SHARP RATES IN ONLINE CALIBRATION: TWO-TEST LOWER BOUNDS AND KERNEL GEOMETRY

Anonymous authors
Paper under double-blind review

ABSTRACT

Calibrated forecasting in online adversarial settings has been widely studied and occupies a central place in online learning theory. Yet, fundamental gaps still remain in our understanding of minimax rates for a number of key notions of online calibration. In particular, tight online rates for the central *smooth calibration* and *calibration distance* notions have remained open, with a substantial gap between the known $\Omega(T^{1/3})$ and $O(\sqrt{T})$ bounds. In this work, we settle the minimax rate for both notions at $\Theta(\sqrt{T})$, with the tight lower bound witnessed by a single *oblivious* adversary. To do so, we develop a simple and general lower-bounding *two-test principle*, which forces the forecaster to incur $\Omega(\sqrt{T})$ calibration error via a single carefully constructed pair of calibration tests. We next demonstrate the wide applicability of our two-test principle. As an example, we settle the minimax rate of the above calibration notions under delayed feedback. Next, we apply our principle to online *kernel calibration*, which admits Vovk's classical $O(\sqrt{T})$ upper bound. We obtain the first matching $\Omega(\sqrt{T})$ lower bounds for a wide class of kernels, including the Gaussian and Laplace ones. Yet, we then show that the broader kernel calibration landscape is far too nuanced to permit kernel-agnostic $\sqrt{T}$-type rates: instead, kernels that permit various degrees of inter-prediction *cancellation* can be exploited by the learner to obtain faster rates of $\Theta(\log T)$ or even $\Theta(1)$. This opens up an investigation of minimax calibration rates for different kernel geometries.

1 INTRODUCTION

Online calibration asks for the intuitive property that probability forecasts should agree with observed frequencies in the long run: in a binary prediction setting, among predictions close to $p \in [0, 1]$, approximately a p fraction of outcomes should be equal to 1. Fix a positive integer horizon T. At each round $t = 1, \ldots, T$, the learner predicts $p_t \in [0, 1]$ before seeing $y_t \in \{0, 1\}$. A test function $f : [0, 1] \to \mathbb{R}$ measures the weighted bias (which we will also refer to as *(test) score*)

$$S_f(T) = \sum_{t=1}^{T} (y_t - p_t) f(p_t). \tag{1}$$

The classical notion of Expected Calibration Error (ECE) controls $S_f(T)$ uniformly over *all* bounded tests f. *Smooth calibration* relaxes this to control $S_f(T)$ uniformly over all bounded 1-Lipschitz tests f: that is, for all tests for which nearby forecasts receive nearby weights (Kakade & Foster, 2004; 2008). Besides these two notions, one can of course control calibration error over other collections of test functions. For example, *kernel calibration* (Vovk, 2007) controls all tests of at most unit norm in a reproducing kernel Hilbert space (RKHS) $\mathcal{H}_K$; it has been studied in the calibration and uncertainty quantification literature for its ability to flexibly support trainable calibration, statistical tests, and downstream decisions (Kumar et al., 2018; Widmann et al., 2019; Marx et al., 2023).

Many other calibration metrics are naturally defined without explicitly referencing collections of test functions; a key such notion is *calibration distance* (Błasiok et al., 2023). For a sequential transcript, it represents the minimum total change $\sum_{t=1}^{T} |p_t - u_t|$ needed to obtain forecasts u_t that are exactly calibrated for the same outcomes: every nonempty group of rounds assigned a common value v must have average outcome v (Qiao & Zheng, 2024). This measures the smallest change to the forecasts, rather than the outcomes, that would make the reported probabilities agree exactly with empirical

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

frequencies. Such an objective avoids the discontinuity of ECE, which treats distinct forecast values separately; their relationship is developed by Błasiok et al. (2023).

Optimal rate gaps. Despite the central significance of smooth calibration and calibration distance in the literature, their online rates have not been determined. The current state of the art (Qiao & Zheng, 2024; Arunachaleswaran et al., 2025) leaves an $\Omega(T^{1/3})$–$O(\sqrt{T})$ gap, which is also most recently highlighted in the excellent survey of Gopalan & Hu (2025, Table 1). For kernel calibration, similarly, $O(\sqrt{T})$ rates are standard (Vovk, 2007), yet matching lower bounds are, to our knowledge, not available. In this work, our main goal is to close the gap for smooth calibration and calibration distance. Our second goal is to develop ways to translate properties of given kernels into tight online kernel calibration rates (which, as we will discover, turns out to be highly nuanced).

Why establishing tight calibration rates is not so simple. To understand the key difficulty (and to gain an intuition for our solution), it helps to contrast with classic lower-bound constructions in online learning, where $\Omega(\sqrt{T})$ regret rates can often be forced by sequences of independent random outcomes. Indeed, the fluctuation due to label variance already scales as $\sqrt{T}$ with the time horizon, implying regret on the same order even for learners that know label means. For online calibration, however, this argument breaks down: here, the learner has the additional option of gaming the metric's prediction-grouping nature, by selecting its predictions not to match label means but to cancel different predictions' accumulated historical fluctuations against each other. (The susceptibility of calibration measures to such gaming has recently also been explored, starting with Haghtalab et al. (2024), from the broader viewpoint of (the learner's) *truthfulness*.) As a formalization of this intuition, Qiao & Zheng (2024, Proposition 4) show that the forecaster can obtain expected calibration distance $O(\log^{3/2} T)$ on independent fair bits, which rules out simple stochastic lower bounds.

Contribution I: A two-test lower bounding principle. This paper's main methodological contribution consists in a simple strategy for preventing such cross-prediction cancellations by the learner: To enforce $\sqrt{T}$-order calibration error, it would suffice to find a pair of test functions whose accumulated residuals the learner cannot keep small simultaneously. We achieve this by asking that the pair of tests satisfy two properties: (i) Whatever distribution of forecasts the learner chooses, Nature should be able to select a single outcome distribution that gives both test scores nonnegative expected increments. (ii) The tests should be separated by a uniform margin, so the random variation in the outcomes affects their scores differently by an amount of order $\sqrt{T}$. Our key argument shows that these two properties yield a $\Theta(\sqrt{T})$ lower bound on the expected maximum of the absolute values of the two scores. We develop the formal argument in Section 2.

Contribution II: Tight rates for smooth, distance, and kernel calibration via two tests. Beyond formulating the above two-test conditions, we establish that they are benign enough to yield tight bounds for a variety of calibration measures of interest. The principle covers both smooth and distance calibration, as well as kernel calibration for a variety of kernels — in all these cases, the resulting $\Omega(\sqrt{T})$ bounds match the long-known $O(\sqrt{T})$ upper bounds and thereby yield the sought tight rates.

1. Tight $\Theta(\sqrt{T})$ smooth calibration and calibration distance rates (Section 3). These results are our central application. Furthermore, our two-test principle gives tight rates even under *delayed feedback* (Farina & Perdomo, 2026b) where outcomes may not arrive immediately.

2. Online kernel calibration and its rate landscape (Section 4). We first use our two-test machinery to settle $\Theta(\sqrt{T})$ rates for a broad class of kernels, including the Gaussian and Laplace ones. However, perhaps surprisingly, for some kernels, we show that a learner can exploit their "cancellation properties", which imply tight *bounded* error even when the kernel is otherwise well-behaved (e.g. having smooth, injective, nonvanishing features). More, we show that not all types of cancellation let the learner obtain bounded calibration error; for some, the tight rate is *logarithmic*, although demonstrating this requires us to develop tools beyond our two-test machinery. To obtain fast kernel calibration rates, we present a general "online minimax multiobjective" Algorithm 1, which to our knowledge is novel and is of independent interest. Compared to prior works (Lee et al., 2022; Haghtalab et al., 2023), it solves a minimax game with an extra second-order correction term, which makes sub-$\sqrt{T}$ rates available for favorable problem structures. Overall, our findings suggest that the landscape of kernel calibration rates is complex; classifying the set of possible rates by geometric properties that cause them is the main question that we leave open.

2 THE TWO-TEST LOWER-BOUND FRAMEWORK

In this section, we develop the two-test lower-bounding framework in its general form. We first define some notation that will be used throughout. We write $x_+ = \max\{x, 0\}$. For $q \in [0,1]$, $\mathrm{Bernoulli}(q)$ assigns probability q to 1 and $1-q$ to 0. We write Rad for the Rademacher distribution, which assigns probability $1/2$ to each of -1 and 1. We use $\mathbf{1}_E$ for an indicator. Given a function f, we write $\|f\|_\infty = \sup_p |f(p)|$, and $\mathrm{Lip}(f)$ for its smallest Lipschitz constant. For a space X, let $\mathcal{P}(X)$ denote its Borel probability measures, also written $\Delta(X)$ for finite X; δ_x is the point mass at x. For a sequence $(z_t)_{t\geq 1}$, we will use the shorthands: $z_{<t} = (z_1, \ldots, z_{t-1})$, $z_{1:t} = (z_1, \ldots, z_t)$.

Preliminaries: online forecasting game The game runs for T rounds between a learner A and an adversary (or Nature) B; their strategies must be measurable and may use randomness.

Definition 1 (Online protocol)**.** At round $t \geq 1$, A selects and announces a forecast law $\nu_t \in \mathcal{P}([0,1])$, using its private state and the preceding public history $H_{t-1} = ((\nu_s, p_s, y_s))_{s<t}$. It then samples $p_t \sim \nu_t$ with fresh private randomness. Next, B sees H_{t-1}, ν_t, but not p_t, and chooses $y_t \in \{0,1\}$ (it may randomize, independently of the fresh draw).[1] Then p_t, y_t become public.

For a nonnegative measurable loss $\mathcal{L}$ of the *transcript* $p = (p_1, \ldots, p_T)$, $y = (y_1, \ldots, y_T)$, define

$$V_T(\mathcal{L}) = \inf_A \sup_B \mathbb{E}_{A,B}[\mathcal{L}(p, y)]. \tag{2}$$

For now, we keep the loss $\mathcal{L}$ abstract; concrete calibration notions will be handled later.

A (randomized) adversary is called *oblivious* if it selects a distribution Q over outcome sequences $\{0,1\}^T$ before play begins, and samples the outcome sequence independently of the learner's randomness. As our framework below shows, oblivious adversaries in fact suffice to obtain tight lower bounds for a variety of calibration notions — adaptivity is not required.

We now present a general lower-bounding principle for online calibration. One test does not suffice. For instance, consider the single test $f \equiv 1$, corresponding to the requirement that the forecasts be unbiased on average. Predicting $p_1 = 0$ and $p_t = y_{t-1}$ for $t \geq 2$ gives $S_f(T) = y_T \leq 1$. Thus, the learner can keep the score of a single test uniformly bounded even against arbitrary adversaries.

Definition 2 (Common response)**.** Fix a continuous test $g : [0,1] \to \mathbb{R}$, a forecast distribution $\nu \in \mathcal{P}([0,1])$, and $q \in [0,1]$. Denote the expected increment of the score of g, when the forecast has distribution ν and the outcome is independent $\mathrm{Bernoulli}(q)$, as

$$r_g(\nu, q) = \mathbb{E}_{P\sim\nu}\big[(q - P)g(P)\big]. \tag{3}$$

Fix two continuous tests $g_+, g_- : [0,1] \to \mathbb{R}$. We call $q \in [0,1]$ a (g_+, g_-)-*common ν-response* if both $r_{g_+}(\nu, q)$ and $r_{g_-}(\nu, q)$ are nonnegative.

Keeping q away from 0 and 1 leaves randomness in the outcome; keeping the two tests separated ensures that their difference detects it.

Theorem 1 (Two-test lower bound)**.** *Fix any $a \in (0, 1/2)$ and $\gamma > 0$. Consider any two continuous tests g_+, g_- that satisfy the following two conditions:*

1. *for every forecast distribution ν, there exists a (g_+, g_-)-*common ν-response $q \in [a, 1-a]$,
2. *for every forecast value p, the half-difference between the tests $d(p) := \frac{g_+(p) - g_-(p)}{2} \geq \gamma$.*

Then, for each horizon T, there exists an oblivious adversary such that for every learner,

$$\mathbb{E}\big[\max\{|S_{g_+}(T)|, |S_{g_-}(T)|\}\big] \geq a\gamma\mathbb{E}\left[\left|\sum\nolimits_{t=1}^{T} \varepsilon_t\right|\right] \geq a\gamma\sqrt{T/2}, \tag{4}$$

where $(\varepsilon_t)_{t\geq 1}$ are i.i.d. with $\varepsilon_t \sim \mathrm{Rad}$.

Proof sketch; full proof in Appendix B.3. Write $\mathcal{L}(p, y) = \max\{|S_{g_+}(T)|, |S_{g_-}(T)|\}$ and $\mathcal{E}_T = a\gamma\,\mathbb{E}[|\sum_{t=1}^T \varepsilon_t|]$. We first fix a randomized learner A and construct an outcome distribution Q_A with $\mathbb{E}_{A,Q_A}[\mathcal{L}] \geq \mathcal{E}_T$, and then apply the minimax theorem to obtain one distribution for all learners.

[1] Thus y_t and p_t are conditionally independent given (H_{t-1}, ν_t).

For each history $y_{<t}$, let $\bar{\nu}_t(y_{<t})$ be the distribution of the learner's next forecast when its preceding outcomes are prescribed to be $y_{<t}$. Choose a common response $q_t(y_{<t}) \in [a, 1-a]$ to $\bar{\nu}_t(y_{<t})$, and define an outcome distribution Q_A by $Y_t \mid Y_{<t} \sim \mathrm{Bernoulli}(q_t(Y_{<t}))$.

Abbreviate $\bar{\nu}_t = \bar{\nu}_t(Y_{<t})$ and $q_t = q_t(Y_{<t})$. We use $\mathbb{E}_A$ to average over the learner's randomness with Y held fixed. For $i \in \{+, -\}$, consider

$$\mathbb{E}_A[S_{g_i}(T)] = \sum_{t=1}^{T} \mathbb{E}_{P\sim\bar{\nu}_t}[(Y_t - P)g_i(P)], \quad M_i(Y) := \sum_{t=1}^{T} (Y_t - q_t)\mathbb{E}_{P\sim\bar{\nu}_t}[g_i(P)].$$

By the definition of common response, $\mathbb{E}_A[S_{g_i}(T)] - M_i(Y) = \sum_{t=1}^{T} r_{g_i}(\bar{\nu}_t, q_t) \geq 0$.

We can also show that $\mathbb{E}_{Q_A}[M_i(Y)] = 0$.

Define $b_t = \mathbb{E}_{P\sim\bar{\nu}_t}[d(P)] \geq \gamma$ and $N = \sum_t b_t(Y_t - q_t) = (M_+(Y) - M_-(Y))/2$. Then we have

$$\begin{aligned}\mathbb{E}_{A,Q_A}[\mathcal{L}] &= \mathbb{E}_{A,Q_A}[\max_i |S_{g_i}(T)|] \geq \mathbb{E}_{Q_A}\left[\max_i \mathbb{E}_A[S_{g_i}(T)]\right] \geq \mathbb{E}_{Q_A}\left[\max_i M_i(Y)\right] \\ &= \mathbb{E}_{Q_A}[M_+(Y) + M_-(Y)]/2 + \mathbb{E}_{Q_A}[|M_+(Y) - M_-(Y)|]/2 = \mathbb{E}_{Q_A}[|N|].\end{aligned}$$

where the second line follows from the identity $\max\{u, v\} = (u + v + |u - v|)/2$.

We compare N with a Rademacher sum by generating the same outcome distribution Q_A in a different way. Independently of the learner's randomness, let $(\varepsilon_t)_{t\geq 1}$ be i.i.d. with $\varepsilon_t \sim \mathrm{Rad}$, and generate the outcomes successively by $\mathbb{P}(Y_t = 1 \mid Y_{<t}, \varepsilon_{1:T}) = q_t + (a\gamma/b_t) \cdot \varepsilon_t$. These probabilities lie in $[0, 1]$, since $a\gamma/b_t \leq a$ and $q_t \in [a, 1-a]$. Earlier outcomes use only earlier Rademacher variables, so ε_t is independent of $Y_{<t}$. Averaging over it gives $\mathbb{P}(Y_t = 1 \mid Y_{<t}) = q_t$, and hence Y still has law Q_A. Moreover,

$$\mathbb{E}[b_t(Y_t - q_t) \mid Y_{<t}, \varepsilon_{1:T}] = b_t(\mathbb{E}[Y_t \mid Y_{<t}, \varepsilon_{1:T}] - q_t) = a\gamma\varepsilon_t,$$

$$\mathbb{E}[N \mid \varepsilon_{1:T}] = \sum\nolimits_t \mathbb{E}[b_t(Y_t - q_t) \mid \varepsilon_{1:T}] = \sum\nolimits_t \mathbb{E}[\mathbb{E}[b_t(Y_t - q_t) \mid Y_{<t}, \varepsilon_{1:T}] \mid \varepsilon_{1:T}] = a\gamma \sum_{t=1}^{T} \varepsilon_t.$$

Consequently,

$$\mathbb{E}[|N|] \geq \mathbb{E}\left[\left|\mathbb{E}[N \mid \varepsilon_{1:T}]\right|\right] = a\gamma\, \mathbb{E}\left[\left|\sum\nolimits_{t=1}^{T} \varepsilon_t\right|\right] = \mathcal{E}_T.$$

Together with the preceding bound $\mathbb{E}_{A,Q_A}[\mathcal{L}] \geq \mathbb{E}[|N|]$, this proves $\mathbb{E}_{A,Q_A}[\mathcal{L}] \geq \mathcal{E}_T$.

Applying the Khinchin inequality (Lemma B.2) with all weights equal to one gives $\mathcal{E}_T \geq a\gamma\sqrt{T/2}$.

Finally, we can verify the conditions of Sion's minimax theorem by representing randomized learners as distributions over deterministic forecasting rules. The theorem then gives

$$\max_Q \inf_A \mathbb{E}_{A,Q}[\mathcal{L}] = \inf_A \max_Q \mathbb{E}_{A,Q}[\mathcal{L}].$$

The above bound implies that a distribution Q_T attaining the maximum on the left guarantees the same lower bound against every randomized learner. Thus Q_T defines the required oblivious adversary. □

3 Tight rates for smooth calibration and calibration distance

In this section, we apply the two-test lower bound to smooth calibration, calibration distance, and lower calibration distance, obtaining tight $\Omega(\sqrt{T})$ lower bounds for them.

The calibration metrics. All of the following calibration metrics are defined with respect to a prediction sequence p and an outcome sequence y; we will omit (p, y) when clear. We first define the classical notion of smooth calibration; we follow the cumulative formulation in Qiao & Zheng (2024) of the Lipschitz-test notion originating in weak calibration (Kakade & Foster, 2004; 2008).

Definition 3 (Smooth calibration error)**.** Let $\mathcal{F}_{\mathrm{sm}}$ be the class of 1-Lipschitz functions $f : [0, 1] \to [-1, 1]$. The *smooth calibration error* is $\mathrm{smCE}_T(p, y) = \sup_{f\in\mathcal{F}_{\mathrm{sm}}} |S_f(T)|$.

We now define calibration distance and its close cousin, lower calibration distance. Calibration distance compares the forecasts with one exactly calibrated sequence; lower calibration distance instead permits a probability distribution over comparison forecasts at each round, subject to an aggregate calibration constraint. We follow the sequential definitions of Qiao & Zheng (2024) that are based upon the original batch definitions of Błasiok et al. (2023).

Definition 4 (Calibration distance)**.** For fixed outcomes (y_t), an *exactly calibrated sequence* is a sequence $u \in [0,1]^T$ with $\sum_{t:u_t=v}(y_t - v) = 0$ for every $v \in [0,1]$. Let $\mathcal{C}(y)$ be the set of such sequences. The *calibration distance* is $\mathrm{CalDist}_T(p,y) = \min_{u\in\mathcal{C}(y)} \sum_t |p_t - u_t|$.

Definition 5 (Lower calibration distance)**.** For each round, let λ_t be a finitely supported probability measure on $[0,1]$. Call the family (λ_t) calibrated if $\sum_t (y_t - u)\lambda_t(\{u\}) = 0$ for every $u \in [0,1]$. The *lower calibration distance* is $\mathrm{LowerCalDist}_T(p,y) = \inf_{(\lambda_t)\text{ calibrated}} \sum_t \mathbb{E}_{U\sim\lambda_t}\big[|p_t - U|\big]$.

Since every calibrated sequence gives a calibrated family of point masses $\lambda_t = \delta_{u_t}$, one has $\mathrm{LowerCalDist}_T \le \mathrm{CalDist}_T$ (Qiao & Zheng, 2024, Remark 7). See Appendix C for technical details on the two notions and Gopalan et al. (2026) for further characterizations.

A pair of clipped test functions. Write $\mathrm{clip}(x,[a,b]) = \min\{b, \max\{a,x\}\}$ for $a \le b$. Consider the pair of tests:

$$f_+(p) = \mathrm{clip}(2/3 - p, [0, 1/3]),\ \ f_-(p) = f_+(p) - 1/3. \quad (5)$$

Both tests are bounded and 1-Lipschitz; see Figure 1. The first is nonnegative and weights smaller forecasts; the second is nonpositive and weights larger forecasts. Their half-difference is $1/6$. For a forecast law with mean $\mu = \mathbb{E}_{P\sim\nu}[P]$, choose $q = \mathrm{clip}(\mu, [1/3, 2/3])$. The probabilities of outcomes 0 and 1 are then both at least $1/3$. For either test,

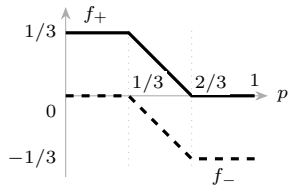


Figure 1: The two tests in (5).

$$r_f(\nu, q) = f(q)(q - \mu) + \mathbb{E}_{P\sim\nu}\big[(f(P) - f(q))(q - P)\big] \ge 0.$$

The expectation term is nonnegative because both tests are nonincreasing. The first term vanishes when $q = \mu$. If $\mu < 1/3$, then $q - \mu > 0$ and both tests are nonnegative at $q = 1/3$; if $\mu > 2/3$, both signs reverse. Thus q is a common response.

Since the two tests are admissible smooth tests, one has $\max\{|S_{f_+}(T)|, |S_{f_-}(T)|\} \le \mathrm{smCE}_T$. In general, the three metrics satisfy $\mathrm{smCE}_T/2 \le \mathrm{LowerCalDist}_T \le \mathrm{CalDist}_T$ (Qiao & Zheng, 2024, Lemma 8 and Remark 7). For our two tests, the fact that $p \mapsto (y-p)f(p)$ is 1-Lipschitz allows us to prove the stronger inequality below; see Appendix C.

Proposition 1 (Two-test simultaneous lower bounds)**.** *On a binary transcript, the tests in* (5) *satisfy*

$$\max\{|S_{f_+}(T)|, |S_{f_-}(T)|\} \le \min\{\mathrm{smCE}_T, \mathrm{LowerCalDist}_T, \mathrm{CalDist}_T\}. \quad (6)$$

Combining Proposition 1 with Theorem 1 at $a = 1/3$ and $\gamma = 1/6$ gives the lower bound below. For the matching upper bounds, Propositions C.1 and G.1 give $\mathrm{smCE}_T \le \sqrt{T}$ for a deterministic forecaster, and Lemma C.2 gives the same $O(\sqrt{T})$ rate for $\mathrm{LowerCalDist}_T$ and $\mathrm{CalDist}_T$.

Theorem 2 (Tight rates for smooth calibration and calibration distance)**.** *Smooth calibration and both calibration distances have minimax rate $\Theta(\sqrt{T})$ against either adaptive or oblivious adversaries. For each T, one oblivious adversary forces every randomized forecaster to satisfy*

$$\mathbb{E}\big[\min\{\mathrm{smCE}_T, \mathrm{LowerCalDist}_T, \mathrm{CalDist}_T\}\big] \ge \sqrt{T}/\left(18\sqrt{2}\right). \quad (7)$$

Tight bounds under *delayed feedback* Using the two-test principle, we can further settle these calibration notions' minimax rate in a *delayed-feedback* extension of the above setting, where the learner may not see the true outcome of a round until several rounds later (Definition 6). It turns out that square-root-type rates are once again optimal, up to the additional (optimal) dependence on the total delay incurred.

Definition 6 (Delayed-feedback protocol)**.** Modify Definition 1 by delaying the learner's outcome observations. A deterministic schedule fixed before play chooses integers $\tau_i \in \{0, \dots, T-i\}$; label (i, y_i) is revealed just before round $i + 1 + \tau_i$. Time $T+1$ follows the last prediction, and $\tau_i = 0$ is ordinary feedback. The learner uses its private state, previous announcements and forecasts, and revealed labels only. Nature knows the schedule and unrevealed past outcomes; the announcement and hidden-fresh-draw rules are unchanged. For a budget $D \ge 0$, define

$$V_{T,D}(\mathcal{L}) = \inf_A \max_{\tau:\sum_i \tau_i \le D} \sup_B \mathbb{E}_{A,B}[\mathcal{L}(p,y)].$$

The learner does not know future reveal times. The upper algorithm needs neither T nor D; the lower bound holds even with the schedule announced.

The next result exhibits the tight bounds in this setting. For the lower bound, we (1) partition the rounds into appropriately sized blocks of consecutive rounds, and (2) hide a repeated outcome across each block. This replaces the two-test fair walk of the above vanilla setting with a *weighted* walk.

Theorem 3 (Tight total-delay calibration rate)**.** *For each* $\mathcal{L} \in \{\mathrm{smCE}_T, \mathrm{LowerCalDist}_T, \mathrm{CalDist}_T\}$, $V_{T,D}(\mathcal{L}) = \Theta\big(\min\{T, \sqrt{T+D}\}\big)$. *For any* T, D*, an oblivious adversary and one schedule within the budget suffice for the lower bound.*

The full proof is in Appendix G; here, we briefly sketch the lower-bounding argument, and note that the upper bound is obtained via defensive forecasting (and written down via the language of kernels; see Section 4 for background). For the lower bound, we repeat one bit in each block of lengths $b_1, \ldots, b_n$, revealing it after the last prediction. Average forecasts over the next block's positions and the learner's randomness; within-block predictions need not be independent. The common response yields the weighted fluctuation $a\gamma \sum_j b_j \varepsilon_j$, for independent fair signs ε_j, $a = 1/3$, $\gamma = 1/6$. Its expected magnitude is at least $(a\gamma/\sqrt{2})\sqrt{\sum_j b_j^2}$. The exact schedule cost is $D_{\mathrm{used}} = \sum_j b_j(b_j - 1)/2$, and $\sum_j b_j^2 = T + 2D_{\mathrm{used}}$. Equal-length blocks and a shorter remainder use at most D and give $\sum_j b_j^2 \geq \frac{1}{4}\min\{T^2, T + D\}$.

4 Tight rates for kernel calibration: Fast and slow

We now show that the two-test machinery gives tight rates for a broad class of online *kernel* calibration problems, where the metric and tests come from a chosen RKHS. While every fixed continuous kernel admits an $O(\sqrt{T})$ forecasting guarantee (Vovk, 2007), its optimality for broad kernel classes has, to our knowledge, remained open. Our method proves a matching lower bound whenever the RKHS contains two tests that force uncancellable $\sqrt{T}$ fluctuations. This criterion covers an important class of radial kernels, including the Gaussian and Laplace kernels studied in the calibration distance literature (Błasiok et al., 2023, Theorems 8.5–8.6).

Yet the story of tight rates only begins with the $\sqrt{T}$ case: natural *cancellation properties* of many kernels can permit logarithmic or even bounded kernel calibration error. To obtain these faster rates, we develop a new algorithm for *online minimax multiobjective optimization* (Lee et al., 2022; Haghtalab et al., 2023), which is of independent interest: compared to standard such algorithms that are limited to $O(\sqrt{T})$ rates, it offers an explicit minimax-based mechanism for exploiting the problem structure to obtain much faster rates. After specializing this algorithm to kernel-calibration objectives, we identify several kernel-geometric criteria that allow the learner to satisfy the algorithm's minimax conditions by forming forecasts that can cancel each other. Using these criteria, we exhibit sharp transitions between the three rate regimes of $\Theta(\sqrt{T})$, $\Theta(\log T)$, and $\Theta(1)$ within concrete kernel families, leaving a general classification open.

4.1 Preliminaries and Tight $\Theta(\sqrt{T})$ Rates

First, we formally define kernel calibration. The point of this concept is simple yet far-reaching: to enable calibration for test families that represent a wide class of kernel-induced *geometries*.

Definition 7 (Kernel and its test space)**.** A *kernel* is a continuous, symmetric function $K : [0,1]^2 \to \mathbb{R}$ that is positive semidefinite: $\sum_{i,j=1}^n c_i c_j K(x_i, x_j) \geq 0$ for every positive integer n, $x_i \in [0,1]$, and $c_i \in \mathbb{R}$. Its reproducing kernel Hilbert space (RKHS) $\mathcal{H}_K$ has canonical features $\phi(p) = K(p, \cdot)$, with $f(p) = \langle f, \phi(p)\rangle_{\mathcal{H}_K}$ and $K(p,q) = \langle \phi(p), \phi(q)\rangle_{\mathcal{H}_K}$. The corresponding norm for this inner product will be denoted $\|\cdot\|_{\mathcal{H}_K}$ or simply $\|\cdot\|$. Also, put $\kappa = \max_{p\in[0,1]} \|\phi(p)\|$. Equivalently, one may use any continuous Hilbert-space feature map with this inner product, restricted to its closed linear span. The kernel has *finite rank* when that span is finite-dimensional.

Definition 8 (Kernel calibration error)**.** For $t \geq 1$, let $\mathbf{S}_K(t) = \sum_{s=1}^t (y_s - p_s)\phi(p_s)$. Then, define

$$\mathrm{kCE}_{K,T} = \|\mathbf{S}_K(T)\|_{\mathcal{H}_K} = \sup_{\|f\|_{\mathcal{H}_K} \leq 1} |S_f(T)|. \tag{8}$$

This is the cumulative RKHS calibration error notion studied in online forecasting by Vovk (2007, Theorems 1–2), following Vovk et al. (2005). Its normalized statistical counterpart is kernel calibration error (Kumar et al., 2018); see Błasiok et al. (2023, Definition 8.1). For the minimax values, write $V_T(K) = V_T(\mathrm{kCE}_{K,T})$ and $V_T^{\mathrm{obl}}(K) = V_T^{\mathrm{obl}}(\mathrm{kCE}_{K,T})$.

A seminal *defensive forecasting* algorithm of Vovk (2007, Theorem 1) guarantees, for all kernels K,

$$\|\mathbf{S}_K(T)\|^2 \le \sum\nolimits_t p_t(1-p_t)K(p_t,p_t) \le \kappa^2 T/4 \implies V_T(K) \le \kappa\sqrt{T}/2. \tag{9}$$

(For completeness, Proposition G.1 gives the strategy and proof.) When is this unavoidable, and when can the learner unlock faster rates?

The radial criterion below uses only two kernel sections. A second criterion extends the argument beyond radial kernels through a stable sign change of an RKHS test. (Note that constants in kernel-dependent rates may depend on K, which stays fixed as T grows.)

Theorem 4 (Square-root kernel criteria)**.** *Each of the following implies $V_T(K) = \Theta_K(\sqrt{T})$.*

(i) $K(p,q) = \rho(|p-q|)$ for a continuous, strictly decreasing $\rho : [0,1] \to (0,\infty)$, with K positive semidefinite.

(ii) There are continuously differentiable tests $m, d \in \mathcal{H}_K$ and $r \in (0,1)$ with $m(r) = 0$, $m'(r) < 0$, $m(p) > 0$ for $p < r$, $m(p) < 0$ for $p > r$, and $d(r) \neq 0$.

For fixed $\sigma > 0$, (i) covers the Gaussian kernel $e^{-(p-q)^2/(2\sigma^2)}$ and the Laplace kernel $e^{-|p-q|/\sigma}$.

For (i), the endpoint sections give $m(p) = \rho(p) - \rho(1-p)$ and $d(p) = \rho(p) + \rho(1-p) > 0$. The strictly decreasing ratio m/d gives two separated tests and a weighted clipped-mean response (Lemma D.2). For (ii), perturbing the crossing test yields two tests that need not remain uniformly separated. We use the extension in Theorem B.1, which allows sufficiently positive expected increments to compensate when their average separation is too small. See Appendix D for the construction. Theorem D.1 gives stationary extensions.

4.2 Fast rates through cancellation

We have just established that for a substantial class of kernels, the tight calibration rate is indeed $\Theta(\sqrt{T})$. However, as we will now show, many kernels have much faster tight rates. The general phenomenon that drives this is the fact that the canonical map ϕ can create feature vectors that admit various cancellations — such that the learner may be able to cancel residuals accumulated at different forecasts by acting strategically, thereby obtaining calibration error rates such as $O(1)$ or $O(\log T)$. For a very simple example of how this could occur, consider the analytic, pointwise positive kernel

$$K_{\mathrm{b}}(p,q) = (1+p)(1+q)\big(2 + \cos(2\pi(p-q))\big). \tag{10}$$

Its features $\phi_{\mathrm{b}}(p) = (1+p)(\sqrt{2}, \cos(2\pi p), \sin(2\pi p))$ are injective and nonvanishing, but $\phi_{\mathrm{b}}(1) = 2\phi_{\mathrm{b}}(0)$. Writing the accumulated error as Zv, where $v = \phi_{\mathrm{b}}(0)/\sqrt{3}$, predict 0 when $Z \le 0$ and 1 otherwise. Either outcome preserves $Z \in [-2\sqrt{3}, \sqrt{3}]$, giving a pathwise bounded error and rate $\Theta(1)$ (Proposition E.4). Now, we will begin to explore deeper ways for cancellations to occur.

Illustration: Sharp rate transitions within kernel families. We now exemplify cancellation phenomena via two concrete kernel classes. In both cases, some kernels in the class have the "default" rate $\Theta(\sqrt{T})$, while the rest admit extra structure leading to a faster rate. Both transitions will be proved from more general proof ingredients that we develop below: an upper-bound meta-algorithm (that gives $O(1)$ or $O(\log T)$ rates depending on the kernel instance), and a log-rate lower bound.

Theorem 5 (The cosine transition)**.** *For $K_\omega(p,q) = \cos(\omega(p-q))$ and fixed $\omega \ge 0$, $V_T(K_\omega) = \Theta_\omega(\sqrt{T})$ for $0 < \omega < 2\pi$, and $V_T(K_\omega) = \Theta_\omega(1)$ for $\omega = 0$ or $\omega \ge 2\pi$.*

In this first family (for which the formal proof is in Appendix E), the features $(\cos(\omega p), \sin(\omega p))$ trace a circular arc. Below one turn, Theorem 4(ii) above supplies a two-test $\Omega(\sqrt{T})$ bound. At one full turn and beyond, cancellations become available, and the below algorithm gives an $O(1)$ rate. Now, we move to our second family, where the dichotomy is such that cancellations cannot be exploited towards an $O(1)$ rate but where an $O(\log T)$ rate can instead be attained.

Definition 9 (Profile kernels and covered points)**.** Let $h : [0,1] \to [0,1]$ be continuous, with $h(0) = 0$ and $h(1) = 1$, and set $\phi(p) = (1+p)(1-h(p), h(p))$ and $K(p,q) = \langle \phi(p), \phi(q) \rangle$. Call an interior point q *covered* if $h(s) = h(t)$ for some $s < q < t$; point 0 is covered if $h(t) = 0$ for some $t > 0$, and point 1 if $h(s) = 1$ for some $s < 1$.

Corollary 1 (The profile transition)**.** *Suppose h is constant or strictly monotone on each piece of a finite partition, and $h = 0$ near 0 or $h = 1$ near 1. Then $V_T(K) = \Theta_K(\log(T+1))$ if every point is covered, and $V_T(K) = \Theta_K(\sqrt{T})$ otherwise.*

Appendix F.3 proves this formally (and an extension to rank-two kernels); the square-root case follows by our two-test Theorem 1, while the logarithmic case uses Theorem 7 and Corollary 2 below.

A generic online minimax algorithm. We now develop an algorithm that operates in a fairly general setting (see e.g. Lee et al. (2022)), which we will then instantiate towards $O(1)$ and $O(\log T)$ kernel calibration. Let $\mathcal{A}, \mathcal{Y}$ be finite nonempty action and outcome sets. A *gain map* is any $g : \mathcal{A} \times \mathcal{Y} \to \mathbb{R}^n$ with $\|g(a,y)\|_2 \le B$ for some $n \ge 1$, $B > 0$. The learner seeks to keep the cumulative vector $Z_t = \sum_{s \le t} g(a_s, y_s)$ near zero. At each round it announces a distribution π_t, privately draws $a_t \sim \pi_t$, and observes y_t. Nature may use the past and π_t, but cannot react to the fresh draw a_t. Now, we define the primitive that the algorithm will use: the multiplicative weights update.

Definition 10 (Multiplicative weights and a score)**.** We denote by $\mathrm{MW}(m, \eta)$ the multiplicative weights update algorithm for m experts, with learning rate $\eta > 0$, and initial weights $w_j = 1$. Given $u \in \mathbb{R}^m$, its operation $\mathrm{update}(u)$ replaces w_j by $w_j e^{\eta u_j}$, and $\mathrm{score}(u) = \sum_{j=1}^m w_j(u_j + \eta u_j^2)$.

The algorithm below uses MW as follows. We use $2n$ experts labelled $(i,+), (i,-)$, responsible for the coordinates of the vectors we seek to control. The resulting weights after round t are $w_i^{\pm} = e^{\pm \eta Z_{t,i}}$. For $z \in \mathbb{R}^n$, define its signed lift by $\widetilde{z}_{i,+} = z_i$ and $\widetilde{z}_{i,-} = -z_i$.

Algorithm 1 A fast-rate online minimax algorithm (given $\mathcal{A}, \mathcal{Y}$, gain map g, and $\eta \in (0, 1/B)$.)

1: Initialize $M \leftarrow \mathrm{MW}(2n, \eta)$ with signed expert labels.
2: **for** $t = 1, 2, \ldots$ **do**
3: Choose $\pi_t \in \arg\min_{\pi \in \Delta(\mathcal{A})} \max_{y \in \mathcal{Y}} \mathbb{E}_{a \sim \pi}\big[M.\mathrm{score}(\widetilde{g}(a,y))\big]$, and draw action $a_t \sim \pi_t$.
4: Observe y_t, chosen by the adversary with knowledge of π_t but not a_t.
5: Call $M.\mathrm{update}(\widetilde{g}(a_t, y_t))$.
6: **end for**

Theorem 6 (Logarithmic and bounded cumulative gains)**.** *Fix $\eta \in (0, 1/B)$. Suppose for $\epsilon \ge 0$,*

$$\forall q \in \Delta(\mathcal{Y})\ \forall x \in \mathbb{R}^n\ \exists a \in \mathcal{A} : \quad \mathbb{E}_{Y \sim q} \sum\nolimits_i \big(x_i g_i(a,Y) + \eta |x_i| g_i(a,Y)^2\big) \le -\epsilon \|x\|_1. \tag{11}$$

Then Algorithm 1 satisfies

$$\mathbb{E}\|Z_T\|_1 \le n/\eta \cdot \log\left(2 + 2\eta^2 B^2/n \cdot \sum\nolimits_{s=0}^{T-1} (1-\eta\epsilon)^s\right). \tag{12}$$

This is $O(\log(T+1))$ when $\epsilon = 0$, and is bounded uniformly in T when $\epsilon > 0$.

This guarantee supplies a condition that we will verify, below, towards fast-rate guarantees; its full proof is in Appendix E.1. For a sketch, the paired weights satisfy $w_i^+ w_i^- = 1$. Subtract their minimum total mass and write $\Phi_t = \sum_i (w_i^+ + w_i^- - 2)$. The estimate $e^u \le 1 + u + u^2$ for $|u| \le 1$, together with the minimax theorem and (11), gives $\mathbb{E}[\Phi_{t+1} \mid \mathcal{F}_t] \le (1 - \eta\epsilon)\Phi_t + 2\eta^2 B^2$, for $\mathcal{F}_t$ being the history through round t. Iterating gives (12).

Controlling kernel calibration error via vector gains. To utilize Algorithm 1 in our kernel calibration setting, in every case we specify a finite action set, where action a produces a forecast $p(a)$, and fixed vectors v_i with $\|v_i\| \le 1$ such that

$$(y - p(a))\phi(p(a)) = \sum\nolimits_i g_i(a,y) v_i. \tag{13}$$

Then $\mathbf{S}_K(T) = \sum_i Z_{T,i} v_i$, so $\mathrm{kCE}_{K,T} \le \|Z_T\|_1$.

A simple sufficient condition for $O(\log T)$ kernel calibration rate. Consider the following:

Definition 11 (Finite cancellation cover)**.** For $0 \le \ell < u \le 1$, let $J(\ell, u)$ be (ℓ, u) together with 0 if $\ell = 0$ and 1 if $u = 1$. A *finite cancellation cover* consists of $N \ge 1$ pairs $0 \le \ell_i < u_i \le 1$, unit vectors v_i, and $\alpha_i, \beta_i > 0$ such that

$$\phi(\ell_i) = \alpha_i v_i, \qquad \phi(u_i) = \beta_i v_i, \qquad \cup_{i=1}^N J(\ell_i, u_i) = [0,1]. \tag{14}$$

This condition embeds a very direct cancellation mechanism that the learner can try to exploit: for a hypothetical outcome mean $q \in (\ell_i, u_i)$, the expected increments $(q - \ell_i)\alpha_i v_i$ and $(q - u_i)\beta_i v_i$ point in opposite directions. A cover supplies such a correcting pair for *every* mean — and in this way, lets us apply the above algorithm: indeed, the minimax theorem removes the need for the learner to know that mean. We now derive a logarithmic-rate upper bound via Algorithm 1. The appendix further extends the construction to positive combinations of feature vectors.

Corollary 2 (A cover gives logarithmic error)**.** *Given a finite cancellation cover, use labelled actions $(i, \ell_i), (i, u_i)$ and gains $g((i, p), y) = r_{i,p}(y - p)\mathbf{e}_i \in \mathbb{R}^N$, where $\mathbf{e}_i$ is the ith coordinate vector, $r_{i,\ell_i} = \alpha_i$, and $r_{i,u_i} = \beta_i$. For some $\eta > 0$, Algorithm 1 gives expected kernel error $O_K(\log(T+1))$.*

The full proof is in Appendix E.3. For a sketch, note that for any outcome mean q, a pair covering q gives nonnegative expected gain in coordinate i at its lower endpoint and nonpositive expected gain at its upper endpoint. We choose the endpoint that makes the weighted gain $x_i \mathbb{E}[g_i]$ nonpositive. The cover ensures that suitable pairs and one sufficiently small fixed $\eta > 0$ make this correction control the quadratic term in Equation (11) as well. Thus Equation (11) holds with $\epsilon = 0$, and Theorem 6 together with Equation (13) gives the bound.

Obtaining $\Theta(1)$ kernel calibration rates. Having seen logarithmic rates, we now define a mechanism that leads to bounded rates: in words, every sufficiently small expected correction can be realized, for every hypothetical mean. The next corollary proves $O(1)$ rates for both this mechanism, and another, even simpler, "endpoint" mechanism, which collectively characterize the above cosine example.

Definition 12 (Strict correction on a finite forecast set)**.** A finite nonempty set $P \subset [0, 1]$ admits *strict correction* if its feature span $\mathcal{H}_P = \operatorname{span}\{\phi(p) : p \in P\}$ is nonzero and some $\delta > 0$ satisfies

$$\{z \in \mathcal{H}_P : \|z\| \le \delta\} \subseteq \operatorname{conv}\{(q - p)\phi(p) : p \in P\} \qquad (q \in [0, 1]).$$

Corollary 3 (Bounded-error kernels)**.** *If a finite forecast set admits strict correction, Algorithm 1, using orthonormal coordinates of $(y - p)\phi(p)$ as gains and a suitable fixed learning rate, has $\sup_T \mathbb{E}[\mathrm{kCE}_{K,T}] < \infty$. If instead $\phi(0) = \alpha v$, $\phi(1) = \beta v$ for a unit vector v and $\alpha, \beta > 0$, use scalar gains αy at forecast 0 and $\beta(y - 1)$ at forecast 1. The same algorithm then has $\mathrm{kCE}_{K,T} \le 3 \max\{\alpha, \beta\}$ on every path, with $\eta = 1/(2 \max\{\alpha, \beta\})$.*

Strict correction verifies (11) with $\epsilon > 0$; endpoint balancing uses a separate scalar invariant (full proofs in Appendix E). In positive finite rank, a strictly correcting set exists if no nonzero RKHS test f has a threshold r with $(r - p)f(p) \ge 0$ for every p (Lemmas E.4 and E.5). This yields the cosine bound beyond one turn, while endpoint balancing covers the boundary case.

An $\Omega(\log T)$ lower bound mechanism. Finally, we give a two-test criterion that certifies the tightness of logarithmic rates. Outcome 1 increases the first test's score, while outcome 0 increases the negative of the second score. A support restriction limits how much the latter choice can undo the first gain. Recall that $V_T^{\mathrm{obl}}(K)$ is the infimum, over learners, of their largest expected error on a fixed outcome sequence (Lemma B.1).

Theorem 7 (Logarithmic lower bound)**.** *Suppose $f, g \in \mathcal{H}_K$ are nonnegative, have norms at most one, and, for some $0 < \epsilon < 1$ and $\eta > 0$, satisfy*

$$f(p) = 0 \quad (p > 1 - \epsilon), \qquad (1 - p)f(p) + pg(p) \ge \eta \quad (p \in [0, 1]). \tag{15}$$

Then for every $T \ge 2$, one can construct an oblivious adversary that witnesses

$$V_T(K) \ge V_T^{\mathrm{obl}}(K) \ge \tfrac{\eta}{4(1+\log(2/\epsilon))} \log T. \tag{16}$$

We give a proof sketch here; the full proof is in Appendix F.1. Fix a learner and let B be its worst-case expected error at time T. After any fixed sequence of $t \le T$ outcomes, both expected test scores lie in $[-B, B]$: otherwise, completing the sequence with all ones or all zeros would contradict the bound at time T. Expectations are over the learner's randomness. Using (15), we can choose each next outcome to increase a certain potential built from these two expected scores: we construct this potential to be linear in the first mean score, and exponential in the negative second score. Under these choices, both scores can remain in $[-B, B]$ for at most $\exp(O(B))$ rounds, giving the logarithmic lower bound. The oblivious minimax lemma then supplies one hard outcome distribution for all learners. Appendix F.1 gives the full proof, including the constant in (16).

REPRODUCIBILITY STATEMENT

This is a theoretical work and includes no empirical evaluations. All theoretical results are rigorously proved in the manuscript; we give proof sketches for the key results in the main text, with the full proofs of all results located in the Appendices.

AI USE STATEMENT

Generative AI was used to assist with the formulation of mathematical claims, proof development, and the writing, checking, and revision of proof arguments, to organize and integrate technical material, to search for related literature, and to draft, edit, illustrate, and typeset the manuscript. We have reviewed all AI-assisted work. The work is theoretical; no empirical dataset or experimental result is reported. This disclosure covers the preparation of the present manuscript. Responsibility for the final submitted text, claims, proofs, and references rests with the authors.

REFERENCES

Eshwar Ram Arunachaleswaran, Natalie Collina, Aaron Roth, and Mirah Shi. An elementary predictor obtaining $2\sqrt{T}+1$ distance to calibration. In *ACM–SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1366–1370, 2025. doi: 10.1137/1.9781611978322.41. URL https://doi.org/10.1137/1.9781611978322.41.

David Blackwell. An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6(1):1–8, 1956. doi: 10.2140/pjm.1956.6.1. URL https://doi.org/10.2140/pjm.1956.6.1.

Jarosław Błasiok, Parikshit Gopalan, Lunjia Hu, and Preetum Nakkiran. A unifying theory of distance from calibration. In *ACM Symposium on Theory of Computing (STOC)*, pp. 1727–1740, 2023. doi: 10.1145/3564246.3585182. URL https://doi.org/10.1145/3564246.3585182.

Nicolò Cesa-Bianchi and Gábor Lugosi. Potential-based algorithms in on-line prediction and game theory. *Machine Learning*, 51(3):239–261, 2003. doi: 10.1023/A:1022901500417. URL https://doi.org/10.1023/A:1022901500417.

Yuval Dagan, Constantinos Daskalakis, Maxwell Fishelson, Noah Golowich, Robert Kleinberg, and Princewill Okoroafor. Breaking the $T^{2/3}$ barrier for sequential calibration. In *ACM Symposium on Theory of Computing (STOC)*, pp. 2007–2018, 2025. doi: 10.1145/3717823.3718178. URL https://doi.org/10.1145/3717823.3718178.

A. Philip Dawid. The well-calibrated Bayesian. *Journal of the American Statistical Association*, 77 (379):605–610, 1982. doi: 10.1080/01621459.1982.10477856. URL https://doi.org/10.1080/01621459.1982.10477856.

Cynthia Dwork, Chris Hays, Nicole Immorlica, Juan C. Perdomo, and Pranay Tankala. From fairness to infinity: Outcome-indistinguishable (omni)prediction in evolving graphs. In *Conference on Learning Theory (COLT)*, volume 291 of *Proceedings of Machine Learning Research*, pp. 1564–1637, 2025. URL https://proceedings.mlr.press/v291/dwork25a.html.

Gabriele Farina and Juan Carlos Perdomo. Defensive generation. In *Conference on Learning Theory (COLT)*, volume 336 of *Proceedings of Machine Learning Research*, pp. 2399–2427, 2026a. URL https://proceedings.mlr.press/v336/farina26a.html.

Gabriele Farina and Juan Carlos Perdomo. An efficient black-box reduction from online learning to multicalibration, and a new route to Φ-regret minimization. arXiv preprint arXiv:2604.19592, 2026b. URL https://arxiv.org/abs/2604.19592.

Dean P. Foster and Rakesh V. Vohra. Asymptotic calibration. *Biometrika*, 85(2):379–390, 1998. doi: 10.1093/biomet/85.2.379. URL https://doi.org/10.1093/biomet/85.2.379.

Parikshit Gopalan and Lunjia Hu. Calibration through the lens of indistinguishability. *ACM SIGecom Exchanges*, 23(1):51–79, 2025. URL https://arxiv.org/abs/2509.02279.

Parikshit Gopalan, Michael P. Kim, Mihir A. Singhal, and Shengjia Zhao. Low-degree multicalibration. In *Conference on Learning Theory (COLT)*, volume 178 of *Proceedings of Machine Learning Research*, pp. 3193–3234, 2022. URL https://proceedings.mlr.press/v178/gopalan22a.html.

Parikshit Gopalan, Konstantinos Stavropoulos, Kunal Talwar, and Pranay Tankala. The importance of being smoothly calibrated. In *Foundations of Responsible Computing (FORC)*, volume 368 of *Leibniz International Proceedings in Informatics*, pp. 21:1–21:22, 2026. doi: 10.4230/LIPIcs.FORC.2026.21. URL https://doi.org/10.4230/LIPIcs.FORC.2026.21.

Nika Haghtalab, Michael Jordan, and Eric Zhao. A unifying perspective on multi-calibration: Game dynamics for multi-objective learning. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 72464–72506. Curran Associates, Inc., 2023. doi: 10.52202/075280-3169. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/e55edcdb01ac45c839a602f96e09fbcb-Paper-Conference.pdf.

Nika Haghtalab, Mingda Qiao, Kunhe Yang, and Eric Zhao. Truthfulness of calibration measures. In *Advances in Neural Information Processing Systems*, volume 37, 2024. URL https://arxiv.org/abs/2407.13979.

Pooria Joulani, András György, and Csaba Szepesvári. Online learning under delayed feedback. In *International Conference on Machine Learning (ICML)*, volume 28(3) of *Proceedings of Machine Learning Research*, pp. 1453–1461, 2013. URL https://proceedings.mlr.press/v28/joulani13.html.

Sham M. Kakade and Dean P. Foster. Deterministic calibration and Nash equilibrium. In John Shawe-Taylor and Yoram Singer (eds.), *Learning Theory (COLT)*, volume 3120 of *Lecture Notes in Computer Science*, pp. 33–48. Springer, 2004. doi: 10.1007/978-3-540-27819-1_3. URL https://doi.org/10.1007/978-3-540-27819-1_3.

Sham M. Kakade and Dean P. Foster. Deterministic calibration and Nash equilibrium. *Journal of Computer and System Sciences*, 74(1):115–130, 2008. doi: 10.1016/j.jcss.2007.04.017. URL https://doi.org/10.1016/j.jcss.2007.04.017.

Aviral Kumar, Sunita Sarawagi, and Ujjwal Jain. Trainable calibration measures for neural networks from kernel mean embeddings. In *International Conference on Machine Learning (ICML)*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2805–2814, 2018. URL https://proceedings.mlr.press/v80/kumar18a.html.

Daniel Lee, Georgy Noarov, Mallesh Pai, and Aaron Roth. Online minimax multiobjective optimization: Multicalibeating and other applications. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 29051–29063. Curran Associates, Inc., 2022. doi: 10.52202/068431-2106. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/ba942323c447c9bbb9d4b638eadefab9-Paper-Conference.pdf.

Charles Marx, Sofian Zalouk, and Stefano Ermon. Calibration by distribution matching: Trainable kernel calibration metrics. In *Advances in Neural Information Processing Systems*, volume 36, 2023. URL https://arxiv.org/abs/2310.20211.

Charles A. Micchelli, Yuesheng Xu, and Haizhang Zhang. Universal kernels. *Journal of Machine Learning Research*, 7:2651–2667, 2006. URL https://www.jmlr.org/papers/v7/micchelli06a.html.

David Oakes. Self-calibrating priors do not exist. *Journal of the American Statistical Association*, 80 (390):339, 1985.

Vianney Perchet. Exponential weight approachability, applications to calibration and regret minimization. *Dynamic Games and Applications*, 5(1):136–153, 2015. doi: 10.1007/s13235-014-0119-x. URL https://doi.org/10.1007/s13235-014-0119-x.

Vianney Perchet and Shie Mannor. Approachability, fast and slow. In *Conference on Learning Theory (COLT)*, volume 30 of *Proceedings of Machine Learning Research*, pp. 474–488, 2013. URL https://proceedings.mlr.press/v30/Perchet13.html.

Mingda Qiao and Gregory Valiant. Stronger calibration lower bounds via sidestepping. In *ACM Symposium on Theory of Computing (STOC)*, pp. 456–466, 2021. doi: 10.1145/3406325.3451050. URL https://doi.org/10.1145/3406325.3451050.

Mingda Qiao and Eric Zhao. Truthfulness of decision-theoretic calibration measures. In *Conference on Learning Theory (COLT)*, volume 291 of *Proceedings of Machine Learning Research*, pp. 4686–4739, 2025. URL https://proceedings.mlr.press/v291/qiao25a.html.

Mingda Qiao and Letian Zheng. On the distance from calibration in sequential prediction. In *Conference on Learning Theory (COLT)*, volume 247 of *Proceedings of Machine Learning Research*, pp. 4307–4357, 2024. URL https://proceedings.mlr.press/v247/qiao24a.html.

Kent Quanrud and Daniel Khashabi. Online learning with adversarial delays. In *Advances in Neural Information Processing Systems*, volume 28, pp. 1270–1278, 2015.

Maurice Sion. On general minimax theorems. *Pacific Journal of Mathematics*, 8(1):171–176, 1958. doi: 10.2140/pjm.1958.8.171. URL https://doi.org/10.2140/pjm.1958.8.171.

S. Szarek. On the best constants in the Khinchin inequality. *Studia Mathematica*, 58(2):197–208, 1976. doi: 10.4064/sm-58-2-197-208.

Vladimir Vovk. Non-asymptotic calibration and resolution. *Theoretical Computer Science*, 387(1): 77–89, 2007. URL https://arxiv.org/abs/cs/0506004.

Vladimir Vovk, Akimichi Takemura, and Glenn Shafer. Defensive forecasting. In *International Workshop on Artificial Intelligence and Statistics (AISTATS)*, pp. 365–372, 2005. URL https://proceedings.mlr.press/r5/vovk05a.html.

David Widmann, Fredrik Lindsten, and Dave Zachariah. Calibration tests in multi-class classification: A unifying framework. In *Advances in Neural Information Processing Systems*, volume 32, 2019. URL https://arxiv.org/abs/1910.11385.

David Widmann, Fredrik Lindsten, and Dave Zachariah. Calibration tests beyond classification. In *International Conference on Learning Representations (ICLR)*, 2021. URL https://arxiv.org/abs/2210.13355. Corrected version, 2022.

Zihan Zhang. Efficient sequential calibration with $O(T^{2/3-\varepsilon})$ error bound. arXiv preprint arXiv:2607.12928, 2026. URL https://arxiv.org/abs/2607.12928.

A Additional related work

A.1 Smooth calibration and calibration distance

Sequential calibration grew out of statistical and game-theoretic work (Dawid, 1982; Oakes, 1985; Foster & Vohra, 1998). The Lipschitz-test notion studied here originates as *weak calibration* in Kakade & Foster (2004; 2008). The *smooth calibration* terminology appears in Gopalan et al. (2022). Smoothness restricts how sharply tests can distinguish nearby forecasts. Ordinary L_1 calibration allows discontinuous tests and has a separate line of lower and upper bounds (Qiao & Valiant, 2021; Dagan et al., 2025; Zhang, 2026).

Błasiok et al. (2023) introduce calibration distance and lower calibration distance and develop a systematic comparison of calibration measures. Qiao & Zheng (2024) formulate these distances for sequential prediction and prove square-root upper bounds, a cubic-root lower bound for smooth calibration error, and an additive $O(\sqrt{T})$ comparison between the two distances. Arunachaleswaran et al. (2025) give a particularly simple forecaster attaining the square-root upper bound for calibration distance. Our contribution is the matching lower bound. One fixed pair of tests witnesses it for all three objectives under a single oblivious adversary at each horizon.

The main obstacle to this lower bound is the learner's ability to offset earlier errors by changing its forecasts. Qiao & Zheng (2024) show that a learner can keep expected calibration distance polylogarithmic on independent fair bits. Thus fluctuations at the truthful forecast alone are insufficient for a minimax lower bound. This freedom to improve measured calibration by changing predictions also connects to work on truthfulness in sequential prediction (Haghtalab et al., 2024; Qiao & Zhao, 2025). Our two-test construction prevents the learner from keeping both accumulated scores small simultaneously.

A.2 Kernel calibration

Defensive forecasting originates with Vovk et al. (2005); Vovk (2007) gives the kernel guarantee used in (9). It bounds the cumulative feature sum and hence controls every test in the RKHS unit ball. Our square-root lower bounds match this guarantee for broad classes of kernels, including the Gaussian and Laplace kernels. The bounded and logarithmic examples then show how the optimal rate can change with the kernel. We note that Vovk (2007) also proves forecast-dependent lower bounds for online kernel calibration (where the adversary is allowed to observe the realized forecasts); moreover, the right-side of those bounds can vanish at endpoint forecasts (i.e. as $p_t \to 0$ or $p_t \to 1$). By contrast, our lower bounds force $\Omega(\sqrt{T})$ kernel calibration error via a single oblivious adversary.

Kernel mean embeddings also underlie trainable calibration measures (Kumar et al., 2018; Marx et al., 2023) and statistical calibration tests (Widmann et al., 2019; 2021). The choice of kernel determines the tests being controlled and the meaning of a small error. For example, Błasiok et al. (2023, Theorems 8.5–8.6) show that Laplace kernel calibration is polynomially related to calibration distance, whereas Gaussian kernel calibration can be much smaller. Our radial theorem establishes the same online exponent for both kernels. These results describe two aspects of kernel choice, its relationship to calibration distance and its attainable online rate.

RKHS representations have also been used to obtain online outcome-indistinguishability and omniprediction guarantees (Dwork et al., 2025; Farina & Perdomo, 2026a). Such guarantees control residuals against a specified function class, which can include functions of side information. Our results concern how the minimax rate depends on a fixed kernel of the predicted binary probability. They identify how the location and direction of features along the probability interval determine which corrections to accumulated error a learner can exploit. The cosine and rank-two families make this dependence explicit through sharp transitions between rates.

A.3 Approachability and fast rates

Blackwell's approachability theorem (Blackwell, 1956) gives a general route to square-root bounds for cumulative calibration error. Exponential weights and potential methods are standard tools in this framework (Cesa-Bianchi & Lugosi, 2003; Perchet, 2015). Lee et al. (2022) develop a general framework for online minimax multiobjective optimization using exponential weights and minimax

responses. Their algorithm competes with the benchmark in which the adversary moves first at each round. Haghtalab et al. (2023) give a unifying treatment of multicalibration through no-regret and best-response dynamics.

Perchet & Mannor (2013) give geometric conditions for $O(1/T)$ expected distance of the average payoff to a target. For the zero target, this gives a uniformly bounded expected norm of the cumulative payoff. Our Algorithm 1 uses paired exponential weights to control both signs of each coordinate and includes a quadratic term in the score used to select an action distribution. Condition (11) requires the expected correction to offset this quadratic term. Under this condition, Theorem 6 bounds the expected norm of the cumulative vector by $O(\log(T+1))$ when $\epsilon = 0$ and by a constant when $\epsilon > 0$, for fixed parameters. For kernel calibration, finite cancellation covers and strict correction verify these respective conditions. The same algorithm therefore exploits different geometric properties to obtain logarithmic or bounded error. Our lower bounds establish sharpness for concrete kernel families.

A.4 Delayed feedback

Delayed online learning is studied by Joulani et al. (2013); Quanrud & Khashabi (2015). Farina & Perdomo (2026b, Section 3.2.1) transfer delayed-learning guarantees to multicalibration, with square-root dependence on total delay. Our main delayed contribution is the matching lower bound for smooth calibration, lower calibration distance, and calibration distance. For each horizon and budget, we construct a deterministic schedule and one oblivious outcome law that give the lower bounds simultaneously against all randomized forecasters. Repeating a hidden bit throughout each block turns the two-test fluctuation into a weighted sum. Choosing the block lengths within the delay budget yields the rate $\Theta(\min\{T, \sqrt{T+D}\})$ proved in Theorem 3.

B Two-test lower bounds

This appendix gives the full proof of Theorem 1 and develops the auxiliary arguments used in this proof and in later lower bounds. We also extend the two-test argument to settings where positive expected increments can compensate for insufficient separation between the tests.

B.1 Randomized forecasters and one oblivious hard law

For fixed horizon T, a pure forecasting strategy is specified by finitely many forecasts, one for each possible outcome history. The resulting strategy space is compact, which allows us to apply a minimax theorem without restricting the learner to finitely many forecast values.

Lemma B.1 (Oblivious finite-horizon minimax)**.** *Let $\mathcal{L}(p, y)$ be nonnegative and continuous in $p \in [0,1]^T$ for each $y \in \{0,1\}^T$. A forecaster may use arbitrary private randomness and all preceding outcomes, forecasts, and announcements. Consider the auxiliary fixed-sequence value:*

$$V_T^{\mathrm{obl}}(\mathcal{L}) := \inf_A \max_{y\in\{0,1\}^T} \mathbb{E}_A[\mathcal{L}(p,y)] \le \inf_A \sup_B \mathbb{E}_{A,B}[\mathcal{L}(p,y)] = V_T(\mathcal{L}). \tag{B.1}$$

Then

$$V_T^{\mathrm{obl}}(\mathcal{L}) = \max_{Q\in\Delta(\{0,1\}^T)} \inf_A \mathbb{E}_{A,Q}[\mathcal{L}(p,Y)]. \tag{B.2}$$

The extrema over mixed strategies are attained. The same statement holds for any fixed deterministic feedback schedule, and for a game with finitely many hidden constant-outcome blocks.

Proof. Once all the learner's randomness is fixed, its earlier forecasts, announcements, and private state are determined recursively by the preceding outcomes. Its rule is therefore specified by one forecast $F(y_{<t}) \in [0,1]$ at each history, so these rules form the compact cube $\mathcal{F}_T := \prod_{t=1}^{T} [0,1]^{\{0,1\}^{t-1}} \cong [0,1]^{2^T-1}$. Every randomized learner induces a distribution $\rho \in \mathcal{P}(\mathcal{F}_T)$; conversely, any such distribution is implementable by sampling a rule and announcing point masses at its prescribed forecasts.

Let $L(F, y) = \mathcal{L}(p^F(y), y)$ be the loss of rule F on outcome sequence y, where $p^F(y) = (F(y_{<t}))_{t=1}^{T}$ is the sequence of forecasts made by F on y. For each fixed y, this is continuous

in F because $\mathcal{L}(p, y)$ is continuous in p and $p^F(y)$ uses only finitely many coordinates of F. The expected payoff is therefore continuous and affine in ρ and in the outcome distribution $Q \in \Delta(\{0,1\}^T)$. Both mixed-strategy spaces are compact and convex. Sion's minimax theorem (Sion, 1958) gives

$$\max_Q \inf_\rho \mathbb{E}_{\rho,Q}[L] = \inf_\rho \max_Q \mathbb{E}_{\rho,Q}[L].$$

The extrema over mixed strategies are attained. For each fixed ρ, the payoff is linear in Q and is maximized by a point mass. Thus the right side equals $V_T^{\mathrm{obl}}(\mathcal{L})$, which proves (B.2).

For a fixed delay schedule, use the finitely many time-indexed histories of revealed labelled outcomes instead of complete outcome prefixes. Fixing the private seed still determines past predictions, and the resulting pure-strategy space is again a finite product of compact intervals. The block game is a special case. The same argument applies. □

In particular, it suffices to construct, for each fixed mixed learner, a response law independent of its *realized* seed. The lemma produces one law hard for every learner, and $V_T \geq V_T^{\mathrm{obl}}$ transfers the bound to the main game.

Remark B.1 (Why the two values are kept separate). They can differ even for a continuous, nonnegative transcript loss. At $T = 2$, take

$$\mathcal{L}(p, y) = 1 - |p_1 - y_2| + p_1(1 - p_1).$$

For fixed y_2, uniformly randomizing p_1 over $\{0, 1\}$ gives expected loss $1/2$. Conversely, averaging $y_2 = 0, 1$ gives $1/2 + p_1(1 - p_1) \geq 1/2$, so $V_2^{\mathrm{obl}}(\mathcal{L}) = 1/2$. In the adaptive game Nature observes p_1 before choosing y_2. The resulting worst loss is

$$\max\{p_1, 1 - p_1\} + p_1(1 - p_1) \geq 3/4,$$

with equality at $p_1 = 1/2$. Thus $V_2(\mathcal{L}) = 3/4$. This example is not a calibration loss; it rules out a general identification of the two protocols.

B.2 Khinchin inequality

The following inequality gives a lower bound on the expected absolute value of a weighted Rademacher sum. We will use the equal-weight case in the proof of Theorem 1 and the general weighted form for delayed feedback.

Lemma B.2 (Khinchin inequality (Szarek, 1976)). *For deterministic real weights $w_1, \ldots, w_n$ and independent fair signs $\varepsilon_1, \ldots, \varepsilon_n$,*

$$\mathbb{E}\left[\left|\sum_{j=1}^n w_j \varepsilon_j\right|\right] \geq \frac{1}{\sqrt{2}} \left(\sum_{j=1}^n w_j^2\right)^{1/2}. \tag{B.3}$$

B.3 Proof of Theorem 1

Write $\mathcal{L}(p, y) = \max\{|S_{g_+}(T)|, |S_{g_-}(T)|\}$ and $\mathcal{E}_T = a\gamma\, \mathbb{E}[|\sum_{t=1}^T \varepsilon_t|]$. We first fix a randomized learner A and construct an outcome distribution Q_A with $\mathbb{E}_{A,Q_A}[\mathcal{L}] \geq \mathcal{E}_T$, and then apply the minimax theorem to obtain one distribution for all learners.

For each history $y_{<t}$, let $\bar{\nu}_t(y_{<t})$ be the distribution of the learner's next forecast when its preceding outcomes are prescribed to be $y_{<t}$. This averages over all the learner's randomness, including its earlier forecasts and private state; it need not equal the law announced on a particular play. Choose a common response $q_t(y_{<t}) \in [a, 1 - a]$ to $\bar{\nu}_t(y_{<t})$, and define an outcome distribution Q_A by $Y_t \mid Y_{<t} \sim \mathrm{Bernoulli}(q_t(Y_{<t}))$. Since there are only finitely many possible histories, choosing one common response for each history defines a measurable rule and hence a valid conditional law for Y_t. This specifies a law on complete outcome sequences using the learner's strategy but not its realized randomness. Sample the sequence independently of that randomness.

Abbreviate $\bar{\nu}_t = \bar{\nu}_t(Y_{<t})$ and $q_t = q_t(Y_{<t})$. We use $\mathbb{E}_A$ to average over the learner's randomness with Y held fixed. For $i \in \{+, -\}$, consider

$$\mathbb{E}_A[S_{g_i}(T)] = \sum_{t=1}^T \mathbb{E}_{P\sim\bar{\nu}_t}[(Y_t - P)g_i(P)], \quad M_i(Y) := \sum_{t=1}^T (Y_t - q_t)\mathbb{E}_{P\sim\bar{\nu}_t}[g_i(P)].$$

Based on the definition of common response, the first quantity is lower bounded by the second:

$$\mathbb{E}_A[S_{g_i}(T)] - M_i(Y) = \sum_{t=1}^{T} r_{g_i}(\bar{\nu}_t, q_t) \geq 0.$$

We can show that $\mathbb{E}_{Q_A}[M_i(Y)] = 0$ as each coefficient is determined by the preceding outcomes, conditional on which $Y_t - q_t$ has mean zero.

Define $b_t = \mathbb{E}_{P\sim\bar{\nu}_t}[d(P)] \geq \gamma$ and $N = \sum_t b_t(Y_t - q_t) = (M_+(Y) - M_-(Y))/2$. Then we have

$$\begin{aligned}\mathbb{E}_{A,Q_A}[\mathcal{L}] &= \mathbb{E}_{A,Q_A}[\max_i |S_{g_i}(T)|] \geq \mathbb{E}_{Q_A}\left[\max_i \mathbb{E}_A[S_{g_i}(T)]\right] \geq \mathbb{E}_{Q_A}\left[\max_i M_i(Y)\right] \\ &= \frac{1}{2}\mathbb{E}_{Q_A}[M_+(Y) + M_-(Y)] + \frac{1}{2}\mathbb{E}_{Q_A}[|M_+(Y) - M_-(Y)|] = \mathbb{E}_{Q_A}[|N|].\end{aligned}$$

where the second line follows from the identity $\max\{u, v\} = (u + v + |u - v|)/2$.

We compare N with a Rademacher sum by generating the same outcome distribution Q_A in a different way. Independently of the learner's randomness, let $(\varepsilon_t)_{t\geq 1}$ be i.i.d. with $\varepsilon_t \sim \mathrm{Rad}$, and generate the outcomes successively by

$$\mathbb{P}(Y_t = 1 \mid Y_{<t}, \varepsilon_{1:T}) = q_t + \frac{a\gamma}{b_t}\varepsilon_t.$$

These probabilities lie in $[0, 1]$, since $a\gamma/b_t \leq a$ and $q_t \in [a, 1-a]$. Earlier outcomes use only earlier Rademacher variables, so ε_t is independent of $Y_{<t}$. Averaging over it gives $\mathbb{P}(Y_t = 1 \mid Y_{<t}) = q_t$, and hence Y still has law Q_A. Moreover,

$$\mathbb{E}[b_t(Y_t - q_t) \mid Y_{<t}, \varepsilon_{1:T}] = b_t(\mathbb{E}[Y_t \mid Y_{<t}, \varepsilon_{1:T}] - q_t) = a\gamma\varepsilon_t,$$

$$\mathbb{E}[N \mid \varepsilon_{1:T}] = \sum_t \mathbb{E}[b_t(Y_t - q_t) \mid \varepsilon_{1:T}] = \sum_t \mathbb{E}[\mathbb{E}[b_t(Y_t - q_t) \mid Y_{<t}, \varepsilon_{1:T}] \mid \varepsilon_{1:T}] = a\gamma \sum_{t=1}^{T} \varepsilon_t.$$

Consequently,

$$\mathbb{E}[|N|] \geq \mathbb{E}\left[\left|\mathbb{E}[N \mid \varepsilon_{1:T}]\right|\right] = a\gamma\,\mathbb{E}\left[\left|\sum_{t=1}^{T} \varepsilon_t\right|\right] = \mathcal{E}_T.$$

Together with the preceding bound $\mathbb{E}_{A,Q_A}[\mathcal{L}] \geq \mathbb{E}[|N|]$, this proves $\mathbb{E}_{A,Q_A}[\mathcal{L}] \geq \mathcal{E}_T$.

Applying the Khinchin inequality (Lemma B.2) with all weights equal to one gives $\mathcal{E}_T \geq a\gamma\sqrt{T/2}$.

For each learner A, the expected loss under Q_A is an average of the expected losses on fixed outcome sequences. Hence

$$\max_{y\in\{0,1\}^T} \mathbb{E}_A[\mathcal{L}(p, y)] \geq \mathbb{E}_{A,Q_A}[\mathcal{L}(p, Y)] \geq \mathcal{E}_T.$$

The loss $\mathcal{L}(p, y)$ is nonnegative and continuous in p for each fixed y because the tests are continuous. Applying Lemma B.1 gives

$$\max_Q \inf_A \mathbb{E}_{A,Q}[\mathcal{L}(p, Y)] = \inf_A \max_y \mathbb{E}_A[\mathcal{L}(p, y)] \geq \mathcal{E}_T.$$

The bound proved above implies that a distribution Q_T attaining the maximum on the left guarantees the same lower bound against every randomized learner. Thus Q_T defines the required oblivious adversary.

B.4 When positive drift replaces test separation

Theorem 1 assumes that, for every forecast value p, the half-difference between the tests satisfies $d(p) = \frac{g_+(p)-g_-(p)}{2} \geq \gamma$. However, some applications have tests whose difference changes sign or becomes small. The same argument continues to work if the expected increments are sufficiently positive when this happens. We will use this extension later for kernels whose relevant tests do not remain uniformly separated. The smooth-calibration lower bound uses only Theorem 1.

We begin with the following inequality. The proof of Theorem 1 implicitly uses the weaker version

$$\mathbb{E}\left[\max\{|s_+|, |s_-|\}\right] \geq \mathbb{E}\left[\left|\frac{M_+ - M_-}{2}\right|\right].$$

Lemma B.3 (Two-coordinate drift inequality). *Suppose $s_i = M_i + D_i$ for $i \in \{+, -\}$, where the variables are integrable, $\mathbb{E}[M_i] = 0$, and $D_i \geq 0$ almost surely. Then*

$$\mathbb{E}\big[\max\{|s_+|, |s_-|\}\big] \geq \mathbb{E}\left[\left|\frac{M_+ - M_-}{2}\right|\right] + \mathbb{E}\big[\min\{D_+, D_-\}\big]. \tag{B.4}$$

Proof. We have

$$\max_i |s_i| \geq \max_i s_i = \max_i (M_i + D_i) \geq \max_i M_i + \min_i D_i.$$

Moreover, the identity $\max\{u, v\} = \frac{u+v}{2} + \frac{|u-v|}{2}$ and the assumptions $\mathbb{E}[M_+] = \mathbb{E}[M_-] = 0$ give

$$\mathbb{E}[\max_i M_i] = \frac{1}{2}\mathbb{E}[M_+ + M_-] + \mathbb{E}\left[\left|\frac{M_+ - M_-}{2}\right|\right] = \mathbb{E}\left[\left|\frac{M_+ - M_-}{2}\right|\right].$$

Taking expectations in the first chain of inequalities and substituting the identity above proves the lemma. □

Using this inequality, we can extend Theorem 1 to cases where the two tests are not uniformly separated.

Theorem B.1 (Compensated two-test lower bound). *For continuous g_+, g_- and $a \in (0, 1/2)$, $\gamma > 0$, the conclusion of Theorem 1 also holds if every forecast distribution ν admits $q \in [a, 1-a]$ such that*

$$\min_{i \in \{+,-\}} r_{g_i}(\nu, q) \geq a\,(\gamma - |\mathbb{E}_{P\sim\nu}[d(P)]|)_+\,, \tag{B.5}$$

where recall the difference notation $d = \frac{g_+ - g_-}{2}$. No sign restriction on d is needed.

The hypotheses of Theorem 1 imply (B.5), since $d \geq \gamma$ makes its right-hand side zero and a common response makes both expected increments nonnegative.

To see why the condition suffices, use the forecast laws $\bar{\nu}_t$ from the two-test construction, write $b_t = \mathbb{E}_{P\sim\bar{\nu}_t}[d(P)]$, and set $c_t = (\gamma - |b_t|)_+$. Inequality (B.4) retains at least $a\mathbb{E}[\sum_t c_t]$ of positive drift. In the coupling, replace $a\gamma/b_t$ by $a\min\{\gamma, |b_t|\}/b_t$, interpreted as zero at $b_t = 0$. The fair-walk comparison loses at most $a\mathbb{E}[\sum_t c_t]$, exactly the amount retained. We now give the full conditional-expectation calculation.

Proof of Theorem B.1. Write $\mathcal{L}(p, y) = \max\{|S_{g_+}(T)|, |S_{g_-}(T)|\}$ and $\mathcal{E}_T = a\gamma\,\mathbb{E}[|\sum_{t=1}^T \varepsilon_t|]$. We first fix a randomized learner A and construct an outcome distribution Q_A with $\mathbb{E}_{A,Q_A}[\mathcal{L}] \geq \mathcal{E}_T$, and then apply the minimax theorem to obtain one distribution for all learners.

For each history $y_{<t}$, let $\bar{\nu}_t(y_{<t})$ be the distribution of the learner's next forecast when its preceding outcomes are prescribed to be $y_{<t}$. This averages over all the learner's randomness, including its earlier forecasts and private state; it need not equal the law announced on a particular play. Choose $q_t(y_{<t}) \in [a, 1-a]$ satisfying (B.5) with $\nu = \bar{\nu}_t(y_{<t})$, and define an outcome distribution Q_A by $Y_t \mid Y_{<t} \sim \mathrm{Bernoulli}(q_t(Y_{<t}))$. Since there are only finitely many possible histories, choosing one such response for each history defines a measurable rule and hence a valid conditional law for Y_t. This specifies a law on complete outcome sequences using the learner's strategy but not its realized randomness. Sample the sequence independently of that randomness.

Abbreviate $\bar{\nu}_t = \bar{\nu}_t(Y_{<t})$ and $q_t = q_t(Y_{<t})$. We use $\mathbb{E}_A$ to average over the learner's randomness with Y held fixed. For $i \in \{+, -\}$, define the centered fluctuation and drift by

$$M_i(Y) = \sum_{t=1}^T (Y_t - q_t)\mathbb{E}_{P\sim\bar{\nu}_t}[g_i(P)], \qquad D_i(Y) = \sum_{t=1}^T r_{g_i}(\bar{\nu}_t, q_t).$$

Then

$$\mathbb{E}_A[S_{g_i}(T)] = M_i(Y) + D_i(Y). \tag{B.6}$$

We can show that $\mathbb{E}_{Q_A}[M_i(Y)] = 0$ as each coefficient is determined by the preceding outcomes, conditional on which $Y_t - q_t$ has mean zero. For the two tests, (B.5) gives nonnegative drift on every path. Define

$$b_t = \mathbb{E}_{P\sim\bar{\nu}_t}[d(P)], \qquad c_t = (\gamma - |b_t|)_+, \qquad N = \sum_t b_t(Y_t - q_t) = \frac{M_+(Y) - M_-(Y)}{2}.$$

Conditional Jensen and Lemma B.3 with $s_i = \mathbb{E}_A[S_{g_i}(T)]$ give

$$\begin{aligned}\mathbb{E}_{A,Q_A}\left[\max_i |S_{g_i}(T)|\right] &\geq \mathbb{E}_{Q_A}\left[\max_i \left|\mathbb{E}_A[S_{g_i}(T)]\right|\right] \\ &\geq \mathbb{E}_{Q_A}[|N|] + \mathbb{E}_{Q_A}\left[\min_i D_i(Y)\right] \\ &\geq \mathbb{E}_{Q_A}[|N|] + a\mathbb{E}_{Q_A}[\sum_t c_t]. \end{aligned} \tag{B.7}$$

The last inequality uses (B.5) and the fact that a minimum of sums is at least the sum of the termwise minima.

Define

$$\beta_t = \begin{cases} a\min\{\gamma, |b_t|\}/b_t, & b_t \neq 0, \\ 0, & b_t = 0. \end{cases}$$

We compare N with a Rademacher sum by generating the same outcome distribution Q_A in a different way. Independently of the learner's randomness, let $(\varepsilon_t)_{t\geq 1}$ be i.i.d. with $\varepsilon_t \sim \mathrm{Rad}$, and generate the outcomes successively by

$$\mathbb{P}(Y_t = 1 \mid Y_{<t}, \varepsilon_{1:T}) = q_t + \beta_t\varepsilon_t. \tag{B.8}$$

These probabilities lie in $[0,1]$, since $|\beta_t| \leq a$ and $q_t \in [a, 1-a]$. Earlier outcomes use only earlier Rademacher variables, so ε_t is independent of $Y_{<t}$. Averaging over it gives $\mathbb{P}(Y_t = 1 \mid Y_{<t}) = q_t$, and hence Y still has law Q_A. Moreover,

$$\mathbb{E}[b_t(Y_t - q_t) \mid Y_{<t}, \varepsilon_{1:T}] = a\min\{\gamma, |b_t|\}\varepsilon_t = a(\gamma - c_t)\varepsilon_t.$$

Taking another conditional expectation gives

$$\mathbb{E}[N \mid \varepsilon_{1:T}] = a\gamma\sum_t \varepsilon_t - a\sum_t \varepsilon_t\mathbb{E}[c_t \mid \varepsilon_{1:T}].$$

Jensen and the triangle inequality, using $c_t \geq 0$, now give

$$\mathbb{E}[|N|] \geq a\gamma\mathbb{E}\left[\left|\sum_t \varepsilon_t\right|\right] - a\mathbb{E}[\sum_t c_t]. \tag{B.9}$$

Combining (B.7) and (B.9) cancels the last terms and proves $\mathbb{E}_{A,Q_A}[\mathcal{L}] \geq \mathcal{E}_T$.

Applying the Khinchin inequality (Lemma B.2) with all weights equal to one gives $\mathcal{E}_T \geq a\gamma\sqrt{T/2}$.

For each learner A, the expected loss under Q_A is an average of the expected losses on fixed outcome sequences. Hence

$$\max_{y\in\{0,1\}^T} \mathbb{E}_A[\mathcal{L}(p,y)] \geq \mathbb{E}_{A,Q_A}[\mathcal{L}(p,Y)] \geq \mathcal{E}_T.$$

The loss $\mathcal{L}(p,y)$ is nonnegative and continuous in p for each fixed y because the tests are continuous. Applying Lemma B.1 gives

$$\max_Q \inf_A \mathbb{E}_{A,Q}[\mathcal{L}(p,Y)] = \inf_A \max_y \mathbb{E}_A[\mathcal{L}(p,y)] \geq \mathcal{E}_T.$$

The bound proved above implies that a distribution Q_T attaining the maximum on the left guarantees the same lower bound against every randomized learner. Thus Q_T defines the required oblivious adversary. □

Corollary B.1 (Conversion to kernel error)**.** *If two tests satisfy Theorem 1 or Theorem B.1, belong to $\mathcal{H}_K$, and $M = \max\{\|g_+\|_{\mathcal{H}_K}, \|g_-\|_{\mathcal{H}_K}\}$, then*

$$V_T(K) \geq V_T^{\mathrm{obl}}(K) \geq \frac{a\gamma}{\sqrt{2}M}\sqrt{T}.$$

Proof. The reproducing property gives $|S_{g_i}(T)| \leq \|g_i\|_{\mathcal{H}_K}\,\mathrm{kCE}_{K,T}$ on every transcript. Here $M > 0$, since a zero pair cannot satisfy either theorem with $\gamma > 0$. Apply Theorem 1 or Theorem B.1, respectively. □

C SMOOTH CALIBRATION AND CALIBRATION DISTANCE

This appendix provides the technical details for Section 3. It gives the comparisons among the three calibration metrics, proves the sharper bound for our pair of tests, and gives the Hilbert-space representation used for the upper bounds.

C.1 LOWER CALIBRATION DISTANCE AND THE COMPARISON USED FOR THE LOWER BOUND

Definition 5 gives the cumulative finite-sample formulation of lower calibration distance from Qiao & Zheng (2024, Definition 6), following Błasiok et al. (2023, Definition 5.3). For arbitrary probability measures, the calibration constraint means

$$\sum_t \int_B (y_t - u)\lambda_t(\,\mathrm{d}u) = 0 \quad \text{for every Borel set } B \subseteq [0,1].$$

Allowing arbitrary probability measures gives the same infimum. For any $\epsilon > 0$, group comparison forecast values into finitely many intervals of width at most ϵ, discard intervals with zero aggregate mass, and replace every value in each remaining interval by its mean under the aggregate measure $T^{-1}\sum_t \lambda_t$. The calibration constraint makes this mean equal to the aggregate outcome mean in the interval; the new family is calibrated and changes the cost by at most $T\epsilon$. An exactly calibrated sequence gives a feasible family by $\lambda_t = \delta_{u_t}$, so $\mathrm{LowerCalDist}_T \le \mathrm{CalDist}_T$.

Lemma C.1 (Comparison with lower calibration distance)**.** *For every binary transcript,*

$$\tfrac{1}{2}\,\mathrm{smCE}_T \le \mathrm{LowerCalDist}_T \le \mathrm{CalDist}_T\,.$$

Proof. For $\|f\|_\infty \le 1$ and $\mathrm{Lip}(f) \le 1$,

$$|(y-p)f(p) - (y-u)f(u)| \le |p-u|\,|f(p)| + |y-u|\,|f(p) - f(u)| \le 2|p-u|.$$

The aggregate score at the comparison forecasts vanishes, since

$$\sum_t \mathbb{E}_{U\sim\lambda_t}\big[(y_t - U)f(U)\big] = \sum_u f(u)\sum_t (y_t - u)\lambda_t(\{u\}) = 0.$$

Taking expectations under $U \sim \lambda_t$ in the first inequality and summing over rounds gives $|S_f(T)| \le 2\sum_t \mathbb{E}_{U\sim\lambda_t}\big[|p_t - U|\big]$. Take the supremum over f and then the infimum over calibrated families (λ_t). □

Lemma C.2 (Known upper comparisons for a completed transcript)**.** *There is an absolute constant C_0 such that, for every binary transcript,*

$$\mathrm{LowerCalDist}_T \le 2\,\mathrm{smCE}_T, \qquad \mathrm{CalDist}_T \le \mathrm{LowerCalDist}_T + C_0\sqrt{T} \le 2\,\mathrm{smCE}_T + C_0\sqrt{T}. \tag{C.1}$$

The first inequality is the smooth-calibration comparison of Błasiok et al. (2023, Theorem 7.3), in cumulative normalization; see also Qiao & Zheng (2024, Lemma 8). The second is Qiao & Zheng (2024, Theorem 2). These are existing deterministic comparison theorems about a fixed completed transcript, not assumptions on how its forecasts were produced. In particular, they apply unchanged with delayed feedback. We use these published comparisons for the distance upper bound; all new lower-bound and forecasting arguments are proved here.

For completeness, calibration distance is a continuous function of p for fixed y. Every distinct value in a calibrated sequence equals the empirical outcome average of its nonempty group, hence belongs to the finite set $\{k/n : 1 \le n \le T,\ 0 \le k \le n\}$. Thus $\mathcal{C}(y)$ is finite and contains $u = y$, and $\mathrm{CalDist}_T$ is the minimum of finitely many continuous distance functions. Smooth error is also continuous: the residual-map estimate above gives $|\,\mathrm{smCE}_T(p,y) - \mathrm{smCE}_T(p',y)| \le 2\sum_t |p_t - p'_t|$.

C.2 A SINGLE EXPLICIT PAIR BOUNDS ALL THREE QUANTITIES

We now prove Proposition 1. The only additional point beyond the general comparison above is that, for our two tests in Equation (5), the function $p \mapsto (y-p)f(p)$ is 1-Lipschitz for each $y \in \{0,1\}$.

For f_+, on $[0, 1/3]$ the absolute slope is $1/3$, and on $[2/3, 1]$ it is zero. On the middle interval the derivatives of $(y-p)f_+(p)$ are $2p-2/3$ for $y=0$ and $2p-5/3$ for $y=1$, each of absolute value at most one. Continuity at the joins proves the assertion. The identity $f_-(p) = -f_+(1-p)$ gives the same conclusion by reflection.

Since the two tests are admissible smooth tests, the bound for smooth calibration error is immediate. Since $h_y(p) := (y-p)f(p)$ is 1-Lipscthiz, every calibrated family (λ_t) in Definition 5 therefore satisfies

$$|S_f(T)| = \left|\sum\nolimits_t \mathbb{E}_{U\sim\lambda_t}\left[h_{y_t}(p_t) - h_{y_t}(U)\right]\right| \le \sum\nolimits_t \mathbb{E}_{U\sim\lambda_t}\left[|p_t - U|\right].$$

Taking the infimum proves the bound for lower calibration distance. Calibrated sequences are the special case $\lambda_t = \delta_{u_t}$.

C.3 A Hilbert space containing every smooth test

We represent a Lipschitz test by its initial value and its almost-everywhere derivative. The resulting norm bound makes one forecasting rule control the whole class.

Proposition C.1 (Smooth-test representation)**.** *Let $\mathcal{H} = \mathbb{R} \oplus L^2([0,1])$ and*

$$\phi(p) = (1, \mathbf{1}_{[0,p]}), \qquad K(p,q) = 1 + \min\{p, q\}.$$

The feature map is norm-continuous, $\|\phi(p)\|^2 \le 2$, and every transcript satisfies

$$\mathrm{smCE}_T \le \sqrt{2}\left\|\sum_t (y_t - p_t)\phi(p_t)\right\|. \tag{C.2}$$

Proof. The squared feature distance is $|p-q|$, proving continuity. Every Lipschitz function is absolutely continuous, so an admissible test has $f(p) = f(0) + \int_0^p f'(u)\,\mathrm{d}u$ with $|f'| \le 1$ almost everywhere. Its representing vector $(f(0), f')$ has squared norm at most two. Therefore

$$S_f(T) = \left\langle (f(0), f'), \sum_t (y_t - p_t)\phi(p_t)\right\rangle.$$

Cauchy–Schwarz and the supremum over f give (C.2). □

Combining this representation with the zero-delay case of Proposition G.1 gives the deterministic bound $\mathrm{smCE}_T \le \sqrt{T}$. Together with Lemma C.2, it also gives an $O(\sqrt{T})$ upper bound for calibration distance and lower calibration distance.

D Square-root lower bounds for kernel calibration

Throughout, K is a continuous real positive-semidefinite kernel. Its canonical feature map is norm-continuous because

$$\|\phi(p) - \phi(q)\|^2 = K(p,p) + K(q,q) - 2K(p,q).$$

Every RKHS test is continuous, and $S_f(T) = \langle f, \mathbf{S}_K(T)\rangle$. In another feature representation we use the closed linear span of the feature vectors; a bounded linear measurement of those vectors defines an RKHS function with norm no greater than the norm of the representing vector.

D.1 When two tests admit a common response

The common-response question has a useful scalar form: it asks whether all mixtures of the tests have their signs ordered around a suitable outcome mean.

Definition D.1 (Sign threshold)**.** A continuous test $g : [0,1] \to \mathbb{R}$ has a *sign threshold* at $r \in [0,1]$ if $(r-p)g(p) \ge 0$ for every p. Thus g is nonnegative below r and nonpositive above it; additional zeros are allowed. The zero function has every point as a threshold.

This condition concerns the signs of a test, not a restriction on the forecasts selected by the learner.

Lemma D.1 (Common-response criterion). *Let g_+, g_- be continuous and let $I \subseteq [0,1]$ be a nonempty closed interval. The following are equivalent:*

(i) For every forecast distribution ν, some $q \in I$ has $r_{g_+}(\nu, q), r_{g_-}(\nu, q) \geq 0$.

(ii) Every convex combination $g_\theta = \theta g_+ + (1-\theta) g_-$, $\theta \in [0,1]$, has a sign threshold in I.

Proof. There are two minimax exchanges, with different variables fixed. First fix ν. Choosing between the two tests is equivalent to minimizing over their mixtures, since $\min_i r_{g_i}(\nu, q) = \min_{\theta \in [0,1]} r_{g_\theta}(\nu, q)$. The latter payoff is continuous and affine separately in θ and q. Minimax on the compact intervals therefore gives

$$\max_{q \in I} \min_i r_{g_i}(\nu, q) = \min_{\theta \in [0,1]} \max_{q \in I} r_{g_\theta}(\nu, q).$$

Now minimize over ν and interchange only the two minima. For each fixed θ, the probability measures on $[0,1]$ form a compact convex set in the weak topology, and the payoff is continuous and affine in ν and q. The second minimax exchange is consequently valid and gives

$$\begin{aligned} \min_\nu \max_{q \in I} \min_i r_{g_i}(\nu, q) &= \min_\theta \min_\nu \max_{q \in I} r_{g_\theta}(\nu, q) \\ &= \min_{\theta \in [0,1]} \max_{q \in I} \min_{p \in [0,1]} (q - p) g_\theta(p). \end{aligned} \tag{D.1}$$

In the last line a minimizing measure can be replaced by a point mass. All extrema are attained: the sets are compact, and the continuous payoffs and their extrema vary continuously in the remaining variables.

For each θ, the inner maximum on the right is at most zero, because $p = q$ is an available choice. It equals zero exactly when some $r \in I$ makes $(r - p) g_\theta(p) \geq 0$ for all p, which is a sign threshold. Thus the right side is nonnegative exactly when every convex combination has a threshold. The left side is nonnegative exactly when each ν admits a common response. No joint affinity in (ν, θ, q) is asserted or required. □

D.2 Proof of the stable sign-change criterion

The sign-change condition is stable under small perturbations. We use the resulting interval of zeros to verify the compensated two-test hypothesis.

Proof of Theorem 4(ii). Replace d by $-d$ if necessary so that $d(r) > 0$. Choose a closed interval $J \subset (0,1)$ containing r in its interior, on which $d > 0$ and $m' < 0$. For sufficiently small $\epsilon > 0$, every function

$$f_\theta = m + \theta d, \qquad -\epsilon \leq \theta \leq \epsilon,$$

has negative derivative on J, is positive to the left of J, and is negative to the right. Indeed, the derivative assertion follows from a bound on $\|d'\|_\infty$; outside the interior of J, the strict signs of m have positive compact margins. Each f_θ consequently has a unique zero $r_\theta \in J$, with strict positive and negative signs before and after it.

The implicit-function theorem gives

$$\frac{\mathrm{d} r_\theta}{\mathrm{d}\theta} = -\frac{d(r_\theta)}{m'(r_\theta) + \theta d'(r_\theta)} > 0.$$

Thus the roots vary continuously and strictly increasingly. Set

$$I = [r_{-\epsilon}, r_\epsilon], \qquad a = \min\{r_{-\epsilon}, 1 - r_\epsilon\} \in (0, 1/2), \qquad d_0 = \min_{q \in I} d(q) > 0.$$

Every convex combination of $m + \epsilon d$ and $m - \epsilon d$ has its sign threshold in I. By Lemma D.1, for every ν there is a $q \in I$ satisfying

$$r_m(\nu, q) \geq \epsilon |r_d(\nu, q)|. \tag{D.2}$$

Let $\mathcal{A}$ be the set of all pairs $(\nu, q) \in \mathcal{P}([0,1]) \times I$ satisfying this inequality. It is compact in the product of the weak topology and the usual topology.

We first determine what happens when $r_m(\nu, q) = 0$ on $\mathcal{A}$. Then $r_d(\nu, q) = 0$. Since the roots fill I, some θ satisfies $r_\theta = q$. Hence

$$0 = r_{f_\theta}(\nu, q) = \mathbb{E}_{P\sim\nu}\big[(q - P)f_\theta(P)\big].$$

The quantity $(q - P)f_\theta(P)$ is nonnegative and vanishes only at $P = q$, so $\nu = \delta_q$. In particular, $\mathbb{E}_{P\sim\nu}[d(P)] = d(q) \geq d_0$.

Put $\gamma = d_0/2$. The compact subset

$$\mathcal{A}_0 = \{(\nu, q) \in \mathcal{A} : |\mathbb{E}_{P\sim\nu}[d(P)]| \leq \gamma\}$$

contains no point with $r_m = 0$. If it is nonempty, let $\zeta = \min_{\mathcal{A}_0} r_m > 0$ and set $L = \epsilon^{-1} + a\gamma/\zeta$. If it is empty, set $L = \epsilon^{-1}$. Take the final two tests to be $g_\pm = Lm \pm d$. At a feasible response from (D.2),

$$\min_i r_{g_i} = Lr_m - |r_d| \geq (L - \epsilon^{-1})r_m \geq 0.$$

If $|\mathbb{E}_{P\sim\nu}[d(P)]| \geq \gamma$, this is the compensation condition. Otherwise the pair lies in $\mathcal{A}_0$, which is then nonempty, and the last expression is at least $a\gamma$, again sufficient for (B.5). The tests are fixed finite-norm RKHS functions. Corollary B.1 proves the lower bound, while (9), proved in Proposition G.1 at zero delay, supplies the upper bound. □

The proof uses strict signs away from the sign threshold to prevent additional zeros under perturbation. It uses C^1 regularity near the sign threshold to obtain an interval of nearby roots. A sign threshold with extra zeros does not suffice: Theorem 5 at $\omega = 2\pi$ is an explicit bounded-error counterexample.

D.3 A DERIVATIVE-FREE CRITERION AND RADIAL KERNELS

For radial kernels, a positive reweighting replaces the local differentiability argument. The next criterion isolates the needed monotonicity.

Lemma D.2 (Ratio criterion)**.** *Let $m, d \in \mathcal{H}_K$ with $d(p) \geq \gamma > 0$. Write $k = m/d$. Suppose that for some $a \in (0, 1/2)$, k is nonincreasing on $[a, 1 - a]$, that*

$$k(p) \geq k(a) \quad (p < a), \qquad k(p) \leq k(1 - a) \quad (p > 1 - a),$$

and that $c = \min\{k(a), -k(1 - a)\} > 0$. Then the tests $g_\pm = m/c \pm d$ have a common nonnegative response in $[a, 1 - a]$. In particular, if $M = \max_i \|g_i\|_{\mathcal{H}_K}$, then

$$V_T(K) \geq V_T^{\mathrm{obl}}(K) \geq \frac{a\gamma}{\sqrt{2}M}\sqrt{T}.$$

Proof. Every convex combination of these tests has the form $d(k/c + s)$ for $s \in [-1, 1]$. The factor d is positive. The other factor is continuous and nonincreasing on the central interval, nonnegative at a, and nonpositive at $1 - a$. A zero in this interval is a sign threshold; the tail inequalities preserve its signs outside the interval. Apply Lemma D.1 and the separated form of Theorem 1. □

The response also has an explicit formula. Reweight the forecast distribution by the positive test d and clip its mean:

$$\widetilde{\nu}(\,\mathrm{d}p) = \frac{d(p)\nu(\,\mathrm{d}p)}{\mathbb{E}_{P\sim\nu}[d(P)]}, \qquad q = \mathrm{clip}\left(\mathbb{E}_{P\sim\widetilde{\nu}}[P], [a, 1 - a]\right).$$

For each of $h = k/c + 1$ and $h = k/c - 1$, the decomposition

$$\begin{aligned}\mathbb{E}_{P\sim\widetilde{\nu}}\big[(q - P)h(P)\big] &= h(q)\left(q - \mathbb{E}_{P\sim\widetilde{\nu}}[P]\right) \\ &\quad + \mathbb{E}_{P\sim\widetilde{\nu}}\big[(q - P)(h(P) - h(q))\big]\end{aligned}$$

has two nonnegative terms, by clipping, the central monotonicity, and the tail inequalities. Multiplying by $\mathbb{E}_{P\sim\nu}[d(P)]$ gives the two responses.

Proof of Theorem 4(i). Take the two endpoint-section combinations

$$m(p) = K(0,p) - K(1,p) = \rho(p) - \rho(1-p), \qquad d(p) = K(0,p) + K(1,p) = \rho(p) + \rho(1-p).$$

They belong to the RKHS. Positivity and continuity give $\min_p d(p) > 0$. As p increases, $A = \rho(p)$ strictly decreases and $B = \rho(1-p)$ strictly increases. The function $(A-B)/(A+B)$ is increasing in A and decreasing in B for positive A, B. Hence m/d is strictly decreasing. It is positive at any $a < 1/2$ and negative at $1-a$, so Lemma D.2 applies. The universal upper bound completes the proof. No differentiability of ρ was used. □

For reference, the test norms are explicit:

$$\|m\|^2_{\mathcal{H}_K} = 2\rho(0) - 2\rho(1), \qquad \|d\|^2_{\mathcal{H}_K} = 2\rho(0) + 2\rho(1), \qquad \langle m, d\rangle_{\mathcal{H}_K} = 0.$$

The lower-bound constant may depend on the fixed kernel parameter; the theorem is not a uniform statement over bandwidths depending on T.

D.4 Stationary kernels with sufficiently rich spectral support

Let μ be a finite nonzero symmetric positive measure on $\mathbb{R}$, write $\operatorname{supp}\mu$ for its support, and let i denote the imaginary unit. Suppose a continuous stationary kernel, restricted to $[0,1]$, has the representation

$$K(p,q) = \int_{\mathbb{R}} e^{i\xi(p-q)}\, \mu(\,\mathrm{d}\xi), \tag{D.3}$$

where symmetry makes the kernel real. The following condition also covers profiles that change sign or are not decreasing. Its approximation argument is related to kernel universality (Micchelli et al., 2006), but we give the required C^1 argument explicitly.

Theorem D.1 (Spectral sufficient conditions)**.** *Either of the following conditions implies* $V_T(K) = \Theta_K(\sqrt{T})$*:*

(i) $\operatorname{supp}\mu$ *contains* ξ_0 *with* $0 < |\xi_0| < 2\pi$*;*

(ii) $\operatorname{supp}\mu$ *has a finite accumulation point.*

In particular, a nonzero absolutely continuous component of μ *suffices.*

Proof. First fix any $\xi_0 \in \operatorname{supp}\mu$. For a bounded interval B around ξ_0 with $\mu(B) > 0$, the complex function

$$F_B(p) = \frac{1}{\mu(B)} \int_B e^{i\xi p}\, \mu(\,\mathrm{d}\xi)$$

is an RKHS function in the complexification, with norm at most $1/\sqrt{\mu(B)}$. This follows directly from the feature representation in $L^2(\mu)$, using coefficient $\mathbf{1}_B/\mu(B)$ and projecting onto the feature span if necessary. By symmetry of μ, complex conjugation preserves the complexified RKHS: reflect a coefficient across $\xi = 0$ and conjugate it. Therefore $(F_B + \overline{F_B})/2$ and $(F_B - \overline{F_B})/(2i)$ belong to its real subspace, which is the real RKHS. If $B \subset [\xi_0 - \epsilon, \xi_0 + \epsilon]$, then for $p \in [0,1]$,

$$|e^{i\xi p} - e^{i\xi_0 p}| \le \epsilon, \qquad |i\xi e^{i\xi p} - i\xi_0 e^{i\xi_0 p}| \le (1 + |\xi_0|)\epsilon.$$

Averaging proves uniform convergence of both F_B and F_B' to the corresponding exponential and derivative. All frequencies in B are bounded, so differentiation under this restricted integral needs no moment assumption on the full measure. Therefore $\cos(\xi_0 p)$ and $\sin(\xi_0 p)$ lie in the C^1 closure of $\mathcal{H}_K \cap C^1$.

Under (i), set $\omega = |\xi_0|$. The functions $m_0(p) = -\sin(\omega(p - 1/2))$ and $d_0(p) = \cos(\omega(p-1/2))$ have a strict sign change and a nonzero second test at $1/2$: m_0 is positive to the left and negative to the right, $m_0'(1/2) = -\omega < 0$, and $d_0(1/2) = 1$. Fixed sufficiently accurate C^1 approximations preserve opposite signs at the endpoints, strict signs away from the zero, a negative derivative near the zero, and a nonzero second test there. Thus actual finite-norm RKHS functions satisfy Theorem 4(ii).

Under (ii), we show that $\mathcal{H}_K \cap C^1$ is dense in $C^1([0,1])$ for the norm $\|f\|_\infty + \|f'\|_\infty$. Let a continuous real linear functional ℓ annihilate this space, extend it complex-linearly, and define

$$F(z) = \ell(p \mapsto e^{izp}), \qquad z \in \mathbb{C}.$$

The exponential power series and its p derivative converge in the C^1 norm locally uniformly in z, so F is entire. The approximation just proved gives $F(\xi) = 0$ for every support frequency. The identity theorem at a finite accumulation point implies $F \equiv 0$. Derivatives at zero show that ℓ annihilates every polynomial. Polynomials are dense in C^1: given $f \in C^1$ and $\epsilon > 0$, choose a polynomial P with $\|P - f'\|_\infty < \epsilon$ and set $Q(p) = f(0) + \int_0^p P(s)\,\mathrm{d}s$. Then $\|Q' - f'\|_\infty < \epsilon$ and $\|Q - f\|_\infty < \epsilon$. Thus $\ell = 0$, and Hahn–Banach separation proves the claimed density. Approximate $1/2 - p$ and 1 closely in C^1 to obtain the required sign-change pair.

Finally, a nonzero absolutely continuous component has infinitely many support points in some bounded interval, which has an accumulation point. The universal upper bound gives the matching rate in all cases. □

The approximations in this proof are fixed once the kernel is fixed; their accuracy does not depend on T. Their norms may be large, affecting the multiplicative constant but not the rate.

E Unified fast upper bounds and bounded-rate kernels

This appendix proves the generic MW guarantee and its kernel instantiations. The same minimax rule gives logarithmic error from cancellation covers and bounded expected error from strict correction. Positive-ray endpoint features admit a stronger pathwise guarantee, proved by a scalar invariant. We then verify the bounded example and the cosine and steadily turning classifications.

E.1 The finite vector game and its potential bound

Throughout this subsection, $\mathcal{A}, \mathcal{Y}$ are finite nonempty sets, $g : \mathcal{A} \times \mathcal{Y} \to \mathbb{R}^n$, $n \geq 1$, and $\|g(a, y)\|_2 \leq B$ for some $B > 0$. Fix $0 < \eta \leq 1/B$. For $z \in \mathbb{R}^n$, its signed lift has entries $\widetilde{z}_{i,+} = z_i$ and $\widetilde{z}_{i,-} = -z_i$. Thus the MW score in Algorithm 1 is

$$\mathrm{score}_w(\widetilde{g}(a, y)) = \sum_{i=1}^{n} \left[(w_i^+ - w_i^-) g_i(a, y) + \eta (w_i^+ + w_i^-) g_i(a, y)^2 \right].$$

Write $w^+ = (w_i^+)_{i=1}^n$ and $w^- = (w_i^-)_{i=1}^n$. The gain condition (11) concerns these same g, η and a fixed margin $\epsilon \geq 0$. A distribution over actions can witness that condition: the scalar expression is linear in the distribution, so at least one action in its support has no larger value.

Lemma E.1 (Minimax score bound)**.** *Assume* (11)*. For any positive signed weights with $w_i^+ w_i^- = 1$ for every i,*

$$\min_{\pi \in \Delta(\mathcal{A})} \max_{y \in \mathcal{Y}} \mathbb{E}_{a \sim \pi}\, \mathrm{score}_w(\widetilde{g}(a, y)) \leq 2\eta B^2 - \epsilon \|w^+ - w^-\|_1. \tag{E.1}$$

Proof. Put $x = w^+ - w^-$. Abbreviating $g_i = g_i(a, y)$, we have

$$\mathrm{score}_w(\widetilde{g}(a, y)) = \sum_i \left(x_i g_i + \eta |x_i| g_i^2 \right) + 2\eta \sum_i \min\{w_i^+, w_i^-\} g_i^2.$$

The product identity implies $\min\{w_i^+, w_i^-\} \leq 1$, so the last sum is at most B^2. For every $q \in \Delta(\mathcal{Y})$, (11) supplies an action whose expected score is at most $2\eta B^2 - \epsilon \|x\|_1$. Applying finite minimax to the action-by-outcome score matrix proves (E.1). □

Proof of Theorem 6. Write $Z_t = \sum_{s=1}^t g(a_s, y_s)$, with $Z_0 = 0$. After round t, the signed updates give $w_i^+ = e^{\eta Z_{t,i}}$ and $w_i^- = e^{-\eta Z_{t,i}}$. In particular, $w_i^+ w_i^- = 1$. Subtract the minimum total weight of the signed pairs and define

$$\Phi_t := \sum_{i=1}^{n} \left(e^{\eta Z_{t,i}} + e^{-\eta Z_{t,i}} - 2 \right) \geq 0, \qquad \Phi_0 = 0. \tag{E.2}$$

For $r > 0$, $|r - r^{-1}| \geq r + r^{-1} - 2$; hence $\|w^+ - w^-\|_1 \geq \Phi_t$.

For $|u| \le 1$, the elementary exponential estimate $e^u \le 1 + u + u^2$ holds. Apply it to each signed coordinate of the next update; $|\eta g_i| \le \eta B \le 1$. With w denoting the weights after round t, this gives the pathwise inequality

$$\Phi_{t+1} - \Phi_t \le \eta \, \mathrm{score}_w(\widetilde{g}(a_{t+1}, y_{t+1})).$$

At every fixed state, the mixture chosen by Algorithm 1 bounds the expected score for each outcome separately. Nature may therefore choose an outcome after seeing that mixture, provided it cannot react to the fresh action draw. Let $\mathcal{F}_t$ contain the history and learner state through round t. Conditioning on this information and using Lemma E.1 yields

$$\mathbb{E}[\Phi_{t+1} \mid \mathcal{F}_t] \le (1 - \eta\epsilon)\Phi_t + 2\eta^2 B^2. \tag{E.3}$$

The statewise argument also permits Nature to see the learner's state. Randomized responses are handled by averaging over the current outcome.

Condition (11), applied to a coordinate vector x, implies $\epsilon \le B$: its left side is always at least $-B$. Consequently $0 \le \eta\epsilon \le 1$. Iterating (E.3) gives

$$\mathbb{E}\Phi_T \le 2\eta^2 B^2 H_T, \qquad H_T := \sum_{s=0}^{T-1} (1 - \eta\epsilon)^s. \tag{E.4}$$

The sum is T when $\epsilon = 0$, and is at most $1/(\eta\epsilon)$ when $\epsilon > 0$.

For every state,

$$e^{\eta \|Z_T\|_\infty} \le 2 + \Phi_T, \qquad e^{\eta \|Z_T\|_1 / n} \le \frac{1}{n} \sum_i e^{\eta |Z_{T,i}|} \le 2 + \frac{\Phi_T}{n}.$$

Taking expectations and using Jensen's inequality proves (12). More explicitly, the two norm bounds are

$$\mathbb{E}\|Z_T\|_\infty \le \frac{1}{\eta} \log(2 + \mathbb{E}\Phi_T), \qquad \mathbb{E}\|Z_T\|_1 \le \frac{n}{\eta} \log\left(2 + \frac{\mathbb{E}\Phi_T}{n}\right). \tag{E.5}$$

These prove both claimed regimes. □

The hypothesis compensates for the quadratic cost of the MW update. A zero margin bounds the additive growth of the centered potential; a positive margin contracts that potential up to a constant. The signed MW construction and minimax step use standard potential-based approachability ideas (Cesa-Bianchi & Lugosi, 2003; Perchet, 2015); it is the gain condition and its verification below that provide these faster rates.

E.2 Implementation, private labels, and conversion to kernel error

For fixed weights, minimizing the maximum expected score is a linear program: introduce a scalar b, minimize b, and impose

$$\sum_a \pi(a) \, \mathrm{score}_w(\widetilde{g}(a, y)) \le b \quad (y \in \mathcal{Y}), \qquad \pi(a) \ge 0, \quad \sum_a \pi(a) = 1.$$

An optimizer exists because the objective is continuous on the finite probability simplex. Only countably many states are reachable in finitely many rounds: each finite sequence of actions and outcomes determines one state. Fixing an optimizer at each reachable state therefore defines a measurable learner. The next paragraph gives a finite implementation with fixed tie-breaking for binary outcomes.

Lemma E.2 (Two-action implementation for binary outcomes)**.** *When $\mathcal{Y} = \{0, 1\}$, an optimizing mixture can be chosen with support of size at most two.*

Proof. For each action let

$$c_a = \big(\mathrm{score}_w(\widetilde{g}(a, 0)), \mathrm{score}_w(\widetilde{g}(a, 1))\big) \in \mathbb{R}^2.$$

The objective minimizes $\max\{c_0, c_1\}$ over the convex hull of these finitely many points. If the hull has dimension two, an interior point cannot minimize the objective: moving it slightly downward in

both coordinates improves the objective while staying in the hull. A boundary point lies on an edge or is a vertex, so it is a mixture of at most two action points. A hull of dimension at most one also has this property.

An explicit finite search evaluates every pure action and the following candidate for every pair a, b. Set $d_a = c_{a,0} - c_{a,1}$, and define d_b similarly. If $d_a \neq d_b$ and $\theta = d_b/(d_b - d_a) \in [0,1]$, evaluate $\theta c_a + (1-\theta)c_b$. Along a segment the objective is the maximum of two affine functions; its minimum is attained at an endpoint or at their intersection. The candidates therefore include an optimizer. Select the first minimizing candidate in a fixed ordering. The scores are continuous functions of the state and the candidate formulas are measurable on their domains, so this rule is measurable. □

For a kernel application, associate a forecast $p(a)$ with each action and exhibit fixed vectors $v_1, \dots, v_n$ of norm at most one such that

$$(y - p(a))\phi(p(a)) = \sum_{i=1}^{n} g_i(a, y) v_i. \tag{E.6}$$

Summing the identity and using the triangle inequality gives

$$\mathbf{S}_K(T) = \sum_i Z_{T,i} v_i, \qquad \|\mathbf{S}_K(T)\| \le \|Z_T\|_1. \tag{E.7}$$

The v_i need not be independent. The coordinates of Z_T are bookkeeping variables, not necessarily test scores or orthogonal projections of $\mathbf{S}_K(T)$.

When distinct labels have the same forecast, announce the forecast marginal $\nu_t = \sum_a \pi_t(a)\delta_{p(a)}$, privately draw the labelled action, and retain its label for the gain update. The generic guarantee holds even if Nature sees the entire labelled mixture. It therefore holds under Definition 1, where Nature sees only the announced forecast law and the past. The fresh labelled draw must remain hidden until Nature commits to the outcome; the one-step score bound then holds conditionally at every private state.

Proposition E.1 (Kernel exponential moments and fixed-time tails)**.** *Suppose* (11) *and* (E.6) *hold. With H_T as in* (E.4)*, for every fixed integer $T \ge 1$ and $\zeta \in (0,1)$,*

$$\mathbb{E} \exp\left(\frac{\eta}{n}\|\mathbf{S}_K(T)\|\right) \le 2 + \frac{2\eta^2 B^2}{n} H_T, \tag{E.8}$$

$$\mathbb{P}\left\{\|\mathbf{S}_K(T)\| > \frac{n}{\eta}\log\frac{2 + (2\eta^2 B^2/n)H_T}{\zeta}\right\} \le \zeta. \tag{E.9}$$

Proof. The pathwise comparison (E.7) and the exponential estimate in the proof of Theorem 6 give $\exp(\eta\|\mathbf{S}_K(T)\|/n) \le 2 + \Phi_T/n$. Take expectations and apply (E.4); Markov's inequality gives the tail bound. □

E.3 Cancellation covers and logarithmic error

Fix a finite cancellation cover as in Definition 11, with N pairs, unit vectors v_i, and positive amplitudes α_i, β_i. We give the gain representation, the learning rate, and the full verification of (11).

Lemma E.3 (A uniform margin for the cover)**.** *There is $\delta \in (0,1]$ such that for every $q \in [0,1]$ some pair i satisfies*

$$\ell_i = 0 \text{ or } q - \ell_i \ge \delta, \qquad u_i = 1 \text{ or } u_i - q \ge \delta. \tag{E.10}$$

Proof. Define continuous functions

$$A_i(q) = \begin{cases} q - \ell_i, & \ell_i > 0, \\ 1, & \ell_i = 0, \end{cases} \qquad B_i(q) = \begin{cases} u_i - q, & u_i < 1, \\ 1, & u_i = 1. \end{cases}$$

The relatively open cover makes $\max_i \min\{1, A_i(q), B_i(q)\}$ positive at every q. Compactness gives

$$\delta_* := \min_{q \in [0,1]} \max_i \min\{1, A_i(q), B_i(q)\} > 0. \tag{E.11}$$

Any $\delta \in (0, \delta_*]$ has the claimed property. □

Use labelled actions $\mathcal{A} = \{(i, \ell_i), (i, u_i) : 1 \le i \le N\}$ and write $r_{i,\ell_i} = \alpha_i$, $r_{i,u_i} = \beta_i$. Set

$$g((i,p),y) = r_{i,p}(y-p)\mathbf{e}_i, \qquad Z_t = Z_{t-1} + g((i_t,p_t),y_t), \tag{E.12}$$

where $\mathbf{e}_i$ is the ith standard basis vector of $\mathbb{R}^N$. Since $\phi(p) = r_{i,p}v_i$ on each action,

$$\mathbf{S}_K(T) = \sum_{i=1}^{N} Z_{T,i}v_i, \qquad \|\mathbf{S}_K(T)\| \le \sum_{i=1}^{N} |Z_{T,i}|. \tag{E.13}$$

Only one gain coordinate changes per round. Thus $\|g(a,y)\|_2 \le B$ for $B = \max_i\{\alpha_i, \beta_i\} > 0$. Choose

$$\eta = \frac{\delta}{2B}. \tag{E.14}$$

Proof of Corollary 2. We verify (11) with $n = N$ and $\epsilon = 0$. Fix $q \in [0,1]$ and $x \in \mathbb{R}^N$, and let $Y \sim \mathrm{Bernoulli}(q)$. Choose a pair i satisfying (E.10). If $x_i = 0$, any action from that pair makes the expression in (11) zero. Otherwise use the upper endpoint when $x_i > 0$ and the lower endpoint when $x_i < 0$. Let c denote the selected amplitude and put $h = \mathrm{sgn}(x_i)g_i(a,Y)$. At an interior endpoint, $\mathbb{E}h \le -c\delta$ and $|h| \le c$, so

$$\mathbb{E}(h + \eta h^2) \le c(-\delta + \eta c) \le -c\delta/2 \le 0.$$

At a boundary endpoint, $h \in [-B, 0]$ for either outcome, whence $h + \eta h^2 = h(1+\eta h) \le 0$ because $\eta B \le 1/2$. Multiplying by $|x_i|$ verifies (11); every other gain coordinate is zero. Apply Theorem 6 and (E.13) to obtain

$$\mathbb{E}\|\mathbf{S}_K(T)\| \le \frac{N}{\eta}\log\left(2 + \frac{2\eta^2 B^2 T}{N}\right) = O_K(\log(T+1)).$$

All cover data and the learning rate are fixed independently of T. □

The cancellation interpretation does not require knowing Nature's mean. For a hypothetical $q \in [\ell_i, u_i]$, assigning probability

$$\frac{\beta_i(u_i - q)}{\alpha_i(q - \ell_i) + \beta_i(u_i - q)}$$

to (i, ℓ_i) and the remaining mass to (i, u_i) makes the expected auxiliary gain zero. The denominator is positive since $\ell_i < u_i$ and the amplitudes are positive. This explains the geometric condition; the actual learner uses the minimax score, and its analysis additionally controls the quadratic term. Cancellation in expectation alone is not the rate argument.

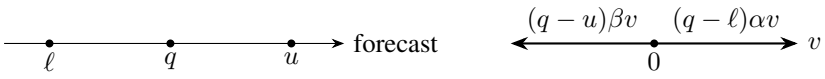


Figure 2: Positive-ray cancellation: the expected increments point oppositely along the same full-feature direction. Their mixture can cancel for a hypothetical outcome mean q; the online learner instead uses the mean-independent minimax score.

Proposition E.2 (Moments and tails for a cancellation cover)**.** *Under the preceding pair-cover construction, for every fixed $T \ge 1$ and $\zeta \in (0,1)$,*

$$\mathbb{E}\exp\left(\frac{\eta}{N}\,\mathrm{kCE}_{K,T}\right) \le 2 + \frac{2\eta^2 B^2 T}{N}, \tag{E.15}$$

$$\mathbb{P}\left\{\mathrm{kCE}_{K,T} > \frac{N}{\eta}\log\frac{2 + 2\eta^2 B^2 T/N}{\zeta}\right\} \le \zeta. \tag{E.16}$$

Proof. The gain-condition verification above has $n = N$ and $\epsilon = 0$, so $H_T = T$. Apply Proposition E.1. □

No finite-rank assumption is needed. Execution with supplied pairs needs only kernel evaluations:

$$\alpha_i = \sqrt{K(\ell_i, \ell_i)}, \qquad \beta_i = \sqrt{K(u_i, u_i)}.$$

The positive-ray condition is equivalent to $K(\ell_i, u_i) = \sqrt{K(\ell_i, \ell_i)K(u_i, u_i)} > 0$ by equality in Cauchy–Schwarz. Finding a cover and a valid margin remains a separate geometric task. A cover is only an upper-bound certificate: the single pair $(0, 1)$ also permits bounded error, as shown below.

E.4 Strict correction and bounded expected error

We first prove the finite-set criterion stated in the main text, then connect it to the absence of sign thresholds in finite rank.

Strict-correction part of Corollary 3. Let P be the supplied finite forecast set, and let its feature span have dimension $n \geq 1$. Choose an orthonormal basis of that span and identify its feature vectors with vectors in $\mathbb{R}^n$. Use actions $p \in P$ and gains $g(p, y) = (y - p)\phi(p)$. Put $B = \max_{p \in P} \|\phi(p)\|_2$. The assumed radius-$\delta$ ball implies $B \geq \delta > 0$. Take

$$\eta = \frac{\delta}{2B^2}, \qquad \epsilon = \frac{\delta}{2\sqrt{n}}. \tag{E.17}$$

Fix $q \in [0, 1]$ and $x \neq 0$. The convex-hull condition supplies a distribution on P whose expected gain against $Y \sim \text{Bernoulli}(q)$ equals $-\delta x/\|x\|_2$. For every gain vector,

$$\sum_i |x_i| g_i^2 \leq \|x\|_\infty \|g\|_2^2 \leq B^2 \|x\|_2.$$

The witnessing distribution consequently satisfies

$$\mathbb{E} \sum_i (x_i g_i + \eta |x_i| g_i^2) \leq (-\delta + \eta B^2)\|x\|_2 = -\frac{\delta}{2}\|x\|_2 \leq -\frac{\delta}{2\sqrt{n}}\|x\|_1.$$

One pure action does at least as well. For $x = 0$ the condition is immediate, and $\eta B \leq 1/2$. Thus (11) holds with the positive margin in (E.17). Apply Theorem 6 and the kernel conversion, using the chosen orthonormal basis as the v_i, to obtain $\sup_T \mathbb{E}\|\mathbf{S}_K(T)\| < \infty$. □

For the global characterization, let $\mathcal{H}$ be the closed feature span and put $\kappa = \max_p \|\phi(p)\|$. A nonzero vector in $\mathcal{H}$ represents a nonzero RKHS test. Recall from Definition D.1 that a test f has a sign threshold at r when $(r - p)f(p) \geq 0$ for all p; additional zeros are allowed. For a unit vector $u \in \mathcal{H}$, set $f_u(p) = \langle u, \phi(p) \rangle$ and

$$v(u) = \min_{\nu \in \mathcal{P}([0,1])} \max_{q \in [0,1]} \mathbb{E}_{P \sim \nu}[(q - P) f_u(P)].$$

Continuity and compactness of the forecast interval permit minimax, giving

$$v(u) = \max_{q \in [0,1]} \min_{p \in [0,1]} (q - p) f_u(p) \leq 0. \tag{E.18}$$

Indeed, the payoff is continuous and affine in each of the distribution ν and the scalar q, on compact convex domains. The inner minimum is a continuous function of q, so the maximum is attained. It equals zero precisely when f_u has a sign threshold.

Lemma E.4 (Finite-rank uniform correction)**.** *Suppose $1 \leq \dim \mathcal{H} < \infty$. The absence of a sign threshold for every nonzero RKHS test is equivalent to the existence of $\delta > 0$ such that every unit direction u admits a forecast distribution ν_u with*

$$\max_{y \in \{0,1\}} \mathbb{E}_{P \sim \nu_u}[(y - P)\langle u, \phi(P) \rangle] \leq -\delta. \tag{E.19}$$

Equivalently, for every $q \in [0, 1]$, the convex set $C(q) = \text{conv}\{(q - p)\phi(p) : p \in [0, 1]\}$ contains the closed radius-δ ball about zero, with the same δ.

Proof. By (E.18), the absence of sign thresholds means $v(u) < 0$ for every unit u. Each scalar payoff is κ-Lipschitz in u, and minima and maxima preserve this bound. Hence v is continuous on the finite-dimensional unit sphere. Compactness gives $\delta = -\max_{\|u\|=1} v(u) > 0$, proving (E.19). Conversely, that inequality makes every $v(u)$ negative, ruling out sign thresholds.

The scalar dual identity makes (E.19) equivalent to $\min_{z \in C(q)} \langle u, z \rangle \le -\delta$ for every q and unit u. Replacing u by $-u$ gives $\max_{z \in C(q)} \langle u, z \rangle \ge \delta$. Each $C(q)$ is compact, since it is the convex hull of a compact set in finite dimension. If it contains the radius-δ ball, these support inequalities follow by evaluating at δu. Conversely, separation of a point of that ball from $C(q)$ would give a unit u with $\max_{z \in C(q)} \langle u, z \rangle < \delta$, a contradiction. This proves the equivalence with the same radius. □

Lemma E.5 (A finite grid preserves strict correction)**.** *Let ϕ be continuous with finite-dimensional feature span $\mathcal{H}$. If $C(q)$ contains the radius-δ ball in $\mathcal{H}$ for every $q \in [0,1]$, then there is a finite $P \subset [0,1]$ such that*

$$\{z \in \mathcal{H} : \|z\| \le \delta/2\} \subseteq \operatorname{conv}\{(q-p)\phi(p) : p \in P\} \quad \textit{for every } q \in [0,1].$$

In particular, the features on P span $\mathcal{H}$.

Proof. Uniform continuity of the two maps $p \mapsto -p\phi(p)$ and $p \mapsto (1-p)\phi(p)$ gives a finite grid P such that every $p \in [0,1]$ has a grid point $p' \in P$ at which both maps differ by at most $\delta/2$. The same p' works for every convex combination of the two maps. Thus, simultaneously for every $q \in [0,1]$,

$$\|(q-p)\phi(p) - (q-p')\phi(p')\| \le \delta/2.$$

For every unit $u \in \mathcal{H}$, the support function of the finite convex hull is at least the support function of $C(q)$ minus $\delta/2$, hence at least $\delta/2$. The compact-convex separation argument from Lemma E.4 now gives the ball inclusion. A set containing a positive-radius ball spans $\mathcal{H}$. □

Combining Lemmas E.4 and E.5 and Corollary 3 proves bounded expected error whenever a continuous finite-rank kernel has no nonzero test with a sign threshold. The finite-set certificate itself does not require the full kernel to have finite rank: forecasts confined to P keep all errors in its finite feature span.

The finite-rank assumption in the global margin equivalence cannot simply be omitted. If the feature span is infinite-dimensional, compactness of the feature curve gives, for each $\tau > 0$, a finite τ-net of feature vectors. A unit vector orthogonal to their span has $\sup_p |\langle u, \phi(p) \rangle| \le \tau$, precluding a positive uniform margin in (E.19). This does not preclude bounded infinite-rank kernels; endpoint balancing or a finite-set certificate can still apply.

E.5 Endpoint balancing and zero features

Lemma E.6 (The MW endpoint invariant)**.** *Suppose $\phi(0) = \alpha v$ and $\phi(1) = \beta v$, with $\|v\| = 1$ and $\alpha, \beta > 0$. Restrict Algorithm 1 to the two endpoint actions and scalar gains*

$g(p,y)$	$y=0$	$y=1$
$p=0$	0	α
$p=1$	$-\beta$	0

and set $B = \max\{\alpha, \beta\}$ and $\eta = 1/(2B)$. Then $\|\mathbf{S}_K(t)\| \le 3B$ on every path for every $t \ge 0$. If instead $\phi(p_*) = 0$, the singleton action p_* gives zero error on every path.

Proof. Write the scalar cumulative gain as z, so $\mathbf{S}_K = zv$. If $z > 2B$, then $w^+/w^- = e^{2\eta z} > e^2 > 3$. For any $0 < c \le B$,

$$\operatorname{score}_w(c, -c) = c[(w^+ - w^-) + \eta c(w^+ + w^-)] > 0,$$

whereas

$$\operatorname{score}_w(-c, c) = c[-(w^+ - w^-) + \eta c(w^+ + w^-)] < 0.$$

The second inequality follows from $\eta c \le 1/2$ and $w^+ > 3w^-$. The pure forecast 1 has worst-case score zero: its score is negative at outcome 0 and zero at outcome 1. Any positive mass on forecast 0 gives positive expected score at outcome 1. Thus the unique minimizer is forecast 1, whose increment is nonpositive. Symmetrically, $z < -2B$ forces forecast 0 and a nonnegative increment. Every jump has magnitude at most B, so starting at zero the interval $[-3B, 3B]$ is invariant. The zero-feature assertion follows because every increment then vanishes. □

This proves the endpoint and zero-feature parts of Corollary 3. The endpoint bound does not follow by taking a positive margin in (11). For the scalar endpoint game, take the outcome law concentrated at 0 and $x < 0$. Both possible values of $xg + \eta|x|g^2$ are nonnegative, so no $\epsilon > 0$ can satisfy that condition. The scalar invariant supplies the stronger conclusion for the same algorithm.

Lemma E.7 (Elementary boundary cases)**.** *If $\phi(p_*) = 0$ at some forecast, then $V_T(K) = 0$ for every T. Otherwise*

$$V_T(K) \geq V_T^{\mathrm{obl}}(K) \geq \frac{1}{2} \min_p \|\phi(p)\| > 0 \qquad (T \geq 1). \tag{E.20}$$

If $\phi(0) = \alpha v$, $\phi(1) = \beta v$, $\|v\| = 1$, and $\alpha, \beta > 0$, a deterministic endpoint strategy has error at most $\max\{\alpha, \beta\}$. In particular, continuous rank-one kernels have either zero error for all horizons or rate $\Theta_K(1)$.

Proof. A zero feature permits the constant prediction p_*. Otherwise, take a fixed outcome prefix and let the final outcome be an independent fair bit. Conditional on the learner's randomness, the two possible final error vectors differ by $\phi(p_T)$. Their average norm is at least $\|\phi(p_T)\|/2$ by the triangle inequality. Continuity and compactness give the strictly positive minimum in (E.20). The outcome law is oblivious, so the bound holds for the oblivious value as well.

For the deterministic strategy write $\mathbf{S}_K = Zv$, predict 0 when $Z \leq 0$, and predict 1 otherwise. A forecast 0 adds either zero or α, while a forecast 1 adds either $-\beta$ or zero. The interval $[-\beta, \alpha]$ is invariant. Finally, a continuous rank-one feature curve can be written as $f(p)v$ for a fixed unit vector and continuous scalar f. Either f has a zero, or it has constant sign and the endpoint features are positive multiples of one another. □

E.6 A BROADER POSITIVE-COMBINATION COVER

Cancellation pairs can be replaced by positive linear identities between separated groups of forecasts.

Definition E.1 (Positive-combination block)**.** A block consists of finite nonempty sets $L, R \subset [0, 1]$ and positive coefficients a_p for $p \in L$, b_p for $p \in R$, such that

$$\ell := \max L < \min R =: u, \qquad \sum_{p \in L} a_p \phi(p) = \sum_{p \in R} b_p \phi(p) \neq 0. \tag{E.21}$$

Its interval is (ℓ, u), including 0 if $\ell = 0$ and 1 if $u = 1$. A finite positive-combination cover is a collection of such blocks whose intervals cover $[0, 1]$.

Proposition E.3 (Positive combinations suffice)**.** *A finite positive-combination cover gives a finite gain representation and a fixed learning rate satisfying (11) with $\epsilon = 0$. Consequently Algorithm 1 achieves $\mathbb{E}\|\mathbf{S}_K(T)\| = O_K(\log(T + 1))$.*

Proof. For each block j, choose an orthonormal basis of the span $\mathcal{H}_j$ of its feature vectors. Concatenate these coordinate blocks, and label every action by its block and forecast. Its gain is the coordinate vector of $(y - p)\phi(p)$ in that block, with zeros elsewhere. The concatenated basis vectors give (E.6), although vectors belonging to different blocks need not be orthogonal. Take B to be the largest norm of a selected feature vector. The nonzero identity guarantees $B > 0$ and that each $\mathcal{H}_j$ has positive dimension.

Set $s_p = 1$ on L and $s_p = -1$ on R. In each block, (E.21) is a strictly positive linear combination of the spanning vectors $s_p\phi(p)$ equal to zero. Their convex hull contains a ball about zero in $\mathcal{H}_j$: otherwise a separating nonzero linear functional would be nonnegative on all these vectors, and their strictly positive zero combination would force it to vanish on every one, contradicting their spanning property. Choose a radius $\gamma > 0$ valid for all blocks. The interval-cover argument of Lemma E.3 supplies a common forecast margin $\delta \in (0, 1]$. Fix

$$0 < \eta \leq \min\{1/B, \gamma\delta/B^2\}.$$

Fix an outcome mean $q \in (0, 1)$ and a direction in the concatenated coordinate space. Select a block with the cover margin, and let x be the restriction of the direction to that block. If $x = 0$, any action

in this block gives zero in (11). Otherwise the feature ball supplies probabilities θ_p on its forecasts with

$$\sum_p \theta_p s_p \phi(p) = -\gamma x/\|x\|_2.$$

Since q lies strictly between the block's two forecast sets, $\operatorname{sgn}(q-p) = s_p$ and $|q-p| > 0$ for every selected forecast. Use action probabilities proportional to $\theta_p/|q-p|$. Before normalizing these probabilities, the linear part of the gain expression has expectation $-\gamma\|x\|_2$.

For nonboundary forecasts the cover margin gives $|q-p| \geq \delta$, so $\mathbb{E}_q(Y-p)^2/|q-p| \leq 1/\delta$. At $p = 0$ or 1 the same ratio is exactly one. These are the only forecasts needing a boundary exception: $\max L = 0$ forces $L = \{0\}$, and $\min R = 1$ forces $R = \{1\}$. Moreover, in the chosen orthonormal coordinates,

$$\sum_i |x_i|\phi_i(p)^2 \leq B^2\|x\|_2.$$

Consequently the unnormalized expected gain expression is at most

$$\sum_p \frac{\theta_p}{|q-p|}\mathbb{E}_q \sum_i \big(x_i g_i(p,Y) + \eta|x_i|g_i(p,Y)^2\big) \leq \left(-\gamma + \frac{\eta B^2}{\delta}\right)\|x\|_2 \leq 0.$$

Normalization preserves the inequality, and the inactive coordinates contribute zero. For $q = 0$ or $q = 1$, the cover supplies the matching endpoint action, whose gain is identically zero under that outcome law. Thus a mixture, and hence a pure action, witnesses (11) with $\epsilon = 0$ in every case. Apply Theorem 6 and (E.7). □

Singleton sets L, R recover cancellation pairs. Larger blocks may have smaller feature margins or narrower forecast gaps; this extension gives a sufficient condition and a fixed learning rate, not a uniform improvement of the constants of a given pair cover.

E.7 The bounded example and the cosine transition

Proposition E.4 (An analytic, pointwise positive bounded-rate kernel)**.** *The kernel K_{b} in (10) is analytic, strictly positive at every pair of arguments, and has rank three. Its feature map $\phi_{\mathrm{b}}(p) = (1+p)(\sqrt{2}, \cos(2\pi p), \sin(2\pi p))$ is injective with nonzero derivative everywhere, and $K_{\mathrm{b}}(p,p) = 3(1+p)^2$. For every $T \geq 1$,*

$$\frac{\sqrt{3}}{2} \leq V_T(K_{\mathrm{b}}) \leq 2\sqrt{3},$$

with a deterministic pathwise upper bound.

Proof. The displayed feature inner product gives K_{b}, proving positive semidefiniteness. Its trigonometric factor $2 + \cos(2\pi(p-q))$ is at least one, proving strict positivity. Analyticity is immediate. The three functions $1, \cos(2\pi p), \sin(2\pi p)$ are linearly independent, and multiplication by $1 + p > 0$ preserves this independence, so the rank is three. The first coordinate $\sqrt{2}(1+p)$ proves injectivity and ensures a nonzero derivative. The diagonal follows from $\cos^2 + \sin^2 = 1$.

For $v = (\sqrt{2}, 1, 0)/\sqrt{3}$, the endpoint features are $\sqrt{3}v$ and $2\sqrt{3}v$. The deterministic strategy of Lemma E.7 gives the upper bound. The minimum feature norm is $\sqrt{3}$, so (E.20) gives the lower bound. □

Thus injectivity does not preclude positive-ray cancellation. This example has no zero feature, no identical features at distinct forecasts, and no negative kernel values.

Proof of Theorem 5. Use $\phi_\omega(p) = (\cos(\omega p), \sin(\omega p))$. For $0 < \omega < 2\pi$, take

$$m(p) = \sin(\omega(1/2 - p)), \qquad d(p) = \cos(\omega(p - 1/2)).$$

These are RKHS tests; m is strictly positive below $1/2$ and strictly negative above it, $m'(1/2) = -\omega < 0$, and $d(1/2) = 1$. Thus Theorem 4(ii) gives the square-root rate. At $\omega = 0$ or $\omega = 2\pi$, the endpoints coincide, so Lemma E.7 gives bounded error.

For $\omega > 2\pi$, every nonzero RKHS test is a nonzero sinusoid traced over an angular interval longer than 2π. It has both a positive-to-negative crossing and a negative-to-positive crossing in the interior, which rules out a sign threshold. By Lemmas E.4 and E.5 and Corollary 3, the kernel has bounded expected error. Its unit diagonal supplies the constant lower bound in every bounded case by Lemma E.7. This proves the theorem. □

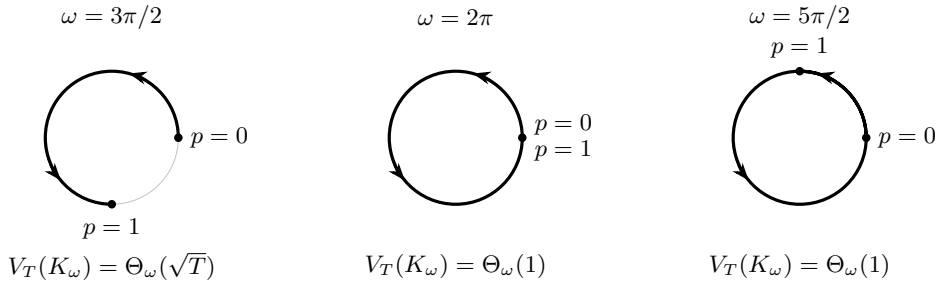


Figure 3: The cosine feature map. Below one full turn, two tests force square-root error. At one turn the endpoints coincide and permit balancing. Beyond one turn every nonzero direction admits strict correction. Arrows indicate increasing forecast labels, not a required order of play.

Corollary E.1 (The threshold for steadily turning planar curves)**.** *Let $a, \theta \in C^1([0,1])$ satisfy $a > 0$ and $\theta' > 0$, and let $\phi(p) = a(p)(\cos\theta(p), \sin\theta(p))$. Write $L = \theta(1) - \theta(0)$. The induced kernel has rate $\Theta_K(\sqrt{T})$ for $0 < L < 2\pi$, and $\Theta_K(1)$ for $L \geq 2\pi$.*

Proof. For $L < 2\pi$, let r be the unique point at the angular midpoint, and put

$$m(p) = a(p)\sin(\theta(r) - \theta(p)), \qquad d(p) = a(p)\cos(\theta(p) - \theta(r)).$$

Both are linear measurements of ϕ, hence RKHS tests. The angular deviations have magnitude below π, so m has strict signs on the two sides of r, $m'(r) = -a(r)\theta'(r) < 0$, and $d(r) = a(r) > 0$. Apply Theorem 4(ii). At $L = 2\pi$ the endpoint vectors are positive multiples, so Lemma E.7 applies. For $L > 2\pi$, positivity of a and strict increase of θ preserve both orientations of the sinusoidal crossings. Thus every nonzero test lacks a sign threshold, and Lemmas E.4 and E.5 and Corollary 3 give bounded expected error. Finally, $\min_p a(p) > 0$ by continuity and compactness, so Lemma E.7 gives the constant lower bound. □

F Logarithmic lower bounds and rank-two families

This appendix proves the logarithmic lower bound independently of the upper-bound algorithm, verifies the smooth four-forecast example, and gives the full rank-two classification underlying Corollary 1. Throughout, errors are cumulative. The upper bounds use Algorithm 1 and Corollary 2; the lower bounds hold against randomized learners and admit one oblivious hard distribution at each fixed horizon.

F.1 Nonnegative tests and the logarithmic lower bound

The argument has two steps. Small expected terminal error would confine two mean test scores at every prescribed outcome history. A linear–exponential potential then shows that a full outcome tree cannot stay in that square for more than exponentially many rounds in its radius.

The role of the two tests. Let $f, g \in \mathcal{H}_K$ be nonnegative and suppose that, for $\epsilon \in (0,1)$ and $\eta > 0$,

$$f(p) = 0 \quad \text{for } p > 1 - \epsilon, \qquad (1-p)f(p) + pg(p) \geq \eta \quad (0 \leq p \leq 1).$$

An outcome one increases S_f by $(1-p)f(p)$, whereas an outcome zero increases $-S_g$ by $pg(p)$. The second inequality gives a positive lower bound on the sum of these two possible gains. The first

condition limits how much an outcome zero can undo the possible gain in S_f: for nonnegative f, it is equivalent to

$$pf(p) \le \frac{1-\epsilon}{\epsilon}(1-p)f(p) \qquad (0 \le p \le 1). \tag{F.1}$$

Indeed, multiplying by ϵ and rearranging gives $(p-(1-\epsilon))f(p) \le 0$. This explains why exact vanishing is needed by this argument. Making f merely small near one does not control the ratio $p/(1-p)$ wherever $f(p) > 0$. No corresponding ratio condition is imposed on g.

We first record the terminal-to-prefix implication with general test norms; this form is also used in Proposition H.1.

Lemma F.1 (Terminal error bounds all mean prefixes)**.** *Fix a randomized learner A for horizon T, and let*

$$B = \max_{y \in \{0,1\}^T} \mathbb{E}_A[\mathrm{kCE}_{K,T}].$$

For a prescribed binary outcome history h of length at most T, let $S_f(h)$ be its accumulated test score. If $f \in \mathcal{H}_K$ is nonnegative, then

$$|\mathbb{E}_A[S_f(h)]| \le \|f\|_{\mathcal{H}_K} B.$$

Here the expectation uses the learner's original private randomness with the outcomes in h forced; it does not condition that randomness on realized forecasts.

Proof. For every fixed terminal sequence,

$$|\mathbb{E}_A[S_f(T)]| \le \mathbb{E}_A[|S_f(T)|] \le \|f\|_{\mathcal{H}_K} \mathbb{E}_A[\mathrm{kCE}_{K,T}] \le \|f\|_{\mathcal{H}_K} B.$$

If a prefix mean exceeded $\|f\|_{\mathcal{H}_K} B$, appending only ones would preserve that inequality, since every added increment $(1-p)f(p)$ is nonnegative for every realization of the learner's randomness. If it were smaller than $-\|f\|_{\mathcal{H}_K} B$, appending only zeros would preserve that inequality, since every added increment $-pf(p)$ is nonpositive. Either continuation contradicts the terminal bound. □

The next statement isolates the deterministic potential argument. Its two coordinates will be scaled mean scores, not realized scores of a random play.

Lemma F.2 (Duration of bounded mean-score control)**.** *Fix $\epsilon \in (0,1)$ and $R \ge 0$. Assign a state (X,Y) to every node of a binary outcome tree of depth T, starting at $(0,0)$. At each nonterminal node, suppose some $a \ge 0$ has the following properties: appending one increases X by a, whereas appending zero decreases X by at most $((1-\epsilon)/\epsilon)a$ and increases Y by at least $1-a$. If every state lies in $[-R,R]^2$, then*

$$T \le (1+4R)\left(\frac{2}{\epsilon}\right)^{4R}. \tag{F.2}$$

No restriction on the change in Y after appending one is required.

Proof. We seek a potential $F(X,Y) = X + w(Y)$ with w positive and increasing. Its two terms play different roles: only a decrease in X is controlled relative to a possible increase, whereas an outcome one may decrease Y without such a comparison.

Require $w(Y) \le 1/4$ for $Y \le R$, and abbreviate its current value by w. If $a \ge 2w$, choose outcome one. Even losing the entire w term leaves a potential increase of at least $a - w \ge w$. If $a < 2w$, choose outcome zero. Then $a < 1/2$, so Y increases by at least $1/2$, while X decreases by at most $2((1-\epsilon)/\epsilon)w$. To pay for this loss and still gain w, it suffices that an increase of $1/2$ in Y multiply $w(Y)$ by $2/\epsilon$.

This motivates the exponential

$$w(Y) = \frac{1}{4}\left(\frac{2}{\epsilon}\right)^{2(Y-R)}.$$

It satisfies $w(Y+1/2) = (2/\epsilon)w(Y)$ and $0 < w(Y) \le 1/4$ on the square. The outcome-one choice above therefore gives an increase of at least w. For the outcome-zero choice, monotonicity of w gives an increase of at least

$$-\frac{2(1-\epsilon)}{\epsilon}w + \left(\frac{2}{\epsilon} - 1\right)w = w.$$

Thus at every node some successor increases F by at least $\frac{1}{4}(2/\epsilon)^{-4R}$, since $Y \geq -R$. Following these successors for T rounds, and using $F(0,0) > 0$ and $F \leq R + 1/4$, yields

$$\frac{T}{4}\left(\frac{2}{\epsilon}\right)^{-4R} \leq F(X_T, Y_T) - F(0,0) \leq R + \frac{1}{4}.$$

Rearranging proves (F.2). □

Proposition F.1 (Quantitative logarithmic lower bound)**.** *Let nonnegative $f, g \in \mathcal{H}_K$ satisfy the conditions above for $\epsilon \in (0,1)$ and $\eta > 0$, and put $M = \max\{\|f\|_{\mathcal{H}_K}, \|g\|_{\mathcal{H}_K}\}$. For every $T \geq 2$,*

$$V_T(K) \geq V_T^{\mathrm{obl}}(K) \geq \frac{\eta}{4M(1 + \log(2/\epsilon))} \log T.$$

For each horizon, one oblivious outcome distribution gives this bound against every randomized learner. In particular, $V_T^{\mathrm{obl}}(K) = \Omega_K(\log(T+1))$ for all $T \geq 1$. Taking tests of norm at most one proves Theorem 7.

Proof. First suppose that both test norms are at most one. Fix a randomized learner and let B be its worst expected terminal error over fixed outcome sequences, as in Lemma F.1. Set $R = B/\eta$. At each prescribed outcome history, define

$$X = \frac{\mathbb{E}_A[S_f]}{\eta}, \qquad Y = -\frac{\mathbb{E}_A[S_g]}{\eta}.$$

Both coordinates lie in $[-R, R]$ by Lemma F.1. Let P be the learner's next random forecast at that prescribed history, with the expectation still taken over its original private randomness, and put

$$a = \frac{\mathbb{E}[(1-P)f(P)]}{\eta} \geq 0.$$

Appending one increases X by a. By (F.1), appending zero decreases X by at most $((1-\epsilon)/\epsilon)a$. The same outcome increases Y by

$$\frac{\mathbb{E}[Pg(P)]}{\eta} \geq 1 - a.$$

These statements concern the two prescribed successors of the same history, so they use the same distribution of P. Thus Lemma F.2 applies to the entire tree of mean scores. Taking logarithms and using $\log(1+4R) \leq 4R$ gives

$$\log T \leq \log(1+4R) + 4R\log(2/\epsilon) \leq 4(1 + \log(2/\epsilon))R.$$

Substituting $R = B/\eta$ proves the claimed worst-fixed-sequence bound for this learner. The path in the duration argument depends on the learner and the distribution of its seed, but never on the realized seed. The kernel error is continuous in the forecasts, so Lemma B.1 supplies one oblivious outcome law hard for all randomized learners at this horizon. The inequality $V_T(K) \geq V_T^{\mathrm{obl}}(K)$ gives the adaptive-game lower bound.

For general test norms, $M > 0$ because $\eta > 0$. Apply the preceding argument to $f/M, g/M$, whose gain lower bound is η/M. This gives the displayed constant.

Finally,

$$\eta \leq (1-p)f(p) + pg(p) \leq M\|\phi(p)\| \qquad (0 \leq p \leq 1).$$

At horizon one, an independent fair outcome therefore gives expected error at least $\eta/(2M)$, since $\mathbb{E}[|Y - p| \mid p] = 1/2$. Combining this with the bound for $T \geq 2$ and $\log(T+1) \leq 2\log T$ proves the all-positive-horizon statement. The hard law may depend on the horizon. □

The exact support restriction is sufficient for this proof, not claimed necessary for every logarithmic-rate kernel. Combining Theorem 7 and Corollary 2 gives a tight logarithmic rate whenever both sets of hypotheses hold; a cancellation cover alone need not rule out bounded error.

F.2 A SMOOTH, INJECTIVE LOGARITHMIC EXAMPLE

We now verify the smooth rank-two example using the same generic learner as the other fast upper bounds. For $t \in \mathbb{R}$, define

$$b(t) = \begin{cases} e^{-1/t}, & t > 0, \\ 0, & t \le 0, \end{cases} \qquad \sigma(t) = \frac{b(t)}{b(t) + b(1-t)},$$

and put

$$h(p) = \sigma(6p-1) - \sigma(5p-2) + \sigma(6p-4), \qquad \phi(p) = (1+p)(1-h(p), h(p)). \tag{F.3}$$

The denominator defining σ never vanishes. This function is C^∞, equals zero for $t \le 0$, equals one for $t \ge 1$, and is strictly increasing between zero and one.

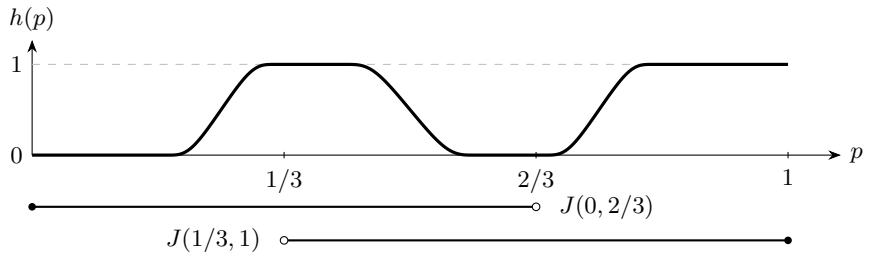


Figure 4: A smooth logarithmic example: $\phi(p) = (1+p)(1-h(p), h(p))$, with h from (F.3). The pairs $(0, 2/3)$ and $(1/3, 1)$ repeat the first and second positive feature directions. Their intervals, shown below the graph, cover every point. The factor $1+p$ makes the full feature map injective, so repeated directions do not require repeated vectors.

Corollary F.1 (Nondegenerate smooth logarithmic kernel)**.** *The kernel $K(p,q) = \langle \phi(p), \phi(q) \rangle$ from (F.3) is C^∞ and has rank two. Its feature map is injective, has nonzero derivative everywhere, and satisfies $K(p,p) \ge 1/2$. Its minimax rate is $\Theta_K(\log(T+1))$. The upper bound is achieved by Algorithm 1 using only the forecasts $0, 1/3, 2/3, 1$.*

Proof. Every derivative of $e^{-1/t}$ tends to zero as $t \downarrow 0$, proving the smoothness of σ at its transition endpoints. On $(0,1)$, strict increase follows by differentiating $\log(b(t)/b(1-t))$. The three transition intervals in (F.3) are disjoint: $(1/6, 1/3)$, $(2/5, 3/5)$, and $(2/3, 5/6)$. Thus h takes values in $[0,1]$, has finitely many constant or strictly monotone pieces, equals zero on $[0, 1/6]$, equals one on $[5/6, 1]$, and satisfies $h(1/3) = 1$, $h(2/3) = 0$.

The following two positive-ray pairs form a cancellation cover:

Pair (ℓ_i, u_i)	Unit direction v_i	α_i	β_i
$(0, 2/3)$	$(1, 0)$	1	$5/3$
$(1/3, 1)$	$(0, 1)$	$4/3$	2

Their relatively open intervals are $[0, 2/3)$ and $(1/3, 1]$, whose union is $[0,1]$. A valid cover margin is $\delta = 1/6$: use the first pair when $q \le 1/2$ and the second when $q \ge 1/2$, with the boundary-endpoint exceptions in the cover condition. In the gain representation there are $n = 2$ coordinates and the gain bound is $B = 2$. The learning rate in Algorithm 1 can therefore be chosen as $\eta_{\mathrm{MW}} = \delta/(2B) = 1/24$. The reduction in Corollary 2 and the bound (12) give

$$\mathbb{E}[\mathrm{kCE}_{K,T}] \le 48 \log\left(2 + \frac{T}{144}\right). \tag{F.4}$$

Only the chosen pair's accumulator is updated on each round, as in the general cancellation-cover reduction.

For the lower bound, use the coordinate tests

$$f(p) = (1+p)(1-h(p)), \qquad g(p) = (1+p)h(p).$$

The feature vectors span $\mathbb{R}^2$, so both tests have RKHS norm one. The first vanishes on $[5/6, 1]$, giving $\epsilon = 1/6$. On $[1/6, 5/6]$,

$$(1-p)f(p) + pg(p) = (1+p)((1-p)(1-h(p)) + ph(p)) \geq (1+p)\min\{p, 1-p\} \geq 1/6.$$

On $[0, 1/6]$ this expression is $1 - p^2 \geq 35/36$, and on $[5/6, 1]$ it is $p(1+p) \geq 55/36$. Thus the lower-bound parameter $\eta = 1/6$ is valid in Theorem 7, giving the matching logarithmic lower bound directly.

Finally, the coordinate sum of $\phi(p)$ is $1 + p$, proving injectivity; the sum of its coordinate derivatives is one, so its derivative never vanishes. Also

$$\|\phi(p)\|^2 = (1+p)^2((1-h(p))^2 + h(p)^2) \geq \tfrac{1}{2}.$$

The endpoint vectors are linearly independent, giving rank exactly two. All feature representations are isometric on their closed feature spans, so these injectivity and derivative properties hold for the canonical feature map as well. □

F.3 The general rank-two classification

We prove a version that allows an arbitrary positive radial factor and any two independent feature directions. The family in Definition 9 is its special case with radial factor $1 + p$ and orthonormal coordinate directions.

Definition F.1 (Rank-two profile and covered points)**.** Let v_0, v_1 be linearly independent vectors in a real Hilbert space. Let $\rho : [0, 1] \to (0, \infty)$ and $h : [0, 1] \to [0, 1]$ be continuous, with $h(0) = 0$ and $h(1) = 1$. Define

$$\phi(p) = \rho(p)((1 - h(p))v_0 + h(p)v_1), \qquad K(p, q) = \langle \phi(p), \phi(q) \rangle. \tag{F.5}$$

An interior point q is *covered* if $h(s) = h(t)$ for some $s < q < t$. Point zero is covered if $h(t) = 0$ for some $t > 0$, and point one is covered if $h(s) = 1$ for some $s < 1$.

Corollary F.2 (Classification with finite monotonicity and an endpoint plateau)**.** *In Definition F.1, suppose h is constant or strictly monotone on each piece of a finite partition, and $h = 0$ on $[0, \epsilon]$ or $h = 1$ on $[1 - \epsilon, 1]$ for some $\epsilon \in (0, 1)$. Then*

$$V_T(K) = \begin{cases} \Theta_K(\log(T+1)), & \text{if every point is covered,} \\ \Theta_K(\sqrt{T}), & \text{otherwise.} \end{cases} \tag{F.6}$$

The same rates hold for $V_T^{\mathrm{obl}}(K)$. The logarithmic upper condition is equivalent to the existence of a finite cancellation cover.

We first identify the two coordinate tests and prove the logarithmic branch. The square-root branch will follow from the two auxiliary statements below.

Coordinate tests and repeated directions. The endpoint vectors span $\mathrm{span}\{v_0, v_1\}$, so the kernel has rank two. Let w_0, w_1 be the dual basis in this span, defined by $\langle w_i, v_j \rangle = \mathbf{1}\{i = j\}$. The associated RKHS tests are

$$f(p) = \rho(p)(1 - h(p)), \qquad g(p) = \rho(p)h(p). \tag{F.7}$$

Their RKHS norms are finite but need not equal one. Also, $\phi(s)$ and $\phi(t)$ lie on the same positive ray exactly when $h(s) = h(t)$: comparing coefficients in the independent basis v_0, v_1 and then summing them proves the forward implication, and the positive factors $\rho(s), \rho(t)$ give the converse.

Logarithmic branch of Corollary F.2. Suppose first that $h = 1$ on $[1 - \epsilon, 1]$. Then f vanishes there. Moreover,

$$(1-p)f(p) + pg(p) = \rho(p)((1-p)(1-h(p)) + ph(p)) > 0 \qquad (0 \leq p \leq 1).$$

At the endpoints this follows from $h(0) = 0, h(1) = 1$; in the interior the bracket is a convex combination of the positive numbers $1 - p$ and p. Continuity and compactness give a uniform lower bound $\eta > 0$. Apply Proposition F.1, with the norm factor $M = \max\{\|f\|_{\mathcal{H}_K}, \|g\|_{\mathcal{H}_K}\}$, to obtain an $\Omega_K(\log(T+1))$ lower bound whether or not every point is covered.

If instead $h = 0$ near zero, reflect the protocol by $p' = 1 - p$, $y' = 1 - y$. The reflected features $\widetilde{\phi}(p') = \phi(1 - p')$ have the same form after replacing h by $\widetilde{h}(p') = 1 - h(1 - p')$, replacing ρ by $\rho(1 - p')$, and exchanging v_0, v_1. The new profile equals one near one. The feature sum changes sign under this transformation, leaving its norm and the game values unchanged, so the same lower bound follows.

If every point is covered, all pairs $s < t$ with $h(s) = h(t)$ give a relatively open cover of $[0, 1]$ by their intervals $J(s, t)$ from the cancellation-cover definition. Compactness gives a finite subcover. The equivalence of equal levels and positive rays makes this a finite cancellation cover, so Corollary 2 gives $O_K(\log(T + 1))$. Conversely, every positive-ray cancellation pair has equal profile levels, so a finite cancellation cover covers all points. This proves the logarithmic branch and the asserted equivalence. □

Lemma F.3 (An uncovered point gives an interval of unique levels)**.** *Let $h : [0, 1] \to [0, 1]$ be continuous, with $h(0) = 0, h(1) = 1$, and constant or strictly monotone on each piece of a finite partition. If some point is not covered, there is a nonempty open interval $J \subset (0, 1)$ such that every $c \in J$ has exactly one preimage under h.*

Proof. First let $q \in (0, 1)$ be uncovered. The sets $A = h([0, q))$ and $B = h((q, 1])$ are disjoint intervals containing zero and one, respectively. Every element of A is less than every element of B: otherwise the interval property and these endpoint values would force an intersection. Set $c_0 = h(q)$. Continuity implies $h(p) \le c_0$ for $p < q$ and $h(p) \ge c_0$ for $p > q$; at least one side does not attain c_0.

Suppose that the right side does not attain c_0, so $h(p) > c_0$ for every $p > q$. Insert q into the finite partition if necessary. On the first partition segment to its right, h cannot be constant or strictly decreasing, and hence is strictly increasing. Choose r inside that segment with $q < r < 1$. On the compact set $[r, 1]$, the minimum $b = \min_{[r,1]} h$ satisfies $b > c_0$. Every level in (c_0, b) occurs exactly once in (q, r), nowhere to the left of q, and nowhere in $[r, 1]$. This is a nonempty interval of levels in $(0, 1)$.

If the left side does not attain c_0, the analogous partition argument gives a strictly increasing segment $[r, q]$ with $0 < r < q$. Set $a = \max_{[0,r]} h < c_0$. Every level in (a, c_0) occurs exactly once in (r, q) and nowhere else.

If zero is uncovered, then $h(p) > 0$ for all $p > 0$. Its initial partition segment must be strictly increasing. Choosing $r > 0$ inside that segment and using the positive minimum of h on $[r, 1]$ gives a nonempty interval of uniquely attained positive levels. The case of an uncovered point one follows by applying this argument to $p \mapsto 1 - h(1 - p)$. □

Proposition F.2 (Unique levels give square-root error)**.** *For the feature map (F.5), suppose that every level in a nonempty open interval $J \subset (0, 1)$ has a unique preimage. Then $V_T(K) = \Theta_K(\sqrt{T})$, with the same rate for $V_T^{\mathrm{obl}}(K)$. This implication requires neither an endpoint plateau nor a finite monotonicity partition.*

Proof. Choose $c_- < c_+$ with $[c_-, c_+] \subset J$, and define the RKHS tests

$$g_+(p) = \rho(p)(c_+ - h(p)), \qquad g_-(p) = \rho(p)(c_- - h(p)).$$

They are linear combinations of (F.7). Their half-difference is bounded below by

$$\gamma = \frac{c_+ - c_-}{2} \min_{p \in [0,1]} \rho(p) > 0.$$

For each $c \in [c_-, c_+]$, let r_c be its unique preimage. The intermediate value theorem and endpoint values imply $h(p) < c$ for $p < r_c$ and $h(p) > c$ for $p > r_c$. Hence the test $\rho(p)(c - h(p))$ has a sign threshold at r_c.

Continuity at zero and one keeps every r_c uniformly away from both endpoints: there is $a \in (0, 1/2)$ such that $h < c_-$ on $[0, a]$ and $h > c_+$ on $[1 - a, 1]$. Thus every convex combination of g_+, g_- has a sign threshold in $[a, 1 - a]$. By Lemma D.1, these tests admit a common response there for every forecast law. Apply Theorem 1 and use

$$\max\{|S_{g_+}(T)|, |S_{g_-}(T)|\} \le \max\{\|g_+\|_{\mathcal{H}_K}, \|g_-\|_{\mathcal{H}_K}\}\ \mathrm{kCE}_{K,T}$$

to obtain the square-root lower bound under one oblivious law at each horizon. The universal upper bound (9) completes the proof. □

Completion of Corollary F.2 and proof of Corollary 1. An uncovered point gives an interval of unique levels by Lemma F.3, and Proposition F.2 proves the remaining branch of Corollary F.2. Taking $\rho(p) = 1 + p$, $v_0 = (1, 0)$, and $v_1 = (0, 1)$ now proves Corollary 1 in full. □

Finite monotonicity is used only to pass from an uncovered point to unique levels; the endpoint plateau is used only for the quantitative logarithmic lower bound. Neither is needed for the general cancellation-cover upper bound.

G TIGHT RATES FOR DELAYED FEEDBACK

We take $T \geq 1$ integer. For each round i, the additional delay τ_i is an integer in $\{0, \dots, T - i\}$, and the label is revealed before prediction $r_i = i + 1 + \tau_i$. Revelation at $T + 1$ means after the last prediction. Labels arriving together are identified by their round indices. A forecaster may remember every earlier prediction and any labels already revealed. The total additional delay of a schedule is $\sum_i \tau_i$, at most $T(T-1)/2$.

G.1 A PATHWISE HILBERT-SPACE UPPER BOUND

The learner applies the same endpoint/root rule to the observed feature sum. We bound the difference from the full sum by counting labels that are still outstanding.

Proposition G.1 (Delayed defensive forecasting)**.** *Let $\phi : [0, 1] \to \mathcal{H}$ be norm-continuous, with $\max_p \|\phi(p)\| \leq \kappa$. There is a deterministic forecaster using only revealed feedback such that every binary transcript and every schedule satisfy*

$$\left\| \sum_{t=1}^{T} (y_t - p_t)\phi(p_t) \right\|^2 \leq \sum_t p_t(1 - p_t)\|\phi(p_t)\|^2 + 2\kappa^2 \sum_i \tau_i \leq \kappa^2 \left(T/4 + 2 \sum_i \tau_i \right). \quad \text{(G.1)}$$

The algorithm need not know future reveal times, the budget, or the horizon.

Proof. Before prediction t, form the observed feature sum

$$O_t = \sum_{i<t : r_i \leq t} (y_i - p_i)\phi(p_i)$$

and the continuous function $H_t(p) = \langle O_t, \phi(p) \rangle + (1/2 - p)\|\phi(p)\|^2$. Choose $p_t = 0$ if $H_t(0) \leq 0$; otherwise choose $p_t = 1$ if $H_t(1) \geq 0$; in the remaining case choose a zero of H_t. Such a zero exists by the intermediate value theorem. This rule ensures

$$(y - p_t)H_t(p_t) \leq 0 \qquad (y = 0, 1). \quad \text{(G.2)}$$

For definiteness one may use the smallest zero, which is a measurable choice for this continuous one-dimensional rule; at a fixed horizon there are also only finitely many possible labelled-feedback histories for a deterministic learner.

For each integer $s \in \{0, \dots, T\}$, let $S_s = \sum_{i=1}^{s} (y_i - p_i)\phi(p_i)$ be the full feature sum, used only in the analysis. Put $U_t = S_{t-1} - O_t$ and $u_t = |\{i < t : r_i > t\}|$. Then $\|U_t\| \leq \kappa u_t$. Expanding the squared norm and using the binary identity $(y - p)^2 - 2(y - p)(1/2 - p) = p(1 - p)$ gives

$$\begin{aligned} \|S_t\|^2 - \|S_{t-1}\|^2 &= 2(y_t - p_t)H_t(p_t) + p_t(1 - p_t)\|\phi(p_t)\|^2 + 2(y_t - p_t)\langle U_t, \phi(p_t) \rangle \\ &\leq p_t(1 - p_t)\|\phi(p_t)\|^2 + 2\kappa^2 u_t. \end{aligned}$$

Label i contributes to u_t at precisely $t = i + 1, \dots, i + \tau_i$, so $\sum_t u_t = \sum_i \tau_i$. Sum the displayed inequality. At zero delay it is exactly the kernel forecasting bound (9) of Vovk (2007). □

With the feature map in Proposition C.1, (G.1) gives

$$\begin{aligned} \mathrm{smCE}_T &\le \sqrt{T + 8\sum_i \tau_i}, \qquad \mathrm{LowerCalDist}_T \le 2\sqrt{T + 8\sum_i \tau_i}, \\ \mathrm{CalDist}_T &\le 2\sqrt{T + 8\sum_i \tau_i} + C_0\sqrt{T}. \end{aligned} \tag{G.3}$$

using Lemma C.2. Smooth error is at most T because each summand has absolute value at most one. Distance is also at most T, since the sequence $u_t = y_t$ is exactly calibrated. This proves the upper half of Theorem 3, uniformly over all schedules in the budget. The proof is pathwise and does not require the outcomes to be oblivious.

G.2 The two-test argument for hidden blocks

Partition the horizon into consecutive blocks $B_1, \dots, B_n$ with positive integer lengths $b_1, \dots, b_n$. Reveal all the labels of a block immediately after its last prediction, before the first prediction in the next block. Nature chooses one bit Y_j for block j and repeats it on all rounds of that block. No label from the current block is available when any of its predictions is made.

Lemma G.1 (Weighted two-test lower bound)**.** *Suppose g_+, g_- satisfy the separated hypotheses of Theorem 1, with parameters a, γ. For the block schedule above there is a randomized oblivious adversary that chooses the block bits in advance, such that every randomized forecaster satisfies, for independent fair signs $\varepsilon_1, \dots, \varepsilon_n$,*

$$\mathbb{E}\big[\max_i |S_{g_i}(T)|\big] \ge a\gamma\mathbb{E}\left[\left|\sum_{j=1}^n b_j\varepsilon_j\right|\right] \ge \frac{a\gamma}{\sqrt{2}}\left(\sum_j b_j^2\right)^{1/2}. \tag{G.4}$$

The forecaster may know the block schedule and the fact that each block has a repeated outcome in advance.

Proof. Write $\mathcal{L}(p, y) = \max\{|S_{g_+}(T)|, |S_{g_-}(T)|\}$ for the loss when each bit of $y \in \{0,1\}^n$ is repeated throughout its block, and set $\mathcal{E} = a\gamma\,\mathbb{E}[|\sum_{j=1}^n b_j\varepsilon_j|]$. We first fix a randomized learner A and construct an outcome distribution Q_A with $\mathbb{E}_{A,Q_A}[\mathcal{L}] \ge \mathcal{E}$, and then apply the minimax theorem to obtain one distribution for all learners.

For each history $y_{<j}$ of completed block bits, let $\bar\nu_j(y_{<j})$ be the equally weighted average, over $t \in B_j$, of the distributions of the learner's forecasts when the preceding block bits are prescribed to be $y_{<j}$. This averages over all the learner's randomness, including its earlier forecasts and private state; it need not equal the law announced on a particular play. This is an average of marginal distributions, not a product distribution on the block. The forecasts may be arbitrarily correlated and may depend on previous forecasts. What matters is that none depends on the still-hidden bit Y_j.

Choose a common response $q_j(y_{<j}) \in [a, 1-a]$ to $\bar\nu_j(y_{<j})$, and define an outcome distribution Q_A by $Y_j \mid Y_{<j} \sim \mathrm{Bernoulli}(q_j(Y_{<j}))$. Since there are only finitely many possible histories, choosing one common response for each history defines a measurable rule and hence a valid conditional law for Y_j. This specifies a law on complete block-bit sequences using the learner's strategy but not its realized randomness. Sample the sequence independently of that randomness.

Abbreviate $\bar\nu_j = \bar\nu_j(Y_{<j})$ and $q_j = q_j(Y_{<j})$. We use $\mathbb{E}_A$ to average over the learner's randomness with Y held fixed. For $i \in \{+,-\}$, consider

$$\begin{aligned} \mathbb{E}_A[S_{g_i}(T)] &= \sum\nolimits_{j=1}^n b_j\mathbb{E}_{P\sim\bar\nu_j}[(Y_j - P)g_i(P)], \\ M_i(Y) &:= \sum\nolimits_{j=1}^n b_j(Y_j - q_j)\mathbb{E}_{P\sim\bar\nu_j}[g_i(P)]. \end{aligned}$$

Based on the definition of common response, the first quantity is lower bounded by the second.

$$\mathbb{E}_A[S_{g_i}(T)] - M_i(Y) = \sum\nolimits_{j=1}^n b_j r_{g_i}(\bar\nu_j, q_j) \ge 0.$$

We can show that $\mathbb{E}_{Q_A}[M_i(Y)] = 0$ as each coefficient is determined by the preceding block bits, conditional on which $Y_j - q_j$ has mean zero.

Define $h_j = \mathbb{E}_{P\sim\bar{\nu}_j}[d(P)] \geq \gamma$ and $N = \sum_j b_j h_j (Y_j - q_j) = (M_+(Y) - M_-(Y))/2$. Then we have

$$\mathbb{E}_{A,Q_A}[\mathcal{L}] = \mathbb{E}_{A,Q_A}[\max_i |S_{g_i}(T)|] \geq \mathbb{E}_{Q_A}\left[\max_i \mathbb{E}_A[S_{g_i}(T)]\right] \geq \mathbb{E}_{Q_A}\left[\max_i M_i(Y)\right]$$
$$= \frac{1}{2}\mathbb{E}_{Q_A}[M_+(Y) + M_-(Y)] + \frac{1}{2}\mathbb{E}_{Q_A}[|M_+(Y) - M_-(Y)|] = \mathbb{E}_{Q_A}[|N|],$$

where the second line follows from the identity $\max\{u, v\} = (u + v + |u - v|)/2$.

We compare N with a Rademacher sum by generating the same outcome distribution Q_A in a different way. Independently of the learner's randomness, let $(\varepsilon_j)_{j=1}^n$ be i.i.d. with $\varepsilon_j \sim \mathrm{Rad}$, and generate the outcomes successively by

$$\mathbb{P}(Y_j = 1 \mid Y_{<j}, \varepsilon_{1:n}) = q_j + \frac{a\gamma}{h_j}\varepsilon_j.$$

These probabilities lie in $[0, 1]$, since $a\gamma/h_j \leq a$ and $q_j \in [a, 1-a]$. Earlier outcomes use only earlier Rademacher variables, so ε_j is independent of $Y_{<j}$. Averaging over it gives $\mathbb{P}(Y_j = 1 \mid Y_{<j}) = q_j$, and hence Y still has law Q_A. Moreover,

$$\mathbb{E}[b_j h_j (Y_j - q_j) \mid Y_{<j}, \varepsilon_{1:n}] = b_j h_j (\mathbb{E}[Y_j \mid Y_{<j}, \varepsilon_{1:n}] - q_j)$$
$$= a\gamma b_j \varepsilon_j.$$

Taking another conditional expectation gives

$$\mathbb{E}[N \mid \varepsilon_{1:n}] = \sum_j \mathbb{E}[\mathbb{E}[b_j h_j (Y_j - q_j) \mid Y_{<j}, \varepsilon_{1:n}] \mid \varepsilon_{1:n}]$$
$$= a\gamma \sum_{j=1}^n b_j \varepsilon_j.$$

Consequently,

$$\mathbb{E}[|N|] \geq \mathbb{E}\left[\left|\mathbb{E}[N \mid \varepsilon_{1:n}]\right|\right] = a\gamma\, \mathbb{E}\left[\left|\sum_{j=1}^n b_j \varepsilon_j\right|\right] = \mathcal{E}.$$

Together with the preceding bound $\mathbb{E}_{A,Q_A}[\mathcal{L}] \geq \mathbb{E}[|N|]$, this proves $\mathbb{E}_{A,Q_A}[\mathcal{L}] \geq \mathcal{E}$.

Applying the Khinchin inequality (Lemma B.2) with weights $b_1, \ldots, b_n$ gives

$$\mathcal{E} \geq (a\gamma/\sqrt{2})(\sum_j b_j^2)^{1/2}.$$

For each learner A, the expected loss under Q_A is an average of the expected losses on fixed block-bit sequences. Hence

$$\max_{y\in\{0,1\}^n} \mathbb{E}_A[\mathcal{L}(p, y)] \geq \mathbb{E}_{A,Q_A}[\mathcal{L}(p, Y)] \geq \mathcal{E}.$$

The loss $\mathcal{L}(p, y)$ is nonnegative and continuous in p for each fixed y because the tests are continuous. Applying the block-game version of Lemma B.1 gives

$$\max_{Q\in\Delta(\{0,1\}^n)} \inf_A \mathbb{E}_{A,Q}[\mathcal{L}(p, Y)] = \inf_A \max_y \mathbb{E}_A[\mathcal{L}(p, y)] \geq \mathcal{E}.$$

The bound proved above implies that a distribution Q_T attaining the maximum on the left guarantees the same lower bound against every randomized learner, including learners informed of the schedule. Thus Q_T defines the required oblivious adversary. □

Corollary G.1 (Simultaneous block lower bound)**.** *For every block partition above, one oblivious repeated-bit law gives*

$$\mathbb{E}\left[\min\{\mathrm{smCE}_T, \mathrm{LowerCalDist}_T, \mathrm{CalDist}_T\}\right] \geq \frac{1}{18\sqrt{2}}\left(\sum_j b_j^2\right)^{1/2} \tag{G.5}$$

against every randomized forecaster.

Proof. Use the explicit pair (5), with $a = 1/3$ and $\gamma = 1/6$, and the pathwise domination in Proposition 1. □

The comparison is a weighted fair walk in the analysis, not a claim that the actual block bits are independent. Their means may depend on the preceding block bits. The full outcome law nevertheless remains independent of the forecaster's realized randomness.

G.3 Choosing blocks within a total-delay budget

It remains to choose a deterministic schedule that uses the allowed budget efficiently. The cost of each hidden block is quadratic in its length, exactly as in the variance of the weighted fair walk.

Lemma G.2 (A feasible block schedule)**.** *For every $T \geq 1$ and $D \geq 0$, there is a partition into consecutive blocks for which the end-of-block reveal schedule uses at most D additional delay and*

$$\sum_j b_j^2 \geq \tfrac{1}{4} \min\{T^2, T + D\}. \tag{G.6}$$

Proof. Put $D_0 = \min\{D, T(T-1)/2\}$ and

$$b = \left\lfloor 1 + \frac{2D_0}{T} \right\rfloor \in \{1, \dots, T\}.$$

Use as many full blocks of length b as possible and, if needed, one final shorter block. In a block of length ℓ, the additional delays are $\ell - 1, \ell - 2, \dots, 0$, so its cost is $\ell(\ell-1)/2$. Since $\ell \leq b$, the total cost is at most

$$\frac{T(b-1)}{2} \leq D_0 \leq D.$$

Let $k = \lfloor T/b \rfloor$. The full blocks contain $kb \geq T/2$ rounds: the remaining block has fewer than b rounds, and $k \geq 1$. Thus

$$\sum_j b_j^2 \geq kb^2 \geq Tb/2 \geq (T + 2D_0)/4,$$

where the last inequality uses $\lfloor x \rfloor \geq x/2$ for $x \geq 1$. If $D \leq T(T-1)/2$, then $T + 2D_0 \geq T + D$. Otherwise $T + 2D_0 = T^2$. In either case $T + 2D_0 \geq \min\{T^2, T + D\}$, proving (G.6). □

Completion of Theorem 3. Choose the schedule in Lemma G.2 and apply Corollary G.1. The expected minimum of all three losses is at least

$$\frac{1}{36\sqrt{2}} \min\{T, \sqrt{T + D}\}.$$

The schedule depends only on T, D; the universal hard law for that schedule does not depend on the learner. The uniform upper bounds were proved in (G.3). This proves all three rates and the simultaneous lower-bound statement. □

Adaptive labels and an optional schedule extension. The value in Definition 6 already allows adaptive labels. Its upper bound is pathwise, and its lower bound uses a fixed schedule and a randomized oblivious adversary, so neither proof requires a minimax exchange for adaptive Nature. The same rate also holds in the following larger schedule game: at round i, after the forecast law is announced but without seeing its fresh draw, Nature commits to y_i and $\tau_i \in \{0, \dots, T - i\}$, subject to $\sum_i \tau_i \leq D$ on every path. The learner observes the label only at its scheduled reveal time. Every realized schedule obeys the upper bound, while the fixed block schedule remains a permissible lower-bound strategy.

Comparison with delayed multicalibration. Farina & Perdomo (2026b, Section 3.2.1) give a general reduction to multicalibration with delayed observations and a square-root total-delay dependence. Their forecast-distribution formulation controls cumulative expected test increments under the issued distributions. The potential calculation here gives a deterministic bound on the realized feature norm for a norm-continuous kernel, and Lemma C.2 transfers it to calibration distance. The matching worst-schedule lower bound is a separate consequence of the hidden-block two-test construction. The upper-bound dependence on total delay is not claimed to originate here.

H FURTHER EXAMPLES AND LIMITS OF THE RATE CRITERIA

The preceding criteria are deliberately not a general kernel classification. We give examples showing why several simpler summaries of a kernel do not determine its rate, and a weaker condition that rules out bounded error without proving a logarithmic lower bound.

H.1 FEATURE IMAGES AND SPECTRAL SUMMARIES DO NOT DETERMINE THE RATE

The feature maps

$$\phi_A(p) = (\cos(\pi p), \sin(\pi p)), \qquad \phi_B(p) = (\cos(\pi \sin^2(\pi p)), \sin(\pi \sin^2(\pi p)))$$

have the same image, the upper semicircle, and both have unit diagonal. The first kernel has rate $\Theta(\sqrt{T})$ by Theorem 5. The second has identical endpoint vectors and therefore rate $\Theta(1)$ by Lemma E.7. The probability labels matter because the increments are $-p\phi(p)$ and $(1-p)\phi(p)$, not just because of the unlabelled image of ϕ.

Even rank, smoothness, endpoint inner products, and integral-operator eigenvalues together do not determine the rate. The kernels K_π and $K_{3\pi}$ are analytic, stationary, rank two, and unit-diagonal. Their endpoint Gram matrices both equal $\left(\begin{smallmatrix} 1 & -1 \\ -1 & 1 \end{smallmatrix}\right)$. Under Lebesgue measure on $[0,1]$, the two feature functions $\cos(n\pi p)$ and $\sin(n\pi p)$, for $n = 1, 3$, are orthogonal and each has squared L^2 norm $1/2$. Thus the associated kernel integral operators have the same two nonzero eigenvalues $1/2, 1/2$. Their online rates nevertheless differ by Theorem 5. These operator eigenvalues are distinct from the frequency support in Theorem D.1.

H.2 A WEAKER CONDITION RULING OUT BOUNDED ERROR

The sufficient condition in Theorem 7 requires a nonzero test to vanish on a terminal interval $(1-\epsilon, 1]$, which a real-analytic test cannot do. This limits that criterion, not the possibility of logarithmic lower bounds by other arguments. The following weaker statement separates unboundedness from a quantitative logarithmic rate.

Proposition H.1 (Nonnegative tests excluding bounded error)**.** *Suppose $f, g \in \mathcal{H}_K$ are nonnegative and satisfy*

$$f(0) > 0, \qquad f(1) = 0, \qquad g(p) > 0 \text{ whenever } f(p) = 0. \tag{H.1}$$

Then $V_T^{\mathrm{obl}}(K) \to \infty$, and consequently $V_T(K) \to \infty$, as $T \to \infty$.

Proof. Suppose that $V_T^{\mathrm{obl}}(K)$ has a bounded subsequence along horizons tending to infinity. Choose learners whose worst-fixed-sequence terminal expected kernel errors are bounded by a common constant (an additive approximation to the infimum suffices). By Lemma F.1, their mean states $U(h) = \mathbb{E}_A[S_f(h)]$, $V(h) = -\mathbb{E}_A[S_g(h)]$ lie in one fixed square $[-R, R]^2$ at every available forced history.

The space $\mathcal{P}([0,1])$ is weakly compact and metrizable. There are countably many finite binary histories. Passing to a diagonal subsequence makes the seed-averaged forecast law converge at each fixed history. The expectations of the four continuous functions defining the mean increments converge under weak convergence of these laws. Therefore the limiting laws define an infinite binary tree of mean states, all in $[-R, R]^2$, with both successors given by the corresponding expectations.

Let C be the closure of the set of these states. It is nonempty and compact. At every $z \in C$, some forecast law has both mean successors in C: approximate z by tree states, take a weakly convergent subsequence of their forecast laws, and pass to the limit in both increments. Choose a state $z = (U, V) \in C$ maximizing U, and among those maximizing V. Let ν be such a law. Its outcome-one increment of U is $\mathbb{E}_{P\sim\nu}\big[(1-P)f(P)\big] \geq 0$, so maximality forces this increment to be zero. Because $f(1) = 0$, the zero set of $(1-p)f(p)$ is exactly $Z_f = \{p : f(p) = 0\}$. Hence ν is supported on Z_f.

This set is compact and disjoint from zero, and g is positive on it. Thus $\min_{p\in Z_f} pg(p) > 0$. The outcome-zero successor has unchanged U, since f vanishes on the support, but a strictly larger V, since its increment is $\mathbb{E}_{P\sim\nu}\big[Pg(P)\big] > 0$. This contradicts the choice of z. No bounded subsequence of $V_T^{\mathrm{obl}}(K)$ can occur along unbounded horizons. The inequality $V_T(K) \geq V_T^{\mathrm{obl}}(K)$ proves the second claimed limit. □

Corollary H.1 (An analytic example with unbounded error and a logarithmic upper bound)**.** *For*

$$K_{\mathrm{an}}(p,q) = \frac{1+\cos(3\pi(p-q))}{2}, \tag{H.2}$$

one has $V_T(K_{\mathrm{an}}) \to \infty$ *and* $V_T(K_{\mathrm{an}}) = O(\log(T+1))$.

Proof. A feature representation is

$$\psi(p) = \left(\frac{1+\cos(3\pi p)}{2}, \frac{1-\cos(3\pi p)}{2}, \frac{\sin(3\pi p)}{\sqrt{2}}\right).$$

Expanding the inner product verifies (H.2); the kernel is analytic, stationary, rank three, and unit-diagonal. The features at $0, 2/3$ both equal the first basis vector, and those at $1/3, 1$ both equal the second. Thus these two pairs give a cancellation cover and, by Corollary 2, an $O(\log(T+1))$ upper bound with Algorithm 1. Its first two coordinate functions are nonnegative, and the first vanishes only at $1/3$ and 1, where the second equals one. They satisfy (H.1), so Proposition H.1 applies. □