Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Budgeted Day-Ahead Reserve Allocation from Conformal Forecasts under Delayed Feedback

Ran Wang

wangran@hdu.edu.cn

Hangzhou Dianzi University

Research Article

Keywords: Conformal prediction, Delayed feedback, Budgeted decision making, Day-ahead load forecasting, Reserve allocation

Posted Date: November 9th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11363552/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Budgeted Day-Ahead Reserve Allocation from Conformal Forecasts under Delayed Feedback

Ran Wang[1*]

[1*]School of Computer Science and Technology, Hangzhou Dianzi University, No. 1158, 2nd Street, Baiyang Subdistrict, Qiantang District, Hangzhou, 310018, Zhejiang, China.

Corresponding author(s). E-mail(s): wangran@hdu.edu.cn;

Abstract

Forecast-driven intelligent systems must allocate limited interventions before the outcomes needed for feedback become available. We study this problem through day-ahead reserve allocation from multivariate conformal load forecasts under a fixed annual budget of committed days. A shortfall-sum rule uses released coverage feedback to place commitments, while a calendar-based contract enforces the two-day feedback delay, handles missing observations, and preserves the non-replenishable budget. Local deficit accounting gives a window bound with an explicit budget-blocking term and connects modeled coverage service to finite zonal adequacy. On seven annual New York origins, the rule raises the lowest seven-day mean modeled service from **0.582** to **0.713** and achieves positive paired mean effects against five comparators sharing the same annual cap, using **44.29** of the **55** available commitment days per year. With extra zonal capacity set to the reference-year 95th percentile of positive corrected residuals, the same action path reduces annual uncovered energy by **25.9%** and the worst-seven-day deficit by **14.8%**, while booking **1.12** million MW·h of additional capacity-time per year. Response and price sensitivity quantify the capacity-price range in which the energy reduction justifies its reservation. The framework links calibrated forecasting uncertainty to reproducible budgeted decisions and measures their local service, physical protection, and capacity cost.

Keywords: Conformal prediction, Delayed feedback, Budgeted decision making, Day-ahead load forecasting, Reserve allocation

1 Introduction

Forecast-driven intelligent systems must translate predictive uncertainty into timely actions under resource constraints. When outcomes arrive after a decision, the system must allocate its remaining resources using released feedback and protect service over the calendar. We study this problem in day-ahead load planning: an operator publishes a joint forecast region and decides which delivery days receive additional reserve capacity under a finite, nonreplenishable annual budget of committed days.

The official New York ISO day-ahead load forecast is published on the morning of day $t-1$ (New York ISO 2026). The commitment for delivery day t follows that publication, but the latest complete outcome is from $t-2$: day $t-1$ is still running. An outcome from day t therefore first informs the decision for $t+2$, a two-day feedback delay that is distinct from the one-day forecast horizon. We apply the same clock to forecast correction, calibration updates and commitment histories, isolating the allocation decision.

We evaluate a commitment on two complementary scales. *Modeled coverage service* assigns the base joint-region coverage to an uncommitted day and full service to a committed day, isolating placement on the forecast domain. *Finite zonal reserve response* instead adds a declared MW headroom to the issued forecast's residual requirement and measures uncovered energy, all-zone adequacy and capacity cost.

Worst-week protection requires attention to the timing of coverage shortfalls. Conformal prediction and its adaptive variants control marginal or time-aggregated miscoverage (Shafer and Vovk 2008; Gibbs and Candès 2021, 2024; Bhatnagar et al. 2023), yet shortfalls can cluster within a rolling window, and a global miscoverage level does not by itself protect a local minimum. A nonreplenishable commitment budget adds a temporal allocation decision to region calibration: which delivery days benefit most from additional protection?

Our contribution lies in the temporal allocation of uncertainty-informed actions and the accounting of their local effects. The framework combines an established learned forecast correction and online conformal calibration with the following elements:

- **A delayed-allocation contract.** We specify the day-ahead clock, the annual commitment cap, and the separate treatment of input gaps known before a decision and missing outcomes revealed later; spent commitments remain booked even when their days later prove unscorable.
- **Local deficit accounting.** A shortfall-sum rule bounds the modeled deficit of a same-segment window with an explicit term for budget blocking, and a coupling proposition relates the same action path to finite physical adequacy.
- **A matched evaluation.** Seven policies share forecasts, information and a cap on seven annual New York origins, reported with modeled service, finite reserve response, capacity and price sensitivity, paired annual effects and leave-one-origin-out checks.

The study quantifies the service and energy benefits of placement together with their capacity trade-offs. Sections 3–4 define the problem and policies; Section 5 gives the bounds; Sections 6–7 present the evaluation.

2 Related Work

2.1 Load forecasting and forecast-informed operation

Machine learning provides flexible demand predictors, including multiscale deep models for short-term load forecasting (Yang et al. 2024b), while sample-dependent prediction intervals quantify uncertainty in learned regression outputs (Akusok et al. 2019). Probabilistic load forecasting represents demand uncertainty through quantiles, predictive distributions or scenarios; day-ahead studies respect the market cutoff and evaluate quantile loss, interval coverage and width (Austnes et al. 2025). Ensemble peak-demand density forecasts have been applied to risk-aware power system scheduling (Yu and Tang 2025), and dynamic reserve is set within distributionally robust integrated-energy scheduling using adaptive copulas (Jiao et al. 2025); joint day-ahead demand and renewable forecasts similarly inform operating-reserve allocation (Terrén-Serrano et al. 2026). Adaptive conformal intervals and online quantile ensembles under distribution shift extend this line (Suresh et al. 2025; Qin et al. 2025). These studies establish the operational relevance of day-ahead uncertainty to scheduling, reserves and adequacy, the setting of this study. We focus on a finite sequence of commitment days, whose operational value depends on *when* the additional capacity is booked.

2.2 Conformal prediction under delay, and online allocation

Adaptive conformal inference and its dynamically tuned extension adjust calibration under distribution shift (Gibbs and Candès 2021, 2024); SAOCP uses strongly adaptive online learning for localized calibration (Bhatnagar et al. 2023), conformal PID uses feedback control (Angelopoulos et al. 2023), and conformal prediction has been applied to wind-speed forecasting (Zuege et al. 2025). For multiple outputs, ellipsoidal conformal methods distinguish joint-region constructions (Messoudi et al. 2022); the coordinate support of these joint regions gives the per-zone MW requirement in Section 5. OAS covariance regularization (Chen et al. 2010; Ledoit and Wolf 2004) and the online threshold algorithm are adopted as established components.

Delayed-feedback online learning studies decisions made before losses from earlier actions arrive (Joulani et al. 2013; Quanrud and Khashabi 2015), and simultaneous absolute coverage violation with prediction-set efficiency has been analyzed over time (Vaze 2026). The formulation studied here targets a complementary regime: complete daily feedback is released on a fixed delay, so a decision may use only outcomes whose release date has passed, and the resource is a nonreplenishable count of days rather than a per-round cost. Online allocation with *replenishable* budgets (Yang et al. 2024a) permits top-ups and targets cumulative utility with competitive ratios; the contract developed here couples a nonreplenishable day-level budget to a local-window service objective with explicit delayed-feedback accounting.

2.3 Decision-focused conformal work and the switching threshold

Conformal uncertainty sets for robust optimization (Johnstone and Cox 2021) and conformal risk control (Angelopoulos et al. 2024) connect uncertainty with decisions

in other ways. Dronavajjala (2026) characterizes when calibrated uncertainty changes a nominal robust optimization decision, with a switching threshold, a pre-deployment safety margin and an energy dispatch example under real market prices; that analysis concerns the geometry of a single-period decision. The regime studied here is temporal rather than single-period: the decision problem is itself budgeted and repeated, the resource is a nonreplenishable count of days, and the feedback that would inform the next decision arrives two days late. Sections 3 and 5 therefore formalize the allocation contract for a finite, delayed, day-level commitment budget whose endpoint is a local window minimum and whose physical response is finite.

3 The Delayed Day-Ahead Allocation Problem

3.1 Decision clock and declared calendar

The declared calendar $\mathcal{D}$ runs from January 2 to December 31, with $T_{\text{planned}} = |\mathcal{D}| \in \{364, 365\}$ days. Day t contains $H_t \in \{23, 24, 25\}$ real hours; both folds of an autumn repeated hour share the day's frozen slot threshold and each contributes one real hour.

The decision for t follows the morning day-ahead forecast release on $t-1$ (New York ISO 2026) and precedes delivery. Only outcomes through $t-2$ are complete. Point training uses the conservative boundary of midnight starting $t-1$, excluding partial day-$t-1$ labels. Inputs are the published forecast, realized-error lags of 72 and 168 elapsed hours, local-clock terms and released daily feedback; lag provenance requires the source instant to be exactly $\tau-72$ or $\tau-168$ on the UTC timeline and its local date to be at least two or six days old, respectively. The archive begins in 2017, which supplies early point-training history; covariance and threshold calibration use the exact prior calendar year $Y-1$.

Base daily coverage is

$$q_t = \frac{1}{H_t} \sum_{h=1}^{H_t} \mathbf{1}\{y_{t,h} \in C_{t,h}\} \in [0, 1].$$

The earlier daily-vector benchmarks instead use a single joint-coverage indicator $q_t \in \{0, 1\}$; their service units are kept separate. Complete day-t feedback is released for the decision at $t+2$, giving lag $\ell = 2$. An *input gap*, known at the cutoff, suppresses that date's issuance and resets the segment immediately. A *missing outcome* for an issued day is discovered at its release date and resets the segment then. A reset discards old-segment pending feedback but returns no budget, and pending records retain their original segment tags. The distinction between the two events matters operationally: if the forecast inputs for Wednesday are absent, that absence is known when Wednesday's commitment would be issued; if Wednesday is issued normally but its actual load is incomplete, its commitment has already been spent, and only Friday's feedback release reveals the missing outcome. Only completely observed, issuable days enqueue daily feedback.

Let $\mathcal{M} \subseteq \mathcal{D}$ contain the complete scored days. Modeled service and physical window metrics use $\mathcal{M}$; annual uncovered energy also includes finite issued zone-hours

from partially observed days; and capacity is booked for every issued hour of a committed day, including hours whose actual later proves missing. Missing values are never imputed, and finite-support shares are retained in the artifacts.

3.2 Commitments, extra capacity, and the declared budget

The binary action $a_t = 1$ commits extra zonal capacity ηH_z throughout day t. Every policy already books the same region-implied base requirement $r^0_{t,h,z}$, to which a commitment adds. The annual cap is fixed before evaluation support is read:

$$\sum_{t \in \mathcal{D}} a_t \le B, \qquad B = \mathrm{round}(0.15\, T_{\mathrm{planned}}) = 55. \tag{1}$$

The resource counts committed days, not generated electricity. A commitment consumes one unit even if its day later becomes unscorable, the cap never replenishes, including at resets, and a trigger refused because the cap is exhausted is recorded as blocked. Policies share the cap, but their executed counts can differ: a pacing rule can consume its units steadily, whereas a reactive rule can leave units unused when recent service is adequate. We therefore report executed commitments and capacity-time alongside the service endpoint, since an improvement achieved with more capacity is a different operating trade-off from an improvement obtained by placing fewer commitments more effectively.

3.3 Modeled coverage service

A committed day scores full modeled service; otherwise it scores base coverage:

$$\begin{aligned} s_t &= a_t + (1 - a_t) q_t \in [0, 1], \\ d_t &= 1 - s_t = (1 - a_t)(1 - q_t). \end{aligned} \tag{2}$$

Higher service is better. This full-service convention isolates commitment placement; Section 5.3 accounts for finite physical response. Daily rescue is $s_t - q_t = a_t(1 - q_t)$, positive only if a committed day has $q_t < 1$. With $T = |\mathcal{M}|$, the modeled service and action loss is

$$L_\lambda = \frac{1}{T} \sum_{t \in \mathcal{M}} (1 - s_t) + \lambda \frac{1}{T} \sum_{t \in \mathcal{M}} a_t, \qquad \lambda \ge 0, \tag{3}$$

Here λ is a shortfall-scale sensitivity weight, not a capacity price. Both sums use scored days, whereas the cap and annual commitment counts use $\mathcal{D}$; thus L_1 is not obtained by adding the all-declared commitment rate to one minus mean service.

3.4 Finite zonal reserve response

The same actions are evaluated on corrected residuals $e_{t,h,z}$, defined as actual load minus issued point in MW. The finite capacity response of Section 4.4 yields uncovered energy in MWh, one-sided all-zone adequacy and booked capacity-time in MW·h.

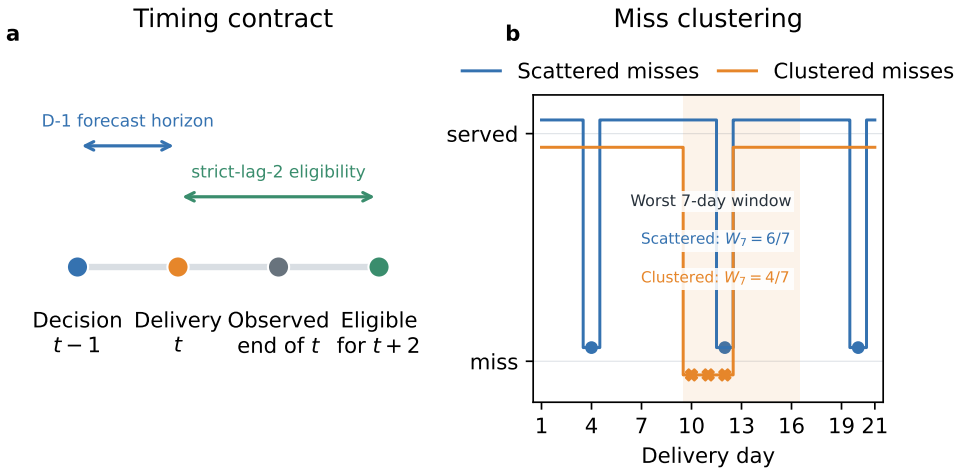


Fig. 1 Day-ahead timing and illustrative worst-window modeled service. (a) The commitment for delivery day t is made after the day-ahead forecast is published on $t-1$ and before delivery; the completed day-t outcome first becomes eligible for the decision for delivery $t+2$, while the implementation stops training labels at midnight at the start of $t-1$. (b) Two illustrative binary streams with three zero-service days and equal mean service have minimum seven-day means 6/7 and 4/7: clustering, not the count of misses, sets the worst window.

Lower uncovered energy and deficit, and higher adequacy, are better. Prices c_R per MW·h and c_D per MWh give

$$J = c_R(\text{base booked} + \text{extra booked}) + c_D \text{ uncovered}.$$

Equation (9) normalizes this comparison. The response is a zonal balancing model, without network flows or market dispatch.

3.5 Window endpoints and their orientation

Let $\mathcal{S}$ partition the calendar into uninterrupted segments. An admissible window contains consecutive delivery dates within one segment, all belonging to $\mathcal{M}$. The modeled endpoint is

$$W_7 = \min_{\substack{S\in\mathcal{S},\ w:\ [w,w+6]\subseteq S \\ \text{and } w,\ldots,w+6\in\mathcal{M}}} \frac{1}{7}\sum_{i=0}^{6} s_{w+i}, \tag{4}$$

Define W_3 and W_{14} analogously; higher W_n is better. Physical worst-seven-day deficit is the maximum summed deficit over such windows, and adequacy is their minimum mean daily adequacy. Windows bridge neither gaps nor resets, including dates whose missing outcome is discovered later.

Figure 1 illustrates why placement matters: two 21-day binary streams with three zero-service days have equal means, but dispersed misses yield $W_7 = 6/7$ and consecutive misses $W_7 = 4/7$.

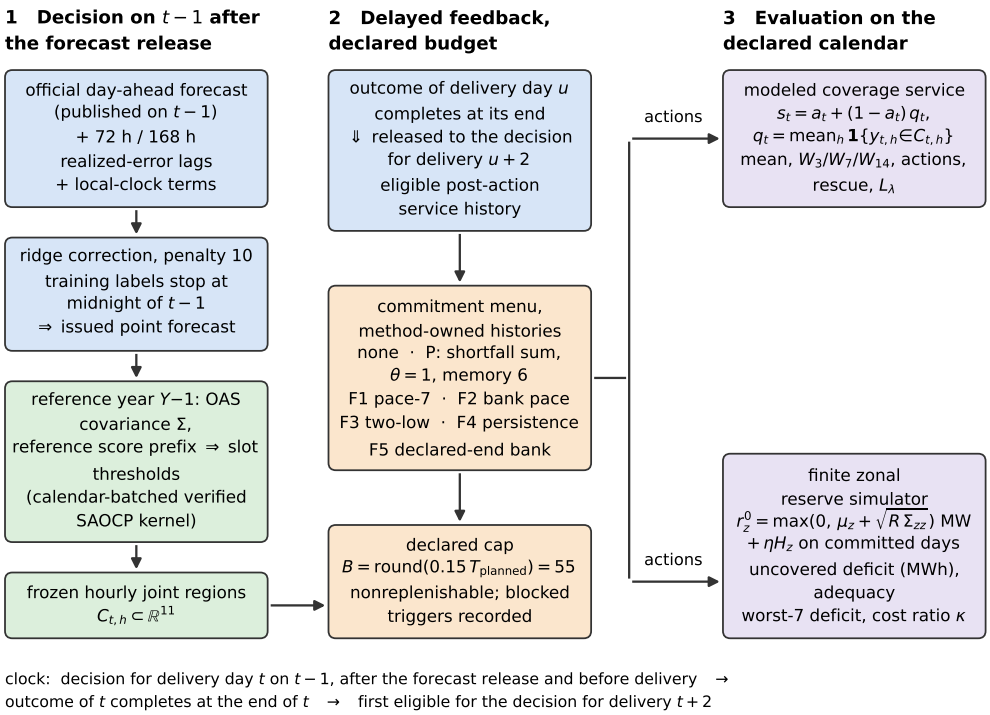


Fig. 2 Design of the end-to-end study. Column 1: everything that is known at the day-ahead cutoff produces the frozen hourly joint regions. Column 2: the completed outcome of day u is released at $u + 2$ and drives a declared-budget commitment menu. Column 3: the same action path is scored on the forecast domain and in the finite zonal reserve accounting. The schematic summarizes the decision and evaluation workflow.

4 Method

Figure 2 connects date-available forecast regions, a commitment decision based on released feedback, and two evaluations of the resulting action path. This section states the implementable rules; Online Resource 1, Section A gives complete event-handling pseudocode.

4.1 Date-available regions and thresholds

Issued point forecast.

The issued point equals the official day-ahead forecast plus a ridge correction of its realized error. Features are the slot's official forecast, 72/168-hour error lags for eleven zones, and sine/cosine terms for local hour and weekday. Training uses the 365 local days before midnight starting $t - 1$, z-scored features, an unpenalized intercept, penalty 10, and at least 1000 eligible rows; non-finite training rows are excluded. The January 2 decision can therefore use December 31 labels. Issuance depends on finite

pre-action features, with actual-load availability revealed later under the event clock of Section 3.1.

All policies share the point model and the exact prior reference year; older history supplies point training only.

Reference geometry.

Corrected residuals from the exact prior year $Y-1$ define the reference mean $\widehat{\mu}$ and OAS covariance $\widehat{\Sigma}$ (Chen et al. 2010). With squared-Mahalanobis score $z(e) = (e-\widehat{\mu})^{\intercal}\widehat{\Sigma}^{-1}(e-\widehat{\mu})$, the residual ellipsoid is $E_{t,h} = \{e : z(e) \le R_{t,h}\}$ (Xu et al. 2024), and the published load region is $C_{t,h} = \widehat{y}_{t,h} + E_{t,h}$. Reserve requirements below use the upper support of this *residual* ellipsoid, with the issued point already accounted for in the residual.

Online thresholds.

A calendar-batched adapter runs one pinned SAOCP kernel (Bhatnagar et al. 2023) per local hour slot, with nominal coverage 0.90 and lifetime parameter 8. Each is warm-started with its finite reference-year prefix and frozen scale $\max|\text{finite reference prefix}|\sqrt{3}$; empty or non-finite prefixes are refused. At decision d, same-segment scores from dates at most $d-2$ are released chronologically before freezing that day's thresholds. Autumn folds share a frozen slot threshold; partial days enqueue no scores; and resets discard pending records tagged with the old segment. The numerical kernel is extracted by abstract syntax tree from the pinned public source and byte-verified. The verification covers source identity and the adapter's delayed interface; Online Resource 1, Section H.4 documents the runtime-parity scope.

4.2 Budgeted commitment policies

All seven policies share base regions, released information, resets and the annual cap. Each owns its service history and walks the calendar once using parameters fixed before evaluation.

F1–F2 allocate by time, F3 counts low-service days, F4 uses the latest low-service day, and F5 defers P's credits. Each updates its own service history from its own actions; the no-commitment policy sets $a_t = 0$.

P (production).

Let $\mathcal{E}_t$ contain eligible same-segment dates at most $t-2$ and $\mathcal{L}_6(\mathcal{E}_t)$ its six most recent records. P commits when $h_t = \sum_{\tau\in\mathcal{L}_6(\mathcal{E}_t)}(1-s_\tau) \ge \theta = 1$ and budget remains. Partial shortfalls accumulate, so one full day of normalized shortfall in memory triggers. Memory $m = 6$ satisfies the seven-day bound's sufficient condition $m \ge n - \ell = 5$.

F1 (pace-7).

After a six-day offset within a segment, commit on every seventh accessible day, subject to the cap. This comparator uses calendar position only.

F2 (bank-pace).

Accumulate 15 credits per issuable day and commit when the bank reaches 100, subject to the cap: about one commitment every 6.7 days, without service feedback.

F3 (two-low-coverage).

Commit if at least two of the six most recent eligible service values are below 0.9. This trigger uses a count of low-service days, complementing P's shortfall-mass trigger.

F4 (persistence).

Commit when the latest eligible service is at most 0.5.

F5 (declared-end bank).

P's shadow path mints credits. Pre-action tension τ_d averages $-\log p$ across slots, where $p = (1+\#\{v \le \text{threshold}\})/(1+\#\{\text{history}\})$ ranks the frozen threshold against that slot's released threshold history. Commit a minted credit immediately if the regularized rank $\rho_d = (1+\#\{v \ge \tau_d\})/(1+N_d) \le 0.5$, where N_d counts prior segment tensions and small ranks indicate high tension, or if the declared end is within $B+1$ days. Otherwise bank it for the next issuable day. At most one credit is banked: an existing credit is committed on that next day regardless of tension, and any credit minted that day is retained in the bank after the old credit is spent. Resets clear both tension histories and the bank. This tests deferring P's commitments toward high-tension days; lost credits can reduce executed counts below B.

The shadow path remains P's own path; tension and the known calendar end use pre-decision information only. Online Resource 1, Section A states the exact regularized rank and transition order.

Every commitment consumes one unit even if its day later becomes unscorable. Unfunded triggers are recorded as blocked, or cap-skipped for pacing rules. Online Resource 1, Section A gives P's calendar pseudocode, all policy transitions and initial states.

4.3 Service, rescue, and modeled cost accounting

Table 1 reports declared-calendar commitment counts and rates, blocked days and rescue precision. Per-origin reproduction records also give scored-support counts, cap-skipped days, direct rescue counts and $\sum_{t\in\mathcal{M}:a_t=1}(1-q_t)$ (Online Resource 1, Section B.5). Rescue precision is the fraction of observed committed days with $q_t < 1$; it is the "Rescue" column in Table 1 and is undefined for no commitments. The declared modeled-loss weights are $\lambda \in \{0.25, 0.5, 1, 2, 4\}$. All policies share regions; operational size is measured by common base and extra booked capacity-time.

4.4 Finite zonal reserve response

For zone z and real hour h, compute the residual against the issued point:

$$e_{t,h,z} = \text{actual load}_{t,h,z} - \text{issued point}_{t,h,z} \quad \text{[MW]}, \tag{5}$$

The nonnegative upper support of the residual ellipsoid gives the base requirement, booked under every policy:

$$r^0_{t,h,z} = \max\left(0,\ \mu_z + \sqrt{\max(R\,\widehat{\Sigma}_{zz}, 0)}\right) \quad \text{[MW]}, \tag{6}$$

Here $R = R_{t,h}$ and $\mu_z = \widehat{\mu}_z$ from the reference geometry; the reference-year mean and covariance are fixed across issued hours. Committed days add ηH_z, with reference-only scale

$$H_z = Q_{0.95}(\max(\tilde{e}_z, 0)) \quad \text{[MW]}, \tag{7}$$

Here $\tilde{e}_z$ contains corrected residuals in $Y - 1$ and $Q_{0.95}$ uses linear interpolation. The multiplier $\eta \in \{0, 0.5, 1, 2\}$ sets capacity, not success probability. Thus $r_{t,h,z} = r^0_{t,h,z} + a_t \eta H_z$, and uncovered energy is

$$\Delta_{t,h} = \sum_z \int_h \max(e_{t,h,z} - r_{t,h,z}, 0)\, dt \quad \text{[MWh]}, \tag{8}$$

Each real hour contributes one hour, including both autumn folds. Annual energy integrates finite issued zone-hours; absent actuals contribute no deficit or adequacy observation. An hour is adequate when $e_{t,h,z} \le r_{t,h,z}$ for every zone, and daily adequacy is its mean on complete observed days. Worst-seven-day deficit sums daily uncovered energy over admissible windows, and the adequacy endpoint minimizes their mean daily adequacy, following Section 3.5. The maximizing deficit window and minimizing adequacy window are selected independently. Extra booked capacity is $a_t \eta \sum_z H_z$ times the day's issued hours, including hours with subsequently missing actuals; this MW·h quantity measures reserved capacity-time, while delivered protection is measured separately by uncovered energy and adequacy.

The capacity price c_R per MW·h and deficit price c_D per MWh define $\kappa = c_R/c_D$. Their normalized costs are

$$\begin{aligned} \frac{J_{\text{full}}}{c_D} &= \kappa\,(\text{base booked} + \text{extra booked}) \\ &\quad + \text{uncovered}, \\ \frac{J_{\text{inc}}}{c_D} &= \kappa\,\text{extra booked} + \text{uncovered}, \end{aligned} \tag{9}$$

in MWh-equivalent units. Common base booking cancels in the incremental comparison. Every policy is evaluated on the full declared nine-point κ grid (declared values in Online Resource 1, Table S2; cost cells for every policy, multiplier and origin in the machine-readable aggregate) and at every capacity multiplier.

4.5 AI-assisted preparation

During the preparation and revision of this work, the author used Claude Code and OpenAI Codex to assist with code development, manuscript drafting, literature

positioning, language editing, formatting, preparation of the study's schematic and data figures with programmatic vector and plotting code, consistency checks, and submission-file preparation. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

5 Local Deficit Accounting

The following results connect placement to modeled shortfall and finite adequacy. The propositions and operating implications are stated here, with complete proofs in Online Resource 1, Section C. Range checks and counterexamples are in Online Resource 1, Section H.8.

5.1 Zonal support of a residual ellipsoid

Proposition 1 (Residual support) *Let $E = \{e : (e-\mu)^{\mathsf{T}}\Sigma^{-1}(e-\mu) \le R\}$ with Σ positive definite and $R \ge 0$. For every $e \in E$ and every zone z,*

$$e_z \le \mu_z + \sqrt{R\,\Sigma_{zz}},$$

and the bound is attained on the boundary of E in the direction Σe_z. For $R < 0$ the set is empty and the clipped requirement below remains well defined.

Consequently, on an uncommitted hour joint-region coverage implies one-sided all-zone adequacy: every residual lies below its zonal upper support. Equation (6) clips that support to a nonnegative requirement.

5.2 Window bound for the shortfall-sum trigger

Proposition 2 (Window shortfall with blocked triggers) *Consider the shortfall-sum policy: state $h_t = \sum_{\tau \in \mathcal{L}_m(\mathcal{E}_t)} d_\tau$ over the m most recent eligible same-segment days, lag $\ell \ge 1$, threshold $\theta > 0$, memory $m \ge n - \ell$, and modeled full-service commitments, so that an executed commitment makes that day's shortfall zero. Let the window $w, \dots, w+n-1$ lie inside one segment, let $P = \{i : d_{w+i} > 0\}$ be its positive-shortfall days, and let $T = \{i \in P :$ the trigger on day $w+i$ was blocked$\}$ be those whose trigger the exhausted budget refused. Then*

$$\sum_{i \in P} d_{w+i} \;<\; \theta + \ell + \sum_{i \in T} d_{w+i} \;\le\; \theta + \ell + |T|. \tag{10}$$

The blocked set T is realized ex post: the bound accounts for blocking that actually occurred in the window rather than assuming that the budget suffices.

Corollary 3 (Window floor when nothing is blocked) *If no trigger in the window is blocked, then T is empty and $\sum_{i=0}^{n-1} d_{w+i} < \theta + \ell$. For the production parameters $(\ell, m, \theta) = (2, 6, 1)$ and a seven-day window this gives a total shortfall strictly below 3, hence a window mean service strictly greater than $4/7$ for every seven-day window inside one segment whose triggers are all served. If every seven-day window of the horizon satisfies that premise, then $W_7 > 4/7$.*

The proof separates the last unblocked positive-shortfall day from the blocked tail: the former contributes less than $\theta+\ell$, while the latter is charged to the window where blocking occurs (Online Resource 1, Section C).

5.3 Coupling to the finite zonal response

Proposition 4 (Physical adequacy shortfall of the same action path) *Fix a binary commitment path and let s_t^{phys} be the daily one-sided adequacy fraction of the reserve response and s_t the modeled service of Eq. (2). Then, for any window, the sum of the daily inadequacy fractions obeys*

$$\sum_{t\in w}\bigl(1-s_t^{\mathrm{phys}}\bigr) \;\le\; \sum_{t\in w}(1-s_t) \;+\; \sum_{t\in w} a_t\bigl(1-s_t^{\mathrm{phys}}\bigr).$$

On uncommitted days, adequacy is at least modeled coverage; on committed days, the final term is the remaining physical shortfall. Combining the propositions therefore bounds window inadequacy by $\theta+\ell$, realized blocked deficit and committed-day capacity shortfall. These are dimensionless adequacy fractions; uncovered MWh remains an empirical endpoint, and lower bounds clipped at zero use a non-strict inequality.

5.4 Assumptions

The window must lie inside one uninterrupted segment on complete scored support, memory must satisfy $m \ge n-\ell$, and executed commitments must give full *modeled* service. The blocked term is realized ex post, and the cap need not prevent blocking. Resets and missing support follow Section 3.1. Finite capacity is handled by the coupling term, not by changing P's feedback to physical deficit. The memory counterexample, reset and partial-response cases, and the conditional chance extension are retained in Online Resource 1, Section H.8.

6 Study Design

The study evaluates seven adjacent annual NYISO origins $Y \in \{2019,\dots,2025\}$, each with reference year $Y-1$ and evaluation dates January 2–December 31. The historical evaluation archive ends in December 2025. Before execution we froze the clock, training boundary, cap $B=55$, policy menu, response and price grids, endpoints, analysis rules and input identities.

Endpoints.

All seven policies, four capacity multipliers, nine price ratios and seven origins are reported. Forecast-domain endpoints are mean and $W_3/W_7/W_{14}$ service, base coverage, declared commitments and rates, blocking, rescue and L_λ. Physical endpoints are uncovered energy, all-zone adequacy, worst-seven-day deficit, booked base/extra capacity-time and normalized cost. W_3 and W_{14} check sensitivity to window scale, with the full grid retained.

Inference and diagnostics.

Paired annual mean differences use 10,000 IID-over-origin bootstrap draws (seed 20261011), leave-one-origin-out mean ranges and counts of strictly positive annual effects. The modeled material-effect benchmark is 0.01 on the W_7 scale. These adjacent, previously exposed historical origins provide descriptive paired evidence. Resampling annual paired effects preserves each year's policy contrast; leave-one-origin-out ranges and positive-effect counts characterize sensitivity to individual origins and consistency across years. The crossing price $\kappa^\star$ and stored-geometry summaries are post-run diagnostics of frozen outputs and do not alter the declared grids.

Execution and evidence retention.

Origins were admitted separately only after artifact-identity and endpoint-completeness checks passed; no origin was selected on endpoint values. Failed candidates and reasons remain archived. CPU-only execution used two numeric threads and averaged 25.3 seconds per origin (range 19.3–43.7, total 176.9), including all policies, reserve grids and saved artifacts but excluding setup, guards and audits; these are shared-server timings. Online Resource 1, Sections A–B and D–H retain the guard records and earlier replay and deep-model studies, whose supports and statistical families remain separate from this end-to-end evaluation.

7 Results

7.1 Modeled coverage service

Across the seven annual origins, all seven declared policies were run on identical regions, identical information, and the same declared cap. Table 1 reports the modeled endpoints and the executed commitments. The no-commitment policy attains an origin-averaged W_7 of 0.5816 and a mean service of 0.9045. The production rule raises W_7 to 0.7134 and mean service to 0.9159 using 44.29 committed days per year, a rate of 12.16% of the declared calendar under the $B = 55$ cap.

Calendar pacing (F1) and bank pacing (F2) use 51.57 and 53.71 days per year and achieve higher mean service (0.9195 and 0.9190), but P protects the worst week better with fewer commitments. Two-low coverage (F3) uses all 55 days yet leaves W_7 at the no-commitment value; P gains 0.1318 while spending 10.7 fewer days. Persistence (F4) reaches 0.6531 with 11.71 commitments. Declared-end banking (F5) reaches 0.7083 with the same executed count as P on these origins; only the cap is matched by design.

Figure 3 and Online Resource 1, Table S4 give paired annual contrasts. Every mean effect, bootstrap interval and leave-one-origin-out range is positive. Effects are +0.1318 against none and F3 (seven of seven origins positive), +0.0672 against F1 (six positive), +0.0723 against F2 (five positive), and +0.0604 against F4 (six positive). F5 is the closest comparator: +0.0051 [0.0009, 0.0119], three positive origins and four ties, below the 0.01 material-effect benchmark.

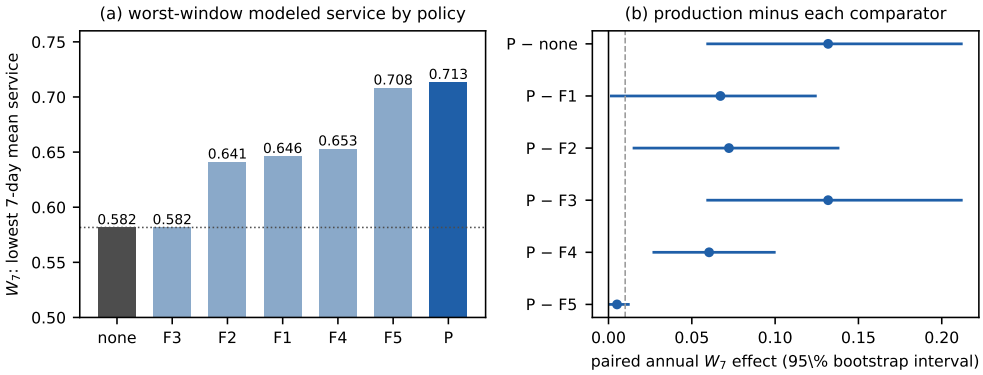


Fig. 3 Modeled worst-window service on the seven annual NYISO origins. (a) Origin-averaged W_7 by policy; the dotted line is the no-commitment value, and higher is better. (b) Paired annual W_7 effect of the production rule against each comparator with its 95% bootstrap interval; the dashed line marks the declared 0.01 material-effect benchmark.

Table 1 Modeled coverage service on the seven annual NYISO origins (2019–2025) for every declared policy. W_3 and W_7 are the lowest 3- and 7-day mean service inside one segment; actions and blocked days are annual means; the horizon cap is $B = 55$ committed days per origin.

Policy	Mean	W_3	W_7	Actions	Rate (%)	Blocked	Rescue	L_1
none	0.9045	0.3393	0.581633	0.00	0.00	0.00	–	0.0955
P (production)	0.9159	0.4087	0.713435	44.29	12.16	0.86	0.4679	0.2063
F1 pace-7	0.9195	0.4325	0.646259	51.57	14.16	0.00	0.5766	0.2228
F2 bank-pace	0.9190	0.3948	0.641156	53.71	14.75	0.00	0.5559	0.2292
F3 two-low	0.9176	0.3393	0.581633	55.00	15.10	113.57	0.5636	0.2342
F4 persistence	0.9077	0.4048	0.653061	11.71	3.22	0.00	0.4510	0.1246
F5 end-bank	0.9158	0.4087	0.708333	44.29	12.16	0.00	0.4558	0.2064

Service is the modeled coverage score of Eq. (2); actions are committed days counted over the declared calendar. "–" marks an undefined ratio (the no-commitment policy has no committed day).

7.2 Finite zonal reserve response

Table 2 evaluates the same actions with the declared finite response $\eta = 1$: each committed zone receives its reference-year 95th-percentile positive corrected residual as extra capacity. Online Resource 1, Section B reports the full $\eta \in \{0, 0.5, 1, 2\}$ grid.

At $\eta = 1$ the no-commitment policy leaves 1305.41 MWh of uncovered energy per year on average and records a worst-seven-day deficit of 606.27 MWh with a minimum seven-day mean adequacy of 0.9549. The production rule reduces annual uncovered energy to 966.98 MWh, a reduction of 338.43 MWh or 25.92%, and reduces the worst-seven-day deficit to 516.67 MWh, a reduction of 89.60 MWh or 14.78%, while raising minimum seven-day mean adequacy to 0.9592. Online Resource 1, Table S6 reports the paired contrasts: the annual uncovered reduction has interval [−868.05, −23.29]

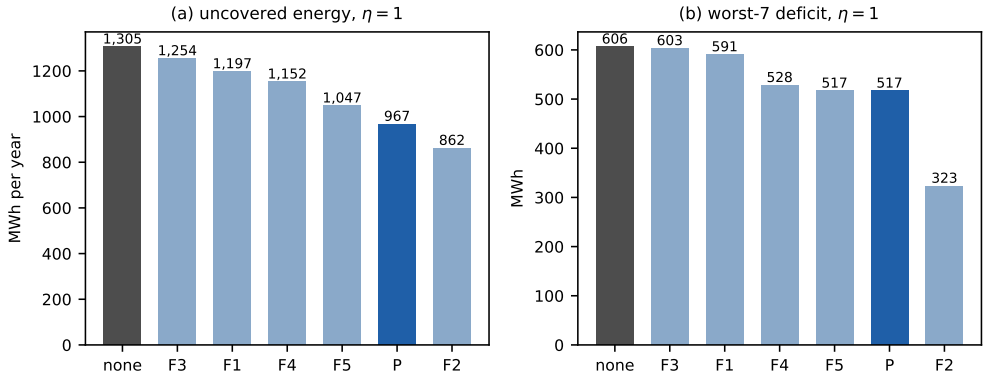


Fig. 4 Finite zonal reserve response at the declared multiplier $\eta = 1$, annual means over seven annual origins: (a) uncovered energy in MWh per year and (b) worst-seven-day deficit in MWh. Lower is better in both panels; the production rule is dark and the comparators light. Bank pacing (F2) shows the strongest worst-window protection with more booked capacity than the production rule. The two panels show how commitment placement and capacity booking shape the physical response.

MWh with a negative leave-one-origin-out range throughout, while the worst-seven-day deficit and adequacy intervals touch zero on their favorable side, so those two are reported as directional descriptive evidence. The corresponding extra booking is 1,123,611 MW·h per year of extra booked capacity-time, reported with its own interval in the same table, and Figure 4 shows the policy ordering on both physical endpoints.

Leave-one-origin-out evaluation retains energy benefits: excluding 2025 gives reductions of 7.82% in annual uncovered energy and 2.81% in the worst-seven-day deficit. The 2025 origin contributes 79.9% of total saved uncovered energy, reflecting protection on several high-deficit events. These percentages are ratios of the corresponding macro means.

Bank pacing (F2) provides the strongest physical worst-window protection: 323.16 MWh, using 53.71 commitments and 1.363 million MW·h/year. The P-minus-F2 deficit contrast is +193.51 MWh [12.03, 430.36], with leave-one-origin-out range [92.67, 225.76] (Online Resource 1, Table S6). Thus F2 buys about 194 MWh less worst-window deficit with 239,504 MW·h/year more capacity and lower modeled W_7 (0.6412 versus 0.7134). F5 matches P's physical worst-window endpoints in all seven origins; F3 books the most capacity but reduces that deficit only to 603.06 MWh. P's annual uncovered-energy reduction grows from 18.27% at $\eta = 0.5$ to 25.92% at $\eta = 1$ and 28.18% at $\eta = 2$; its worst-seven-day reduction stays at 14.78% between the latter two response sizes.

Table 2 Finite zonal reserve response at the declared multiplier $\eta = 1$ (dimensionless; it multiplies the dimensional reference headroom H_z in MW). Annual means over the seven origins; booked capacity is measured as capacity-time in MW·h.

Policy	Uncovered (MWh/yr)	W_7 deficit (MWh)	Min. 7-day adequacy	Extra booked (10^6 MW·h/yr)	Committed days/yr
none	1305.4	606.3	0.9549	0.0000	0.00
P (production)	967.0	516.7	0.9592	1.1236	44.29
F1 pace-7	1197.2	590.5	0.9592	1.3090	51.57
F2 bank-pace	861.6	323.2	0.9626	1.3631	53.71
F3 two-low	1254.2	603.1	0.9549	1.3947	55.00
F4 persistence	1152.1	528.0	0.9549	0.2994	11.71
F5 end-bank	1047.3	516.7	0.9592	1.1236	44.29

Uncovered energy integrates all finite issued zone-hour rows; the worst-window endpoints use complete observed days. Lower uncovered energy and lower deficit are better, higher adequacy is better. The no-commitment row books none of the extra capacity.

7.3 Capacity price

The two measurements meet in the normalized cost of Eq. (9), where $\kappa = c_R/c_D$ is the capacity price per MW·h relative to the uncovered-energy price per MWh. Table 3 reports the complete declared grid, and Online Resource 1, Figure S1 shows the two cost curves with the near-zero crossing magnified. The crossing price that equates the production rule with no commitments solves $\kappa^\star = \Delta\text{uncovered}/\Delta\text{booked}$ and equals 4.25×10^{-4} at $\eta = 0.5$, 3.01×10^{-4} at $\eta = 1$, and 1.64×10^{-4} at $\eta = 2$. This algebraic diagnostic of the fixed cost model complements the declared price-grid results.

Below its crossing price, P's uncovered-energy reduction outweighs the extra reservation in this objective. At every positive declared grid price ($\kappa \geq 0.001$), the no-extra-commitment comparator has lower normalized cost. These thresholds describe the declared price model; they quantify the capacity-to-deficit price ratio required to justify protection.

7.4 Supporting pipeline evidence

Online Resource 1, Sections D–H retain the earlier observation-index pipeline's recovery evaluations, forecasting-base comparisons and transfer studies, with their original cohorts and statistical contracts. They provide supporting context; the declared-calendar study above supplies the primary results of this article.

Table 3 Normalized annual cost $J/c_D = \kappa\,(\text{base} + \text{extra booked}) + \text{uncovered}$ in MWh-equivalent units at $\eta = 1$, over the complete declared price grid. $\kappa = c_R/c_D$ is the capacity price per MW·h relative to the uncovered-energy price per MWh.

κ	none	P (production)	$P-$none
0	1,305	967	-338
0.001	29,856	30,641	+785
0.005	144,060	149,339	+5,280
0.01	286,814	297,712	+10,898
0.025	715,077	742,829	+27,752
0.05	1,428,848	1,484,690	+55,842
0.1	2,856,391	2,968,414	+112,023
0.25	7,139,020	7,419,584	+280,564
0.5	14,276,734	14,838,201	+561,467

The crossing $\kappa^\star$ that equates the two policies solves $\kappa^\star = \Delta\text{uncovered}/\Delta\text{booked}$: 0.0004245663 at $\eta = 0.5$, 0.0003011944 at $\eta = 1$ and 0.0001636909 at $\eta = 2$. This algebraic diagnostic identifies the price range that supports additional protection; every declared positive price in the grid above favors the no-extra-commitment comparator in this normalized objective.

8 Discussion

The operating choice has three parts: placement, response size and price. The production rule places commitments in response to accumulated recent shortfall and achieves higher modeled worst-week service while spending fewer committed days than the pacing comparators; bank pacing favors physical tail protection at higher capacity booking. Increasing response size gives diminishing uncovered-energy reductions, and under the declared price model the crossing price $\kappa^\star$ and the declared price grid state when those reductions justify the additional reservation.

The paired historical evaluation characterizes performance across seven adjacent years. The 15% cap and response multipliers define a reproducible zonal capacity-accounting study, with the local bounds applying under Section 5.4. Earlier recovery evaluations retain their cohort-specific analysis and prior-outcome-access disclosure in Online Resource 1.

The framework provides a basis for combining operational response and price data with commitment planning. Prospective evaluation on later outcomes or another system can assess transfer; network-constrained dispatch, market-calibrated reserve products, physical-deficit feedback and shared multi-party budgets are natural extensions.

9 Conclusion

The proposed rule allocates a finite day-ahead commitment budget using two-day-delayed conformal feedback, with explicit accounting for budget blocking and finite capacity. It turns calibrated uncertainty into a scheduled sequence of uncertainty-informed actions. Across seven annual New York origins, it raises modeled worst-seven-day service from 0.582 to 0.713 and leads the five comparators that share the

declared cap on that worst-window endpoint. At $\eta = 1$, annual uncovered energy falls 25.9% and worst-seven-day deficit 14.8%, with 1.12 million MW·h/year of extra booked capacity, and the price analysis identifies the capacity price range in which that protection is worth its reservation. The response and price grids provide a reproducible basis for choosing the desired service, physical protection and capacity cost.

Statements and Declarations

Data and code availability

Source archives for the primary study are the NYISO forecasts and actual loads (New York Independent System Operator 2026a,b), with publication timing documented by New York ISO (2026). GEFCom2014 (Hong et al. 2016; Xie and Hong 2016); the processed CAISO benchmark (Ibebuchi 2025); and NOAA GEFS (NOAA National Centers for Environmental Information 2026) support the supplementary evaluations in Sections D–H. Provider links, archive identities, preparation steps and terms are recorded in the acquisition manifests. Original provider archives and full legacy training dependencies are obtained separately under their applicable terms.

Code and numerical artifacts are supplied in Online Resource 2 and fixed release https://github.com/ripeconan/dayahead-reserve-allocation-segan/releases/tag/segan-v1.0.1-20261011 (`segan-v1.0.1-20261011`). It separates the frozen end-to-end protocol, guard records, checkpoints, complete grids and publication generators from historical replay code and audits. Included derived records contain aligned forecast, actual and residual rows, rather than copies of provider archives. SHA-256 manifests bind inputs, outputs and external artifacts; acquisition manifests specify access routes for externally hosted files. Analytic proofs are in Online Resource 1, Section C.

Author contributions

Ran Wang: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Writing – review & editing, Visualization, Resources, Project administration.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Competing interests

The author declares that there are no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Ethics approval and consent

This study uses public energy-system time series and does not involve human participants or animals. Ethics approval and consent to participate are not applicable.

Supplementary information

Online Resource 1 contains the exact event contract and policies, complete results, analytic proofs, and supporting pipeline evaluations. Online Resource 2 provides the fixed reproduction archive with code, derived numerical artifacts, and provenance manifests.

References

Akusok A, Miche Y, Björk KM, et al (2019) Per-sample prediction intervals for extreme learning machines. International Journal of Machine Learning and Cybernetics 10(5):991–1001. https://doi.org/10.1007/s13042-017-0777-2

Angelopoulos A, Candès E, Tibshirani RJ (2023) Conformal PID control for time series prediction. In: Advances in Neural Information Processing Systems, vol 36. Curran Associates, Inc., pp 23047–23074, https://doi.org/10.52202/075280-1000

Angelopoulos AN, Bates S, Fisch A, et al (2024) Conformal risk control. In: The Twelfth International Conference on Learning Representations (ICLR 2024), URL https://openreview.net/forum?id=33XGfHLtZg

Austnes PF, García-Pareja C, Nobile F, et al (2025) Probabilistic load forecasting of distribution power systems based on empirical copulas. Sustainable Energy, Grids and Networks 42:101708. https://doi.org/10.1016/j.segan.2025.101708

Bhatnagar A, Wang H, Xiong C, et al (2023) Improved online conformal prediction via strongly adaptive online learning. In: Proceedings of the 40th International Conference on Machine Learning, Proceedings of Machine Learning Research, vol 202. PMLR, pp 2337–2363, URL https://proceedings.mlr.press/v202/bhatnagar23a.html

Chen Y, Wiesel A, Eldar YC, et al (2010) Shrinkage algorithms for MMSE covariance estimation. IEEE Transactions on Signal Processing 58(10):5016–5029. https://doi.org/10.1109/TSP.2010.2053029

Dronavajjala C (2026) When is conformal coverage free? switching thresholds for predict-then-optimize. In: Ahlberg E, Johansson U, Boström H, et al (eds) Proceedings of the Fifteenth Symposium on Conformal and Probabilistic Prediction with Applications, Proceedings of Machine Learning Research, vol 329. PMLR, pp 548–575, URL https://proceedings.mlr.press/v329/dronavajjala26a.html

Gibbs I, Candès E (2021) Adaptive conformal inference under distribution shift. In: Advances in Neural Information Processing Systems, vol 34. Curran Associates, Inc., pp 1660–1672, URL https://proceedings.neurips.cc/paper/2021/hash/0d441de75945e5acbc865406fc9a2559-Abstract.html

Gibbs I, Candès EJ (2024) Conformal inference for online prediction with arbitrary distribution shifts. Journal of Machine Learning Research 25(162):1–36. URL https://jmlr.org/papers/v25/22-1218.html

Hong T, Pinson P, Fan S, et al (2016) Probabilistic energy forecasting: Global energy forecasting competition 2014 and beyond. International Journal of Forecasting 32(3):896–913. https://doi.org/10.1016/j.ijforecast.2016.02.001

Ibebuchi C (2025) California Independent System Operator (CAISO) Energy Market Data. Dataset, Zenodo, https://doi.org/10.5281/zenodo.15034924, processed hourly dataset.

Jiao P, Chen J, Wang L, et al (2025) Distributionally robust optimal scheduling of integrated energy system considering adaptive copula function and dynamic reserve. International Journal of Electrical Power & Energy Systems 172:111073. https://doi.org/10.1016/j.ijepes.2025.111073

Johnstone C, Cox B (2021) Conformal uncertainty sets for robust optimization. In: Carlsson L, Luo Z, Cherubin G, et al (eds) Proceedings of the Tenth Symposium on Conformal and Probabilistic Prediction and Applications, Proceedings of Machine Learning Research, vol 152. PMLR, pp 72–90, URL https://proceedings.mlr.press/v152/johnstone21a.html

Joulani P, Gyorgy A, Szepesvari C (2013) Online learning under delayed feedback. In: Proceedings of the 30th International Conference on Machine Learning, Proceedings of Machine Learning Research, vol 28. PMLR, pp 1453–1461, URL https://proceedings.mlr.press/v28/joulani13.html

Ledoit O, Wolf M (2004) A well-conditioned estimator for large-dimensional covariance matrices. Journal of Multivariate Analysis 88(2):365–411. https://doi.org/10.1016/S0047-259X(03)00096-4

Messoudi S, Destercke S, Rousseau S (2022) Ellipsoidal conformal inference for multi-target regression. In: Johansson U, Boström H, Nguyen KA, et al (eds) Proceedings of the Eleventh Symposium on Conformal and Probabilistic Prediction with Applications, Proceedings of Machine Learning Research, vol 179. PMLR, pp 294–306, URL https://proceedings.mlr.press/v179/messoudi22a.html

New York Independent System Operator (2026a) ISO Load Forecast (P-7) Public Archive. NYISO Market Information System, URL https://mis.nyiso.com/public/P-7list.htm, official portal; accessed 16 July 2026.

New York Independent System Operator (2026b) Real-Time Actual Load (PAL) CSV Archive. NYISO Market Information System, URL https://mis.nyiso.com/public/P-58Blist.htm, official archive landing page; accessed 16 July 2026.

New York ISO (2026) Market participants user's guide. New York Independent System Operator, Rensselaer, NY, URL https://www.nyiso.com/documents/20142/3625950/mpug.pdf, section 5.6.1: the ISO load forecast is posted each day at approximately 08:00

NOAA National Centers for Environmental Information (2026) Global Ensemble Forecast System (GEFS). Weather and Climate Models product portal, URL https://www.ncei.noaa.gov/products/weather-climate-models/global-ensemble-forecast, official product and access portal; accessed 22 July 2026.

Qin D, Wu X, Sun D, et al (2025) Load forecasting under distribution shift: An online quantile ensembling approach. Applied Energy 401:126812. https://doi.org/10.1016/j.apenergy.2025.126812

Quanrud K, Khashabi D (2015) Online learning with adversarial delays. In: Advances in Neural Information Processing Systems, vol 28. Curran Associates, Inc., pp 1270–1278, URL https://proceedings.neurips.cc/paper_files/paper/2015/hash/72da7fd6d1302c0a159f6436d01e9eb0-Abstract.html

Shafer G, Vovk V (2008) A tutorial on conformal prediction. Journal of Machine Learning Research 9(12):371–421. URL https://www.jmlr.org/papers/v9/shafer08a.html

Suresh V, Swain A, Revathi BS, et al (2025) Mamba based adaptive conformal inference for probabilistic short-term load forecasting. Knowledge-Based Systems 328:114222. https://doi.org/10.1016/j.knosys.2025.114222

Terrén-Serrano G, Deshmukh R, Martínez-Ramón M (2026) Probabilistic day-ahead forecasting of system-level renewable energy and electricity demand. Nature Communications 17(1):3307. https://doi.org/10.1038/s41467-026-69015-w

Vaze R (2026) Simultaneous coverage and efficiency guarantee in online conformal prediction. https://doi.org/10.48550/arXiv.2607.26577, URL https://arxiv.org/abs/2607.26577, arXiv:2607.26577

Xie J, Hong T (2016) GEFCom2014 probabilistic electric load forecasting: An integrated solution with forecast combination and residual simulation. International Journal of Forecasting 32(3):1012–1016. https://doi.org/10.1016/j.ijforecast.2015.11.005

Xu C, Jiang H, Xie Y (2024) Conformal prediction for multi-dimensional time series by ellipsoidal sets. In: Proceedings of the 41st International Conference on Machine Learning, Proceedings of Machine Learning Research, vol 235. PMLR, pp 55076–55099, URL https://proceedings.mlr.press/v235/xu24m.html

Yang J, Li P, Islam MJ, et al (2024a) Online allocation with replenishable budgets: Worst case and beyond. https://doi.org/10.48550/arXiv.2401.04340, URL https://

arxiv.org/abs/2401.04340, accepted at ACM SIGMETRICS 2024, arXiv:2401.04340

Yang Y, Gao Y, Wang Z, et al (2024b) Multiscale-integrated deep learning approaches for short-term load forecasting. International Journal of Machine Learning and Cybernetics 15(12):6061–6076. https://doi.org/10.1007/s13042-024-02302-4

Yu B, Tang W (2025) Ensemble-based peak demand probability density forecasting with application to risk-aware power system scheduling. International Journal of Electrical Power & Energy Systems 173:111452. https://doi.org/10.1016/j.ijepes.2025.111452

Zuege CV, Stefenon SF, Yamaguchi CK, et al (2025) Wind speed forecasting approach using conformal prediction and feature importance selection. International Journal of Electrical Power & Energy Systems 168:110700. https://doi.org/10.1016/j.ijepes.2025.110700

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- ESM1.pdf
- ESM2.zip