Supplementary Methods and Results

Forecast Calibration and Charging Feedback in PV–Battery Scheduling across Three Environments

Reference numbers follow the main manuscript. This supplement preserves the complete model, frozen selections, secondary results and reproducibility records underlying the compressed article. The reported outcomes come from the existing fixed-horizon run with the uniformly amended numerical policy; editorial relocation adds no experiments or design changes. Earlier hourly outcomes and the superseded numerical attempt remain historical records and are not pooled with these results. Uncalibrated identifies the forecast treatment without residual correction, whereas S0 identifies the reference tariff and export scenario. The data, source code and supporting records named below are available from the first author, Xuan-You Liu (liuxuanyou@cu.ac.kr), upon request; local evidence links are shown as file identifiers rather than public archive links.

S1. Complete model and implementation

Table S1. Parameter values and their basis. Cited precedents support model forms or motivation, not the exact numerical choices unless stated. The equations and selection rules below define the implemented policy.

Parameter	Value or candidate set	Basis and selection	Reference boundary
State update	Charge/discharge losses retained	AC-power convention; Equation S7	Model form [17,22]
SOC bounds; initial SOC	0.20–0.80; 0.50 per episode	Equal assumed initialization; no vendor calibration	Bounded SOC [15,22]
One-way efficiency	√0.85 each	Assumed round-trip efficiency 0.85	Sensitivity rationale [17]
Scale P*	Training native load peak, kW	No test-based hardware sizing	Source data [37,38]
PV amplitude	Training PV peak = 1.5P*	Controlled normalization; native-PV diagnostic retained	Study assumption
Rated energy; power	1.5 h × P*; energy / 4 h	Both multiplied by 0.5 or 2 in sensitivities	Sizing context [23,35]
Horizon and steps	12 h; hourly plans; A 5 min, B/C 15 min	Fixed horizon; native execution resolution	Rolling principle [1,22]; data [37,38]
Import price	0.08 / 0.14 / 0.24 CU/kWh; S2 peak 0.40	Hypothetical low/shoulder/peak prices	Price dependence [11,35]
Export price; cap	0.04 CU/kWh; 0.2P*	S1 cap zero; relaxed cap 1.0P*	Surplus context [23]
Import limit	10P*	Common planning limit and execution check	Study assumption
Discharge proxy κ	199/6000 CU/kWh; 0.5× / 2×	Illustrative throughput charge; no aging validation	Proxy interpretation [17]

Parameter	Value or candidate set	Basis and selection	Reference boundary
Curtailment preference γ	0.01 CU/kWh	Planning only; excluded from cash	Study assumption
Planning terminal value ν	0.08, 0.14, 0.20 CU/kWh	Joint validation selection with D	Terminal accounting [17,23]
Ledger energy credit ξ	0.08 CU/kWh	Common state-adjusted diagnostic; not revenue	Study accounting rule
Median residual window D	14, 30, 60 calendar days	Mean validation adjusted cost, original four cells	Residual context [15,17]
Weighted correction	60-day history; half-life 3/7/14/30 d; gain 0/.25/.5/.75/1	First-hour validation component MAE sum	Prespecified diagnostic
Charging intensity α	0, 0.5, 1	Blend only when planned discharge inactive	Rolling-operation context [22]
Fixed margin	0/.025/.05/.1/.2 of P*	Validation mean across median endpoint cells	Study diagnostic
Adaptive margin	q=.5/.8/.9; ϱ=.1/.25/.5	Sum of component nearest-rank quantiles	Study diagnostic; no coverage claim
Action gate; objective budget	10^{-7}; 10^{-8} max(1, absolute optimum)	Normalized units; current-throughput selection	Solver family [39]; study tolerances
Paired resampling	Blocks 3/7/14 d; 2000 draws	Observation-contiguous circular blocks; seed 20260929	Conditional statistical sensitivity

S1.1. Notation, data eligibility and normalization

The notation distinguishes information, action stage and physical state. The complete data splits and frozen choices appear in Tables S3 and S5. Hourly decisions are indexed by k, forecast targets by h, native execution intervals by t, historical hours by n, and UTC days by d. The map ψ returns local clock hour; τ returns UTC clock hour. Local time determines tariffs and residual groups; UTC determines availability, elapsed time and eligible days. Repeated daylight-saving hours retain distinct UTC identities, and absent local hours are not synthesized. The planning step ΔT is one hour, the horizon H is 12 steps, and the execution step Δt is 1/12 hour in A and 1/4 hour in B/C.

The vector x=(g,l) contains normalized PV and load powers. Physical source powers G and L are in kW. The symbols a, b and v denote normalized grid import, grid export and PV curtailment. The two components of u are charging (+) and discharging (−); the superscript r identifies planning (0), requests (1) and realized action (2). Forecast treatments i=0,1,2 denote uncalibrated, median and exponentially weighted mean correction. Execution intensity α takes values 0, 0.5 and 1; j=0,1 is reserved for the endpoint comparison. The LP decision vector is y. Stored energy s is divided by training peak load and therefore has units of hours; $z=s/\bar{E}$ is dimensionless SOC. A tilde identifies a common-state shadow diagnostic. These conventions apply throughout the supplement and main text.

Only complete native hours and complete valid UTC days enter evaluation. Missingness remains on the time grid, execution truth is not imputed, and energy state is not propagated across excluded days. Every contiguous evaluation episode starts at 50% SOC. A uses PERFORM solar433 model-generated PV and ERCOT Coast demand scaled by 0.00025; these are not co-located installation measurements. B/C use OPSD CoSSMic DE_KN_residential3/4, release 15 April 2020. Exact endpoint differences of cumulative PV and import/export meters are divided by interval duration, and load is reconstructed as PV plus import minus export. Source-interpolated intervals, missing endpoints, negative differences and negative reconstructed load are excluded. Source flag fractions of 5.86%/5.43% overlap the final invalid fractions of 5.95%/6.75%. Original capacities, wiring, unobserved storage and unflagged regularization cannot be verified. The residential choices used eligible coverage and meter interpretability; residential6 remains in the screening record.

C includes 43 single-day test episodes. These are evaluation fragments, not independent installations. Training references use eligible native intervals in the set $\mathcal{T}_0$. Let ζ be the PV amplitude factor, E rated battery energy, P rated battery power, θ the energy-to-load ratio and ω the rated duration:

$$\begin{gathered} P_* = max_{t\in\mathcal{T}_0} L_t, \zeta = \frac{1.5P_*}{max_{t\in\mathcal{T}_0} G_t}, \\ x_t = \left(\frac{\zeta G_t}{P_*}, \frac{L_t}{P_*}\right), \\ E = \theta P_*, P = \frac{E}{\omega}, \overline{E} = \frac{E}{P_*}, \overline{P} = \frac{P}{P_*}. \end{gathered} \tag{S1}$$

The values θ=1.5 hours and ω=4 hours are simulation assumptions. Rated energy is in kWh and rated power in kW. The usable 20–80% SOC window is 60% of rated energy and provides 2.4 hours at rated power. PV peak is scaled to 1.5 times training peak load. The native-PV diagnostic reverses that PV amplitude scaling; it is another scenario for the same source. Capacity sensitivities multiply both rated energy and power by 0.5 or 2. Hourly powers are complete-hour averages of native powers.

S1.2. Causal information and forecast correction

Every decision uses targets k through k+11, including near episode ends; future episode endpoints are not provided to the optimizer. Define $\mathcal{V}$ as valid observed hours, A_{kh} as a binary source-availability indicator, and f_h as the archived two-component source product. A_{kh} is one only when both products are finite and both issue times are no later than the decision. Source products are used in A; B/C use persistence. Define the eligible same-UTC-hour history $\mathcal{P}$ and the fallback p:

$$\begin{gathered} \mathcal{P}_{kh} = \{n \in \mathcal{V}: n < k - \ell,\ \tau(n) = \tau(h)\}, \\ p_{kh} = \begin{cases} x_{max\mathcal{P}_{kh}}, \mid \mathcal{P}_{kh} \mid > 0, \\ 0, \mid \mathcal{P}_{kh} \mid = 0, \end{cases} \\ \hat{x}_{h|k}^{(0)} = \begin{cases} f_h, A_{kh} = 1, \\ p_{kh}, A_{kh} = 0. \end{cases} \end{gathered} \tag{S2}$$

The additional completion lag ℓ is zero in A and one hour in B/C. Zero is the declared no-history fallback. Fallback target counts are recorded. The observed-future-information diagnostic substitutes future truth only for targets with valid truth and otherwise retains this causal forecast; its coverage is reported. It is neither deployable nor an unconditional optimum.

Let $f°_n$ denote the archived source/persistence forecast associated with historical target n. This historical target definition is retained: the calibration is not fitted separately to every horizon lead. Residential historical persistence uses the most recent valid

observation of the same UTC hour no later than the target minus 24 hours. The residual vector and the local-hour group available to decision k are

$$\begin{gathered} r_n = x_n - f_n^{\circ}, \\ \mathcal{H}_{kv}(D) = \{n \in \mathcal{V} : k - 24D \le n < k - \ell, \\ \psi(n) = v,\ r_n \in \mathbb{R}^2\}. \end{gathered} \tag{S3}$$

The symbol v in the index of $\boldsymbol{\mathcal{H}}$ is a local clock-hour value and is distinct from curtailment v with a time index. D is a calendar-day window, not a count of available days. Both component residuals must be finite. Define $\boldsymbol{\mathcal{M}}$ as the componentwise sample median, taking the middle order statistic for odd size and the mean of the two central order statistics for even size. Empty histories give zero correction. With β denoting gain and τ_1 the half-life in days, the corrections and forecasts are

$$\begin{gathered} \mathbf{b}_{kv}^{(0)} = 0, \\ \mathbf{b}_{kv}^{(1)} = \mathcal{M}(\{r_n : n \in \mathcal{H}_{kv}(D)\}), \\ w_{kn} = 2^{-(k-n)/(24\tau_1)}, \\ \mathbf{b}_{kv}^{(2)} = \beta \frac{\sum_{n \in \mathcal{H}_{kv}(60)} w_{kn}\, r_n}{\sum_{n \in \mathcal{H}_{kv}(60)} w_{kn}}, \\ \hat{x}_{h|k}^{(i)} = \left[\hat{x}_{h|k}^{(0)} + \mathbf{b}_{k,\psi(h)}^{(i)}\right]_+, i \in \{1,2\}. \end{gathered} \tag{S4}$$

The weights w decay exponentially with historical age. The positive-part operation acts componentwise and enforces nonnegative corrected PV and load. The uncalibrated forecast is defined by Equation S2. Median and decaying residual correction are experimental treatments, not new forecast algorithms. The weighted mean uses a fixed 60-day history. Its half-life and gain minimize the sum of normalized component MAEs on the first planned hour of the original validation partition $\mathcal{T}_1$:

$$(\tau_1^*, \beta^*) \in \underset{\substack{\tau_1 \in \{3,7,14,30\} \\ \beta \in \{0,.25,.5,.75,1\}}}{\operatorname{arg\,min}} \frac{1}{|\mathcal{T}_1|} \sum_{k \in \mathcal{T}_1} \| \hat{x}_{k|k}^{(2)} - x_k \|_1. \tag{S5}$$

Scores within 10^{-12} are tied; smaller gain, then longer half-life wins. Zero gain is admissible and reproduces the uncalibrated forecast on the prepared data. D is selected from 14, 30 and 60 days and terminal value ν from 0.08, 0.14 and 0.20 CU/kWh by minimum mean state-adjusted validation cost over the original four uncalibrated/median × endpoint cells. Cost scores within the larger of 10^{-8} CU and 10^{-10} times the absolute minimum are tied; smaller D, then smaller ν wins. All choices are frozen before test evaluation and are not retuned by scenario. Historical test outcomes were known before this revision protocol was fixed.

S1.3. Planning model and charging margins

Import and export prices are λ and μ; κ is the discharged-energy proxy, γ the curtailment preference and ν terminal energy value, each in CU/kWh. The normalized objective is

$$J_k(y) = \sum_{h=k}^{k+H-1} \left(\lambda_h a_h - \mu b_h + \kappa u_h^{(0),-} + \gamma v_h\right) \Delta T - \nu s_{k+H}. \tag{S6}$$

Multiplication by P* converts the objective to CU. The fixed values are μ=0.04, κ=199/6000 and γ=0.01 CU/kWh. The throughput coefficient is an illustrative economic proxy, not measured aging or validated replacement cost. Curtailment preference and

terminal value guide planning but are not electricity transactions. Constraints use forecast powers exclusively:

$$\begin{gathered}\hat{g}_{h|k} - v_h + a_h + u_h^{(0),-} = \hat{l}_{h|k} + u_h^{(0),+} + b_h,\\ s_{h+1} = s_h + \eta_+ u_h^{(0),+}\Delta T - \frac{u_h^{(0),-}\Delta T}{\eta_-},\\ 0.2\overline{E} \le s_h \le 0.8\overline{E}, s_k = s_k^\circ,\\ 0 \le u_h^{(0),+} \le min\{\overline{P}, [\hat{g}_{h|k} - \hat{l}_{h|k} - m_{h|k}]_+\},\\ 0 \le u_h^{(0),-} \le min\{\overline{P}, [\hat{l}_{h|k} - \hat{g}_{h|k}]_+\},\\ 0 \le a_h \le 10, 0 \le b_h \le \overline{B}, 0 \le v_h \le \hat{g}_{h|k}.\end{gathered} \tag{S7}$$

The hats carry the selected treatment, suppressed here for legibility. Bounds apply to all planned actions and states through the terminal state. The starting energy s° is the simulated state entering the hour. Conversion efficiencies $\eta_+=\eta_-=\sqrt{0.85}$ give a round-trip product of 0.85. The normalized export cap is $\bar{B}$. Direction-specific bounds forbid grid charging and battery export and structurally prevent simultaneous positive charging and discharging. Import is limited to ten times training peak load in planning and checked against that same limit in execution.

Margins are zero in the primary design. The fixed diagnostic uses a constant from 0, 0.025, 0.05, 0.1 and 0.2 in normalized power units. For the adaptive diagnostic, Qq is the nearest-rank empirical quantile: an ascending sample of n values uses rank max(1,⌈qn⌉). With $r_{n,}c$ and $b_{kv,}c$ denoting component c of the residual and median correction, the adaptive margin is

$$m_{h|k} = \rho \sum_{c=1}^{2} Q_q\left(\{|\, r_{n,c} - \mathbf{b}^{(1)}_{k,\psi(h),c} \,|: n \in \mathcal{H}_{k,\psi(h)}(D)\}\right). \tag{S8}$$

An empty history gives zero margin. The component quantiles are added; this is not a quantile of net-load residuals. Candidates are q=0.5, 0.8 or 0.9 and ϱ=0.1, 0.25 or 0.5. Both margin diagnostics use median-calibrated forecasts and minimum mean state-adjusted validation cost over the two execution endpoints. The same cost tie tolerance as for D and ν applies; smaller ϱ then smaller q wins for adaptive margins, and the smaller constant wins for fixed margins. These diagnostics alter charging headroom and make no chance-constraint coverage claim. An adaptive margin combined with native-PV rescaling is not part of the registered implementation.

Let $\boldsymbol{\mathcal{F}}_k$ be the feasible set defined by Equation S7. Each hourly plan first minimizes economic cost, then minimizes current charge plus discharge subject to a budget on the primary objective:

$$\begin{gathered}J_k^* = min_{y\in\mathcal{F}_k} J_k(y),\\ y_k^* \in \arg min_{\substack{y\in\mathcal{F}_k \\ J_k(y)\le J_k^*+\epsilon_k}} \left(u_k^{(0),+} + u_k^{(0),-}\right),\\ \epsilon_k = 10^{-8} max(1, |\, J_k^* \,|).\end{gathered} \tag{S9}$$

This rule defers current throughput among economically near-equivalent schedules; it does not guarantee a unique full schedule. The opposing sensitivity maximizes current throughput under the same objective budget. No chronological perturbation is added to the economic objective.

S1.4. Requests, physical clipping and settlement

Only the first planned pair is held as the reference during the current hour. The delayed surplus f_t is zero when the delayed interval is invalid; otherwise it is the positive part of delayed PV minus load, capped at $\bar{P}$. The native delay ℓ_e is one interval in A and

two in B/C. Thus residential feedback includes an additional fifteen-minute latency beyond the preceding completed interval. With action threshold $\delta=10^{-7}$ in normalized power units, the request is

$$\begin{aligned} u_t^{(1)} &= (0, u_k^{(0),-}), \quad \text{if } u_k^{(0),-} > \delta, \\ u_t^{(1)} &= ((1-\alpha)u_k^{(0),+} + \alpha f_t, 0), \quad \text{if } u_k^{(0),-} \le \delta. \end{aligned} \tag{S10}$$

The delayed surplus is computed using x at $t-\ell_e$. At α=0 the charging reference is retained; at α=1 delayed surplus replaces it; α=0.5 blends them. The discharge gate and discharge request are unchanged. This is restricted charging feedback, not a bidirectional optimizer. Current interval-average truth enters the subsequent physical replay:

$$\begin{aligned} u_t^{(2),+} &= \left[\min\left\{u_t^{(1),+}, \overline{P}, [g_t - l_t]_+, \frac{0.8\overline{E} - s_t}{\eta_+ \Delta t}\right\}\right]_+, \\ u_t^{(2),-} &= \left[\min\left\{u_t^{(1),-}, \overline{P}, [l_t - g_t]_+, \frac{(s_t - 0.2\overline{E})\eta_-}{\Delta t}\right\}\right]_+. \end{aligned} \tag{S11}$$

This is an interval energy-feasibility simulation; five- and fifteen-minute averages do not validate instantaneous device control. Remaining surplus n determines settlement, and realized actions update stored energy:

$$\begin{aligned} n_t &= g_t + u_t^{(2),-} - l_t - u_t^{(2),+}, \\ a_t = [-n_t]_+, b_t &= min\{\overline{B}, [n_t]_+\}, v_t = [n_t - b_t]_+, \\ s_{t+1} &= s_t + \eta_+ u_t^{(2),+} \Delta t - \frac{u_t^{(2),-} \Delta t}{\eta_-}, \\ g_t - v_t + a_t + u_t^{(2),-} &= l_t + u_t^{(2),+} + b_t. \end{aligned} \tag{S12}$$

Matched treatments share feasible days, episode boundaries, initialization, physical constraints and prices. Every replay audits balance residuals, SOC limits and action overlap. A raw constraint flag records a request reaching or exceeding a limit. An attributed binding event additionally requires positive clipping and agreement between the realized action and the reconstructed limiting bound within the action tolerance. Limits can bind jointly. Individual binding counts are not mutually exclusive energy-loss fractions; combination-based ledger partitions retain all jointly active limits.

S1.5. Economic ledger, contrasts and experimental matrix

Let $\mathcal{E}$ index evaluated episodes and let $s_{e,0}$ and $s_{e,1}$ denote their starting and ending normalized energies. Cash cost C and the separately reported state-adjusted quantity C† are

$$\begin{aligned} C &= P_* \sum_t \Delta t \left(\lambda_t a_t - \mu b_t + \kappa u_t^{(2),-}\right), \\ C^\dagger &= C - \xi P_* \sum_{e \in \mathcal{E}} (s_{e,1} - s_{e,0}). \end{aligned} \tag{S13}$$

The common terminal asset credit ξ is 0.08 CU/kWh, distinct from the validation-selected planning value ν. It is not cash revenue. Equivalent full cycles divide total AC-side discharged kWh by nominal battery kWh; throughput sums charged and discharged kWh. Neither quantity estimates electrochemical aging or divides by usable-window capacity. Terminal adjustment is also reported under daily resets.

For a fixed scenario and environment, C_{ij} uses forecast i and execution endpoint j. Define ΔC_{ij} as calibrated minus uncalibrated cost at the same endpoint. The median endpoint interaction and percentages are

$$\begin{gathered} \Delta C_{ij} = C_{ij} - C_{0j}, \\ I = (C_{11} - C_{01}) - (C_{10} - C_{00}), \\ \delta_{ij} = 100\frac{\Delta C_{ij}}{C_{0j}}, \iota = 100\frac{I}{C_{00}}. \end{gathered} \tag{S14}$$

The denominator for a calibration percentage is its own uncalibrated execution cell. Only the interaction percentage uses C_{00}. Negative I means calibration's cost effect is lower with replacement; this can attenuate harm without making calibration beneficial. Weighted-mean and intermediate-intensity contrasts use the same pairing with their treatment identified explicitly. Absolute CU and CU per MWh of served load accompany percentages; the three environments are not pooled as independent site replicates.

For a paired calibration contrast, let ΔA_ν, ΔB_t and ΔD_t be differences in realized imported, exported and discharged interval energy in kWh. The cash ledger closes as

$$\Delta C = \sum_t \lambda_t \Delta A_t - \mu \sum_t \Delta B_t + \kappa \sum_t \Delta D_t. \tag{S15}$$

Import prices remain inside the sum. This accounting identity identifies transaction contributions, not independent causal mediators. State-adjusted differences add the separate change in the terminal credit.

The primary matrix crosses three environments, three forecasts, three intensities and three scenarios for 81 configurations. In B and C, the selected weighted-mean gain is zero, so its 18 environment × intensity × scenario cells duplicate the corresponding uncalibrated outcomes and add no distinct forecast treatment. S0 uses import prices 0.08/0.14/0.24 CU/kWh for low/shoulder/peak periods and export cap 0.2P*. S1 changes only the export cap to zero; S2 changes only peak import price to 0.40 CU/kWh. Export is credited at 0.04 throughout. Low periods are local 00:00–06:59 and 22:00–23:59; peak is 16:00–20:59; other hours are shoulder. These common illustrative tariffs are not country-specific bills or investment estimates.

The registered diagnostic matrix includes half/double capacity and power, export cap 1.0P*, half/double κ, a flat import price of 0.14, native PV amplitude, daily state resets, alternative current-throughput selection, hourly cold starts and retained episode-boundary bases. Ordinary sensitivities preserve the original four cells. A-flat-price and C-double-capacity use all nine forecast/intensity combinations, retaining previously examined flat-price and storage-reversal cases regardless of their revised outcome. Transfer uses the source-selected D and ν jointly in another environment, without target retuning; it cannot isolate their individual influences. Observed-future-information and fixed/adaptive-margin diagnostics use both endpoints. The frozen matrix comprises 265 test configurations and 12 additional common-SOC shadows; these counts are not independent systems.

S1.6. Action transmission and common-state shadows

Actions are compared as two-component hourly charge/discharge energy vectors in kWh. Let $\mathcal{N}_k$ contain the native intervals of hour k. With forecast and execution indices shown only when needed, define

$$\begin{gathered} U_k^{(0)} = P_* \Delta T u_k^{(0)}, \\ U_k^{(r)} = P_* \sum_{t \in \mathcal{N}_k} u_t^{(r)} \Delta t, r \in \{1,2\}, \\ \Delta U_k^{(r)} = U_{k,ij}^{(r)} - U_{k,0j}^{(r)}. \end{gathered} \tag{S16}$$

The energy transmission ratio T_1 and count transmission ratio T_0 are

$$\begin{aligned} T_1 &= \frac{\sum_k \|\Delta U_k^{(2)}\|_1}{\sum_k \|\Delta U_k^{(0)}\|_1}, \\ T_0 &= \frac{\sum_k \mathbf{1}\{\|\Delta U_k^{(2)}\|_\infty > \varepsilon_E\}}{\sum_k \mathbf{1}\{\|\Delta U_k^{(0)}\|_\infty > \varepsilon_E\}}, \\ \varepsilon_E &= 10^{-6} P_* \Delta T. \end{aligned} \tag{S17}$$

These correspond to the reported energy and count ATR measures. L1 adds absolute charge and discharge changes and must not be replaced by absolute net-power change. The infinity norm tests whether either component differs above a common energy threshold. A zero denominator makes a ratio undefined. Ratios may exceed one because gating and state histories can amplify differences; they are descriptive, not conversion efficiencies or causal shares. The intermediate requested vector is retained. Comparisons are stratified by whether neither, one or both forecasts activate planned discharge. Hourly sums can conceal within-hour timing differences, so native traces remain available.

The common-state diagnostic reinitializes every hour from the recorded starting energy under uncalibrated forecasts and planned execution, $s_{k,00}$. The map $\mathcal{G}$ comprises planning, requests, clipping and within-hour state updates using the same exogenous information:

$$\begin{aligned} \tilde{s}_{k,ij} &= s_{k,00}, \\ \tilde{U}_{k,ij}^{(2)} &= \mathcal{G}\big(\hat{x}_{k:k+H-1|k}^{(i)}, j, s_{k,00}; x_{\mathcal{N}_k}\big). \end{aligned} \tag{S18}$$

The shadow endpoint is discarded before the next hour, which starts from that next hour's recorded state on the uncalibrated planned trajectory. Both plans and realized actions can therefore change relative to fully propagated trajectories. This is a one-hour diagnostic of state-path dependence, not a feasible long-run policy, an operating bill or a formal mediation decomposition. Equality under shared non-discharge gates, full request replacement and equal states follows from the execution rule; occurrence and economic weight on the propagated paths remain empirical quantities.

S1.7. Regimes and conditional uncertainty

Daily surplus share R, within-hour net-load variation V and near-bound SOC fraction S are

$$\begin{aligned} R_d &= \frac{\sum_{t\in d}[g_t - l_t]_+ \Delta t}{\sum_{t\in d} g_t \Delta t}, \\ V_d &= \frac{1}{24}\sum_{k\in d}\frac{1}{N-1}\sum_{q=2}^{N} \left| (l_{kq} - g_{kq}) - (l_{k,q-1} - g_{k,q-1}) \right|, \\ S_d &= \frac{1}{T_d}\sum_{t\in d} \mathbf{1}\{z_{t+1} \in (-\infty, .21] \cup [.79, \infty)\}\Delta t, z_{t+1} = s_{t+1}/\overline{E}. \end{aligned} \tag{S19}$$

N is the native intervals per hour (12 or 4), and T^d is evaluated duration in hours. Surplus share is undefined for zero daily PV energy. Ramps exclude boundaries between hours, then average the 24 hourly values. The near-bound SOC fraction uses end-of-interval states, matching the saved implementation. Forecast-error, surplus and ramp strata split at their environment-specific training medians; equality belongs to the high stratum, and missing descriptors form an undefined stratum. The forecast-error descriptor is daily mean absolute archived uncalibrated net-load error. The common uncalibrated forecast with α=0 supplies the near-bound SOC fraction within each scenario; a fraction at or above

0.25 is high. The near-bound SOC fraction remains a post-treatment descriptor. Bins below ten days are flagged sparse, and native resolution limits ramp comparisons. Extreme positive and negative interaction days are retained and shown with preceding 24-hour context without being presented as representative days.

Daily contrasts are paired on identical UTC dates. Circular block resampling uses lengths 3, 7 and 14 days, 2000 draws and seed 20260929 within observation-contiguous segments. For a segment of n days, use block width min(block length,n), sample enough circular block starts to cover n positions, concatenate the blocks and truncate to exactly n days. Segment sizes and composition remain fixed. A segment no longer than the block contributes its exact total to every draw; fixed day counts are reported. Artificial daily SOC resets do not split observation segments. The 2.5th and 97.5th percentiles of summed contrasts form the reported nominal 95% bands. Seasonal frequentist coverage is not validated. These bands describe sensitivity conditional on observed segments; they omit missing-period and between-site uncertainty, and a band crossing zero is not evidence of superiority. Daily sign counts treat absolute differences no greater than 10^{-8} CU as tied.

S1.8. Numerical implementation and scope

The LP has six H-length blocks: charging, discharging, import, export, curtailment and end-of-step energy, with H balance rows and H state-transition rows. HiGHS dual simplex uses one thread, primal and dual feasibility tolerances of 10^{-9} and seed 20260929. The secondary LP adds the primary objective-bound row and disables presolve uniformly. Within-episode bases are reused, and both LP instances are cleared at episode boundaries. Cold-hourly and retained-boundary-basis diagnostics test dependence on these choices. Every hour logs the primary optimum, selected objective, deviation and budget. The primal feasibility allowance is 2×10^{-7} in normalized units; the objective-bound audit adds 2×10^{-8} numerical slack. Native replay checks import and SOC with 10^{-7} allowance and records overlap, balance and extrema.

An earlier cold-start check returned secondary infeasibility despite a feasible primary witness. Disabling secondary presolve was applied uniformly, followed by full revalidation and rerunning under the unchanged economic selection rule. The superseded attempt remains archived. This statement records the existing study history; editorial compression adds no simulation, selection or analysis. Prepared inputs, frozen settings, forecast coverage, interval/hour/day ledgers, diagnostic outputs, hashes and saved-output/full-replay audit records remain the reproducibility evidence.

The system has fixed load, no grid charging, no battery export and restricted charging recourse. It omits demand charges, network/inverter constraints, bidirectional discharge recourse and detailed calendar, thermal or state-dependent aging. Model-generated A and the common residential acquisition source limit environmental independence; fragmented C limits long-state-path inference. These limits apply to every equation and contrast above.

S2. Closest-study comparison and evidence boundaries

Table S2. Eight closest studies retained from the full comparison matrix. The final column describes a contrast not established by the examined methods and comparisons, rather than a global absence claim or a flaw in work with another objective.

Study	Information and execution	Environment	Evaluated evidence	Present contrast beyond that evidence
Schmitt et al. [8]	PV-error variation; current-step perfect information versus forecast-only handling and compensation	German commercial model; 30 days in 2020	Energy/peak costs and SOC paths vary with information handling	Cross the same causal calibration with fixed-latency charging-request rules
Ma et al. [22]	Scenario-based ultra-short-term rolling operation under forecast uncertainty	Daqing 2022 PV with specified regional demand	Load shifting and operational risk	Matched calibration × execution operating-cost contrast
Xu et al. [23]	Scenario-oriented uncertainty; joint forecasting and declaration strategy	Daqing PV; weather, storage and DC/AC cases	Generation and deviation penalties as two objectives	Execution-rule intervention at fixed hardware and multiple source profiles
Cheng et al. [15]	Point/fixed-conformal/adaptive-conformal net load; common first-action MPC	Modeled PV–battery–EV station; two-month test and stress events	Cost/reserve/coverage, SOC and throughput	Separate point-calibration interaction from risk-margin comparisons
Zhang et al. [17]	Load forecast; day-ahead scheduling versus closed-loop MPC with measured load and SOC	Australian load/price case; annual comparison	Common-ledger economics and non-additive information/feedback effects	Full calibration × request-rule crossing plus stage and numerical-selection audits
Pereira et al. [2]	Rolling XGBoost and adaptive MPC; shrinking horizon and charging constraints	Metered South African commercial facility	TOU costs and charging-peak comparisons	Orthogonal forecast/execution treatments with paired action traces
Langtry et al. [34]	Forecast data/model reuse under common MPC first-action implementation	Multiple measured building loads; simulated energy-system operation	Accuracy versus storage-MPC cost; model reuse	Vary execution permission independently of the forecast/data treatment

Study	Information and execution	Environment	Evaluated evidence	Present contrast beyond that evidence
Karttunen et al. [35]	Rule-based/MPC/RL policies under tariffs and battery usage cost	Four measured household loads; common PV; 2020–2022 prices	Bills, profitability and price-dependent rankings	Matched calibration × execution under one policy architecture

The full comparison matrix CSV (literature/comparison_matrix.csv) retains the original seven comparison fields. The evidence-anchor record (literature/comparison_evidence_anchors.md) identifies the examined sections, equations and figures for every row. Table S2 retains eight studies of the existing ten-study comparison; the two studies of general PV forecasting and plant-level stochastic control that are no longer cited in the main text remain in the full matrix file. Reference numbers follow the revised main bibliography. In particular, [17] already reports non-additive information and feedback effects, while [22,23] already evaluate rolling and declaration-based PV–storage operation under forecast uncertainty; the present contribution is the narrower controlled comparison documented in the main text.

S3. Data registry, forecast coverage and frozen choices

The full data registry (data_registry.csv) retains source/site identity, licenses, chronological partitions, exclusion fractions, scaling, hardware references, interval conventions and input hashes. A combines the PERFORM solar433 series with ERCOT Coast demand multiplied by 0.00025 before the stated training-based normalization; it is a constructed model-generated benchmark [37]. B/C are distinct meters within the same OPSD/CoSSMic project, region and acquisition process, with reconstructed demand. They are not three independent metering campaigns. Public dataset licenses are CC BY 4.0 and require attribution [37,38]. Original residential nameplates, wiring and absence of behind-meter storage remain unverified.

Table S3. Chronological eligible-day spans and test fragmentation. Gaps within spans remain excluded.

Environment	Training eligible span and days	Validation eligible span and days	Test eligible span	Test days / episodes
A	1 Jan–28 Feb 2018; 59	1–31 Mar 2018; 31	1 Apr–31 Dec 2018	275 / 1
B	29 Feb–28 Apr 2016; 60	29 Apr–5 Jun 2016; 30	6 Jun 2016–7 Jul 2017	368 / 6
C	27 Nov 2015–15 Feb 2016; 60	16 Feb–12 Aug 2016; 30	14 Aug 2016–3 Feb 2018	213 / 77

Table S4. Training normalization and retained coverage. Invalid fractions reflect overlapping source and physical exclusions and must not be added to source-interpolation rates.

Environment	Train / validation / test days	P* (kW)	PV scale	Invalid (%)
A	59 / 31 / 275	3845.4358	0.8248	0.0000
B	60 / 30 / 368	4.1520	1.6035	5.9465
C	60 / 30 / 213	6.7840	1.4340	6.7516

The source-processed residential records regularize approximately one-minute irregular meter readings. Released-file residual missingness is zero, but source interpolation affects 5.8622% of B and 5.4327% of C intervals, and the original raw missingness cannot be recovered. Physical-invalidity fractions are 0.8465% and 1.4687%; these overlap other exclusions. No counter resets were detected in the retained preprocessing record. The full registry preserves precision and those qualifications. A has a modeled source rating of 7500 kW; residential rated capacity and module/inverter specifications are unknown. Intervals represent average power over the recorded start-to-end interval and become available after the endpoint. A uses America/Chicago local time; B/C use Europe/Berlin, with elapsed time and eligibility in UTC.

A has 66 causal fallback targets among 79,200 test planning targets (0.0833%). B/C use persistence throughout. Target-level forecast audits retain 246,528 horizon records with information timestamps and treatment forecasts. The twelve-hour planning horizon is retained even at archive boundaries; missing execution truth ends scoring and the next eligible episode is reinitialized. Forecast availability is independent of any future episode endpoint.

Table S5. Frozen validation selections. Weighted-mean gain zero reproduces the uncalibrated forecast, and zero fixed margin reproduces the calibrated no-margin reference.

Environment	D (d) / terminal (CU/kWh)	EW half-life (d) / gain	Fixed margin	Adaptive q / rho
A	30 / 0.14	14 / 1	0	0.5 / 0.1
B	60 / 0.08	30 / 0	0	0.5 / 0.1
C	60 / 0.08	30 / 0	0	0.5 / 0.1

Validation comprises 108 original four-cell runs and 84 margin runs. The distinct weighted-mean accuracy grid and tie rules are specified in Section S1.2. Complete weighted-mean scores (results/ew_validation_scores.csv) and frozen parameter selections (selected_parameters_LOCKED.json) retain the outcome of every selection. Joint source-to-target transfer freezes D and ν; weighted-mean fields are present in stored tuples but inactive in the median/uncalibrated transfer runs.

Table S6. Observed-future-information references under S0. Cost saving is relative to the uncalibrated cost under the same execution rule. Valid future truth replaces the causal input only where it exists; elsewhere causal fallback remains. These references retain the finite horizon and execution restrictions and are neither deployable policies nor guaranteed global cost bounds. Cost savings are given in CU with percentages in parentheses; truth-target coverage remains a percentage.

Environment	Feedback alpha	Truth targets (%)	Cost saving CU (%)
A	0.0000	99.9167	25287.004 (1.7550%)
A	1.0000	99.9167	2555.006 (0.1807%)
B	0.0000	99.8755	36.26657 (11.3701%)
B	1.0000	99.8755	19.65812 (6.4029%)
C	0.0000	99.8549	28.93997 (9.1103%)
C	1.0000	99.8549	13.80833 (4.5152%)

S4. Complete operating contrasts and action records

Supplementary Data S1, primary_operating_metrics.csv, contains all 81 primary policies with cash components, import/export, curtailment, charge/discharge, throughput, equivalent full cycles, peak import and SOC-bound exposure. Supplementary Data S2, primary_normalized_metrics.csv, provides CU per MWh served load. Primary effect curves (analysis/primary_effect_curves.csv), interaction effects (analysis/interaction_effects.csv) and all 265 summaries (results/all_test_summary.csv) retain the paired effects and supplementary configurations. No pooled site average replaces environment-specific results.

Table S7. S0 cash costs, including the discharged-energy proxy and excluding planning curtailment preference and terminal energy value.

Environment	Uncalibrated planned (CU)	Median planned (CU)	Uncalibrated feedback (CU)	Median feedback (CU)
A	1,440,817.348	1,445,112.527	1,413,690.258	1,414,362.573
B	318.96534	319.37322	307.01728	307.14580
C	317.66317	318.30546	305.81604	306.08696

Table S8. Median-calibration cost effects and interactions. Positive effects mean higher cost. Costs and intervals are in CU, with percentages in parentheses. Calibration percentages use each execution cell's own uncalibrated cost; the interaction uses the uncalibrated planned cost. Intervals use seven-day circular blocks conditional on observed segments. Subtracting the two calibration percentages does not recover the defined interaction. These test windows had been inspected before the revision protocol was frozen; the intervals do not adjust for that prior inspection.

Environment / scenario	Planned calibration effect CU (%)	Calibration effect under feedback CU (%)	Interaction CU (%)	Conditional 95% interval CU (%)
A / S0	+4295.179 (+0.298%)	+672.315 (+0.048%)	−3622.864 (-0.251%)	[−6031.847, −1600.115] ([-0.419%, -0.111%])
A / S1	+2786.546 (+0.187%)	+822.988 (+0.057%)	−1963.559 (-0.132%)	[−3919.563, −273.118] ([-0.264%, -0.018%])

Environment / scenario	Planned calibration effect CU (%)	Calibration effect under feedback CU (%)	Interaction CU (%)	Conditional 95% interval CU (%)
A / S2	+8282.722 (+0.459%)	+1778.705 (+0.102%)	−6504.017 (-0.361%)	[−11081.923, −2619.284] ([-0.614%, -0.145%])
B / S0	+0.40788 (+0.128%)	+0.12852 (+0.042%)	−0.27936 (-0.088%)	[−0.55989, −0.02126] ([-0.176%, -0.007%])
B / S1	+0.32659 (+0.082%)	+0.09796 (+0.026%)	−0.22863 (-0.057%)	[−0.54684, +0.07166] ([-0.137%, +0.018%])
B / S2	+0.64660 (+0.150%)	+0.01964 (+0.005%)	−0.62696 (-0.146%)	[−1.14948, −0.15869] ([-0.267%, -0.037%])
C / S0	+0.64229 (+0.202%)	+0.27092 (+0.089%)	−0.37137 (-0.117%)	[−0.50817, −0.22680] ([-0.160%, -0.071%])
C / S1	+0.61347 (+0.163%)	+0.25727 (+0.072%)	−0.35619 (-0.094%)	[−0.53588, −0.16985] ([-0.142%, -0.045%])
C / S2	+1.15328 (+0.283%)	+0.49860 (+0.129%)	−0.65468 (-0.161%)	[−0.91462, −0.37624] ([-0.225%, -0.092%])

Table S9. Stagewise L1 charge/discharge-energy differences under S0. Units are kWh; α identifies execution intensity. The final ratio is descriptive, may exceed one and is undefined for zero planned difference.

Environment / alpha	Plan (kWh)	Request (kWh)	Realized (kWh)	Energy ATR
A / 0.0	325,203.99	325,203.98	196,218.13	0.603
A / 0.5	235,397.76	164,395.19	99,693.76	0.424
A / 1.0	199,576.88	97,326.80	62,384.22	0.313
B / 0.0	140.19	140.19	85.99	0.613
B / 0.5	102.35	95.35	60.42	0.590
B / 1.0	95.93	87.70	52.95	0.552
C / 0.0	117.84	117.84	65.57	0.556
C / 0.5	96.45	81.18	45.63	0.473
C / 1.0	85.83	66.72	38.20	0.445

Under full feedback, identical requests in the shared non-discharge branch still leave realized differences of 12,778.02, 13.44 and 8.57 kWh in A/B/C through different accumulated energy headroom. Both request and realized differences in that branch are exactly zero in the common-state shadows. A's 191 hours with different discharge-gate states contribute 437.844 CU of the 672.315 CU feedback calibration penalty. Other gate groups close the remaining ledger; they are not independent causal mediators.

Table S10. Full-trajectory and common-state action contrasts under full feedback. Every shadow hour is replanned from SOC on the uncalibrated planned trajectory. Its plan denominator and gate membership can change, so differences between ratios are not an SOC causal share or attenuation of identical plans.

Environment	Trajectory	Plan (kWh)	Realized (kWh)	Energy ATR
A	Full trajectory	199576.8833	62384.2218	0.3126
A	Common-state shadow	282461.4385	22131.8244	0.0784
B	Full trajectory	95.9321	52.9468	0.5519
B	Common-state shadow	124.5712	34.4812	0.2768
C	Full trajectory	85.8302	38.2048	0.4451
C	Common-state shadow	100.1087	29.0487	0.2902

The complete transmission data (analysis/transmission.csv), shadow records (analysis/shadow_transmission.csv), gate contrasts (analysis/gate_transmission.csv), constraint combinations (analysis/constraint_components.csv) and cash-ledger closure (analysis/cost_component_closure.csv) preserve all stages. Individual binding counts may overlap, whereas combination groups partition intervals after reconstructing tight limits with positive clipping.

S5. Conditional uncertainty, fragmentation and regimes

Table S11. Variable and fixed day counts for the S0 median interaction. Intervals are nominal 95% conditional resampling bands in CU, not validated seasonal coverage or uncertainty across independent sites. The test windows had been inspected before the revision protocol was frozen; these bands do not adjust for prior inspection.

Environment / block (d)	Total days	Variable / fixed days	Lower (CU)	Upper (CU)
A / 3	275	275 / 0	-5586.5545	-1753.0285
A / 7	275	275 / 0	-6031.8472	-1600.1146
A / 14	275	275 / 0	-6389.1972	-1240.0455
B / 3	368	367 / 1	-0.5557	0.0026
B / 7	368	362 / 6	-0.5599	-0.0213
B / 14	368	362 / 6	-0.5669	-0.0093
C / 3	213	123 / 90	-0.5580	-0.1611
C / 7	213	59 / 154	-0.5082	-0.2268
C / 14	213	59 / 154	-0.4634	-0.2837

The test windows had been inspected before the revision protocol was frozen. The reported bands quantify conditional variation within previously examined data and do not adjust for prior inspection. The conditional circular algorithm is given in Section S1.7. A three-day block leaves B's S0 interval crossing zero, while its seven-day interval is below zero. C's seven-day procedure varies only 59 of 213 days; 154 days contribute fixed segment sums. Those intervals do not cover missing periods, segment occurrence or independent-site uncertainty. Daily state-reset boundaries are not additional observation gaps.

The complete bootstrap record (analysis/bootstrap_intervals.csv), monthly contributions (analysis/monthly_effects.csv), episode contributions (analysis/episode_effects.csv) and regime counts/effects (analysis/regimes.csv) retain all block lengths and occupied groups. Error, surplus and native within-hour ramp cutoffs are training medians. Zero-PV days retain undefined surplus groups; the near-bound SOC fraction is post-treatment context, not an exogenous exposure. No influential day, unfavorable season or adverse intervention is removed.

The exploratory forecast-error cost elasticity (FECE) is omitted when first-hour relative net-load MAE improvement is below 1% or uncalibrated cost is nonpositive. Otherwise it divides relative cash saving by relative MAE improvement and retains negative values when accuracy improves but cost worsens. It is neither an efficiency nor a physical constant and supports no main contribution. Source values and exclusion reasons (supplementary/FECE_exploratory.csv) remain available.

S6. Extreme events, transfer and secondary boundaries

S6.1. Prespecified maximum and minimum event days

Both daily interaction extrema are selected over all retained test days and remain in aggregate totals. Preceding context is shown only where observed; C's minimum day begins a new episode and has no preceding observed 24 hours. The plots explain contributions rather than estimating a representative day.

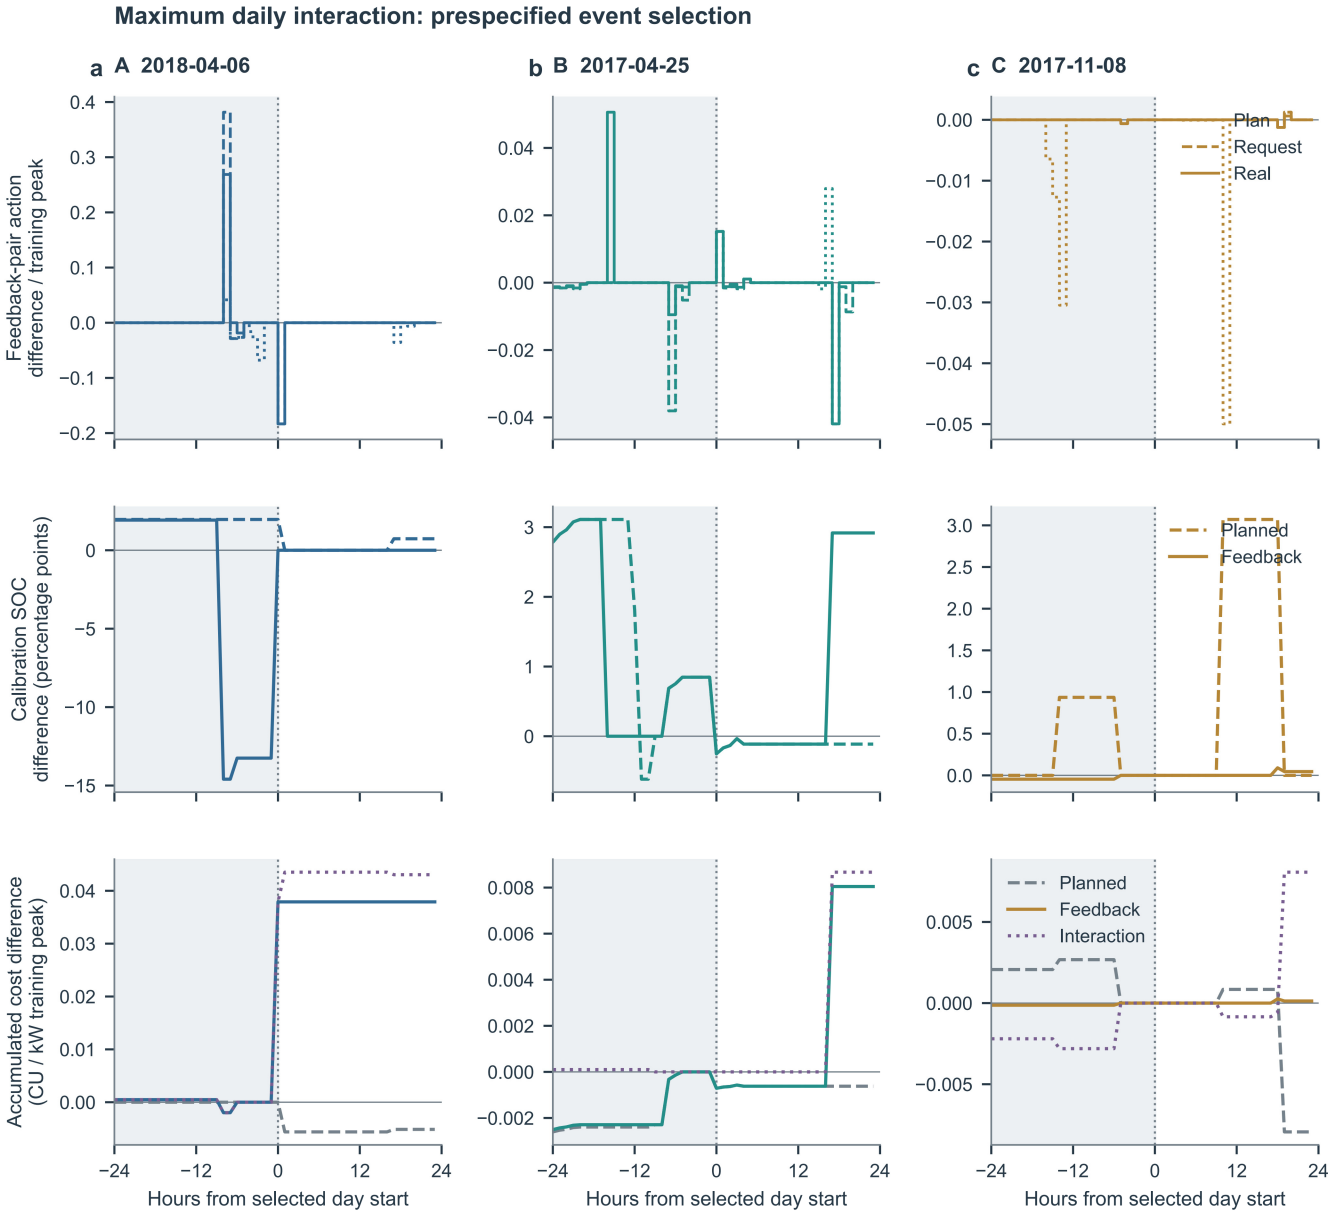


Figure S1. Maximum daily S0 median-calibration interaction in each environment, selected over all retained test days without excluding negative events. Top: feedback-pair signed discharge-minus-charge action differences at plan, request and realized stages, normalized by the training load peak. Middle: calibrated-minus-uncalibrated SOC differences under planned and feedback execution. Bottom: corresponding cumulative cost differences and their interaction, divided by the training load peak; cumulative values are anchored to zero at the selected-day boundary. Grey background denotes up to 24 hours of preceding observations. Missing history is not interpolated and may belong to a different initialized episode, as recorded in events.csv. These extreme events explain contributions and are not representative effect estimates. The paired minimum/maximum counterpart is retained in the companion figure and all events remain in aggregate analysis.

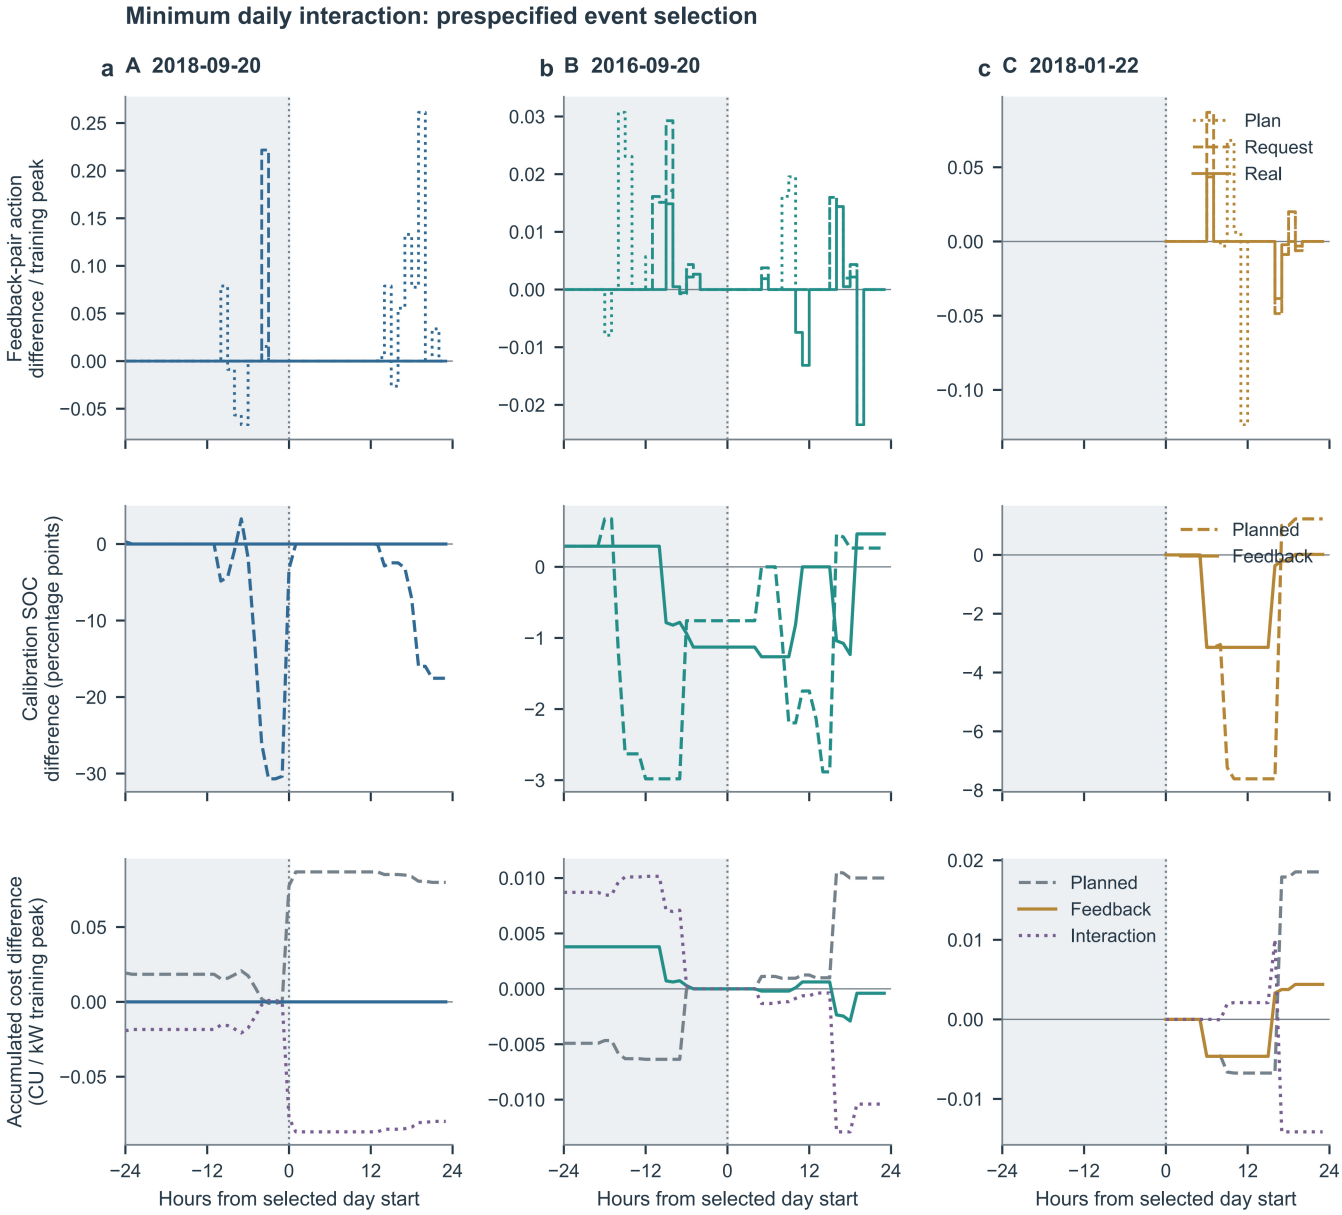


Figure S2. Minimum daily S0 median-calibration interaction in each environment, selected over all retained test days without excluding negative events. Top: feedback-pair signed discharge-minus-charge action differences at plan, request and realized stages, normalized by the training load peak. Middle: calibrated-minus-uncalibrated SOC differences under planned and feedback execution. Bottom: corresponding cumulative cost differences and their interaction, divided by the training load peak; cumulative values are anchored to zero at the selected-day boundary. Grey background denotes up to 24 hours of preceding observations. Missing history is not interpolated and may belong to a different initialized episode, as recorded in events.csv. These extreme events explain contributions and are not representative effect estimates. The paired minimum/maximum counterpart is retained in the companion figure and all events remain in aggregate analysis.

The event selection table (analysis/events.csv) and hourly event traces (analysis/event_hourly.csv) retain all extrema, episode identities and observed prehistory. Signed discharge-minus-charge is used only to display event paths; the ATR statistic continues to use the two separate charge/discharge energy components in Equation S17.

S6.2. Joint controller-setting transfer

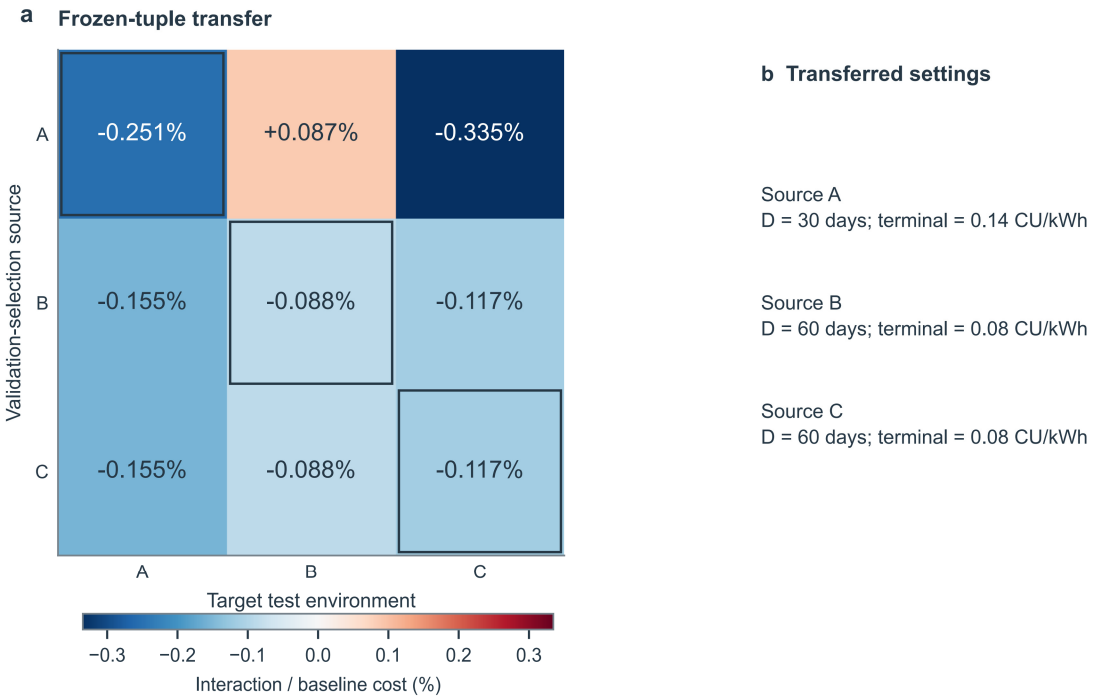


Figure S3. Frozen validation-selected median-window and terminal-value transfer under S0. Rows identify the selection source and columns the target environment; outlined diagonal cells use local validation selection. Each cell uses four matched policies and its own uncalibrated planned cost for normalization. B/C select identical tuples and give duplicated rows, not independent evidence of learned transfer. A-selected settings reverse B's interaction sign. Values are descriptive totals; intervals remain in the analysis data.

Transferring A's tuple to B changes the interaction from −0.27936 CU (−0.0876%) to +0.28135 CU (+0.0868%). The transferred interaction has a conditional 95% interval of [−0.19502, +0.86107] CU ([−0.0601%, +0.2656%]), spanning zero. Median calibration then reduces cost relative to the corresponding transferred-setting uncalibrated references, while the transferred policies still cost more than the locally selected counterparts. Both residual window and terminal value change, including the valuation used by uncalibrated cells. The transfer audit (analysis/transfer_audit.csv) records this joint-setting intervention; it is not isolated forecast-model transfer. B-to-C and C-to-B have no discrimination power because their selected tuples coincide. No target is retuned after inspection.

S6.3. Storage, tariff, margin and numerical outcomes

The C double-storage diagnostic changes both energy and power at fixed four-hour rated duration. Its interaction changes from −0.37137 to +0.37041 CU (+0.1242%); the terminal-adjusted interaction is +0.43410 CU. Purchase, export-revenue and proxy contributions are +0.44744, −0.04217 and −0.03485 CU. Total imports decrease by 1.340 kWh under planned execution and 0.289 kWh under full feedback, yet purchase costs rise by 0.1937 and 0.6411 CU. Peak-price purchase penalties rise from 0.48169 to 1.02361 CU; different-gate penalties rise from 0.02481 to 0.23263 CU and both-discharge penalties from 0.22487 to 0.35148 CU. Energy ATR still falls from 0.521 to 0.409. Transaction timing therefore matters to the economic sign even when total imports and action differences decline.

A's flat-price interaction is +338.080 CU (+0.0210%), with a seven-day interval of [−407.095, +1034.618] CU ([−0.0253%, +0.0643%]). Contributions of +942.175 CU in

purchases, −380.890 CU from export revenue and −223.206 CU in proxy cost nearly cancel. Energy ATR changes from 0.690 to 0.700, and the both-discharge calibration penalty rises from 3154.46 to 3947.61 CU. This evaluated unresolved aggregate effect does not establish that flat tariffs always eliminate forecast value.

Zero fixed margins reproduce calibrated no-margin outcomes. Adaptive margins change planned-execution cost by +1007.437, +1.375 and +0.615 CU in A/B/C and full-feedback cost by −28.849, −0.069 and −0.331 CU. The mixed effects do not support universal adaptive-margin superiority. Cold-hourly and retained-boundary-basis checks reproduce primary interactions to reported precision; advancing instead of deferring current throughput changes B's interaction to +0.03716 CU (+0.0121%) and makes all three S0 seven-day intervals cross zero. This is action-selection dependence within the declared economic budget. C's flat-price cash interaction is positive while the terminal-adjusted interaction is negative, preserving inventory-valuation sensitivity.

All values remain in the diagnostic comparisons (analysis/diagnostic_comparisons.csv), numerical-policy audit (analysis/numerical_policy_audit.csv), matched reversal ledgers (analysis/reversal_matched_operating_deltas.csv) and gate/tariff components (analysis/gate_tariff_components.csv). The complete frozen 265-configuration matrix (test_matrix_LOCKED.json) also retains native-PV, daily-reset, relaxed-export and throughput-proxy variants; these are scenarios, not new datasets.

S7. Traceability and reproducibility scope

Python 3.12.14 and HiGHS 1.15.1 on Windows 11 produced the recorded runs. The final matrix contains 1,804,800 hourly plans across 265 policies, with two optimization stages per plan, plus 12 common-state shadows. Saved-result accounting independently recomputes 11,971,200 native intervals. A separate rerun of all four C/S0 endpoint policies solves 20,448 hourly plans and reproduces saved core outcomes within declared tolerances. Removing future truth and validity indicators preserves current forecasts and first actions in 28 sampled decisions, supplemented by archive-boundary fallback checks. Saved-CSV recomputation and optimizer reruns are distinct checks; neither validates meter wiring or hardware dynamics.

The release README (README.md) identifies the supported fresh-output workflow and the source-preprocessing boundary. This existing revision began with preserved prepared inputs and reran validation and dispatch; raw historical preparation was retained as provenance, not recertified as original telemetry. Source input documentation (source_inputs/README.md), the data-quality audit (source_inputs/External_Data_Quality_Audit.md), frozen protocol (PROTOCOL.md), pre-validation hashes (pre_validation_lock.json) and analysis manifest (analysis/analysis_manifest.json) identify the underlying evidence. Full-replay verification is recorded in C/S0 rerun verification (verification/fresh_core_C/verification.json).

Historical material indexing (historical_materials.json) retains the earlier simplified hourly study. The numerical amendment (audit/NUMERICAL_AMENDMENT.md) documents the secondary-presolve failure and uniform repair; the failed attempt remains archived. Current results use the amended engine identified in the validation lock.

The public source datasets remain subject to CC BY 4.0 attribution [37,38]. The derived data, source code and supporting records, including Supplementary Data S1 and S2, are not publicly archived or supplied as separate attachments. They are available from the first author, Xuan-You Liu (liuxuanyou@cu.ac.kr), upon request. This supplement reports existing analyses and does not add experiments or certify original telemetry.