# A Black-Box Reduction from Regret to Multi-Level Coverage

**Anonymous Authors**[1]

## Abstract

Online Conformal Prediction (OCP) sequentially constructs prediction sets whose empirical coverage converges to a target level on arbitrary, possibly adversarial streams. In the multi-level setting, prediction sets are produced at $d$ coverage levels simultaneously, and must be properly nested across levels. Prior work (`MultiQT`) wraps online gradient descent in a lazy isotonic projection and achieves per-level miscoverage $\mathcal{O}(\sqrt{d/T})$. We both generalize and improve this result. We show that the lazy isotonic projection is a generic wrapper: *any* online learner with a linearized regret bound is converted into a multi-level OCP algorithm with per-level coverage control governed by the Fenchel conjugate of the regret. Instantiating with $p$-norm OMD and coordinate-wise Universal Portfolio sharpens the dependence on $d$ from $\sqrt{d}$ to $\sqrt{\log d}$. We further prove that whenever the projection acts it strictly reduces both the sharpness and miscoverage components of the Weighted Interval Score, explaining the empirical Pareto-frontier improvement. Experiments on stocks, electricity demand, and synthetic streams confirm consistent gains.

## 1. Introduction

Conformal prediction provides a distribution-free framework for uncertainty quantification: it constructs a prediction set that contains the realized outcome with a user-specified probability, while making few assumptions about the underlying predictive model (Vovk et al., 2005; Shafer & Vovk, 2008; Lei et al., 2018; Barber et al., 2021). The classical exchangeability assumption can be violated, however, in sequential decision-making and time-series forecasting. Online Conformal Prediction (OCP) (Gibbs & Candes, 2021) addresses this by operating in a *fully adversarial sequential setting*. A growing body of work has developed single-level OCP algorithms maintaining long-run coverage at a target $1-\alpha$ (Gibbs & Candes, 2021; Zaffran et al., 2022; Angelopoulos et al., 2023; Bhatnagar et al., 2023; Gibbs & Candès, 2024; Podkopaev et al., 2024; Zhang et al., 2024; Liu et al., 2026); recently, Liu et al. (2026) established a black-box reduction showing that *any* online learner with a linearized regret guarantee yields a single-level OCP algorithm with provable coverage.

Many applications, however, benefit from prediction sets with simultaneous validity at *multiple* coverage levels, so that users can inspect them *ex post facto* and pick the one fitting their downstream cost. Ding et al. (2025) introduced this multi-level setting. Observed data $((Z_t, Y_t))_{t\geq 1}$ arrive over time $t = 1, 2, \ldots$; at each round $t$, a base model produces a point forecast $\hat{Y}_t$. A prediction set $\hat{C}_t(\hat{Y}_t, x)$, parameterized by a threshold $x$, covers $Y_t$ iff $Y_t \in \hat{C}_t(\hat{Y}_t, x)$. We fix $d$ target miscoverage levels $\mathcal{A} = \{\alpha_1, \ldots, \alpha_d\}$, with $0 < \alpha_1 < \cdots < \alpha_d < 1$, and require the predictor to announce a threshold vector $\boldsymbol{x}_t = (x_{t,1}, \ldots, x_{t,d})$ before $Y_t$ is revealed. Two properties are essential:

1. *coverage.* For any (even adversarial) data sequence, $\lim_{T\to\infty} \mathrm{MisCov}_{T,i} = 0$ at every level, where

$$\mathrm{MisCov}_{T,i} := \left| \frac{1}{T} \sum_{t=1}^{T} \mathbf{1}\{S_t \leq x_{t,i}\} - (1-\alpha_i) \right|.$$

2. *distributional consistency.* Forecasts are ordered: $x_{t,1} \geq \cdots \geq x_{t,d}$, for all $t$.

We refer to the two properties jointly as *multi-level coverage*.

Ding et al. (2025) proposed `MultiQT`, which wraps Online Gradient Descent (OGD) in a lazy isotonic projection and provably achieves multi-level coverage with per-level miscoverage $\mathcal{O}(\sqrt{d/T})$. However, their result applies only to a single base algorithm, and the single-level literature offers a much richer toolbox. It remains unclear how to migrate these algorithms to the multi-level setting via a black-box reduction.

**Contributions.** (1) In Section 3, we present a *black-box reduction*: any online learner with a linearized regret bound, properly wrapped with an isotonic projection layer, yields a multi-level OCP algorithm with provable multi-level coverage; the coverage-error region is characterized by the

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

Fenchel conjugate of the regret. (2) Instantiating with $p$-norm Online Mirror Descent (Gentile, 2003) and coordinate-wise Universal Portfolio (Cover & Ordentlich, 2002) yields per-level miscoverage scaling as $\sqrt{\log d}$, exponentially improving over the $\sqrt{d}$ rate of Ding et al. (2025). (3) In Section 4, we show that whenever the projection acts, it strictly and simultaneously reduces both the *sharpness reward* and the *miscoverage penalty* of the *Weighted Interval Score* (WIS) (Winkler, 1972; Gneiting & Raftery, 2007; Bracher et al., 2021), formalizing the Pareto-frontier improvement. (4) In Section 5, we evaluate the framework on synthetic and real-world streams, confirming consistent gains. A discussion of related work is deferred to Appendix A.

## 2. Problem Setup

**Notation.** For $f : \mathbb{R}^d \to \mathbb{R} \cup \{+\infty\}$, $\partial f(\boldsymbol{x})$ denotes the subdifferential and $f^\star(\boldsymbol{\theta}) := \sup_{\boldsymbol{x}} \langle \boldsymbol{\theta}, \boldsymbol{x} \rangle - f(\boldsymbol{x})$ the Fenchel conjugate. We write $\Pi_{\mathcal{C}}$ for the Euclidean projection onto a nonempty closed convex set $\mathcal{C}$, and $[d] := \{1, \dots, d\}$. Additional preliminaries on convex analysis are revisited in Appendix A.

**Multi-Level OCP.** We focus on symmetric prediction sets $\hat{C}_t(\hat{Y}_t, x) := [\hat{Y}_t - x, \hat{Y}_t + x]$ (Lei et al., 2018); our methodology extends to other nested constructions (Romano et al., 2019; Gupta et al., 2022). Let $S_t := \inf\{x \geq 0 : Y_t \in \hat{C}_t(\hat{Y}_t, x)\} = |Y_t - \hat{Y}_t|$ be the non-conformity score, so that the coverage event at level $1 - \alpha_i$ reduces to $\mathbf{1}\{S_t \leq x_{t,i}\}$. Distributional consistency requires $\boldsymbol{x}_t \in \mathcal{K} := \left\{\boldsymbol{x} \in \mathbb{R}^d : x_1 \geq \cdots \geq x_d\right\}$ at every round, where $\mathcal{K}$ is the (decreasing) *isotonic cone*[1].

**Pinball Loss and its Aggregate.** The *pinball (or quantile) loss* for online quantile learning (Gibbs & Candes, 2021) is

$$\ell^{(1-\alpha)}(x, S) := \max\left\{(1-\alpha)(S - x),\, \alpha(x - S)\right\}, \quad (1)$$

convex and $L$-Lipschitz in $x$ with $L = \max\{1-\alpha, \alpha\}$. To handle the $d$ levels jointly, we use the *aggregate pinball loss* $\mathcal{L}_t(\boldsymbol{x}) := \sum_{i=1}^d \ell^{(1-\alpha_i)}(x_i, S_t)$. We select the subgradient $\boldsymbol{g}_t \in \partial \mathcal{L}_t(\boldsymbol{x}_t)$ with coordinates

$$g_{t,i} = \mathbf{1}\{x_{t,i} \geq S_t\} - (1 - \alpha_i), \quad i \in [d], \quad (2)$$

which satisfies $\|\boldsymbol{g}_t\|_2 \leq \sqrt{d}$ and $\|\boldsymbol{g}_t\|_\infty \leq 1$. Against a fixed comparator $\boldsymbol{u} \in \mathbb{R}^d$, the regret of an algorithm producing $(\boldsymbol{x}_t)_{t \geq 1}$ is $\mathrm{Regret}_T(\boldsymbol{u}) := \sum_{t=1}^T \mathcal{L}_t(\boldsymbol{x}_t) - \sum_{t=1}^T \mathcal{L}_t(\boldsymbol{u})$. The subgradient (2) admits a direct interpretation as per-level miscoverage:

$$\mathrm{MisCov}_{T,i} = \frac{1}{T}\left|\sum_{t=1}^T g_{t,i}\right|, \quad i \in [d]. \quad (3)$$

[1]Larger sets correspond to smaller miscoverage levels by convention.

We use a single, mild assumption on the score sequence.

**Assumption 2.1.** There exist $D > 0$ and $\gamma \geq 0$ such that $S_t \leq D T^\gamma$ for every $t \in [T]$.

This generalizes the bounded-score assumption $S_t \leq D$ (recovered at $\gamma = 0$) and accommodates non-stationary streams whose magnitudes grow with the horizon (Liu et al., 2026).

## 3. From Regret to Multi-Level Coverage: A Black-Box Reduction

We design a black-box reduction that (1) keeps iterates $\boldsymbol{x}_t$ in the isotonic cone $\mathcal{K}$ at every round, and (2) secures per-level coverage by controlling the cumulative subgradient $\boldsymbol{G}_T := \sum_{t=1}^T \boldsymbol{g}_t$ via (3). A *lazy isotonic projection layer* applied as a black-box wrapper around any unconstrained online learner Alg with a linearized-regret guarantee resolves both.

**The Lazy Isotonic Projection Wrapper.** Alg is a black box producing a hidden iterate $\tilde{\boldsymbol{x}}_t \in \mathbb{R}^d$ at each round and accepting an arbitrary subgradient.

**Algorithm 1** Lazy projection wrapper Alg $\to$ multi-level OCP

1: **Input:** base learner Alg on $\mathbb{R}^d$; isotonic cone $\mathcal{K}$
2: **for** $t = 1, 2, \dots$ **do**
3: Receive hidden iterate $\tilde{\boldsymbol{x}}_t \in \mathbb{R}^d$ from Alg
4: Play projected iterate $\boldsymbol{x}_t = \Pi_{\mathcal{K}}(\tilde{\boldsymbol{x}}_t)$
5: Observe $S_t$, compute $\boldsymbol{g}_t \in \partial \mathcal{L}_t(\boldsymbol{x}_t)$ via (2)
6: Feed $\boldsymbol{g}_t$ to Alg, which produces $\tilde{\boldsymbol{x}}_{t+1}$
7: **end for**

The defining feature is the asymmetry between the played iterate $\boldsymbol{x}_t$ (projected) and the gradient $\boldsymbol{g}_t$ fed to Alg (evaluated at the projected iterate, but routed back into the unconstrained dynamics): a pattern known as *lazy projection* (Zinkevich, 2003). Since $\boldsymbol{x}_t \in \mathcal{K}$, the empirical coverages are automatically ordered, $\mathrm{Coverage}_{T,1} \geq \cdots \geq \mathrm{Coverage}_{T,d}$.

**Linear Regret Preservation via Inward Flow.** Two features of the wrapper could inflate the regret: (1) the projection moves $\tilde{\boldsymbol{x}}_t$ inside $\mathcal{K}$, and (2) Alg is updated with $\boldsymbol{g}_t$ at the played iterate, not at its own iterate. Both are reconciled by the *inward flow* property of Ding et al. (2025).

**Definition 3.1** (Inward flow (Ding et al., 2025))**.** A function $f$ with a closed convex set $\mathcal{C} \subseteq \mathbb{R}^d$ satisfies the *inward flow* property if there exists $\boldsymbol{g}(\boldsymbol{x}) \in \partial f(\boldsymbol{x})$ with $-\boldsymbol{g}(\boldsymbol{x}) \in T_{\mathcal{C}}(\boldsymbol{x})$ for every $\boldsymbol{x} \in \mathrm{bd}(\mathcal{C})$, where $T_{\mathcal{C}}(\boldsymbol{x})$ is the tangent cone.

Ding et al. (2025, Lemma 3) verifies inward flow for $(\mathcal{L}_t, \mathcal{K})$ under the choice (2); we adopt this throughout. The structural consequence we use is the following.

**Lemma 3.2** (Projection inequality (Ding et al., 2025))**.** *Let $f$ with subgradient $\boldsymbol{g}$ and a closed convex set $\mathcal{C} \subseteq \mathbb{R}^d$*

*satisfy inward flow. For every* $\boldsymbol{x} \in \mathbb{R}^d$, $\langle \boldsymbol{g}(\Pi_{\mathcal{C}}(\boldsymbol{x})), \boldsymbol{x}\rangle \geq \langle \boldsymbol{g}(\Pi_{\mathcal{C}}(\boldsymbol{x})), \Pi_{\mathcal{C}}(\boldsymbol{x})\rangle$.

A self-contained proof is in Appendix C.1. Lemma 3.2 converts a regret bound on hidden iterates into one on played iterates.

**Proposition 3.3** (Linear regret preservation)**.** *Suppose* Alg *admits a linearized regret bound* $F_T(\boldsymbol{u})$*, i.e.,* $\sum_{t=1}^{T} \langle \boldsymbol{g}_t, \tilde{\boldsymbol{x}}_t - \boldsymbol{u}\rangle \leq F_T(\boldsymbol{u})$ *for all* $\boldsymbol{u} \in \mathbb{R}^d$*. Then Algorithm 1 satisfies, for every* $\boldsymbol{u} \in \mathbb{R}^d$*,*

$$\mathrm{LinRegret}_T(\boldsymbol{u}) := \sum_{t=1}^{T} \langle \boldsymbol{g}_t, \boldsymbol{x}_t - \boldsymbol{u}\rangle \leq F_T(\boldsymbol{u}). \quad (4)$$

Notably, (4) holds for *every* $\boldsymbol{u} \in \mathbb{R}^d$, not just $\boldsymbol{u} \in \mathcal{K}$: the projection is free against any comparator the base learner could compete with. The proof (Appendix C.2) combines Lemma 3.2, the linearized-regret assumption, and convexity of $\mathcal{L}_t$.

**Characterization of Miscoverage Level Sets.** We now translate the regret guarantee into a finite-time bound on $\boldsymbol{G}_T$. By regret–reward duality (Orabona, 2019), (4) is equivalent to $F_T^\star(-\boldsymbol{G}_T) \leq -\sum_t \langle \boldsymbol{g}_t, \boldsymbol{x}_t\rangle$. Convexity and nonnegativity of $\mathcal{L}_t$ give $-\langle \boldsymbol{g}_t, \boldsymbol{x}_t\rangle \leq \mathcal{L}_t(\boldsymbol{0}) - \mathcal{L}_t(\boldsymbol{x}_t) \leq \mathcal{L}_t(\boldsymbol{0})$. Setting $B_T := \sum_{t=1}^{T} \mathcal{L}_t(\boldsymbol{0})$ and using $\mathcal{L}_t(\boldsymbol{0}) = \sum_i (1-\alpha_i) S_t \leq d S_t$, Assumption 2.1 yields $B_T \leq d D T^{1+\gamma}$.

**Theorem 3.4** (Miscoverage level set)**.** *Suppose* Alg *admits the linearized regret bound* $F_T$ *on the gradient stream produced by Algorithm 1, and that* $(\mathcal{L}_t, \mathcal{K})$ *satisfies inward flow. Then*

$$F_T^\star(-\boldsymbol{G}_T) \leq B_T, \quad (5)$$

*equivalently,*

$$\boldsymbol{G}_T \in \mathcal{Z}_T(F_T) := \left\{ \boldsymbol{z} \in \mathbb{R}^d : F_T^\star(-\boldsymbol{z}) \leq B_T \right\}.$$

Through (3), the geometry of $F_T^\star$ entirely characterizes how each $\mathrm{MisCov}_{T,i}$ is controlled. One implication is immediate: *smaller regret implies tighter coverage*; optimizing the regret-bound dependence on problem parameters directly tightens the coverage-error region (Appendix C.6). Ding et al. (2025) showed via counterexamples that naive combinations of single-level coverage with ordering can break per-level coverage; Algorithm 1 and Theorem 3.4 resolve this by converting *any* regret bound for Alg into a deterministic constraint on every $\mathrm{MisCov}_{T,i}$ while guaranteeing distributional consistency.

### 3.1. $p$-Norm Online Mirror Descent

Fix $p \in (1, 2]$ with $q := p/(p-1) \in [2, \infty)$. Take Alg to be Online Mirror Descent with regularizer $\psi(\tilde{\boldsymbol{x}}) = \frac{1}{2}\|\tilde{\boldsymbol{x}}\|_p^2$ and constant step size $\eta > 0$, initialized at $\tilde{\boldsymbol{x}}_1 = \boldsymbol{0}$. The wrapped algorithm has the standard mirror-descent dual form with an isotonic projection at prediction time (Algorithm 3 in Appendix E.2). Standard analysis (Orabona, 2019, Sec. 6.7) gives $F_T(\boldsymbol{u}) = \|\boldsymbol{u}\|_p^2/(2\eta) + \eta \sum_t \|\boldsymbol{g}_t\|_q^2/(2(p-1))$, with $F_T^\star(\boldsymbol{\theta}) = (\eta/2)\|\boldsymbol{\theta}\|_q^2 - \eta \sum_t \|\boldsymbol{g}_t\|_q^2/(2(p-1))$. Substituting into Theorem 3.4 and using $\|\boldsymbol{g}_t\|_q \leq d^{1/q}$ yields:

**Corollary 3.5** ($\ell_q$ coverage-error ball)**.** *Algorithm 3 with step size* $\eta > 0$ *satisfies*

$$\max_{i\in[d]} \mathrm{MisCov}_{T,i} \leq \frac{\|\boldsymbol{G}_T\|_q}{T} \leq \frac{1}{T}\sqrt{\frac{2B_T}{\eta} + \frac{d^{2/q}T}{p-1}}. \quad (6)$$

Assuming $d \geq 3$ and setting $q = 2\log d$ (so $p - 1 = 1/(2\log d - 1)$ and $d^{2/q} = e$), then balancing with $\eta = 2B_T/(e(2\log d - 1)T)$,

$$\max_{i\in[d]} \mathrm{MisCov}_{T,i} \leq \sqrt{\frac{2e(2\log d - 1)}{T}} = \mathcal{O}\left(\sqrt{\log d/T}\right).$$

This is an exponential improvement over the $\sqrt{d}$ rate of `MultiQT` (recovered at $p = 2$). The optimal $\eta$ in $p$-norm OMD, as in `MultiQT`, depends on the unknown growth parameters $D, \gamma$.

### 3.2. Coordinate-Wise Universal Portfolio

To remove the dependence on $B_T$ in the step-size choice, replace OMD with a parameter-free algorithm (Orabona, 2019, Sec. 10). Since $\mathcal{L}_t$ decomposes coordinate-wise, we run an independent copy of the scalar universal-portfolio learner `UP-OCP` (Liu et al., 2026) at each level $\alpha_j$ and apply Algorithm 1 on the joint vector (pseudocode in Appendix E.2). The total regret is $F_T(\boldsymbol{u}) = \sum_{j=1}^{d} F_T^{(j)}(|u_j|)$, with $F_T^{(j)}$ as in (Liu et al., 2026, Theorem 3.4); the conjugate factorizes, $F_T^\star(\boldsymbol{\theta}) = \sum_{j=1}^{d} (F_T^{(j)})^\star(|\theta_j|)$, and Theorem 3.4 becomes $\sum_{j=1}^{d} (F_T^{(j)})^\star(|G_{T,j}|) \leq B_T$ (cf. (5)). Since $(F_T^{(j)})^\star(0) \geq -1$, dropping all but one term yields a per-coordinate budget of $B_T + d - 1$; inverting the resulting Bernstein-type inequality gives:

**Theorem 3.6** (Per-level coverage of coordinate-wise `UP-OCP`)**.** *For every* $j \in [d]$*,*

$$\mathrm{MisCov}_{T,j} \leq \hat{\varepsilon}_{T,j} + \sqrt{2\alpha_j(1-\alpha_j)\hat{\varepsilon}_{T,j}}, \quad (7)$$

*where* $\hat{\varepsilon}_{T,j} := (1/T)\left[\ln(B_T + d) + (1/2)\ln(2(T+1))\right]$*. Under Assumption 2.1,*

$$\max_j \mathrm{MisCov}_{T,j} = \mathcal{O}\left(\sqrt{((1+\gamma)\log T + \log d)/T}\right).$$

The rate matches Corollary 3.5 in $d$, up to a $\log T$ factor that is the cost of being parameter-free. Crucially, no quantity

used by the algorithm requires knowledge of $D$, $\gamma$, or $T$, and the $\alpha_j(1-\alpha_j)$ factor delivers a faster rate at extreme levels. An extension of the framework from a finite collection $\mathcal{A}$ to a compact interval $\mathcal{X} \subseteq (0,1)$ (in the Hilbert space $L^2(\mathcal{X},\mu)$) is collected in Appendix C.8.

## 4. Sharpness and Coverage Improvement of Isotonic Projection

In Section 3 the projection's role was structural. We now show it does strictly more: whenever $\tilde{\boldsymbol{x}}_t$ violates isotonicity, projecting onto $\mathcal{K}$ *strictly* reduces a measure of prediction width while *weakly* reducing a measure of miscoverage. Neither is traded for the other, so the empirical width-vs-coverage Pareto frontier shifts monotonically toward the ideal corner.

We frame the analysis via the *Weighted Interval Score* (WIS) (Bracher et al., 2021), the standard scoring rule for multi-level interval forecasts, extending the classical Interval Score (Winkler, 1972; Gneiting & Raftery, 2007). For symmetric prediction sets $\hat{C}_t(\hat{Y}_t, x) = [\hat{Y}_t - x, \hat{Y}_t + x]$, a direct calculation (Appendix D.1) gives

$$\begin{aligned}\mathcal{L}_t(\boldsymbol{x}) - S_t \sum_{i=1}^{d} \alpha_i &= \frac{1}{2}\,\mathrm{WIS}_{\mathcal{A}} \\ &= \underbrace{\sum_{i=1}^{d} \alpha_i x_i}_{\textit{sharpness reward}} + \underbrace{\sum_{i=1}^{d} \max(S_t - x_i, 0)}_{\textit{coverage penalty}}. \end{aligned} \tag{8}$$

That is, up to a constant in $\boldsymbol{x}_t$, the aggregate pinball loss *is* half the WIS. The sharpness term is an $\alpha$-weighted total width (higher miscoverage levels weighted more); the coverage penalty is a one-sided $\ell_1$ residual penalizing under-coverage at every level.

**Theorem 4.1** (Per-round reduction)**.** *For every round $t$, $\tilde{\boldsymbol{x}}_t \in \mathbb{R}^d$ and $\boldsymbol{x}_t = \Pi_{\mathcal{K}}(\tilde{\boldsymbol{x}}_t)$ satisfy $\mathcal{L}_t(\boldsymbol{x}_t) \le \mathcal{L}_t(\tilde{\boldsymbol{x}}_t)$. Moreover, both the $\alpha$-weighted prediction sum and the total coverage penalty are individually reduced:*

$$\sum_{i=1}^{d} \alpha_i x_{t,i} \le \sum_{i=1}^{d} \alpha_i \tilde{x}_{t,i}, \tag{9}$$

$$\sum_{i=1}^{d} \max(S_t - x_{t,i}, 0) \le \sum_{i=1}^{d} \max(S_t - \tilde{x}_{t,i}, 0). \tag{10}$$

*When $\tilde{\boldsymbol{x}}_t \notin \mathcal{K}$, the sharpness and loss reductions are strict.*

The proof relies on the Pool-Adjacent-Violators (PAVA) characterization of $\Pi_{\mathcal{K}}$ (Robertson et al., 1988); details are in Appendix D.2. Ding et al. (2025) obtained a decomposition analogous to (8) but only used it as motivation; they observed better empirical performance of `MultiQT` over naive forecasters but did not explain it. Theorem 4.1 provides this missing explanation.

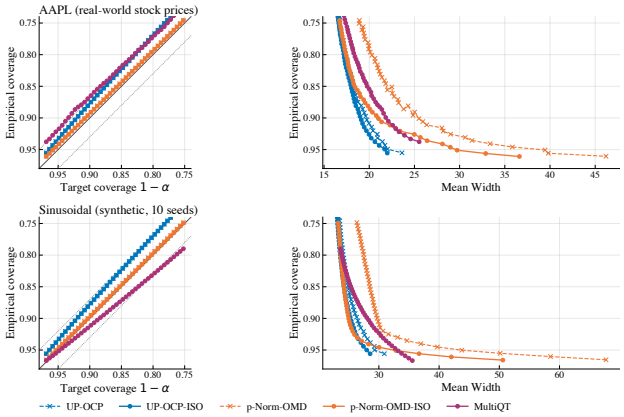


*Figure 1.* Calibration (left) and mean prediction-set width (right) vs. target coverage $1-\alpha$, for AAPL (top) and a sinusoidal synthetic stream (bottom, mean over 10 seeds). Dotted lines: $\pm 0.03$ band. Lower-left is better.

## 5. Experiments

We evaluate the framework on three categories of streams: US-stock prices (AAPL, AMZN, GOOGL) (Nguyen, 2018); New South Wales electricity demand (Harries, 1999); and synthetic environments (Angelopoulos et al., 2023, App. F.5). Nonconformity scores are absolute residuals $S_t = \left|Y_t - \hat{Y}_t\right|$ from daily-refit Prophet (Taylor & Letham, 2018) on stocks and AR(3) on electricity; full details are in Appendix E.1.

**Methods.** We compare `UP-OCP` (parameter-free), $p$-`norm OMD` ($p = 2\log d/(2\log d - 1)$, $\eta$ from oracle grid search), `KT` (parameter-free), and `MultiQT` (Ding et al., 2025); the first three are run both vanilla and wrapped with the lazy projection (-ISO), while `MultiQT` is by construction projected. We calibrate $d = 100$ levels uniformly in $[0.001, 0.5]$ with $T_{\text{burn-in}} = 100$, reporting on $\alpha \in [0.03, 0.25]$. Hyperparameters, additional datasets, and width-distribution analyses are in Appendix E.

**Calibration and Pareto Frontiers.** Figure 1 shows two observations. *Calibration is preserved by the projection*: -ISO and unprojected curves are visually indistinguishable inside the $\pm 0.03$ band, matching Proposition 3.3. *Sharpness is strictly improved*: every -ISO curve lies on or below its unprojected counterpart, with the gap widening at high coverage where isotonicity violations concentrate; unprojected $p$-`norm OMD` widths can be 30–50% larger near $1-\alpha = 0.97$. Coordinate-wise `UP-OCP`-ISO achieves the smallest mean width on AAPL while remaining parameter-free, in contrast to oracle-tuned $p$-`norm OMD`. Tables with `KT` and tail statistics, width distributions (projection removes heavy right-tail mass), and the remaining streams are in Appendix E.

## References


Angelopoulos, A., Candes, E., and Tibshirani, R. J. Conformal PID control for time series prediction. *Advances in neural information processing systems*, 36:23047–23074, 2023.

Angelopoulos, A. N., Bates, S., Fisch, A., Lei, L., and Schuster, T. Conformal risk control. *arXiv preprint arXiv:2208.02814*, 2022.

Angelopoulos, A. N., Jordan, M. I., and Tibshirani, R. J. Gradient equilibrium in online learning: Theory and applications. *arXiv preprint arXiv:2501.08330*, 2025.

Barber, R. F., Candes, E. J., Ramdas, A., and Tibshirani, R. J. Predictive inference with the jackknife+. *The Annals of Statistics*, 49(1):486–507, 2021.

Bastani, O., Gupta, V., Jung, C., Noarov, G., Ramalingam, R., and Roth, A. Practical adversarial multivalid conformal prediction. *Advances in neural information processing systems*, 35:29362–29373, 2022.

Bhatnagar, A., Wang, H., Xiong, C., and Bai, Y. Improved online conformal prediction via strongly adaptive online learning. In *International Conference on Machine Learning*, pp. 2337–2363. PMLR, 2023.

Boucheron, S., Lugosi, G., and Massart, P. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.

Bracher, J., Ray, E. L., Gneiting, T., and Reich, N. G. Evaluating epidemic forecasts in an interval format. *PLoS computational biology*, 17(2):e1008618, 2021.

Cai, Y., Daskalakis, C., Luo, H., Wei, C.-Y., and Zheng, W. On tractable $\phi$-equilibria in non-concave games. *Advances in Neural Information Processing Systems*, 37: 140366–140404, 2024.

Cover, T. M. and Ordentlich, E. Universal portfolios with side information. *IEEE Transactions on Information Theory*, 42(2):348–363, 2002.

Ding, T., Gibbs, I., and Tibshirani, R. J. Calibrated multi-level quantile forecasting. *arXiv preprint arXiv:2512.23671*, 2025.

Gauthier, E., Bach, F., and Jordan, M. I. E-values expand the scope of conformal prediction. *arXiv preprint arXiv:2503.13050*, 2025.

Gentile, C. The robustness of the $p$-norm algorithms. *Machine Learning*, 53(3):265–299, 2003.

Gibbs, I. and Candes, E. Adaptive conformal inference under distribution shift. *Advances in Neural Information Processing Systems*, 34:1660–1672, 2021.

Gibbs, I. and Candès, E. J. Conformal inference for online prediction with arbitrary distribution shifts. *Journal of Machine Learning Research*, 25(162):1–36, 2024.

Gneiting, T. and Raftery, A. E. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477):359–378, 2007.

Gupta, C., Kuchibhotla, A. K., and Ramdas, A. Nested conformal prediction and quantile out-of-bag ensemble methods. *Pattern Recognition*, 127:108496, 2022.

Harries, M. Splice-2 comparative evaluation: Electricity pricing. Technical report, University of New South Wales, School of Computer Science and Engineering, 1999.

Lee, Y., Dobriban, E., and Tchetgen, E. T. Simultaneous conformal prediction of missing outcomes with propensity score $\varepsilon$-discretization. *arXiv preprint arXiv:2403.04613*, 2024.

Lei, J., G'Sell, M., Rinaldo, A., Tibshirani, R. J., and Wasserman, L. Distribution-free predictive inference for regression. *Journal of the American Statistical Association*, 113(523):1094–1111, 2018.

Liu, T., Dobriban, E., and Orabona, F. Online conformal prediction via universal portfolio algorithms. *arXiv preprint arXiv:2602.03168*, 2026.

Nguyen, C. S&P 500 stock data, 2018. URL https://www.kaggle.com/datasets/camnugent/sandp500.

Orabona, F. A modern introduction to online learning. *arXiv preprint arXiv:1912.13213*, 2019.

Park, S., Dobriban, E., Lee, I., and Bastani, O. PAC prediction sets under covariate shift. In *International Conference on Learning Representations*, 2022.

Podkopaev, A. and Ramdas, A. Distribution-free uncertainty quantification for classification under label shift. *arXiv preprint arXiv:2103.03323*, 2021.

Podkopaev, A., Xu, D., and Lee, K.-c. Adaptive conformal inference by betting. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 40886–40907, 2024.

Qiu, H., Dobriban, E., and Tchetgen Tchetgen, E. Prediction sets adaptive to unknown covariate shift. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, pp. qkad069, 07 2023.

Robertson, T., Wright, F. T., and Dykstra, R. L. *Order Restricted Statistical Inference*. Wiley, 1988.

Rockafellar, R. T. *Convex analysis*, volume 28. Princeton university press, 1997.

Romano, Y., Patterson, E., and Candes, E. Conformalized quantile regression. *Advances in neural information processing systems*, 32, 2019.

Shafer, G. and Vovk, V. A tutorial on conformal prediction. *Journal of Machine Learning Research*, 9(3):371–421, 2008.

Shalev-Shwartz, S. *Online Learning: Theory, Algorithms, and Applications*. PhD thesis, Hebrew University of Jerusalem, 2007.

Si, W., Park, S., Lee, I., Dobriban, E., and Bastani, O. PAC prediction sets under label shift. *International Conference on Learning Representations*, 2024.

Srinivas, V. Online conformal prediction with efficiency guarantees. In *Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 6696–6726. SIAM, 2026.

Taylor, S. J. and Letham, B. Forecasting at scale. *The American Statistician*, 72(1):37–45, 2018.

Tibshirani, R. J., Foygel Barber, R., Candes, E., and Ramdas, A. Conformal prediction under covariate shift. *Advances in neural information processing systems*, 32, 2019.

Vovk, V., Gammerman, A., and Shafer, G. *Algorithmic Learning in a Random World*. Springer Science & Business Media, 2005.

Winkler, R. L. A decision-theoretic approach to interval estimation. *Journal of the American Statistical Association*, 67(337):187–191, 1972.

Xu, C. and Xie, Y. Conformal prediction interval for dynamic time-series. In *International Conference on Machine Learning*, pp. 11559–11569, 2021.

Zaffran, M., Féron, O., Goude, Y., Josse, J., and Dieuleveut, A. Adaptive conformal predictions for time series. In *International conference on machine learning*, pp. 25834–25866. PMLR, 2022.

Zhang, Z., Bombara, D., and Yang, H. Discounted adaptive online learning: towards better regularization. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 58631–58661, 2024.

Zinkevich, M. Online convex programming and generalized infinitesimal gradient ascent. In *Proc. of the International Conference on Machine Learning*, pp. 928–936, 2003.

# Appendix


**A Related Work** **8**

**B Preliminaries of Convex Analysis** **8**

**C Deferred Proofs and Supporting Material for Section 3** **8**

C.1 Projection Inequality (Lemma 3.2) . . . 8

C.2 Regret Preservation and the Central Inequality . . . 9

C.3 Inward Flow for the Isotonic Cone . . . 9

C.4 Detailed Derivation for Corollary 3.5: $p$-Norm OMD . . . 9

C.5 Beyond the Isotonic Cone . . . 10

C.6 Tighter Regret Implies Tighter Coverage . . . 10

C.7 Coordinate-Wise UP-OCP: Per-Level and Mean-Absolute Proofs . . . 11

C.8 Continuous Quantile Levels . . . 12

**D Deferred Proofs for Section 4** **14**

D.1 Pinball Loss as a Weighted Interval Score (Equation (8)) . . . 14

D.2 Per-Round Structural Reduction (Theorem 4.1) . . . 15

**E Experimental Details and Additional Results** **16**

E.1 Datasets and Forecasting Pipelines . . . 16

E.2 Detailed Update Rules . . . 16

E.2.1 Coordinate-Wise Universal Portfolio (`UP-OCP`) . . . 17

E.2.2 $p$-Norm Online Mirror Descent . . . 17

E.2.3 Coordinate-Wise Krichevsky–Trofimov Bettor (`KT`) . . . 18

E.2.4 `MultiQT` . . . 18

E.3 Width Distribution Analysis . . . 19

E.4 Additional Results . . . 20

E.4.1 AAPL (Full Results) . . . 21

E.4.2 AMZN . . . 22

E.4.3 GOOGL . . . 23

E.4.4 Electricity Demand (`elec2`) . . . 24

E.4.5 Sinusoidal (Synthetic) . . . 25

E.4.6 Stationary Trend with Random Waves (Synthetic) . . . 26

E.4.7 Quadratic Drift with Random Waves (`mixed`, Synthetic) . . . 27

# A. Related Work

**Related Work.** Our work sits at the intersection of online conformal prediction, online learning, and multi-level coverage. Adaptive conformal inference and its variants treat the conformal radius as a sequentially updated decision variable, often optimized through the pinball loss (Gibbs & Candes, 2021; Zaffran et al., 2022; Gibbs & Candès, 2024). Strongly adaptive and control-based methods improve local responsiveness or empirical stability, but require method-specific choices such as expert grids, gain parameters, or range assumptions (Angelopoulos et al., 2023; Bhatnagar et al., 2023). Parameter-free approaches obtain optimal regret guarantees uniformly over all comparators, by avoiding explicit stepsize tuning (Podkopaev et al., 2024; Zhang et al., 2024; Liu et al., 2026). Among these works, the closest to ours is Liu et al. (2026), which proves a black-box linearized-regret-to-coverage implication, but only for scalar OCP, i.e., a single coverage level. Related analyses distinguish standard regret from the stronger conditions needed for coverage, including gradient-equilibrium and proximal-regret viewpoints (Angelopoulos et al., 2025; Cai et al., 2024).

The multi-level setting is connected to, but distinct from, simultaneous and multi-valid forms of conformal inference, where one controls empirical coverage across groups or threshold families (Bastani et al., 2022), while simultaneous conformal and conformal risk-control methods handle many inference targets or losses at once (Angelopoulos et al., 2022; Lee et al., 2024). Instead, our focus is an adversarial online protocol for an ordered family of conformal sets indexed by miscoverage levels. Finally, our sharpness interpretation is aligned with proper interval and quantile scoring rules (Winkler, 1972; Gneiting & Raftery, 2007; Bracher et al., 2021).

Conformal methods have also been developed for covariate shift, label shift, general distribution shift, and time series (Tibshirani et al., 2019; Qiu et al., 2023; Park et al., 2022; Podkopaev & Ramdas, 2021; Si et al., 2024; Gauthier et al., 2025; Xu & Xie, 2021); efficiency guarantees through competitive-ratio formulations are studied by Srinivas (2026).

# B. Preliminaries of Convex Analysis

We collect a few standard tools from convex analysis in $\mathbb{R}^d$ (see, e.g., Rockafellar, 1997) for reference. For a function $f : \mathbb{R}^d \to \mathbb{R} \cup \{+\infty\}$, a vector $\boldsymbol{g} \in \mathbb{R}^d$ is a *subgradient* of $f$ at $\boldsymbol{x} \in \mathbb{R}^d$ if $f(\boldsymbol{y}) \geq f(\boldsymbol{x}) + \langle \boldsymbol{g}, \boldsymbol{y} - \boldsymbol{x} \rangle$ for every $\boldsymbol{y} \in \mathbb{R}^d$, and the set of all such $\boldsymbol{g}$ is the *subdifferential* $\partial f(\boldsymbol{x})$. The *Fenchel conjugate* of $f$ is $f^\star(\boldsymbol{\theta}) := \sup_{\boldsymbol{x} \in \mathbb{R}^d} \langle \boldsymbol{\theta}, \boldsymbol{x} \rangle - f(\boldsymbol{x})$, which is always closed and convex. The *Euclidean projection* onto a nonempty closed convex set $\mathcal{C} \subseteq \mathbb{R}^d$ is $\Pi_{\mathcal{C}}(\boldsymbol{x}) := \arg\min_{\boldsymbol{z} \in \mathcal{C}} \|\boldsymbol{x} - \boldsymbol{z}\|_2$, well-defined and unique. The *normal cone* of $\mathcal{C}$ at $\boldsymbol{x} \in \mathcal{C}$ is $N_{\mathcal{C}}(\boldsymbol{x}) := \{\boldsymbol{v} \in \mathbb{R}^d : \langle \boldsymbol{v}, \boldsymbol{y} - \boldsymbol{x} \rangle \leq 0 \text{ for all } \boldsymbol{y} \in \mathcal{C}\}$, the polar of the tangent cone $T_{\mathcal{C}}(\boldsymbol{x}) := \mathrm{cl}\,\{\boldsymbol{v} : \exists \delta > 0, \boldsymbol{x} + \varepsilon \boldsymbol{v} \in \mathcal{C} \text{ for } \varepsilon \in (0, \delta]\}$ (Rockafellar, 1997, Sec. 6).

# C. Deferred Proofs and Supporting Material for Section 3

This appendix collects the proofs and auxiliary results deferred from Section 3. We restate each statement before its proof. Throughout, we write

$$\boldsymbol{\delta}_T \;:=\; \frac{\boldsymbol{G}_T}{T} \;\in\; \mathbb{R}^d$$

for the time-averaged cumulative subgradient, so that $\delta_{T,i} = G_{T,i}/T$ is the per-level coverage gap and (3) reads $\mathrm{MisCov}_{T,i} = |\delta_{T,i}|$.

## C.1. Projection Inequality (Lemma 3.2)

We first record an immediate consequence of inward flow, then prove Lemma 3.2. The argument is in $\mathbb{R}^d$, but every step extends verbatim to a Hilbert space, which is exploited in Appendix C.8. Recall the *normal cone* of a closed convex set $C$ at $\boldsymbol{x} \in C$, $N_C(\boldsymbol{x}) := \{\boldsymbol{v} \in \mathbb{R}^d : \langle \boldsymbol{v}, \boldsymbol{y} - \boldsymbol{x} \rangle \leq 0 \text{ for all } \boldsymbol{y} \in C\}$, is the polar of the tangent cone $T_C(\boldsymbol{x})$, i.e., $\langle \boldsymbol{v}, \boldsymbol{w} \rangle \leq 0$ for every $\boldsymbol{v} \in N_C(\boldsymbol{x})$ and $\boldsymbol{w} \in T_C(\boldsymbol{x})$ (Rockafellar, 1997, Sec. 6).

**Lemma C.1.** *Let $\ell$ with subgradient $\boldsymbol{g}$ and a closed convex set $C \subseteq \mathbb{R}^d$ satisfy inward flow. For every $\boldsymbol{x} \in \mathrm{bd}(C)$ and $\boldsymbol{v} \in N_C(\boldsymbol{x})$, $\langle \boldsymbol{v}, \boldsymbol{g}(\boldsymbol{x}) \rangle \geq 0$.*

*Proof.* By inward flow, $-\boldsymbol{g}(\boldsymbol{x}) \in T_C(\boldsymbol{x})$. By tangent–normal cone polarity, $\langle \boldsymbol{v}, -\boldsymbol{g}(\boldsymbol{x}) \rangle \leq 0$ for every $\boldsymbol{v} \in N_C(\boldsymbol{x})$, i.e., $\langle \boldsymbol{v}, \boldsymbol{g}(\boldsymbol{x}) \rangle \geq 0$. □

*Proof of Lemma 3.2.* If $\tilde{\boldsymbol{x}} \in C$, then $\Pi_C(\tilde{\boldsymbol{x}}) = \tilde{\boldsymbol{x}}$ and the inequality is an equality. Otherwise, let $\boldsymbol{x}_0 = \Pi_C(\tilde{\boldsymbol{x}})$. By the

variational characterization of Euclidean projection, there exists $\boldsymbol{v} \in N_C(\boldsymbol{x}_0)$ such that $\tilde{\boldsymbol{x}} = \boldsymbol{x}_0 + \boldsymbol{v}$. Then

$$\langle \boldsymbol{g}(\boldsymbol{x}_0), \tilde{\boldsymbol{x}} \rangle = \langle \boldsymbol{g}(\boldsymbol{x}_0), \boldsymbol{x}_0 \rangle + \langle \boldsymbol{g}(\boldsymbol{x}_0), \boldsymbol{v} \rangle \geq \langle \boldsymbol{g}(\boldsymbol{x}_0), \boldsymbol{x}_0 \rangle ,$$

where the inequality is Lemma C.1 with $\boldsymbol{x} = \boldsymbol{x}_0 \in \mathrm{bd}(C)$. □

### C.2. Regret Preservation and the Central Inequality

*Proof of Proposition 3.3.* By convexity of $\mathcal{L}_t$ and $\boldsymbol{g}_t \in \partial \mathcal{L}_t(\boldsymbol{x}_t)$, $\mathcal{L}_t(\boldsymbol{x}_t) - \mathcal{L}_t(\boldsymbol{u}) \leq \langle \boldsymbol{g}_t, \boldsymbol{x}_t - \boldsymbol{u} \rangle$. Summing over $t$ gives the first inequality of (4). By Lemma 3.2 applied to $C = \mathcal{K}$, $\langle \boldsymbol{g}_t, \boldsymbol{x}_t \rangle \leq \langle \boldsymbol{g}_t, \tilde{\boldsymbol{x}}_t \rangle$ for every $t$, so

$$\sum_{t=1}^{T} \langle \boldsymbol{g}_t, \boldsymbol{x}_t - \boldsymbol{u} \rangle \leq \sum_{t=1}^{T} \langle \boldsymbol{g}_t, \tilde{\boldsymbol{x}}_t - \boldsymbol{u} \rangle \leq F_T(\boldsymbol{u}),$$

where the last step is the regret bound of Alg. □

*Proof of Theorem 3.4.* By regret-reward duality (Orabona, 2019, Sec. 6.4), the regret bound of Alg is equivalent to $F_T^\star\Big(-\sum_{t=1}^{T} \boldsymbol{g}_t\Big) \leq -\sum_{t=1}^{T} \langle \boldsymbol{g}_t, \tilde{\boldsymbol{x}}_t \rangle$ . Inward flow of $(\mathcal{L}_t, \mathcal{K})$ together with Lemma 3.2 gives $\langle \boldsymbol{g}_t, \boldsymbol{x}_t \rangle \leq \langle \boldsymbol{g}_t, \tilde{\boldsymbol{x}}_t \rangle$, hence $-\langle \boldsymbol{g}_t, \tilde{\boldsymbol{x}}_t \rangle \leq -\langle \boldsymbol{g}_t, \boldsymbol{x}_t \rangle$. Convexity of $\mathcal{L}_t$ at $\boldsymbol{x}_t$ with comparator $\mathbf{0}$ yields

$$\mathcal{L}_t(\mathbf{0}) \geq \mathcal{L}_t(\boldsymbol{x}_t) + \langle \boldsymbol{g}_t, \mathbf{0} - \boldsymbol{x}_t \rangle \quad \Longrightarrow \quad -\langle \boldsymbol{g}_t, \boldsymbol{x}_t \rangle \leq \mathcal{L}_t(\mathbf{0}) - \mathcal{L}_t(\boldsymbol{x}_t) \leq \mathcal{L}_t(\mathbf{0}),$$

where we used $\mathcal{L}_t(\boldsymbol{x}_t) \geq 0$ since each pinball loss is non-negative. Summing over $t$ gives

$$F_T^\star(-\boldsymbol{G}_T) \leq -\sum_{t=1}^{T} \langle \boldsymbol{g}_t, \tilde{\boldsymbol{x}}_t \rangle \leq \sum_{t=1}^{T} \mathcal{L}_t(\mathbf{0}) = B_T.$$ □

### C.3. Inward Flow for the Isotonic Cone

*Verification of inward flow for the isotonic cone.* Fix $t$ and $\boldsymbol{x} \in \mathrm{bd}(\mathcal{K})$. Let $I(\boldsymbol{x}) := \{i \in [d-1] : x_i = x_{i+1}\}$ be the set of tight constraints; $I(\boldsymbol{x})$ is non-empty since $\boldsymbol{x}$ is on the boundary. Because $\mathcal{K}$ is polyhedral with constraints $x_i - x_{i+1} \geq 0$, the tangent cone at $\boldsymbol{x}$ is

$$T_{\mathcal{K}}(\boldsymbol{x}) = \left\{ \boldsymbol{v} \in \mathbb{R}^d : v_i - v_{i+1} \geq 0 \text{ for every } i \in I(\boldsymbol{x}) \right\}.$$

Inward flow asks $-\boldsymbol{g}_t \in T_{\mathcal{K}}(\boldsymbol{x})$, equivalently $g_{t,i} \leq g_{t,i+1}$ for every $i \in I(\boldsymbol{x})$. With the subgradient (2), we have $g_{t,i} = \mathbf{1}\{x_i \geq S_t\} - (1 - \alpha_i)$. Take $i \in I(\boldsymbol{x})$ so that $x_i = x_{i+1}$.

*Case 1:* $x_i = x_{i+1} \geq S_t$. Then $g_{t,i} = \alpha_i$ and $g_{t,i+1} = \alpha_{i+1}$. Since $\alpha_i < \alpha_{i+1}$, $g_{t,i} < g_{t,i+1}$.

*Case 2:* $x_i = x_{i+1} < S_t$. Then $g_{t,i} = -(1 - \alpha_i)$ and $g_{t,i+1} = -(1 - \alpha_{i+1})$. Since $\alpha_i < \alpha_{i+1}$, $1 - \alpha_i > 1 - \alpha_{i+1}$, hence $g_{t,i} < g_{t,i+1}$.

In either case $g_{t,i} \leq g_{t,i+1}$, completing the proof. □

### C.4. Detailed Derivation for Corollary 3.5: $p$-Norm OMD

The regret bound for $p$-norm OMD with regulariser $\psi(\boldsymbol{x}) = (1/2)\|\boldsymbol{x}\|_p^2$ at constant step size $\eta > 0$ is standard (Orabona, 2019, Theorem 6.8): $\psi$ is $(p-1)$-strongly convex with respect to $\|\cdot\|_p$ by (Shalev-Shwartz, 2007, Lemma 6.36), yielding

$$F_T(\boldsymbol{u}) = \frac{\|\boldsymbol{u}\|_p^2}{2\eta} + \frac{\eta}{2(p-1)} \sum_{t=1}^{T} \|\boldsymbol{g}_t\|_q^2 .$$

The Fenchel conjugate of $a\|\cdot\|_p^2$ is $(1/(4a))\|\cdot\|_q^2$. Applying this with $a = 1/(2\eta)$ and absorbing the additive $T$-dependent term gives

$$F_T^\star(\boldsymbol{\theta}) = \frac{\eta}{2}\|\boldsymbol{\theta}\|_q^2 - \frac{\eta}{2(p-1)} \sum_{t=1}^{T} \|\boldsymbol{g}_t\|_q^2 .$$

The central inequality (5) then becomes $(\eta/2)\,\|\boldsymbol{G}_T\|_q^2 - (\eta/(2(p-1)))\sum_{t=1}^T \|\boldsymbol{g}_t\|_q^2 \le B_T$, i.e.,

$$\|\boldsymbol{G}_T\|_q^2 \le \frac{2B_T}{\eta} + \frac{1}{p-1}\sum_{t=1}^T \|\boldsymbol{g}_t\|_q^2 .$$

Since $\|\boldsymbol{g}_t\|_\infty \le 1$ and $\|\boldsymbol{g}_t\|_q \le d^{1/q}\,\|\boldsymbol{g}_t\|_\infty \le d^{1/q}$ by Hölder, $\sum_{t=1}^T \|\boldsymbol{g}_t\|_q^2 \le d^{2/q}\,T$.

For the optimization, set $q = 2\log d$. Then $p-1 = 1/(q-1) = 1/(2\log d - 1)$ and $d^{2/q} = d^{1/\log d} = e$. Plugging in,

$$\|\boldsymbol{G}_T\|_q^2 \le \frac{2B_T}{\eta} + e(2\log d - 1)T.$$

Setting $\eta := 2B_T/(e(2\log d - 1)T)$ balances the two terms and gives $\|\boldsymbol{G}_T\|_q^2 \le 2e(2\log d - 1)T$. This choice of $\eta$ is well-defined whenever $B_T > 0$; the trivial case $B_T = 0$ corresponds to $S_t = 0$ for all $t$, where any positive $\eta$ already yields zero miscoverage. Under Assumption 2.1, the upper bound $B_T \le dDT^{1+\gamma}$ gives $\eta \le 2dDT^\gamma/(e(2\log d - 1)) = O(dDT^\gamma/\log d)$ as a function of $T$.

The per-level miscoverage bound follows from $\max_{i\in[d]} |\delta_{T,i}| = \max_i |G_{T,i}|/T \le \|\boldsymbol{G}_T\|_q/T$, since $\|\cdot\|_\infty \le \|\cdot\|_q$.

### C.5. Beyond the Isotonic Cone

The lazy-projection framework rests on a single geometric assumption—inward flow—and is therefore agnostic to the details of $\mathcal{K}$. We record here a composition principle that allows additional polyhedral constraints to be layered on top of the isotonic cone without redoing the analysis.

**Proposition C.2** (Inward flow under intersection)**.** *Let $C_1, C_2 \subseteq \mathbb{R}^d$ be closed convex sets and $\ell$ a loss with subgradient selection $\boldsymbol{g}(\cdot)$. Suppose that*

*(i) $(\ell, C_1)$ and $(\ell, C_2)$ both satisfy inward flow with the same selection $\boldsymbol{g}$;*
*(ii) for every $\boldsymbol{x} \in \mathrm{bd}(C_1 \cap C_2)$, $T_{C_1\cap C_2}(\boldsymbol{x}) = T_{C_1}(\boldsymbol{x}) \cap T_{C_2}(\boldsymbol{x})$.*

*Then $(\ell, C_1 \cap C_2)$ satisfies inward flow with $\boldsymbol{g}$. Condition (ii) holds whenever $C_1$ and $C_2$ are polyhedral.*

*Proof.* Fix $\boldsymbol{x} \in \mathrm{bd}(C_1 \cap C_2)$. For $i \in \{1,2\}$, either $\boldsymbol{x} \in \mathrm{bd}(C_i)$, in which case inward flow gives $-\boldsymbol{g}(\boldsymbol{x}) \in T_{C_i}(\boldsymbol{x})$, or $\boldsymbol{x} \in \mathrm{int}(C_i)$, in which case $T_{C_i}(\boldsymbol{x}) = \mathbb{R}^d$. Either way, $-\boldsymbol{g}(\boldsymbol{x}) \in T_{C_1}(\boldsymbol{x}) \cap T_{C_2}(\boldsymbol{x}) = T_{C_1\cap C_2}(\boldsymbol{x})$ by (ii). □

**Example: Bounded Isotonic Cone.** Suppose Assumption 2.1 holds with $\gamma = 0$, so $S_t \le D$ for all $t$, and let $R \ge D$. The set $\mathcal{K}_R := \mathcal{K} \cap \{\boldsymbol{x} \in \mathbb{R}^d : x_j \le R \text{ for all } j\}$ enforces both isotonicity and an upper bound. To verify inward flow on $\mathcal{K}_R$, by Proposition C.2 it suffices to verify it on the upper-bound polyhedron $C_R := \{\boldsymbol{x} : x_j \le R\}$. At a boundary point with $x_j = R \ge D \ge S_t$, the subgradient coordinate is $g_{t,j} = \alpha_j > 0$, so $-g_{t,j} < 0$, which lies in the half-space $v_j \le 0$ defining the tangent cone of $C_R$ along $x_j = R$. Both $\mathcal{K}$ and $C_R$ are polyhedral, so (ii) holds, and Proposition C.2 applies. All bounds in Section 3 carry over with $\mathcal{K}$ replaced by $\mathcal{K}_R$, with the additional benefit that the played radii are clipped at the largest plausible score.

### C.6. Tighter Regret Implies Tighter Coverage

The following monotonicity formalizes the second corollary of Theorem 3.4; it generalizes the scalar argument of Liu et al. (2026, Sec. 3) to the vector setting and underlies the design of the case studies in Sections 3.1–3.2.

**Proposition C.3.** *Let $F_T^A, F_T^B : \mathbb{R}^d \to \mathbb{R}$ be two valid linearized regret bounds for two base learners $\mathsf{Alg}_A, \mathsf{Alg}_B$ that satisfy $F_T^A(\boldsymbol{u}) \le F_T^B(\boldsymbol{u})$ for all $\boldsymbol{u} \in \mathbb{R}^d$. Then their Fenchel conjugates satisfy $(F_T^A)^\star(\boldsymbol{z}) \ge (F_T^B)^\star(\boldsymbol{z})$ for all $\boldsymbol{z} \in \mathbb{R}^d$, and the coverage-error level sets $\mathcal{Z}_T$ are nested:*

$$\mathcal{Z}_T(F_T^A) \;\subseteq\; \mathcal{Z}_T(F_T^B).$$

*Consequently, if $\mathsf{Alg}_A$ admits a uniformly tighter regret bound than $\mathsf{Alg}_B$, the worst-case miscoverage region delivered by Theorem 3.4 for $\mathsf{Alg}_A$ is contained in that for $\mathsf{Alg}_B$.*

*Proof.* Fix $\boldsymbol{z} \in \mathbb{R}^d$. By definition of the Fenchel conjugate,

$$(F_T^A)^\star(\boldsymbol{z}) = \sup_{\boldsymbol{u}\in\mathbb{R}^d} \big[\langle \boldsymbol{z}, \boldsymbol{u}\rangle - F_T^A(\boldsymbol{u})\big] \geq \sup_{\boldsymbol{u}\in\mathbb{R}^d} \big[\langle \boldsymbol{z}, \boldsymbol{u}\rangle - F_T^B(\boldsymbol{u})\big] = (F_T^B)^\star(\boldsymbol{z}),$$

using $-F_T^A(\boldsymbol{u}) \geq -F_T^B(\boldsymbol{u})$ pointwise. For every $\boldsymbol{z} \in \mathcal{Z}_T(F_T^A)$, $(F_T^B)^\star(-\boldsymbol{z}) \leq (F_T^A)^\star(-\boldsymbol{z}) \leq B_T$, so $\boldsymbol{z} \in \mathcal{Z}_T(F_T^B)$. □

This makes precise the operational principle invoked in the main text: minimizing the linearized regret of the base learner directly tightens the coverage-error region of the wrapped algorithm.

### C.7. Coordinate-Wise UP-OCP: Per-Level and Mean-Absolute Proofs

We first record the analytic facts about the Bernstein-type function used by UP-OCP, then prove Theorem 3.6 and Theorem C.6. Throughout, the Fenchel conjugate of the UP-OCP regret bound at level $\alpha$ has the form $(F_T^{(\alpha)})^\star(|\theta|) = \exp(h_\alpha(\theta)) - 1$ with

$$h_\alpha(\theta) := \frac{\theta^2}{2\alpha(1-\alpha)T + (2/3)\theta} - \frac{1}{2}\log(2(T+1));$$

see (Liu et al., 2026, Theorem 3.4).

**Lemma C.4** (Bernstein inversion)**.** *Let $\sigma, \theta \geq 0$ and $c > 0$. If $\theta^2/\big(2\sigma T + (2/3)\theta\big) \leq c$, then $\theta \leq \sqrt{2\sigma T c} + 2c/3$.*

*Proof.* Rearranging gives $\theta^2 - (2c/3)\theta - 2\sigma T c \leq 0$. The quadratic formula yields $\theta \leq c/3 + \sqrt{c^2/9 + 2\sigma T c} \leq c/3 + c/3 + \sqrt{2\sigma T c} = 2c/3 + \sqrt{2\sigma T c}$. □

**Lemma C.5.** *For every fixed $\sigma > 0$ and $T > 0$, the map $\theta \mapsto \theta^2/(2\sigma T + (2/3)\theta)$ is convex and non-decreasing on $[0, \infty)$.*

*Proof.* Write $f(\theta) = \theta^2/(a + b\theta)$ with $a = 2\sigma T$, $b = 2/3$. Direct calculation gives $f'(\theta) = \theta(2a + b\theta)/(a + b\theta)^2 \geq 0$ and $f''(\theta) = 2a^2/(a + b\theta)^3 > 0$. □

*Proof of Theorem 3.6.* Substituting the form of $(F_T^{(j)})^\star$ into the coordinate-wise instance of (5),

$$\sum_{j=1}^{d} \exp(h_{\alpha_j}(|G_{T,j}|)) \leq B_T + d. \tag{11}$$

Each term is non-negative, so we may drop all but the $j$-th: $\exp(h_{\alpha_j}(|G_{T,j}|)) \leq B_T + d$, equivalently

$$\frac{|G_{T,j}|^2}{2\alpha_j(1-\alpha_j)T + (2/3)\,|G_{T,j}|} \leq \log(B_T + d) + \tfrac{1}{2}\log(2(T+1)) = T\hat{\varepsilon}_{T,j}.$$

Lemma C.4 with $\sigma = \alpha_j(1-\alpha_j)$ and $c = T\hat{\varepsilon}_{T,j}$ gives $|G_{T,j}| \leq T\sqrt{2\alpha_j(1-\alpha_j)\hat{\varepsilon}_{T,j}} + (2/3)T\hat{\varepsilon}_{T,j}$. Dividing by $T$ and using $2/3 < 1$ yields (7). The asymptotic claim uses $B_T \leq dDT^{1+\gamma}$ to give $\hat{\varepsilon}_{T,j} = \mathcal{O}((\log d + \log T)/T)$ and hence $\max_j |\delta_{T,j}| = \mathcal{O}(\sqrt{(\log d + \log T)/T})$. □

The per-level guarantee (7) can be aggregated into a mean-absolute coverage bound across levels. By dropping the variance budgets at individual levels in favour of their envelope $\sigma_{\mathcal{A}}$ and applying Jensen's inequality to the resulting sum, one obtains a single $\ell_1$-style bound whose rate depends on the level set only through $\sigma_{\mathcal{A}}$ and $\log d$. The bound below is a vector analogue of Liu et al. (2026, Theorem 3.4) and, by norm equivalence, improves any $\ell_2$ bound by a factor of $1/\sqrt{d}$ on the same quantity (cf. Remark C.7 below).

**Theorem C.6** ($\ell_1$ miscoverage error of coordinate-wise `UP-OCP`)**.** *Run Algorithm 2, and let $\sigma_{\mathcal{A}} := \max_{i\in[d]} \alpha_i(1-\alpha_i) \in [0, 1/4]$ and $C_T := B_T + d$. Then*

$$\frac{1}{d}\sum_{i=1}^{d} |\delta_{T,i}| \;\leq\; \sqrt{\frac{2\sigma_{\mathcal{A}}\Lambda_T}{T}} + \frac{2\Lambda_T}{3T}, \qquad \Lambda_T \;:=\; \log\frac{C_T}{d} + \frac{1}{2}\log\big(2(T+1)\big). \tag{12}$$

*Under Assumption 2.1, $\Lambda_T = \mathcal{O}((1+\gamma)\log T)$, hence $\frac{1}{d}\sum_{i=1}^{d} |\delta_{T,i}| = \mathcal{O}\big(\sqrt{\sigma_{\mathcal{A}}(1+\gamma)\log T/T}\big)$.*

*Remark* C.7 (Comparison with the $\ell_2$-bound of Ding et al. (2025)). By norm equivalence, $\frac{1}{d}\|\boldsymbol{\delta}_T\|_1 \le \frac{1}{\sqrt{d}}\|\boldsymbol{\delta}_T\|_2$, so (12) delivers a $1/\sqrt{d}$ improvement over any $\ell_2$ bound on the mean-absolute miscoverage. Combined with Theorem 3.6, this gives a strictly tighter aggregate guarantee than the OGD-based `MultiQT`.

*Proof of Theorem C.6.* Substituting the form of $(F_T^{(j)})^\star$ into (11),

$$\sum_{j=1}^{d} \exp\big(h_{\alpha_j}(|G_{T,j}|)\big) \le B_T + d = C_T.$$

By definition of $\sigma_{\mathcal{A}}$, $\alpha_j(1-\alpha_j) \le \sigma_{\mathcal{A}}$ for every $j$, so $h_{\alpha_j}(\theta) \ge h_{\sigma_{\mathcal{A}}}(\theta)$ for every $\theta \ge 0$ (where we abuse notation and set $h_{\sigma_{\mathcal{A}}}(\theta) := \theta^2/(2\sigma_{\mathcal{A}}T + (2/3)\theta) - \frac{1}{2}\log(2(T+1))$). Hence

$$\sum_{j=1}^{d} \exp\big(h_{\sigma_{\mathcal{A}}}(|G_{T,j}|)\big) \le C_T.$$

By Lemma C.5, $\theta \mapsto h_{\sigma_{\mathcal{A}}}(\theta) + \frac{1}{2}\log(2(T+1))$ is convex; the exponential is increasing and convex; so the composition $\theta \mapsto \exp(h_{\sigma_{\mathcal{A}}}(\theta))$ is convex. Jensen's inequality applied to the uniform measure on $[d]$ gives

$$d\,\exp\big(h_{\sigma_{\mathcal{A}}}(\bar{G}_T)\big) \le \sum_{j=1}^{d} \exp\big(h_{\sigma_{\mathcal{A}}}(|G_{T,j}|)\big) \le C_T, \qquad \bar{G}_T := \frac{1}{d}\sum_{j=1}^{d} |G_{T,j}|\,.$$

Taking logarithms and rearranging,

$$\frac{\bar{G}_T^2}{2\sigma_{\mathcal{A}}T + (2/3)\bar{G}_T} \le \log\frac{C_T}{d} + \frac{1}{2}\log(2(T+1)) = \Lambda_T.$$

Lemma C.4 with $\sigma = \sigma_{\mathcal{A}}$, $c = \Lambda_T$ gives $\bar{G}_T \le \sqrt{2\sigma_{\mathcal{A}}T\Lambda_T} + 2\Lambda_T/3$. Dividing by $T$ and noting $\frac{1}{d}\sum_j |\delta_{T,j}| = \bar{G}_T/T$ yields (12). For the asymptotic claim, $C_T/d = B_T/d + 1 \le DT^{1+\gamma} + 1$ by Assumption 2.1, hence $\Lambda_T = \mathcal{O}(\log T)$, and the right-hand side of (12) is $\mathcal{O}(\sqrt{\sigma_{\mathcal{A}}\log T/T})$. □

### C.8. Continuous Quantile Levels

We extend the framework from a finite collection $\mathcal{A}$ to a compact interval $\mathcal{X} \subseteq (0,1)$. Fix a finite Borel reference measure $\mu$ with full support and work in the Hilbert space $\mathcal{H} := L^2(\mathcal{X},\mu)$ with inner product $\langle f,h\rangle := \int_{\mathcal{X}} f(\alpha)h(\alpha)\,d\mu(\alpha)$. The radius forecast is now a function $b_t \in \mathcal{H}$, and the score $S_t$ is a scalar with $S_t \le DT^\gamma$ as in Assumption 2.1. Define the integrated pinball loss

$$\mathcal{L}_t(b) := \int_{\mathcal{X}} \ell^{(1-\alpha)}(b(\alpha), S_t)\,d\mu(\alpha),$$

which is convex on $\mathcal{H}$ and has subgradient $g_t(\alpha) := \mathbf{1}\{b_t(\alpha) \ge S_t\} - (1-\alpha)$ in $L^\infty(\mathcal{X},\mu) \subseteq \mathcal{H}$. The isotonic constraint becomes the cone of non-increasing functions:

$$\mathcal{K} := \{b \in \mathcal{H} : b \text{ admits a non-increasing representative on } \mathcal{X}\}\,.$$

$\mathcal{K}$ is a closed convex cone in $\mathcal{H}$, and the $L^2$-projection $\Pi_{\mathcal{K}}$ is well-defined by the Hilbert projection theorem. Continuous PAVA characterizes $\Pi_{\mathcal{K}}(b)$ as $-\frac{d\hat{F}}{d\mu}$, where $\hat{F}$ is the greatest concave majorant of $F(x) := -\int_{\alpha_{\min}}^{x} b(\alpha)\,d\mu(\alpha)$ (Robertson et al., 1988).

**Inward Flow, Regret Preservation, Central Inequality.** The proofs of Lemma 3.2, Proposition 3.3, and Theorem 3.4 use only the inner-product structure of $\mathcal{H}$ and so transfer verbatim. Inward flow of $(\mathcal{L}_t, \mathcal{K})$ on $\mathcal{H}$ is verified by the same case analysis as in Appendix C.3, applied pointwise. The central inequality reads

$$F_T^\star(-G_T) \le B_T, \qquad G_T := \sum_{t=1}^{T} g_t \in \mathcal{H}, \qquad B_T := \sum_{t=1}^{T} \mathcal{L}_t(0). \tag{13}$$

Under Assumption 2.1, $\mathcal{L}_t(0) \le \mu(\mathcal{X}) \cdot S_t \le \mu(\mathcal{X})\,DT^\gamma$, so $B_T \le \mu(\mathcal{X})\,DT^{1+\gamma}$, depending on $\mathcal{X}$ only through $\mu(\mathcal{X})$ and *not* through any notion of dimensionality.

**$L^2$ Coverage via OGD.** Take Alg to be OGD on $\mathcal{H}$ with regularizer $\psi(f) = \frac{1}{2}\|f\|_{\mathcal{H}}^2$ and step size $\eta > 0$. The regret bound is $F_T(u) = \|u\|_{\mathcal{H}}^2/(2\eta) + (\eta/2)\sum_t \|g_t\|_{\mathcal{H}}^2$, with conjugate $F_T^\star(\theta) = (\eta/2)\|\theta\|_{\mathcal{H}}^2 - (\eta/2)\sum_t \|g_t\|_{\mathcal{H}}^2$. Since $\|g_t\|_{\mathcal{H}}^2 = \int_{\mathcal{X}} g_t(\alpha)^2\, d\mu(\alpha) \le \mu(\mathcal{X})$, Theorem 3.4 gives $\|G_T\|_{\mathcal{H}}^2 \le 2B_T/\eta + \mu(\mathcal{X})\,T$. Choosing $\eta = 2B_T/(\mu(\mathcal{X})T)$ to balance the two terms yields

$$\|\delta_T\|_{L^2(\mathcal{X},\mu)} \le \sqrt{\frac{2\mu(\mathcal{X})}{T}}, \qquad \delta_T := G_T/T. \tag{14}$$

The bound is dimension-free: it depends on the size of the level set only through $\mu(\mathcal{X})$.

**$L^q$ Coverage via OMD.** Take Alg to be OMD with regularizer $\psi(f) = \frac{1}{2}\|f\|_{L^p}^2$ for $p \in (1,2]$, $q = p/(p-1)$. The regularizer is $(p-1)$-strongly convex with respect to $\|\cdot\|_{L^p}$ (Shalev-Shwartz, 2007, Sec. 4), by the uniform convexity of $L^p$ in this range. The regret bound is $F_T(u) = \|u\|_{L^p}^2/(2\eta) + (\eta/(2(p-1)))\sum_t \|g_t\|_{L^q}^2$, with conjugate $F_T^\star(\theta) = (\eta/2)\|\theta\|_{L^q}^2 - (\eta/(2(p-1)))\sum_t \|g_t\|_{L^q}^2$. Since $\|g_t\|_{L^q} \le \mu(\mathcal{X})^{1/q}$,

$$\|\delta_T\|_{L^q} \le \frac{1}{T}\sqrt{\frac{2B_T}{\eta} + \frac{\mu(\mathcal{X})^{2/q}\,T}{p-1}}. \tag{15}$$

Balancing the two terms gives $\|\delta_T\|_{L^q} \le \sqrt{2(q-1)\mu(\mathcal{X})^{2/q}/T}$.

*Remark* C.8 (Optimal $q$ depends on $\mu(\mathcal{X})$). Differentiating $g(q) := (q-1)\mu(\mathcal{X})^{2/q}$ in $q \in [2,\infty)$ shows that $g$ is monotone in $q$ unless $\mu(\mathcal{X}) > e$, in which case the unique stationary point satisfies $q^2 = 2(q-1)\log\mu(\mathcal{X})$ and is approximately $2\log\mu(\mathcal{X})$ for large $\mu(\mathcal{X})$. Since $\mathcal{X} \subseteq (0,1)$, $\mu(\mathcal{X}) \le 1 < e$, $g$ is increasing on $[2,\infty)$ and the optimum is $q = 2$, recovering the dimension-free $L^2$ bound (14). So OGD on $\mathcal{H}$ already delivers the optimal rate among $L^q$ bounds.

This corrects the optimization in the discrete-to-continuous translation: the discrete-case choice $q = 2\log d$ is optimal only when the analogue $\mu(\mathcal{X}) > e$ holds, which fails under Lebesgue and any normalized probability measure on $\mathcal{X}$. In those cases, the $L^2$ bound (14) *is* the right answer.

**Coordinate-Wise Universal Portfolio on a Continuum.** Running an independent UP-OCP instance at each level $\alpha \in \mathcal{X}$ with initial wealth 1, the regret factorizes as $F_T(u) = \int_{\mathcal{X}} F_T^{(\alpha)}(|u(\alpha)|)\, d\mu(\alpha)$, and similarly $F_T^\star(\theta) = \int_{\mathcal{X}} (F_T^{(\alpha)})^\star(|\theta(\alpha)|)\, d\mu(\alpha)$. The central inequality (13) becomes

$$\int_{\mathcal{X}} \exp\big(h_\alpha(|G_T(\alpha)|)\big)\, d\mu(\alpha) \le C_T := B_T + \mu(\mathcal{X}), \tag{16}$$

where $h_\alpha$ is as defined in Appendix C.7.

To control miscoverage along an arbitrary direction in the level set, we record the Donsker–Varadhan change-of-measure inequality and then derive a coverage bound for any probability measure $\nu$ on $\mathcal{X}$.

**Lemma C.9** (Donsker–Varadhan change-of-measure; Boucheron et al. 2013, Sec. 2.1)**.** *Let $(\mathcal{X},\mathcal{B})$ be a measurable space, $\bar\mu$ a probability measure on $\mathcal{X}$, and $\nu$ a probability measure with $\nu \ll \bar\mu$. Then for every bounded measurable $\Phi : \mathcal{X} \to \mathbb{R}$,*

$$\int_{\mathcal{X}} \Phi\, d\nu \le \mathrm{KL}(\nu\|\bar\mu) + \log\int_{\mathcal{X}} e^{\Phi}\, d\bar\mu.$$

Fix any probability measure $\nu \ll \bar\mu$ on $\mathcal{X}$ (with $\bar\mu := \mu/\mu(\mathcal{X})$), and define the $\nu$-variance envelope $\sigma_\nu := \sup_{\alpha\in\mathrm{supp}(\nu)} \alpha(1-\alpha)$ and the $\nu$-mean miscoverage $\bar G_\nu := \int_{\mathcal{X}} |G_T(\alpha)|\, d\nu(\alpha)$. Lemma C.9 with $\Phi(\alpha) = h_\alpha(|G_T(\alpha)|)$, combined with (16) (so $\log\int e^{\Phi}\, d\bar\mu \le \log(C_T/\mu(\mathcal{X}))$), yields

$$\int_{\mathcal{X}} h_\alpha(|G_T(\alpha)|)\, d\nu(\alpha) \le \mathrm{KL}(\nu\|\bar\mu) + \log\frac{C_T}{\mu(\mathcal{X})}. \tag{17}$$

Bounding $h_\alpha \ge h_{\sigma_\nu}$ $\nu$-almost everywhere and applying Jensen's inequality to the convex function $h_{\sigma_\nu}$ (Lemma C.5) under the probability measure $\nu$ gives $h_{\sigma_\nu}(\bar G_\nu) \le \int_{\mathcal{X}} h_{\sigma_\nu}(|G_T(\alpha)|)\, d\nu(\alpha) \le \mathrm{KL}(\nu\|\bar\mu) + \log(C_T/\mu(\mathcal{X}))$, i.e., $\bar G_\nu^2/(2\sigma_\nu T + (2/3)\bar G_\nu) \le R_\nu$ with $R_\nu := \mathrm{KL}(\nu\|\bar\mu) + \log(C_T/\mu(\mathcal{X})) + \frac{1}{2}\log(2(T+1))$. Lemma C.4 then yields the general coverage bound

$$\int_{\mathcal{X}} |\delta_T(\alpha)|\, d\nu(\alpha) \le \sqrt{\frac{2\sigma_\nu R_\nu}{T}} + \frac{2R_\nu}{3T}. \tag{18}$$

**Sub-Interval Coverage as a Special Case.** Take $\mathcal{I} \subseteq \mathcal{X}$ measurable with $\mu(\mathcal{I}) > 0$ and set $\nu = \mu|_{\mathcal{I}}/\mu(\mathcal{I})$. Then $\mathrm{KL}(\nu\|\bar{\mu}) = \log(\mu(\mathcal{X})/\mu(\mathcal{I}))$, $\sigma_\nu = \sigma_{\mathcal{I}} := \sup_{\alpha\in\mathcal{I}} \alpha(1-\alpha)$, and $\bar{G}_\nu = \bar{G}_{\mathcal{I}} := \frac{1}{\mu(\mathcal{I})}\int_{\mathcal{I}} |G_T(\alpha)|\, d\mu(\alpha)$, so (18) reads

$$\frac{1}{\mu(\mathcal{I})}\int_{\mathcal{I}} |\delta_T(\alpha)|\, d\mu(\alpha) \le \sqrt{\frac{2\sigma_{\mathcal{I}} R_{\mathcal{I}}}{T}} + \frac{2R_{\mathcal{I}}}{3T}, \qquad R_{\mathcal{I}} := \log\frac{C_T}{\mu(\mathcal{I})} + \tfrac{1}{2}\log(2(T+1)). \tag{19}$$

Specializing further to $\mathcal{I} = \mathcal{X}$ recovers the global mean-absolute miscoverage of Theorem C.6; specializing to a small interval near $\alpha = 0$ or $\alpha = 1$ exploits the smaller variance envelope $\sigma_{\mathcal{I}}$, mirroring the level-adaptive behavior of the discrete case.

**Reduction to the Discrete Setting.** Specializing $\mathcal{X} = \mathcal{A} = \{\alpha_1, \dots, \alpha_d\}$ and $\mu$ the counting measure recovers $\mu(\mathcal{X}) = d$ and reduces every continuous bound to its finite analogue: (14) becomes the OGD bound, (15) becomes Corollary 3.5, (16) becomes (11), and (19) with $\mathcal{I} = \mathcal{X}$ recovers Theorem C.6. In particular, the dichotomy in Remark C.8 explains why the discrete optimization $q = 2\log d$ *is* valid (since the counting measure makes $\mu(\mathcal{X}) = d \ge 3 > e$) but does not transfer to Lebesgue.

## D. Deferred Proofs for Section 4

This appendix collects the proofs deferred from Section 4: the WIS–pinball identity (8) and the per-round structural reduction (Theorem 4.1).

### D.1. Pinball Loss as a Weighted Interval Score (Equation (8))

We verify the identity (8): for symmetric central intervals $[\hat{Y}_t - x_{t,i}, \hat{Y}_t + x_{t,i}]$ at miscoverage levels $\alpha_1 < \cdots < \alpha_d$ evaluated at the realized response $Y_t$,

$$\mathrm{WIS}_{\mathcal{A}}\Big(\Big\{[\hat{Y}_t - x_{t,i}, \hat{Y}_t + x_{t,i}]\Big\}_{i=1}^{d}, Y_t\Big) \;=\; 2\mathcal{L}_t(\boldsymbol{x}_t) + 2S_t\sum_{i=1}^{d}\alpha_i.$$

First of all, observe that $\mathrm{WIS}_{\mathcal{A}}\big(\{[l_{\alpha_i}, u_{\alpha_i}]\}_{i=1}^{d}, y\big) = \sum_{i=1}^{d}\alpha_i\,\mathrm{IS}_{\alpha_i}([l_{\alpha_i}, u_{\alpha_i}], y)$, where $\mathrm{IS}_{\alpha_i}([l_{\alpha_i}, u_{\alpha_i}], y) := (u_{\alpha_i} - l_{\alpha_i}) + \frac{2}{\alpha_i}\,\mathrm{dist}\,(y, [l_{\alpha_i}, u_{\alpha_i}])$.

With $D := Y_t - \hat{Y}_t$ and $S_t = |D|$, we have

$$\max(l_{\alpha_i} - Y_t, 0) = \max(-D - x_{t,i}, 0), \quad \text{and} \quad \max(Y_t - u_{\alpha_i}, 0) = \max(D - x_{t,i}, 0).$$

Because $D$ has a definite sign and $x_{t,i} \ge 0$, exactly one of these two non-negative quantities can be non-zero, and it equals $\max(|D| - x_{t,i}, 0) = \max(S_t - x_{t,i}, 0)$. Hence

$$\mathrm{IS}_{\alpha_i}([\hat{Y}_t - x_{t,i}, \hat{Y}_t + x_{t,i}], Y_t) = 2x_{t,i} + \frac{2}{\alpha_i}\max(S_t - x_{t,i}, 0),$$

and therefore $\alpha_i\,\mathrm{IS}_{\alpha_i} = 2\alpha_i x_{t,i} + 2\max(S_t - x_{t,i}, 0)$. Summing over $i \in [d]$,

$$\mathrm{WIS}_{\mathcal{A}} = 2\sum_{i=1}^{d}\alpha_i x_{t,i} + 2\sum_{i=1}^{d}\max(S_t - x_{t,i}, 0).$$

The pointwise identity $\ell^{(1-\alpha)}(x, S) = \alpha(x - S) + \max(S - x, 0)$ (a direct case analysis on $x \gtreqless S$) gives, after summation,

$$\mathcal{L}_t(\boldsymbol{x}_t) = \sum_{i=1}^{d}\alpha_i x_{t,i} + \sum_{i=1}^{d}\max(S_t - x_{t,i}, 0) - S_t\sum_{i=1}^{d}\alpha_i,$$

which is (8). Combining the two displays yields (8).

### D.2. Per-Round Structural Reduction (Theorem 4.1)

The proof of Theorem 4.1 relies on two structural lemmas about the PAVA decomposition of the isotonic projection. Throughout, fix $\tilde{\boldsymbol{x}} \in \mathbb{R}^d$, let $\boldsymbol{x} := \Pi_{\mathcal{K}}(\tilde{\boldsymbol{x}})$ with PAVA blocks $B_1, \dots, B_p$ (each $B_k = \{l_k, l_k+1, \dots, r_k\}$), and let $\bar{\tilde{x}}_{B_k} := |B_k|^{-1} \sum_{i \in B_k} \tilde{x}_i$ denote the block average, so $x_j = \bar{\tilde{x}}_{B_k}$ for all $j \in B_k$.

**Lemma D.1** (Partial-average property of PAVA)**.** *For every block $B_k = \{l, \dots, r\}$ and every $j \in [l, r]$,*

$$\frac{1}{j-l+1} \sum_{i=l}^{j} \tilde{x}_i \;\le\; \bar{\tilde{x}}_{B_k}.$$

*Proof.* Recall that $\boldsymbol{x}$ solves the quadratic program $\min_{\boldsymbol{z} \in \mathcal{K}} \frac{1}{2} \|\boldsymbol{z} - \tilde{\boldsymbol{x}}\|_2^2$, where $\mathcal{K} = \{\boldsymbol{z} : z_1 \ge z_2 \ge \cdots \ge z_d\}$. Introducing non-negative Lagrange multipliers $\lambda_i \ge 0$ for the constraints $z_i \ge z_{i+1}$, $i \in [d-1]$ (with the convention $\lambda_0 = \lambda_d = 0$), the KKT stationarity condition at the optimum reads, for $i \in [d]$,

$$\tilde{x}_i - x_i = \lambda_{i-1} - \lambda_i.$$

Within block $B_k = \{l, \dots, r\}$, the projected entries are constant, so the constraints $x_{l-1} > x_l$ and $x_r > x_{r+1}$ are strict at the boundary of the block (with the convention that $x_0 = +\infty$ and $x_{d+1} = -\infty$). Complementary slackness then forces $\lambda_{l-1} = \lambda_r = 0$. Summing the stationarity equation over $i = l, \dots, j$ and telescoping,

$$\sum_{i=l}^{j} (\tilde{x}_i - \bar{\tilde{x}}_{B_k}) = \sum_{i=l}^{j} (\lambda_{i-1} - \lambda_i) = \lambda_{l-1} - \lambda_j = -\lambda_j \;\le\; 0,$$

where the last inequality uses dual feasibility $\lambda_j \ge 0$. Dividing by $j - l + 1$ gives the claim. □

**Lemma D.2** (Weighted-sum reduction within a block)**.** *Let $w_l \le w_{l+1} \le \cdots \le w_r$ be non-decreasing weights on a block $B_k = \{l, \dots, r\}$. Then*

$$\sum_{i \in B_k} w_i \, x_i \;=\; \bar{\tilde{x}}_{B_k} \sum_{i \in B_k} w_i \;\le\; \sum_{i \in B_k} w_i \, \tilde{x}_i.$$

*The inequality is strict whenever $\{\tilde{x}_i\}_{i \in B_k}$ is non-constant and $w$ is strictly increasing on $[l, r]$.*

*Proof.* Let $c_i := \tilde{x}_i - \bar{\tilde{x}}_{B_k}$, so $\sum_{i=l}^{r} c_i = 0$ and, by Lemma D.1, $\sum_{i=l}^{j} c_i \le 0$ for every $j \in [l, r]$. Abel summation yields

$$\sum_{i=l}^{r} w_i c_i = w_r \cdot \underbrace{\sum_{i=l}^{r} c_i}_{=0} - \sum_{j=l}^{r-1} (w_{j+1} - w_j) \sum_{i=l}^{j} c_i = -\sum_{j=l}^{r-1} \underbrace{(w_{j+1} - w_j)}_{\ge 0} \underbrace{\sum_{i=l}^{j} c_i}_{\le 0} \;\ge\; 0,$$

which gives $\sum_{i=l}^{r} w_i \tilde{x}_i \ge \bar{\tilde{x}}_{B_k} \sum_{i=l}^{r} w_i = \sum_{i=l}^{r} w_i x_i$. For strictness, observe that if $\{\tilde{x}_i\}_{i \in B_k}$ is non-constant, then $\sum_{i=l}^{j} c_i < 0$ for some $j \in [l, r-1]$; if in addition $w_{j+1} > w_j$, the corresponding term in the Abel sum is strictly positive. □

For convenience, we restate Theorem 4.1 below, attaching labels to its individual reduction inequalities so that the proof can refer to them.

**Theorem D.3** (Per-round reduction; restatement of Theorem 4.1)**.** *For every round $t$, the hidden iterate $\tilde{\boldsymbol{x}}_t \in \mathbb{R}^d$ and the played iterate $\boldsymbol{x}_t = \Pi_{\mathcal{K}}(\tilde{\boldsymbol{x}}_t) \in \mathcal{K}$ satisfy the per-round sharpness and coverage reductions*

$$\sum_{i=1}^{d} \alpha_i x_{t,i} \le \sum_{i=1}^{d} \alpha_i \tilde{x}_{t,i}, \tag{20}$$

$$\sum_{i=1}^{d} \max(S_t - x_{t,i}, 0) \le \sum_{i=1}^{d} \max(S_t - \tilde{x}_{t,i}, 0), \tag{21}$$

*and consequently the per-round loss reduction*

$$\mathcal{L}_t(\boldsymbol{x}_t) \;\le\; \mathcal{L}_t(\tilde{\boldsymbol{x}}_t). \tag{22}$$

*Both* (20) *and* (22) *are strict whenever* $\tilde{\boldsymbol{x}}_t \notin \mathcal{K}$.

*Proof of Theorem 4.1.* By the WIS decomposition (8), the loss reduction (22) is implied by the sharpness reduction (20) and the coverage reduction (21), since the constant $S_t \sum_i \alpha_i$ cancels. We prove each block by block.

*Sharpness.* On any block $B_k$, the miscoverage levels $\alpha_l < \alpha_{l+1} < \cdots < \alpha_r$ are strictly increasing (by our convention). Lemma D.2 with $w_i = \alpha_i$ then gives $\sum_{i\in B_k} \alpha_i x_{t,i} \le \sum_{i\in B_k} \alpha_i \tilde{x}_{t,i}$, with strict inequality whenever $\{\tilde{x}_{t,i}\}_{i\in B_k}$ is non-constant. PAVA pools coordinates only when consecutive (sub-)block averages violate the order, which forces a strict difference between elements; hence any block of size $|B_k| \ge 2$ has non-constant entries. Summing over $k$ gives (20), with strict inequality whenever $\tilde{\boldsymbol{x}}_t \notin \mathcal{K}$ (equivalently, whenever some block has size $\ge 2$).

*Coverage.* The map $\phi(x) := \max(S_t - x, 0)$ is convex on $\mathbb{R}$. Jensen's inequality applied to $\tilde{\boldsymbol{x}}_t$ restricted to $B_k$ gives

$$\sum_{i\in B_k} \max(S_t - x_{t,i}, 0) \;=\; |B_k|\, \phi(\bar{\tilde{x}}_{B_k}) \;\le\; \sum_{i\in B_k} \phi(\tilde{x}_{t,i}) \;=\; \sum_{i\in B_k} \max(S_t - \tilde{x}_{t,i}, 0).$$

Summing over $k$ gives (21).

*Loss.* Summing (20) and (21) and applying (8) gives (22), which is strict whenever (20) is strict, i.e., whenever $\tilde{\boldsymbol{x}}_t \notin \mathcal{K}$. □

## E. Experimental Details and Additional Results

This appendix collects the dataset descriptions, forecasting-pipeline configurations, and complete results referenced from Section 5.

### E.1. Datasets and Forecasting Pipelines

**Synthetic Streams.** All synthetic scores are non-negative trajectories with i.i.d. Gaussian noise truncated at zero. Each trajectory is generated from a fixed base seed; the 10 seeds we use are $42, 43, \ldots, 51$.

- *Sinusoidal*: $S_t = \max(0, m + a\sin(2\pi t/p) + \varepsilon_t)$ with center $m = 11$ and amplitude $a = 9$ (so that $S_t$ oscillates between $\{2, 20\}$), period $p = 200$, $\varepsilon_t \sim \mathcal{N}(0, 0.3^2)$, $T = 3{,}000$. Stresses cyclic non-stationarity.
- *Mixed regime*: a polynomial trend overlaid with rare exponential spikes, smoothed by a 25-step rolling maximum (Appendix E.4). Stresses heavy tails.

**Real-World Streams.**

- *AAPL, AMZN, GOOGL*: daily close-prices of the three stocks, from Nguyen (2018), covering the period 2006-01-03 to 2017-12-29 ($T = 3{,}019$ trading days). Forecasting model: Prophet (Taylor & Letham, 2018) re-fit at every step on log-prices using all data up to time $t-1$ to produce a one-step-ahead forecast $\hat{Y}_t$. Nonconformity score $S_t = \left|Y_t - \hat{Y}_t\right|$.
- *Electricity (`elec2`)*: half-hourly New South Wales electricity demand from (Harries, 1999), restricted to a non-stationary regime; one-step-ahead forecasts produced by an autoregressive model.

### E.2. Detailed Update Rules

This subsection records the explicit update rules of the four method families compared in Section 5: coordinate-wise universal portfolio, $p$-norm online mirror descent, coordinate-wise Krichevsky–Trofimov bettor, and `MultiQT`. The first three are base online learners that we run *both* in their vanilla form and through the lazy projection wrapper of Algorithm 1; `MultiQT`, which is by construction a projected algorithm, is presented in its single canonical form. All wrapped variants share the same outer loop: receive the hidden iterate $\tilde{\boldsymbol{x}}_t$ from the base learner, project to obtain $\boldsymbol{x}_t = \Pi_{\mathcal{K}}(\tilde{\boldsymbol{x}}_t)$, observe the score $S_t$, compute the subgradient $\boldsymbol{g}_t$ at $\boldsymbol{x}_t$ via (2), and feed $\boldsymbol{g}_t$ back to the base learner. The unprojected variants are obtained by replacing the projected play $\boldsymbol{x}_t = \Pi_{\mathcal{K}}(\tilde{\boldsymbol{x}}_t)$ by $\boldsymbol{x}_t = \tilde{\boldsymbol{x}}_t$ throughout, leaving the rest of the update untouched.

#### E.2.1. Coordinate-Wise Universal Portfolio (UP-OCP)

**Algorithm.** Coordinate-wise UP-OCP runs an independent scalar UP-OCP instance (Liu et al., 2026) at each miscoverage level $\alpha_j$ and synchronizes only at prediction time through the projection. At level $\alpha_j$, the scalar instance treats the betting game as a two-asset *conformal market* (Definition 4.1 of Liu et al. 2026) with returns $w_{t,1}^{(j)} = -g_{t,j}/\alpha_j + 1$ and $w_{t,2}^{(j)} = 1 + g_{t,j}/(1-\alpha_j)$, and chooses portfolio weight $\lambda_t^{(j)}$ as the wealth-weighted average over $[0,1]$ under the Beta$(1/2, 1/2)$ prior. Specializing the universal-portfolio update to this market admits the closed form

$$\lambda_t^{(j)} \;=\; \frac{1}{t}\left(\sum_{s=1}^{t-1} \mathbf{1}\left\{g_{s,j} = -(1-\alpha_j)\right\} + \tfrac{1}{2}\right) \;=\; \frac{N_{t-1}^{(j)} + 1/2}{t},$$

where $N_{t-1}^{(j)}$ counts past per-level miscoverage events. The hidden iterate is the wealth-to-radius mapping $\tilde{x}_{t,j} = W_{t-1}^{(j)}\,(\lambda_t^{(j)} - \alpha_j)/(\alpha_j(1-\alpha_j))$, and the wealth recursion $W_t^{(j)} = W_{t-1}^{(j)} \cdot (\lambda_t^{(j)} w_{t,1}^{(j)} + (1-\lambda_t^{(j)}) w_{t,2}^{(j)})$ simplifies, on the binary outcome $g_{t,j} \in \{-(1-\alpha_j), \alpha_j\}$, to the multiplicative updates in Algorithm 2. The coordinate-wise nature is essential: the per-level wealth and miscoverage counter make every level adapt to its own variance budget $\alpha_j(1-\alpha_j)$, which is what ultimately produces the $\sqrt{\log d}$ rate of Theorem 3.6.

**Hyperparameters.** UP-OCP is parameter-free; the initial wealth is fixed at 1 for every dataset and every level (see Section 5), matching the theoretical analysis. No tuning is performed.

**Projected vs. Unprojected.** The projected variant UP-OCP-Iso plays $\boldsymbol{x}_t = \Pi_{\mathcal{K}}(\tilde{\boldsymbol{x}}_t)$ on Line 2 of Algorithm 2 below; the unprojected variant simply plays $\boldsymbol{x}_t = \tilde{\boldsymbol{x}}_t$. Both feed the gradient evaluated at the played iterate back into the per-level wealth recursions.

**Algorithm 2** Coordinate-wise UP-OCP with lazy projection.

1: **Input:** miscoverage levels $\alpha_1 < \cdots < \alpha_d$
2: Initialize $W_0^{(j)} \leftarrow 1$ and miscoverage count $N^{(j)} \leftarrow 0$ for each $j \in [d]$
3: **for** $t = 1, 2, \ldots$ **do**
4:  **for** $j = 1, \ldots, d$ (in parallel) **do**
5:  $\lambda_t^{(j)} \leftarrow (N^{(j)} + 1/2)/t$
6:  $\tilde{x}_{t,j} \leftarrow W_{t-1}^{(j)} \cdot (\lambda_t^{(j)} - \alpha_j)/(\alpha_j(1-\alpha_j))$
7:  **end for**
8:  $\boldsymbol{x}_t \leftarrow \Pi_{\mathcal{K}}(\tilde{\boldsymbol{x}}_t)$  (*unprojected:* $\boldsymbol{x}_t \leftarrow \tilde{\boldsymbol{x}}_t$); observe $S_t$
9:  **for** $j = 1, \ldots, d$ (in parallel) **do**
10:  **if** $S_t > x_{t,j}$ **then**
11:  $N^{(j)} \leftarrow N^{(j)} + 1$; $W_t^{(j)} \leftarrow W_{t-1}^{(j)} \cdot \lambda_t^{(j)}/\alpha_j$
12:  **else**
13:  $W_t^{(j)} \leftarrow W_{t-1}^{(j)} \cdot (1-\lambda_t^{(j)})/(1-\alpha_j)$
14:  **end if**
15:  **end for**
16: **end for**

#### E.2.2. $p$-Norm Online Mirror Descent

**Algorithm.** $p$-norm OMD is online mirror descent on $\mathbb{R}^d$ with regularizer $\psi(\boldsymbol{x}) = \frac{1}{2}\|\boldsymbol{x}\|_p^2$ for $p \in (1, 2]$ and conjugate exponent $q = p/(p-1)$. Maintaining a dual iterate $\boldsymbol{\theta}_t \in \mathbb{R}^d$ initialized at $\mathbf{0}$, the primal mirror map is the gradient of the conjugate regularizer (Shalev-Shwartz, 2007; Orabona, 2019), which yields the explicit closed form

$$\tilde{x}_{t,j} \;=\; \|\boldsymbol{\theta}_t\|_q^{2-q}\,\operatorname{sign}(\theta_{t,j})\,|\theta_{t,j}|^{q-1}, \qquad j = 1, \ldots, d,$$

followed by the standard dual update $\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \eta \boldsymbol{g}_t$. The choice $p = 2\log d/(2\log d - 1)$ recommended by Corollary 3.5 interpolates between OGD ($p = 2$) and $\ell_\infty$-style updates as $d$ grows: it makes the regularizer $(p-1)$-strongly convex with respect to $\|\cdot\|_p$ while keeping the conjugate gradient norm $\|\boldsymbol{g}_t\|_q^2$ at most $d^{2/q} = e$, eliminating the polynomial dimension dependence of $\ell_2$ regret. The case $p = 2$ recovers projected OGD on the isotonic cone and is the closest direct comparison with MultiQT.

**Hyperparameters and Tuning.** The constant step size $\eta$ is the only tuned knob. Following standard practice for adaptive conformal baselines (Podkopaev et al., 2024; Ding et al., 2025), we run a per-dataset grid search over $\eta \in \{2^{-2}, 2^{-1}, 1, 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024\}$ and report the value whose mean-width vs. empirical-coverage Pareto frontier dominates the others on $\alpha \in [0.03, 0.25]$. For both AAPL and the sinusoidal stream the selected $\eta$ is 64. We emphasize that this oracle-style selection is a best-case, not realistic, deployment of the method: a practitioner without access to the full stream cannot select $\eta$ this way. The exponent $p$ is fixed at $2\log d/(2\log d - 1)$ from theory and not tuned. Initialization is $\boldsymbol{\theta}_1 = \mathbf{0}$.

**Projected vs. Unprojected.** The projected variant $p$-norm-OMD-ISO plays $\boldsymbol{x}_t = \Pi_{\mathcal{K}}(\tilde{\boldsymbol{x}}_t)$ on Line 3 below; the unprojected variant plays $\boldsymbol{x}_t = \tilde{\boldsymbol{x}}_t$. The dual update is identical in both cases.

**Algorithm 3** $p$-norm online mirror descent with lazy projection.

1: **Input:** $p \in (1, 2]$ (with conjugate exponent $q := p/(p-1)$), step size $\eta > 0$. Set $\boldsymbol{\theta}_1 = \mathbf{0}$
2: **for** $t = 1, 2, \ldots$ **do**
3: $\quad \tilde{x}_{t,j} \leftarrow \|\boldsymbol{\theta}_t\|_q^{2-q} \operatorname{sign}(\theta_{t,j}) \cdot |\theta_{t,j}|^{q-1}$ for $j = 1, \ldots, d$
4: $\quad \boldsymbol{x}_t \leftarrow \Pi_{\mathcal{K}}(\tilde{\boldsymbol{x}}_t)$ (*unprojected:* $\boldsymbol{x}_t \leftarrow \tilde{\boldsymbol{x}}_t$); observe $S_t$; $\boldsymbol{g}_t \in \partial\mathcal{L}_t(\boldsymbol{x}_t)$ via (2)
5: $\quad \boldsymbol{\theta}_{t+1} \leftarrow \boldsymbol{\theta}_t - \eta\, \boldsymbol{g}_t$
6: **end for**

### E.2.3. Coordinate-Wise Krichevsky–Trofimov Bettor (KT)

**Algorithm.** Coordinate-wise KT (Podkopaev et al., 2024; Liu et al., 2026) runs an independent KT bettor at each level $\alpha_j$, treating the per-level subgradients as the outcomes of a coin-betting game. At level $\alpha_j$, the bettor maintains wealth $W_t^{(j)}$ and a betting fraction $\beta_t^{(j)} = -\Sigma_{t-1}^{(j)}/t$ where $\Sigma_{t-1}^{(j)} = \sum_{s=1}^{t-1} g_{s,j}$ is the cumulative subgradient (Podkopaev et al., 2024); the hidden iterate is $\tilde{x}_{t,j} = \beta_t^{(j)} W_{t-1}^{(j)}$ and the wealth follows $W_t^{(j)} = W_{t-1}^{(j)} \cdot (1 - \beta_t^{(j)} g_{t,j})$. Unlike UP-OCP, the KT bettor uses a closed-form betting rule that is min-max optimal for symmetric coin-betting games (i.e., $\alpha_j = 1/2$); it does *not* adapt to the asymmetry of the per-level pinball loss when $\alpha_j$ is small. This is the source of the under-performance of KT at high target coverage in our experiments, and is the reason UP-OCP achieves the strictly tighter rate $\sqrt{\alpha_j(1-\alpha_j)\log T/T}$.

**Hyperparameters.** KT is parameter-free; the initial wealth is fixed at 1 across all datasets and levels, matching the theoretical analysis. No tuning is performed.

**Projected vs. Unprojected.** The projected variant KT-ISO plays $\boldsymbol{x}_t = \Pi_{\mathcal{K}}(\tilde{\boldsymbol{x}}_t)$ on Line 4; the unprojected variant plays $\boldsymbol{x}_t = \tilde{\boldsymbol{x}}_t$.

**Algorithm 4** Coordinate-wise Krichevsky–Trofimov bettor with lazy projection.

1: **Input:** miscoverage levels $\alpha_1 < \cdots < \alpha_d$
2: Initialize $W_0^{(j)} \leftarrow 1$ and cumulative gradient $\Sigma^{(j)} \leftarrow 0$ for each $j \in [d]$
3: **for** $t = 1, 2, \ldots$ **do**
4: $\quad$ **for** $j = 1, \ldots, d$ (in parallel) **do**
5: $\quad\quad$ betting fraction $\beta_t^{(j)} \leftarrow -\Sigma^{(j)}/t$
6: $\quad\quad \tilde{x}_{t,j} \leftarrow \beta_t^{(j)} \cdot W_{t-1}^{(j)}$
7: $\quad$ **end for**
8: $\quad \boldsymbol{x}_t \leftarrow \Pi_{\mathcal{K}}(\tilde{\boldsymbol{x}}_t)$ (*unprojected:* $\boldsymbol{x}_t \leftarrow \tilde{\boldsymbol{x}}_t$); observe $S_t$; $\boldsymbol{g}_t \in \partial\mathcal{L}_t(\boldsymbol{x}_t)$ via (2)
9: $\quad$ **for** $j = 1, \ldots, d$ (in parallel) **do**
10: $\quad\quad W_t^{(j)} \leftarrow W_{t-1}^{(j)} \cdot \big(1 - \beta_t^{(j)} g_{t,j}\big)$; $\Sigma^{(j)} \leftarrow \Sigma^{(j)} + g_{t,j}$
11: $\quad$ **end for**
12: **end for**

### E.2.4. MultiQT

**Algorithm.** MultiQT (Ding et al., 2025) is the lazy-projection-wrapped instance of online gradient descent: in our notation it coincides with Algorithm 3 restricted to $p = 2$ (so $q = 2$), composed with the isotonic projection. Unlike $p$-norm OMD with constant step size, MultiQT uses the data-driven adaptive step-size schedule from Ding et al. (2025,

Section 5): at each round $t$,

$$\eta_t \;=\; \max(0.1 \cdot Q_{0.9}(\{S_s : s \in [\max(1, t-50),\, t-1]\}),\; 0.1)\,,$$

i.e. a multiplicative factor of the empirical 90th-percentile of the most recent 50 scores, lower-bounded by 0.1 to prevent collapse early in the stream. This schedule effectively scales the step with the magnitude of recent residuals, eliminating sensitivity to the score scale at the cost of introducing a sliding window and a multiplicative constant; both are inherited verbatim from Ding et al. (2025). Because `MultiQT` is by design a projected algorithm, we report only its projected form and use it as the prior baseline against which we compare every other (base learner, projection) combination.

**Hyperparameters and Tuning.** The window length (50), tail quantile (0.9), multiplicative factor (0.1), and lower bound (0.1) are all taken from the public implementation of Ding et al. (2025) without modification. The dual iterate is initialized at $\boldsymbol{\theta}_1 = \mathbf{0}$, matching $p$`-norm OMD`. As with $p$`-norm OMD`, this adaptive schedule is best understood as a tuned baseline: its constants reflect choices made by Ding et al. (2025) on their own benchmarks rather than on ours.

### E.3. Width Distribution Analysis

The mean and tail-quantile statistics in Section 5 already give a rough picture of the projection's effect on sharpness, but they do not capture how it reshapes the underlying *distribution* of widths. Figure 2 overlays the histograms of per-round predicted widths on the sinusoidal dataset at two representative levels, $\alpha = 0.03$ and $\alpha = 0.25$, for the two base learners `UP-OCP` and $p$`-norm OMD` that have a meaningful projected/unprojected pair.

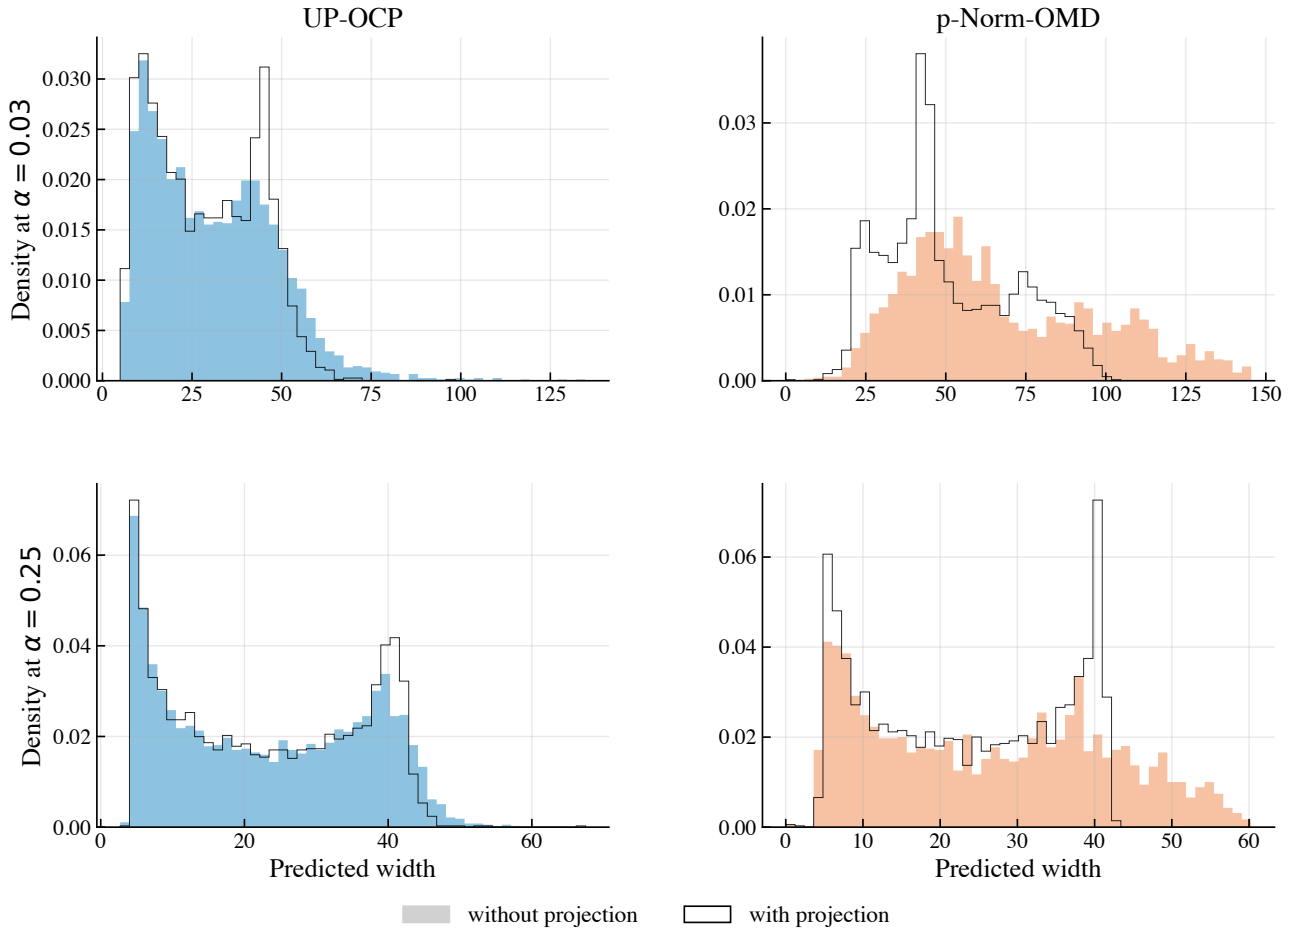


*Figure 2.* Distribution of per-round prediction-set widths on the sinusoidal stream, at $\alpha = 0.03$ (top row) and $\alpha = 0.25$ (bottom row). Each panel overlays the histogram of widths for the unprojected base learner (filled, solid color) against the projected variant (outlined, black border), at the same $\alpha$ and base learner.

The qualitative pattern is consistent across the two base learners. The projection consistently concentrates the distribution

onto a tighter range: the heaviest right-tail mass is removed, and the overall distribution conforms closer to two narrower modes (the lower and higher modes correspond to the extrema pattern of the sinusoidal wave). From a downstream-decision perspective, this reshaping is more useful than a small mean improvement, because the cost of an unusually wide set is non-linear for most users (e.g., capacity reservation, contract pricing).

### E.4. Additional Results

This subsection collects, for every dataset except AXP, the full set of artifacts referenced from Section 5: the calibration curve and three sharpness–coverage Pareto frontiers (using mean, median, and $Q_{75}$ widths) on a single $2\times2$ grid, and a coverage–width summary table at three target levels $1-\alpha \in \{0.95, 0.9, 0.8\}$ including `KT` and its projected variant. The conclusions of Section 5 carry over qualitatively to every dataset; we record only the deviations from the AAPL baseline.

For each tuned baseline ($p$-`norm OMD` with $p$ from Corollary 3.5), the constant step size $\eta$ is selected by per-dataset grid search over $\eta \in \left\{2^{-2}, 2^{-1}, 1, 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024\right\}$, exactly as in Section 5. The per-dataset selected $\eta$ is reported in each subsubsection.

### E.4.1. AAPL (Full Results)

This is the same dataset and forecasting pipeline as in Section 5; we extend Figure 1 with `KT` and `KT`-Iso (omitted there for scale) and add median and $Q_{75}$ width Pareto frontiers (Figure 3). Tuned $\eta = 64$ for $p$-`norm OMD` with $p$ optimal. Per-coverage summary statistics are reported in Table 1.

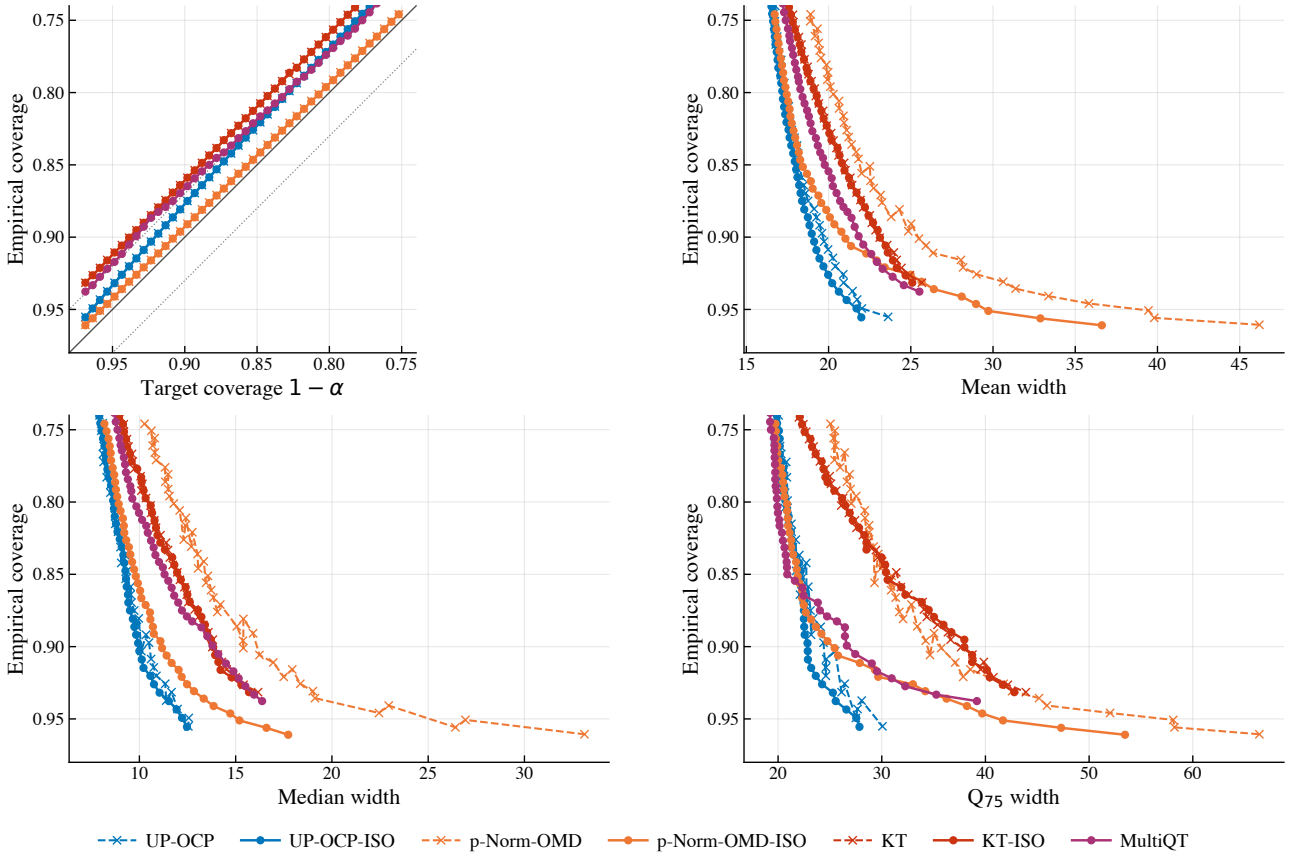


*Figure 3.* AAPL: calibration (top-left) and three width Pareto frontiers (mean, median, $Q_{75}$) versus empirical coverage. `KT` and `KT`-Iso are now included, exposing their wider mean and median widths but competitive $Q_{75}$ tail.

*Table 1.* AAPL: empirical coverage and prediction-set widths at three representative miscoverage levels. **Bold** marks the best value in each column (closest empirical coverage to the target; smallest width).

| Method | $1-\alpha=0.95$ | | | | $1-\alpha=0.9$ | | | | $1-\alpha=0.8$ | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | Cov. | Mean width | Median width | $Q_{95}$ width | Cov. | Mean width | Median width | $Q_{95}$ width | Cov. | Mean width | Median width | $Q_{95}$ width |
| UP-OCP | 0.93 | 20.92 | 11.67 | 73.35 | 0.88 | 18.75 | 9.73 | 69.41 | 0.77 | 16.98 | **8.14** | 66.16 |
| UP-OCP-Iso | 0.93 | **20.22** | **11.05** | **70.53** | 0.88 | **18.39** | **9.51** | **67.17** | 0.77 | **16.85** | 8.26 | 65.07 |
| $p$-Norm-OMD | 0.94 | 33.37 | 22.95 | 99.63 | 0.89 | 25.02 | 15.90 | 81.06 | 0.80 | 20.11 | 11.56 | 69.90 |
| $p$-Norm-OMD-Iso | **0.94** | 28.10 | 13.85 | 101.58 | **0.89** | 20.35 | 10.75 | 71.39 | **0.80** | 17.38 | 8.85 | **64.99** |
| KT | 0.91 | 23.97 | 14.28 | 75.22 | 0.86 | 21.33 | 12.34 | 70.26 | 0.76 | 18.20 | 9.44 | 66.24 |
| KT-Iso | 0.91 | 23.60 | 14.10 | 72.86 | 0.86 | 21.24 | 12.40 | 70.02 | 0.76 | 18.11 | 9.36 | 65.63 |
| MultiQT | 0.92 | 22.94 | 14.88 | 71.40 | 0.86 | 20.26 | 11.81 | 67.97 | 0.77 | 17.85 | 9.26 | 65.64 |

### E.4.2. AMZN

Same Prophet pipeline as AAPL (Taylor & Letham (2018), log-prices, daily refit). Tuned $\eta = 256$ for $p$-`norm OMD` with $p$ optimal.

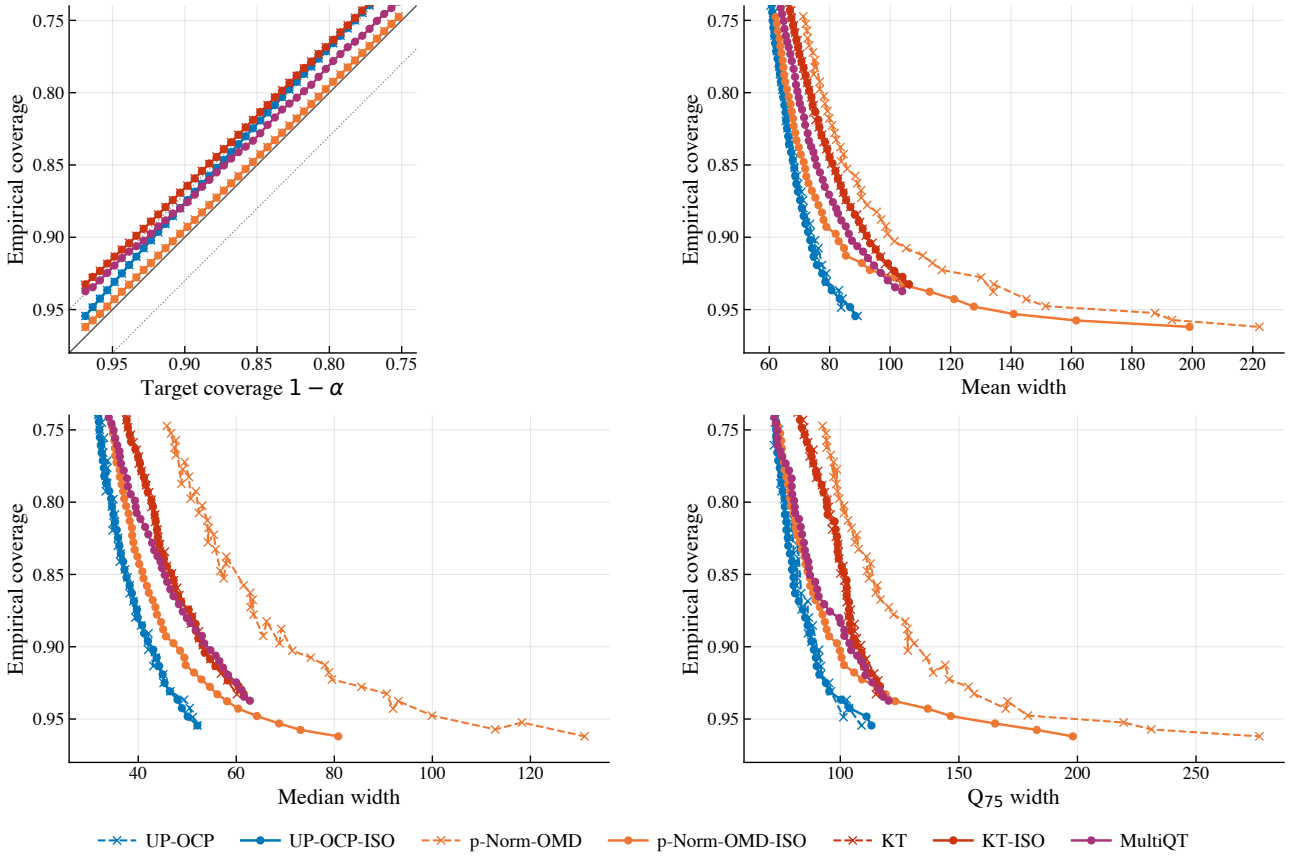


*Figure 4.* AMZN: calibration and three width Pareto frontiers. Behavior is qualitatively identical to AAPL.

*Table 2.* AMZN: empirical coverage and prediction-set widths at three representative miscoverage levels. **Bold** marks the best value in each column (closest empirical coverage to the target; smallest width).

| Method | $1-\alpha = 0.95$ | | | | $1-\alpha = 0.9$ | | | | $1-\alpha = 0.8$ | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | Cov. | Mean width | Median width | $Q_{95}$ width | Cov. | Mean width | Median width | $Q_{95}$ width | Cov. | Mean width | Median width | $Q_{95}$ width |
| UP-OCP | 0.93 | 79.01 | 46.54 | 285.27 | 0.87 | 71.19 | **39.58** | 267.82 | 0.77 | 62.54 | 33.70 | **240.78** |
| UP-OCP-ISO | 0.93 | **78.57** | **46.45** | **277.33** | 0.87 | **70.21** | 39.73 | **261.89** | 0.77 | **62.25** | **32.77** | 241.45 |
| $p$-Norm-OMD | **0.94** | 145.04 | 91.97 | 514.42 | 0.89 | 98.53 | 65.50 | 324.27 | **0.80** | 76.74 | 50.70 | 266.98 |
| $p$-Norm-OMD-ISO | 0.94 | 121.20 | 60.39 | 475.34 | **0.89** | 78.89 | 45.65 | 287.50 | 0.80 | 65.85 | 37.14 | 253.93 |
| KT | 0.91 | 97.00 | 55.96 | 328.28 | 0.86 | 83.57 | 48.39 | 287.34 | 0.77 | 69.39 | 40.16 | 248.53 |
| KT-ISO | 0.91 | 96.60 | 55.74 | 325.48 | 0.86 | 83.66 | 48.10 | 288.93 | 0.77 | 69.58 | 39.97 | 246.44 |
| MULTIQT | 0.92 | 94.68 | 58.32 | 326.20 | 0.88 | 81.05 | 49.08 | 298.66 | 0.78 | 68.05 | 37.70 | 263.80 |

### E.4.3. GOOGL

Same Prophet pipeline as AAPL. Tuned $\eta = 256$ for $p$-`norm OMD` with $p$ optimal. The qualitative pattern matches AAPL/AMZN; the projection consistently improves both tail and mean widths for every base learner without affecting calibration.

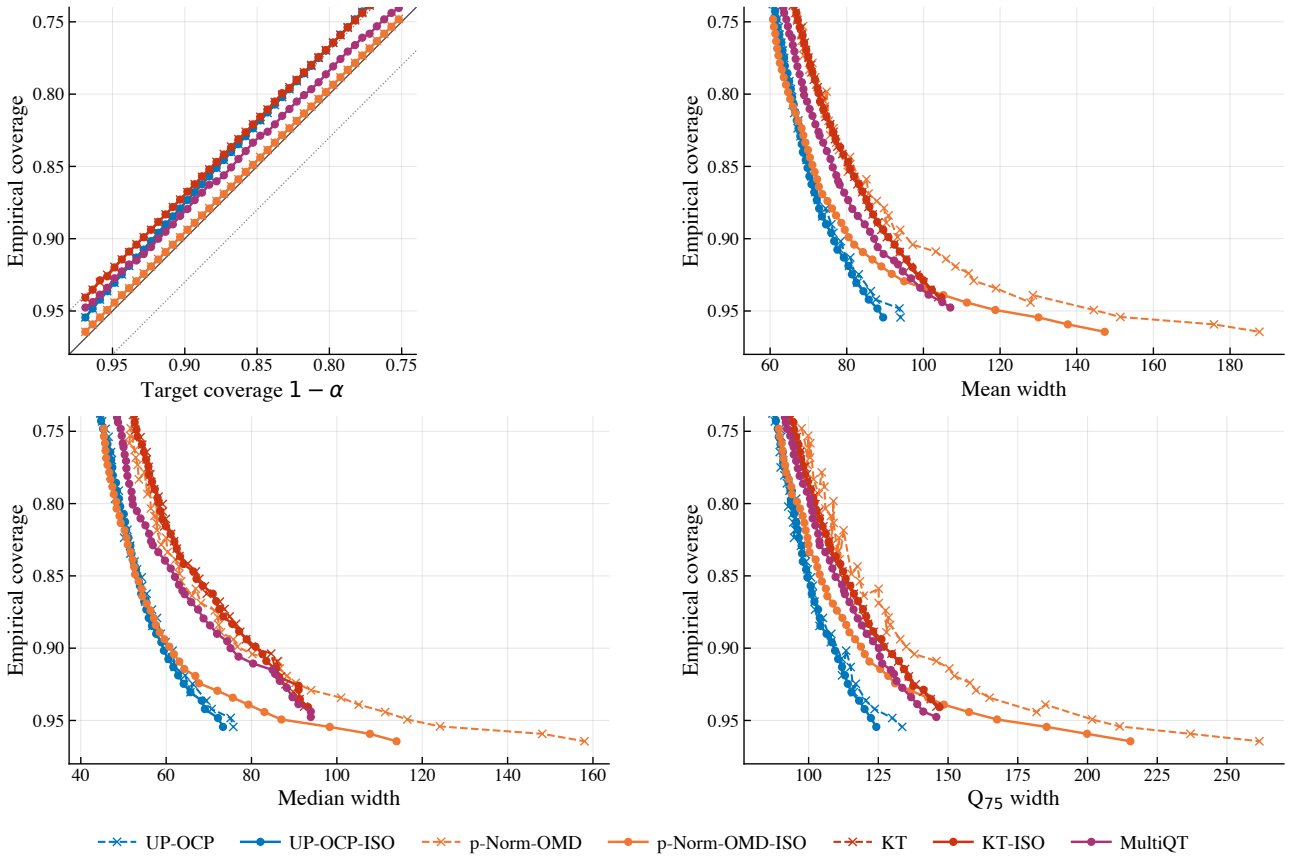


*Figure 5.* GOOGL: calibration and three width Pareto frontiers.

*Table 3.* GOOGL: empirical coverage and prediction-set widths at three representative miscoverage levels. **Bold** marks the best value in each column (closest empirical coverage to the target; smallest width).

| Method | $1-\alpha=0.95$ | | | | $1-\alpha=0.9$ | | | | $1-\alpha=0.8$ | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | Cov. | Mean width | Median width | $Q_{95}$ width | Cov. | Mean width | Median width | $Q_{95}$ width | Cov. | Mean width | Median width | $Q_{95}$ width |
| UP-OCP | 0.93 | 82.74 | **65.69** | 191.80 | 0.87 | 72.47 | 56.68 | 176.64 | 0.77 | 63.29 | 47.28 | 160.16 |
| UP-OCP-ISO | 0.93 | **82.49** | 65.69 | **187.07** | 0.87 | **72.17** | **55.36** | 173.32 | 0.77 | **63.24** | **47.12** | 159.41 |
| $p$-Norm-OMD | **0.94** | 127.93 | 111.31 | 286.65 | **0.89** | 93.94 | 75.78 | 228.67 | 0.80 | 74.74 | 57.80 | 191.89 |
| $p$-Norm-OMD-ISO | 0.94 | 111.35 | 83.02 | 279.91 | 0.89 | 79.47 | 59.60 | 195.80 | **0.80** | 64.69 | 48.29 | 163.53 |
| KT | 0.92 | 96.45 | 86.90 | 187.69 | 0.87 | 84.03 | 72.58 | 173.19 | 0.77 | 69.33 | 55.41 | 153.74 |
| KT-ISO | 0.92 | 97.17 | 86.45 | 192.08 | 0.87 | 84.06 | 71.74 | 173.09 | 0.77 | 69.43 | 55.29 | 154.01 |
| MULTIQT | 0.93 | 96.73 | 88.28 | 189.76 | 0.88 | 81.45 | 68.76 | **172.88** | 0.79 | 67.69 | 51.35 | **153.22** |

### E.4.4. Electricity Demand (elec2)

Half-hourly New South Wales electricity demand from (Harries, 1999), with one-step-ahead forecasts produced by an autoregressive model (lag 3, refit at every step). Scores are absolute residuals on the demand variable, normalized to lie roughly in $[0, 1]$ — hence the smaller numerical scale relative to the stock datasets. Tuned $\eta = 0.25$ for $p$-norm OMD with $p$ optimal.

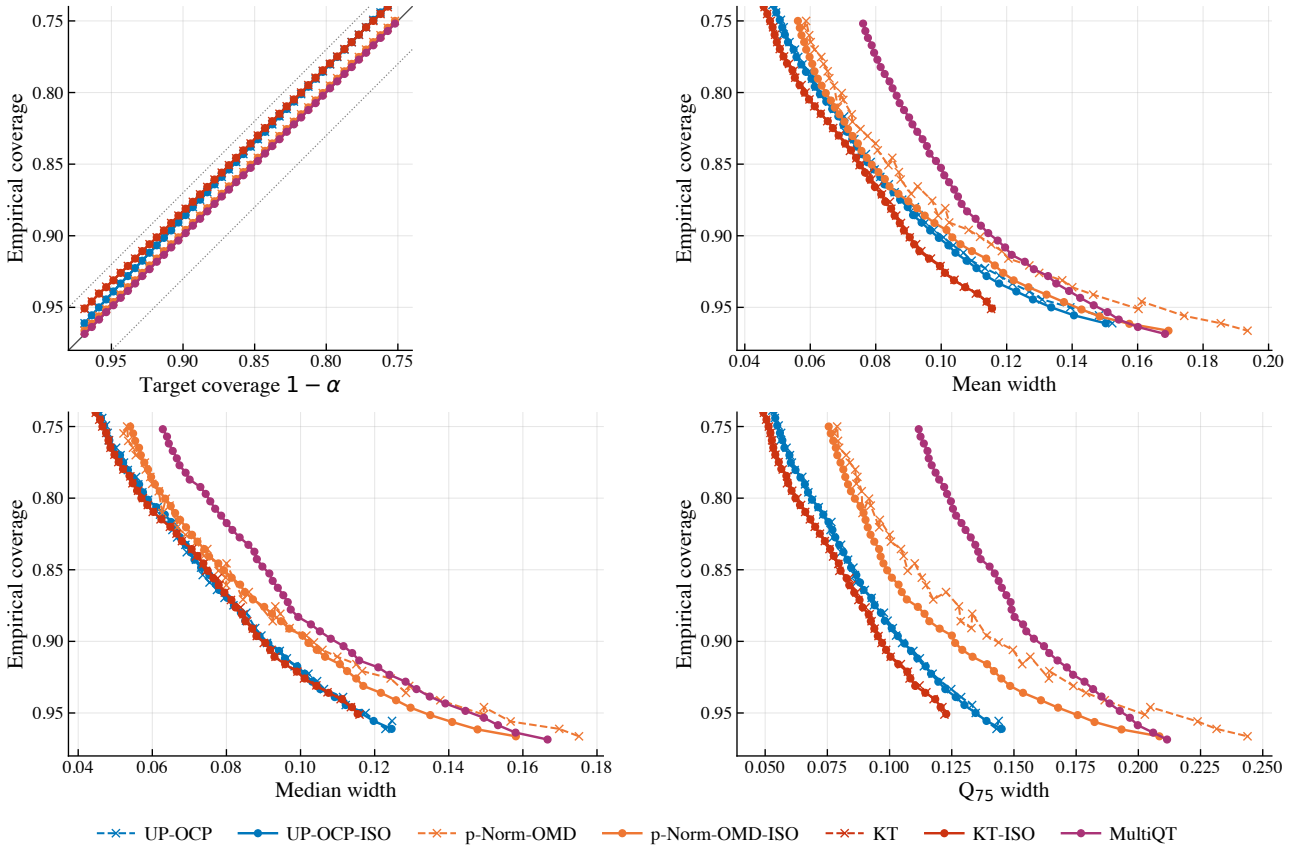


*Figure 6.* elec2: calibration and three width Pareto frontiers.

*Table 4.* elec2: empirical coverage and prediction-set widths at three representative miscoverage levels. **Bold** marks the best value in each column (closest empirical coverage to the target; smallest width).

| Method | $1-\alpha=0.95$ | | | | $1-\alpha=0.9$ | | | | $1-\alpha=0.8$ | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | Cov. | Mean width | Median width | $Q_{95}$ width | Cov. | Mean width | Median width | $Q_{95}$ width | Cov. | Mean width | Median width | $Q_{95}$ width |
| UP-OCP | 0.94 | 0.13 | 0.11 | 0.16 | 0.89 | 0.09 | 0.09 | 0.12 | 0.79 | 0.06 | 0.06 | 0.08 |
| UP-OCP-Iso | 0.94 | 0.12 | 0.11 | 0.15 | 0.89 | 0.09 | 0.09 | 0.12 | 0.79 | 0.06 | 0.05 | 0.08 |
| $p$-Norm-OMD | 0.95 | 0.16 | 0.15 | 0.29 | 0.90 | 0.11 | 0.10 | 0.20 | 0.80 | 0.07 | 0.06 | 0.14 |
| $p$-Norm-OMD-Iso | 0.95 | 0.14 | 0.13 | 0.23 | 0.90 | 0.10 | 0.10 | 0.16 | 0.80 | 0.06 | 0.06 | 0.11 |
| KT | 0.93 | 0.10 | **0.10** | 0.12 | 0.88 | 0.08 | 0.08 | 0.10 | 0.78 | 0.05 | 0.05 | 0.07 |
| KT-Iso | 0.93 | **0.10** | 0.10 | **0.12** | 0.88 | **0.08** | **0.08** | **0.10** | 0.78 | **0.05** | **0.05** | **0.07** |
| MultiQT | **0.95** | 0.15 | 0.14 | 0.25 | **0.90** | 0.11 | 0.11 | 0.23 | **0.80** | 0.09 | 0.08 | 0.20 |

On this stream the relative ordering shifts slightly: KT-Iso attains the smallest mean, median, and $Q_{95}$ widths in every column at $1-\alpha \in \{0.9, 0.8\}$, while UP-OCP-Iso stays within a few percent. The reason is the small score range and the absence of fat tails on this stream: the variance budget $\alpha_j(1-\alpha_j)$ that UP-OCP exploits is dominated here by the constant per-step quantization, which is precisely the regime where the simpler KT update rule is competitive with the universal-portfolio update.

### E.4.5. Sinusoidal (Synthetic)

We generate the nonconformity scores $S_t$ as a sinusoidal wave distorted by Gaussian noise. With $T = 3{,}000$, $\alpha_{\min} = 2$, $\alpha_{\max} = 20$, period $P = 200$, noise scale $\sigma = 0.3$, and $\varepsilon_t \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2)$, define the centre and amplitude $c := (\alpha_{\min} + \alpha_{\max})/2$ and $a := (\alpha_{\max} - \alpha_{\min})/2$ and set

$$S_t \;=\; \max\big(0,\; c + a\sin(2\pi t/P) + \varepsilon_t\big), \qquad t = 1, \ldots, T.$$

Concretely $c = 11$, $a = 9$. Reported quantities are means over 10 independent seeds (base seed 42, seeds $42, 43, \ldots, 51$) with standard errors. Tuned $\eta = 64$ for $p$-`norm OMD` with $p$ optimal. This is the same stream as in Section 5; we extend Figure 1 with `KT` and `KT`-Iso and the median and $Q_{75}$ frontiers (Figure 7), and Table 5 reports the per-method coverage and width summary at three target levels.

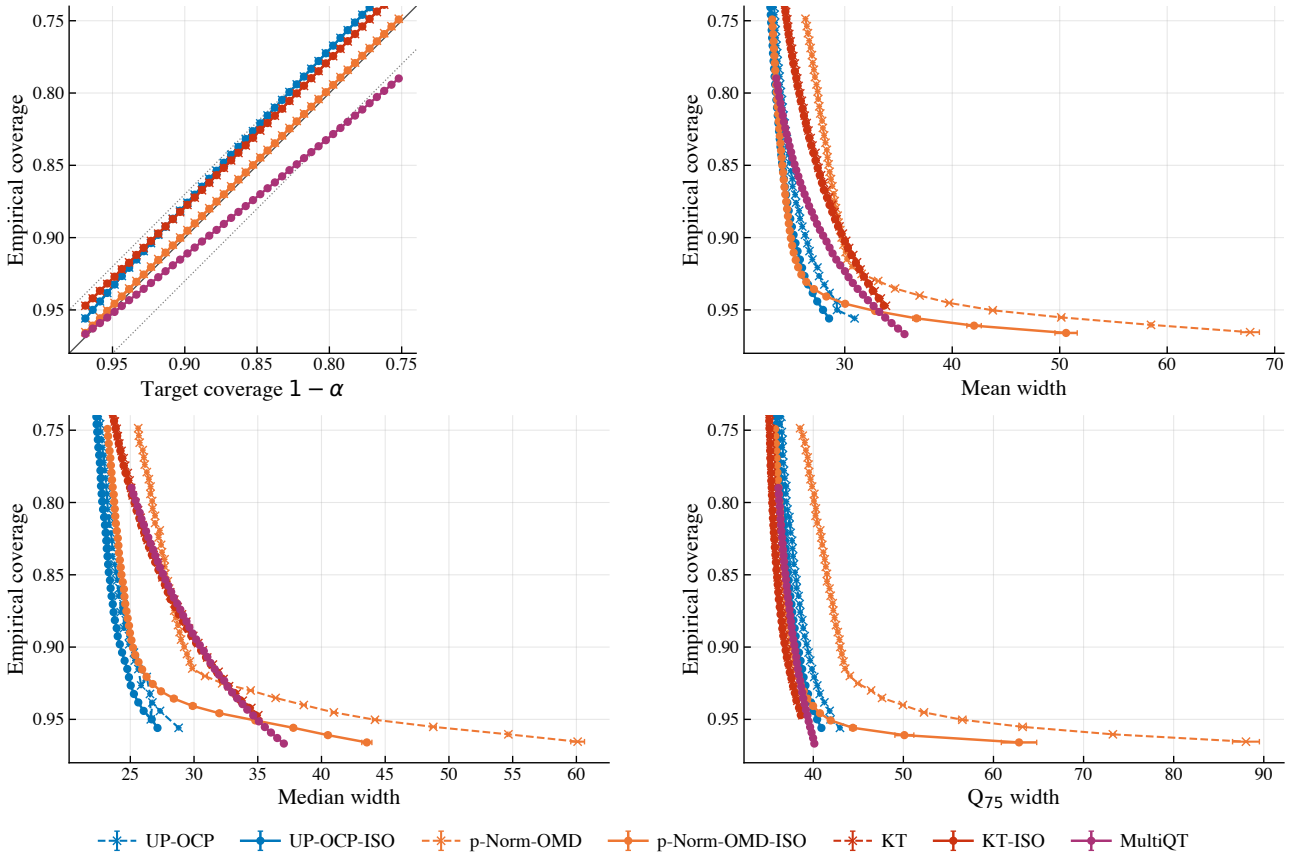


*Figure 7.* Sinusoidal: calibration and three width Pareto frontiers. `UP-OCP`-Iso dominates in mean and median width across the entire $\alpha$ range, while `KT` and `MultiQT` lead on $Q_{75}$, mirroring the bounded twin-extrema structure of the wave.

*Table 5.* sinusoidal: empirical coverage and prediction-set widths at three representative miscoverage levels. **Bold** marks the best value in each column (closest empirical coverage to the target; smallest width).

| Method | $1-\alpha = 0.95$ | | | | $1-\alpha = 0.9$ | | | | $1-\alpha = 0.8$ | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | Cov. | Mean width | Median width | $Q_{95}$ width | Cov. | Mean width | Median width | $Q_{95}$ width | Cov. | Mean width | Median width | $Q_{95}$ width |
| UP-OCP | $0.93_{\pm.00}$ | $28.10_{\pm.02}$ | $26.54_{\pm.01}$ | $54.42_{\pm.03}$ | $0.88_{\pm.00}$ | $25.47_{\pm.02}$ | $24.24_{\pm.03}$ | $49.18_{\pm.06}$ | $0.77_{\pm.00}$ | $23.72_{\pm.01}$ | $22.91_{\pm.03}$ | $45.26_{\pm.04}$ |
| UP-OCP-Iso | $0.93_{\pm.00}$ | $\mathbf{26.58}_{\pm.05}$ | $\mathbf{25.28}_{\pm.09}$ | $50.12_{\pm.19}$ | $0.88_{\pm.00}$ | $\mathbf{24.66}_{\pm.02}$ | $\mathbf{23.69}_{\pm.05}$ | $46.42_{\pm.08}$ | $0.77_{\pm.00}$ | $\mathbf{23.32}_{\pm.01}$ | $\mathbf{22.61}_{\pm.03}$ | $43.84_{\pm.07}$ |
| $p$-Norm-OMD | $0.95_{\pm.00}$ | $39.72_{\pm.20}$ | $40.97_{\pm.19}$ | $67.35_{\pm.35}$ | $0.89_{\pm.00}$ | $29.59_{\pm.02}$ | $28.98_{\pm.05}$ | $58.67_{\pm.08}$ | $0.80_{\pm.00}$ | $27.48_{\pm.02}$ | $26.72_{\pm.03}$ | $55.22_{\pm.05}$ |
| $p$-Norm-OMD-Iso | $0.95_{\pm.00}$ | $30.04_{\pm.14}$ | $31.99_{\pm.23}$ | $49.55_{\pm.43}$ | $\mathbf{0.90}_{\pm.00}$ | $24.89_{\pm.03}$ | $25.09_{\pm.04}$ | $42.73_{\pm.08}$ | $\mathbf{0.80}_{\pm.00}$ | $23.67_{\pm.02}$ | $23.73_{\pm.03}$ | $41.69_{\pm.05}$ |
| KT | $0.93_{\pm.00}$ | $31.86_{\pm.04}$ | $32.75_{\pm.04}$ | $41.38_{\pm.04}$ | $0.88_{\pm.00}$ | $28.82_{\pm.02}$ | $29.05_{\pm.03}$ | $41.09_{\pm.02}$ | $0.78_{\pm.00}$ | $25.35_{\pm.01}$ | $24.80_{\pm.03}$ | $40.72_{\pm.03}$ |
| KT-Iso | $0.93_{\pm.00}$ | $31.78_{\pm.05}$ | $32.68_{\pm.05}$ | $\mathbf{41.24}_{\pm.03}$ | $0.88_{\pm.00}$ | $28.58_{\pm.02}$ | $28.81_{\pm.03}$ | $40.94_{\pm.02}$ | $0.78_{\pm.00}$ | $25.21_{\pm.02}$ | $24.62_{\pm.03}$ | $40.60_{\pm.02}$ |
| MultiQT | $\mathbf{0.95}_{\pm.00}$ | $33.21_{\pm.04}$ | $35.09_{\pm.05}$ | $41.34_{\pm.05}$ | $0.91_{\pm.00}$ | $28.87_{\pm.04}$ | $31.29_{\pm.05}$ | $\mathbf{40.89}_{\pm.04}$ | $0.83_{\pm.00}$ | $24.73_{\pm.02}$ | $26.79_{\pm.04}$ | $\mathbf{40.52}_{\pm.03}$ |

### E.4.6. Stationary Trend with Random Waves (Synthetic)

Nonconformity scores follow a constant baseline distorted by sparse, heavy-tailed multiplicative spikes and smoothed by a centred rolling maximum. With baseline $C = 10$, spike probability $p = 0.1$, exponential scale $\sigma = 10$, rolling window $W = 25$, and $T = 3{,}000$, we draw

$$B_t \sim \text{Bernoulli}(p), \qquad E_t \sim \text{Exponential}(1/\sigma), \qquad \tilde{S}_t \;=\; C\left(1 + B_t E_t\right),$$

and obtain $S_t$ by a centered rolling-maximum smoother of half-width $\lfloor W/2 \rfloor$:

$$S_t \;=\; \max_{\tau \in [t-\lfloor W/2 \rfloor,\, t+\lfloor W/2 \rfloor]} \tilde{S}_\tau .$$

This stresses heavy tails without an underlying trend. We report means over 10 seeds (base seed 42). Tuned $\eta = 1024$ for $p$-`norm OMD` with $p$ optimal.

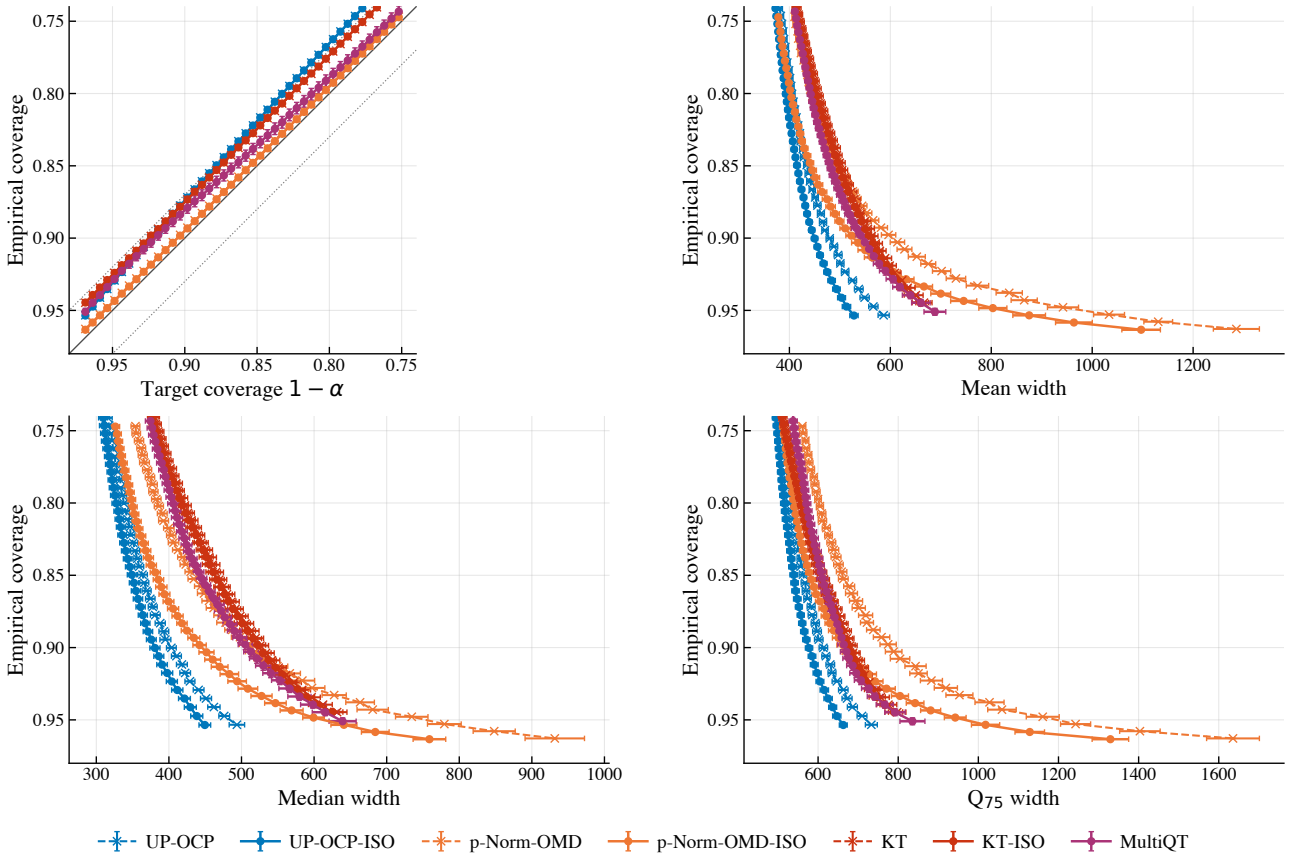


*Figure 8.* Stationary trend with random waves: calibration and three width Pareto frontiers.

*Table 6.* stationary: empirical coverage and prediction-set widths at three representative miscoverage levels. **Bold** marks the best value in each column (closest empirical coverage to the target; smallest width).

| Method | $1-\alpha = 0.95$ | | | | $1-\alpha = 0.9$ | | | | $1-\alpha = 0.8$ | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | Cov. | Mean width | Median width | $Q_{95}$ width | Cov. | Mean width | Median width | $Q_{95}$ width | Cov. | Mean width | Median width | $Q_{95}$ width |
| UP-OCP | $0.93_{\pm.00}$ | $525.10_{\pm7.99}$ | $440.92_{\pm7.39}$ | $1172_{\pm29.01}$ | $0.87_{\pm.00}$ | $451.12_{\pm6.20}$ | $377.72_{\pm6.14}$ | $1026_{\pm27.17}$ | $0.77_{\pm.00}$ | $391.88_{\pm6.07}$ | $325.61_{\pm6.29}$ | $912.26_{\pm23.90}$ |
| UP-OCP-Iso | $0.93_{\pm.00}$ | $\mathbf{483.63}_{\pm\mathbf{7.43}}$ | $\mathbf{411.50}_{\pm\mathbf{7.77}}$ | $\mathbf{1059}_{\pm\mathbf{26.10}}$ | $0.87_{\pm.00}$ | $\mathbf{427.91}_{\pm\mathbf{6.05}}$ | $\mathbf{360.99}_{\pm\mathbf{6.36}}$ | $967.28_{\pm24.87}$ | $0.77_{\pm.00}$ | $\mathbf{381.22}_{\pm\mathbf{5.34}}$ | $\mathbf{317.14}_{\pm\mathbf{5.90}}$ | $885.35_{\pm22.26}$ |
| $p$-Norm-OMD | $0.94_{\pm.00}$ | $865.25_{\pm25.90}$ | $680.72_{\pm21.55}$ | $2301_{\pm102.11}$ | $0.89_{\pm.00}$ | $581.46_{\pm12.87}$ | $497.91_{\pm11.83}$ | $1347_{\pm50.22}$ | $0.80_{\pm.00}$ | $441.37_{\pm5.32}$ | $382.66_{\pm6.51}$ | $1005_{\pm20.14}$ |
| $p$-Norm-OMD-Iso | $\mathbf{0.94}_{\pm\mathbf{.00}}$ | $745.34_{\pm25.86}$ | $569.02_{\pm15.81}$ | $2088_{\pm130.20}$ | $\mathbf{0.89}_{\pm\mathbf{.00}}$ | $510.48_{\pm10.38}$ | $434.35_{\pm8.26}$ | $1191_{\pm48.89}$ | $\mathbf{0.80}_{\pm\mathbf{.00}}$ | $401.20_{\pm5.85}$ | $348.20_{\pm5.97}$ | $915.13_{\pm25.25}$ |
| KT | $0.92_{\pm.00}$ | $608.70_{\pm12.49}$ | $571.92_{\pm10.53}$ | $1079_{\pm41.79}$ | $0.87_{\pm.00}$ | $527.99_{\pm12.18}$ | $491.64_{\pm10.79}$ | $964.65_{\pm37.28}$ | $0.78_{\pm.00}$ | $439.18_{\pm8.64}$ | $402.11_{\pm8.33}$ | $825.50_{\pm23.61}$ |
| KT-Iso | $0.92_{\pm.00}$ | $601.46_{\pm11.81}$ | $565.92_{\pm10.37}$ | $1067_{\pm38.87}$ | $0.87_{\pm.00}$ | $523.56_{\pm11.66}$ | $488.15_{\pm10.63}$ | $\mathbf{952.76}_{\pm\mathbf{34.55}}$ | $0.78_{\pm.00}$ | $435.85_{\pm8.20}$ | $398.80_{\pm8.15}$ | $\mathbf{817.11}_{\pm\mathbf{21.49}}$ |
| MultiQT | $0.93_{\pm.00}$ | $605.17_{\pm13.54}$ | $566.08_{\pm14.59}$ | $1102_{\pm46.29}$ | $0.88_{\pm.00}$ | $512.82_{\pm9.02}$ | $476.98_{\pm9.11}$ | $995.40_{\pm33.23}$ | $0.79_{\pm.00}$ | $436.87_{\pm6.50}$ | $400.76_{\pm7.38}$ | $902.39_{\pm23.39}$ |

The width scale on this stream is large ($\sim 10^2$), reflecting the rolling-maximum smoothing of exponential spikes. The projection still strictly improves the unprojected base learner on every cell of the table, and the parameter-free `UP-OCP`-Iso achieves the smallest mean, median, and $Q_{95}$ widths at every target.

### E.4.7. Quadratic Drift with Random Waves (mixed, Synthetic)

Same multiplicative-spike-plus-rolling-max structure as the stationary stream, now overlaid on a quadratic drift. With $T = 3{,}000$, baseline trajectory $T_t = (20/T^2)\, t^2$ on $t \in [0, T-1]$, spike probability $p = 0.1$, exponential scale $\sigma = 10$, and rolling window $W = 25$, we draw

$$B_t \sim \text{Bernoulli}(p), \qquad E_t \sim \text{Exponential}(1/\sigma), \qquad \tilde{S}_t \;=\; T_t\,\big(1 + B_t E_t\big),$$

and again take $S_t = \max_{\tau \in [t-\lfloor W/2 \rfloor, t+\lfloor W/2 \rfloor]} \tilde{S}_\tau$. The deterministic trend $T_t$ has quadratic shape in $t$ and ranges over $[0, 20]$ across the horizon. We report means over 10 seeds (base seed 42). Tuned $\eta = 1024$ for $p$-norm OMD with $p$ optimal.

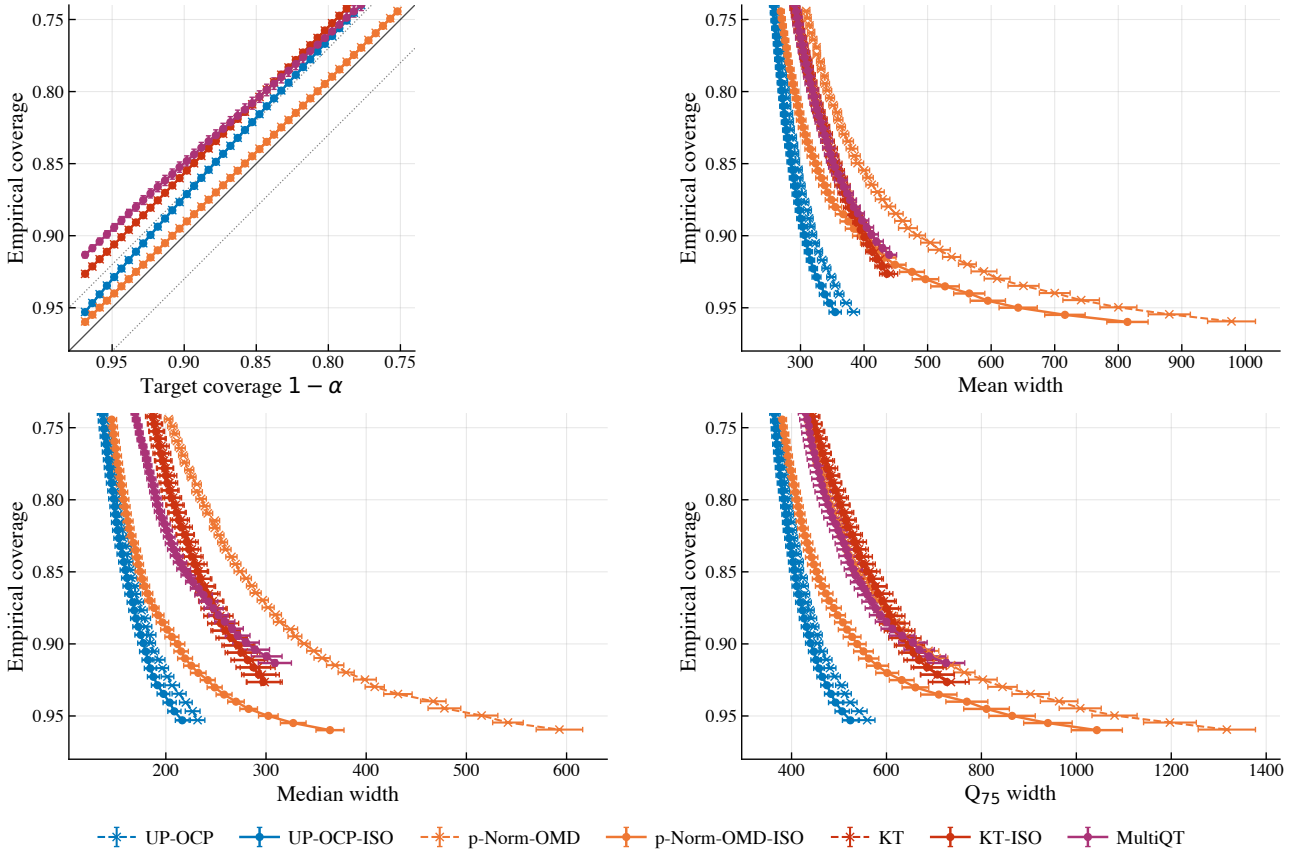


*Figure 9.* Quadratic drift with random waves: calibration and three width Pareto frontiers.

*Table 7.* mixed: empirical coverage and prediction-set widths at three representative miscoverage levels. **Bold** marks the best value in each column (closest empirical coverage to the target; smallest width).

| Method | $1-\alpha = 0.95$ | | | | $1-\alpha = 0.9$ | | | | $1-\alpha = 0.8$ | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | Cov. | Mean width | Median width | $Q_{95}$ width | Cov. | Mean width | Median width | $Q_{95}$ width | Cov. | Mean width | Median width | $Q_{95}$ width |
| UP-OCP | $0.93_{\pm.00}$ | $348.44_{\pm7.14}$ | $206.36_{\pm7.19}$ | $1158_{\pm40.08}$ | $0.87_{\pm.00}$ | $303.10_{\pm5.77}$ | $172.50_{\pm5.22}$ | $1037_{\pm34.24}$ | $0.77_{\pm.00}$ | $267.23_{\pm5.44}$ | $143.56_{\pm4.55}$ | $943.16_{\pm35.66}$ |
| UP-OCP-Iso | $0.93_{\pm.00}$ | $\mathbf{325.17}_{\pm\mathbf{7.00}}$ | $\mathbf{191.73}_{\pm\mathbf{6.11}}$ | $\mathbf{1079}_{\pm\mathbf{37.51}}$ | $0.87_{\pm.00}$ | $\mathbf{293.20}_{\pm\mathbf{6.40}}$ | $\mathbf{166.90}_{\pm\mathbf{5.51}}$ | $\mathbf{1002}_{\pm\mathbf{36.87}}$ | $0.77_{\pm.00}$ | $\mathbf{262.85}_{\pm\mathbf{5.55}}$ | $\mathbf{141.34}_{\pm\mathbf{4.66}}$ | $930.76_{\pm35.18}$ |
| $p$-Norm-OMD | $0.94_{\pm.00}$ | $700.07_{\pm24.17}$ | $466.97_{\pm12.68}$ | $2249_{\pm119.77}$ | $0.89_{\pm.00}$ | $460.82_{\pm12.52}$ | $324.56_{\pm6.93}$ | $1443_{\pm65.93}$ | $0.79_{\pm.00}$ | $338.58_{\pm6.98}$ | $231.17_{\pm5.27}$ | $1070_{\pm33.82}$ |
| $p$-Norm-OMD-Iso | $\mathbf{0.94}_{\pm\mathbf{.00}}$ | $565.52_{\pm24.49}$ | $270.02_{\pm8.23}$ | $2122_{\pm129.98}$ | $\mathbf{0.89}_{\pm\mathbf{.00}}$ | $375.04_{\pm12.77}$ | $201.36_{\pm6.52}$ | $1294_{\pm64.43}$ | $\mathbf{0.79}_{\pm\mathbf{.00}}$ | $288.22_{\pm6.59}$ | $156.79_{\pm5.13}$ | $984.98_{\pm35.75}$ |
| KT | $0.91_{\pm.00}$ | $410.94_{\pm11.68}$ | $277.76_{\pm16.07}$ | $1283_{\pm62.65}$ | $0.85_{\pm.00}$ | $358.83_{\pm10.14}$ | $235.56_{\pm10.34}$ | $1096_{\pm43.06}$ | $0.76_{\pm.00}$ | $299.00_{\pm7.45}$ | $192.79_{\pm7.62}$ | $926.47_{\pm33.12}$ |
| KT-Iso | $0.91_{\pm.00}$ | $407.09_{\pm10.74}$ | $275.23_{\pm16.03}$ | $1267_{\pm53.10}$ | $0.85_{\pm.00}$ | $356.09_{\pm9.65}$ | $234.22_{\pm10.27}$ | $1086_{\pm41.75}$ | $0.76_{\pm.00}$ | $297.40_{\pm7.65}$ | $191.74_{\pm7.54}$ | $\mathbf{921.67}_{\pm\mathbf{33.54}}$ |
| MultiQT | $0.89_{\pm.00}$ | $404.10_{\pm10.88}$ | $271.70_{\pm12.32}$ | $1234_{\pm43.23}$ | $0.85_{\pm.00}$ | $350.65_{\pm9.81}$ | $217.81_{\pm6.87}$ | $1111_{\pm31.78}$ | $0.76_{\pm.00}$ | $300.23_{\pm7.98}$ | $177.17_{\pm3.82}$ | $996.18_{\pm34.54}$ |