The cost of enforcing long-run coverage: a phase transition in online conformal prediction

Yu Gui[1] Rina Foygel Barber[2] Cong Ma[2]

[1]Department of Statistics and Data Science, The Wharton School, University of Pennsylvania
[2]Department of Statistics, University of Chicago

November 14, 2026

Abstract

Consider an online prediction setting, where at each time, we are given a new feature vector and asked to provide a prediction interval or prediction set that will likely contain the corresponding response. In this setting, if the data points are i.i.d. (or more generally, exchangeable), conformal prediction provides a mechanism for building valid prediction sets around the output of any trained model, offering coverage probability $\geq 1-\alpha$ at each time t. In practice, however, issues such as distribution drift, changepoints, or other challenges may cause violations of the i.i.d. assumption. In this adversarial setting, online conformal prediction methods can still offer guarantees on long-run empirical coverage. However, there is a trade-off between these two types of goals: enforcing tight constraints on long-run empirical coverage to protect against the adversarial setting can lead to unstable oscillations in the instantaneous training-conditional coverage in the i.i.d. setting.

In this work, we seek *best-of-both-worlds* algorithms that aim to minimize cumulative error in instantaneous training-conditional coverage when in the i.i.d. setting, while also offering long-run empirical coverage guarantees in the adversarial setting. Specifically, consider a condition that requires long-run empirical coverage to deviate from the target level $1-\alpha$ by at most $O(T^{-\beta})$, for some $\beta \in [0,1]$. We prove that this family of conditions can be described with a phase transition. For $\beta \in [0, \frac{1}{2}]$, it is possible to construct algorithms satisfying such a constraint in the adversarial setting, while offering the same instantaneous coverage performance as split conformal prediction if the data is in fact i.i.d. On the other hand, for $\beta \in (\frac{1}{2}, 1]$, this type of condition comes at a cost, and does not achieve best-of-both-worlds behavior.

1 Introduction

Consider an online uncertainty quantification problem in which each data point consists of a feature $X_t \in \mathcal{X}$ and a response $Y_t \in \mathcal{Y}$. For each time $t \geq 1$, let C_t denote any prespecified construction that maps from an observed data stream to a prediction set for the data point at time t. Concretely, we write

$$C_t(X_t) = C_t\big(X_t; (X_1, Y_1), \ldots, (X_{t-1}, Y_{t-1})\big) \subseteq \mathcal{Y} \tag{1}$$

to denote the prediction set at time t, when trained on data points $(X_1, Y_1), \ldots, (X_{t-1}, Y_{t-1})$. The current outcome Y_t is not available when the set is constructed. Throughout this paper, we use the notation $C = (C_t)_{t \geq 1}$ to denote the online algorithm producing the sequence of prediction sets $(C_t(X_t))_{t \geq 1}$.

After Y_t is revealed, the performance of the prediction set can be summarized by a binary miscoverage indicator

$$\mathsf{e}_t = \mathbf{1}\{Y_t \notin C_t(X_t; (X_1, Y_1), \ldots, (X_{t-1}, Y_{t-1}))\}.$$

We expect the prediction sets to frequently cover the true response Y_t, quantified by the nominal coverage level $1 - \alpha \in (0, 1)$. To quantify this, we may require that the probability of coverage at each time point is calibrated—that is, that $\mathbb{P}(\mathsf{e}_t \neq 1)$ or the coverage rate conditional on the history is $\approx 1 - \alpha$, at each time t. This type of guarantee can only be achieved with some assumptions on the data stream, however. More generally, we can instead aim for a weaker guarantee, requiring that the empirical *average* coverage rate is $\approx 1 - \alpha$ over a period of time. We formally define and discuss these two guarantees in the following subsections.

1.1 The i.i.d. setting: instantaneous training-conditional coverage

When the data points (X_t, Y_t) are independent and identically distributed (or exchangeable), split conformal prediction uses past scores to calibrate prediction sets. The score is fixed or fitted to independent data, and higher scores indicate greater deviation from the fitted model. Concretely, with a fixed score function $s : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$, the split conformal prediction set can be defined by

$$C_t^{\mathrm{CP}}(x) = \left\{ y \in \mathcal{Y} : \frac{1 + \sum_{r<t} \mathbf{1}\{s(X_r, Y_r) \geq s(x, y)\}}{t} > \alpha \right\}. \tag{2}$$

Exchangeability gives a finite-sample marginal coverage guarantee, i.e., $\mathbb{P}(Y_t \in C_t^{\mathrm{CP}}(X_t)) \geq 1 - \alpha$. A more local requirement asks whether coverage is close to $1 - \alpha$ given the observed past at each time t. For i.i.d. observations and a nonatomic score distribution,[1] the split conformal prediction set in (2) satisfies the following bound (Vovk, 2012; Angelopoulos et al., 2024b):

$$\mathbb{E}_{\mathbb{P}} \left| \mathbb{P}\{Y_t \in C_t^{\mathrm{CP}}(X_t) \mid (X_s, Y_s)_{s<t}\} - (1 - \alpha) \right| \leq \frac{1}{2\sqrt{t+1}} + \frac{1}{t}. \tag{3}$$

Thus, under these assumptions, the instantaneous training-conditional coverage of split conformal prediction, i.e., its coverage at time t given the history $(X_s, Y_s)_{s<t}$, deviates from the target level $1 - \alpha$ by $O(t^{-1/2})$ in expectation. To measure this deviation over T rounds, for any online algorithm C we define

$$\Delta_T^{\mathrm{inst}}(C) := \sum_{t=1}^{T} \mathbb{E}_{\mathbb{P}} \left| \mathbb{P}\{Y_t \in C_t(X_t) \mid (X_s, Y_s)_{s<t}\} - (1 - \alpha) \right|. \tag{4}$$

With $C = C^{\mathrm{CP}}$, summing (3) over t gives $\Delta_T^{\mathrm{inst}} = O(\sqrt{T})$.

[1] A distribution is nonatomic if its probability measure μ satisfies that for any measurable set A with $\mu(A) > 0$, there exists a measurable subset B of A such that $\mu(A) > \mu(B) > 0$.

1.2 The adversarial setting: long-run empirical coverage

When exchangeability is not assumed, one widely adopted quantity that measures the *adversarial* performance of online algorithms is the long-run cumulative coverage deviation (Gibbs and Candès, 2021):

$$\sum_{t=1}^{T} [\mathbf{1}\{Y_t \in C_t(X_t)\} - (1-\alpha)] = \sum_{t=1}^{T} (\alpha - \mathsf{e}_t). \tag{5}$$

Here, we denote $\mathtt{E}_T = \sum_{t=1}^{T}(\mathsf{e}_t - \alpha)$, where $|\mathtt{E}_T|/T$ tracks how far the long-run empirical coverage rate departs from $1-\alpha$. A well-calibrated algorithm keeps $|\mathtt{E}_T|$ small relative to T. This guarantee can be achieved by a large class of algorithms without imposing distributional assumptions on the data sequence $\{(X_t, Y_t), t \geq 1\}$. To quantify the long-run empirical coverage guarantee in the adversarial setting, we say the algorithm satisfies a (β, κ)-*constraint* if

$$|\mathtt{E}_T|/T \leq \kappa T^{-\beta} \quad \text{for all } T \geq 1. \tag{6}$$

Throughout this paper, $\beta \in [0,1]$ and $\kappa \geq 1/2$, so the interval has width at least one and a feasible error outcome exists at every round. Note that if $\beta = 0$ and $\kappa \geq \max\{\alpha, 1-\alpha\}$, then the constraint holds trivially for any sequence $(\mathsf{e}_t)_{t\geq 1}$. For $\beta > 0$, this implies that the long-run empirical coverage rate converges to $1-\alpha$. A larger β requires tighter control over the long-run empirical coverage rate. We use $\mathcal{A}_{\beta,\kappa}$ to collect all algorithms that satisfy (6). [2]

Our central question is how close an algorithm in $\mathcal{A}_{\beta,\kappa}$ can stay to exact coverage given the history when the data are i.i.d. By (3), the instantaneous training-conditional coverage of split conformal prediction is within $O(t^{-1/2})$ of $1-\alpha$ in expectation at round t, so its cumulative deviation sum to $O(\sqrt{T})$ over T rounds, i.e., $\Delta_T^{\text{inst}} = O(\sqrt{T})$ for $C = C^{\text{CP}}$. For a nontrivial constraint, split conformal prediction, however, need not belong to $\mathcal{A}_{\beta,\kappa}$. We show that the smallest cumulative error in the instantaneous training-conditional coverage that an algorithm in $\mathcal{A}_{\beta,\kappa}$ can guarantee for i.i.d. data is of order

$$T^{\max\{1/2,\beta\}} = \begin{cases} \sqrt{T}, & 0 \leq \beta \leq 1/2, \\ T^{\beta}, & 1/2 < \beta \leq 1. \end{cases} \tag{7}$$

Thus, the constraint incurs no cost in rate when $\beta \leq 1/2$, whereas for $\beta > 1/2$ the rate is dominated by the constraint, yielding a phase transition at $\beta = 1/2$.

1.3 Adaptive conformal inference (ACI) revisited

Adaptive Conformal Inference (ACI), introduced by Gibbs and Candès (2021), is one way to enforce long-run empirical coverage on every data sequence. We consider its quantile-threshold

[2] For simplicity of presentation, we present the main theoretical results under symmetric constraints and give their asymmetric counterparts in Appendix A. We can likewise define asymmetric constraints: we say the algorithm satisfies a (b, d, c_-, c_+)-*constraint* if $\mathtt{E}_T/T \in [-c_- T^{-d},\ c_+ T^{-b}]$. We use $\mathcal{A}_{b,d,c_-,c_+}$ to collect such algorithms. The same phase transition holds in the asymmetric setting, with $\min\{b, d\}$ playing the role of β: the minimax oracle disagreement has order $T^{\max\{1/2,\min\{b,d\}\}}$.

form (Angelopoulos et al., 2024a,b), which uses a fixed score $s : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$ and constructs $C_t^{\mathrm{ACI}}(x) = \{y : s(x, y) \le q_t\}$. After observing the outcome, ACI updates its threshold by

$$q_{t+1} \;=\; q_t + \eta_t\big(\mathsf{e}_t - \alpha\big), \tag{8}$$

where $\eta_t > 0$ is the step size. We note that the original ACI procedure (Gibbs and Candès, 2021) updates the nominal coverage level. Here we analyze the quantile-threshold update (8). The threshold increases after miscoverage and decreases after coverage. The step sizes determine how tightly ACI controls the cumulative coverage deviation in the adversarial case. More concretely, when $s(x, y) \in [0, B]$, $\eta_t \equiv \eta$ (i.e., a constant step size), and $q_1 \in [0, B]$, ACI has the long-run empirical coverage guarantee in the adversarial setting (Gibbs and Candès, 2021; Angelopoulos et al., 2024a)

$$|\mathrm{E}_T| = \Big|\sum_{t=1}^{T}(\mathsf{e}_t - \alpha)\Big| \le \frac{B + \eta}{\eta},$$

indicating that ACI satisfies the constraint with $\beta = 1$ and a constant sufficiently large to dominate $(B + \eta)/\eta$.

1.4 Best-of-both-worlds algorithms?

Recall that split conformal prediction satisfies (3) but can violate (6) on some data sequences. Constant-step ACI satisfies (6), but its threshold moves by at least $\eta \min\{\alpha, 1 - \alpha\}$ after each observation, so its instantaneous training-conditional coverage continues to fluctuate even for i.i.d. data. Since it is rarely known in advance whether the data are i.i.d., we seek a single algorithm that satisfies the long-run constraint for every data sequence and, when the data are i.i.d., keeps its instantaneous training-conditional coverage as close as possible to $1 - \alpha$. We call such an algorithm *best-of-both-worlds*.

For ACI, this goal suggests preserving its long-run empirical coverage bound while reducing the size of each threshold update, that is, specifying decaying step sizes (Angelopoulos et al., 2024a). Concretely, when $s(x, y) \in [0, B]$, $q_1 \in [0, B]$, and $\eta_t = ct^{-(1-\beta)}$ with $c > 0$ and $0 < \beta < 1$, the bound of Angelopoulos et al. (2024a, Theorem 1) gives

$$\frac{|\mathrm{E}_T|}{T} \le \frac{B + c}{c}\, T^{-\beta} \qquad \text{for every } T \ge 1. \tag{9}$$

Thus ACI satisfies a (β, κ)-constraint with $\kappa = (B + c)/c$ on every data sequence. Faster decay of the step sizes corresponds to smaller β and a weaker long-run empirical coverage guarantee. Under i.i.d. sampling with a fixed bounded score whose CDF is continuous and whose $(1 - \alpha)$-quantile is unique, the steps $\eta_t = ct^{-(1-\beta)}$ with $0 < \beta < 1/2$ make the working quantile q_t converge almost surely to the population $(1 - \alpha)$-quantile (Angelopoulos et al., 2024a, Theorem 3). Instantaneous training-conditional coverage then converges to $1 - \alpha$ as well. Together with Equation (9), this gives both guarantees asymptotically.

Each new observation changes ACI's threshold by an amount proportional to η_t, even as the threshold converges to the population quantile. A larger β yields tighter long-run empirical

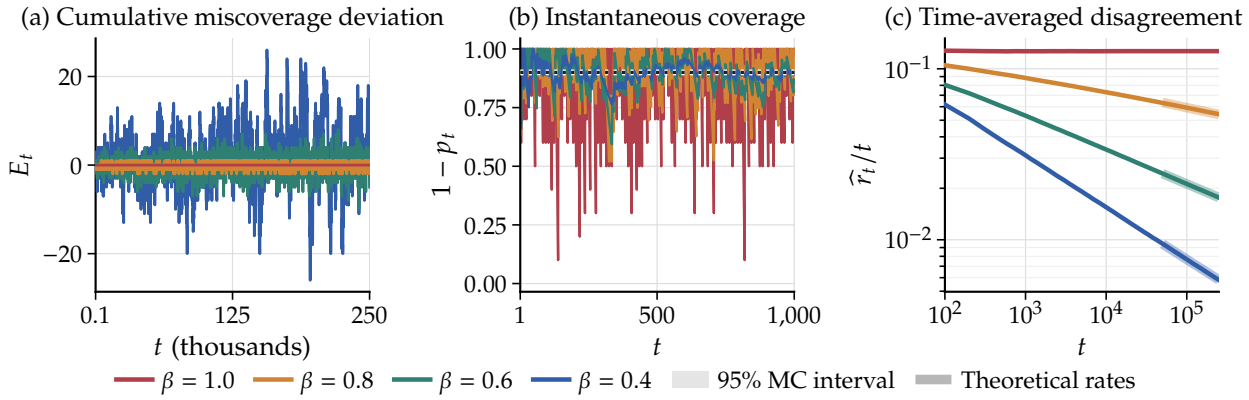


Figure 1: **Long-run empirical coverage versus oracle disagreement in ACI.** We use the regression model in Section 4, with $\alpha = 0.1$, $q_1 = 0.9$, and step sizes $\eta_t = t^{-(1-\beta)}$. We define $p_t = \mathbb{P}(Y_t \notin C_t^{\mathtt{ACI}}(X_t) \mid (X_s, Y_s)_{s<t})$ and compute $\widehat{r}_t$ by averaging $\sum_{s \le t} |p_s - \alpha|$. Dashed guides are proportional to $t^{-(1-\beta)/2}$ for $\beta < 1$, as predicted by Proposition 3.1.

coverage but causes these updates to decay more slowly, raising the question of how much error in the instantaneous training-conditional coverage accumulates over time. More concretely, for i.i.d. data under the score-density assumptions in Section 3, Proposition 3.1 quantifies this cost. With $\eta_t = ct^{-(1-\beta)}$ and $0 < \beta < 1$, it gives

$$\sum_{t=1}^{T} \mathbb{E}_{\mathbb{P}} \left| \mathbb{P}\{Y_t \in C_t^{\mathtt{ACI}}(X_t) \mid (X_s, Y_s)_{s<t}\} - (1-\alpha) \right| \asymp T^{(1+\beta)/2}. \tag{10}$$

For every $0 < \beta < 1$, this rate exceeds the $O(\sqrt{T})$ bound for split conformal prediction in (3). Split conformal prediction, however, does not itself enforce the long-run adversarial constraint on every sequence. Then it motivates us to ask the question: *what is the smallest cumulative error in instantaneous training-conditional coverage achievable under the long-run empirical coverage constraints?*

Figure 1 illustrates the tradeoff in a synthetic dataset (Section 4), where the fixed score, by construction, is uniformly distributed on $[0, 1]$. We compare $\eta_t = t^{-(1-\beta)}$ for $\beta \in \{1, 0.8, 0.6, 0.4\}$, including constant steps at $\beta = 1$. In panel (c), we focus on a metric that quantifies the expected disagreement between ACI prediction sets and the population score oracle; a formal definition is presented in Section 2. In this experiment, a larger β keeps the cumulative coverage deviation closer to zero, but instantaneous training-conditional coverage fluctuates more, and the time-averaged oracle disagreement decreases more slowly. These observations leave open the extent to which ACI's statistical cost is required by the coverage constraint and how much arises from its own algorithmic design.

Core question. Motivated by the aforementioned tradeoff, it remains an open question how much instantaneous training-conditional coverage error is unavoidable under the long-run empirical coverage constraint. Our goal is to find an algorithm that offers the *best-of-both-worlds* guarantee: it

satisfies the prescribed long-run empirical coverage constraint on every data sequence and achieves the best statistical performance when the data are i.i.d.

To make this objective precise, we compare an algorithm with an oracle that knows the data distribution and has coverage exactly $1-\alpha$ given the past. Their cumulative disagreement counts the rounds on which the observed response belongs to exactly one of their prediction sets. Our central question is:

> How closely can a single algorithm match an exact oracle under i.i.d. sampling while satisfying a prescribed long-run empirical coverage constraint on every data sequence?

In particular, can the $\sqrt{T}$ cumulative rate of split conformal prediction be retained, and when does tighter long-run empirical coverage force a higher cost? We seek the smallest possible cost and an algorithm that attains it without knowing the data distribution or whether the observations are i.i.d.

Paper organization. Section 2 formalizes oracle disagreement and separates the difficulty of approximating an oracle from the cost of the long-run empirical coverage constraint. For the family of all exact oracles, under i.i.d. laws with a nonatomic distribution for the fixed score, we prove a minimax rate of $T^{\max\{1/2,\beta\}}$ and achieve it by truncating split conformal prediction only when needed to enforce the constraint. Thus the $\sqrt{T}$ rate is retained for $\beta \le 1/2$, while tighter constraints necessarily increase the rate. Section 3 shows that fixed-score ACI with steps proportional to $t^{\beta-1}$ incurs a larger cost for $0 < \beta < 1$ under the stated density assumptions.

2 Oracle disagreement under an adversarial constraint

Recall the setting introduced in Section 1. We use $\mathcal{A}$ to denote the class of measurable online algorithms in (1) that are independent of the unknown data distribution. At round t, an online algorithm C uses the past observations $(X_r, Y_r)_{r<t}$ and the current feature X_t to construct a prediction set $C_t(X_t) \subseteq \mathcal{Y}$ before observing Y_t. Write $\mathrm{e}_t^C = \mathbf{1}\{Y_t \notin C_t(X_t)\}$ and $\mathrm{E}_T^C = \sum_{t=1}^T (\mathrm{e}_t^C - \alpha)$. With a fixed joint distribution $\mathbb{P}$ for the data sequence $(X_t, Y_t)_{t\in\mathbb{N}}$ and to make the conditioning clear, we define $\mathcal{G}_{t-1} = \sigma((X_r, Y_r)_{r<t})$.[3]

In this section, we formally quantify the instantaneous training-conditional performance of an online algorithm. To define this metric, we compare each prediction set $C_t(x)$ at time t with an *oracle* whose instantaneous training-conditional coverage is exactly calibrated at each t. Concretely, an oracle $C^*_{\mathbb{P}}$ may know $\mathbb{P}$, uses the same available information as C, and satisfies

$$\text{Definition of an oracle:} \qquad \mathbb{P}\big(Y_t \in C^*_{\mathbb{P},t}(X_t) \mid \mathcal{G}_{t-1}\big) = 1-\alpha \qquad \text{almost surely for every } t \ge 1. \tag{11}$$

As an example, under i.i.d. sampling with a measurable score function $s : \mathcal{X}\times\mathcal{Y} \to \mathbb{R}$, if the distribution of $s(X_1, Y_1)$ is nonatomic, let $F_{\mathbb{P},s}$ be its CDF and set

$$q_{\mathbb{P},s} = \inf\{q : F_{\mathbb{P},s}(q) \ge 1-\alpha\}, \qquad C^*_{\mathbb{P},s,t}(x) = \{y : s(x,y) \le q_{\mathbb{P},s}\}. \tag{12}$$

[3]When needed, we use $\zeta_{1:t} = (\zeta_1, \ldots, \zeta_t)$, which is independent of the data, to collect auxiliary randomization up to time t. Throughout this paper, all algorithms and oracles considered here are deterministic measurable functions of the observed data.

Independence of the data sequence and continuity of $F_{\mathbb{P},s}$ give the oracle condition (11). We note that, in general, the existence of $C^*_{\mathbb{P}}$ requires regularity conditions on the joint distribution $\mathbb{P}$, and for simplicity, we assume $\mathbb{P}$ is nonatomic throughout this paper unless otherwise specified.

To evaluate two online algorithms, we measure their cumulative disagreement on the same data sequence by

$$\ell_T(C, C') = \sum_{t=1}^{T} |\mathsf{e}_t^C - \mathsf{e}_t^{C'}|. \tag{13}$$

This counts responses that belong to exactly one of the two sets. Then, the expected disagreement with an oracle $C^*_{\mathbb{P}}$, i.e., $\mathbb{E}_{\mathbb{P}}[\ell_T(C, C^*_{\mathbb{P}})]$, quantifies the instantaneous training-conditional performance of an algorithm $C = (C_t)_{t\in\mathbb{N}}$ under a given data distribution $\mathbb{P}$. For $p_t^C = \mathbb{P}(\mathsf{e}_t^C = 1 \mid \mathcal{G}_{t-1})$, conditional Jensen's inequality further gives

$$\mathbb{E}_{\mathbb{P}}\ell_T(C, C^*_{\mathbb{P}}) \geq \sum_{t=1}^{T} \mathbb{E}_{\mathbb{P}}|p_t^C - \alpha| = \Delta_T^{\mathrm{inst}} \qquad \text{for every exact oracle } C^*_{\mathbb{P}}. \tag{14}$$

Equality holds when two prediction sets are coupled, i.e., the sets are defined by two thresholds on the same score, whose cutoffs are determined by past data.

2.1 Hardness of matching an oracle under long-run empirical coverage constraints

Let $\widetilde{\mathcal{Q}}$ be the class of all i.i.d. joint laws on the data sequence $(X_t, Y_t)_{t\in\mathbb{N}}$. Fix a measurable score $s : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$, and let

$$\mathcal{Q} = \{\mathbb{P} \in \widetilde{\mathcal{Q}} : s(X_1, Y_1) \text{ has a nonatomic distribution under } \mathbb{P}\}.$$

We assume that $\mathcal{Q}$ is nonempty. For each $\mathbb{P} \in \mathcal{Q}$, let $\mathcal{O}_\alpha(\mathbb{P})$ denote the set of algorithms that satisfy the oracle condition (11). Particularly, this family includes the population score oracle defined in (12).

We measure an algorithm's performance by

$$\sup_{\mathbb{P}\in\mathcal{Q}} \inf_{C^*_{\mathbb{P}}\in\mathcal{O}_\alpha(\mathbb{P})} \mathbb{E}_{\mathbb{P}}\ell_T(C, C^*_{\mathbb{P}}). \tag{15}$$

The infimum allows a separate oracle for each algorithm and law, without giving the algorithm access to the oracle or the distribution. Throughout, $\lesssim$, $\gtrsim$, and $\asymp$ compare quantities up to positive constants for all sufficiently large T.

Matching an oracle is difficult for two reasons. The algorithm must learn from an unknown distribution and keep its realized errors within the prescribed bounds. The following result separates these two costs.

Theorem 2.1. *Under the setup above, every $C \in \mathcal{A}_{\beta,\kappa}$ satisfies*

$$\sup_{\mathbb{P}\in\mathcal{Q}} \inf_{C^*_{\mathbb{P}}\in\mathcal{O}_\alpha(\mathbb{P})} \mathbb{E}_{\mathbb{P}}\ell_T(C, C^*_{\mathbb{P}}) \geq C_{\alpha,\beta,\kappa} T^{\max\{1/2,\beta\}}, \tag{16}$$

for sufficiently large T, where $C_{\alpha,\beta,\kappa} > 0$ depends only on α, β, κ.

More generally, the lower bound can be decomposed as

$$\sup_{\mathbb{P}\in\mathcal{Q}} \inf_{C^*_{\mathbb{P}}\in O_\alpha(\mathbb{P})} \mathbb{E}_{\mathbb{P}}\ell_T(C, C^*_{\mathbb{P}}) \gtrsim \underbrace{\inf_{C^0\in\mathcal{A}} \sup_{\mathbb{P}\in\mathcal{Q}} \inf_{C^*_{\mathbb{P}}\in O_\alpha(\mathbb{P})} \mathbb{E}_{\mathbb{P}}\ell_T(C^0, C^*_{\mathbb{P}})}_{\asymp\sqrt{T}} + \begin{cases} 0, & 0 \le \beta < 1/2, \\ T^\beta, & 1/2 \le \beta \le 1. \end{cases} \tag{17}$$

The first term is the unavoidable cost of learning from past data. Even without a long-run empirical coverage constraint, no algorithm can achieve disagreement smaller than order $\sqrt{T}$ uniformly over $\mathcal{Q}$. This remains true when comparison is allowed with every exact oracle. Split conformal prediction attains this order by estimating the population score quantile.

The second term is the cost of enforcing the long-run empirical coverage constraint. Exact calibration does not prevent the oracle's realized errors from fluctuating: its centered error sum has standard deviation $\sqrt{\alpha(1-\alpha)T}$. However, the constraint allows a deviation of only $\kappa T^{1-\beta}$. When $\beta \ge 1/2$, keeping the errors within these bounds requires repeated departures from the oracle. These departures force a cumulative disagreement with the order of at least T^β, even if the distribution is known.

The higher cost determines the lower-bound rate, which exhibits a sharp phase transition at $\beta = 1/2$. The cost of learning the unknown distribution dominates when $\beta < 1/2$, both costs have order $\sqrt{T}$ at $\beta = 1/2$, and the long-run empirical coverage constraint dominates when $\beta > 1/2$. In the next subsection, we construct an algorithm that achieves this rate in both regimes.

If distributions are not nonatomic, there may be no oracle that satisfies the condition (11), but we can approach the problem differently by considering relaxed oracles; see Section E.

2.2 Truncation attains the minimax rate

In this section, we present an explicit online construction of the online algorithm that has a matching upper bound in Theorem 2.1. Fix a measurable score $s : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$. Recall that $O_\alpha(\mathbb{P})$ denotes the set of oracle algorithms that satisfy (11) and under $\mathbb{P} \in \mathcal{Q}$, the distribution of score $s(X_t, Y_t)$ is nonatomic.

To meet the long-run empirical coverage constraints, a straightforward way is to truncate the cumulative miscoverage error $\mathtt{E}_T$ and enforce the adversarial constraints at each step. Concretely, for any baseline C^0, define $\widehat{C} = \Gamma_{\kappa,\beta} \circ C^0$ recursively, starting at $\mathtt{E}_0^{\widehat{C}} = 0$, by

$$(\Gamma_{\kappa,\beta} \circ C^0)_t(x) = \widehat{C}_t(x) = \begin{cases} \mathcal{Y}, & \mathtt{E}_{t-1}^{\widehat{C}} + 1 - \alpha > \kappa t^{1-\beta}, \\ \varnothing, & \mathtt{E}_{t-1}^{\widehat{C}} - \alpha < -\kappa t^{1-\beta}, \\ C_t^0(x), & \text{otherwise.} \end{cases} \tag{18}$$

If a potential miscoverage at round t would push the cumulative deviation above its upper bound, we output $\mathcal{Y}$ to ensure coverage. Similarly, if coverage would push $\mathtt{E}_t^{\widehat{C}}$ below its lower bound, we output $\varnothing$ to force a miscoverage. Otherwise, either outcome respects the bounds, so we keep the

baseline prediction set. $\mathrm{E}_t^{\widehat{C}}$ is computed from the errors of $\widehat{C}$, whereas the baseline C^0 continues to use every observation, including those from rounds in which its set is replaced by its truncated counterpart. For $\kappa \geq 1/2$, the two forcing conditions cannot hold simultaneously, and the resulting prediction sets satisfy (6) on every data sequence.

With split conformal prediction as the baseline, this truncation attains the minimax rate of expected oracle disagreement, as the following theorem shows.

Theorem 2.2. *The single algorithm $\Gamma_{\kappa,\beta} \circ C^{\mathrm{CP}} \in \mathcal{A}_{\beta,\kappa}$ attains the upper bound:*

$$\sup_{\mathbb{P}\in\mathcal{Q}} \inf_{C^*_{\mathbb{P}}\in O_\alpha(\mathbb{P})} \mathbb{E}_{\mathbb{P}}\ell_T(\Gamma_{\kappa,\beta} \circ C^{\mathrm{CP}}, C^*_{\mathbb{P}}) \lesssim T^{\max\{1/2,\beta\}}. \tag{19}$$

As an immediate corollary of Theorem 2.1 and Theorem 2.2, given the score, class $\mathcal{Q}$, and full oracle families defined above, we obtain the minimax rate

$$\inf_{C\in\mathcal{A}_{\beta,\kappa}} \sup_{\mathbb{P}\in\mathcal{Q}} \inf_{C^*_{\mathbb{P}}\in O_\alpha(\mathbb{P})} \mathbb{E}_{\mathbb{P}}\ell_T(C, C^*_{\mathbb{P}}) \asymp T^{\max\{1/2,\beta\}}. \tag{20}$$

The rate in (20) reflects two distinct costs. The first is the cost of approximating an exact oracle from past data. For split conformal prediction, disagreement with the population score oracle equals the expected error in instantaneous training-conditional coverage. Summing (3) therefore yields an $O(\sqrt{T})$ upper bound. The lower-bound argument for Theorem 2.1 shows that this order is unavoidable over $\mathcal{Q}$, even without a long-run empirical coverage constraint and after taking the infimum over all exact oracles.

The second cost comes from enforcing the constraint, even if the population prediction set is already known. As we already see from Theorem 2.1, even exactly calibrated prediction sets are forced to deviate from the oracle by the long-run empirical coverage constraint, and the resulting disagreement is at least of order T^β when $\beta \geq 1/2$.

Truncation attains the combined rate by controlling both costs. By Lemma C.1, the disagreement between the truncated baseline and the truncated oracle is at most twice the baseline's disagreement with the oracle. To compare with the original oracle, the triangle inequality adds the oracle's own truncation cost. Since an exact oracle has i.i.d. $\mathsf{Bern}(\alpha)$ errors, Lemma C.2 bounds this additional cost in expectation. Thus, for any exact oracle $C^*_{\mathbb{P}}$ and baseline $C^0 \in \mathcal{A}$, truncation gives

$$\mathbb{E}_{\mathbb{P}}\ell_T(\Gamma_{\kappa,\beta} \circ C^0, C^*_{\mathbb{P}}) \leq 2\mathbb{E}_{\mathbb{P}}\ell_T(C^0, C^*_{\mathbb{P}}) + K_{\alpha,\beta,\kappa} \begin{cases} 1, & 0 \leq \beta < 1/2, \\ T^\beta, & 1/2 \leq \beta \leq 1. \end{cases} \tag{21}$$

The bound holds for every $T \geq 1$. For split conformal prediction and the population score oracle, the first term is $O(\sqrt{T})$. When $\beta < 1/2$, the additional term is bounded uniformly in T, so enforcing long-run empirical coverage preserves the unconstrained statistical rate. At $\beta = 1/2$, the enforcement cost becomes $\sqrt{T}$, matching the cost of learning the target. For $\beta > 1/2$, the enforcement cost T^β dominates, and the optimal disagreement rate increases with the tightness of the constraint. This leads to the sharp transition in (20): the optimal rate remains at $\sqrt{T}$ up to $\beta = 1/2$, and becomes T^β beyond that threshold.

3 Comparing to ACI with decaying step sizes

We next compare truncation with ACI with decaying step sizes and a fixed score function, which adjusts its threshold after each observation. Fix $B > 0$, a measurable score $s : \mathcal{X} \times \mathcal{Y} \to [0, B]$, a deterministic initial threshold $q_1 \in [0, B]$, and positive deterministic step sizes $(\eta_t)_{t \geq 1}$. ACI constructs

$$C_t^{\mathrm{ACI}}(x) = \{y : s(x, y) \leq q_t\}, \qquad q_{t+1} = q_t + \eta_t(\mathrm{e}_t^{C^{\mathrm{ACI}}} - \alpha). \tag{22}$$

We study the adversarial guarantee of this update and its disagreement with the closest conditionally exact oracle defined in (11).

Adversarial coverage guarantee. For nonincreasing step sizes, the coverage bound of Angelopoulos et al. (2024a, Theorem 1) gives $|\mathrm{E}_T^{C^{\mathrm{ACI}}}| \leq (B + \eta_1)/\eta_T$ on every sequence of scores in $[0, B]$. In particular, for $\eta_t = ct^{-(1-\beta)}$ with $c > 0$ and $\beta \in (0, 1]$,

$$\frac{|\mathrm{E}_T^{C^{\mathrm{ACI}}}|}{T} \leq \frac{B + c}{c} T^{-\beta} \qquad \text{for every } T \geq 1. \tag{23}$$

Thus ACI has a (β, κ) adversarial guarantee with $\kappa = (B + c)/c$, indicating that, for a prescribed $\kappa > 1$, choosing $c = B/(\kappa - 1)$ yields the long-run coverage constraint.

Disagreement with the oracle. Now suppose the observations are i.i.d. under $\mathbb{P}$ and the score has a density satisfying

$$0 < \lambda_- \leq f_s(u) \leq \lambda_+ < \infty \qquad \text{for almost every } u \in [0, B]. \tag{24}$$

Then we have the following result, where the infimum over $C_{\mathbb{P}}^*$ below ranges over all oracles satisfying (11). The smallest expected disagreement is achieved by the population prediction set that uses the same score, as defined in (12).

Proposition 3.1 (Disagreement of ACI)**.** *With $\mathbb{P} \in \mathcal{Q}$ satisfying (24), if $\eta_t \to 0$ and $|\eta_{t+1}^{-1} - \eta_t^{-1}| \to 0$, then*

$$\inf_{C_{\mathbb{P}}^* \in O_\alpha(\mathbb{P})} \mathbb{E}_{\mathbb{P}}\left[\ell_T(C^{\mathrm{ACI}}, C_{\mathbb{P}}^*)\right] \asymp \sum_{t=1}^{T} \sqrt{\eta_t}. \tag{25}$$

ACI's threshold fluctuates around the population quantile on the scale $\sqrt{\eta_t}$, which creates disagreement even as the step sizes decay. Specifically, for $\eta_t = ct^{-(1-\beta)}$ with $0 < \beta < 1$, we have the following corollary.

Corollary 3.1. *With $\mathbb{P} \in \mathcal{Q}$ satisfying (24), for $\eta_t = ct^{-(1-\beta)}$ with $0 < \beta < 1$,*

$$\inf_{C_{\mathbb{P}}^* \in O_\alpha(\mathbb{P})} \mathbb{E}_{\mathbb{P}}\left[\ell_T(C^{\mathrm{ACI}}, C_{\mathbb{P}}^*)\right] \asymp T^{(1+\beta)/2}. \tag{26}$$

Since $(1 + \beta)/2 > \max\{1/2, \beta\}$, this grows faster than the disagreement of truncated split conformal prediction in Theorem 2.2 under the same adversarial guarantee with $\kappa > 1$. Thus, with the same fixed score function, for these step sizes, ACI's cumulative disagreement exceeds the minimax rate for every $\beta \in (0, 1)$. Truncated split conformal prediction attains the optimal rate by changing the baseline set only when the constraint forces coverage or miscoverage.

4 Numerical experiments

We study the cost of truncation and compare truncated split conformal prediction with ACI in a heteroscedastic regression model. Independently over time, let

$$X_t \sim \mathrm{Unif}[0,1], \qquad Y_t = f(X_t) + \sigma(X_t)\varepsilon_t, \qquad \varepsilon_t \sim \mathrm{Unif}[-1,1],$$

where X_t and ε_t are independent, $f(x) = \sin(2\pi x)$, and $\sigma(x) = 0.2 + 0.3x$. We use the fixed score $s(x,y) = \min\{|y - f(x)|/\sigma(x), 1\}$, which is uniform on $[0,1]$ under this model. We set $\alpha = 0.1$ and compare all algorithms with the population prediction set $C^*_{\mathbb{P},t}(x) = \{y : s(x,y) \le 0.9\}$.

We use the split-conformal baseline C^{CP} defined in (2). Each experiment in this section uses 500 independent replications over $T = 250{,}000$ time steps. We evaluate disagreement with the oracle using the known score distribution, which is not used to choose the conformal thresholds or update ACI. Since the prediction sets are nested score thresholds, their disagreement probability, conditional on past observations, equals the absolute difference between their instantaneous training-conditional miscoverage probabilities. We sum these probabilities within each replication and average the resulting totals to estimate expected cumulative disagreement.

Long-run empirical coverage and oracle disagreement. We first truncate the split conformal baseline with $\kappa = 0.5$. The baseline continues to use all past scores, including observations from rounds in which truncation forces coverage or miscoverage. Figure 2 shows paths of the cumulative deviation and estimates of expected disagreement with the population oracle. All simulated paths satisfy the long-run empirical coverage constraint at every round, up to numerical precision. For $\beta < 1/2$, disagreement remains close to that of the unconstrained baseline and follows the $\sqrt{T}$ rate predicted by theory. For $\beta \ge 1/2$, the curves are consistent with the T^{β} rate in Theorem 2.2.

The cost of truncation. To measure how much truncation changes the baseline's predictions, we compare $\Gamma_{\kappa,\beta} \circ C^{\mathrm{CP}}$ with C^{CP} on the same observations. Let $p_t^{\mathrm{CP}} = \mathbb{P}(Y_t \notin C_t^{\mathrm{CP}}(X_t) \mid \mathcal{G}_{t-1})$ be the baseline's instantaneous training-conditional miscoverage probability. The training-conditional disagreement probability is p_t^{CP} when truncation forces coverage, $1 - p_t^{\mathrm{CP}}$ when it forces miscoverage, and zero otherwise. Summing these probabilities and averaging across replications estimates $\mathbb{E}_{\mathbb{P}}[\ell_T(\Gamma_{\kappa,\beta} \circ C^{\mathrm{CP}}, C^{\mathrm{CP}})]$.

Figure 3 compares this truncation cost with the baseline and total disagreement for $\beta = 0.3, 0.5, 0.7$. The baseline disagreement is consistent with $\sqrt{T}$ growth in all three cases. For $\beta = 0.3$, the truncation cost remains nearly zero over the plotted horizons, so the total stays close to the baseline. For $\beta = 0.5$ and $\beta = 0.7$, the truncation cost grows approximately along the $\sqrt{T}$ and $T^{0.7}$ rates, respectively.

Comparison with ACI. We compare ACI with truncated split conformal prediction under the same $(\beta, 2)$ long-run empirical coverage guarantee, for $\beta \in \{0.3, 0.5, 0.7\}$. We initialize ACI at the population threshold $q_1 = 0.9$ and use step sizes $\eta_t = t^{-(1-\beta)}$. Since the score lies in $[0,1]$, (23) gives $\kappa = 2$ for ACI; we use the same value for truncation. Within each replication, both algorithms receive the same observations and are evaluated against the same population oracle.

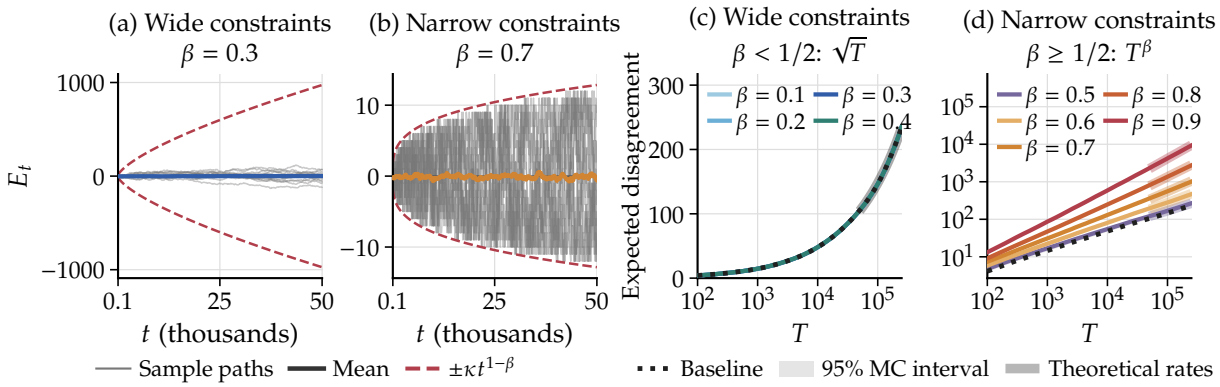


Figure 2: **Long-run empirical coverage and oracle disagreement of truncated split conformal prediction.** We use $\alpha = 0.1$ and $\kappa = 0.5$. Panels (a)–(b) show the cumulative deviations for $\beta = 0.3$ and $\beta = 0.7$, respectively: the first twelve paths, the mean over 500 replications, and the bounds $\pm\kappa t^{1-\beta}$. Paths are displayed every 100 rounds through $t = 50{,}000$. Panels (c)–(d) compare estimated expected cumulative disagreement with the population oracle for the truncated and unconstrained conformal algorithms. Both panels have logarithmic horizontal axes; panel (c) has a linear vertical axis and panel (d) a logarithmic one. The theoretical rates are proportional to $\sqrt{T}$ for $\beta < 1/2$ and T^β for $\beta \geq 1/2$.

Under this model, ACI accumulates disagreement faster than truncated split conformal prediction in all three panels of Figure 4. The ACI curves are consistent with the $T^{(1+\beta)/2}$ rate in Corollary 3.1.

5 Related work

Conformal prediction. Conformal prediction constructs prediction sets with finite-sample marginal coverage under exchangeability (Vovk et al., 2005; Shafer and Vovk, 2008; Angelopoulos et al., 2024b). Inductive or split conformal prediction separates model fitting from calibration (Papadopoulos et al., 2002; Lei et al., 2018). Conformalized quantile regression uses fitted conditional quantiles to adapt interval widths to heteroscedasticity (Romano et al., 2019), while distributional conformal prediction builds scores from an estimated conditional distribution (Chernozhukov et al., 2021).

Extensions address covariate shift (Tibshirani et al., 2019), label shift (Podkopaev and Ramdas, 2021), and coverage loss when exchangeability fails (Barber et al., 2023). Cauchois et al. (2024) construct prediction sets with finite-sample coverage bounds for test distributions within an f-divergence ball. Related work on robust risk evaluation incorporates isotonic constraints on density ratios (Gui et al., 2026). Applications include inference for counterfactuals and individual treatment effects (Lei and Candès, 2021; Jin et al., 2023), prediction with censored survival outcomes (Candès et al., 2023; Gui et al., 2024), missing covariates (Zaffran et al., 2023), and unobserved matrix entries (Gui et al., 2023).

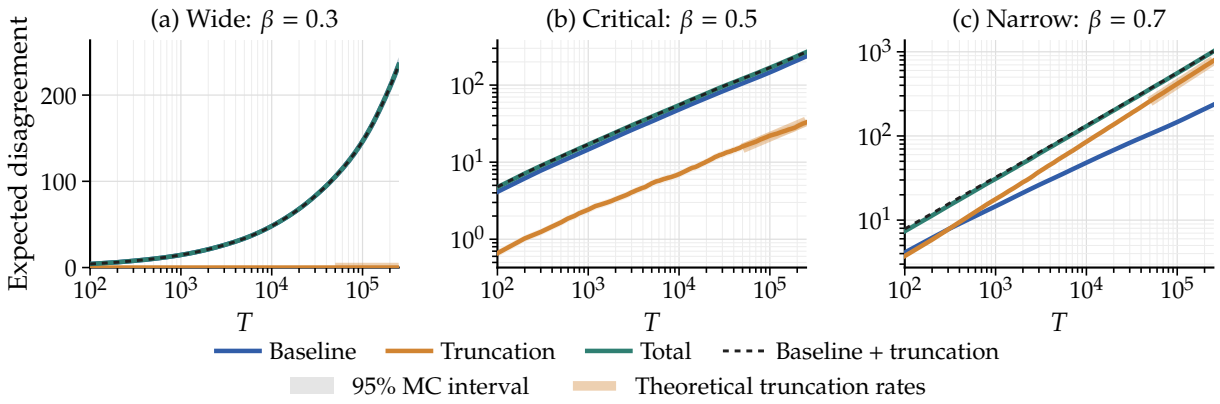


Figure 3: **Baseline disagreement, truncation cost, and total disagreement.** For $\kappa = 0.5$, the three panels show $\beta = 0.3, 0.5, 0.7$, respectively. Each compares $\mathbb{E}_{\mathbb{P}}[\ell_T(C^{\mathrm{CP}}, C^*_{\mathbb{P}})]$, $\mathbb{E}_{\mathbb{P}}[\ell_T(\Gamma_{\kappa,\beta} \circ C^{\mathrm{CP}}, C^{\mathrm{CP}})]$, and $\mathbb{E}_{\mathbb{P}}[\ell_T(\Gamma_{\kappa,\beta} \circ C^{\mathrm{CP}}, C^*_{\mathbb{P}})]$. The sum curve adds the first two estimates from paired replications and bounds the corresponding total estimate from above. All horizontal axes are logarithmic. Panel (a) uses a linear vertical axis to show values near zero; panels (b)–(c) use logarithmic vertical axes.

In the batch setting, Areces et al. (2024) establish a minimax lower bound of order $n^{-1/2}$ on the expected training-conditional coverage error after n i.i.d. observations, for prediction sets defined by thresholds or intervals on a fixed score. Duchi (2025) develops sample-conditional coverage bounds for split conformal prediction and relates them to these lower bounds. We study cumulative error in online prediction and the additional cost of a long-run empirical coverage constraint.

Online conformal prediction and adaptive calibration. Adaptive calibration uses past coverage errors to adjust prediction sets as the data distribution changes. Adaptive conformal inference (ACI) adjusts the nominal coverage level and controls long-run empirical coverage on arbitrary data sequences (Gibbs and Candès, 2021). Subsequent methods adapt learning rates with local dynamic-regret guarantees (Gibbs and Candès, 2024), or aggregate ACI predictors with different learning rates (Zaffran et al., 2022). Betting-based methods use parameter-free online optimization to avoid choosing a learning rate (Podkopaev et al., 2024).

Direct threshold updates connect long-run empirical coverage control with quantile estimation. Conformal PID control combines quantile tracking, correction for errors in cumulative coverage, and forecasts of future score quantiles (Angelopoulos et al., 2023). For fixed bounded scores, appropriately decaying step sizes preserve long-run empirical coverage and yield quantile convergence under i.i.d. sampling and a unique target quantile (Angelopoulos et al., 2024a).

Time-series methods use past residuals and forecasts to adapt prediction intervals to temporal dependence. EnbPI uses bootstrap ensemble residuals and updates their empirical quantiles as new observations arrive (Xu and Xie, 2021). Sequential predictive conformal inference instead estimates conditional quantiles of future residuals from past residuals (Xu and Xie, 2023). Bellman conformal inference uses multi-step forecasts and dynamic programming to choose prediction intervals (Yang et al., 2024). Jiang et al. (2026) document severe undercoverage of the leave-one-out

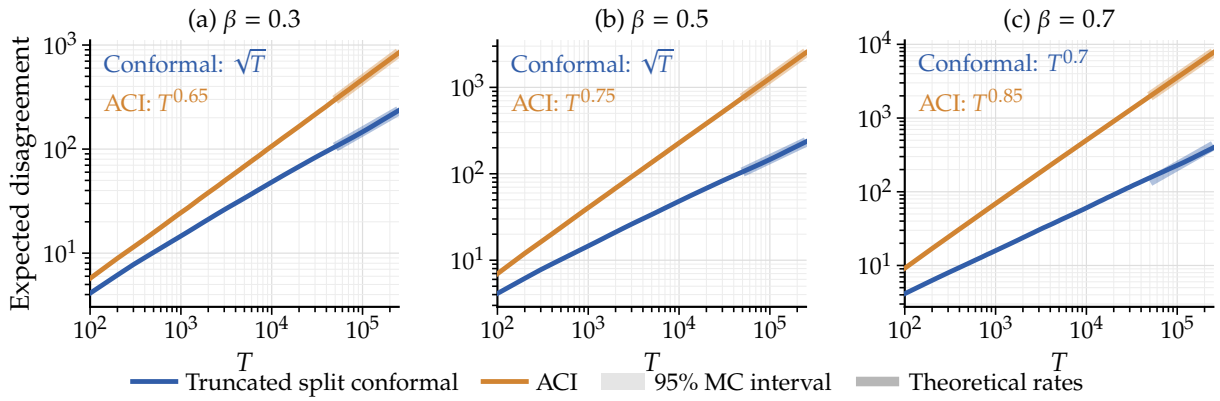


Figure 4: **ACI and truncated split conformal prediction under the same long-run empirical coverage guarantee.** Both algorithms use $\alpha = 0.1$ and satisfy the $(\beta, 2)$ constraint. Each panel uses logarithmic axes and shows estimated expected cumulative disagreement with the population oracle over 500 paired replications. ACI uses $q_1 = 0.9$ and $\eta_t = t^{-(1-\beta)}$. The theoretical rates are proportional to $T^{(1+\beta)/2}$ for ACI and $T^{\max\{1/2,\beta\}}$ for truncated split conformal prediction.

jackknife under short-range temporal dependence. Their leave-a-window-out modification omits blocks of observations and gives approximate marginal-coverage guarantees under stability of the fitting algorithm and quantitative bounds on temporal dependence.

Srinivas (2026) studies the trade-off between empirical coverage and average interval length, measured relative to the shortest fixed interval that achieves the target coverage in hindsight. Areces et al. (2025) formulate calibration as an online optimization problem, obtaining the long-run empirical coverage in the adversarial setting and convergence of instantaneous training-conditional coverage under assumptions that include a linear model for conditional score quantiles. When the data law changes, a central goal is to keep training-conditional coverage close to its target. Liang et al. (2026) study Δ_T^{inst}, defined in (4), for independent observations with unknown change points or bounded cumulative distribution drift. Their methods combine drift detection and recalibration and attain minimax rates up to logarithmic factors within the algorithm classes they study, but do not impose a long-run empirical coverage constraint that holds on every data sequence.

Regret, local coverage, and group coverage. Strongly adaptive online methods bound quantile-loss regret over all time intervals (Bhatnagar et al., 2023), where, however, quantile-loss regret alone does not directly imply coverage guarantees. Multivalid methods control empirical coverage across groups and threshold ranges under smoothness conditions on the score distributions (Bastani et al., 2022). Ramalingam et al. (2025) distinguish regret from coverage, relate swap regret to multivalidity, and develop group-conditional extensions of ACI. Other work studies exponentially discounted regret and coverage (Zhang et al., 2024), selection among candidate prediction models (Hajihashemi and Shen, 2024), and gradient equilibrium, which relates vanishing average loss gradients to online calibration (Angelopoulos et al., 2025).

6 Discussion

We seek a *best-of-both-worlds* online algorithm whose prediction sets satisfy the long-run empirical coverage constraint (6) on every data sequence and disagree as little as possible with oracles whose instantaneous training-conditional coverage is exactly $1 - \alpha$ for i.i.d. data. We show that, over the class $\mathcal{Q}$ of i.i.d. distributions with a nonatomic score, the minimax expected disagreement with the closest such oracle is of order $T^{\max\{1/2,\beta\}}$ under this constraint. The constraint therefore leaves the unconstrained rate $\sqrt{T}$ unchanged for $\beta \leq 1/2$ and raises it to T^{β} for $\beta > 1/2$. Truncated split conformal prediction attains this rate. With the same fixed score, ACI with step sizes proportional to $t^{-(1-\beta)}$ satisfies the same constraint, but its expected disagreement is of order $T^{(1+\beta)/2}$ when the score has a density bounded away from zero and infinity. This rate exceeds $T^{\max\{1/2,\beta\}}$ for every $0 < \beta < 1$. This paper also leaves the following questions open. Since the minimax-optimal construction may not be unique, it may be of interest to refine the optimality criterion by quantifying the accuracy of such prediction sets via the volume. Open questions include how to retain informative prediction sets and adaptive instantaneous training-conditional coverage under distribution drift while preserving a prescribed adversarial constraint.

Acknowledgement

R.F.B. was partially supported by the Office of Naval Research via grant N00014-24-1-2544. C.M. was partially supported by the National Science Foundation via the CAREER Award DMS2443867.

Disclosure of AI usage

GPT-6-Astra and Claude-Fable-5.1 models were used to produce numerical experiments, polish writing, suggest literature, and audit existing proofs; the authors take full responsibility for the contents of the paper.

References

Angelopoulos, A., Candès, E., and Tibshirani, R. J. (2023). Conformal PID control for time series prediction. In *Advances in Neural Information Processing Systems*, volume 36, pages 23047–23074.

Angelopoulos, A. N., Barber, R., and Bates, S. (2024a). Online conformal prediction with decaying step sizes. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 1616–1630. PMLR.

Angelopoulos, A. N., Barber, R. F., and Bates, S. (2024b). Theoretical foundations of conformal prediction. *arXiv preprint arXiv:2411.11824*.

Angelopoulos, A. N., Jordan, M. I., and Tibshirani, R. J. (2025). Gradient equilibrium in online learning: Theory and applications. *Journal of Machine Learning Research*, 26(305):1–68.

Areces, F., Cheng, C., Duchi, J., and Kuditipudi, R. (2024). Two fundamental limits for uncertainty

quantification in predictive inference. In *The Thirty Seventh Annual Conference on Learning Theory*, pages 186–218. PMLR.

Areces, F., Mohri, C., Hashimoto, T., and Duchi, J. (2025). Online conformal prediction via online optimization. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 1604–1649. PMLR.

Barber, R. F., Candès, E. J., Ramdas, A., and Tibshirani, R. J. (2023). Conformal prediction beyond exchangeability. *The Annals of Statistics*, 51(2):816–845.

Bastani, O., Gupta, V., Jung, C., Noarov, G., Ramalingam, R., and Roth, A. (2022). Practical adversarial multivalid conformal prediction. In *Advances in Neural Information Processing Systems*, volume 35, pages 29362–29373.

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

Candès, E., Lei, L., and Ren, Z. (2023). Conformalized survival analysis. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 85(1):24–45.

Cauchois, M., Gupta, S., Ali, A., and Duchi, J. C. (2024). Robust validation: Confident predictions even when distributions shift. *Journal of the American Statistical Association*, 119(548):3033–3044.

Chernozhukov, V., Wüthrich, K., and Zhu, Y. (2021). Distributional conformal prediction. *Proceedings of the National Academy of Sciences*, 118(48):e2107794118.

Chung, K. L. (1954). On a stochastic approximation method. *The Annals of Mathematical Statistics*, pages 463–483.

Duchi, J. C. (2025). A few observations on sample-conditional coverage in conformal prediction. *arXiv preprint arXiv:2503.00220*.

Durrett, R. (2019). *Probability: theory and examples*, volume 49. Cambridge university press.

Fabian, V. (1968). On asymptotic normality in stochastic approximation. *The Annals of Mathematical Statistics*, pages 1327–1332.

Gibbs, I. and Candès, E. J. (2021). Adaptive conformal inference under distribution shift. In *Advances in Neural Information Processing Systems*, volume 34, pages 1660–1672.

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

Gui, Y., Barber, R., and Ma, C. (2023). Conformalized matrix completion. *Advances in Neural Information Processing Systems*, 36:4820–4844.

Gui, Y., Barber, R. F., and Ma, C. (2026). Distributionally robust risk evaluation with an isotonic constraint. *Information and Inference: A Journal of the IMA*, 15(3):iaag025.

Gui, Y., Hore, R., Ren, Z., and Barber, R. F. (2024). Conformalized survival analysis with adaptive cut-offs. *Biometrika*, 111(2):459–477.

Hajihashemi, E. and Shen, Y. (2024). Multi-model ensemble conformal prediction in dynamic environments. In *Advances in Neural Information Processing Systems*, volume 37, pages 118678–

118700.

Hoeffding, W. (1963). Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30.

Jiang, H., Barber, R. F., Pananjady, A., and Xie, Y. (2026). Leave a window out: Modifying the jackknife for predictive inference in time series. *arXiv preprint arXiv:2605.30292*.

Jin, Y., Ren, Z., and Candès, E. J. (2023). Sensitivity analysis of individual treatment effects: A robust conformal inference approach. *Proceedings of the National Academy of Sciences*, 120(6):e2214889120.

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

Lei, L. and Candès, E. J. (2021). Conformal inference of counterfactuals and individual treatment effects. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 83(5):911–938.

Liang, J., Ren, Z., and Chen, Y. (2026). Optimal training-conditional regret for online conformal prediction. *arXiv preprint arXiv:2602.16537*.

Papadopoulos, H., Proedrou, K., Vovk, V., and Gammerman, A. (2002). Inductive confidence machines for regression. In *Machine Learning: ECML 2002*, volume 2430 of *Lecture Notes in Computer Science*, pages 345–356. Springer.

Podkopaev, A. and Ramdas, A. (2021). Distribution-free uncertainty quantification for classification under label shift. In *Uncertainty in artificial intelligence*, pages 844–853. PMLR.

Podkopaev, A., Xu, D., and Lee, K.-C. (2024). Adaptive conformal inference by betting. *arXiv preprint arXiv:2412.19318*.

Ramalingam, R., Kiyani, S., and Roth, A. (2025). The relationship between no-regret learning and online conformal prediction. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 51060–51078. PMLR.

Ramsey, F. P. (1930). On a problem of formal logic. *Proceedings of the London Mathematical Society*, s2-30(1):264–286.

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

Shafer, G. and Vovk, V. (2008). A tutorial on conformal prediction. *Journal of Machine Learning Research*, 9(12):371–421.

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

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

Van der Vaart, A. W. (2000). *Asymptotic statistics*, volume 3. Cambridge university press.

Vovk, V. (2012). Conditional validity of inductive conformal predictors. In *Asian conference on machine learning*, pages 475–490. PMLR.

Vovk, V., Gammerman, A., and Shafer, G. (2005). *Algorithmic learning in a random world*. Springer.

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

Xu, C. and Xie, Y. (2023). Sequential predictive conformal inference for time series. In *International Conference on Machine Learning*, pages 38707–38727. PMLR.

Yang, Z., Huang, Y. M., Candès, E., and Lei, L. (2024). Bellman conformal inference: Calibrating prediction intervals for time series. *arXiv preprint arXiv:2402.05203*.

Zaffran, M., Dieuleveut, A., Josse, J., and Romano, Y. (2023). Conformal prediction with missing values. In *International Conference on Machine Learning*, pages 40578–40604. PMLR.

Zaffran, M., Féron, O., Goude, Y., Josse, J., and Dieuleveut, A. (2022). Adaptive conformal predictions for time series. In *International Conference on Machine Learning*, pages 25834–25866. PMLR.

Zhang, Z., Bombara, D., and Yang, H. (2024). Discounted adaptive online learning: Towards better regularization. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 58631–58661. PMLR.

Contents

A Results under asymmetric constraints **20**
A.1 Setting and notation . . . 20
A.2 Unconstrained minimax rate . . . 21
A.3 Lower bounds . . . 22
A.4 Truncation and the minimax rate . . . 22
A.5 Proof of Theorem A.3 . . . 22

B Proof of the asymmetric lower bound **23**
B.1 Proof of Theorem A.1 . . . 24
B.2 Proof of Lemma B.1 . . . 25
B.3 Proof of Lemma B.2 . . . 25
B.4 Justification of Lemma B.3 . . . 25

C Proof of the truncation upper bound **26**
C.1 Proof of Theorem A.2 . . . 26
C.2 Proof of Lemma C.1 . . . 27
C.3 Proof of Lemma C.2 . . . 27
C.4 Proof of Lemma C.3 . . . 28
C.5 Proof of Lemma C.4 . . . 29
C.6 Proof of Lemma C.5 . . . 29

D Unconstrained minimax rate **31**
D.1 Proof of Lemma A.1 . . . 31
D.2 Proof of Lemma D.1 . . . 32
D.3 Proof of Lemma D.2 . . . 33
D.4 Proof of Lemma D.3 . . . 35

E Results with a relaxed oracle condition **35**
E.1 Proof of Proposition E.1 . . . 36
E.2 Proof of Lemma E.1 . . . 37

F Proofs for ACI **37**
F.1 Proof of Proposition 3.1 . . . 39
F.2 Proof of Corollary 3.1 . . . 40
F.3 Proofs of Lemmas F.3 and F.4 . . . 40
F.4 Proof of Lemma F.5 . . . 41
F.5 Proof of Lemma F.6 . . . 42
F.6 Proof of Lemma F.7 . . . 43

A Results under asymmetric constraints

The main text uses the symmetric constraint (6). Here we allow its two sides to grow at different rates and prove the main results in this setting. The symmetric statements follow by taking $b = d = \beta$ and $c_- = c_+ = \kappa$. The ACI result is proved separately in Appendix F.

A.1 Setting and notation

Fix $\alpha \in (0,1)$, $b, d \in [0,1]$, and $c_-, c_+ > 0$ with $c_- + c_+ \geq 1$. Write

$$L_t = c_- t^{1-d}, \qquad U_t = c_+ t^{1-b}, \qquad \underline{\beta} = \min\{b,d\}, \qquad \bar{\beta} = \max\{b,d\}, \qquad t \geq 1.$$

Set $L_0 = U_0 = 0$.

Notations. We write $\sigma = \sqrt{\alpha(1-\alpha)}$, $\nu = \min\{\alpha, 1-\alpha\}$, and $x_+ = \max\{x, 0\}$. Write $Z_t = (X_t, Y_t) \in \mathcal{Z} = \mathcal{X} \times \mathcal{Y}$ and $\mathcal{G}_{t-1} = \sigma((Z_r)_{r<t})$. For a fixed measurable score s, put $s_t = s(Z_t)$. Write $\widetilde{\mathcal{Q}}$ for all i.i.d. distributions on the observation sequence and $\mathcal{Q}$ for the nonempty subclass under which s_1 is atomless. For an i.i.d. distribution $\mathbb{P}$, denote its one-observation marginal by $\mu_{\mathbb{P}}$, so $\mathbb{P} = \mu_{\mathbb{P}}^{\otimes\mathbb{N}}$. For fixed $\mathbb{P}$ and s, abbreviate the score CDF $F_{\mathbb{P},s}$ to F and retain $q_{\mathbb{P},s}$ for the population quantile in (12).

Write $e^C = (\mathsf{e}_t^C)_{t\geq 1}$ for the error sequence of C and $p_t^C = \mathbb{P}(\mathsf{e}_t^C = 1 \mid \mathcal{G}_{t-1})$ for its instantaneous training-conditional miscoverage probability. The error on a history h and fresh observation z is a product-measurable function $e_t(h, z)$. Use e^0 and e^* for the error sequences of the baseline and comparison oracle, respectively. When a single algorithm is fixed, we may omit its superscript. For binary paths u, v, write

$$\mathtt{E}_t(u) = \sum_{r=1}^{t} (u_r - \alpha), \qquad D_T(u,v) = \sum_{r=1}^{T} |u_r - v_r|, \qquad \mathtt{E}_0(u) = 0.$$

Thus $\mathtt{E}_t^C = \mathtt{E}_t(e^C)$ and $\ell_T(C, C') = D_T(e^C, e^{C'})$. When comparing a constrained path u to a reference path v, abbreviate their centered sums to $\mathtt{E}_t = \mathtt{E}_t(u)$ and $S_t = \mathtt{E}_t(v)$, respectively.

The coverage constraint. The class $\mathcal{A}_{b,d,c_-,c_+}$ consists of algorithms whose centered errors satisfy

$$-L_t \leq \mathtt{E}_t^C \leq U_t \qquad \text{for every } t \geq 1 \text{ and every data sequence.} \tag{27}$$

Thus b controls the upper side and d the lower side. All algorithms and oracles in this appendix are deterministic, as in the main text.

Every binary path u satisfies $-\alpha t \leq \mathtt{E}_t(u) \leq (1-\alpha)t$. The first horizon at which the constraint excludes a binary sequence is

$$T_0 = \inf\{t \in \mathbb{N}_{\geq 1} : U_t < (1-\alpha)t \text{ or } L_t < \alpha t\}, \qquad \inf \varnothing = \infty. \tag{28}$$

We call the constraint nontrivial when $T_0 < \infty$. Otherwise $b = d = 0$, $c_+ \geq 1-\alpha$, and $c_- \geq \alpha$, so $\mathcal{A}_{b,d,c_-,c_+} = \mathcal{A}$.

For nontrivial constraints, the cost of imposing the constraint on an exact oracle has three regimes:

$$g_{b,d}(T) = \begin{cases} 1, & \bar{\beta} < 1/2, \\ \sqrt{T}, & \underline{\beta} \le 1/2 \le \bar{\beta}, \\ T^{\underline{\beta}}, & \underline{\beta} > 1/2. \end{cases} \tag{29}$$

Both boundaries grow faster than $\sqrt{t}$ in the first regime. At least one boundary grows at most as fast as $\sqrt{t}$ in the second, and both grow more slowly in the third. In the symmetric case, $g_{\beta,\beta}(T)$ is 1 for $\beta < 1/2$ and T^{β} for $\beta \ge 1/2$.

Asymmetric truncation. For a baseline C^0, let $\widehat{C} = \Gamma_{\text{asym}} \circ C^0$ denote the truncated algorithm for the fixed parameters (b, d, c_-, c_+), which is defined by

$$\widehat{C}_t(x) = \begin{cases} \mathcal{Y}, & \mathrm{E}_{t-1}^{\widehat{C}} + 1 - \alpha > U_t, \\ \varnothing, & \mathrm{E}_{t-1}^{\widehat{C}} - \alpha < -L_t, \\ C_t^0(x), & \text{otherwise}, \end{cases} \qquad \mathrm{E}_0^{\widehat{C}} = 0. \tag{30}$$

For a binary input e, write $\widehat{e} = \Gamma_{\text{asym}}(e)$ for the output on errors. The first two conditions in (30), evaluated at the centered sum of $\widehat{e}$, specify rounds forced from above and below, respectively. In particular,

$$e^{\Gamma_{\text{asym}} \circ C^0} = \Gamma_{\text{asym}}(e^{C^0}), \qquad \widehat{e}_t = \begin{cases} 0, & \text{forced from above}, \\ 1, & \text{forced from below}, \\ e_t, & \text{otherwise}. \end{cases} \tag{31}$$

The two conditions are disjoint since $L_t + U_t \ge 1$. The baseline uses every observation, including observations from forced rounds. Under $b = d = \beta$ and $c_- = c_+ = \kappa$, Γ_{asym} specializes to $\Gamma_{\kappa,\beta}$ in (18).

Let N_T^+, N_T^- count the rounds forced from above and below through T, and M_T^+, M_T^- the rounds on which those respective branches change the input error. Write $N_T = N_T^+ + N_T^-$ and use subscript ∞ for the corresponding limits. Appendix C proves feasibility and stability.

Oracle risks. When $\mathbb{P}$ is fixed, write $C_{\mathbb{P}}^*$ for an oracle in the exact-oracle family $\mathcal{O}_\alpha(\mathbb{P})$ of (11). Define

$$\begin{aligned} r_T(C; \mathbb{P}) &= \inf_{C_{\mathbb{P}}^* \in \mathcal{O}_\alpha(\mathbb{P})} \mathbb{E}_{\mathbb{P}} \ell_T(C, C_{\mathbb{P}}^*), \\ R_T(C) &= \sup_{\mathbb{P} \in \mathcal{Q}} r_T(C; \mathbb{P}), \qquad R_T^0 = \inf_{C \in \mathcal{A}} R_T(C). \end{aligned} \tag{32}$$

All constants in the constraint bounds below depend only on α, b, d, c_-, c_+.

A.2 Unconstrained minimax rate

Without a long-run empirical coverage constraint, the cost of learning from past observations is of order $\sqrt{T}$, even when the comparison ranges over all exact oracles.

Lemma A.1. *For every $T \geq 1$,*

$$\frac{\nu^2}{32}\sqrt{T} \leq R_T^0 \leq 3\sqrt{T}.$$

Appendix D proves the lemma and shows that split conformal prediction attains its upper bound.

A.3 Lower bounds

Theorem A.1 (Disagreement forced by an asymmetric constraint)**.** *Suppose that $T_0 < \infty$. For every i.i.d. law $\mathbb{P}$ admitting an exact oracle, every such oracle $C^*_{\mathbb{P}}$, and every online algorithm C satisfying (27),*

$$\mathbb{E}_{\mathbb{P}}\ell_T(C, C^*_{\mathbb{P}}) \geq c\, g_{b,d}(T) \qquad (T \geq T_1), \tag{33}$$

where $c > 0$ and T_1 are uniform over the law, algorithm, and oracle. This statement allows C to know $\mathbb{P}$. Consequently, every $C \in \mathcal{A}_{b,d,c_-,c_+}$ satisfies

$$R_T(C) \gtrsim R_T^0 + g_{b,d}(T). \tag{34}$$

Appendix B proves the constraint bound and combines the two costs. Under the symmetric specialization, this gives Theorem 2.1.

A.4 Truncation and the minimax rate

Theorem A.2 (Stability of truncation)**.** *For every baseline $C^0 \in \mathcal{A}$, the algorithm $\Gamma_{\text{asym}} \circ C^0$ belongs to $\mathcal{A}_{b,d,c_-,c_+}$. For every i.i.d. law $\mathbb{P}$ and every exact oracle $C^*_{\mathbb{P}}$ under that law,*

$$\mathbb{E}_{\mathbb{P}}\ell_T(\Gamma_{\text{asym}} \circ C^0, C^*_{\mathbb{P}}) \leq 2\mathbb{E}_{\mathbb{P}}\ell_T(C^0, C^*_{\mathbb{P}}) + K\, g_{b,d}(T) \qquad (T \geq 1), \tag{35}$$

where $K < \infty$ is independent of the baseline, law, and oracle.

This result does not require a continuous score distribution or calibrated baseline errors. It uses only the exact oracle condition of the comparison algorithm. Appendix C proves it, and the symmetric case gives (21).

Theorem A.3 (Minimax disagreement under asymmetric constraints)**.** *For the fixed score s, nonempty class $\mathcal{Q}$, and full exact-oracle families of the main text,*

$$\inf_{C \in \mathcal{A}_{b,d,c_-,c_+}} R_T(C) \asymp T^{\max\{1/2,\beta\}}. \tag{36}$$

The single algorithm $\Gamma_{\text{asym}} \circ C^{\text{CP}}$ attains the upper bound.

Appendix A.5 combines the learning and constraint bounds to prove this theorem. Taking $b = d = \beta$ and $c_- = c_+ = \kappa$ gives Theorem 2.2.

A.5 Proof of Theorem A.3

Proof. We first compare the three possible constraint costs with the learning cost. We then establish matching lower and upper bounds.

Step 1: identify the resulting rate. The definition of $g_{b,d}$ gives

$$\sqrt{T} + g_{b,d}(T) \asymp T^{\max\{1/2,\underline{\beta}\}}. \tag{37}$$

Indeed, $g_{b,d}(T)$ is bounded when $\bar{\beta} < 1/2$ and is $\sqrt{T}$ when $\underline{\beta} \le 1/2 \le \bar{\beta}$. In both cases the sum has order $\sqrt{T}$. When $\underline{\beta} > 1/2$, the term $T^{\underline{\beta}}$ is the larger term.

Step 2: bound every constrained algorithm from below. Suppose first that the constraint is nontrivial. By Theorem A.1 and Lemma A.1, every $C \in \mathcal{A}_{b,d,c_-,c_+}$ satisfies

$$R_T(C) \gtrsim \sqrt{T} + g_{b,d}(T) \asymp T^{\max\{1/2,\underline{\beta}\}}.$$

The constants and initial horizon are uniform over C, so taking the infimum over $\mathcal{A}_{b,d,c_-,c_+}$ preserves this bound. If the constraint is trivial, then $\mathcal{A}_{b,d,c_-,c_+} = \mathcal{A}$ and $\underline{\beta} = 0$. In that case Lemma A.1 gives the same conclusion directly.

Step 3: attain the rate with truncated conformal prediction. Take $C^0 = C^{\mathrm{CP}}$ in (30). By Theorem A.2, this algorithm belongs to $\mathcal{A}_{b,d,c_-,c_+}$ on every data sequence. For each $\mathbb{P} \in \mathcal{Q}$, the population score oracle is exact, so

$$\begin{aligned}\inf_{C^*_{\mathbb{P}} \in O_\alpha(\mathbb{P})} \mathbb{E}_{\mathbb{P}}\ell_T(\widehat{C}, C^*_{\mathbb{P}}) &\le \mathbb{E}_{\mathbb{P}}\ell_T(\widehat{C}, C^*_{\mathbb{P},s}) \\ &\le 2\mathbb{E}_{\mathbb{P}}\ell_T(C^{\mathrm{CP}}, C^*_{\mathbb{P},s}) + Kg_{b,d}(T) \\ &\le 6\sqrt{T} + Kg_{b,d}(T).\end{aligned}$$

The last bound is uniform over $\mathbb{P}$, and (37) gives the desired upper rate. Neither the conformal baseline nor truncation depends on T, so the same algorithm works for all horizons.

Finally, setting $b = d = \beta$ and $c_- = c_+ = \kappa$ turns Γ_{asym} into the truncation rule (18). The upper and lower bounds then give (20) and (19). □

B Proof of the asymmetric lower bound

The proof is based on the following facts.

Lemma B.1. *If a deterministic online oracle satisfies* (11)*, its errors are i.i.d.* $\mathsf{Bern}(\alpha)$.

Lemma B.2 (Comparing a constrained path with a reference path)**.** *Suppose that* $-L_t \le \mathtt{E}_t \le U_t$ *for every* $t \le T$*. For* $0 \le r < r' \le T$*,*

$$\sum_{t=r+1}^{r'} |e_t - e^*_t| \ge |S_{r'} - S_r| - (c_- + c_+)T^{1-\underline{\beta}}. \tag{38}$$

At the terminal horizon,

$$D_T(e, e^*) \ge (S_T - U_T)_+, \qquad D_T(e, e^*) \ge (-S_T - L_T)_+. \tag{39}$$

Lemma B.3 (Fluctuations of the Bernoulli walk)**.** *Suppose the reference errors are i.i.d.* $\mathsf{Bern}(\alpha)$*. There is* $a_\alpha^{\mathrm{rw}} > 0$ *such that* $\mathbb{E}|S_n| \geq a_\alpha^{\mathrm{rw}}\sqrt{n}$ *for every* $n \geq 1$*. For every fixed* $a > 0$*,*

$$\lim_{n\to\infty} \mathbb{P}(S_n \geq a\sqrt{n}) = \lim_{n\to\infty} \mathbb{P}(-S_n \geq a\sqrt{n}) = 1 - \Phi(a/\sigma) > 0,$$

where Φ *is the standard normal distribution function.*

The first two lemmas are elementary facts specific to our comparison. The last uses the classical central limit theorem and convergence of moments under uniform integrability; we recall the short application below rather than prove these standard results.

B.1 Proof of Theorem A.1

Proof. Fix $\mathbb{P}$, C, and $C_{\mathbb{P}}^*$ as in the theorem. Apply Lemma B.2 to $e = e^C$ and $e^* = e^{C_{\mathbb{P}}^*}$. By Lemma B.1, S_t has the law in Lemma B.3, independently of the choices of law, algorithm, and oracle.

Step 1: two tight boundaries. Suppose $\underline{\beta} > 1/2$ and write $C = c_- + c_+$. On any block, the constrained path can move by at most $CT^{1-\underline{\beta}}$. Choose blocks long enough that the oracle's expected displacement is at least twice this amount:

$$A = \left(\frac{2C}{a_\alpha^{\mathrm{rw}}}\right)^2, \qquad n = \left\lceil AT^{2-2\underline{\beta}} \right\rceil, \qquad m = \lfloor T/n \rfloor.$$

For each of the m complete blocks, Lemmas B.2 and B.3 give expected disagreement at least $a_\alpha^{\mathrm{rw}}\sqrt{n}/2$. Thus

$$\mathbb{E}D_T(e, e^*) \geq \frac{a_\alpha^{\mathrm{rw}}}{2} m\sqrt{n}.$$

Since $2-2\underline{\beta} < 1$, we have $n \leq T/2$ for all sufficiently large T, including when $\underline{\beta} = 1$. Then $m \geq T/(2n)$ and $n \leq (A+1)T^{2-2\underline{\beta}}$, giving

$$\mathbb{E}D_T(e, e^*) \geq \frac{a_\alpha^{\mathrm{rw}} T}{4\sqrt{n}} \geq \frac{a_\alpha^{\mathrm{rw}}}{4\sqrt{A+1}} T^{\underline{\beta}}. \tag{40}$$

Step 2: at least one tight boundary. Suppose $\underline{\beta} \leq 1/2 \leq \bar{\beta}$. If $b \geq 1/2$, then $U_T \leq c_+\sqrt{T}$. On $\{S_T \geq (c_+ + 1)\sqrt{T}\}$, (39) gives $D_T(e, e^*) \geq \sqrt{T}$. Consequently,

$$\mathbb{E}D_T(e, e^*) \geq \sqrt{T}\,\mathbb{P}\{S_T \geq (c_+ + 1)\sqrt{T}\} \gtrsim \sqrt{T}$$

by Lemma B.3. If $d \geq 1/2$, apply the same argument to $-S_T$ and c_-.

Step 3: two wide boundaries. Suppose $\bar{\beta} < 1/2$. Since $T_0 < \infty$, at time T_0 the constraint excludes either the all-one or all-zero error path. The oracle follows these paths with probabilities α^{T_0} and $(1-\alpha)^{T_0}$, respectively. On the excluded path, the constrained algorithm must disagree at least once. As $D_T(e, e^*)$ is nondecreasing,

$$\mathbb{E}D_T(e, e^*) \geq \nu^{T_0} > 0 \qquad (T \geq T_0).$$

Together, Steps 1–3 prove (33), with constants and starting horizons depending only on the constraint parameters and α.

Step 4: add the learning cost. Now let $C \in \mathcal{A}_{b,d,c_-,c_+}$ be distribution-independent. The pointwise bound survives the infimum over exact oracles and the supremum over Q, so $R_T(C) \geq c\, g_{b,d}(T)$ for large T. Also $R_T(C) \geq R_T^0$ because $C \in \mathcal{A}$. Taking the larger of these two bounds gives

$$R_T(C) \geq \frac{\min\{1, c\}}{2}\left(R_T^0 + g_{b,d}(T)\right).$$

Finally, Lemma A.1 gives $R_T^0 \asymp \sqrt{T}$ for the class and oracle family of the main text. This proves the theorem. □

B.2 Proof of Lemma B.1

Proof. Because the oracle is deterministic and online, $e_1^*, \ldots, e_{t-1}^*$ are measurable with respect to $\mathcal{G}_{t-1}$. The exact oracle condition gives $\mathbb{P}(e_t^* = 1 \mid \mathcal{G}_{t-1}) = \alpha$ almost surely. For any $a_1, \ldots, a_t \in \{0, 1\}$, conditioning on $\mathcal{G}_{t-1}$ therefore gives

$$\mathbb{P}(e_1^* = a_1, \ldots, e_t^* = a_t) = \alpha^{a_t}(1-\alpha)^{1-a_t}\mathbb{P}(e_1^* = a_1, \ldots, e_{t-1}^* = a_{t-1}).$$

Iteration proves the product Bernoulli law. In fact, this argument only needs exact instantaneous training-conditional coverage, not independence of the original observations. □

B.3 Proof of Lemma B.2

Proof. For $0 \leq r < r' \leq T$, the triangle inequality gives

$$\sum_{t=r+1}^{r'} |e_t - e_t^*| \geq \left|(S_{r'} - S_r) - (\mathtt{E}_{r'} - \mathtt{E}_r)\right|.$$

Both $\mathtt{E}_r$ and $\mathtt{E}_{r'}$ lie in $[-c_- T^{1-\underline{\beta}}, c_+ T^{1-\underline{\beta}}]$; this includes $\mathtt{E}_0 = 0$. Their difference has absolute value at most $(c_- + c_+)T^{1-\underline{\beta}}$, proving (38). At $r = 0$, the bounds $\mathtt{E}_T \leq U_T$ and $\mathtt{E}_T \geq -L_T$ give $D_T(e, e^*) \geq S_T - U_T$ and $D_T(e, e^*) \geq -S_T - L_T$. Since $D_T(e, e^*) \geq 0$, the positive-part bounds follow. □

B.4 Justification of Lemma B.3

The classical i.i.d. central limit theorem gives $S_n/\sqrt{n} \Rightarrow N(0, \sigma^2)$. Also $\mathbb{E}(S_n/\sqrt{n})^2 = \sigma^2$, so $|S_n|/\sqrt{n}$ is uniformly integrable. Convergence of first moments therefore gives

$$\frac{\mathbb{E}|S_n|}{\sqrt{n}} \longrightarrow \sigma\sqrt{2/\pi} > 0;$$

see Van der Vaart (2000, Chapter 2, Theorem 2.20) for these standard convergence results. Every finite n also has $\mathbb{E}|S_n| > 0$, so the infimum of these normalized expectations is positive. Finally, convergence of the two tail probabilities follows from the central limit theorem because the normal distribution has no atoms.

C Proof of the truncation upper bound

This section proves Theorem A.2, and we first present the following lemmas. In the notation of Appendix A.1, take $(u, v) = (\widehat{e}, e)$, the truncated output and its input.

Lemma C.1 (Feasibility and stability)**.** *For every binary input e, truncation satisfies*

$$-L_t \le \mathtt{E}_t \le U_t \qquad \text{for every } t \ge 1.$$

For every two binary inputs e, e' and every $T \ge 1$,

$$D_T\big(\Gamma_{\mathrm{asym}}(e), \Gamma_{\mathrm{asym}}(e')\big) \le 2D_T(e, e'). \tag{41}$$

Lemma C.2 (Expected cost on Bernoulli errors)**.** *Let the input errors e_t be i.i.d. $\mathsf{Bern}(\alpha)$. There is a finite constant K such that*

$$\mathbb{E}D_T\big(\Gamma_{\mathrm{asym}}(e), e\big) \le \mathbb{E}N_T \le Kg_{b,d}(T) \qquad \text{for every } T \ge 1. \tag{42}$$

If $\bar{\beta} < 1/2$, then $\mathbb{E}N_\infty < \infty$.

Lemma C.1 shows that truncating two nearby error sequences keeps them nearby. Lemma C.2 bounds the additional disagreement needed to truncate an exact oracle, independently of its data law. We first combine these facts and then prove the lemmas.

C.1 Proof of Theorem A.2

Proof. Compare the baseline and oracle on the same observations.

Step 1: enforce the constraint. The first part of Lemma C.1 applies to every realization of e^0. Hence $\Gamma_{\mathrm{asym}} \circ C^0$ satisfies the asymmetric coverage constraint on every data sequence.

Step 2: compare the two truncated sequences. By the triangle inequality and (41),

$$\begin{aligned}
\ell_T(\Gamma_{\mathrm{asym}} \circ C^0, C^*_{\mathbb{P}}) &= D_T\big(\Gamma_{\mathrm{asym}}(e^0), e^*\big) \\
&\le D_T\big(\Gamma_{\mathrm{asym}}(e^0), \Gamma_{\mathrm{asym}}(e^*)\big) + D_T\big(\Gamma_{\mathrm{asym}}(e^*), e^*\big) \\
&\le 2\ell_T(C^0, C^*_{\mathbb{P}}) + D_T\big(\Gamma_{\mathrm{asym}}(e^*), e^*\big).
\end{aligned}$$

This comparison is pathwise and requires no independence between the baseline and the oracle.

Step 3: bound the oracle's truncation cost. By Lemma B.1, e^* is i.i.d. $\mathsf{Bern}(\alpha)$. Taking expectations and applying Lemma C.2 gives

$$\mathbb{E}_{\mathbb{P}}\ell_T(\Gamma_{\mathrm{asym}} \circ C^0, C^*_{\mathbb{P}}) \le 2\mathbb{E}_{\mathbb{P}}\ell_T(C^0, C^*_{\mathbb{P}}) + Kg_{b,d}(T) \qquad (T \ge 1).$$

The constant is independent of the baseline, the oracle, and the law, as required. □

C.2 Proof of Lemma C.1

Proof. We prove the lemma in the following steps.

Step 1: write the allowed error counts. Let $K_t = \sum_{r \le t} \widehat{e}_r$, with $K_0 = 0$, and define

$$a_t = \lceil \alpha t - L_t \rceil, \qquad z_t = \lfloor \alpha t + U_t \rfloor \qquad (t \ge 1), \qquad a_0 = z_0 = 0.$$

The coverage constraint is equivalent to $a_t \le K_t \le z_t$. The interval $[\alpha t - L_t, \alpha t + U_t]$ has length $L_t + U_t \ge c_- + c_+ \ge 1$, so $a_t \le z_t$. Also, for every $t \ge 1$,

$$z_t \ge z_{t-1}, \qquad a_t \le a_{t-1} + 1. \tag{43}$$

For $t \ge 2$, these inequalities follow because U_t is nondecreasing and $\alpha t - L_t$ increases by at most $\alpha < 1$ in one round. For $t = 1$, they follow from $z_1 \ge 0$ and $a_1 \le 1$.

Step 2: identify truncation with clipping. Define $\mathrm{clip}_t(x) = \min\{z_t, \max\{a_t, x\}\}$. We prove by induction that

$$K_t = \mathrm{clip}_t(K_{t-1} + e_t), \qquad a_t \le K_t \le z_t. \tag{44}$$

Assume $a_{t-1} \le K_{t-1} \le z_{t-1}$. The upper trigger in (30) is equivalent to $K_{t-1} + 1 > z_t$. By (43), this can occur only when $K_{t-1} = z_t$ and forcing error zero then gives the clipped value. The lower trigger is equivalent to $K_{t-1} < a_t$. Again, by (43), this can occur only when $K_{t-1} + 1 = a_t$, which forces the error to be 1, giving the clipped value. Both triggers cannot hold simultaneously, since it would imply $z_t = K_{t-1} < a_t$. If neither holds, both K_{t-1} and $K_{t-1} + 1$ lie in $[a_t, z_t]$, and the input error is unchanged. This proves (44) and feasibility.

Step 3: change one input at a time. Suppose first that e and e' differ only at round r, with $e_r = 1$ and $e'_r = 0$. Let K_t, K'_t be their clipped counts. The function clip_t is nondecreasing and satisfies

$$0 \le \mathrm{clip}_t(x) - \mathrm{clip}_t(y) \le x - y \qquad (x \ge y).$$

Thus $K_t - K'_t$ is zero before r, is either zero or one at r, and is nonnegative and nonincreasing afterward because the later inputs agree. The gap, which takes integer values, can change at most twice: from zero to one at r, and from one back to zero at a later round. The difference of the output errors equals the increment of this gap, so $D_T(\Gamma_{\mathrm{asym}}(e), \Gamma_{\mathrm{asym}}(e')) \le 2$. The case $e_r = 0, e'_r = 1$ follows by exchanging the two inputs. For arbitrary inputs, replace their differing coordinates in the first T rounds, one at a time, and apply the triangle inequality. The first T outputs depend only on those coordinates, giving (41). □

C.3 Proof of Lemma C.2

Counts and three bounds. For the i.i.d. $\mathsf{Bern}(\alpha)$ input in Lemma C.2, the counts of Appendix A.1 give

$$D_T(\Gamma_{\mathrm{asym}}(e), e) = M_T^+ + M_T^- \le N_T, \qquad \mathtt{E}_t = S_t - M_t^+ + M_t^-. \tag{45}$$

The trigger at round t depends only on $e_1, \dots, e_{t-1}$ and is independent of e_t. It follows that

$$\mathbb{E}M_T^+ = \alpha \mathbb{E}N_T^+, \qquad \mathbb{E}M_T^- = (1-\alpha)\mathbb{E}N_T^-. \tag{46}$$

We use three bounds, depending on whether a boundary grows faster than $\sqrt{t}$.

Lemma C.3. *If $d < 1/2$, then $\mathbb{E}N_\infty^- < \infty$. If $b < 1/2$, then $\mathbb{E}N_\infty^+ < \infty$.*

Lemma C.4. *If $d < 1/2 \le b$, then $\mathbb{E}N_T^+ \le C\sqrt{T}$ for every $T \ge 1$. If $b < 1/2 \le d$, then $\mathbb{E}N_T^- \le C\sqrt{T}$ for every $T \ge 1$.*

Lemma C.5. *If $\underline{\beta} \ge 1/2$, then $\mathbb{E}N_T \le CT^{\underline{\beta}}$ for every $T \ge 1$.*

Proof of Lemma C.2. The first inequality in (42) follows from (45).

- If $\bar{\beta} < 1/2$, Lemma C.3 bounds $\mathbb{E}N_\infty^+$ and $\mathbb{E}N_\infty^-$, so $\mathbb{E}N_T$ is bounded uniformly in T.
- If $\underline{\beta} < 1/2 \le \bar{\beta}$, Lemmas C.3 and C.4 give $\mathbb{E}N_T \le C\sqrt{T}$.
- If $\underline{\beta} = 1/2$, Lemma C.5 gives the same bound.
- Finally, if $\underline{\beta} > 1/2$, Lemma C.5 gives $\mathbb{E}N_T \le CT^{\underline{\beta}}$.

These are exactly the three rates in $g_{b,d}(T)$. □

C.4 Proof of Lemma C.3

Proof. Assume $d < 1/2$ and that round t is forced from below. Let $r < t$ be the most recent upper forced round, or set $r = 0$ if there is none. If $r > 0$, its upper trigger gives $\mathtt{E}_{r-1} + 1 - \alpha > U_r \ge 0$, and hence $\mathtt{E}_r = \mathtt{E}_{r-1} - \alpha > -1$. If $r = 0$, the same lower bound holds because $\mathtt{E}_0 = 0$. There are no upper changes between $r+1$ and $t-1$, so (45) gives

$$\mathtt{E}_{t-1} \ge \mathtt{E}_r + S_{t-1} - S_r > -1 + S_{t-1} - S_r.$$

The lower trigger gives $\mathtt{E}_{t-1} < -L_t + \alpha$. Thus some $0 \le r < t$ satisfies

$$S_{t-1} - S_r < -L_t + \alpha + 1.$$

For all sufficiently large t, $L_t \ge 2(\alpha+1)$, so the right-hand side is at most $-L_t/2$. The choice $r = t-1$ is then impossible. For each fixed $r \le t-2$, the increment is the sum of $t-1-r$ independent, centered variables, each taking values in an interval of length 1. Then, Hoeffding's inequality (Hoeffding, 1963, Theorem 2, Eq. (2.6), p. 16), applied to the negative increment, and a union bound give

$$\begin{aligned}\mathbb{P}\{t \text{ is forced from below}\} &\le \sum_{r=0}^{t-2} \exp\left(-\frac{L_t^2}{2(t-1-r)}\right) \\ &\le t \exp\left(-\frac{c_-^2}{2} t^{1-2d}\right).\end{aligned}$$

Since $1 - 2d > 0$, the last expression is summable in t. Summing the trigger probabilities, including the finitely many initial rounds, proves $\mathbb{E}N_\infty^- < \infty$.

For the other side, replace e_t by $1 - e_t$, α by $1 - \alpha$, and exchange (U_t, L_t). The two triggers are mutually exclusive by Lemma C.1. Induction in (30) therefore shows that the transformed output is $1 - \widehat{e}_t$ and its centered sum is $-\mathtt{E}_t$. Its lower forced rounds are the original upper forced rounds. The result already proved gives $\mathbb{E}N_\infty^+ < \infty$ when $b < 1/2$. □

C.5 Proof of Lemma C.4

Proof. Assume $d < 1/2 \le b$. We bound the upper changes by the maximum of the free walk and the changes in the opposite direction. If $M_T^+ > 0$, let $t \le T$ be its last upper change. Then $M_T^+ = M_{t-1}^+ + 1$, while the upper trigger gives $\mathtt{E}_{t-1} > U_t - (1 - \alpha) \ge -1$. Using (45), we obtain

$$M_T^+ = S_{t-1} + M_{t-1}^- - \mathtt{E}_{t-1} + 1 < S_{t-1} + M_T^- + 2.$$

The case $M_T^+ = 0$ is immediate, so in all cases

$$M_T^+ \le \max_{0 \le r \le T} |S_r| + M_T^- + 2. \tag{47}$$

The free walk is a square-integrable martingale with $\mathbb{E}S_T^2 = \sigma^2 T$. Doob's L^2 maximal inequality (Durrett, 2019, Theorem 4.4.4, p. 236) and Cauchy–Schwarz give

$$\mathbb{E} \max_{0 \le r \le T} |S_r| \le \left(\mathbb{E} \max_{0 \le r \le T} |S_r|^2 \right)^{1/2} \le 2\sigma\sqrt{T}.$$

Also, $M_T^- \le N_\infty^-$ and Lemma C.3 gives $\mathbb{E}N_\infty^- < \infty$. Taking expectations in (47) and using (46),

$$\mathbb{E}N_T^+ \le \frac{2\sigma\sqrt{T} + \mathbb{E}N_\infty^- + 2}{\alpha} \le C\sqrt{T} \qquad (T \ge 1).$$

Reversing the upper and lower boundaries, as in the proof of Lemma C.3, gives the other case. □

C.6 Proof of Lemma C.5

Proof. If $\underline{\beta} = 1$, then $N_T \le T$ proves the claim. Assume $1/2 \le \underline{\beta} < 1$.

Step 1: bound the width and the midpoint's movement. Define

$$m_t = \frac{U_t - L_t}{2}, \qquad w_t = \frac{U_t + L_t}{2}, \qquad H_t = \mathtt{E}_t - m_t, \qquad \rho_t = m_t - m_{t-1}.$$

Set $m_0 = w_0 = H_0 = 0$. By feasibility, $|H_t| \le w_t$. There are constants $0 < c \le C < \infty$ such that

$$c t^{1-\underline{\beta}} \le w_t \le C t^{1-\underline{\beta}}, \qquad |\rho_t| \le C t^{-\underline{\beta}}, \qquad w_{t-1}|\rho_t| \le C \qquad (t \ge 1). \tag{48}$$

Indeed, one of b, d equals $\underline{\beta}$, so the first bound holds with lower constant $\min\{c_-, c_+\}/2$ and upper constant $(c_- + c_+)/2$. For $t \ge 2$ and $\gamma \in [0, 1]$, the mean value theorem gives

$$0 \le t^\gamma - (t-1)^\gamma \le 2t^{\gamma - 1}.$$

Applying this to $1 - b$ and $1 - d$ proves the second bound, provided the constant is increased to cover $t = 1$. The first two bounds give $w_{t-1}|\rho_t| \le C t^{1-2\underline{\beta}} \le C$. Also, w_t is nondecreasing and tends to infinity, while ρ_t tends to zero. Choose $t_\star$ so that

$$w_t \ge 2, \qquad |\rho_t| \le \nu/2 \qquad (t \ge t_\star). \tag{49}$$

Step 2: bound the increase on unforced rounds. Let I_t^0 indicate that neither trigger holds. On this event,

$$H_t = H_{t-1} + e_t - \alpha - \rho_t.$$

Both I_t^0 and H_{t-1} are determined by the previous input errors, whereas $e_t - \alpha$ is independent of them, centered, and has variance σ^2. Thus

$$\mathbb{E}\big[(H_t^2 - H_{t-1}^2)I_t^0\big] = \mathbb{E}\big[(-2H_{t-1}\rho_t + \sigma^2 + \rho_t^2)I_t^0\big] \le C_1, \tag{50}$$

where the last bound follows from $|H_{t-1}| \le w_{t-1}$ and (48).

Step 3: bound the decrease on forced rounds. Fix $t \ge t_\star$ and put $h = \mathtt{E}_{t-1} - m_t$, so $H_{t-1} = h + \rho_t$. If the upper trigger holds, then $h > w_t - (1-\alpha)$ and $H_t = h - \alpha$. Consequently,

$$H_t^2 - H_{t-1}^2 = -(\alpha + \rho_t)(2h - \alpha + \rho_t) \le -\frac{\alpha}{2} w_t.$$

Here (49) gives $\alpha + \rho_t \ge \alpha/2$ and $2h - \alpha + \rho_t \ge 2w_t - 2 \ge w_t$. If the lower trigger holds, then $h < -w_t + \alpha$ and $H_t = h + 1 - \alpha$. In this case,

$$H_t^2 - H_{t-1}^2 = (1 - \alpha - \rho_t)(2h + 1 - \alpha + \rho_t) \le -\frac{1-\alpha}{2} w_t,$$

because the first factor is at least $(1-\alpha)/2$ and the second is at most $-2w_t + 2 \le -w_t$. Writing $I_t = 1 - I_t^0$ for the indicator of a forced round, both cases give

$$(H_t^2 - H_{t-1}^2)I_t \le -\frac{\nu}{2} w_t I_t \qquad (t \ge t_\star). \tag{51}$$

Step 4: count the forced rounds in a block. For an integer $n \ge t_\star$, sum (50) and the expectation of (51) over $n < t \le 2n$. Since $w_t \ge w_n$ on the block,

$$\mathbb{E}H_{2n}^2 - \mathbb{E}H_n^2 \le C_1 n - \frac{\nu}{2} w_n \mathbb{E}(N_{2n} - N_n).$$

The left side is at least $-w_n^2$. Rearranging and using (48),

$$\mathbb{E}(N_{2n} - N_n) \le \frac{2}{\nu}\left(\frac{C_1 n}{w_n} + w_n\right) \le C\left(n^{\underline{\beta}} + n^{1-\underline{\beta}}\right) \le C n^{\underline{\beta}}.$$

The last inequality uses $\underline{\beta} \ge 1/2$.

Step 5: sum the blocks. Choose $j_0 \ge 0$ with $2^{j_0} \ge t_\star$. For $i > j_0$, Step 4 and $N_{2^{j_0}} \le 2^{j_0}$ give

$$\mathbb{E}N_{2^i} \le 2^{j_0} + C \sum_{j=j_0}^{i-1} 2^{j\underline{\beta}} \le C' 2^{i\underline{\beta}}.$$

After increasing C', this holds also for $0 \le i \le j_0$. For any $T \ge 1$, choose $i \ge 0$ with $2^{i-1} < T \le 2^i$. Monotonicity of N_T then gives $\mathbb{E}N_T \le \mathbb{E}N_{2^i} \le C'(2T)^{\underline{\beta}}$, which proves the claim. □

D Unconstrained minimax rate

We use the notation from Appendix A.1. We first consider the following lemmas. For a fixed distribution, the first lemma turns disagreement with the closest exact oracle into a sum of instantaneous training-conditional errors.

Lemma D.1. *Fix $\mathbb{P} \in \mathcal{Q}$ and a deterministic measurable online algorithm C. There is a deterministic measurable oracle $C_{\mathbb{P}}^{\mathrm{rep}} \in \mathcal{O}_\alpha(\mathbb{P})$ such that, for every $T \geq 1$,*

$$r_T(C; \mathbb{P}) = \mathbb{E}_{\mathbb{P}} \ell_T(C, C_{\mathbb{P}}^{\mathrm{rep}}) = \sum_{t=1}^{T} \mathbb{E}_{\mathbb{P}} |p_t^C - \alpha|.$$

The same oracle works for every horizon.

Let $B_t \sim \mathrm{Binomial}(t, \alpha)$. For t i.i.d. observations, B_t has the distribution of the number of miscoverages by the population score oracle. If the realized count is b, declaring miscoverage on the k largest scores disagrees with the population oracle on exactly $|b - k|$ observations. The following lemma shows why the same count yields a lower bound even for algorithms that need not threshold scores.

Lemma D.2 (Comparison with a binomial count)**.** *For every $C \in \mathcal{A}$ and every $T \geq 1$,*

$$R_T(C) \geq \sum_{t=1}^{T} \frac{1}{t} \min_{k \in \{0,\dots,t\}} \mathbb{E}|B_t - k|.$$

The last lemma supplies the upper bound using the unsmoothed conformal set in the main text.

Lemma D.3 (Conformal disagreement with the population oracle)**.** *For every $\mathbb{P} \in \mathcal{Q}$ and every $t \geq 1$,*

$$\mathbb{E}_{\mathbb{P}} \left| \mathbb{P}\{Y_t \in C_t^{\mathrm{CP}}(X_t) \mid \mathcal{G}_{t-1}\} - (1 - \alpha) \right| \leq \frac{1}{2\sqrt{t+1}} + \frac{1}{t}.$$

The population score oracle attains the smallest expected disagreement with C^{CP} among all exact oracles. Consequently, for every $T \geq 1$,

$$R_T(C^{\mathrm{CP}}) \leq 3\sqrt{T}.$$

D.1 Proof of Lemma A.1

Proof. We first match the binomial lower bound by a prediction set with an integer cutoff at each round. The rate follows from the size of the binomial fluctuations.

Step 1. Match the lower bound at every horizon. Fix deterministic integers $0 \leq k_t \leq t$ and form the prediction set

$$C_t^{(k)}(x) = \left\{ y : 1 + \sum_{r<t} \mathbf{1}\{s_r \geq s(x,y)\} > k_t \right\}.$$

Under any $\mathbb{P} \in \mathcal{Q}$, the t scores are almost surely distinct. Hold out each observation in turn and use the other $t - 1$ scores in this formula. The algorithm declares miscoverage on the k_t largest

scores. The population oracle declares miscoverage on the scores above $q_{\mathbb{P},s}$, whose number has distribution B_t. The two subsets are nested. By exchangeability, the disagreement probability at round t equals the expected fraction of these t decisions that disagree. Thus, for every $\mathbb{P} \in \mathcal{Q}$,

$$\mathbb{E}_{\mathbb{P}}\ell_T(C^{(k)}, C^*_{\mathbb{P},s}) = \sum_{t=1}^{T} \frac{\mathbb{E}|B_t - k_t|}{t}.$$

Choose an integer median of B_t for k_t. It minimizes the summand because $\mathbb{E}|B_t - (k+1)| - \mathbb{E}|B_t - k| = 2\Pr(B_t \le k) - 1$. Comparing with the population oracle gives an upper bound uniform over $\mathcal{Q}$. Together with Lemma D.2, this gives the identity

$$R_T^0 = \sum_{t=1}^{T} \frac{1}{t} \min_{k\in\{0,\dots,t\}} \mathbb{E}|B_t - k|. \tag{52}$$

Step 2. Bound the binomial fluctuations. Put $X = B_t - t\alpha$ and let $\xi \sim \text{Bernoulli}(\alpha)$. Independence and centering give

$$\mathbb{E}X^2 = t\sigma^2, \qquad \mathbb{E}X^4 = t\mathbb{E}(\xi-\alpha)^4 + 3t(t-1)\sigma^4 \le 2t^2\sigma^2,$$

using $|\xi - \alpha| \le 1$ and $\sigma^2 \le 1/4$. For any k, $|t\alpha - k| \le \mathbb{E}|B_t - k|$ by Jensen's inequality. The triangle inequality and Hölder's inequality therefore give

$$\mathbb{E}|B_t - k| \ge \tfrac{1}{2}\mathbb{E}|X| \ge \frac{(\mathbb{E}X^2)^{3/2}}{2(\mathbb{E}X^4)^{1/2}} \ge \frac{\sigma^2}{2\sqrt{2}}\sqrt{t}.$$

Since $\sigma^2 \ge \nu^2$ and $\sum_{t=1}^{T} t^{-1/2} \ge \sqrt{T}$, (52) implies $R_T^0 \ge (\nu^2/32)\sqrt{T}$.

The usual split conformal set in (2) corresponds to $k_t = \lfloor \alpha t \rfloor$. For this choice, $\mathbb{E}|B_t - k_t| \le \sigma\sqrt{t} + 1$. Since $\sigma \le 1/2$ and $\sum_{t=1}^{T} t^{-1/2} \le 2\sqrt{T}$, Step 1 gives $R_T^0 \le R_T(C^{\text{CP}}) \le 3\sqrt{T}$. □

D.2 Proof of Lemma D.1

Proof. We present the lemma in the following steps.

Step 1. Bound the disagreement of any exact oracle. For the error e_t^* of any exact oracle, conditional Jensen's inequality gives

$$\mathbb{E}_{\mathbb{P}}[|\mathsf{e}_t^C - e_t^*| \mid \mathcal{G}_{t-1}] \ge \left|\mathbb{E}_{\mathbb{P}}[\mathsf{e}_t^C - e_t^* \mid \mathcal{G}_{t-1}]\right| = |p_t^C - \alpha|.$$

It suffices to attain this bound at each round with the same oracle.

Step 2. Construct the oracle and compute its loss. Fix a round t and history h. For a fresh $Z \sim \mu_{\mathbb{P}}$, put

$$A_h(z) = e_t(h,z) + F(s(z)), \qquad G_h(a) = \int \mathbf{1}\{A_h(z') \le a\}\, d\mu_{\mathbb{P}}(z').$$

Since $F(s(Z))$ is uniform on $(0,1)$, the algorithm covers exactly when $A_h(Z) \le 1$, almost surely. Also $\{A_h(Z) = a\} \subseteq \{F(s(Z)) \in \{a, a-1\}\}$, so $A_h(Z)$ has no atoms and G_h is continuous.

Define

$$C^{\mathrm{rep}}_{\mathbb{P},t}(x;h) = \{y : G_h(A_h(x,y)) \le 1-\alpha\}.$$

Given h, the probability integral transform makes $G_h(A_h(Z))$ uniformly distributed, so the oracle has instantaneous training-conditional coverage of $1-\alpha$.

Both coverage events are lower intervals in $A_h(Z)$ and hence are nested. If $p(h) = \int e_t(h,z)\, d\mu_{\mathbb{P}}(z)$, their probabilities are $G_h(1) = 1 - p(h)$ and $1-\alpha$. Their conditional disagreement is therefore

$$\big|G_h(1) - (1-\alpha)\big| = |p(h) - \alpha|.$$

At the observed history $p(h) = p_t^C$ almost surely. Taking expectations and summing attains the lower bound from Step 1 for every T, since the construction is independent of the horizon. □

D.3 Proof of Lemma D.2

Proof. Fix C and T. We first compute the calibration error of a decision that depends only on the relative order of the scores. We then choose one distribution under which such decisions approximate C at every $t \le T$.

Step 1. Count disagreements for a rule based on ranks. Let $V_1, \dots, V_t$ be i.i.d. uniform on $(0,1)$. Write R_i for the rank of V_i among all t values, and let $B_t = \sum_{i=1}^t \mathbf{1}\{V_i \le \alpha\}$. For $0 \le k \le t$, the two indicators below select the first k and the first B_t ranks. Hence

$$\sum_{i=1}^{t} \big|\mathbf{1}\{R_i \le k\} - \mathbf{1}\{V_i \le \alpha\}\big| = |B_t - k|.$$

Let $V_{(1)}, \dots, V_{(t-1)}$ be the sorted past values, with $V_{(0)} = 0$ and $V_{(t)} = 1$. Since $R_t \le k$ exactly when $V_t \le V_{(k)}$ almost surely, conditioning on the past and then using exchangeability gives

$$\mathbb{E}|V_{(k)} - \alpha| = \mathbb{E}\big|\mathbf{1}\{R_t \le k\} - \mathbf{1}\{V_t \le \alpha\}\big| = \frac{1}{t}\mathbb{E}|B_t - k|.$$

Fix the past order ρ. A rank-based rule declares miscoverage on some set $J(\rho)$. Given the past values, its miscoverage probability is

$$q_t = \sum_{j \in J(\rho)} G_j, \qquad G_j = V_{(j)} - V_{(j-1)}.$$

Given ρ, the sorted values have constant density $(t-1)!$, so the gaps are uniform on the simplex $\{g_j \ge 0,\ \sum_j g_j = 1\}$ and are exchangeable. For $t = 1$ the sole gap is $G_1 = 1$. For $k = |J(\rho)|$, the sum has the same distribution as $\sum_{j=1}^k G_j = V_{(k)}$. Thus for each ρ,

$$\mathbb{E}[|q_t - \alpha| \mid \rho] = \frac{1}{t}\mathbb{E}|B_t - k| \ge \frac{1}{t} \min_{0 \le j \le t} \mathbb{E}|B_t - j|.$$

Step 2. Approximate C using only score ranks. Fix an integer $m \geq T$. We will give equal mass to m cells ordered by score. On each tuple of distinct cells, a function of the score ranks will approximate the miscoverage indicator $e_t = \mathbf{1}\{Y_t \notin C_t(X_t)\}$ with error probability below $1/m$.

Choose the cells. Let N be large enough. Take μ_0 with $\mu_0^{\otimes \mathbb{N}} \in \boldsymbol{Q}$. Since the score is nonatomic, we can split $\mathcal{Z}$ into N consecutive score bins $I_1, \ldots, I_N$, each of mass $1/N$. For each $t \leq T$, approximate $\{e_t = 1\}$ in product measure by a finite union of rectangles with indicator f_t. Refine the bins by the sides of all these rectangles. Then f_t is constant on each product cell, and we can make the errors $E_t = \{e_t \neq f_t\}$ small enough that

$$\sum_{t=1}^{T} N^t \mu_0^{\otimes t}(E_t) < 1/m.$$

Independently choose one partition cell D_i in each bin I_i, with $\Pr(D_i = D) = N\mu_0(D)$ for each cell $D \subseteq I_i$. Set $\mu_i = \mu_0(\cdot \mid D_i)$. Averaging μ_i over this choice gives $\mu_0(\cdot \mid I_i)$. For each ordered tuple $\mathbf{i} = (i_1, \ldots, i_t)$ of distinct labels, write $\mu_{\mathbf{i}} = \bigotimes_{r=1}^{t} \mu_{i_r}$ and $I_{\mathbf{i}} = \prod_{r=1}^{t} I_{i_r}$. Independence and disjointness of the bin products at each t give

$$\mathbb{E} \sum_{t=1}^{T} \sum_{\mathbf{i}} \mu_{\mathbf{i}}(E_t) = \sum_{t=1}^{T} N^t \sum_{\mathbf{i}} \mu_0^{\otimes t}(E_t \cap I_{\mathbf{i}}) \leq \sum_{t=1}^{T} N^t \mu_0^{\otimes t}(E_t) < 1/m.$$

Fix a choice of cells for which the sum is below $1/m$. Then $\mu_{\mathbf{i}}(E_t) < 1/m$ for every tuple in the sum.

Select cells on which ranks suffice. The color of a set of t labels records which orders give $f_t = 1$. For $i_1 < \cdots < i_t$, define

$$c_t(i_1, \ldots, i_t) = \left\{\pi \in S_t : \ f_t = 1 \text{ on } D_{i_{\pi(1)}} \times \cdots \times D_{i_{\pi(t)}}\right\},$$

where S_t is the set of permutations of $\{1, \ldots, t\}$. There are at most $2^{t!}$ colors. Applying the finite Ramsey theorem (Ramsey, 1930, Theorem B, p. 267) successively for $t = 1, \ldots, T$ gives a set H of m labels such that, for each t, every t-subset of H has the same color Π_t.

For distinct scores, let $\pi(r)$ be the rank of s_r among $s_1, \ldots, s_t$ and set $\widetilde{e}_t = \mathbf{1}\{\pi \in \Pi_t\}$. On distinct cells from H, the scores have the same relative order as their cell labels. Thus $f_t = \widetilde{e}_t$ on each such product cell, and

$$\mu_{\mathbf{i}}\{e_t \neq \widetilde{e}_t\} = \mu_{\mathbf{i}}(E_t) < 1/m \qquad (i_1, \ldots, i_t \text{ distinct in } H).$$

Step 3. Sum the bound under one distribution. Relabel the retained cells $1, \ldots, m$ and set $\mu = m^{-1} \sum_{i=1}^{m} \mu_i$. Since $\mu \ll \mu_0$, the score remains nonatomic and $\mathbb{P} = \mu^{\otimes \mathbb{N}} \in \boldsymbol{Q}$. Set $\widetilde{p}_t = \mathbb{E}_{\mathbb{P}}[\widetilde{e}_t \mid \mathcal{G}_{t-1}]$. The cell labels are i.i.d. uniform on $\{1, \ldots, m\}$, so Step 2 and a union bound give

$$\mathbb{E}_{\mathbb{P}}|p_t^C - \widetilde{p}_t| \leq \mathbb{P}(\mathsf{e}_t^C \neq \widetilde{e}_t) \leq \frac{1}{m} + \mathbb{P}(\text{two of the first } t \text{ labels coincide}) \leq \frac{1 + \binom{t}{2}}{m}.$$

The uniforms $V_r = 1 - F(s_r)$ rank the scores in decreasing order, and V_t is independent of $\mathcal{G}_{t-1}$. Thus $\widetilde{p}_t$ is the gap sum in Step 1. By Lemma D.1 and the triangle inequality,

$$R_T(C) \geq r_T(C; \mathbb{P}) = \sum_{t=1}^{T} \mathbb{E}_{\mathbb{P}}|p_t^C - \alpha| \geq \sum_{t=1}^{T} \frac{1}{t} \min_{0 \leq k \leq t} \mathbb{E}|B_t - k| - \frac{1}{m} \sum_{t=1}^{T} \left(1 + \binom{t}{2}\right).$$

Let $m \to \infty$ with T fixed. □

D.4 Proof of Lemma D.3

Proof. Fix $\mathbb{P} \in \mathcal{Q}$. Let $k_t = \lceil (1-\alpha)t \rceil$. If $k_t \le t-1$, the unsmoothed conformal set in (2) is, up to a fresh-score null event, the set whose cutoff is the k_t-th smallest of the $t-1$ past scores.

The transformed scores $F(s(Z_r))$ are independent uniform random variables. Consequently, the instantaneous training-conditional coverage probability $Q_t = 1 - p_t^{\mathrm{CP}}$ of the conformal set has distribution $\mathrm{Beta}(k_t, t-k_t)$ when $k_t < t$, and equals one when $k_t = t$. The beta mean and variance, together with the case $k_t = t$, give

$$\mathbb{E}Q_t = \frac{k_t}{t}, \qquad \mathrm{Var}(Q_t) \le \frac{1}{4(t+1)}, \qquad \left| \frac{k_t}{t} - (1-\alpha) \right| \le \frac{1}{t}.$$

The triangle inequality and Cauchy–Schwarz prove

$$\mathbb{E}|Q_t - (1-\alpha)| \le \frac{1}{2\sqrt{t+1}} + \frac{1}{t}.$$

The conformal and population prediction sets are thresholds of the same score, with both cutoffs determined before the fresh outcome. Their conditional disagreement probability is therefore exactly $|Q_t - (1-\alpha)|$. Conditional Jensen's inequality gives this same quantity as a lower bound for every exact oracle, so the population oracle attains the infimum. Finally,

$$\sum_{t=1}^{T} \left(\frac{1}{2\sqrt{t+1}} + \frac{1}{t} \right) \le \frac{3}{2} \sum_{t=1}^{T} t^{-1/2} \le 3\sqrt{T}.$$

The bound is uniform over $\mathbb{P} \in \mathcal{Q}$, which proves the lemma. □

E Results with a relaxed oracle condition

An exact oracle condition may be impossible when the joint distribution has atoms. We therefore also consider relaxed oracles whose instantaneous training-conditional coverage is at least $1-\alpha$:

$$\mathbb{P}\{Y_t \in C^*_{\mathbb{P},t}(X_t) \mid \mathcal{G}_{t-1}\} \ge 1-\alpha \quad \text{almost surely for every } t \ge 1. \tag{oracle'}$$

Let $O^+_\alpha(\mathbb{P})$ be a nonempty family of such oracles, fixed independently of the algorithm and horizon. The population quantile set (12) satisfies this condition even when the score distribution has atoms. For a specified nonempty family $O^+_\alpha(\mathbb{P})$ satisfying (oracle'), write

$$R^+_T(C) = \sup_{\mathbb{P} \in \widetilde{\mathcal{Q}}} \inf_{C^*_{\mathbb{P}} \in O^+_\alpha(\mathbb{P})} \mathbb{E}_{\mathbb{P}} \ell_T(C, C^*_{\mathbb{P}}), \qquad R^{0,+}_T = \inf_{C \in \mathcal{A}} R^+_T(C).$$

For this relaxed formulation, every $C \in \mathcal{A}_{b,d,c_-,c_+}$ satisfies the lower-bound decomposition over the class $\mathcal{Q}$ of all i.i.d. laws:

$$\begin{aligned} &\sup_{\mathbb{P} \in \widetilde{\mathcal{Q}}} \inf_{C^*_{\mathbb{P}} \in O^+_\alpha(\mathbb{P})} \mathbb{E}_{\mathbb{P}} \ell_T(C, C^*_{\mathbb{P}}) \\ &\qquad \gtrsim \inf_{C^0 \in \mathcal{A}} \sup_{\mathbb{P} \in \widetilde{\mathcal{Q}}} \inf_{C^*_{\mathbb{P}} \in O^+_\alpha(\mathbb{P})} \mathbb{E}_{\mathbb{P}} \ell_T(C^0, C^*_{\mathbb{P}}) + \begin{cases} 0, & 0 \le \beta < 1/2, \\ T^\beta, & 1/2 \le \beta \le 1. \end{cases} \end{aligned} \tag{53}$$

The unconstrained term need not have order $\sqrt{T}$. For example, if $O_\alpha^+(\mathbb{P})$ contains all relaxed oracles, this term is zero: the algorithm that always predicts $\mathcal{Y}$ agrees perfectly with the oracle that does the same. A relaxed oracle may therefore be a poor target for statistical accuracy, since matching it need not yield informative predictions or coverage close to $1-\alpha$.

Formally, we have the following result.

Proposition E.1 (Lower bounds with relaxed oracles). *Every $C \in \mathcal{A}_{b,d,c_-,c_+}$ satisfies, for every $T \geq 1$,*

$$R_T^+(C) \geq \tfrac{1}{2}\left(R_T^{0,+} + (\alpha T - L_T)_+\right). \tag{54}$$

If $d > 0$, or if $d = 0$ and $c_- < \alpha$, the constrained minimax risk has order T. If the oracle family contains all relaxed oracles and $d = 0$, $c_- \geq \alpha$, the constrained minimax risk is zero.

Only the lower boundary forces a prediction set to miss a response that an always-covering oracle covers. This explains why $g_{b,d}$ is not the appropriate universal lower bound for relaxed oracles.

To prove the lower bound above, we first consider the following lemma.

Lemma E.1 (Disagreement under a point mass). *Fix any $z_0 = (x_0, y_0) \in \mathcal{X} \times \mathcal{Y}$ and let $\mathbb{P}_{\mathrm{pt}}$ be the i.i.d. law concentrated on z_0. For every $C \in \mathcal{A}_{b,d,c_-,c_+}$, every deterministic relaxed oracle $C_{\mathbb{P}}^*$ under $\mathbb{P}_{\mathrm{pt}}$, and every $T \geq 1$,*

$$\mathbb{E}_{\mathbb{P}_{\mathrm{pt}}} \ell_T(C, C_{\mathbb{P}}^*) = \sum_{t=1}^{T} \mathsf{e}_t^C \geq (\alpha T - L_T)_+.$$

E.1 Proof of Proposition E.1

Proof. We prove the proposition in the following steps.

Step 1. Combine the two lower bounds. Fix $C \in \mathcal{A}_{b,d,c_-,c_+}$ and $T \geq 1$. Since $C \in \mathcal{A}$, we have $R_T^+(C) \geq R_T^{0,+}$. Since $\mathbb{P}_{\mathrm{pt}} \in \widetilde{\mathcal{Q}}$, Lemma E.1 gives

$$R_T^+(C) \geq \inf_{C_{\mathbb{P}}^* \in O_\alpha^+(\mathbb{P}_{\mathrm{pt}})} \mathbb{E}_{\mathbb{P}_{\mathrm{pt}}} \ell_T(C, C_{\mathbb{P}}^*) \geq (\alpha T - L_T)_+.$$

Taking the larger of the bounds yields

$$R_T^+(C) \geq \max\{R_T^{0,+}, (\alpha T - L_T)_+\} \geq \tfrac{1}{2}\big(R_T^{0,+} + (\alpha T - L_T)_+\big),$$

which proves (54) for every horizon.

Step 2. Determine when the forced cost is linear. Since $L_T = c_- T^{1-d}$,

$$\frac{(\alpha T - L_T)_+}{T} = (\alpha - c_- T^{-d})_+ \longrightarrow \begin{cases} \alpha, & d > 0, \\ (\alpha - c_-)_+, & d = 0. \end{cases}$$

The limit is positive if $d > 0$ or if $d = 0$ and $c_- < \alpha$. Truncation supplies a feasible algorithm by Lemma C.1, and $\ell_T \leq T$ on every path. Hence in either case, for all sufficiently large T,

$$cT \leq \inf_{C \in \mathcal{A}_{b,d,c_-,c_+}} R_T^+(C) \leq T$$

for some $c > 0$.

Step 3. Attain zero risk when the lower boundary permits full coverage. Suppose $d = 0$, $c_- \geq \alpha$, and the comparison family contains all relaxed oracles. The algorithm $C_t^{\mathrm{all}}(x) = \mathcal{Y}$ satisfies

$$-L_t = -c_- t \leq \mathrm{E}_t^{C^{\mathrm{all}}} = -\alpha t \leq U_t \qquad (t \geq 1),$$

so it belongs to $\mathcal{A}_{b,d,c_-,c_+}$. Its instantaneous training-conditional coverage is one under every distribution, so it is also an admissible comparison oracle for each $\mathbb{P} \in \widetilde{\mathcal{Q}}$. Comparing the algorithm with itself gives

$$0 \leq \inf_{C \in \mathcal{A}_{b,d,c_-,c_+}} R_T^+(C) \leq R_T^+(C^{\mathrm{all}}) \leq \sup_{\mathbb{P} \in \widetilde{\mathcal{Q}}} \mathbb{E}_{\mathbb{P}} \ell_T(C^{\mathrm{all}}, C^{\mathrm{all}}) = 0.$$

□

E.2 Proof of Lemma E.1

Proof. On the constant sequence $Z_t = z_0$, both the history $h = (z_0, \ldots, z_0)$ through round $t-1$ and the fresh response y_0 are fixed. For a deterministic relaxed oracle, (oracle′) becomes

$$1 - \alpha \leq \mathbb{P}_{\mathrm{pt}}\{Y_t \in C_t^*(X_t) \mid \mathcal{G}_{t-1}\} = \mathbf{1}\{y_0 \in C_t^*(x_0; h)\} \in \{0, 1\}.$$

Since $1 - \alpha > 0$, the indicator is one at every round. The oracle's error sequence is therefore zero, and on the unique data path

$$\mathbb{E}_{\mathbb{P}_{\mathrm{pt}}} \ell_T(C, C_{\mathbb{P}}^*) = \sum_{t=1}^{T} |\mathrm{e}_t^C - 0| = \alpha T + \mathrm{E}_T^C \geq \alpha T - L_T.$$

The error count is also nonnegative, giving the positive part in the claim. □

Take $b = d = \beta$ and $c_- = c_+ = \kappa$. If $\beta \geq 1/2$, then $\beta > 0$, so Step 2 gives $R_T^+(C) \gtrsim T \geq T^\beta$ for every constrained algorithm. Together with $R_T^+(C) \geq R_T^{0,+}$, this proves the second case of (53). If $\beta < 1/2$, that display adds zero to the unconstrained risk, so class inclusion alone proves its first case. In fact, the point-mass argument gives the stronger linear lower bound whenever $\beta > 0$. For asymmetric constraints, a tight upper bound alone does not yield this lower bound: Step 3 shows that zero risk is possible if the lower bound permits the always-covering algorithm.

F Proofs for ACI

This section proves Proposition 3.1 and corollary 3.1. Here $s_t \in [0, B]$ and $\mathrm{e}_t = \mathbf{1}\{s_t > q_t\}$ is the ACI error. The update (22) is $q_{t+1} = q_t + \eta_t(\mathrm{e}_t - \alpha)$, where $q_1 \in [0, B]$ and the positive step sizes are deterministic. Write $\bar{\eta}_t = \max_{1 \leq r \leq t} \eta_r$, with $\bar{\eta}_0 = 0$. The following two bounds hold for every score sequence, without any distributional assumption. We cite their proofs from Angelopoulos et al. (2024a).

Lemma F.1 (Angelopoulos et al., 2024a, Lemma 1)**.** *For every sequence of scores in* $[0, B]$ *and every* $t \geq 1$,

$$-\alpha\bar{\eta}_{t-1} \leq q_t \leq B + (1 - \alpha)\bar{\eta}_{t-1}.$$

Lemma F.2 (Angelopoulos et al., 2024a, Theorems 1 and 2). *For every sequence of scores in* $[0, B]$ *and every* $T \geq 1$,

$$|\mathrm{E}_T^{C^{\mathrm{ACI}}}| \leq (B + \bar{\eta}_T)\left(\frac{1}{\eta_1} + \sum_{t=2}^{T}\left|\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}}\right|\right).$$

For nonincreasing step sizes, this becomes $|\mathrm{E}_T^{C^{\mathrm{ACI}}}| \leq (B + \eta_1)/\eta_T$. *In particular,* $\eta_t = ct^{-(1-\beta)}$ *with* $c > 0$ *and* $\beta \in (0, 1]$ *gives the* (β, κ) *adversarial guarantee with* $\kappa = (B + c)/c$.

The cited theorems bound the long-run empirical coverage rate, with the difference from $1 - \alpha$ given by $-\mathrm{E}_T^{C^{\mathrm{ACI}}}/T$. For nonincreasing steps, $1/\eta_1 + \sum_{t=2}^{T}|1/\eta_t - 1/\eta_{t-1}| = 1/\eta_T$ and $\bar{\eta}_T = \eta_1$. This gives (23) for the stated power steps.

The score distribution and step sizes. For the rest of this section, fix an i.i.d. law $\mathbb{P}$ under which the score has a density satisfying (24), and suppose

$$\eta_t \longrightarrow 0, \qquad \left|\eta_{t+1}^{-1} - \eta_t^{-1}\right| \longrightarrow 0. \tag{55}$$

In particular, $\bar{\eta} := \sup_{t\geq 1}\eta_t < \infty$. The score CDF F equals 0 below 0 and 1 above B. The density bounds give

$$\lambda_-(u - v) \leq F(u) - F(v) \leq \lambda_+(u - v) \qquad (0 \leq v \leq u \leq B). \tag{56}$$

Thus F is continuous and strictly increasing on $[0, B]$, so $q_{\mathbb{P},s} \in (0, B)$ and $F(q_{\mathbb{P},s}) = 1 - \alpha$.

Cumulative quantile deviation and instantaneous training-conditional miscoverage. The threshold q_t is determined by the past observations, and the fresh score s_t is independent of that past. Hence, almost surely,

$$p_t = p_t^{C^{\mathrm{ACI}}} = 1 - F(q_t). \tag{57}$$

Define

$$\Delta_t = q_t - q_{\mathbb{P},s}, \qquad \chi(u) = F(q_{\mathbb{P},s} + u) - F(q_{\mathbb{P},s}). \tag{58}$$

Then $p_t - \alpha = -\chi(\Delta_t)$, and the recursion becomes

$$\Delta_{t+1} = \Delta_t - \eta_t\chi(\Delta_t) + \eta_t\xi_t, \qquad \xi_t = \mathsf{e}_t - p_t. \tag{59}$$

Here $-\eta_t\chi(\Delta_t)$ pulls the threshold toward the population quantile, while $\eta_t\xi_t$ is noise satisfying

$$\mathbb{E}_{\mathbb{P}}[\xi_t \mid \mathcal{G}_{t-1}] = 0, \qquad \mathbb{E}_{\mathbb{P}}[\xi_t^2 \mid \mathcal{G}_{t-1}] = p_t(1 - p_t), \qquad |\xi_t| \leq 1. \tag{60}$$

We suppress the subscript $\mathbb{P}$ on expectations below.

From disagreement to threshold deviations. The next lemma identifies the fixed-distribution risk in (32) and shows that the population prediction set attains it.

Lemma F.3 (Disagreement with the closest exact oracle). *For every $T \geq 1$,*

$$r_T(C^{\mathtt{ACI}}; \mathbb{P}) = \mathbb{E} \sum_{t=1}^{T} |p_t - \alpha| = \mathbb{E} \sum_{t=1}^{T} |\chi(\Delta_t)|. \tag{61}$$

The population prediction set in (12) attains the infimum.

The density bounds make the change in instantaneous training-conditional miscoverage comparable to the threshold deviation, including when the threshold leaves $[0, B]$.

Lemma F.4 (Comparison with the threshold deviation). *Let $\underline{\lambda} = \min\{\lambda_-, \nu/(B + \bar{\eta})\} > 0$. Then $|\Delta_t| \leq B + \bar{\eta}$ for every $t \geq 1$. For every u with $|u| \leq B + \bar{\eta}$, the value $\chi(u)$ has the same sign as u and*

$$\underline{\lambda}|u| \leq |\chi(u)| \leq \lambda_+|u|.$$

The last lemma gives the expected size of the deviation. Its proof balances the contraction of the second moment, of order $\eta_t \mathbb{E}\Delta_t^2$, against noise of order η_t^2. A fourth-moment bound ensures that the second moment does not come only from rare large deviations.

Lemma F.5 (Expected absolute deviation). *Under (55), there are constants $0 < c_0 \leq C_0 < \infty$ and a time $t_\star \geq 1$ such that*

$$c_0\sqrt{\eta_t} \leq \mathbb{E}|\Delta_t| \leq C_0\sqrt{\eta_t} \qquad (t \geq t_\star).$$

We first combine these lemmas to prove the proposition and derive the corollary. The supporting lemma proofs follow.

F.1 Proof of Proposition 3.1

Proof. The oracle identity and the density bounds reduce expected disagreement to the sum of expected absolute threshold deviations. We apply the deviation estimate and show that finitely many initial rounds are negligible. Write $A_T^\eta = \sum_{t=1}^T \sqrt{\eta_t}$, with $A_0^\eta = 0$.

Step 1: from the risk to the deviations. By Lemma F.3 and Lemma F.4,

$$\underline{\lambda} \sum_{t=1}^{T} \mathbb{E}|\Delta_t| \;\leq\; r_T(C^{\mathtt{ACI}}; \mathbb{P}) \;\leq\; \lambda_+ \sum_{t=1}^{T} \mathbb{E}|\Delta_t|. \tag{62}$$

Step 2: the size of the deviations. By Lemma F.5, $c_0\sqrt{\eta_t} \leq \mathbb{E}|\Delta_t| \leq C_0\sqrt{\eta_t}$ for $t \geq t_\star$. For $t < t_\star$ we use $0 \leq \mathbb{E}|\Delta_t| \leq B + \bar{\eta}$ from Lemma F.4. Inserting these bounds into (62) gives, for $T \geq t_\star$,

$$\underline{\lambda}\, c_0 \big(A_T^\eta - A_{t_\star - 1}^\eta\big) \;\leq\; r_T(C^{\mathtt{ACI}}; \mathbb{P}) \;\leq\; \lambda_+ C_0\, A_T^\eta + \lambda_+ (B + \bar{\eta})\, t_\star.$$

Step 3: the first rounds are negligible. It remains to show that $A_T^\eta \to \infty$, so that the constants $A_{t_\star - 1}^\eta$ and $\lambda_+(B + \bar{\eta})t_\star$ do not matter. Condition (55) says that the increments of $1/\eta_t$ tend to zero. The average of these increments then tends to zero as well, that is,

$$\frac{1}{t\,\eta_t} = \frac{1}{t\,\eta_1} + \frac{1}{t} \sum_{r<t} \Big(\frac{1}{\eta_{r+1}} - \frac{1}{\eta_r} \Big) \longrightarrow 0.$$

Hence $t\eta_t \to \infty$, so $\sqrt{\eta_t} \geq t^{-1/2}$ for all large t, and $A_T^\eta \to \infty$. With Step 2, this gives

$$0 < \liminf_{T\to\infty} \frac{r_T(C^{\mathtt{ACI}};\mathbb{P})}{A_T^\eta} \leq \limsup_{T\to\infty} \frac{r_T(C^{\mathtt{ACI}};\mathbb{P})}{A_T^\eta} < \infty, \tag{63}$$

which is (25). □

F.2 Proof of Corollary 3.1

Proof. We check the step-size assumptions of Proposition 3.1 and then evaluate the sum in its conclusion.

Step 1: check the step-size assumptions. Let $c > 0$ and $\beta \in (0, 1)$. The steps $\eta_t = ct^{-(1-\beta)}$ tend to zero because $\beta < 1$, and

$$0 \leq \frac{1}{\eta_{t+1}} - \frac{1}{\eta_t} = \frac{(t+1)^{1-\beta} - t^{1-\beta}}{c} \leq \frac{(1-\beta)t^{-\beta}}{c} \longrightarrow 0$$

because $\beta > 0$. Thus Proposition 3.1 applies.

Step 2: evaluate the cumulative rate. For $0 \leq \gamma < 1$, comparison with an integral gives

$$T^{1-\gamma} \leq \sum_{t=1}^{T} t^{-\gamma} \leq \int_0^T u^{-\gamma}\, du = \frac{T^{1-\gamma}}{1-\gamma}.$$

Taking $\gamma = (1-\beta)/2$ yields

$$\sum_{t=1}^{T} \sqrt{\eta_t} = \sqrt{c} \sum_{t=1}^{T} t^{-(1-\beta)/2} \asymp T^{(1+\beta)/2}.$$

Substitution into (25) proves (26). □

By Lemma F.2, these steps give the adversarial guarantee with $\kappa = (B + c)/c$. For any prescribed $\kappa > 1$, choose $c = B/(\kappa - 1)$. Since $(1 + \beta)/2 > \max\{1/2, \beta\}$ for $0 < \beta < 1$, the corollary gives a larger disagreement rate than truncated split conformal prediction in Theorem 2.2 under the same guarantee.

F.3 Proofs of Lemmas F.3 and F.4

Proof of Lemma F.3. The second equality in (61) holds because $p_t - \alpha = -\chi(\Delta_t)$. For the first equality, we prove two inequalities.

Lower bound. Let $C' \in O_\alpha(\mathbb{P})$ be any deterministic exact oracle. Its oracle condition gives $\mathbb{E}[\mathsf{e}_t^{C'} \mid \mathcal{G}_{t-1}] = \alpha$ almost surely, so conditional Jensen's inequality gives

$$\mathbb{E}_{\mathbb{P}}\left[|\mathsf{e}_t - \mathsf{e}_t^{C'}| \,\middle|\, \mathcal{G}_{t-1}\right] \geq \left|\mathbb{E}_{\mathbb{P}}[\mathsf{e}_t - \mathsf{e}_t^{C'} \mid \mathcal{G}_{t-1}]\right| = |p_t - \alpha|.$$

Summing over t and taking expectations, $\mathbb{E}_{\mathbb{P}}[\ell_T(C^{\mathtt{ACI}}, C')] \geq \mathbb{E}_{\mathbb{P}} \sum_{t \leq T} |p_t - \alpha|$ for every oracle in $\mathcal{O}_\alpha(\mathbb{P})$.

Upper bound. The population prediction set in (12) is deterministic and measurable. Since the fresh score is independent of the past and $F(q_{\mathbb{P},s}) = 1 - \alpha$, it belongs to $\mathcal{O}_\alpha(\mathbb{P})$. The ACI set and the population set are thresholds of the same score, at the levels q_t and $q_{\mathbb{P},s}$. Their errors differ exactly when s_t lies between the two levels, which has conditional probability $|F(q_t) - F(q_{\mathbb{P},s})| = |p_t - \alpha|$ given the past. So this oracle attains the lower bound. □

Proof of Lemma F.4. By Lemma F.1 and $q_{\mathbb{P},s} \in (0, B)$, we have $-\bar{\eta}\alpha - B < \Delta_t < B + \bar{\eta}(1 - \alpha)$, so $|\Delta_t| \leq B + \bar{\eta}$. The function χ is nondecreasing with $\chi(0) = 0$, so $\chi(u) \geq 0$ for $u > 0$ and $\chi(u) \leq 0$ for $u < 0$. It remains to prove the two bounds on $|\chi(u)|$.

For the upper bound, the extended distribution function F has a density bounded by λ_+ on the whole real line, so $|\chi(u)| = |F(q_{\mathbb{P},s} + u) - F(q_{\mathbb{P},s})| \leq \lambda_+ |u|$ for every u.

For the lower bound, let $|u| \leq B + \bar{\eta}$ and consider three cases. If $q_{\mathbb{P},s} + u \in [0, B]$, then (56) gives $|\chi(u)| \geq \lambda_- |u| \geq \underline{\lambda}|u|$. If $q_{\mathbb{P},s} + u > B$, then $\chi(u) = 1 - (1 - \alpha) = \alpha$, and $\underline{\lambda}|u| \leq \nu \leq \alpha$ because $|u| \leq B + \bar{\eta}$. If $q_{\mathbb{P},s} + u < 0$, then $\chi(u) = -(1 - \alpha)$, and $\underline{\lambda}|u| \leq \nu \leq 1 - \alpha$. □

F.4 Proof of Lemma F.5

The proof has two ingredients. The first is a comparison lemma that turns a one-step inequality for a sequence into a bound on its size. The second applies this lemma to the moments of Δ_t. The comparison is related to classical stochastic-approximation arguments (Chung, 1954; Fabian, 1968). We prove the form needed here for the general step sizes in (55).

Lemma F.6 (A comparison lemma for recursions)**.** *Let (η_t) be positive step sizes satisfying (55), let $m \geq 1$ be an integer, and let $A > 0$ and $D \geq 0$. Let (z_t) be a nonnegative sequence. Below, $o(\eta_t)$ and $o(1)$ stand for sequences of numbers whose ratio to η_t (resp. whose value) tends to zero as $t \to \infty$.*

(a) *If $z_{t+1} \leq \big(1 - A\eta_t + o(\eta_t)\big) z_t + \big(D + o(1)\big) \eta_t^{m+1}$ for all large t, then $\limsup_{t\to\infty} z_t/\eta_t^m \leq D/A$.*

(b) *If $z_{t+1} \geq \big(1 - A\eta_t + o(\eta_t)\big) z_t + \big(D + o(1)\big) \eta_t^{m+1}$ for all large t, then $\liminf_{t\to\infty} z_t/\eta_t^m \geq D/A$.*

The level $(D/A)\eta_t^m$ is where contraction and added noise balance. The two parts give the corresponding upper and lower comparisons.

Lemma F.7 (Moments of the deviation)**.** *Suppose that the step sizes satisfy (55). There exist constants $0 < c_\Delta \leq C_\Delta < \infty$ and a time $t_\star \geq 1$ such that, for every $t \geq t_\star$,*

$$c_\Delta \eta_t \leq \mathbb{E}\Delta_t^2 \leq C_\Delta \eta_t, \qquad \mathbb{E}\Delta_t^4 \leq C_\Delta \eta_t^2.$$

Proof of Lemma F.5. We deduce the size of $\mathbb{E}|\Delta_t|$ from the second and the fourth moment in Lemma F.7. Let $t \geq t_\star$. For the upper bound, the Cauchy–Schwarz inequality gives

$$\mathbb{E}|\Delta_t| \leq (\mathbb{E}\Delta_t^2)^{1/2} \leq \sqrt{C_\Delta}\,\sqrt{\eta_t}.$$

For the lower bound, we need that Δ_t is not concentrated on rare large values, and this is what the

fourth moment controls. Hölder's inequality gives $\mathbb{E}\Delta_t^2 \le (\mathbb{E}|\Delta_t|)^{2/3}(\mathbb{E}\Delta_t^4)^{1/3}$, and therefore

$$\mathbb{E}|\Delta_t| \ge \frac{(\mathbb{E}\Delta_t^2)^{3/2}}{(\mathbb{E}\Delta_t^4)^{1/2}} \ge \frac{(c_\Delta \eta_t)^{3/2}}{(C_\Delta \eta_t^2)^{1/2}} = \frac{c_\Delta^{3/2}}{C_\Delta^{1/2}}\sqrt{\eta_t}.$$

This proves the lemma with $C_0 = C_\Delta^{1/2}$ and $c_0 = c_\Delta^{3/2}/C_\Delta^{1/2}$. □

F.5 Proof of Lemma F.6

Proof. We first divide the sequence by η_t^m, and then compare the result with the level at which the two effects balance.

Step 1: normalization. Let $\widetilde{z}_t = z_t/\eta_t^m$. By (55),

$$\frac{\eta_t}{\eta_{t+1}} = 1 + \eta_t\Big(\frac{1}{\eta_{t+1}} - \frac{1}{\eta_t}\Big) = 1 + o(\eta_t),$$

and, since the step sizes are bounded, also $(\eta_t/\eta_{t+1})^m = 1 + o(\eta_t)$. So the change of the normalization from η_t^m to η_{t+1}^m only contributes a factor of this form. Dividing the inequality in (a) by η_{t+1}^m and multiplying out, we get

$$\widetilde{z}_{t+1} \le \big(1 - A\eta_t + o(\eta_t)\big)\,\widetilde{z}_t + \big(D + o(1)\big)\,\eta_t \qquad \text{for all large } t. \tag{64}$$

Only the deterministic coefficients are expanded here; no bound on $\widetilde{z}_t$ is assumed. In the same way, the inequality in (b) gives $\widetilde{z}_{t+1} \ge \big(1 - A\eta_t + o(\eta_t)\big)\,\widetilde{z}_t + \big(D + o(1)\big)\,\eta_t$ for all large t.

Step 2: the step sizes are not summable. The signed increments of $1/\eta_t$ tend to zero, so their Cesàro averages give

$$\frac{1}{t\eta_t} = \frac{1}{t\eta_1} + \frac{1}{t}\sum_{r<t}(\eta_{r+1}^{-1} - \eta_r^{-1}) \longrightarrow 0.$$

Thus $t\eta_t \to \infty$, and $\sum_t \eta_t = \infty$.

Step 3: part (a). Fix $\varepsilon \in (0, A)$. For all large t, the two o-terms in (64) are at most $\varepsilon\eta_t$ and ε. Since $\widetilde{z}_t \ge 0$, this gives, for all large t,

$$\widetilde{z}_{t+1} \le \big(1 - (A-\varepsilon)\eta_t\big)\,\widetilde{z}_t + (D+\varepsilon)\,\eta_t.$$

Let $L^+ = (D+\varepsilon)/(A-\varepsilon)$ be the level at which the two terms balance. Subtracting L^+ from both sides gives $\widetilde{z}_{t+1} - L^+ \le (1 - (A-\varepsilon)\eta_t)(\widetilde{z}_t - L^+)$. The factor lies in $(0,1)$ for large t, so the positive part satisfies

$$(\widetilde{z}_{t+1} - L^+)_+ \le \big(1 - (A-\varepsilon)\eta_t\big)(\widetilde{z}_t - L^+)_+.$$

Iterating from a large time t_1,

$$(\widetilde{z}_T - L^+)_+ \le (\widetilde{z}_{t_1} - L^+)_+ \prod_{t=t_1}^{T-1}\big(1 - (A-\varepsilon)\eta_t\big) \longrightarrow 0 \qquad \text{as } T \to \infty,$$

because $\sum_t \eta_t = \infty$. Hence $\limsup_t \widetilde{z}_t \le L^+$, and letting $\varepsilon \downarrow 0$ gives $\limsup_t \widetilde{z}_t \le D/A$.

Step 4: part (b). $D = 0$ directly implies $\widetilde{z}_t \geq 0$. Otherwise fix $\varepsilon \in (0, D)$ and let $L^- = (D-\varepsilon)/(A+\varepsilon)$. For all large t, the two o-terms are at least $-\varepsilon\eta_t$ and $-\varepsilon$, so the lower inequality gives $\widetilde{z}_{t+1} \geq (1-(A+\varepsilon)\eta_t)\,\widetilde{z}_t + (D-\varepsilon)\,\eta_t$. As in Step 3, this gives $(L^- - \widetilde{z}_{t+1})_+ \leq (1-(A+\varepsilon)\eta_t)(L^- - \widetilde{z}_t)_+$ for all large t. Hence $(L^- - \widetilde{z}_T)_+ \to 0$, so $\liminf_t \widetilde{z}_t \geq L^-$, and letting $\varepsilon \downarrow 0$ gives $\liminf_t \widetilde{z}_t \geq D/A$. □

F.6 Proof of Lemma F.7

Proof. The idea is to raise the recursion (59) to the second and the fourth power and to take expectations. The noise has mean zero given the past, so the terms that are linear in the noise vanish. What remains is a one-step inequality for each moment, to which Lemma F.6 applies.

We write $\Delta_{t+1/2} = \Delta_t - \eta_t \chi(\Delta_t)$ for the deviation, so that $\Delta_{t+1} = \Delta_{t+1/2} + \eta_t \xi_t$. The variable $\Delta_{t+1/2}$ is determined by the past. The range bound and $\bar{\eta} < \infty$ make Δ_t, $\Delta_{t+1/2}$, and all the moments below bounded.

Step 1: contraction. By Lemma F.4, $\underline{\lambda}\Delta_t^2 \leq \Delta_t \chi(\Delta_t) \leq \lambda_+ \Delta_t^2$ and $\chi(\Delta_t)^2 \leq \lambda_+^2 \Delta_t^2$. Expanding $\Delta_{t+1/2}^2 = \Delta_t^2 - 2\eta_t \Delta_t \chi(\Delta_t) + \eta_t^2 \chi(\Delta_t)^2$ therefore gives

$$(1 - 2\lambda_+\eta_t)\,\Delta_t^2 \;\leq\; \Delta_{t+1/2}^2 \;\leq\; (1 - 2\underline{\lambda}\eta_t + \lambda_+^2\eta_t^2)\,\Delta_t^2. \tag{65}$$

Step 2: the second moment, upper bound. Squaring $\Delta_{t+1} = \Delta_{t+1/2} + \eta_t \xi_t$ and using (60),

$$\mathbb{E}\big[\Delta_{t+1}^2 \mid \mathcal{G}_{t-1}\big] = \Delta_{t+1/2}^2 + \eta_t^2\, p_t(1-p_t). \tag{66}$$

Taking expectations and using the upper bound in (65) and $p_t(1-p_t) \leq 1/4$,

$$\mathbb{E}\Delta_{t+1}^2 \leq (1 - 2\underline{\lambda}\eta_t + \lambda_+^2\eta_t^2)\,\mathbb{E}\Delta_t^2 + \tfrac{1}{4}\eta_t^2.$$

Lemma F.6(a) with $z_t = \mathbb{E}\Delta_t^2$, $m = 1$, $A = 2\underline{\lambda}$, and $D = 1/4$ gives $\limsup_t \mathbb{E}\Delta_t^2/\eta_t \leq 1/(8\underline{\lambda})$. In particular $\mathbb{E}\Delta_t^2 \to 0$.

Step 3: the second moment, lower bound. For the lower bound we need that the variance of the noise does not vanish. For $p \in [0,1]$ we have $|p(1-p) - \alpha(1-\alpha)| = |p-\alpha|\,|1-p-\alpha| \leq |p-\alpha|$, and $|p_t - \alpha| = |\chi(\Delta_t)| \leq \lambda_+ |\Delta_t|$. Hence

$$\big|\mathbb{E}[p_t(1-p_t)] - \sigma^2\big| \leq \lambda_+ \mathbb{E}|\Delta_t| \leq \lambda_+ \big(\mathbb{E}\Delta_t^2\big)^{1/2} \longrightarrow 0$$

by Step 2. Taking expectations in (66) and using the lower bound in (65),

$$\mathbb{E}\Delta_{t+1}^2 \geq (1 - 2\lambda_+\eta_t)\,\mathbb{E}\Delta_t^2 + \big(\sigma^2 + o(1)\big)\,\eta_t^2.$$

Lemma F.6(b) with $z_t = \mathbb{E}\Delta_t^2$, $m = 1$, $A = 2\lambda_+$, and $D = \sigma^2$ gives $\liminf_t \mathbb{E}\Delta_t^2/\eta_t \geq \sigma^2/(2\lambda_+) > 0$. Steps 2 and 3 prove the two bounds on $\mathbb{E}\Delta_t^2$.

Step 4: the fourth moment. We expand $(\Delta_{t+1/2} + \eta_t \xi_t)^4$ and take the conditional expectation given the past. The term that is linear in ξ_t vanishes. For the other terms we use $\mathbb{E}[|\xi_t|^j \mid \mathcal{G}_{t-1}] \le \mathbb{E}[\xi_t^2 \mid \mathcal{G}_{t-1}] \le 1/4$ for $j \ge 2$, which holds because $|\xi_t| \le 1$. This gives

$$\mathbb{E}\big[\Delta_{t+1}^4 \mid \mathcal{G}_{t-1}\big] \le \Delta_{t+1/2}^4 + \tfrac{3}{2}\eta_t^2 \Delta_{t+1/2}^2 + \eta_t^3 |\Delta_{t+1/2}| + \tfrac{1}{4}\eta_t^4.$$

Squaring the upper bound in (65) gives $\Delta_{t+1/2}^4 \le (1 - 4\underline{\lambda}\eta_t + C\eta_t^2)\Delta_t^4$. Moreover, $\mathbb{E}\Delta_{t+1/2}^2 \le C\eta_t$ for all large t by Step 2, and $|\Delta_{t+1/2}|$ is bounded. So the last three terms have expectation at most $C\eta_t^3$, and for all large t,

$$\mathbb{E}\Delta_{t+1}^4 \le (1 - 4\underline{\lambda}\eta_t + C\eta_t^2)\,\mathbb{E}\Delta_t^4 + C\eta_t^3.$$

Lemma F.6(a) with $z_t = \mathbb{E}\Delta_t^4$, $m = 2$, $A = 4\underline{\lambda}$, and $D = C$ gives $\limsup_t \mathbb{E}\Delta_t^4/\eta_t^2 < \infty$.

The three limits in Steps 2–4 give the lemma, with suitable constants c_Δ and C_Δ and a suitable time $t_\star$. □