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

Tuo Liu
King Abdullah University of Science and Technology
tuo.liu@kaust.edu.sa

Edgar Dobriban
University of Pennsylvania
dobriban@wharton.upenn.edu

Francesco Orabona
King Abdullah University of Science and Technology
francesco.orabona@kaust.edu.sa

Abstract

Online Conformal Prediction (OCP) sequentially constructs prediction sets whose average empirical coverage converges to a target level over time, for arbitrary, potentially adversarial datasets. In the multi-level setting, prediction sets are produced at d coverage levels at once. Prior work argued for the use of online gradient descent with isotonic projections, which produces sets nested across levels, but with a $\sqrt{d}$ dependence in the per-level coverage convergence rate. Here, we both generalize and improve prior work. First, we show that isotonic projections can be used for the construction of multi-level online conformal prediction sets from *any* online learning algorithm (not just gradient descent). Then, through a black-box reduction, we show that the linearized regret of the online algorithm itself automatically gives control over both per-level coverage and distributional consistency of multi-level predictions via Fenchel duality. Finally, by using the Weighted Interval Score, a standard scoring rule for assessing multi-level interval forecasts, we prove that these projections reduce both the size of the sets and their under-coverage. We study two instantiations of our theory, p-norm Online Mirror Descent and coordinate-wise Universal Portfolio methods, both improving the dependence on the number of levels to $\sqrt{\log d}$ for the multi-level OCP problem. Experiments on several datasets show favorable performance compared to baselines.

1 Introduction

Constructing prediction sets with reliable coverage guarantees is a central goal of uncertainty quantification. Conformal prediction provides a distribution-free framework for this task: it constructs a prediction set that contains the realized outcome with a user-specified probability, while making few assumptions about the underlying predictive model [28, 34, 35, 39, 40]. In the classical exchangeable setting, conformal methods have been developed for regression, classification, and structured prediction [see, e.g., 4, 23, 24, 33]. This assumption, however, can be violated in settings where the data distribution shifts, such as sequential decision-making and time-series forecasting.

Online conformal prediction (OCP) addresses these limitations by operating in a *fully adversarial sequential setting*, where the data may be generated by an arbitrary and possibly adversarial process. A growing body of work has developed algorithms for single-level OCP, in which the goal is to maintain long-run coverage at a single target level $1 - \alpha$ for any data sequence. These include Adaptive Conformal Inference (`ACI`) [16], an FTRL-based approach [43], conformal PID control [1], meta-algorithms [5, 7, 17, 42], and betting-based methods [25, 29]. Recently, Liu et al. [25] established a

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

black-box reduction showing that *any* online learning algorithm with a linearized regret guarantee can be converted into a single-level OCP algorithm with provable coverage. This reveals a rich landscape of available algorithms for the single-level problem.

However, many real-world applications can benefit from prediction sets that have simultaneous validity at multiple different levels of coverage. This can allow users to inspect the prediction sets *ex post facto* and choose one of them that fits their needs in terms of coverage, length, as well as any potential downstream risk calculations.

Motivated by this, we consider the setting of multi-level online conformal prediction recently proposed by Ding et al. [14]. 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 $\hat{Y}_t$ for the response Y_t, using all the previously observed data points, as well as the current features Z_t. We say coverage is attained at t for the prediction set $\hat{C}_t(\hat{Y}_t, x)$, parameterized by a threshold x, iff it contains Y_t, i.e., $Y_t \in \hat{C}_t(\hat{Y}_t, x)$. We fix d target miscoverage levels $\mathcal{A} = \{\alpha_1, \ldots, \alpha_d\}$, where $0 < \alpha_1 < \alpha_2 < \cdots < \alpha_d < 1$. The conformal predictor announces a vector of thresholds $\boldsymbol{x}_t = (x_{t,1}, \ldots, x_{t,d})$ before Y_t is revealed, where $x_{t,i}$ targets a long-run coverage rate of $1 - \alpha_i$. Two properties are essential for multi-level forecasts:

1. Coverage: for any sequence of data $((Z_t, Y_t))_{t\geq 1}$, even adversarially chosen, the long-run observed coverage at each level should converge to the nominal target as the time horizon T grows:

$$\lim_{T\to\infty} \mathrm{MisCov}_{T,i} = 0,\ \forall i \in [d],\ \text{where}\ \mathrm{MisCov}_{T,i} := \left| \frac{1}{T} \sum_{t=1}^{T} \mathbf{1}\left\{S_t \leq x_{t,i}\right\} - (1 - \alpha_i) \right|.$$

2. Distributional consistency. Forecasts should be ordered across quantile levels, i.e., larger levels should give larger sets: $x_{t,1} \geq x_{t,2} \geq \cdots \geq x_{t,d}$, for all $t = 1, 2, \ldots$.

We will refer to these two properties as *multi-level coverage*.

Ding et al. [14] proposed the `MultiQT` algorithm, which wraps Online Gradient Descent (OGD) with a lazy isotonic projection step to enforce ordering. The algorithm provably achieves asymptotic multi-level coverage. However, their result applies to a single algorithm instance. Meanwhile, as said above, the single-level OCP literature offers a diverse toolbox of algorithms, provably better than OGD. It remains unclear how to migrate these algorithms to the multi-level setting (ideally, with a black-box reduction), while preserving multi-level coverage.

Contributions. In this work, we answer the above open questions with the following contributions:

- In Section 3, we present a general *black-box reduction*: any online learner satisfying a linearized regret bound properly wrapped with an isotonic projection layer yields a multi-level OCP algorithm that provably guarantees multi-level coverage.
- Instantiating our framework with p-norm Online Mirror Descent [15] and coordinate-wise Universal Portfolio algorithms [13] yields per-level miscoverage rates that depend on $\sqrt{\log d}$, improving over the $\sqrt{d}$ dependence of the general guarantee in Ding et al. [14], and, for the latter, with no parameter tuning under polynomially growing scores. The improvement is in the dependence on d (rather than on the horizon T; see Table 1 for the regime where `MultiQT` attains a faster rate in T).
- In Section 4, we show that when the isotonic projection acts, it strictly reduces the *sharpness reward* and weakly reduces the *coverage penalty* of the *Weighted Interval Score* (WIS) [9, 18, 41]. This gives a precise theoretical understanding for the empirical observation that projection improves the Pareto frontier of coverage-sharpness trade-offs.
- In Section 5, we evaluate our algorithms on synthetic and real-world datasets, demonstrating consistent improvements in the Pareto frontiers and in the WIS over existing baselines.

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 [16, 17, 42]. 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 [1, 7]. Parameter-free approaches obtain optimal regret guarantees uniformly over all comparators, by avoiding explicit stepsize tuning [25, 29, 43], and can be extended to group-conditional coverage at a single level [6]. Among these works, the closest one to our work is Liu et al. [25], 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 [3, 10].

The multi-level setting is connected to, but distinct from, simultaneous and multivalid forms of conformal inference, where one controls empirical coverage across groups or threshold families [5], while simultaneous conformal and conformal risk-control methods handle many inference targets or losses at once [2, 22]. 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 [9, 18, 41]. Table 1 in Appendix A.9 compares our results with the two works we build on closely.

2 Problem setup

In this section, we formalize the multi-level OCP problem outlined in Section 1.

Notation. We collect a few standard tools from convex analysis in $\mathbb{R}^d$ [see, e.g., 31] for reference. For a function $f\colon \mathbb{R}^d \to \mathbb{R} \cup \{+\infty\}$, 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 *indicator function* $\mathbf{1}\{A\}$ is equal to 1 if A is true and 0 otherwise. The *Fenchel conjugate* of $f\colon \mathbb{R}^d \to [-\infty, +\infty]$ is $f^\star(\boldsymbol{\theta}) \coloneqq \sup_{\boldsymbol{x} \in \mathbb{R}^d} \langle \boldsymbol{\theta}, \boldsymbol{x} \rangle - f(\boldsymbol{x})$, which is always convex and closed, i.e., its epigraph is a closed set (equivalently, $f^\star$ is lower semicontinuous). The *Euclidean projection* onto a nonempty closed convex set $\mathcal{C} \subseteq \mathbb{R}^d$ is $\Pi_{\mathcal{C}}(\boldsymbol{x}) \coloneqq \arg\min_{\boldsymbol{z} \in \mathcal{C}} \|\boldsymbol{x} - \boldsymbol{z}\|_2$, which is well-defined and unique. We denote by $[d] \coloneqq \{1, \dots, d\}$.

Multi-level OCP. We now describe the sequential protocol of OCP with multiple quantile levels. We choose to focus on symmetric prediction sets of the form $\hat{C}_t(\hat{Y}_t, x) \coloneqq [\hat{Y}_t - x, \hat{Y}_t + x]$, defined as empty if $x < 0$ [24]. Our methodology is applicable beyond this setting: it only needs a nested family of sets parameterized by one scalar per level; e.g., it can be used to adjust the prediction intervals based on conditional quantile regression models [32], but we avoid the details for brevity. We refer the reader to Gupta et al. [19] for more versions of the nested prediction sets for which our methodology is applicable. Let $S_t \geq 0$ denote the radius of the smallest prediction set that contains the true response Y_t, i.e., $S_t \coloneqq \inf\{x \in [0, \infty) : Y_t \in \hat{C}_t(\hat{Y}_t, x)\} = |Y_t - \hat{Y}_t|$. With this choice, $Y_t \in \hat{C}_t(\hat{Y}_t, x) \iff S_t \leq x$ for $x > 0$, and the coverage event at level $1 - \alpha_i$ collapses to the single threshold comparison $\mathbf{1}\{S_t \leq x_{t,i}\}$. We will also refer to S_t as the non-conformity score [39].

Note that a negative radius does not describe a valid interval, so our algorithms output the empty set instead, represented by a 0 radius, the Lebesgue measure of $\emptyset$ (Section 3). Accordingly, a played radius of 0 is read as $\emptyset$, which never covers Y_t, even when $S_t = 0$.

As we said in Section 1, we require per-level *coverage* and *distributional consistency*. For the second one, the increasing miscoverage levels $\alpha_1 < \cdots < \alpha_d$ should be matched by non-increasing thresholds $x_{t,1} \geq \cdots \geq x_{t,d}$. Equivalently, $\boldsymbol{x}_t \in \mathcal{K}$ for every t, where $\mathcal{K} \coloneqq \{\boldsymbol{x} \in \mathbb{R}^d : x_1 \geq x_2 \geq \cdots \geq x_d\}$ is the *isotonic cone*.[1]

Pinball loss and its aggregate. As in the single-level case, we introduce the *pinball (or quantile) loss* for online quantile learning [16, 29], defined as

$$\ell^{(1-\alpha)}(x, S) \coloneqq \max\{(1 - \alpha)(S - x), \alpha(x - S)\}, \tag{1}$$

where S is the non-conformity score and x is the radius of the prediction interval. This loss is convex and L-Lipschitz in the first argument, where $L \coloneqq \max\{1 - \alpha, \alpha\}$. Intuitively, the pinball loss is a proper surrogate for the $(1 - \alpha_i)$-th quantile, as under any distribution of S, the expectation of $\ell^{(1-\alpha_i)}(x, S)$ in x is minimized exactly at the $(1 - \alpha_i)$-th quantile of S [18].

To handle the d levels jointly, we work with the *aggregate pinball loss* $\mathcal{L}_t(\boldsymbol{x}) \coloneqq \sum_{i=1}^d \ell^{(1-\alpha_i)}(x_i, S_t)$, for $\boldsymbol{x} \in \mathbb{R}^d$. We select the subgradient $\boldsymbol{g}_t \in \partial \mathcal{L}_t(\boldsymbol{x}_t)$ with coordinates

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

Note that $\mathcal{L}_t$ is Lipschitz in both the ℓ_2 and ℓ_1 (dual of ℓ_∞) norms: $\|\boldsymbol{g}_t\|_2 \leq \sqrt{d}$ and $\|\boldsymbol{g}_t\|_\infty \leq \max\{1 - \alpha_1, \alpha_d\} \leq 1$. These are favorable conditions for online learnability [see, e.g., 12, 21, 27].

[1] By convention, we choose the non-increasing cone.

Performance on the aggregate loss is measured by the standard notion of *regret*. Against a fixed comparator $\boldsymbol{u} \in \mathbb{R}^d$, the regret of an algorithm producing $(\boldsymbol{x}_t)_{t\geq 1}$ is $\mathrm{Regret}_T(\boldsymbol{u}) \coloneqq \sum_{t=1}^T \mathcal{L}_t(\boldsymbol{x}_t) - \sum_{t=1}^T \mathcal{L}_t(\boldsymbol{u})$.

Finally, we note that the subgradient (2) admits a direct interpretation as a miscoverage error:

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

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

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

This assumption generalizes the bounded-score assumption $S_t \leq D$ used in much of the OCP literature (recovered by $\gamma = 0$), and accommodates non-stationary streams whose magnitudes grow with the horizon (e.g., financial returns with volatility clustering [25]).

3 From regret to multi-level coverage: a black-box reduction

Our goal is to design a black-box reduction for an online algorithm: (1) to have iterates $\boldsymbol{x}_t$ lie in the isotonic cone $\mathcal{K}$ at every round; (2) to secure per-level coverage by controlling the cumulative subgradient $\boldsymbol{G}_T \coloneqq \sum_{t=1}^T \boldsymbol{g}_t$, via (3). Linearized regret is the natural interface for (2): $G_{T,i}$ is T times the signed coverage error at level i, and a bound on $\sum_t \langle \boldsymbol{g}_t, \boldsymbol{x}_t - \boldsymbol{u}\rangle$ holding for every comparator $\boldsymbol{u}$ is, by duality, a bound on $\boldsymbol{G}_T$. We resolve both issues through a single mechanism: a *lazy isotonic projection layer, applied as a black-box wrapper around any unconstrained online learner* Alg *with a linearized-regret guarantee*.

The lazy isotonic projection wrapper. Let Alg be a generic online learner. We make no assumption on its internal structure: Alg is a black box that, at each round, produces a vector of radii $\tilde{\boldsymbol{x}}_t \in \mathbb{R}^d_+$ and accepts an arbitrary subgradient. The wrapper inserts Alg with a single Euclidean projection step:

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

1: **Input:** base learner Alg with outputs in $\mathbb{R}^d_+$; isotonic cone $\mathcal{K}$
2: **for** $t = 1, 2, \ldots$ **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 and 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 of this protocol is the asymmetry between the played iterate $\boldsymbol{x}_t$ (projected onto $\mathcal{K}$) and the gradient input $\boldsymbol{g}_t$ to Alg (evaluated at the projected iterate, then routed back into the unconstrained dynamics). This pattern is known as *lazy projection* [44]. Despite being embedded in `MultiQT` [14], the generality of such a structure was largely unrecognized. Here, we treat the projection as a strictly external layer.

Any online learner on $\mathbb{R}^d$ becomes a base learner with outputs in $\mathbb{R}^d_+$ through the convention of Section 2: whenever its raw radius $z_{t,i}$ is negative, it outputs the empty set at level i, represented as $\tilde{x}_{t,i} = \max(z_{t,i}, 0) = 0$. The learner's update rule and internal state are unchanged. It also makes the input of the isotonic layer a valid vector of radii, which Section 4 exploits.

To start off, observe that, since $\boldsymbol{x}_t \in \mathcal{K}$ at every round, the empirical coverage rates inherit the same monotonicity: $\mathrm{Coverage}_{T,1} \geq \mathrm{Coverage}_{T,2} \geq \cdots \geq \mathrm{Coverage}_{T,d}$, where $\mathrm{Coverage}_{T,i} \coloneqq (1/T)\sum_{t=1}^T \mathbf{1}\{S_t \leq x_{t,i}\}$. This is an automatic property from the algorithm, coherent to multi-level coverage, with no extra analysis.

Linearized regret preservation through inward flow. Two features of the wrapper need to be reconciled with the regret guarantee: (1) The projection step moves $\tilde{\boldsymbol{x}}_t$ inside $\mathcal{K}$ before predicting, and (2) Alg is updated with $\boldsymbol{g}_t \in \partial\mathcal{L}_t(\boldsymbol{x}_t)$ evaluated at the played iterate rather than at the iterate Alg

produced. Both could in principle inflate the regret. The *inward flow* property of Ding et al. [14] ensures they do not. We restate it for any abstract loss-set pair, since this generality is essential below.

Definition 3.1 (Inward flow [14])**.** A function f paired with a closed convex set $\mathcal{C} \subseteq \mathbb{R}^d$ satisfies the *inward flow* property if there exists a subgradient $\boldsymbol{g}(\boldsymbol{x}) \in \partial f(\boldsymbol{x})$ such that $-\boldsymbol{g}(\boldsymbol{x}) \in T_{\mathcal{C}}(\boldsymbol{x})$ for every boundary point $\boldsymbol{x} \in \mathrm{bd}(\mathcal{C})$, where $T_{\mathcal{C}}(\boldsymbol{x}) \coloneqq \mathrm{cl}\{\boldsymbol{v} : \exists \delta > 0, \boldsymbol{x} + \varepsilon \boldsymbol{v} \in \mathcal{C} \text{ for } \varepsilon \in (0, \delta]\}$ is the tangent cone of $\mathcal{C}$ at $\boldsymbol{x}$, $\mathrm{bd}(\mathcal{S})$ denotes the boundary of a set $\mathcal{S}$, and $\mathrm{cl}(\mathcal{S})$ denotes the closure of $\mathcal{S}$.

Geometrically, the chosen subgradient at the boundary points outward, so the descent direction points inward, and a small step opposite $\boldsymbol{g}(\boldsymbol{x})$ stays feasible. For the isotonic cone, the meaning is clear: two tied levels $x_i = x_{i+1}$ share the same coverage indicator, so $g_i = \mathbf{1}\{x_i \geq S_t\} - (1 - \alpha_i) \leq g_{i+1}$, and a descent step raises the radius at level i at least as much as at level $i + 1$, never creating a crossing. Ding et al. [14, Lemma 2] verifies inward flow for the aggregated pinball $\mathcal{L}_t$ paired with the (decreasing) isotonic cone $\mathcal{K}$ under the choice of subgradient as in (2); we adopt this throughout. The structural consequence relevant to our analysis is the following projection-and-gradient inequality.

Lemma 3.2 (Projection inequality [14])**.** *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, which leverages the duality between the tangent cone (containing $-\boldsymbol{g}_t$) and the normal cone (containing the projection direction $\tilde{\boldsymbol{x}}_t - \boldsymbol{x}_t$), is in Appendix A.1. We state it for a general $\mathcal{C}$ because inward flow composes under intersection with polyhedral sets, so further polyhedral constraints can be layered on top of $\mathcal{K}$ (Appendix A.5).

Lemma 3.2 converts a regret bound on the hidden iterates of Alg into a regret bound on the played iterates and allows us to prove the following regret guarantee for Algorithm 1.

Proposition 3.3 (Linearized regret preservation)**.** *Suppose* Alg *admits a linearized regret bound $F_T(\boldsymbol{u})$ on the gradient sequence $(\boldsymbol{g}_t)_{t \in [T]}$, that is, $\sum_{t=1}^T \langle \boldsymbol{g}_t, \tilde{\boldsymbol{x}}_t - \boldsymbol{u} \rangle \leq F_T(\boldsymbol{u})$ for every $\boldsymbol{u} \in \mathbb{R}^d$. The played iterates $\boldsymbol{x}_t = \Pi_{\mathcal{K}}(\tilde{\boldsymbol{x}}_t)$ produced by Algorithm 1 satisfy*

$$\mathrm{LinRegret}_T(\boldsymbol{u}) \coloneqq \sum_{t=1}^T \langle \boldsymbol{g}_t, \boldsymbol{x}_t - \boldsymbol{u} \rangle \leq F_T(\boldsymbol{u}), \quad \forall \boldsymbol{u} \in \mathbb{R}^d. \tag{4}$$

A consequence that is worth highlighting is that (4) holds for every $\boldsymbol{u} \in \mathbb{R}^d$, not just for $\boldsymbol{u} \in \mathcal{K}$, so the projection is free of charge against any comparator the base learner could compete with. The proof in Appendix A.2 is a direct combination of Lemma 3.2, the assumption on linearized regret, and convexity of $\mathcal{L}_t$.

Characterization of miscoverage level sets. We now translate the regret guarantee into a finite-time bound on the cumulative subgradient. We show a sketch of the reasoning, while the full argument is given in Appendix A.2.

For notational convenience, define $\boldsymbol{G}_T \coloneqq \sum_{t=1}^T \boldsymbol{g}_t \in \mathbb{R}^d$. First, note that (4), by regret–reward duality [27], is equivalent to

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

where in the second and third inequality we used the convexity and non-negativity of $\mathcal{L}_t$. Define $B_T \coloneqq \sum_{t=1}^T \mathcal{L}_t(\boldsymbol{0})$. Since $S_t \geq 0$, $\mathcal{L}_t(\boldsymbol{0}) = \sum_{i=1}^d (1 - \alpha_i) S_t \leq d S_t$, so under Assumption 2.1, $B_T \leq d D T^{1+\gamma}$. Putting all together, we have the following theorem.

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 each $(\mathcal{L}_t, \mathcal{K})$ satisfies inward flow. Then*

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

where $F_T^\star$ is the Fenchel conjugate of F_T on $\mathbb{R}^d$. Equivalently, $\boldsymbol{G}_T$ lies in the level set

$$\mathcal{Z}_T(F_T) \coloneqq \left\{\boldsymbol{z} \in \mathbb{R}^d : F_T^\star(-\boldsymbol{z}) \leq B_T\right\}. \tag{6}$$

The proof follows that of Liu et al. [25, Thm. 3.1] with Lemma 3.2 and convexity of $\mathcal{L}_t$ in place of their Lemma A.1; the direct benefit of this combination is the generality of our results (Table 1). The

geometry of the coverage-error region is characterized entirely in $F_T^\star$, and, through (3), it dictates how the values $\mathrm{MisCov}_{T,i}$ for $i \in [d]$ are controlled. One implication is immediate: *Smaller regret implies tighter coverage.* In particular, optimizing the *dependence of a regret bound on problem parameters* directly tightens the coverage-error region. We show it formally in Appendix A.6.

Remark 3.5. Ding et al. [14] demonstrated that naive approaches to combining single-level coverage with ordering constraints can fail: both post-hoc isotonic projection of independently coverage-controlled quantiles and (non-lazy) projected OGD on the aggregate pinball loss can break per-level coverage. Our results do not contradict their result. Instead, Algorithm 1 and Theorem 3.4 generalize their result by showing that any generic online learning algorithm can be used directly to obtain *multi-level coverage*. In fact, we can convert *any* regret bound for Alg into a deterministic constraint on the miscoverage components $\mathrm{MisCov}_{T,i}$, while guaranteeing distributional consistency too.

The next two subsections instantiate algorithms based on Algorithm 1 and Theorem 3.4.

3.1 Online mirror descent in the p-norm

Fix $p \in (1, 2]$ and let $q = p/(p-1) \in [2, \infty)$. 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 $\tilde{\boldsymbol{x}}_1 = \mathbf{0}$. The resulting wrapped algorithm has a standard mirror-descent dual-update form with an isotonic projection at prediction time; we defer the explicit pseudocode (Algorithm 3) to Appendix C.2.

Standard mirror-descent analysis gives the regret bound and its conjugate [see, e.g., 27, Sec. 6.7]

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

where the conjugate uses $(\frac{1}{2}\|\cdot\|_p^2)^\star = \frac{1}{2}\|\cdot\|_q^2$. Substituting into Theorem 3.4 and bounding $\|\boldsymbol{g}_t\|_q \le d^{1/q}$ (since $\|\boldsymbol{g}_t\|_\infty \le 1$) yields the following ℓ_q coverage-error ball.

Corollary 3.6 (ℓ_q coverage-error ball)**.** *Algorithm 3 with step size $\eta > 0$ satisfies*

$$\max_{i\in[d]} \mathrm{MisCov}_{T,i} \le \frac{\|\boldsymbol{G}_T\|_q}{T} \le \frac{1}{T}\sqrt{\frac{2B_T}{\eta} + \frac{d^{2/q}T}{p-1}}. \tag{7}$$

The bound (7) interpolates between the Euclidean case ($q = 2$, dependence $\sqrt{d/T}$) and the geometry of the unit hypercube ($q \to \infty$). We optimize the parameter pair (p, η). Assuming $d \ge 3$ and choosing $q := 2\log d$ (so $p-1 = 1/(2\log d - 1)$ and $d^{2/q} = e$), the second term inside the square root becomes $e(2\log d - 1)T$. The right-hand side of (7) is strictly decreasing in η, so (7) alone would favor an arbitrarily large step size. We instead balance its two terms by setting $\eta = 2B_T/(e(2\log d - 1)T)$, which gives the $\eta \to \infty$ limit up to a factor $\sqrt{2}$ at a finite step size, and yields

$$\max_{i\in[d]} \mathrm{MisCov}_{T,i} \le \frac{\|\boldsymbol{G}_T\|_{2\ln d}}{T} \le \sqrt{\frac{2e(2\log d - 1)}{T}} = \mathcal{O}\left(\sqrt{\frac{\log d}{T}}\right).$$

A detailed derivation, including verification that the balanced η is well-defined under Assumption 2.1, is in Appendix A.4.

We emphasize that the $\sqrt{\log d}$ dimension dependence represents an exponential improvement over the $\sqrt{d}$ dependence of OGD, i.e., of `MultiQT`, which our analysis recovers as the $p = 2$ case. The improvement lies in d, not T: for a common point forecast with bounded scores, (our setting with $\gamma = 0$), Ding et al. [14, Cor. 1] prove a fast $\mathcal{O}(1/T)$ rate for a *fixed* step size, for which the regret bound grows linearly in T, as for any constant step size; balancing the two yields $\mathcal{O}(T^{-1/2})$ for both [14, Sec. 4.3]. Appendix A.9 states both regimes with constants.

The optimal step size η in p-norm OMD, as in `MultiQT`, depends on B_T, hence on the unknown polynomial-growth parameters D and γ. In Appendix C.5, we give a heuristic and a doubling-trick variant with provable guarantees.

3.2 Coordinate-wise Universal Portfolio

The cleanest way to remove the dependence on B_T in the choice of step size is to replace OMD by a parameter-free algorithm [27, Sec. 10]. Since the aggregate pinball loss decomposes coordinate-wise, $\mathcal{L}_t(\boldsymbol{x}) = \sum_{j=1}^d \ell^{(1-\alpha_j)}(x_j, S_t)$, a natural instantiation is to run an independent copy of the scalar universal-portfolio learner `UP-OCP` [25] at each level α_j, with no shared state between levels, and to apply Algorithm 1 on the joint vector. We refer to this composed algorithm as coordinate-wise `UP-OCP` and defer the complete pseudocode (Algorithm 2) to Appendix C.2.

The total regret of d independent instances is the sum of their per-coordinate regrets [27, Sec. 10.3]: $F_T(\boldsymbol{u}) = \sum_{j=1}^d F_T^{(j)}(|u_j|)$, where $F_T^{(j)}$ is the linearized-regret bound of `UP-OCP` at level α_j [25, Theorem 4.3]. The Fenchel conjugate factorizes accordingly, $F_T^\star(\boldsymbol{\theta}) = \sum_{j=1}^d (F_T^{(j)})^\star(|\theta_j|)$, and Theorem 3.4 becomes

$$\sum_{j=1}^d \left(F_T^{(j)}\right)^\star\left(\,|G_{T,j}|\,\right) \le B_T. \tag{8}$$

The Fenchel conjugate of `UP-OCP` is non-decreasing in $|\theta_j|$ and $(F_T^{(j)})^\star(0) \ge -1$, so dropping all but one term in (8) yields a per-coordinate budget of $B_T + d - 1$. Inverting the resulting Bernstein-type inequality [25, App. F.1] gives the following theorem.

Theorem 3.7 (Per-level coverage of coordinate-wise `UP-OCP`)**.** *Run Algorithm 2. For every $j \in [d]$,*

$$\mathrm{MisCov}_{T,j} \le \hat{\varepsilon}_{T,j} + \sqrt{2\alpha_j(1-\alpha_j)\hat{\varepsilon}_{T,j}}, \quad \hat{\varepsilon}_{T,j} \coloneqq \frac{1}{T}\left[\ln(B_T + d) + \tfrac{1}{2}\ln\left(\pi(T+1)\right)\right]. \tag{9}$$

In particular, with Assumption 2.1, $\max_{j\in[d]}\ \mathrm{MisCov}_{T,j} = \mathcal{O}\left(\sqrt{((1+\gamma)\log T + \log d)/T}\right)$.

The per-level rate matches the p-norm bound of Corollary 3.6 in its dependence on d, up to an additional $\log T$ factor that is the price of being parameter-free. Crucially, no quantity used by the algorithm requires knowledge of D, γ, or T. The factor $\alpha_j(1-\alpha_j)$ also guarantees a faster rate for extreme quantile levels.

Extensions to continuous quantile levels. An important advantage of our framework is that it can be easily generalized from a finite collection $\mathcal{A} = \{\alpha_1, \dots, \alpha_d\}$ to a compact interval $\mathcal{X} \subseteq (0,1)$ by replacing $\mathbb{R}^d$ with the Hilbert space $\mathcal{H} = L^2(\mathcal{X}, \mu)$ for a finite Borel reference measure μ with full support, and replacing the isotonic cone by the cone of non-increasing functions. We collect the formal statements, including a coordinate-wise `UP-OCP` analogue obtained via Donsker–Varadhan change-of-measure, in Appendix A.8.

4 Sharpness and coverage improvement of isotonic projection

Section 3 used the lazy projection structurally, to make the played iterate feasible without inflating the coverage bound. We now show that it does strictly more: whenever the hidden iterate $\tilde{\boldsymbol{x}}_t$ violates isotonicity, projecting it onto $\mathcal{K}$ *strictly* reduces a measure of prediction width while *weakly* reducing a measure of miscoverage, so the empirical Pareto frontier of width versus coverage shifts monotonically toward the ideal corner of small width and accurate coverage.

We frame our analysis in terms of the Weighted Interval Score (WIS) [9], the standard scoring rule for assessing multi-level interval forecasts. WIS is an extension of the classical Interval Score [18, 41] to the multi-level setting. For a collection of interval forecasts $[l_\alpha, u_\alpha]$, $\alpha \in \mathcal{A}$, the WIS is defined as

$$\mathrm{WIS}_{\mathcal{A}}\left(\{[l_\alpha, u_\alpha]\}_{\alpha\in\mathcal{A}}, y\right) \coloneqq \underbrace{\sum_{\alpha\in\mathcal{A}} \alpha\,(u_\alpha - l_\alpha)}_{\textit{sharpness reward}} + \underbrace{\sum_{\alpha\in\mathcal{A}} 2\cdot \mathrm{dist}\,(y, [l_\alpha, u_\alpha])}_{\textit{coverage penalty}}, \tag{10}$$

where $\mathrm{dist}(y, [l_\alpha, u_\alpha]) = \inf_{x\in[l_\alpha,u_\alpha]} |y - x|$. For symmetric prediction sets $\hat{C}_t(\hat{Y}_t, x) = [\hat{Y}_t - x, \hat{Y}_t + x]$ at miscoverage levels $\mathcal{A} = \{\alpha_1, \dots, \alpha_d\}$, a direct calculation (Appendix B.1) gives

$$\mathcal{L}_t(\boldsymbol{x}) + S_t \sum_{i=1}^d \alpha_i = \frac{1}{2}\,\mathrm{WIS}_{\mathcal{A}}\left(\left\{\hat{C}_t(\hat{Y}_t, x_i)\right\}_i, Y_t\right) = \underbrace{\sum_{i=1}^d \alpha_i x_i}_{\textit{sharpness reward}} + \underbrace{\sum_{i=1}^d \max(S_t - x_i, 0)}_{\textit{coverage penalty}}. \tag{11}$$

This means that, up to an additive term that is constant in $\boldsymbol{x}_t$, the aggregate pinball loss *is* half the WIS of our prediction sets. The sharpness term is an α-weighted measure of total predicted width, with higher miscoverage levels (smaller intervals) weighted more; the coverage penalty $\sum_i \max(S_t - x_i, 0)$ is a one-sided ℓ_1 residual that penalizes under-coverage at every level.

Our main result in this section is that the projection step in Algorithm 1 reduces both terms of (11) at every round (and hence WIS itself), and strictly so whenever the hidden iterate is not already isotonic.

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

$$\sum_{i=1}^{d} \alpha_i x_{t,i} \leq \sum_{i=1}^{d} \alpha_i \tilde{x}_{t,i} \quad \textit{and} \quad \sum_{i=1}^{d} \max(S_t - x_{t,i}, 0) \leq \sum_{i=1}^{d} \max(S_t - \tilde{x}_{t,i}, 0). \tag{12}$$

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

The proof hinges on the characterization of Pool-Adjacent-Violators (PAVA) algorithm used for the isotonic projection [30]. The full proof, including the supporting lemmas, is in Appendix B.2. The inequalities hold for any $\tilde{\boldsymbol{x}}_t \in \mathbb{R}^d$; with valid hidden radii (Section 3), both sides are WIS components of genuine prediction intervals, since PAVA averages non-negative radii.

Ding et al. [14] obtained a decomposition analogous to (11) under the assumption that $\mathcal{A}$ is symmetric about $1/2$, but they only used it as a motivation to minimize the regret with $\mathcal{L}_t$. In fact, they observed a better empirical performance of `MultiQT` compared to the *naive forecasters*, but they did not provide any explanation for it.Our Theorem 4.1 provides this missing explanation.

5 Experiments

Here, we evaluate our framework on both synthetic and real-world datasets.

Datasets. We follow the OCP literature [1, 16, 17, 29] and use three categories of streams: (1) daily opening prices of major US stocks (AAPL, AMZN, GOOGL) from 2006 to 2017 [26]; (2) electricity demand records from New South Wales [20]; and (3) synthetic environments following Angelopoulos et al. [1, App. F.5] that bypass forecasting and directly simulate the nonconformity scores S_t under three patterns (sinusoidal with Gaussian noise; constant and quadratic trends with sparse bumps), allowing repeated runs for error bars. Forecasting pipelines (daily-refit Prophet for stocks [38], AR(3) for electricity) and full synthetic-stream specifications are in Appendix C.1.

Methods. We compare four method families: `UP-OCP` (Algorithm 2, parameter-free), p-`norm OMD` (Algorithm 3, with $p = 2\log d/(2\log d - 1)$ from Corollary 3.6 and a constant step size chosen per dataset by oracle grid search on a base-2 logarithmic grid), `KT` (Algorithm 4, parameter-free reference), and the OGD-based prior baseline `MultiQT` [14] (adaptive; uses a data-dependent learning-rate schedule). The first three are run both as vanilla learners and wrapped with the lazy projection of Algorithm 1 (suffix -ISO); `MultiQT` is by construction projected. Update rules, hyperparameter grids, and selected η values are in Appendix C.2.

Metrics. We track, per method, (1) the *empirical coverage* $\mathrm{Coverage}_{T,j}$ as a function of the target $1-\alpha$; (2) the per-round *prediction-set width* $|\hat{C}_t(\hat{Y}_t, x_{t,j})|$, summarized by its mean (the standard sharpness measure), median, and 95th percentile; (3) the *full distribution* of widths at selected levels; and (4) the *WIS* (10), summed over all d levels and averaged over rounds, with its two components.

The scores are the absolute residuals $S_t = |Y_t - \hat{Y}_t|$; none is exactly zero (the smallest is about 10^{-5}), so the empty-set convention never changes a coverage indicator. In all cases we calibrate $d = 100$ levels uniformly spaced in $[0.001, 0.5]$ for each method, with warm-up periods of $T_{\text{burn-in}} = 100$ rounds (300 for `elec2`) during which the predictor outputs sets but is not updated. Results are reported on the practically interesting window $\alpha \in [0.03, 0.25]$ (target coverage $1-\alpha \in [0.75, 0.97]$). We focus on the AAPL dataset and the synthetic sinusoidal stream; full results for the other datasets are in Appendix C.

Calibration and Pareto frontiers. Higher coverage requires larger sets, but this trade-off varies across algorithms. Following Liu et al. [25], we assess it through a *Pareto frontier*: Figure 1 shows,

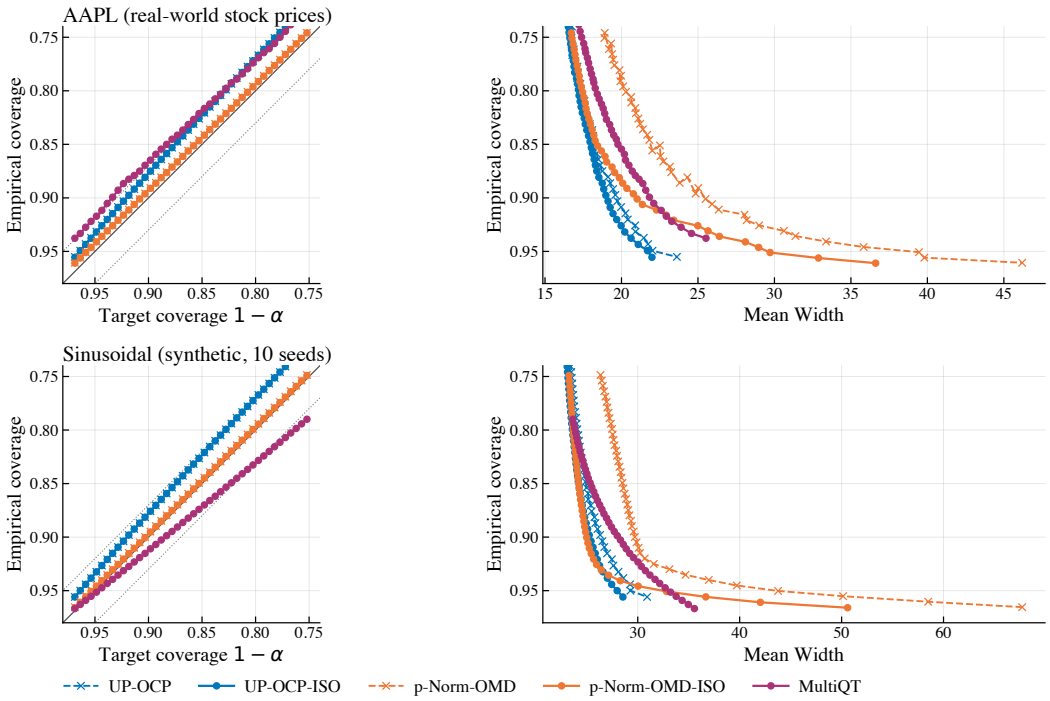


Figure 1: Calibration (left column: empirical versus target coverage $1-\alpha$) and Pareto frontiers (right column: empirical coverage versus mean prediction-set width), for AAPL (top row) and the sinusoidal dataset (bottom row, mean across 10 seeds). The solid diagonal implies perfect calibration $y = x$, the dotted lines are the ± 0.03 error band. Better performance is closer to the bottom-left corner.

per dataset, a calibration curve (left) and the empirical coverage against the mean width (right) for $\alpha \in [0.03, 0.25]$. We prioritize the mean width, which reflects the cumulative cost of uncertainty downstream and, unlike quantiles of the width, captures unreasonably large sets. We exclude `KT`, whose width is substantially larger at almost all levels and would compress every other curve; it is included in Figure 2 and Appendix C. Several observations follow.

Calibration is preserved by the projection. In each calibration panel, the projected (`ISO`) and unprojected curves are visually indistinguishable; the p-norm curves lie inside the ± 0.03 tolerance band over the entire α window, and the `UP-OCP` curves leave it only slightly, at the lowest target coverages. This is consistent with Proposition 3.3: the projection does not consume any regret, so it does not consume any coverage.

Sharpness is strictly improved by the projection, for almost all α. On the right-hand panels, every solid (-ISO) curve sits on or to the left of its dashed (unprojected) counterpart, i.e., it is narrower at the same coverage, with the gap widening at small α (high coverage), where the hidden iterate is most likely to violate isotonicity. The reduction is largest for p-`norm OMD`, whose unprojected mean width is up to 33% (AAPL) and 39% (sinusoidal) larger than the projected one near $1-\alpha = 0.97$; for `UP-OCP` it is about 7–8%.

Coordinate-wise `UP-OCP` uniformly improves over p-`norm OMD` and `MultiQT` in width. On AAPL, `UP-OCP-ISO` is narrower than both p-`norm OMD` variants and `MultiQT` at every α. On AAPL, tuned p-norm OMD with the optimal p achieves the highest empirical coverage but pays for it with much wider sets. On both datasets, the constant step size selected from our grid by Pareto dominance is $\eta = 64$. Oracle tuning gives p-`norm OMD` an advantage that is unavailable in practice, while `UP-OCP` and `KT` are parameter-free; still, the results are not fragile to it: every $\eta \in [16, 128]$ stays within about 10% of the best WIS on both datasets, and a plug-in step size using only revealed scores comes within 4% (Appendix C.5).

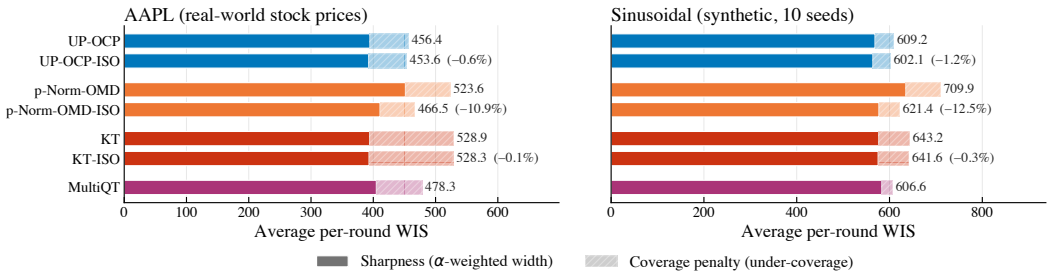


Figure 2: Average per-round WIS over all $d = 100$ levels (post burn-in; sinusoidal: mean over 10 seeds), split into the sharpness (solid) and coverage-penalty (hatched) components of (10). Percentages give the change of each projected variant relative to its unprojected counterpart.

Remark 5.1 (Slight miscoverage of the parameter-free and adaptive instances). There is a fundamental trade-off between marginal coverage and prediction set size in OCP, as shown in recent lower bounds arguments [37]: any algorithm whose prediction sets have size at most μ times the optimal must incur $\Omega\left(\log(1/\varepsilon_{\min})/\log\mu \cdot \alpha T\right)$ coverage errors, where $\varepsilon_{\min}$ is the smallest relevant scale. As suggested by Liu et al. [25], empirically, this can be fixed by a simple α-correction technique.

Tables in Appendix C.7 report the same comparison at three representative target-coverage levels $1-\alpha \in \{0.95, 0.9, 0.8\}$ for every dataset, including `KT` and the median and 95th-percentile widths (Table 6 for AAPL).

Weighted Interval Score. The projection lowers the total WIS (Figure 2) of every base learner on both datasets, most for p-`norm OMD` (-10.9% on AAPL, -12.5% on the sinusoidal stream), and `UP-OCP-Iso` attains the smallest WIS on both. Since Figure 2 compares runs with different trajectories, we also checked Theorem 4.1 in its exact within-run form: the hidden iterate violates isotonicity in 99.6% of the rounds, and in every one of them the projection strictly reduced the WIS and its sharpness component without increasing the coverage penalty (Appendix C.4, which also reports the WIS on every dataset).

Further results. The projection also removes the heavy right tail of the per-round width distribution (Appendix C.3). As d grows from 5 to 1,000 on the sinusoidal stream, the realized worst-level errors grow at most slowly, so the worst-case $\sqrt{d}$ and $\sqrt{\log d}$ scalings are not visible there; the methods designed through the tighter analysis are nonetheless more accurate at every d (Appendix C.6).

6 Conclusion, limitations, and future work

We presented a black-box reduction that turns any online learner with a linearized regret bound into a multi-level OCP algorithm with *multi-level coverage*, by wrapping the learner in a lazy isotonic projection layer. The reduction preserves regret (and therefore coverage) while strictly improving sharpness, formalized through the Weighted Interval Score. Instantiating the framework with p-norm OMD and coordinate-wise Universal Portfolio replaces the $\sqrt{d}$ dependence in the per-level miscoverage rate of `MultiQT` with $\sqrt{\log d}$, with the coordinate-wise Universal Portfolio instantiation being parameter-free. Extensive experiments confirm that the projection wrapper is a free (in terms of order of computational complexity) per-round improvement on every base learner we tested, and that `UP-OCP` delivers smaller and better-behaved prediction sets than the tuned baselines.

Limitations and future work. Our guarantees concern long-run *marginal* coverage, the standard target in adversarial OCP. Coverage conditional on groups (attainable at a single level [6]) or on time intervals remains open in the multi-level setting and would likely require multi-validity [5] or strongly adaptive online learning [7].

Acknowledgments and Disclosure of Funding

ED's work was supported in part by the US NSF, ARO, ONR, and the Sloan Foundation.

References

[1] Anastasios Angelopoulos, Emmanuel Candes, and Ryan J Tibshirani. Conformal PID control for time series prediction. *Advances in neural information processing systems*, 36:23047–23074, 2023.

[2] Anastasios N Angelopoulos, Stephen Bates, Adam Fisch, Lihua Lei, and Tal Schuster. Conformal risk control. *arXiv preprint arXiv:2208.02814*, 2022.

[3] Anastasios N Angelopoulos, Michael I Jordan, and Ryan J Tibshirani. Gradient equilibrium in online learning: Theory and applications. *arXiv preprint arXiv:2501.08330*, 2025.

[4] Rina Foygel Barber, Emmanuel J Candes, Aaditya Ramdas, and Ryan J Tibshirani. Predictive inference with the jackknife+. *The Annals of Statistics*, 49(1):486–507, 2021.

[5] Osbert Bastani, Varun Gupta, Christopher Jung, Georgy Noarov, Ramya Ramalingam, and Aaron Roth. Practical adversarial multivalid conformal prediction. *Advances in neural information processing systems*, 35:29362–29373, 2022.

[6] Beepul Bharti, Ambar Pal, Jacopo Teneggi, and Jeremias Sulam. Parameter-free and group conditional online conformal prediction. *arXiv preprint arXiv:2606.00419*, 2026.

[7] Aadyot Bhatnagar, Huan Wang, Caiming Xiong, and Yu Bai. Improved online conformal prediction via strongly adaptive online learning. In *International Conference on Machine Learning*, pages 2337–2363. PMLR, 2023.

[8] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.

[9] Johannes Bracher, Evan L Ray, Tilmann Gneiting, and Nicholas G Reich. Evaluating epidemic forecasts in an interval format. *PLoS computational biology*, 17(2):e1008618, 2021.

[10] Yang Cai, Constantinos Daskalakis, Haipeng Luo, Chen-Yu Wei, and Weiqiang Zheng. On tractable ϕ-equilibria in non-concave games. *Advances in Neural Information Processing Systems*, 37:140366–140404, 2024.

[11] Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.

[12] Nicolò Cesa-Bianchi and Francesco Orabona. Online learning algorithms. *Annual Review of Statistics and Its Application*, 8:165–190, 2021.

[13] Thomas M Cover and Erik Ordentlich. Universal portfolios with side information. *IEEE Transactions on Information Theory*, 42(2):348–363, 2002.

[14] Tiffany Ding, Isaac Gibbs, and Ryan J Tibshirani. Calibrated multi-level quantile forecasting. *arXiv preprint arXiv:2512.23671*, 2025.

[15] Claudio Gentile. The robustness of the p-norm algorithms. *Machine Learning*, 53(3):265–299, 2003.

[16] Isaac Gibbs and Emmanuel Candes. Adaptive conformal inference under distribution shift. *Advances in Neural Information Processing Systems*, 34:1660–1672, 2021.

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

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

[19] Chirag Gupta, Arun K Kuchibhotla, and Aaditya Ramdas. Nested conformal prediction and quantile out-of-bag ensemble methods. *Pattern Recognition*, 127:108496, 2022.

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

[21] Elad Hazan. Introduction to online convex optimization. *Foundations and Trends in Optimization*, 2(3-4):157–325, 2016.

[22] Yonghoon Lee, Edgar Dobriban, and Eric Tchetgen Tchetgen. Simultaneous conformal prediction of missing outcomes with propensity score ε-discretization. *arXiv preprint arXiv:2403.04613*, 2024.

[23] Jing Lei, James Robins, and Larry Wasserman. Distribution-free prediction sets. *Journal of the American Statistical Association*, 108(501):278–287, 2013.

[24] Jing Lei, Max G'Sell, Alessandro Rinaldo, Ryan J Tibshirani, and Larry Wasserman. Distribution-free predictive inference for regression. *Journal of the American Statistical Association*, 113(523):1094–1111, 2018.

[25] Tuo Liu, Edgar Dobriban, and Francesco Orabona. Online conformal prediction via universal portfolio algorithms. In *Proceedings of the 43rd International Conference on Machine Learning*, volume 306 of *Proceedings of Machine Learning Research*, pages 75891–75952. PMLR, 2026.

[26] Cam Nguyen. S&P 500 stock data, 2018. URL `https://www.kaggle.com/datasets/camnugent/sandp500`.

[27] Francesco Orabona. A modern introduction to online learning. *arXiv preprint arXiv:1912.13213*, 2019.

[28] Harris Papadopoulos, Kostas Proedrou, Volodya Vovk, and Alex Gammerman. Inductive confidence machines for regression. In *European conference on machine learning*, pages 345–356. Springer, 2002.

[29] Aleksandr Podkopaev, Darren Xu, and Kuang-chih Lee. Adaptive conformal inference by betting. In *Proceedings of the 41st International Conference on Machine Learning*, pages 40886–40907, 2024.

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

[31] R Tyrrell Rockafellar. *Convex analysis*, volume 28. Princeton university press, 1997.

[32] Yaniv Romano, Evan Patterson, and Emmanuel Candes. Conformalized quantile regression. *Advances in neural information processing systems*, 32, 2019.

[33] Yaniv Romano, Matteo Sesia, and Emmanuel Candes. Classification with valid and adaptive coverage. *Advances in neural information processing systems*, 33:3581–3591, 2020.

[34] Craig Saunders, Alexander Gammerman, and Volodya Vovk. Transduction with confidence and credibility. In *IJCAI*, 1999.

[35] Glenn Shafer and Vladimir Vovk. A tutorial on conformal prediction. *Journal of Machine Learning Research*, 9(3):371–421, 2008.

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

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

[38] Sean J Taylor and Benjamin Letham. Forecasting at scale. *The American Statistician*, 72(1): 37–45, 2018.

[39] Vladimir Vovk, Alexander Gammerman, and Glenn Shafer. *Algorithmic Learning in a Random World*. Springer Science & Business Media, 2005.

[40] Volodya Vovk, Alexander Gammerman, and Craig Saunders. Machine-learning applications of algorithmic randomness. In *Proceedings of the Sixteenth International Conference on Machine Learning*, pages 444–453, 1999.

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

[42] Margaux Zaffran, Olivier Féron, Yannig Goude, Julie Josse, and Aymeric Dieuleveut. Adaptive conformal predictions for time series. In *International conference on machine learning*, pages 25834–25866. PMLR, 2022.

[43] Zhiyu Zhang, David Bombara, and Heng Yang. Discounted adaptive online learning: towards better regularization. In *Proceedings of the 41st International Conference on Machine Learning*, pages 58631–58661, 2024.

[44] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proc. of the International Conference on Machine Learning*, pages 928–936, 2003.

Appendix

A Deferred Proofs and supporting material for Section 3 **14**

A.1 Projection inequality (Lemma 3.2) . . . 15
A.2 Regret preservation and the central inequality . . . 15
A.3 Inward flow for the isotonic cone . . . 15
A.4 Detailed derivation for Corollary 3.6: p-norm OMD . . . 16
A.5 Beyond the isotonic cone . . . 16
A.6 Tighter regret implies tighter coverage . . . 17
A.7 Coordinate-wise UP-OCP: per-level and mean-absolute proofs . . . 17
A.8 Continuous quantile levels . . . 19
A.9 Detailed comparison with prior work . . . 21

B Deferred proofs for Section 4 **22**

B.1 Pinball loss as a Weighted Interval Score (Equation (11)) . . . 22
B.2 Per-round structural reduction (Theorem 4.1) . . . 23

C Experimental details and additional results **24**

C.1 Datasets and forecasting pipelines . . . 24
C.2 Detailed update rules . . . 25
C.2.1 Coordinate-wise Universal Portfolio (`UP-OCP`) . . . 25
C.2.2 p-norm Online Mirror Descent . . . 25
C.2.3 Coordinate-wise Krichevsky–Trofimov bettor (`KT`) . . . 26
C.2.4 `MultiQT` . . . 27
C.3 Width distribution analysis . . . 27
C.4 Weighted Interval Score . . . 28
C.5 Step-size sensitivity and adaptive step sizes for p-norm OMD . . . 29
C.6 Varying the number of levels . . . 30
C.7 Additional results . . . 31
C.7.1 AAPL (full results) . . . 32
C.7.2 AMZN . . . 33
C.7.3 GOOGL . . . 34
C.7.4 Electricity demand (`elec2`) . . . 35
C.7.5 Sinusoidal (synthetic) . . . 36
C.7.6 Stationary trend with random waves (synthetic) . . . 37
C.7.7 Quadratic drift with random waves (`mixed`, synthetic) . . . 38

A 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

$$\boldsymbol{\delta}_T \coloneqq \frac{\boldsymbol{G}_T}{T} \in \mathbb{R}^d$$

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

A.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 A.8. Recall the *normal cone* of a closed convex set C at $\boldsymbol{x} \in C$, $N_C(\boldsymbol{x}) \coloneqq \{\boldsymbol{v} \in \mathbb{R}^d : \langle \boldsymbol{v}, \boldsymbol{y} - \boldsymbol{x}\rangle \leq 0 \text{ for all } \boldsymbol{y} \in C\}$, 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 $\boldsymbol{w} \in T_C(\boldsymbol{x})$ [31, Sec. 6].

Lemma A.1. *Let ℓ 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 A.1 with $\boldsymbol{x} = \boldsymbol{x}_0 \in \mathrm{bd}(C)$. □

A.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 $\mathrm{Regret}_T(\boldsymbol{u}) \leq \mathrm{LinRegret}_T(\boldsymbol{u})$, so the bound (4) proved below also bounds the regret. 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 [27, 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

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

A.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 $I(\boldsymbol{x}) \coloneqq \{i \in [d-1] : x_i = x_{i+1}\}$ 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 $g_{t,i} = \mathbf{1}\{x_i \geq S_t\} - (1 - \alpha_i)$, where the coverage indicator follows the empty-set convention of Section 2; in particular, it depends on x_i only through its value, so tied coordinates share it. Take $i \in I(\boldsymbol{x})$ so that $x_i = x_{i+1}$.

Case 1: both levels cover, i.e., $x_i = x_{i+1} \geq S_t$ and $x_i > 0$. 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: neither level covers, i.e., $x_i = x_{i+1} < S_t$, or $x_i = x_{i+1} = 0$ (the empty set). 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. □

A.4 Detailed derivation for Corollary 3.6: p-norm OMD

The regret bound for p-norm OMD with regularizer $\psi(\boldsymbol{x}) = (1/2)\|\boldsymbol{x}\|_p^2$ at constant step size $\eta > 0$ is standard [27, Theorem 6.8]: ψ is $(p-1)$-strongly convex with respect to $\|\cdot\|_p$ by [36, 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 $(\eta/2)\|\boldsymbol{G}_T\|_q^2 - (\eta/(2(p-1)))\sum_{t=1}^{T}\|\boldsymbol{g}_t\|_q^2 \leq B_T$, i.e.,

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

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

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,

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

Setting $\eta \coloneqq 2B_T/(e(2\log d - 1)T)$ balances the two terms and gives $\|\boldsymbol{G}_T\|_q^2 \leq 2e(2\log d - 1)T$. This choice of η is well-defined whenever $B_T > 0$; the trivial case $B_T = 0$ corresponds to $S_t = 0$ for all t, where (7) holds for every $\eta > 0$ and already gives $\|\boldsymbol{G}_T\|_q^2 \leq e(2\log d - 1)T$. Under Assumption 2.1, the upper bound $B_T \leq dDT^{1+\gamma}$ gives $\eta \leq 2dDT^\gamma/(e(2\log d - 1)) = O(dDT^\gamma/\log d)$ as a function of T.

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

On the choice of η. The right-hand side of (7) is strictly decreasing in η, with infimum $\frac{1}{T}\sqrt{d^{2/q}T/(p-1)}$ as $\eta \to \infty$; the balanced η above attains it up to a factor $\sqrt{2}$, at a finite step size. Small miscoverage alone is a weak target: the data-free predictor that outputs $\mathbb{R}$ on the first $\lceil(1-\alpha_i)T\rceil$ rounds and $\emptyset$ afterwards attains $\mathrm{MisCov}_{T,i} \leq 1/T$ (see [25, App. H.5] for a horizon-free variant). The regret bound is what controls efficiency; for instance, when the scores are i.i.d. (under regularity conditions), sublinear regret drives the average radius to the oracle quantile [25, Thm. D.3]. Finally, there is no contradiction with Proposition A.3: for $\eta' > \eta$, the two OMD bounds satisfy $F_T^{\eta'}(\boldsymbol{u}) > F_T^{\eta}(\boldsymbol{u})$ if and only if $\|\boldsymbol{u}\|_p^2 < \eta\eta'\sum_t \|\boldsymbol{g}_t\|_q^2/(p-1)$, so neither dominates the other uniformly and the proposition does not apply.

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

(i) (ℓ, C_1) and (ℓ, 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 $\mathcal{K}_R \coloneqq \mathcal{K} \cap \{\boldsymbol{x} \in \mathbb{R}^d : x_j \le R \text{ for all } j\}$ enforces both isotonicity and an upper bound. To verify inward flow on $\mathcal{K}_R$, by Proposition A.2 it suffices to verify it on the upper-bound polyhedron $C_R \coloneqq \{\boldsymbol{x} : x_j \le R\}$. 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 A.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.

A.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. [25, Sec. 4] to the vector setting and underlies the design of the case studies in Sections 3.1–3.2.

Proposition A.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* (6) *are nested:*

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

Consequently, if Alg_A admits a uniformly tighter regret bound than Alg_B, the worst-case miscoverage region delivered by Theorem 3.4 for Alg_A is contained in that for Alg_B.

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

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

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

A.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 3.7 and Theorem A.6. Throughout, the Fenchel conjugate of the UP-OCP regret bound at level α has the form $(F_T^{(\alpha)})^\star(|\theta|) = \exp(h_\alpha(\theta)) - 1$ with

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

see [25, Theorem 4.3 and App. E.4].

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

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

Lemma A.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 $f'(\theta) = \theta(2a + b\theta)/(a + b\theta)^2 \ge 0$ and $f''(\theta) = 2a^2/(a + b\theta)^3 > 0$. □

Proof of Theorem 3.7. Substituting the form of $(F_T^{(j)})^\star$ into (8), $\sum_{j=1}^d \exp(h_{\alpha_j}(|G_{T,j}|)) \le B_T + d$. Each term is non-negative, so we may drop all but the j-th: $\exp(h_{\alpha_j}(|G_{T,j}|)) \le B_T + d$, equivalently

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

Lemma A.4 with $\sigma = \alpha_j(1-\alpha_j)$ and $c = T\hat{\varepsilon}_{T,j}$ gives $|G_{T,j}| \le T\sqrt{2\alpha_j(1-\alpha_j)\hat{\varepsilon}_{T,j}} + (2/3)T\hat{\varepsilon}_{T,j}$. Dividing by T and using $2/3 < 1$ yields (9). The asymptotic claim uses $B_T \le dDT^{1+\gamma}$ 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 (9) can be aggregated into a mean-absolute coverage bound across levels. By dropping the variance budgets at individual levels in favor of their envelope $\sigma_{\mathcal{A}}$ and applying Jensen's inequality to the resulting sum, one obtains a single ℓ_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. [25, Theorem 4.4] and, by norm equivalence, improves any ℓ_2 bound by a factor of $1/\sqrt{d}$ on the same quantity (cf. Remark A.7 below).

Theorem A.6 (ℓ_1 miscoverage error of coordinate-wise `UP-OCP`)**.** *Run Algorithm 2, and let $\sigma_{\mathcal{A}} \coloneqq \max_{i\in[d]} \alpha_i(1-\alpha_i) \in [0, 1/4]$ and $C_T \coloneqq B_T + d$. Then*

$$\frac{1}{d}\sum_{i=1}^d |\delta_{T,i}| \;\le\; \sqrt{\frac{2\sigma_{\mathcal{A}}\Lambda_T}{T}} + \frac{2\Lambda_T}{3T}, \qquad \Lambda_T \;\coloneqq\; \log\frac{C_T}{d} + \frac{1}{2}\log\big(\pi(T+1)\big). \tag{13}$$

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 A.7 (Comparison with the ℓ_2-bound of Ding et al. [14]). By norm equivalence, $\frac{1}{d}\|\boldsymbol{\delta}_T\|_1 \le \frac{1}{\sqrt{d}}\|\boldsymbol{\delta}_T\|_2$, so (13) delivers a $1/\sqrt{d}$ improvement over any ℓ_2 bound on the mean-absolute miscoverage. Combined with Theorem 3.7, this improves the dependence on d over the ℓ_2 guarantee of `MultiQT`, at the price of a $\log T$ factor; it is not uniformly tighter in T, since for a fixed step size the bounds of `MultiQT` carry no $\log T$ factor and, for a common point forecast, decay as $\mathcal{O}(1/T)$ (Appendix A.9).

Proof of Theorem A.6. Substituting the form of $(F_T^{(j)})^\star$ into (8),

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

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 $h_{\sigma_{\mathcal{A}}}(\theta) \coloneqq \theta^2/(2\sigma_{\mathcal{A}}T + (2/3)\theta) - \frac{1}{2}\log(\pi(T+1))$). Hence

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

By Lemma A.5, $\theta \mapsto h_{\sigma_{\mathcal{A}}}(\theta) + \frac{1}{2}\log(\pi(T+1))$ 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

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

Taking logarithms and rearranging,

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

Lemma A.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 (13). 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 (13) is $\mathcal{O}(\sqrt{\sigma_{\mathcal{A}} \log T/T})$. □

A.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 μ 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 $F(x) := -\int_{\alpha_{\min}}^{x} b(\alpha)\, d\mu(\alpha)$ [30].

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 A.3, applied pointwise. The central inequality reads

$$F_T^\star(-G_T) \le B_T, \qquad G_T := \sum_{t=1}^{T} g_t \in \mathcal{H}, \qquad B_T := \sum_{t=1}^{T} \mathcal{L}_t(0). \tag{14}$$

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

$$\|\delta_T\|_{L^2(\mathcal{X},\mu)} \le \sqrt{\frac{2\mu(\mathcal{X})}{T}}, \qquad \delta_T := G_T/T. \tag{15}$$

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 $\|\cdot\|_{L^p}$ [36, 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 $\|g_t\|_{L^q} \le \mu(\mathcal{X})^{1/q}$,

$$\|\delta_T\|_{L^q} \le \frac{1}{T}\sqrt{\frac{2B_T}{\eta} + \frac{\mu(\mathcal{X})^{2/q}\, T}{p-1}}. \tag{16}$$

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

Remark A.8 (Optimal q depends on $\mu(\mathcal{X})$). Differentiating $g(q) \coloneqq (q-1)\mu(\mathcal{X})^{2/q}$ in $q \in [2,\infty)$ shows that g is non-decreasing on $[2,\infty)$ unless $\mu(\mathcal{X}) > e^2$ (since $q^2/(q-1) \geq 4$ there), in which case the unique stationary point in $[2,\infty)$ satisfies $q^2 = 2(q-1)\log\mu(\mathcal{X})$ and is approximately $2\log\mu(\mathcal{X})$ for large $\mu(\mathcal{X})$. For the Lebesgue measure or any probability measure on $\mathcal{X} \subseteq (0,1)$, $\mu(\mathcal{X}) \leq 1 < e^2$, so g is increasing on $[2,\infty)$ and the optimum is $q = 2$, recovering the dimension-free L^2 bound (15). 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 beneficial only when the analogue $\mu(\mathcal{X}) > e^2$ holds, which fails under Lebesgue and any normalized probability measure on $\mathcal{X}$. In those cases, the L^2 bound (15) *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 (14) becomes

$$\int_{\mathcal{X}} \exp\big(h_\alpha(|G_T(\alpha)|)\big)\, d\mu(\alpha) \leq C_T \coloneqq B_T + \mu(\mathcal{X}), \tag{17}$$

where h_α is as defined in Appendix A.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 ν on $\mathcal{X}$.

Lemma A.9 (Donsker–Varadhan change-of-measure; [8, Sec. 2.1])**.** *Let $(\mathcal{X}, \mathcal{B})$ be a measurable space, $\bar{\mu}$ a probability measure on $\mathcal{X}$, and ν a probability measure with $\nu \ll \bar{\mu}$. Then for every bounded measurable $\Phi\colon \mathcal{X} \to \mathbb{R}$,*

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

Fix any probability measure $\nu \ll \bar{\mu}$ on $\mathcal{X}$ (with $\bar{\mu} \coloneqq \mu/\mu(\mathcal{X})$), and define the ν-variance envelope $\sigma_\nu \coloneqq \sup_{\alpha \in \mathrm{supp}(\nu)} \alpha(1-\alpha)$ and the ν-mean miscoverage $\bar{G}_\nu \coloneqq \int_{\mathcal{X}} |G_T(\alpha)|\, d\nu(\alpha)$. Lemma A.9 with $\Phi(\alpha) = h_\alpha(|G_T(\alpha)|)$, combined with (17) (so $\log \int e^{\Phi}\, d\bar{\mu} \leq \log(C_T/\mu(\mathcal{X}))$), yields

$$\int_{\mathcal{X}} h_\alpha(|G_T(\alpha)|)\, d\nu(\alpha) \leq \mathrm{KL}(\nu\|\bar{\mu}) + \log\frac{C_T}{\mu(\mathcal{X})}. \tag{18}$$

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

$$\int_{\mathcal{X}} |\delta_T(\alpha)|\, d\nu(\alpha) \leq \sqrt{\frac{2\sigma_\nu R_\nu}{T}} + \frac{2R_\nu}{3T}. \tag{19}$$

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 $\mathrm{KL}(\nu\|\bar{\mu}) = \log(\mu(\mathcal{X})/\mu(\mathcal{I}))$, $\sigma_\nu = \sigma_{\mathcal{I}} \coloneqq \sup_{\alpha\in\mathcal{I}} \alpha(1-\alpha)$, and $\bar{G}_\nu = \bar{G}_{\mathcal{I}} \coloneqq \frac{1}{\mu(\mathcal{I})}\int_{\mathcal{I}} |G_T(\alpha)|\, d\mu(\alpha)$, so (19) reads

$$\frac{1}{\mu(\mathcal{I})}\int_{\mathcal{I}} |\delta_T(\alpha)|\, d\mu(\alpha) \leq \sqrt{\frac{2\sigma_{\mathcal{I}} R_{\mathcal{I}}}{T}} + \frac{2R_{\mathcal{I}}}{3T}, \qquad R_{\mathcal{I}} \coloneqq \log\frac{C_T}{\mu(\mathcal{I})} + \tfrac{1}{2}\log(\pi(T+1)). \tag{20}$$

Specializing further to $\mathcal{I} = \mathcal{X}$ recovers the global mean-absolute miscoverage of Theorem A.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 $\mathcal{X} = \mathcal{A} = \{\alpha_1, \ldots, \alpha_d\}$ and μ the counting measure recovers $\mu(\mathcal{X}) = d$ and reduces every continuous bound to its finite analogue: (15) becomes the OGD bound, (16) becomes Corollary 3.6, (17) becomes (8), and (20) with $\mathcal{I} = \mathcal{X}$ recovers Theorem A.6. In particular, the dichotomy in Remark A.8 explains why the discrete choice $q = 2\log d$ helps for many levels (the counting measure makes $\mu(\mathcal{X}) = d$, and $d > e^2 \approx 7.4$) but does not transfer to Lebesgue.

A.9 Detailed comparison with prior work

Table 1: Comparison with the closest prior work. Rates are for $\max_i \mathrm{MisCov}_{T,i}$; $d_{\mathcal{A}} := \min_i \min\{\alpha_i, 1-\alpha_i\}$ ($= \alpha_1$ when $\alpha_1 + \alpha_d \le 1$).

	Liu et al. [25]	Ding et al. [14] (`MultiQT`)	This paper
Scores	$S_t \le Dt^{\gamma}$	$S_t \le D$	$S_t \le DT^{\gamma}$ (Ass. 2.1)
Algorithms	any scalar learner with linearized regret	OGD with lazy isotonic projection	any learner on $\mathbb{R}^d$ with linearized regret (Alg. 1, Prop. 3.3)
Guarantee	scalar Fenchel level set	ℓ_2 ball, via restorativity	level set $F_T^\star(-\boldsymbol{G}_T) \le B_T$ (Thm. 3.4)
Rate, general	single level only	$\mathcal{O}\big(\sqrt{d/T} + d\sqrt{D/(\eta\, d_{\mathcal{A}} T)}\big)$	$\mathcal{O}(\sqrt{\log d/T})$ (Cor. 3.6); parameter-free: $\mathcal{O}\big(\sqrt{(\log d + (1+\gamma)\log T)/T}\big)$ (Thm. 3.7)
Rate, common point forecast	n/a	$\mathcal{O}\big(d^{3/2}(1 + D/\eta)/(d_{\mathcal{A}} T)\big)$, fixed η, linear regret	not recovered by our analysis
Other results	none	sharpness gains (empirical)	WIS reduction (Thm. 4.1); continuous levels (App. A.8)

Table 1 compares our results with the two works we build on, and this subsection details it. We state the results of Ding et al. [14] in our notation: with symmetric sets, every level shares the base forecast $\hat{Y}_t$, so in score space `MultiQT` runs on the scores (S_t) with the zero base forecast at every level, quantile levels $1-\alpha_i$, and initialization $\tilde{\boldsymbol{x}}_1 = \mathbf{0}$. We write $d_{\mathcal{A}} := \min_i \min\{\alpha_i, 1-\alpha_i\}$, which equals α_1 whenever $\alpha_1 + \alpha_d \le 1$.

What the combination adds. The ingredients of Section 3 are known, and Proposition 3.3 and Theorem 3.4 are short combinations of them (see the discussion after Theorem 3.4). The combination nonetheless changes what can be proved.

1. *Black-box generality.* The coverage proof of Ding et al. [14] is not a regret argument: it relies on the OGD update and on the restorativity of the aggregated pinball loss, which they derive from bounded residuals [14, Lemma 1], so it does not apply to other learners. The duality argument of Liu et al. [25] is scalar and unconstrained. Proposition 3.3 gives the played, projected sequence the regret bound of the base learner against every $\boldsymbol{u} \in \mathbb{R}^d$, which is exactly what the duality argument needs, so any learner with a linearized-regret bound can be used.
2. *Naive combinations fail.* Running independent single-level learners and sorting or projecting their outputs post hoc can break coverage, and so can (non-lazy) projected OGD [14, Props. 2 and 4]. Algorithm 1 avoids both failure modes because the projection is lazy and the subgradient is evaluated at the played iterate.
3. *Two steps behind $\sqrt{d} \to \sqrt{\log d}$.* The geometry of the coverage-error region is a design choice ($q = 2\log d$ in Corollary 3.6), and for coordinate-wise `UP-OCP`, a single aggregate budget B_T splits into d simultaneous per-level bounds at a $\log d$ price (Theorem 3.7).
4. *Continuous levels.* Regret preservation transfers verbatim to a Hilbert space (Appendix A.8), which gives coverage guarantees for a continuum of levels, an extension that Ding et al. [14] leave to future work.
5. *Sharpness.* Theorem 4.1 explains the empirical advantage of `MultiQT` over naive forecasters reported by Ding et al. [14].

The two regimes of `MultiQT`. Under Assumption 2.1 with $\gamma = 0$ (residual bound D) and a constant step size η, the general guarantee of Ding et al. [14, Thm. 1] reads

$$\max_{i\in[d]} \mathrm{MisCov}_{T,i} \le \sqrt{\frac{d}{T} + \frac{2Dd^2}{\eta\, d_{\mathcal{A}} T}},$$

and, since all levels share the same point forecast, their fast rate for point forecasts [14, Lemma 3 and Cor. 1] applies and reads

$$\max_{i\in[d]} \mathrm{MisCov}_{T,i} \le \frac{\sqrt{d}}{T} + \frac{d^{3/2}}{2d_{\mathcal{A}} T} + \frac{Dd^{3/2}}{d_{\mathcal{A}}\eta T} + \frac{d^{3/2}}{\sqrt{3}\, T},$$

an $\mathcal{O}(1/T)$ rate for fixed η, d, $d_{\mathcal{A}}$, and D (both bounds hold for the ℓ_2 norm of the coverage-error vector, hence for its largest coordinate). Our analysis does not recover this fast rate. Three qualifications put it in context.

- *Coupling with regret.* The regret bound of Ding et al. [14, Thm. 3] reads $D^2 d/(\eta T) + \eta d/2$ for the average regret. For a fixed η it does not vanish, i.e., the cumulative regret bound grows linearly in T, and with it the control of the WIS, hence of the width, is lost. The same holds for any step size that does not decay with T, including the balanced one of Corollary 3.6 when $\gamma = 0$. Balancing calibration against regret, Ding et al. [14, Sec. 4.3] select $\eta = \mathcal{O}(T^{-1/2})$ and obtain $\mathcal{O}(T^{-1/2})$ for both; coordinate-wise `UP-OCP` attains the same rate in T up to a $\sqrt{\log T}$ factor together with sublinear regret, without any tuning (Theorem 3.7).
- *Dependence on the levels.* The constants scale as $d^{3/2}/d_{\mathcal{A}}$. At our configuration ($d = 100$ levels down to $\alpha_1 = 10^{-3}$, $T = 3{,}000$), the second term alone equals $10^3/6 \approx 167$, so the bound is vacuous, whereas Corollary 3.6 with $q = 2\log d$ and the balanced η gives $\sqrt{2e(2\log d - 1)/T} \approx 0.12$.
- *Score growth.* Under Assumption 2.1 with $\gamma > 0$, the residual bound becomes DT^{γ}, so for an η fixed independently of T the dominant terms degrade to $\mathcal{O}(T^{\gamma-1})$: slower than $T^{-1/2}$ for $\gamma > 1/2$, and no longer vanishing for $\gamma \geq 1$. Scaling $\eta \propto T^{\gamma}$ restores the rates but requires knowing D and γ, as does the balanced step size of Corollary 3.6; coordinate-wise `UP-OCP` needs no such knowledge.

Our improvements are therefore in the dependence on d, parameter-free operation under polynomially growing scores, and black-box generality, not in the dependence on T.

B Deferred proofs for Section 4

This appendix collects the proofs deferred from Section 4: the WIS–pinball identity (11) and the per-round structural reduction (Theorem 4.1).

B.1 Pinball loss as a Weighted Interval Score (Equation (11))

We verify the identity (11). The WIS is a function of the interval endpoints; a zero radius, which encodes the empty set, has width (Lebesgue measure) 0 and is evaluated at the degenerate endpoints $l = u = \hat{Y}_t$. For symmetric central intervals $[\hat{Y}_t - x_{t,i}, \hat{Y}_t + x_{t,i}]$, $x_{t,i} \geq 0$, 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 $\mathrm{IS}_{\alpha_i}([l_{\alpha_i}, u_{\alpha_i}], y) \coloneqq (u_{\alpha_i} - l_{\alpha_i}) + \frac{2}{\alpha_i}\,\mathrm{dist}\,(y, [l_{\alpha_i}, u_{\alpha_i}])$.

With $D \coloneqq Y_t - \hat{Y}_t$ 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 $x_{t,i} \geq 0$ (all played and hidden radii are valid by the convention of Section 3), 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 $x \gtrless S$) gives, after summation,

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

that is, $\mathcal{L}_t(\boldsymbol{x}_t) + S_t \sum_{i=1}^d \alpha_i = \sum_{i=1}^d \alpha_i x_{t,i} + \sum_{i=1}^d \max(S_t - x_{t,i}, 0)$, which, combined with the previous display, is (11).

B.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 $\tilde{\boldsymbol{x}} \in \mathbb{R}^d$ (the lemmas are algebraic and need no sign condition), let $\boldsymbol{x} \coloneqq \Pi_{\mathcal{K}}(\tilde{\boldsymbol{x}})$ with PAVA blocks $B_1, \dots, B_p$ (each $B_k = \{l_k, l_k + 1, \dots, r_k\}$), and let $\overline{\tilde{x}}_{B_k} \coloneqq |B_k|^{-1} \sum_{i \in B_k} \tilde{x}_i$ denote the block average, so $x_j = \overline{\tilde{x}}_{B_k}$ for all $j \in B_k$.

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

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

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,

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

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

Lemma B.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*

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

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 $c_i \coloneqq \tilde{x}_i - \overline{\tilde{x}}_{B_k}$, so $\sum_{i=l}^r c_i = 0$ and, by Lemma B.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 $\sum_{i=l}^r w_i \tilde{x}_i \ge \overline{\tilde{x}}_{B_k} \sum_{i=l}^r w_i = \sum_{i=l}^r w_i x_i$. 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 B.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 $\boldsymbol{x}_t = \Pi_{\mathcal{K}}(\tilde{\boldsymbol{x}}_t) \in \mathcal{K} \cap \mathbb{R}^d_+$ satisfy the per-round sharpness and*

coverage reductions

$$\sum_{i=1}^{d} \alpha_i x_{t,i} \leq \sum_{i=1}^{d} \alpha_i \tilde{x}_{t,i}, \tag{21}$$

$$\sum_{i=1}^{d} \max(S_t - x_{t,i}, 0) \leq \sum_{i=1}^{d} \max(S_t - \tilde{x}_{t,i}, 0), \tag{22}$$

and consequently the per-round loss reduction

$$\mathcal{L}_t(\boldsymbol{x}_t) \;\leq\; \mathcal{L}_t(\tilde{\boldsymbol{x}}_t). \tag{23}$$

Both (21) *and* (23) *are strict whenever* $\tilde{\boldsymbol{x}}_t \notin \mathcal{K}$.

Proof of Theorem 4.1. By the WIS decomposition (11), the loss reduction (23) is implied by the sharpness reduction (21) and the coverage reduction (22), 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 B.2 with $w_i = \alpha_i$ then gives $\sum_{i \in B_k} \alpha_i x_{t,i} \leq \sum_{i \in B_k} \alpha_i \tilde{x}_{t,i}$, 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 $|B_k| \geq 2$ has non-constant entries. Summing over k gives (21), with strict inequality whenever $\tilde{\boldsymbol{x}}_t \notin \mathcal{K}$ (equivalently, whenever some block has size ≥ 2).

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

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

Summing over k gives (22).

Loss. Summing (21) and (22) and applying (11) gives (23), which is strict whenever (21) is strict, i.e., whenever $\tilde{\boldsymbol{x}}_t \notin \mathcal{K}$. □

C Experimental details and additional results

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

C.1 Datasets and forecasting pipelines

Synthetic streams. All synthetic scores are non-negative trajectories, 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.
- *Stationary and mixed regimes*: a constant baseline or a quadratic trend overlaid with rare exponential spikes, smoothed by a 25-step rolling maximum (Appendices C.7.6 and C.7.7). Stresses heavy tails.

Real-world streams.

- *AAPL, AMZN, GOOGL*: daily opening prices of the three stocks, from Nguyen [26], covering the period 2006-01-03 to 2017-12-29 ($T = 3{,}020$ rounds: the 3,019 trading days of each stock plus one date of the panel missing for them, filled by interpolation). Forecasting model: Prophet [38] 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 [20], restricted to a non-stationary regime; one-step-ahead forecasts produced by an autoregressive model.

C.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 $\tilde{\boldsymbol{x}}_t \in \mathbb{R}^d_+$ 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. Every base learner, including `MultiQT`, outputs valid radii following Section 3: its raw radius $z_{t,j}$ is replaced by $\tilde{x}_{t,j} = \max(z_{t,j}, 0)$, where 0 encodes the empty set, while its internal state and update rule are left unchanged [25, App. F.2].

C.2.1 Coordinate-wise Universal Portfolio (`UP-OCP`)

Algorithm. Coordinate-wise `UP-OCP` runs an independent scalar UP-OCP instance [25] at each miscoverage level α_j and synchronizes only at prediction time through the projection. At level α_j, the scalar instance treats the betting game as a two-asset *conformal market* [25, Def. 4.1] with returns $w^{(j)}_{t,1} = -g_{t,j}/\alpha_j + 1$ and $w^{(j)}_{t,2} = 1 + g_{t,j}/(1-\alpha_j)$, and chooses portfolio weight $\lambda^{(j)}_t$ 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

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

where $N^{(j)}_{t-1}$ counts past per-level miscoverage events. The hidden iterate is the wealth-to-radius mapping $\tilde{x}_{t,j} = \max\left(W^{(j)}_{t-1}\,(\lambda^{(j)}_t - \alpha_j)/(\alpha_j(1-\alpha_j)), 0\right)$, clipped as in the scalar `UP-OCP` [25, App. F.2], and the wealth recursion $W^{(j)}_t = W^{(j)}_{t-1} \cdot (\lambda^{(j)}_t w^{(j)}_{t,1} + (1-\lambda^{(j)}_t) w^{(j)}_{t,2})$ 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 3.7.

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 8 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.

C.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 [27, 36], which yields the explicit closed form of the raw radius (clipped as $\tilde{x}_{t,j} = \max(z_{t,j}, 0)$, Section 3)

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

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 3.6 interpolates between OGD ($p = 2$) and ℓ_∞-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 ℓ_2 regret. The case $p = 2$ recovers OGD with lazy isotonic projection, i.e., `MultiQT` with a constant step size, and is the closest direct comparison with it.

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

Input: miscoverage levels $\alpha_1 < \cdots < \alpha_d$
Initialize $W_0^{(j)} \leftarrow 1$ and miscoverage count $N^{(j)} \leftarrow 0$ for each $j \in [d]$
for $t = 1, 2, \ldots$ **do**
  for $j = 1, \ldots, d$ (in parallel) **do**
    $\lambda_t^{(j)} \leftarrow (N^{(j)} + 1/2)/t$
    $\tilde{x}_{t,j} \leftarrow \max\left(W_{t-1}^{(j)} \cdot (\lambda_t^{(j)} - \alpha_j)/(\alpha_j(1-\alpha_j)),\, 0\right)$ (0 *encodes* $\emptyset$)
  end for
  $\boldsymbol{x}_t \leftarrow \Pi_{\mathcal{K}}(\tilde{\boldsymbol{x}}_t)$ (*unprojected:* $\boldsymbol{x}_t \leftarrow \tilde{\boldsymbol{x}}_t$); observe S_t
  for $j = 1, \ldots, d$ (in parallel) **do**
    if $\hat{C}_t(\hat{Y}_t, x_{t,j})$ does not cover Y_t, i.e., $S_t > x_{t,j}$ or $x_{t,j} = 0$ **then**
      $N^{(j)} \leftarrow N^{(j)} + 1$; $W_t^{(j)} \leftarrow W_{t-1}^{(j)} \cdot \lambda_t^{(j)}/\alpha_j$
    else
      $W_t^{(j)} \leftarrow W_{t-1}^{(j)} \cdot (1 - \lambda_t^{(j)})/(1 - \alpha_j)$
    end if
  end for
end for

Hyperparameters and tuning. The constant step size η is the only tuned knob. Following standard practice for adaptive conformal baselines [14, 29], 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 η 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 η 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 4 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.

Input: $p \in (1, 2]$ (with conjugate exponent $q := p/(p-1)$), step size $\eta > 0$. Set $\boldsymbol{\theta}_1 = \mathbf{0}$
for $t = 1, 2, \ldots$ **do**
  $\tilde{x}_{t,j} \leftarrow \max\left(\, \|\boldsymbol{\theta}_t\|_q^{2-q} \operatorname{sign}(\theta_{t,j}) \cdot |\theta_{t,j}|^{q-1},\, 0\right)$ for $j = 1, \ldots, d$ (0 *encodes* $\emptyset$)
  $\boldsymbol{x}_t \leftarrow \Pi_{\mathcal{K}}(\tilde{\boldsymbol{x}}_t)$ (*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)
  $\boldsymbol{\theta}_{t+1} \leftarrow \boldsymbol{\theta}_t - \eta\, \boldsymbol{g}_t$
end for

C.2.3 Coordinate-wise Krichevsky–Trofimov bettor (KT)

Algorithm. Coordinate-wise KT [25, 29] runs an independent KT bettor at each level α_j, treating the per-level subgradients as the outcomes of a coin-betting game. At level α_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 [29]; the hidden iterate is $\tilde{x}_{t,j} = \max(\beta_t^{(j)} W_{t-1}^{(j)}, 0)$ 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 α_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 8; 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: **for** $j = 1, \ldots, d$ (in parallel) **do**
5: betting fraction $\beta_t^{(j)} \leftarrow -\Sigma^{(j)}/t$
6: $\tilde{x}_{t,j} \leftarrow \max\left(\beta_t^{(j)} \cdot W_{t-1}^{(j)},\, 0\right)$ (0 *encodes* $\emptyset$)
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; $\boldsymbol{g}_t \in \partial \mathcal{L}_t(\boldsymbol{x}_t)$ via (2)
9: **for** $j = 1, \ldots, d$ (in parallel) **do**
10: $W_t^{(j)} \leftarrow W_{t-1}^{(j)} \cdot \left(1 - \beta_t^{(j)}\, g_{t,j}\right)$; $\Sigma^{(j)} \leftarrow \Sigma^{(j)} + g_{t,j}$
11: **end for**
12: **end for**

C.2.4 `MultiQT`

Algorithm. `MultiQT` [14] 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. [14, 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. [14]. 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. [14] 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. [14] on their own benchmarks rather than on ours.

C.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 3 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.

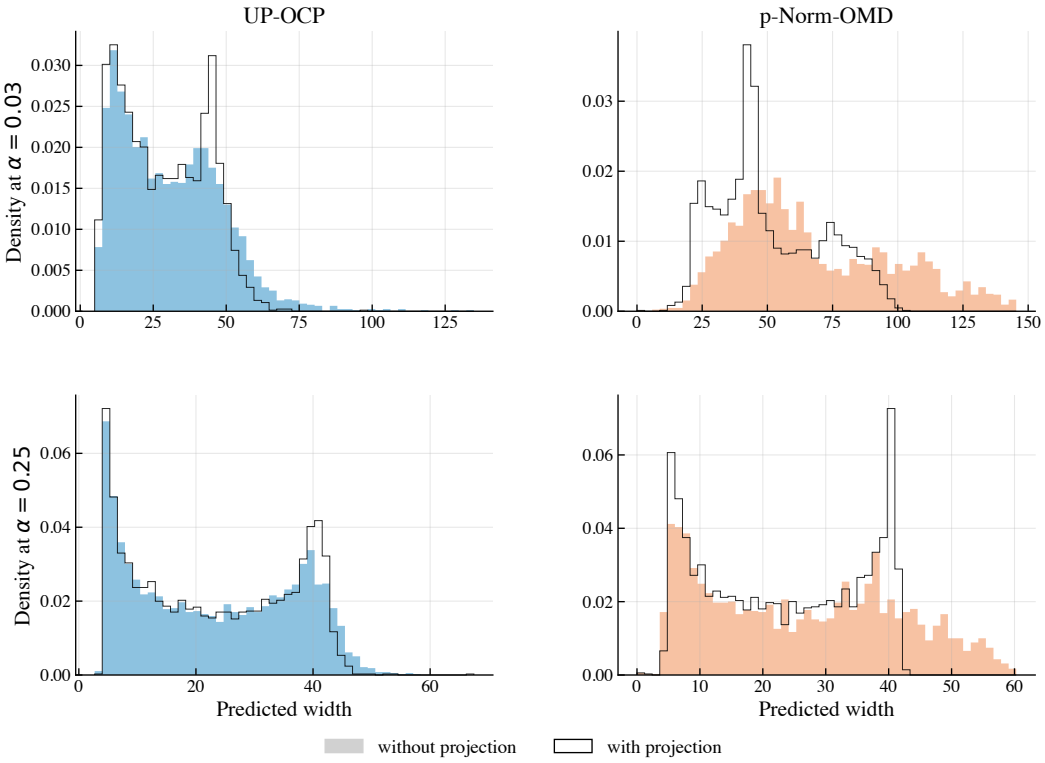


Figure 3: Distribution of per-round prediction-set widths on the sinusoidal stream, at $\alpha = 0.03$ (top row) and $\alpha = 0.25$ (bottom row). Each panel overlays the histogram of widths for the unprojected base learner (filled, solid color) against the projected variant (outlined, black border), at the same α 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).

C.4 Weighted Interval Score

Table 2 extends Figure 2 to every dataset. The projection lowers the WIS of `UP-OCP` and p-`norm OMD` on every dataset; for `KT`, whose isotonicity violations are small (Table 3), the two variants are within 0.3% of each other. These comparisons are across two different runs, whose trajectories differ, so they are not covered by Theorem 4.1. Table 3 checks the theorem in its exact, within-run form: the hidden iterate violates isotonicity on $127{,}206$ of the $127{,}680$ post-burn-in rounds (99.6%) across both datasets and the four projected methods, and on every one of them the projection strictly reduced both the WIS and its sharpness component without increasing the coverage penalty.

Table 2: Average per-round WIS over all $d = 100$ levels (post burn-in; synthetic streams: mean over 10 seeds), and its sharpness and coverage-penalty components. **Bold** marks the smallest WIS on each dataset.

Dataset		UP-OCP	UP-OCP-ISO	p-Norm-OMD	p-Norm-OMD-ISO	KT	KT-ISO	MULTIQT
AAPL	WIS	456.4	**453.6**	523.6	466.5	528.9	528.3	478.3
	Sharpness	394.5	392.4	451.7	411.0	394.1	393.4	405.0
	Penalty	61.8	61.2	71.9	55.5	134.8	134.9	73.3
AMZN	WIS	1746.6	**1734.3**	2026.0	1798.3	2011.7	2010.5	1859.3
	Sharpness	1417.4	1412.4	1689.6	1520.6	1483.0	1482.0	1496.5
	Penalty	329.2	321.9	336.4	277.7	528.7	528.5	362.8
GOOGL	WIS	1777.3	1771.6	1944.6	**1756.4**	2063.4	2064.9	1915.3
	Sharpness	1393.7	1387.4	1596.8	1473.6	1450.8	1451.5	1463.3
	Penalty	383.6	384.2	347.8	282.9	612.7	613.4	451.9
elec2	WIS	2.04	2.03	2.12	2.06	1.98	**1.98**	2.38
	Sharpness	1.06	1.05	1.30	1.25	1.01	1.00	1.63
	Penalty	0.97	0.97	0.82	0.81	0.97	0.98	0.75
Sinusoidal	WIS	609.2	**602.1**	709.9	621.4	643.2	641.6	606.6
	Sharpness	569.0	564.1	635.2	577.4	576.7	575.7	583.3
	Penalty	40.2	38.0	74.7	44.0	66.6	65.9	23.3
Stationary	WIS	11863.2	**11753.0**	12551.5	11810.6	14208.3	14202.9	13238.4
	Sharpness	8770.9	8671.8	9867.2	9267.3	9128.4	9114.4	9669.0
	Penalty	3092.3	3081.2	2684.3	2543.3	5079.9	5088.4	3569.3
Mixed	WIS	8271.5	**8232.0**	9378.0	8484.8	10008.3	10010.9	9212.5
	Sharpness	6006.8	5961.4	7271.2	6490.6	6150.6	6149.1	6637.6
	Penalty	2264.7	2270.6	2106.8	1994.1	3857.7	3861.8	2574.8

Table 3: Within-run check of Theorem 4.1. For each projected method, the hidden iterate $\tilde{\boldsymbol{x}}_t$ and the played iterate $\boldsymbol{x}_t = \Pi_{\mathcal{K}}(\tilde{\boldsymbol{x}}_t)$ are logged at every post-burn-in round (sinusoidal: all 10 seeds). Columns: share of rounds with $\tilde{\boldsymbol{x}}_t \notin \mathcal{K}$; share of those rounds on which the WIS and its sharpness component are strictly reduced, and the coverage penalty is not increased (up to 10^{-12} relative rounding); median and maximum per-round relative WIS reduction $1 - \mathrm{WIS}(\boldsymbol{x}_t)/\,\mathrm{WIS}(\tilde{\boldsymbol{x}}_t)$; average WIS of the hidden and of the played iterates.

Dataset	Method	$\tilde{\boldsymbol{x}}_t \notin \mathcal{K}$	WIS, sharpness $\downarrow$, penalty γ	median red.	max red.	avg. WIS hidden	avg. WIS played
AAPL	UP-OCP-ISO	99.97%	100%	0.93%	39.9%	467.0	453.6
	p-Norm-OMD-ISO	99.01%	100%	22.43%	80.1%	584.9	466.5
	KT-ISO	99.93%	100%	0.03%	24.4%	529.1	528.3
	MULTIQT	99.86%	100%	0.28%	48.0%	482.2	478.3
Sinusoidal	UP-OCP-ISO	99.97%	100%	4.51%	50.6%	644.9	602.1
	p-Norm-OMD-ISO	99.10%	100%	23.45%	64.6%	828.4	621.4
	KT-ISO	99.76%	100%	0.17%	27.1%	652.3	641.6
	MULTIQT	99.66%	100%	1.74%	31.6%	619.7	606.6

C.5 Step-size sensitivity and adaptive step sizes for p-norm OMD

Corollary 3.6 favors $\eta \to \infty$, whereas $F_T(\boldsymbol{u})$ diverges as $\eta \to \infty$ for every $\boldsymbol{u}$. The two measure different things: small miscoverage alone is a weak target, attained even by data-free predictors [25, App. H.5], while the regret bound controls efficiency, since it bounds half the WIS (Section 4) of the played sets against every comparator. Our step-size sweep confirms the trade-off(Appendices A.4 and C.5).

Sensitivity. Table 4 reports, for projected p-norm OMD with $p = 2\log d/(2\log d - 1)$, the coverage error (mean over all $d = 100$ levels of $\left|\mathrm{Coverage}_{T,i} - (1-\alpha_i)\right|$, post burn-in) and the average per-round WIS over the whole step-size grid. The coverage error decreases monotonically in η, as Corollary 3.6 suggests, while the WIS is minimized at a finite step size ($\eta = 64$ on AAPL and $\eta = 32$ on the sinusoidal stream) and roughly doubles between $\eta = 64$ and $\eta = 1024$. Every $\eta \in [16, 128]$ stays within about 10% of the best WIS on both datasets, and the same $\eta = 64$ is within about 2.2%

of the best WIS on both streams, which are of very different nature (stock residuals and a synthetic wave).

Table 4: Projected p-norm OMD: coverage error (%, mean over all levels) / average per-round WIS, for each constant step size η of the grid and for the plug-in step size (24).

η	0.25	0.5	1	2	4	8	16	32	64	128	256	512	1024	plug-in
AAPL	36.86	28.16	21.01	14.88	9.52	5.46	2.80	1.38	0.64	0.31	0.14	0.06	0.05	3.31
	1075.8	969.8	853.5	736.0	639.1	559.7	500.0	468.3	**466.5**	500.0	568.8	706.6	934.9	482.9
Sinusoidal	48.12	34.29	20.27	10.49	5.40	2.64	1.26	0.61	0.30	0.15	0.08	0.04	0.02	0.29
	1467.1	1234.4	1060.6	940.9	813.9	688.3	625.5	**608.0**	621.4	669.9	767.6	930.9	1231.4	629.3

A plug-in step size. The balanced step size of Section 3.1, $\eta^\star = 2(p-1)B_T/(d^{2/q}T)$, depends on the unknown average budget B_T/T. Replacing it by its running estimate gives the step size used for the update after round $t \geq 2$ (in our experiments, updates start after the warm-up),

$$\eta_t = \frac{2(p-1)}{d^{2/q}} \cdot \frac{1}{t-1} \sum_{\tau=1}^{t-1} \mathcal{L}_\tau(\mathbf{0}), \qquad \mathcal{L}_\tau(\mathbf{0}) = S_\tau \sum_{i=1}^{d} (1-\alpha_i), \tag{24}$$

which uses only revealed scores. Its WIS is about 3.5% above the grid optimum on both datasets (Table 4); its coverage error on the sinusoidal stream (0.29%) matches that of the grid-selected $\eta = 64$ (0.30%), while on AAPL it shows a mild under-coverage of 3.3%, comparable to the parameter-free methods. This schedule is only a heuristic: (η_t) need not be monotone, so the standard analysis of OMD with non-increasing step sizes does not apply, and we claim no bound for it.

A doubling-trick variant with guarantees. Run Algorithm 3 in epochs $k = 0, 1, \ldots,$ restarting the base learner at the beginning of each epoch (the projection layer is stateless), with the constant step size $\eta_k = 2(p-1)b_k/d^{2/q}$ and $b_k = 2^k b_0$, where $b_0 = s_0 \sum_{i=1}^{d}(1-\alpha_i)$ for some $s_0 > 0$. Epoch k ends after the first round at which the average of $\mathcal{L}_\tau(\mathbf{0})$ over the rounds of the epoch exceeds $2b_k$. On the first $n_k - 1$ of its n_k rounds, the epoch's budget is at most $2b_k(n_k - 1)$, so Theorem 3.4 applied to them, as in Corollary 3.6 with $q = 2\log d$, gives $\left\|\boldsymbol{G}^{(k)}\right\|_q \leq \sqrt{3e(2\log d - 1)n_k} + \sqrt{e}$, where the last term accounts for the last round. Summing over the K epochs and applying Cauchy–Schwarz,

$$\max_{i\in[d]} \mathrm{MisCov}_{T,i} \leq \sqrt{\frac{3e(2\log d - 1)K}{T}} + \frac{\sqrt{e}\,K}{T}.$$

Since $\mathcal{L}_\tau(\mathbf{0}) = S_\tau \sum_i (1-\alpha_i) \leq DT^\gamma \sum_i (1-\alpha_i)$ under Assumption 2.1, no epoch ends once $2b_k \geq DT^\gamma \sum_i (1-\alpha_i)$, hence $K \leq 1 + \log_2(DT^\gamma/s_0) = \mathcal{O}(\gamma \log T + \log(D/s_0))$, independently of d [cf. 11, Sec. 2.3]. For bounded scores ($\gamma = 0$), this is the balanced rate up to a factor $\sqrt{3K/2}$ that depends only on D/s_0; under polynomial growth it pays a $\sqrt{\log T}$ factor, multiplicatively, whereas Theorem 3.7 pays it additively. In both cases the step size stays bounded, so the regret bound does not diverge.

C.6 Varying the number of levels

Table 5 reports the realized worst-level miscoverage error $\max_i \mathrm{MisCov}_{T,i}$ on the sinusoidal stream ($T = 3{,}000$, levels uniformly spaced in $[0.001, 0.5]$, 10 seeds; unlike Section 5, these runs use no warm-up period, so the $d = 100$ column differs slightly from the main experiments) as the number of levels grows from $d = 5$ to $d = 1{,}000$, for projected p-norm OMD with the balanced step size of Section 3.1, projected coordinate-wise `UP-OCP`, and `MultiQT`. The errors of `UP-OCP-ISO` and `MultiQT` are flat in d, and that of p-norm OMD grows slowly while staying an order of magnitude below that of `MultiQT`. The $\sqrt{d}$ and $\sqrt{\log d}$ rates are worst-case guarantees, and on this non-adversarial stream the ablation exhibits neither scaling; we do not claim that it does. What it shows is that the methods designed through the tighter analysis are more accurate at every d: p-norm OMD by an order of magnitude over `MultiQT`, and `UP-OCP` by about 40%, with no degradation over more than two orders of magnitude in d. Small coverage error is only half of the picture, since it can be bought with width: the mean widths of `UP-OCP-ISO` and `MultiQT` are within 1% of each other

for $d \geq 100$, while p-norm OMD with the balanced step size pays 23–36% more mean width than `MultiQT` for its smaller error.

Table 5: Worst-level miscoverage error $\max_i \mathrm{MisCov}_{T,i}$ on the sinusoidal stream as a function of the number of levels d (mean over 10 seeds; standard errors are at most 1.5×10^{-4}).

d	5	10	20	50	100	200	500	1000
p-Norm-OMD-Iso	0.0010	0.0017	0.0020	0.0030	0.0034	0.0040	0.0047	0.0053
UP-OCP-Iso	0.0377	0.0379	0.0379	0.0380	0.0380	0.0380	0.0378	0.0381
`MultiQT`	0.0635	0.0639	0.0640	0.0643	0.0644	0.0645	0.0646	0.0648

C.7 Additional results

This subsection collects, for every dataset, 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 2×2 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 3.6), the constant step size η is selected by per-dataset grid search over $\eta \in \{2^{-2}, 2^{-1}, 1, 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024\}$, exactly as in Section 5. The per-dataset selected η is reported in each subsubsection.

Table 6: 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

C.7.1 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 4). Tuned $\eta = 64$ for p-norm OMD with p optimal. Table 6 reports coverage and widths at three target levels.

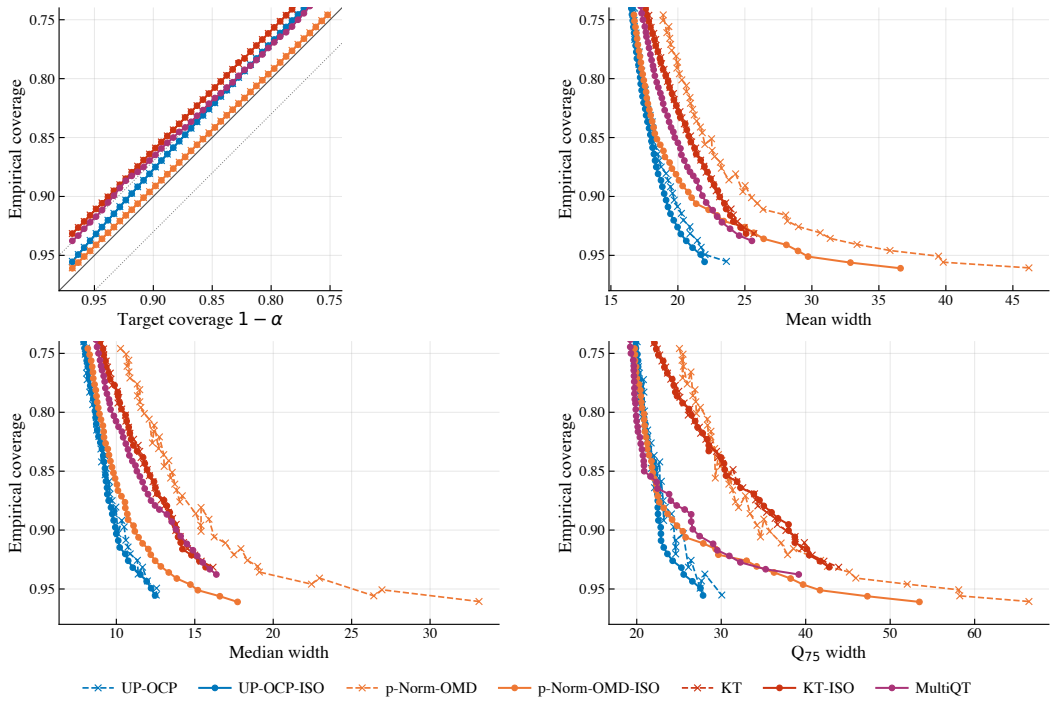


Figure 4: AAPL: calibration (top-left) and three width Pareto frontiers (mean, median, Q_{75}) versus empirical coverage. KT and KT-ISO are now included, exposing their wider mean and median widths but competitive Q_{75} tail.

Table 7: 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

C.7.2 AMZN

Same Prophet pipeline as AAPL (Taylor and Letham [38], log-prices, daily refit). Tuned $\eta = 256$ for p-`norm OMD` with p optimal.

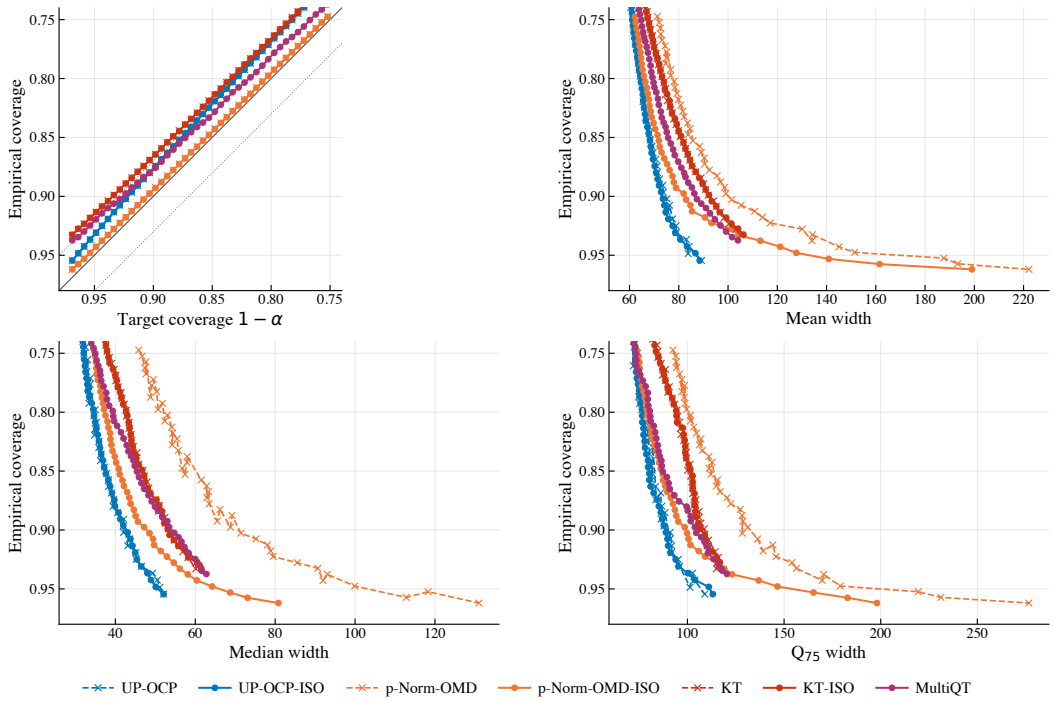


Figure 5: AMZN: calibration and three width Pareto frontiers. Behavior is qualitatively identical to AAPL.

Table 8: 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**

C.7.3 GOOGL

Same Prophet pipeline as AAPL. Tuned $\eta = 256$ for p-`norm OMD` with p optimal. The qualitative pattern matches AAPL/AMZN; the projection improves both tail and mean widths of `UP-OCP` and p-`norm OMD` without affecting calibration, while `KT-ISO` is marginally wider than `KT` in several cells.

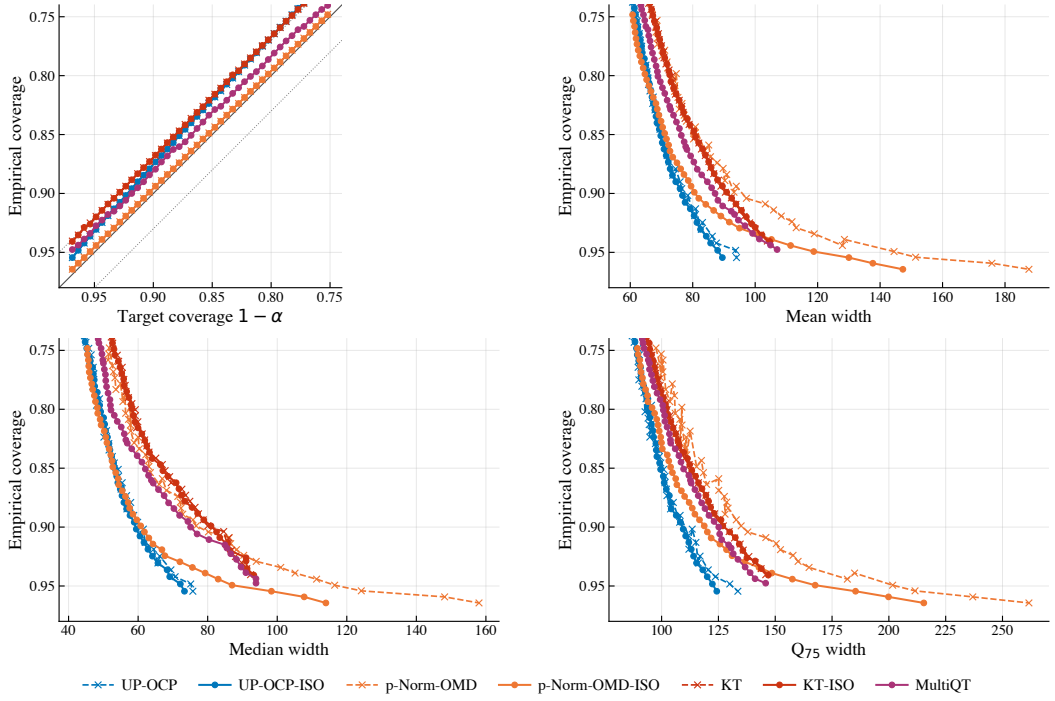


Figure 6: GOOGL: calibration and three width Pareto frontiers.

Table 9: 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

C.7.4 Electricity demand (`elec2`)

Half-hourly New South Wales electricity demand from [20], 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.

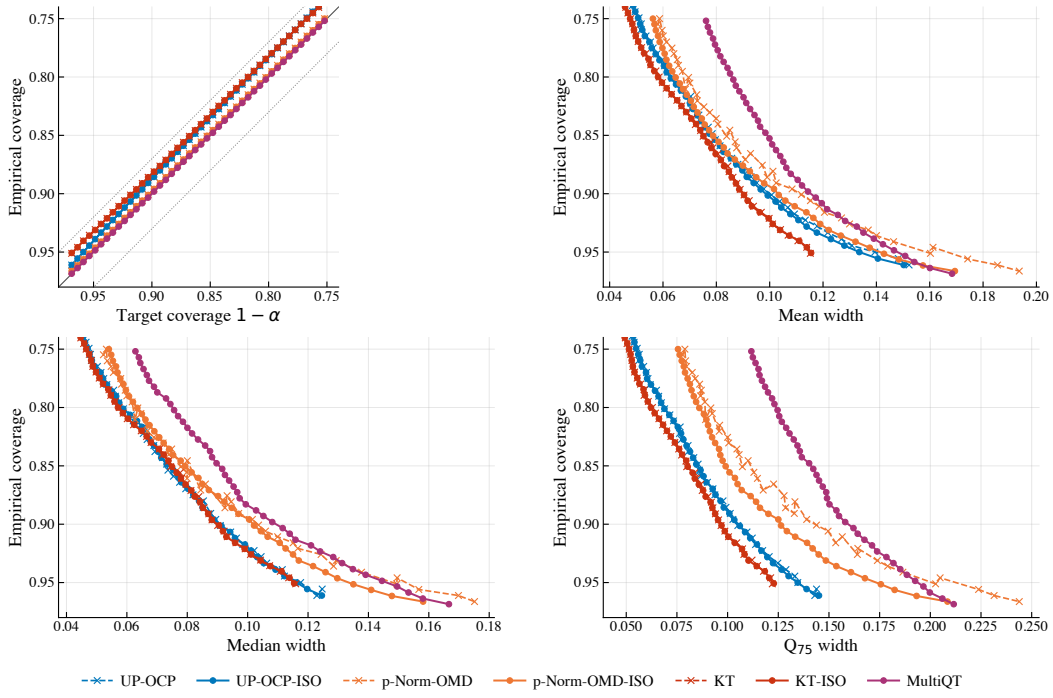


Figure 7: `elec2`: calibration and three width Pareto frontiers.

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` is within about 2% in median width but about 8% wider in mean and up to 22% wider in Q_{95} width. 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.

Table 10: 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	$\mathbf{0.95}_{\pm.00}$	$33.21_{\pm.04}$	$35.09_{\pm.05}$	$41.34_{\pm.05}$	$0.91_{\pm.00}$	$28.87_{\pm.04}$	$31.29_{\pm.05}$	$\mathbf{40.89}_{\pm.04}$	$0.83_{\pm.00}$	$24.73_{\pm.02}$	$26.79_{\pm.04}$	$\mathbf{40.52}_{\pm.03}$

C.7.5 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

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

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 8), and Table 10 reports the per-method coverage and width summary at three target levels.

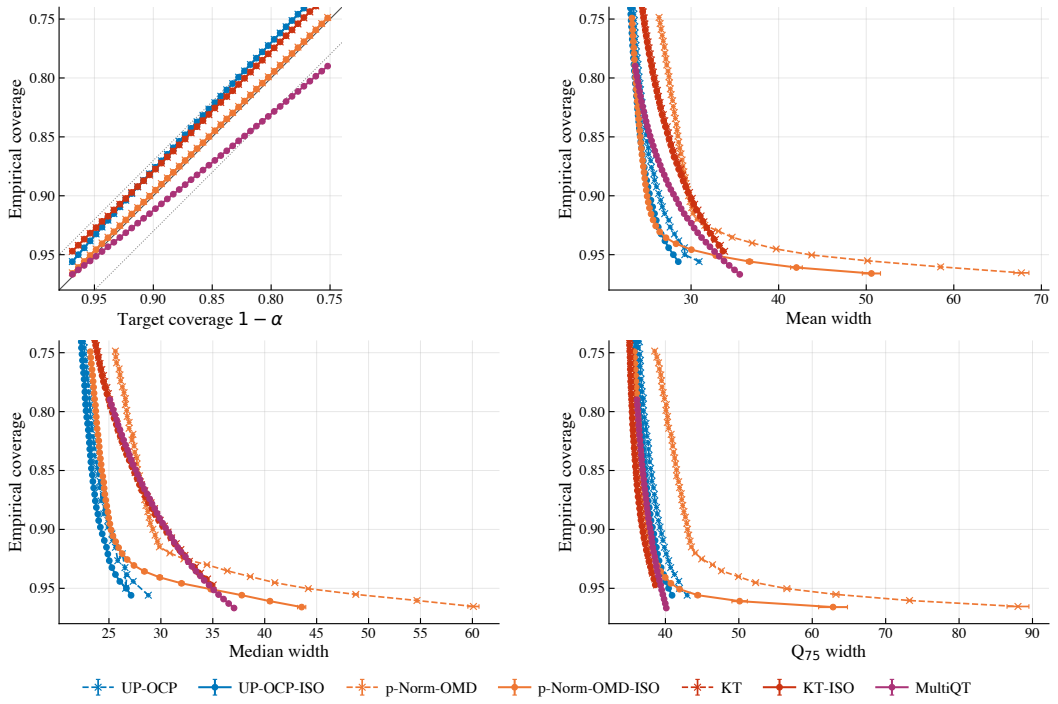


Figure 8: Sinusoidal: calibration and three width Pareto frontiers. `UP-OCP-ISO` dominates in mean and median width across the entire α range, while `KT` and `MultiQT` lead on Q_{75}, mirroring the bounded twin-extrema structure of the wave.

Table 11: 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	$0.93_{\pm.00}$	$\mathbf{483.63}_{\pm7.43}$	$\mathbf{411.50}_{\pm7.77}$	$\mathbf{1059}_{\pm26.10}$	$0.87_{\pm.00}$	$\mathbf{427.91}_{\pm6.05}$	$\mathbf{360.99}_{\pm6.36}$	$967.28_{\pm24.87}$	$0.77_{\pm.00}$	$\mathbf{381.22}_{\pm5.34}$	$\mathbf{317.14}_{\pm5.90}$	$885.35_{\pm22.26}$
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	$\mathbf{0.94}_{\pm.00}$	$745.34_{\pm25.86}$	$569.02_{\pm15.81}$	$2088_{\pm130.20}$	$\mathbf{0.89}_{\pm.00}$	$510.48_{\pm10.38}$	$434.35_{\pm8.26}$	$1191_{\pm48.89}$	$\mathbf{0.80}_{\pm.00}$	$401.20_{\pm5.85}$	$348.20_{\pm5.97}$	$915.13_{\pm25.25}$
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	$0.92_{\pm.00}$	$601.46_{\pm11.81}$	$565.92_{\pm10.37}$	$1067_{\pm38.87}$	$0.87_{\pm.00}$	$523.56_{\pm11.66}$	$488.15_{\pm10.63}$	$\mathbf{952.76}_{\pm34.55}$	$0.78_{\pm.00}$	$435.85_{\pm8.20}$	$398.80_{\pm8.15}$	$\mathbf{817.11}_{\pm21.49}$
MULTIQT	$0.93_{\pm.00}$	$605.17_{\pm13.54}$	$566.08_{\pm14.59}$	$1102_{\pm46.29}$	$0.88_{\pm.00}$	$512.82_{\pm9.02}$	$476.98_{\pm9.11}$	$995.40_{\pm33.23}$	$0.79_{\pm.00}$	$436.87_{\pm6.50}$	$400.76_{\pm7.38}$	$902.39_{\pm23.39}$

C.7.6 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

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

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.

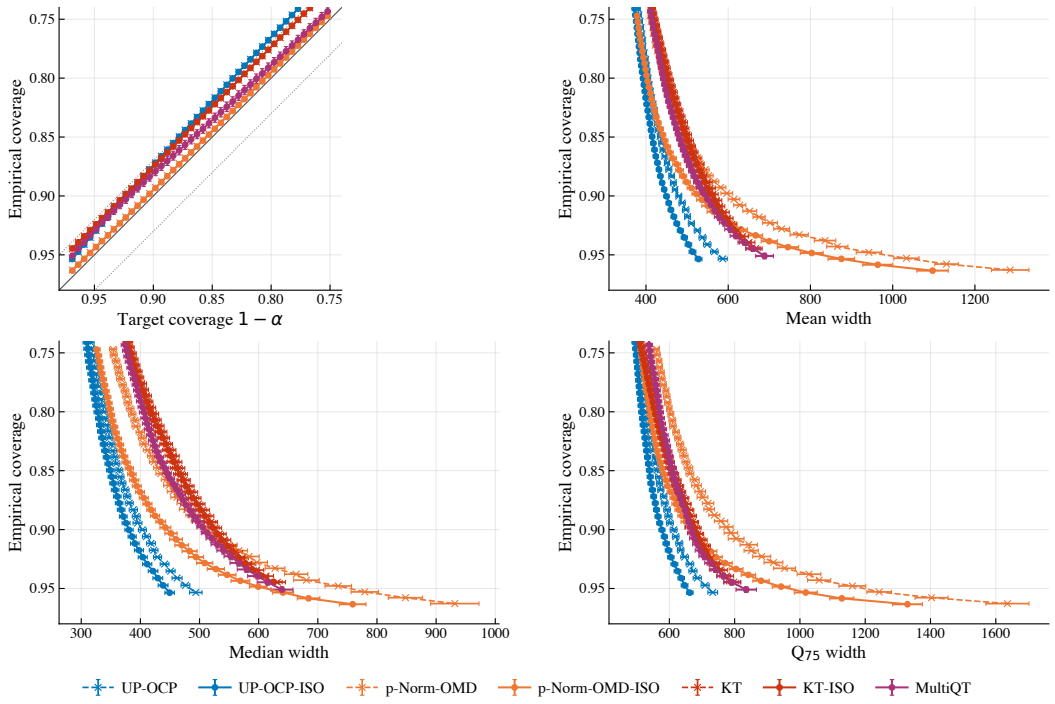


Figure 9: Stationary trend with random waves: calibration and three width Pareto frontiers.

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 and median widths at every target (KT-ISO has the smallest Q_{95} width at $1-\alpha \in \{0.9, 0.8\}$).

Table 12: 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	$0.89_{\pm.00}$	$404.10_{\pm10.88}$	$271.70_{\pm12.32}$	$1234_{\pm43.23}$	$0.85_{\pm.00}$	$350.65_{\pm9.81}$	$217.81_{\pm6.87}$	$1111_{\pm31.78}$	$0.76_{\pm.00}$	$300.23_{\pm7.98}$	$177.17_{\pm3.82}$	$996.18_{\pm34.54}$

C.7.7 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.

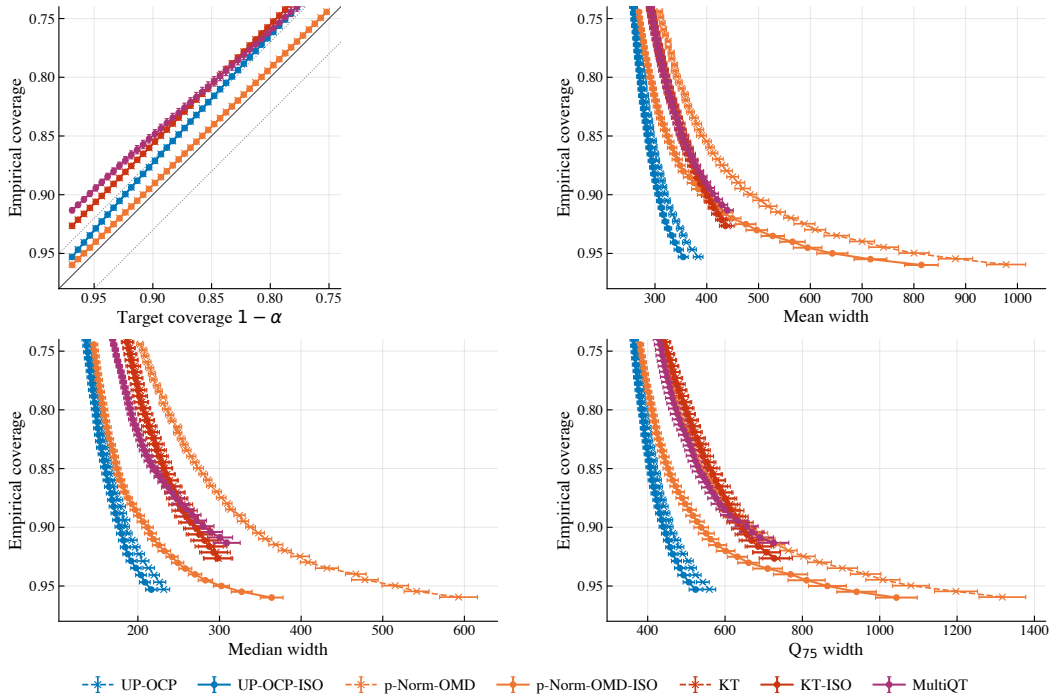


Figure 10: Quadratic drift with random waves: calibration and three width Pareto frontiers.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and introduction (Section 1) state the three contributions—a black-box reduction from linearized regret to multi-level coverage (Theorem 3.4), the $\sqrt{\log d}$ rate via p-norm OMD and coordinate-wise UP (Corollary 3.6 and Theorem 3.7), and the per-round WIS reduction from isotonic projection (Theorem 4.1). Each claim is supported by a theorem (proved in the appendix) or by experiments in Section 5.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: Limitations are flagged where they arise. Section 6 states that our guarantees concern long-run *marginal* coverage only (no conditional or interval-adaptive guarantees) and points to potential extensions via multi-validity meta-algorithms or strongly adaptive online learning. Section 5 flags a slight finite-sample miscoverage of the parameter-free and adaptive instances and notes a known α-correction fix [25]. Appendix A.9 states the regime (common point forecast, bounded scores, fixed step size) in which `MultiQT` attains a faster rate in T than our analysis, and Appendix C.6 notes that the $\sqrt{d}$ versus $\sqrt{\log d}$ scaling is a worst-case statement that our experiments do not exhibit. The analysis throughout relies on the inward-flow geometric condition (Definition 3.1).

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.

- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

   Justification: All theorems state their assumptions inline (notably Assumption 2.1 on polynomial score growth and the inward-flow condition). Short proof sketches appear in the main text where space permits; full proofs are in Appendix A (regret-to-coverage, p-norm and UP rates) and Appendix B (per-round WIS reduction).

   Guidelines:

   - The answer [N/A] means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.
   - The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
   - Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
   - Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: Section 5 and Appendix C list datasets, base forecasters, score definition, number of levels, warm-up length, target window, and grid for tuned methods. Algorithm pseudocode and update rules are given in Appendix C.2. All datasets are public; synthetic seeds are listed.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example

(a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
(b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
(c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
(d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: The code release accompanying the supplementary material reproduces every figure and table in this paper; instructions and exact commands are included in the package. The datasets are publicly available (the S&P 500 stock data of Nguyen [26]; `elec2` for electricity); synthetic streams are generated by the included scripts with the seeds listed in Appendix C.1.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: The protocol is purely online (no train/test split) and is described in Section 5. Hyperparameter grids and selection rules for tuned baselines (p-`norm OMD`, `MultiQT`) are given in Appendix C.2. Parameter-free methods (`UP-OCP`, `KT`) have no tuned hyperparameters.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.

- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: For synthetic streams we report the mean and standard error of the mean over 10 seeds (Table 10 and figures in Appendix C.7). The real-world series (stocks, electricity) are single deterministic sequences and thus do not admit error bars; figures and the WIS tables report means over seeds for synthetic streams, and the summary tables of Appendix C.7 also report standard errors.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.
   - It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
   - For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
   - If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: All experiments are CPU-only; the heaviest base learner (coordinate-wise UP at $d = 100$) runs each $T \approx 3{,}000$ stream in well under a minute on a single laptop core. Total compute for all reported and ablation experiments is on the order of CPU-hours.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

Answer: [Yes]

Justification: The work is theoretical and methodological; it uses only public benchmark data and involves no human subjects, sensitive attributes, or model release.

Guidelines:

- The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [N/A]

Justification: The contribution is foundational research on uncertainty quantification with no direct path to a specific application or deployment; tighter prediction sets with valid coverage are broadly beneficial when used responsibly, and we are not aware of negative societal impacts particular to this work.

Guidelines:

- The answer [N/A] means that there is no societal impact of the work performed.
- If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

Answer: [N/A]

Justification: The paper does not release any data or model with such risks. Code released alongside the paper implements an online-conformal wrapper that calls user-supplied base forecasters and online learners; it carries no high-risk content.

Guidelines:

- The answer [N/A] means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: The S&P 500 stock data [26], the `elec2` electricity dataset [20], and the `Prophet` forecaster [38] are publicly available; their original sources are cited in Appendix C.1. Baseline algorithms (`KT`, `MultiQT`) are also cited where they appear.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
    - If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [Yes]

    Justification: The supplementary code release documents the algorithms (Algorithms 1, 3, 2, 4) and includes a README with usage instructions; no new datasets are introduced.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [N/A]

Justification: The paper does not involve crowdsourcing or human subjects.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [N/A]

Justification: The paper does not involve human subjects.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: LLMs are not used as part of the core methodology of this research.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.