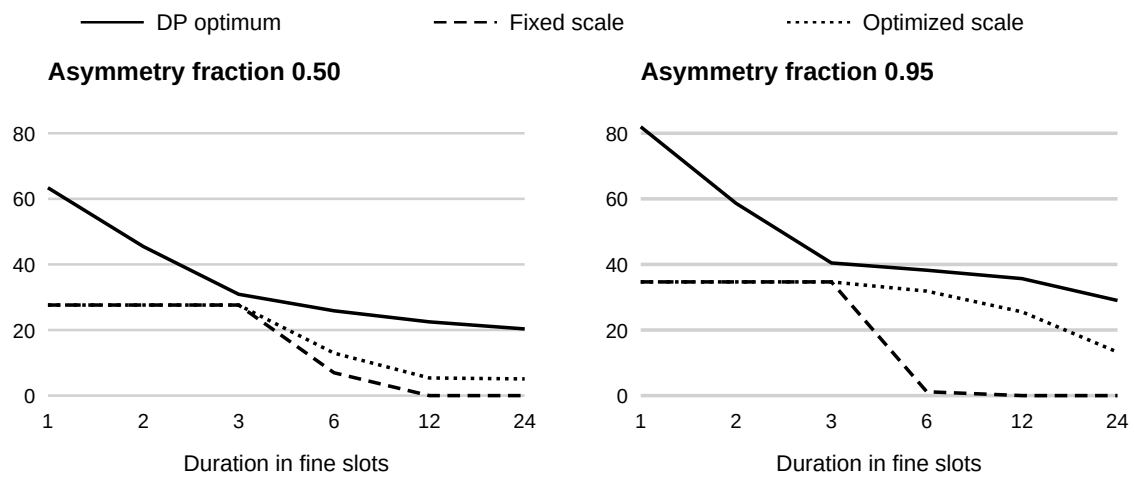

Figure 1: Percentage gain from switching to fine dispatch at $a = 0.80$. The solid line is the optimal gain. The other curves are certified lower gains. Duration labels are equally spaced categories.

in 17 cases. The fixed-scale lower certificate proves a strict gain in 109 cases; scale optimization raises the count to 118. Across the grid the median share of the actual premium certified by the optimized bound is 61 percent. This is a summary of the specified designs, with equal weight for each cell. It does not estimate a market frequency.

Figure 1 shows how the method can succeed or fail at the same nominal persistence. At $u = 0.50$ and duration three, a 27.62 percent premium is certified against a 30.90 percent optimum. At duration twelve, the true gain is 22.49 percent, the fixed-scale bound certifies only nesting, and scale optimization certifies 5.38 percent. At $u = 0.95$ and duration six the corresponding figures are 38.24, 1.17, and 31.85 percent. In the last example, fixing the nominal scale conceals most of the gain that the same conditional forecasts can establish.

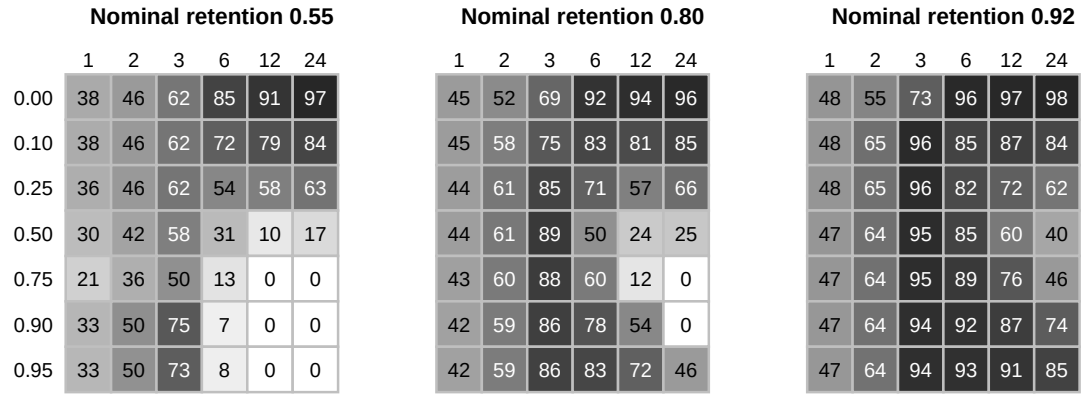

Figure 2: Percentage of the true speed premium certified by the optimized lower bound, rounded to the nearest integer. Darker cells certify a larger share. The eight cells where nesting alone survives are all at durations twelve or twenty-four.

For durations at most three, the coarse store can move its entire inventory in one block. The

one-block holding rule and the inventory upper bound coincide, and every one of these 63 cases certifies a strict gain. Equality of the coarse support bounds does not force equality with the full fine-control value. Conversely, for longer stores the certificates can leave substantial value unaccounted for: the mean conditional prices are known in this experiment, so this last gap comes from the policy bounds themselves.

7 What can be measured from settlement prices?

The preceding experiment specifies the transition law. A trader ordinarily observes neither the true conditional block price nor the optimal stationary revenue. There are two separate empirical questions. Can realized prices alone put a useful upper allowance on misspecification? Can information available before dispatch show that some of the discrepancy is predictable? We address the first with a full year of NSW prices and the second in the forecast experiment below.

7.1 A conditional-mean check

We use AEMO's five-minute regional reference prices (RRP) for New South Wales. The public monthly series contains 105,408 intervals in 2024 and 105,120 in 2025 [14]. Checks on region, interval-end timestamp, five-minute continuity and record type precede analysis. All negative prices and spikes are retained. The archived final RRP serves as a settlement proxy. A record of the final value does not in itself establish what a participant saw before a dispatch deadline.

A single exponential decay rate implies the ratio $b = (1 + a + a^2)/3$ between the expected average deviation over the next three slots and the expected deviation over the first slot, with $a = e^{-\kappa h}$. This is a restriction on conditional means. To test it on interval-average observations, estimate a from the ratio of the lag-two and lag-one autocovariances. For positive lag, integrating an exponential covariance over two nonoverlapping intervals preserves its geometric decay across lags. Dividing the lag-one covariance by the variance would treat an interval average as an instantaneous observation and needlessly change the implied multiplier.

As a diagnostic, estimate a mean profile in 2024 by calendar quarter and half-hour of day, then subtract the frozen profile from both years. The training estimate is $a = 0.94698$ and $b = 0.94792$. At aligned quarter-hour starts, group decisions by the preceding completed price deviation using 2024 cutoffs. The clock-aligned test omits the first decision of each day so all outcomes and the lagged signal lie within the same day. Bootstrap intervals below use 2,000 draws of circular seven-day blocks in the held-out year, conditional on the fitted profile and parameters.

Table 2: Frozen single-rate prediction in NSW, 2025. The error is the next three-slot average deviation less b times the first-slot deviation. Intervals are pointwise and expressed in AUD/MWh.

Previous-price group	Decisions	Mean error	95% interval
Bottom training decile	3,781	−16.27	[−31.07, −6.39]
Middle 80%	28,464	−0.32	[−1.87, 1.33]
90th–99th percentile	2,213	5.61	[−6.40, 17.32]
Upper training percentile	217	−42.91	[−436.23, 200.51]

In the low group the average first-slot deviation was −238.70 AUD/MWh and the three-slot average was −242.55. Weak prices persisted into delivery, while the fitted model forecast reversion. With a constant training mean the low-group error was smaller at −4.62 AUD/MWh, with

interval $[-7.98, -1.17]$. The discrepancy is relevant to charging, but its size alone is insufficient to determine a revenue correction. The other groups in Table 2 are compatible with zero mean error. In particular, the 217 observations in the highest training-price group give little precision. The tests neither establish the global form of the price process nor demonstrate that every single-rate calibration fails.

7.2 How conservative are realized-error allowances?

To value physical trading on raw prices, remove only the 2024 grand mean. The fitted multiplier is then $M = 1.052811$. This choice avoids treating a predictable intraday seasonal pattern as a cash flow that can be subtracted for free. A constant level is immaterial to a stationary round trip. For fine-slot prices Y_t define

$$C_t = \frac{Y_t + Y_{t+1} + Y_{t+2}}{3}, \qquad T_t = \frac{C_t + C_{t+1} + C_{t+2}}{3}, \qquad L_t = Y_t - MC_t, \quad U_t = Y_t - M^2 T_t.$$

The true discrepancies ℓ_t and u_t equal the conditional means of L_t and U_t. Every admissible policy earns the same *expected* cash flow on its conditional signal and corresponding realized outcome. A controller that observes the entire realized path can earn more. This gives conservative upper comparisons in expectation; their sample analogues are sensitive to innovation noise, extreme prices and the period observed.

We evaluated the full 2025 tape over 105,114 fine starts balanced across the three phase offsets. The hindsight optimizer respects the physical power and energy limits while allowing free initial inventory and flexible terminal inventory. For one MW and one hour of capacity, its coarse value, averaged over phases, was AUD 28.22 per operating hour. The implied nominal lower gain was at most 1.49 on that tape, whereas the hindsight correction on $-L$ reached 18.23. For four hours the corresponding amounts were 51.23, 2.71 and 19.14. Thus the simple observable upper allowance erases the positive lower premium even if the slow asset gets the benefit of hindsight. These are sample diagnostics of the procedure. They do not imply that the true premium is zero or that a causal policy could earn the hindsight correction.

A first causal check used only clock time and lagged realized prices. A rule fitted on 2024 prices assigned the next change in the artificial residual price to one of sixty bins and charged on a predicted rise. For the full 2025 four-hour store, delaying the lag by one additional five-minute slot gave AUD 0.90 per MW-hour, with a seven-day block-bootstrap interval of $[-0.49, 2.48]$. The other durations and latency choices also gave intervals containing zero. These exploratory 2025 NSW results motivated a test using archived forecast vintages. We distinguish this exploratory year from the regional transfer check that follows.

7.3 What did an archived system forecast add?

A public forecast permits a decision-time check of the lower-bound error allowance. For a quarter-hour starting at t, let $p_{t,1}, p_{t,2}, p_{t,3}$ denote its three settled five-minute prices. Define its realized synthetic payoff per MWh as

$$Z_t = M(p_{t,1} + p_{t,2} + p_{t,3})/3 - p_{t,1}, \qquad M = 1.0528110113383673.$$

We ask whether a simple storage policy can exploit changes in Z_t using a regional forecast available before dispatch. This payoff is constructed from settled prices and has no separately traded contract. Estimated revenues measure exposure to the discrepancy in the calendar comparison.

The protocol fixed January, April, July and October 2025 before opening the forecast archives. New South Wales was the exploratory region. Victoria was reserved for confirmation under the same policies on the same 123 dates. AEMO produces twelve projected five-minute regional prices every five minutes [15, 16]. At each decision we construct $\widehat{Z}_t$ and $\widehat{Z}_{t+15\,\mathrm{min}}$ from the same forecast vintage. The policy charges when the latter exceeds the former, discharges when it falls below the former, and holds on a tie, subject to inventory limits. The comparison policy predicts the same change from two-hour time-of-day groups and lagged price bins. We fit its bin cutoffs and conditional means to 2024 New South Wales prices and fix them for both test regions. Its latest settled input is the price of the interval ending five minutes before the decision.

Each forecast uses the latest admissible run with a nominal timestamp five, ten or fifteen minutes before dispatch and six distinct, finite prices. Each contributing record must have a last-change timestamp at least one minute before dispatch and an inactive intervention flag. The timestamp rule approximates availability from recorded changes to the source data. Publication and receipt lags at individual desks are unobserved. All 35,404 eligible decision starts per region across the three calendar phases had qualifying prices from the five-minute-lagged run. The data contained no duplicate regional run–delivery keys or flagged intervention rows in these months. We use AEMO's consolidated regional prices for scoring [14]. Revisions to the settlement series can also affect the lagged policy's historical input.

Both policies act at up to one MW, skip the first fifteen minutes of each month, and begin and end each month half full. We average daily cash over the three feasible quarter-hour calendar offsets and divide by the calendar hours in the sample. To estimate sampling variability we resample paired seven-day blocks within each month, circularly, preserving month lengths, and report the 2.5 and 97.5 percentiles over 5,000 draws.

Table 3: Synthetic residual revenue in AUD per rated MW per calendar hour. The interval applies to the paired difference.

Region	Duration (h)	Forecast	Lagged price	Difference	95% interval
New South Wales	1	2.79	0.00	2.79	$[1.52, 4.18]$
	4	3.33	1.14	2.20	$[1.10, 3.66]$
Victoria	1	0.84	0.10	0.75	$[0.57, 0.92]$
	4	0.79	0.05	0.73	$[0.53, 0.93]$

Table 3 shows positive paired increments in both regions. Each region and duration also had a positive increment in each of the four months. Calendar alignment mattered in Victoria: for the one-hour store the three phase-specific increments were 0.09, 0.20 and 1.94. For four hours they were 0.05, 0.13 and 2.02. The reported averages weight these alignments equally. The regions share the same market dates and may experience common shocks. The intervals describe variation within four preselected months. These comparisons show that public forecasts can make the synthetic discrepancy exploitable in this sample. Estimating the five-minute premium would require a separate comparison of optimal feasible physical dispatch policies. Assessing executable net revenue would further require a market implementation and its trading and operating costs.

8 Interpretation and limits

Two prices can produce similar unconditional spread distributions and assign different values to the right to change physical output. The relevant distinction is the timing of the predictable

price changes relative to the constraint on inventory. With exponential conditional means this dependence has an especially compact summary: a single number determined by half-life and the dispatch intervals bounds the value ratio, and the battery's energy capacity determines which endpoint it can approach. Beyond that model, the conditional changes in the fine and coarse quotes give sufficient tests that can be evaluated before solving the control problem.

The empirical results also distinguish the two directions of evidence about the error allowance. Hindsight gives an upper comparison that may be too loose to protect the nominal gain. A forecast-based feasible policy works in the opposite direction. It demonstrates that part of the error could be exploited under the specified timing assumptions. Better forecasting can therefore *weaken* confidence in a single-rate valuation by revealing more of its predictable misspecification. The two approaches answer different questions.

To place the archive replay on a common scale we additionally computed endpoint-matched hindsight optima on those same months, regions and phase offsets. Both artificial and ordinary price paths start and end at half capacity each month. Table 4 compares the estimated nominal margin with the correction. This is a descriptive comparison made after the regional replay. The first column deliberately favors the nominal margin by assigning the slow store its full hindsight revenue. The last two columns concern the artificial error price.

Table 4: Same-window comparison, AUD per operating hour at one MW. Hindsight and forecast columns share the same endpoints. Regional results refer only to the four selected months.

Region	Duration	$(M-1)\times$ coarse hindsight	Forecast rule	Error hindsight
NSW	1 h	1.17	2.79	14.70
NSW	4 h	2.08	3.33	15.39
Victoria	1 h	0.74	0.84	7.42
Victoria	4 h	1.68	0.79	8.12

For NSW the archive-informed policy's realized correction exceeded even the optimistic nominal margin at both durations. In Victoria the ordering of these two columns differed by duration. The four-hour Victorian result illustrates the identification problem particularly clearly: the observed feasible rule earned less than the margin, yet the hindsight ceiling for the correction was considerably higher. Neither observation pins down the population correction. The comparison does not estimate an optimal speed premium or demonstrate a sign reversal of the gain. A fine store can always repeat feasible coarse instructions.

Several features delimit the result. Stationarity is a modeling assumption; the Australian sample has seasonality and exceptional days. Aggregating across phases in the data is an explicit design choice, whereas equivalence of phases in the theorem is a consequence of fine-grid stationarity. Changing the decision clock and the observation clock together could introduce a distinct information advantage. A real battery has asymmetric efficiencies, age-dependent costs and often additional commitments. Those can break both the account-splitting construction and the exact conservation argument. The data compare hypothetical decisions under public time stamps, with no order book, bids or market impact and no observation of the true optimum. Contract prices also require costs, risk preferences and other entitled revenues.

The defensible valuation result is a test. Given a price law or trustworthy conditional forecasts, a trader can ask whether the asset can physically exploit offsetting expected movements within a coarse block, and compute a lower premium without a full inventory solve. An inconclusive result is informative about the method: a longer-duration store or a large predictable model discrepancy can leave the simple bounds too wide. In the Australian application, the archives establish that

part of the discrepancy has decision-time content. An enforceable price for changing a real battery's dispatch schedule would require additional evidence on attainable cash flows.

References

[1] S. D. Hodges (2004). The value of a storage facility. Warwick Business School, working paper 04-215. Warwick working paper series.

[2] R. A. Carmona and M. Ludkovski (2010). Valuation of energy storage: An optimal switching approach. *Quantitative Finance* 10, 359–374. doi:10.1080/14697680902946514.

[3] Z. Chen and P. A. Forsyth (2007–2008). A semi-Lagrangian approach for natural gas storage valuation and optimal operation. *SIAM Journal on Scientific Computing* 30, 339–368. doi:10.1137/060672911.

[4] A. Faghih, M. Roozbehani and M. A. Dahleh (2013). On the value and price-responsiveness of ramp-constrained storage. *Energy Conversion and Management* 76, 472–482. doi:10.1016/j.enconman.2013.07.072.

[5] E. R. Jakobsen, A. Picarelli and C. Reisinger (2019). Improved order 1/4 convergence for piecewise constant policy approximation of stochastic control problems. *Electronic Communications in Probability* 24, paper 59. doi:10.1214/19-ECP256.

[6] D. Metz and J. T. Saraiva (2018). Use of battery storage systems for price arbitrage operations in the 15- and 60-min German intraday markets. *Electric Power Systems Research* 160, 27–36. doi:10.1016/j.epsr.2018.01.020.

[7] M. C. Soini, D. Parra and M. K. Patel (2019). Disaggregation of energy storage operation by timescales. *Journal of Energy Storage* 23, 480–494. doi:10.1016/j.est.2019.04.005.

[8] D. Schaurecker, D. Wozabal, N. Löhndorf and T. Staake (2026). High-frequency intraday trading for battery storages. *Applied Energy* 425, article 128550. doi:10.1016/j.apenergy.2026.128550.

[9] E. Seneta (1979). Coefficients of ergodicity: Structure and applications. *Advances in Applied Probability* 11(3), 576–590. doi:10.2307/1426955.

[10] G. Ch. Pflug and A. Pichler (2012). A distance for multistage stochastic optimization models. *SIAM Journal on Optimization* 22, 1–23. doi:10.1137/110825054.

[11] D. Mukherjee and A. A. Panda (2026). Forecast projection and decision regret: Theory and evidence from battery arbitrage. *Preprint*, arXiv:2611.14030.

[12] B. Le Roux Tardif (2026). Price information is not enough: Ordering and decision rules in storage bidding. *Preprint*, arXiv:2608.08377.

[13] Australian Energy Market Commission (2017). Five minute settlement. Rule change ERC0201, final determination. AEMC rule-change record.

[14] Australian Energy Market Operator. Aggregated price and demand data. Historical NEM regional data and revision policy, accessed November 2026. AEMO data documentation.

[15] Australian Energy Market Operator. Pre-dispatch. MMS documentation for regional five-minute pre-dispatch forecasts, accessed November 2026. AEMO pre-dispatch documentation.

[16] Australian Energy Market Operator. NEMWeb wholesale electricity MMS Data Model monthly archives, 2025. Accessed November 2026. NEMWeb archive.