Minimax Calibration Distance under Simultaneous Play

Abstract

We establish the $\Theta(\sqrt{T})$ minimax rate of calibration distance in sequential binary prediction under simultaneous play. This closes the lower-bound exponent gap posed by Qiao and Zheng [13], from $\Omega(T^{1/3})$ to $\Omega(\sqrt{T})$, for behavioral-kernel randomized forecasters against deterministic past-adaptive adversaries. The adversary may use past forecasts and outcomes but cannot observe the current private forecast. Calibration distance is the total ℓ_1 change needed to make the forecast sequence perfectly calibrated in hindsight. For every $T \geq 2$, every such forecaster incurs expected distance at least $\sqrt{T}/4320$, while a deterministic algorithm guarantees distance below $\sqrt{T}$ on every outcome path. For deterministic forecasters, one fixed outcome string also forces $\Omega(\sqrt{T})$ distance. The upper bound improves the explicit $2\sqrt{T}+1$ guarantee of Arunachaleswaran et al. [1]. Its shifted-grid construction admits exact integer optimization and $O(\log(T+1))$ worst-case time per round in the word-RAM model. Restarting gives a horizon-free bound below $(2+\sqrt{2})\sqrt{T}$ for every prefix, using $O(\sqrt{T})$ memory. The lower bound combines a mean-reverting adversary, Lipschitz calibration witnesses, and martingale fluctuations; averaging then fixes a deterministic adversary. Together, these results determine the minimax exponent in the stated simultaneous-play model. The randomized-forecaster rate against a fixed outcome string remains open.

Note. This paper was generated entirely by AI, including MiMo, using an automated research pipeline developed by Chenghua Liu and Hanyu Li.

1 Introduction

Calibration asks that among rounds assigned probability z, the event occur with frequency z. This foundational criterion in sequential prediction [5, 8] is commonly measured by cumulative unbinned ℓ_1 calibration error (ECE). ECE is discontinuous in the forecasts: an arbitrarily small perturbation can split a bin and change the score sharply. Deterministic self-calibration is impossible in the revealed-forecast setting [10], whereas randomized calibration follows from approachability and minimax arguments [6, 8].

Błasiok, Gopalan, Hu, and Nakkiran proposed the ℓ_1 distance to a nearest perfectly calibrated sequence as a continuous calibration measure [2]. Qiao and Zheng initiated its sequential study and proved expected upper and lower bounds of $O(\sqrt{T})$ and $\Omega(T^{1/3})$, respectively [13]. Arunachaleswaran, Collina, Roth, and Shi later gave an explicit deterministic pathwise bound $2\sqrt{T}+1$ [1].

Our results. We close the exponent gap for the behavioral-kernel randomized model against a deterministic past-adaptive adversary. The two players act simultaneously: the adversary knows the forecasting strategy and public history but not the current private draw. Every forecaster incurs expected calibration distance at least $\sqrt{T}/4320$. For deterministic forecasters, specializing the construction gives the companion result that one fixed outcome string, chosen after the rule, forces distance at least $\sqrt{T}/1200$.

Our upper bound is deterministic and holds on every outcome path, so it matches the order of both lower bounds. It also improves the leading constant in the known explicit construction from 2 to 1. We optimize its finite-horizon grid exactly, reduce its implementation time from a naive $O(\sqrt{T})$

scan to $O(\log(T+1))$ worst-case time per round. Building on the restart framework of Collina et al. [3], block subadditivity yields a horizon-free guarantee with the explicit constant $2+\sqrt{2}$ for all prefixes of the same run.

Proof ideas. For the lower bound, the witness $f_0(u) = 1/2 - u$ charges the mean deviation and conditional forecast variance under an auxiliary mean-reverting Bernoulli adversary. Low expected energy leaves a square-root martingale fluctuation after controlling the drift; averaging then fixes one adversary tape. For the upper bound, shifted grid biases control an outcome-dependent proxy, and Lipschitz stability transfers this control to the announced forecasts.

2 Protocols and main results

This section defines the loss and the two information structures used in the paper, then states the matching minimax bounds. All measurable spaces below carry their Borel sigma-fields.

2.1 Calibration distance

Fix $T \geq 1$. For predictions $p = (p_1, \ldots, p_T) \in [0,1]^T$ and binary outcomes $y = (y_1, \ldots, y_T) \in \{0,1\}^T$, let

$$\mathcal{C}(y) := \left\{ r \in [0,1]^T : \sum_{t: r_t = z} (y_t - z) = 0 \text{ for every } z \in [0,1] \right\}. \tag{2.1}$$

Thus r is perfectly calibrated in hindsight. Empty sums are zero. The set is nonempty: the constant sequence $r_t = T^{-1} \sum_{s=1}^T y_s$ belongs to $\mathcal{C}(y)$. Define the unnormalized calibration distance by

$$\mathrm{CalDist}(p, y) := \inf_{r \in \mathcal{C}(y)} \sum_{t=1}^{T} |p_t - r_t| \,. \tag{2.2}$$

The infimum is in fact attained. Every value used by a calibrated sequence is the empirical mean of a nonempty subset of the outcomes, so only finitely many calibrated vectors are possible for fixed y.

For any $u \in [0,1]^T$, let $\mathrm{Im}(u) := \{u_t : 1 \leq t \leq T\}$ denote its set of occupied forecast values. The cumulative unbinned ℓ_1 calibration error is

$$\mathrm{ECE}(u, y) := \sum_{z \in \mathrm{Im}(u)} \left| \sum_{t: u_t = z} (y_t - z) \right|. \tag{2.3}$$

We use cumulative losses throughout. Their normalized versions are

$$\mathrm{CalDist}_{\mathrm{norm}}(p, y) := \frac{\mathrm{CalDist}(p, y)}{T}, \qquad \mathrm{ECE}_{\mathrm{norm}}(u, y) := \frac{\mathrm{ECE}(u, y)}{T}.$$

2.2 Simultaneous and oblivious protocols

Before round t, the public history belongs to

$$\mathsf{H}_{t-1} := ([0,1] \times \{0,1\})^{t-1}, \qquad h_{t-1} = ((p_s, y_s))_{s<t}.$$

A behavioral-kernel forecaster is a collection of Borel probability kernels $K_{t,T}(\cdot \mid h)$ from H_{t-1} to $[0,1]$ and draws $P_t \sim K_{t,T}(\cdot \mid h_{t-1})$. These public-history kernels, without additional persistent private state, are the randomized strategy class considered in this paper. A deterministic adaptive adversary consists of Borel-measurable maps

$$A_{t,T} : \mathsf{H}_{t-1} \to \{0,1\}, \qquad Y_t = A_{t,T}(h_{t-1}).$$

The two actions are simultaneous: $A_{t,T}$ does not observe the realized P_t. For fixed K and A, the transition kernel

$$K_{t,T}(dp \mid h)\, \delta_{A_{t,T}(h)}(dy)$$

recursively induces a unique trajectory law, denoted $\mathbb{P}^{K,A}$, with expectation $\mathbb{E}^{K,A}$. This is the finite-horizon Ionescu–Tulcea construction. The loss $\mathrm{CalDist}(P,Y)$ is bounded and Borel measurable, so the expectation below is well defined. We set

$$V_T^{\mathrm{rand,ad}} := \inf_K \sup_A \mathbb{E}^{K,A}[\mathrm{CalDist}(P,Y)]\,. \tag{2.4}$$

The expectation is over the forecaster's private randomization. Both players may know T, and the adversary knows K but not the current private draw. The lower bound below exhibits a deterministic adversary in the stated class.

For the companion oblivious result, a deterministic forecaster consists of maps

$$F_{t,T} : \{0,1\}^{t-1} \to [0,1], \qquad p_t = F_{t,T}(y_1, \ldots, y_{t-1}).$$

Write $p^F(y)$ for its path on y and set

$$V_T^{\mathrm{det,obl}} := \inf_F \sup_{y \in \{0,1\}^T} \mathrm{CalDist}\big(p^F(y), y\big). \tag{2.5}$$

Because the current prediction is determined by the past, allowing a past-adaptive adversary does not change the worst-case value in (2.5). For later comparison, if y is fixed and K is randomized, then $P^K(y)$ denotes the random path obtained recursively by sampling $P_t \sim K_{t,T}(\cdot \mid ((P_s, y_s))_{s<t})$; its law contains only the forecaster's randomization.

2.3 Main minimax theorem

The next theorem closes the exponent gap in both stated protocols. The same pathwise construction supplies their upper bounds.

For $m \in \mathbb{Z}_{\geq 1}$, write

$$g_T(m) := \frac{T}{2m} + \frac{m-1}{2}, \qquad U_T := \min_{m \in \mathbb{Z}_{\geq 1}} g_T(m). \tag{2.6}$$

Theorem 2.1 (Square-root minimax rate)**.** *For every integer $T \geq 2$,*

$$\frac{1}{4320}\sqrt{T} \leq V_T^{\mathrm{rand,ad}} \leq U_T < \sqrt{T}, \tag{2.7}$$

$$\frac{1}{1200}\sqrt{T} \leq V_T^{\mathrm{det,obl}} \leq U_T < \sqrt{T}. \tag{2.8}$$

An integer minimizer in (2.6) *is*

$$m_T^* = \left\lceil \frac{\sqrt{4T+1}-1}{2} \right\rceil. \tag{2.9}$$

For $T = 1$, both minimax values equal $1/2$.

Consequently, $V_T^{\mathrm{rand,ad}}/T, V_T^{\mathrm{det,obl}}/T = \Theta(T^{-1/2})$. The upper algorithm is deterministic and pathwise, and hence works against arbitrary outcome generation, including an adversary that sees the current forecast.

Restarting this finite-horizon construction gives a single algorithm for all prefixes. The following corollary records the constant and implementation cost proved in Section 7.

Corollary 2.2 (Horizon-free guarantee)**.** *There is a single deterministic forecaster, independent of the terminal horizon, such that for every infinite outcome sequence and every prefix $T \geq 1$,*

$$\mathrm{CalDist}(p_{1:T}, y_{1:T}) < (2+\sqrt{2})\sqrt{T}. \tag{2.10}$$

In the word-RAM model of Proposition 6.5, with word size $\Omega(\log(T+1))$, it uses $O(\log(T+1))$ time per round and $O(\sqrt{T})$ memory at time T.

Remark 2.3 (The remaining quantifier order)*.* For a randomized forecaster against a fixed string, define

$$V_T^{\mathrm{rand,obl}} := \inf_K \sup_{y\in\{0,1\}^T} \mathbb{E}\left[\mathrm{CalDist}(P^K(y), y)\right].$$

Fix jointly Borel sampling maps that realize the kernels by inverse transforms, and let $W = (W_1, \ldots, W_T)$ have independent $\mathrm{Unif}[0,1]$ coordinates. For each fixed w, let F^w be the resulting deterministic rule and set $L(w,y) := \mathrm{CalDist}(p^{F^w}(y), y)$. The deterministic result bounds $\mathbb{E}_W[\sup_y L(W,y)]$. The displayed value instead requires control of $\sup_y \mathbb{E}_W[L(W,y)]$, and these two operations cannot be exchanged in this direction. Conversely, the adversary for (2.7) depends on past realized forecasts and is not a fixed string. Thus $V_T^{\mathrm{rand,obl}}$ remains open.

3 Related work

We separate the closest work on calibration distance from results for the discontinuous ECE objective. This distinction is essential because the two losses support different lower-bound mechanisms.

Classical sequential calibration. Dawid's formulation of calibrated probability assessment [5] was followed by Oakes's deterministic impossibility [10] and randomized calibration via approachability [6, 8]. In our simultaneous protocol, a deterministic rule's current output is known from the public history; choosing $y_t = \mathbf{1}\{p_t < 1/2\}$ gives $\mathrm{ECE}(p,y) = \sum_t |y_t - p_t| \geq T/2$. This finite-horizon calculation is ours, not the literal statement of Oakes's historical theorem.

Kakade and Foster developed deterministic weak calibration [9]; Foster and Hart established deterministic smooth calibration in the leaky protocol, where Nature observes the current forecast [7]. Both differ from the distance studied here.

Distance to calibration. Błasiok et al. introduced exact ℓ_1 distance to calibration and compared it with continuous surrogates [2]. Qiao and Zheng transferred the metric to sequential prediction and proved an expected $O(\sqrt{T})$ upper bound and an $\Omega(T^{1/3})$ lower bound using i.i.d. Bernoulli(1/2) outcomes with adaptive stopping [13]. Without stopping, the same paper gives expected $O(\log^{3/2} T)$ distance against i.i.d. fair bits, so ordinary sampling noise alone cannot prove a square-root lower bound. Their work also records the standard comparison $\mathrm{CalDist} \geq \mathrm{smCE}/2$ used below.

Arunachaleswaran et al. gave an elementary deterministic predictor with pathwise distance at most $2\sqrt{T}+1$ [1]. Collina et al. already used doubling for a horizon-free pathwise $O(\sqrt{T})$ guarantee [3]. Our upper-bound contributions are the improved leading constant, exact integer-grid optimization, explicit anytime constant, and worst-case implementation. Post-hoc computation and estimation are studied by Qiao [11].

The discontinuous ECE objective. The sequential ECE landscape is different. Qiao and Valiant introduced sidestepping and proved an expected $\Omega(T^{0.528})$ lower bound in the revealed-forecast protocol [12]. Dagan et al. gave a randomized expected $O(T^{2/3-\varepsilon})$ upper bound for some $\varepsilon > 0$ and an $\Omega(T^{0.54389})$ lower bound for an oblivious distribution over outcome strings [4]; Zhang later gave an efficient randomized forecaster with expected $O(T^{2/3-\varepsilon})$ ECE for some $\varepsilon > 0$ [14]. These results neither imply nor contradict ours: CalDist $\leq$ ECE on every path, so an ECE lower bound does not lower-bound calibration distance. Lipschitz stability enables our proxy comparison, while the shifted-bias invariant supplies the pathwise proxy control.

4 Calibration witnesses

The following elementary inequality is the interface between calibration distance and the lower-bound process.

Lemma 4.1 (Witness comparison)**.** *For every integer $T \geq 1$, $p \in [0,1]^T$, $y \in \{0,1\}^T$, and Lipschitz $f : [0,1] \to \mathbb{R}$, let $L := \mathrm{Lip}(f)$ be its least Lipschitz constant and $B := \|f\|_\infty$. Then*

$$\left|\sum_{t=1}^{T} f(p_t)(y_t - p_t)\right| \leq (L + B)\,\mathrm{CalDist}(p, y). \tag{4.1}$$

Consequently,

$$\mathrm{CalDist}(p, y) \geq \left|\sum_{t=1}^{T} (y_t - p_t)\right|, \tag{4.2}$$

$$\mathrm{CalDist}(p, y) \geq \frac{2}{3}\left|\sum_{t=1}^{T} (1/2 - p_t)(y_t - p_t)\right|. \tag{4.3}$$

Proof. Fix $r \in \mathcal{C}(y)$. Calibration and grouping by the distinct values of r imply

$$\sum_{t=1}^{T} f(r_t)(y_t - r_t) = 0.$$

Consequently,

$$\sum_{t=1}^{T} f(p_t)(y_t - p_t) = \sum_{t=1}^{T} \big(f(p_t) - f(r_t)\big)(y_t - p_t) + \sum_{t=1}^{T} f(r_t)(r_t - p_t).$$

Because $|y_t - p_t| \leq 1$, the absolute value of the right-hand side is at most

$$(L + B)\sum_{t=1}^{T} |p_t - r_t|\,.$$

Taking the infimum over $r \in \mathcal{C}(y)$ proves (4.1). The choice $f \equiv 1$ has $(L, B) = (0, 1)$ and gives (4.2); the choice $f(u) = 1/2 - u$ has $(L, B) = (1, 1/2)$ and gives (4.3). ☐

For comparison with prior work, let

$$\mathcal{F}_1 := \{f : [0,1] \to [-1,1] : \mathrm{Lip}(f) \leq 1\}$$

and define the cumulative smooth calibration error

$$\operatorname{smCE}(p, y) := \sup_{f \in \mathcal{F}_1} \sum_{t=1}^{T} f(p_t)(y_t - p_t). \tag{4.4}$$

The class is closed under negation. Thus Lemma 4.1 recovers the standard comparison [13]

$$\operatorname{CalDist}(p, y) \geq \frac{1}{2} \operatorname{smCE}(p, y). \tag{4.5}$$

The witness-specific constants in (4.2) and (4.3), rather than the coarser factor in (4.5), materially improve the finite-horizon lower bounds below.

5 The square-root lower bounds

We first prove the randomized-adaptive statement by analyzing an auxiliary randomized adversary and then fixing its tape. Specializing the construction gives the companion fixed-string result for deterministic forecasters.

5.1 Randomized forecasters against adaptive adversaries

The next theorem supplies the lower half of (2.7). Its adversary uses the public history and the known kernel, but never the current realized forecast.

Theorem 5.1 (Randomized simultaneous lower bound)**.** *For every integer $T \geq 2$ and every Borel behavioral-kernel forecaster K defined in Section 2, there is a Borel-measurable deterministic past-adaptive adversary A that does not observe the current forecast and satisfies*

$$\mathbb{E}^{K,A} \operatorname{CalDist}(P, Y) \geq \frac{1}{4320} \sqrt{T}. \tag{5.1}$$

Proof. The proof first obtains the bound under an auxiliary randomized adversary, using an energy split and a martingale residual, and then fixes the auxiliary tape. Fix K. For every $h \in \mathsf{H}_{t-1}$, define the history functions

$$\begin{aligned} \bar{\mu}_t(h) &:= \int_{[0,1]} p \, K_{t,T}(dp \mid h), \\ \bar{v}_t(h) &:= \int_{[0,1]} (p - \bar{\mu}_t(h))^2 K_{t,T}(dp \mid h). \end{aligned} \tag{5.2}$$

Set $\operatorname{sgn}(0) := 0$ and

$$\begin{aligned} a &:= \left\lceil \frac{1}{2} \log_4 T \right\rceil, & \bar{\epsilon}_t(h) &:= \bar{\mu}_t(h) - \frac{1}{2}, \\ \bar{d}_t(h) &:= \operatorname{sgn}(\bar{\epsilon}_t(h)) \left| \bar{\epsilon}_t(h) \right|^a, & \bar{q}_t(h) &:= \bar{\mu}_t(h) - \bar{d}_t(h). \end{aligned} \tag{5.3}$$

The logarithmic choice of a balances the mean-reversion drift and the martingale fluctuation at the square-root scale. Since $a \geq 1$, $\bar{q}_t(h)$ lies between $\bar{\mu}_t(h)$ and $1/2$; all these history functions are Borel measurable.

For the analysis, use the joint transition kernel

$$K_{t,T}(dp \mid h) \ \operatorname{Bernoulli}(\bar{q}_t(h))(dy). \tag{5.4}$$

Write $\mathbb{P}^\star$ and $\mathbb{E}^\star$ for the induced law and expectation. Under this law, P_t and Y_t are conditionally independent given the pre-round history. Define the history random variable, its filtration, and the realized quantities by

$$H_{t-1} := ((P_s, Y_s))_{s<t}, \quad \mathcal{H}_{t-1} := \sigma(H_{t-1}),$$

$$\mu_t := \bar{\mu}_t(H_{t-1}), \quad v_t := \bar{v}_t(H_{t-1}), \quad \epsilon_t := \bar{\epsilon}_t(H_{t-1}), \quad d_t := \bar{d}_t(H_{t-1}), \quad q_t := \bar{q}_t(H_{t-1}).$$

Thus $\mu_t = \mathbb{E}^\star[P_t \mid \mathcal{H}_{t-1}]$ and $v_t = \mathrm{Var}^\star(P_t \mid \mathcal{H}_{t-1})$. Define the sum of predictable energy increments by

$$G := \sum_{t=1}^{T} \left(|\epsilon_t|^{a+1} + v_t \right). \tag{5.5}$$

Energy identity. For $f_0(u) = 1/2 - u$, conditional independence and (5.3) give

$$\begin{aligned} \mathbb{E}^\star[f_0(P_t)(Y_t - P_t) \mid \mathcal{H}_{t-1}] &= \mathbb{E}^\star[(1/2 - P_t)(q_t - P_t) \mid \mathcal{H}_{t-1}] \\ &= (1/2 - \mu_t)(q_t - \mu_t) + v_t \\ &= |\epsilon_t|^{a+1} + v_t. \end{aligned} \tag{5.6}$$

Applying (4.3), Jensen's inequality, and the tower property yields

$$\mathbb{E}^\star \operatorname{CalDist}(P, Y) \geq \frac{2}{3} \mathbb{E}^\star \left| \sum_{t=1}^{T} f_0(P_t)(Y_t - P_t) \right| \geq \frac{2}{3} \mathbb{E}^\star G. \tag{5.7}$$

Put

$$R := T^{(a-1)/(2a)} = \sqrt{T}\, T^{-1/(2a)}, \qquad \frac{1}{4}\sqrt{T} \leq R \leq \sqrt{T}. \tag{5.8}$$

With $c_0 := 1/720$, the case $\mathbb{E}^\star G \geq c_0 R$ follows from (5.7) and (5.8) because

$$\mathbb{E}^\star \operatorname{CalDist}(P, Y) \geq \frac{2}{3} c_0 R \geq \frac{1}{4320} \sqrt{T}. \tag{5.9}$$

It remains to assume $\mathbb{E}^\star G < c_0 R$. Call round t bad when $|\epsilon_t| > 1/4$, and let N_{bad} be the number of such rounds. Since $N_{\mathrm{bad}} \leq 4^{a+1} \sum_t |\epsilon_t|^{a+1}$ and $4^{a+1} \leq 16\sqrt{T}$, (5.8) gives

$$\mathbb{E}^\star N_{\mathrm{bad}} < T/45. \tag{5.10}$$

On every other round, $q_t(1 - q_t) \geq 3/16$.

Define

$$\zeta_t := (Y_t - q_t) - (P_t - \mu_t), \qquad Z := \sum_{t=1}^{T} \zeta_t, \qquad D := \sum_{t=1}^{T} d_t. \tag{5.11}$$

Conditional independence gives $\mathbb{E}^\star[\zeta_t \mid \mathcal{H}_{t-1}] = 0$, so the ζ_t are martingale differences. It also gives

$$\mathbb{E}^\star[\zeta_t^2 \mid \mathcal{H}_{t-1}] = q_t(1 - q_t) + v_t. \tag{5.12}$$

Martingale orthogonality and (5.10) imply

$$\mathbb{E}^\star Z^2 \geq \frac{3}{16} (T - \mathbb{E}^\star N_{\mathrm{bad}}) \geq \frac{11}{60} T. \tag{5.13}$$

Each of the two centered summands in ζ_t has conditional support in an interval of length one. Conditional Hoeffding, tail integration, and moment interpolation therefore give

$$\mathbb{E}^{\star}|Z| \geq \frac{1}{2}\left(\frac{11}{60}\right)^{3/2}\sqrt{T}. \tag{5.14}$$

Hölder and Jensen, together with the exponent identity $1/(a+1)+(a/(a+1))(a-1)/(2a) = 1/2$, give

$$\mathbb{E}^{\star}|D| \leq \frac{1}{\sqrt{720}}\sqrt{T}. \tag{5.15}$$

The complete moment and drift calculations appear in Appendix A.1.

Finally, $q_t - \mu_t = -d_t$, so

$$\sum_{t=1}^{T}(Y_t - P_t) = Z - D. \tag{5.16}$$

By (4.2), (5.14), and (5.15),

$$\begin{aligned}\mathbb{E}^{\star}\operatorname{CalDist}(P,Y) &\geq \mathbb{E}^{\star}|Z-D| \\ &\geq \left[\frac{1}{2}\left(\frac{11}{60}\right)^{3/2} - \frac{1}{\sqrt{720}}\right]\sqrt{T} > \frac{1}{4320}\sqrt{T}.\end{aligned} \tag{5.17}$$

An exact check of the final strict inequality is given in Appendix A.2.

Removing adversary randomization. Realize (5.4) with independent uniforms $W_1, \ldots, W_T$ by setting $Y_t = \mathbf{1}\{W_t \leq \bar{q}_t(H_{t-1})\}$. For every fixed tape $w \in [0,1]^T$, the maps

$$A^{w}_{t,T}(h) := \mathbf{1}\{w_t \leq \bar{q}_t(h)\}$$

form a Borel-measurable deterministic past-adaptive adversary and do not use the current forecast. By Fubini's theorem, the average over w of $\mathbb{E}^{K,A^w}\operatorname{CalDist}(P,Y)$ is the auxiliary expected loss just bounded. Hence some fixed w satisfies (5.1). This is an existence argument and does not give an efficient procedure for finding the good tape. □

The construction respects simultaneous play because it uses the known conditional law of P_t, never its current realization. The variance term in (5.6) is what permits this distinction.

5.2 A companion deterministic fixed-string bound

When the forecaster is deterministic, forecast innovation vanishes. The same construction can then use a larger energy threshold and yields a better constant, although the two theorems concern different strategy classes.

Theorem 5.2 (Deterministic oblivious lower bound)**.** *For every integer $T \geq 2$ and deterministic nonanticipating forecaster F, there is a fixed string $y \in \{0,1\}^T$ such that*

$$\operatorname{CalDist}\big(p^F(y), y\big) \geq \frac{1}{1200}\sqrt{T}. \tag{5.18}$$

Proof. We specialize the preceding energy split and then use the probabilistic method to fix one outcome string. Run the auxiliary construction from the proof of Theorem 5.1, retaining the definitions of $a, R, G, D, N_{\text{bad}}, \epsilon_t, d_t$, and q_t. Under its auxiliary law, $P_t = F_{t,T}(Y_{<t}) = \mu_t$ almost surely and $v_t = 0$. Use the larger threshold $c_0 = 1/200$. If $\mathbb{E}^\star G \geq c_0 R$, then (5.7) and (5.8) give

$$\mathbb{E}^\star \operatorname{CalDist}\big(p^F(Y), Y\big) \geq \frac{2}{3} c_0 R \geq \frac{1}{1200}\sqrt{T}. \tag{5.19}$$

Suppose instead that $\mathbb{E}^\star G < c_0 R$, and put $Z := \sum_t (Y_t - q_t)$. The bad-round calculation now gives

$$\mathbb{E}^\star N_{\text{bad}} < \frac{2T}{25}, \quad \mathbb{E}^\star Z^2 \geq \frac{69}{400} T, \quad \mathbb{E}^\star Z^4 \leq T^2. \tag{5.20}$$

Thus interpolation gives $\mathbb{E}^\star |Z| \geq (69/400)^{3/2}\sqrt{T}$, while the drift calculation gives $\mathbb{E}^\star |D| \leq \sqrt{T}/\sqrt{200}$. Therefore

$$\mathbb{E}^\star \operatorname{CalDist}\big(p^F(Y), Y\big) \geq \left[\left(\frac{69}{400}\right)^{3/2} - \frac{1}{\sqrt{200}}\right]\sqrt{T} > \frac{1}{1200}\sqrt{T}. \tag{5.21}$$

The calculations and exact strict inequality are verified in Appendix A.3.

For fixed F, the auxiliary recursion defines a probability law on the finite set $\{0,1\}^T$. In either energy branch,

$$\max_{y \in \operatorname{supp}(Y)} \operatorname{CalDist}\big(p^F(y), y\big) \geq \mathbb{E}^\star \operatorname{CalDist}\big(p^F(Y), Y\big).$$

Thus one fixed string satisfies (5.18). This string is chosen after F but before play, exactly as required by (2.5). □

For $T = 1$, perfect calibration forces the comparison sequence to equal the single outcome. A deterministic forecast p therefore has worst-case distance $\max\{p, 1-p\}$. For a randomized forecast P, the two deterministic adversary choices give worst-case expected distance $\max\{\mathbb{E}P, 1 - \mathbb{E}P\} \geq 1/2$. The constant forecast $P = 1/2$ attains equality in both models, proving the $T = 1$ assertion in Theorem 2.1.

6 An exactly optimized deterministic upper bound

We now give the common upper bound in Theorem 2.1. The construction is a shifted-bias refinement of the grid-proxy method of Arunachaleswaran et al. [1]. The proxy $\widetilde{p}_t$ is assigned only after observing y_t. It is internal state, not the forecast announced on round t.

Fix $m \in \mathbb{Z}_{\geq 1}$ and let $x_j = j/m$ for $j = 0, \ldots, m$. The algorithm maintains post-outcome proxies $\widetilde{p}_s$; initially all proxy biases are zero. After rounds $1, \ldots, t-1$, define

$$b_{t-1,j} := \sum_{s<t:\, \widetilde{p}_s = x_j} (x_j - y_s), \qquad \beta_{t-1,j} := b_{t-1,j} + x_j - \frac{1}{2}. \tag{6.1}$$

The algorithm chooses the first nonnegative shifted bias, announces the midpoint of the adjacent grid points, and then assigns the proxy according to the outcome:

$$j_t := \min\{j \in \{0, \ldots, m\} : \beta_{t-1,j} \geq 0\}, \qquad i_t := j_t - 1, \tag{6.2}$$

$$p_t := \frac{x_{i_t} + x_{j_t}}{2}, \tag{6.3}$$

$$\widetilde{p}_t := \begin{cases} x_{i_t}, & y_t = 0, \\ x_{j_t}, & y_t = 1. \end{cases} \tag{6.4}$$

The unshifted endpoint biases remain zero: coordinate 0 is updated only after outcome 0, and coordinate m only after outcome 1. Hence $\beta_{t-1,0} = -1/2$, $\beta_{t-1,m} = 1/2$, so $1 \le j_t \le m$ and

$$\beta_{t-1,i_t} < 0 \le \beta_{t-1,j_t}. \tag{6.5}$$

No monotonicity of the bias vector is assumed. Finally, all state before round t is a deterministic function of $y_{<t}$, so the announced forecast is nonanticipating.

The following invariant bounds every proxy-bin residual and will therefore control the proxy ECE.

Lemma 6.1 (Shifted-bias invariant)**.** *Under the construction* (6.1)–(6.4)*, for every* $t \in \{0, \dots, T\}$ *and* $j \in \{0, \dots, m\}$*,*

$$x_j - 1 \le \beta_{t,j} \le x_j, \qquad \textit{equivalently} \qquad -\frac{1}{2} \le b_{t,j} \le \frac{1}{2}. \tag{6.6}$$

Proof. Initially $\beta_{0,j} = x_j - 1/2 \in [x_j - 1, x_j]$. Suppose the invariant holds before round t. If $y_t = 0$, only coordinate i_t changes, and its shifted bias increases by x_{i_t}. By (6.5), its new value lies in

$$[2x_{i_t} - 1, x_{i_t}] \subseteq [x_{i_t} - 1, x_{i_t}].$$

If $y_t = 1$, only coordinate j_t changes, and its shifted bias increases by $x_{j_t} - 1$. Its new value lies in

$$[x_{j_t} - 1, 2x_{j_t} - 1] \subseteq [x_{j_t} - 1, x_{j_t}].$$

All other coordinates are unchanged, completing the induction. □

We next transfer the proxy's ECE control to the forecasts announced online. The first inequality is the Lipschitz stability of distance to a fixed set; the second supplies the required calibrated comparator.

Lemma 6.2 (Proxy comparison)**.** *For every integer* $T \ge 1$*,* $u, p \in [0,1]^T$*, and* $y \in \{0,1\}^T$*,*

$$|\mathrm{CalDist}(p, y) - \mathrm{CalDist}(u, y)| \le \sum_{t=1}^{T} |p_t - u_t|\,, \tag{6.7}$$

$$\mathrm{CalDist}(p, y) \le \sum_{t=1}^{T} |p_t - u_t| + \mathrm{ECE}(u, y). \tag{6.8}$$

Proof. Distance to a fixed nonempty subset of a normed space is 1-Lipschitz. Since $\mathrm{CalDist}(p, y)$ is the ℓ_1-distance from p to $\mathcal{C}(y)$, the triangle inequality gives (6.7).

For the second claim, for every $z \in \mathrm{Im}(u)$ define

$$\bar{y}_z := \frac{1}{|\{t : u_t = z\}|} \sum_{t : u_t = z} y_t, \qquad r_t := \bar{y}_{u_t}.$$

Within every original u-bin, the residuals $y_t - r_t$ sum to zero. If two or more such bins receive the same value of r_t, their zero residual sums simply merge, so $r \in \mathcal{C}(y)$. Moreover,

$$\begin{aligned} \sum_{t=1}^{T} |u_t - r_t| &= \sum_{z \in \mathrm{Im}(u)} |\{t : u_t = z\}|\, |z - \bar{y}_z| \\ &= \sum_{z \in \mathrm{Im}(u)} \left| \sum_{t : u_t = z} (z - y_t) \right| = \mathrm{ECE}(u, y). \end{aligned}$$

Thus $\mathrm{CalDist}(u, y) \le \mathrm{ECE}(u, y)$; combining this inequality with (6.7) proves (6.8). □

Combining the invariant with the proxy comparison yields the fixed-grid guarantee used by both upper bounds.

Proposition 6.3 (Finite-horizon pathwise bound)**.** *For every integer $T \geq 1$ and $m \in \mathbb{Z}_{\geq 1}$, the construction* (6.1)–(6.4) *satisfies, on every outcome sequence,*

$$\mathrm{CalDist}(p, y) \leq g_T(m) = \frac{T}{2m} + \frac{m-1}{2}. \tag{6.9}$$

Proof. The invariant gives $\mathrm{ECE}(\widetilde{p}, y) = \sum_{j=0}^{m} |b_{T,j}| \leq (m-1)/2$, since both unshifted endpoint biases vanish. Also $\sum_t |p_t - \widetilde{p}_t| = T/(2m)$, so (6.8) proves the claim. □

Optimizing the grid over positive integers gives the advertised leading constant without changing the pathwise guarantee.

Corollary 6.4 (Exact grid optimization)**.** *For every integer $T \geq 1$, the integer m_T^* in* (2.9) *minimizes g_T. Let $\bar{m} := \lceil \sqrt{T} \rceil$ and $\delta := \bar{m} - \sqrt{T} \in [0, 1)$. Then*

$$U_T \leq g_T(\bar{m}) = \sqrt{T} - \frac{1}{2} + \frac{\delta^2}{2\bar{m}} < \sqrt{T}. \tag{6.10}$$

Proof. For $m \geq 1$,

$$g_T(m+1) - g_T(m) = \frac{m(m+1) - T}{2m(m+1)}.$$

Hence g_T decreases until the first positive integer m with $m(m+1) \geq T$ and increases thereafter. This proves the formula for m_T^* in (2.9). If $m_T^*(m_T^* + 1) = T$, then m_T^* and $m_T^* + 1$ tie; otherwise m_T^* is the unique minimizer.

Let $s := \sqrt{T}$, $\bar{m} := \lceil s \rceil$, and $\delta := \bar{m} - s$. Direct substitution gives

$$g_T(\bar{m}) = s - \frac{1}{2} + \frac{\delta^2}{2\bar{m}}.$$

Since $0 \leq \delta < 1$ and $\bar{m} \geq 1$,

$$s - g_T(\bar{m}) = \frac{\bar{m} - \delta^2}{2\bar{m}} > 0,$$

including when $T = 1$, where $\delta = 0$. □

The finite-horizon construction also admits an exact implementation with worst-case, rather than amortized, per-round guarantees.

Proposition 6.5 (Efficient exact implementation)**.** *For a fixed $m \geq 1$, the algorithm in* (6.1)–(6.4) *admits $O(m)$ initialization time and memory, followed by $O(\log(m+1))$ worst-case time per round. Its forecasts are represented exactly as rationals. All stored integers have $O(\log(m+1))$ bits, so these bounds hold in a word-RAM with word size $\Omega(\log(m+1))$.*

The data-structure and exact-arithmetic details are given in Appendix B.1.

The upper inequalities in Theorem 2.1 follow from Proposition 6.3 and Corollary 6.4. With $m = m_T^* = \Theta(\sqrt{T})$, Proposition 6.5 uses $O(\log(T+1))$ worst-case time per round and $O(\sqrt{T})$ memory.

7 Removing knowledge of the horizon

We combine the required one-sided block property with dyadic restarts, including a truncated final epoch. For an index set I, write p_I and y_I for the corresponding restrictions.

Lemma 7.1 (Block subadditivity)**.** *For every integer $T \geq 1$, $p \in [0,1]^T$, and $y \in \{0,1\}^T$, let $I_1, \dots, I_K$ be a partition of $\{1, \dots, T\}$ into nonempty consecutive blocks. Then*

$$\mathrm{CalDist}(p, y) \leq \sum_{k=1}^{K} \mathrm{CalDist}(p_{I_k}, y_{I_k}). \tag{7.1}$$

Proof. Concatenate blockwise minimizers; their calibration residuals sum to zero and their ℓ_1 costs add. □

Proof of Corollary 2.2. Use dyadic epochs $I_k := \{2^k, \dots, 2^{k+1} - 1\}$ of length $L_k := 2^k$, resetting the finite-horizon algorithm with $m = m^*_{L_k}$. For a prefix T, let $k_T := \lfloor \log_2 T \rfloor$ and let the last block be $J_{k_T} := \{2^{k_T}, \dots, T\}$, of length $n \leq L_{k_T}$. Applying the fixed-grid proof to this truncated block and Lemma 7.1 to the full prefix gives

$$\begin{aligned} \mathrm{CalDist}(p_{1:T}, y_{1:T}) &\leq \sum_{k<k_T} \sqrt{L_k} + \frac{n}{2m^*_{L_{k_T}}} + \frac{m^*_{L_{k_T}} - 1}{2} \\ &\leq \sum_{k=0}^{k_T} \sqrt{2^k} \\ &= (2+\sqrt{2})2^{k_T/2} - (1+\sqrt{2}) < (2+\sqrt{2})\sqrt{T}. \end{aligned} \tag{7.2}$$

For the runtime bound, during an epoch of length L the algorithm schedules the $O(m^*_{2L}) = O(\sqrt{L})$ initialization work for the next epoch across the current L rounds. At the boundary it swaps the prepared roots in $O(1)$ time and retires the old epoch arena without traversing it. The arena sizes form a geometric sum of order $O(\sqrt{T})$, while current-epoch operations have the worst-case cost in Proposition 6.5. The node-reuse and retirement details are in Appendix B.2. □

8 Discussion and open directions

Two limitations remain after closing the exponent gap. The leading minimax constants are open, and the randomized-oblivious value in Remark 2.3 is unresolved because our averaging argument does not produce one seed-independent hard string. The result is otherwise scoped to binary, context-free forecasting and the behavioral-kernel class of Section 2; its randomized lower bound is in expectation and uses a kernel-dependent adversary. It concerns continuous distance to an exactly calibrated sequence in hindsight, not discontinuous ECE, and therefore does not resolve the separate ECE exponent problem.

References

[1] E. R. Arunachaleswaran, N. Collina, A. Roth, and M. Shi. An elementary predictor obtaining $2\sqrt{T}+1$ distance to calibration. In *Proceedings of the 2025 Annual ACM–SIAM Symposium on Discrete Algorithms (SODA)*, pages 1366–1370, 2025. doi:10.1137/1.9781611978322.41. Extended version: arXiv:2402.11410.

[2] J. Błasiok, P. Gopalan, L. Hu, and P. Nakkiran. A unifying theory of distance from calibration. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC)*, pages 1727–1740, 2023. doi:10.1145/3564246.3585182. Extended version: arXiv:2211.16886.

[3] N. Collina, S. Goel, V. Gupta, and A. Roth. Tractable agreement protocols. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing (STOC)*, pages 1532–1543, 2025. doi:10.1145/3717823.3718222. Extended version: arXiv:2411.19791.

[4] Y. Dagan, C. Daskalakis, M. Fishelson, N. Golowich, R. Kleinberg, and P. Okoroafor. Breaking the $T^{2/3}$ barrier for sequential calibration. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing (STOC)*, pages 2007–2018, 2025. doi:10.1145/3717823.3718178.

[5] A. P. Dawid. The well-calibrated Bayesian. *Journal of the American Statistical Association*, 77(379):605–610, 1982. doi:10.1080/01621459.1982.10477856.

[6] D. P. Foster. A proof of calibration via Blackwell's approachability theorem. *Games and Economic Behavior*, 29(1–2):73–78, 1999. doi:10.1006/game.1999.0719.

[7] D. P. Foster and S. Hart. Smooth calibration, leaky forecasts, finite recall, and Nash dynamics. *Games and Economic Behavior*, 109:271–293, 2018. doi:10.1016/j.geb.2017.12.022.

[8] D. P. Foster and R. V. Vohra. Asymptotic calibration. *Biometrika*, 85(2):379–390, 1998. doi:10.1093/biomet/85.2.379.

[9] S. M. Kakade and D. 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.

[10] D. Oakes. Self-calibrating priors do not exist. *Journal of the American Statistical Association*, 80(390):339, 1985. doi:10.1080/01621459.1985.10478117.

[11] M. Qiao. Computational and statistical hardness of calibration distance. arXiv:2603.18391, 2026.

[12] M. Qiao and G. Valiant. Stronger calibration lower bounds via sidestepping. In *Proceedings of the 53rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 456–466, 2021. doi:10.1145/3406325.3451050.

[13] M. Qiao and L. Zheng. On the distance from calibration in sequential prediction. In *Proceedings of the 37th Conference on Learning Theory (COLT)*, volume 247 of *Proceedings of Machine Learning Research*, pages 4307–4357, 2024. Proceedings version; arXiv:2402.07458.

[14] Z. Zhang. Efficient sequential calibration with $O(T^{2/3-\varepsilon})$ error bound. arXiv:2607.12928, 2026.

A Deferred lower-bound details

This appendix supplies the deferred moment calculations and elementary exact comparisons used for the lower-bound constants.

A.1 Moment and drift calculations for the randomized lower bound

We work in the small-energy branch $\mathbb{E}^\star G < c_0 R$ with $c_0 = 1/720$. The definition of a bad round gives the pathwise inequality

$$N_{\mathrm{bad}} \le 4^{a+1} \sum_{t=1}^{T} |\epsilon_t|^{a+1}\,.$$

The choice of a implies $4^{a+1} \le 16\sqrt{T}$. Therefore, using $R \le \sqrt{T}$,

$$\mathbb{E}^\star N_{\mathrm{bad}} < 16\sqrt{T}\, c_0 R \le 16 c_0 T = \frac{T}{45}. \tag{A.1}$$

On a nonbad round, $\mu_t \in [1/4, 3/4]$; since q_t lies between μ_t and $1/2$, this gives $q_t(1-q_t) \ge 3/16$. Martingale orthogonality and (5.12) now yield

$$\begin{aligned}\mathbb{E}^\star Z^2 &= \sum_{t=1}^{T} \mathbb{E}^\star[q_t(1-q_t) + v_t] \\ &\ge \frac{3}{16}\,(T - \mathbb{E}^\star N_{\mathrm{bad}}) \ge \frac{11}{60} T. \end{aligned}\tag{A.2}$$

Conditionally on $\mathcal{H}_{t-1}$, each of $Y_t - q_t$ and $-(P_t - \mu_t)$ is centered and supported on an interval of length one, and the two variables are independent. Conditional Hoeffding's lemma therefore gives, for every $\lambda \in \mathbb{R}$,

$$\mathbb{E}^\star\left[e^{\lambda \zeta_t} \mid \mathcal{H}_{t-1}\right] \le e^{\lambda^2/8} e^{\lambda^2/8} = e^{\lambda^2/4}.$$

Iteration and the exponential Markov inequality imply

$$\mathbb{P}^\star(|Z| \ge u) \le 2e^{-u^2/T} \qquad (u \ge 0). \tag{A.3}$$

Tail integration then gives

$$\mathbb{E}^\star Z^4 = 4\int_0^\infty u^3 \mathbb{P}^\star(|Z| \ge u)\, du \le 8 \int_0^\infty u^3 e^{-u^2/T}\, du = 4T^2. \tag{A.4}$$

The interpolation inequality

$$\mathbb{E}^\star Z^2 \le \left(\mathbb{E}^\star |Z|\right)^{2/3} \left(\mathbb{E}^\star Z^4\right)^{1/3}$$

combined with (A.2) and (A.4) proves

$$\mathbb{E}^\star |Z| \ge \frac{1}{2}\left(\frac{11}{60}\right)^{3/2} \sqrt{T},$$

which is (5.14).

For the drift, Hölder's inequality gives pathwise

$$|D| \le \sum_{t=1}^{T} |\epsilon_t|^a \le T^{1/(a+1)} \left(\sum_{t=1}^{T} |\epsilon_t|^{a+1}\right)^{a/(a+1)}.$$

Jensen's inequality and $\sum_t |\epsilon_t|^{a+1} \le G$ hence give

$$\begin{aligned}\mathbb{E}^\star |D| &\le T^{1/(a+1)} (\mathbb{E}^\star G)^{a/(a+1)} \\ &< c_0^{a/(a+1)} T^{1/(a+1)} R^{a/(a+1)} \le \sqrt{c_0 T}. \end{aligned}\tag{A.5}$$

Here $a/(a+1) \ge 1/2$, $c_0 < 1$, and

$$\frac{1}{a+1} + \frac{a}{a+1}\frac{a-1}{2a} = \frac{1}{2}.$$

This proves (5.15).

A.2 Exact randomized constant

The bracket in (5.17) equals

$$\frac{11\sqrt{165}-60\sqrt{5}}{3600}.$$

Since $\sqrt{165} > 64/5$ and $\sqrt{5} < 9/4$, it is strictly larger than $29/18000$, which in turn is larger than $1/4320$.

A.3 Details for the deterministic companion bound

In the small-energy branch with $c_0 = 1/200$, the calculation in (A.1) instead gives

$$\mathbb{E}^\star N_{\mathrm{bad}} < 16c_0T = \frac{2T}{25}. \tag{A.6}$$

Here $P_t = \mu_t$ almost surely, so $Z = \sum_t (Y_t - q_t)$ and the conditional variance of its tth increment is $q_t(1-q_t)$. Thus

$$\mathbb{E}^\star Z^2 \geq \frac{3}{16}\left(1-\frac{2}{25}\right)T = \frac{69}{400}T. \tag{A.7}$$

Each increment is centered and has conditional support in an interval of length one. Conditional Hoeffding, iteration, and tail integration give

$$\mathbb{P}^\star(|Z| \geq u) \leq 2e^{-2u^2/T}, \qquad \mathbb{E}^\star Z^4 \leq 8\int_0^\infty u^3 e^{-2u^2/T}\,du = T^2.$$

The same interpolation inequality therefore yields

$$\mathbb{E}^\star|Z| \geq \left(\frac{69}{400}\right)^{3/2}\sqrt{T}. \tag{A.8}$$

The drift proof in Appendix A.1, now with $c_0 = 1/200$, gives $\mathbb{E}^\star|D| \leq \sqrt{T/200}$. Finally,

$$\left(\frac{69}{400}\right)^{3/2} - \frac{1}{\sqrt{200}} = \frac{207\sqrt{69}-1200\sqrt{2}}{24000}.$$

The elementary bounds $\sqrt{69} > 83/10$ and $\sqrt{2} < 99/70$ make the numerator strictly larger than $207(83/10) - 1200(99/70) = 1467/70 > 20$. Hence the displayed constant is strictly larger than $20/24000 = 1/1200$, as used in (5.21).

B Implementation details for the upper bounds

The proofs below supply the data-structure details deferred from the main text.

B.1 Proof of the exact finite-horizon implementation

Proof of Proposition 6.5. Maintain

$$S_t := \{j \in \{0,\dots,m\} : \beta_{t,j} \geq 0\}$$

in a worst-case balanced search tree and maintain the nonzero unshifted biases $b_{t,j}$ in a worst-case balanced dictionary. Initially the latter is empty and

$$S_0 = \{\lceil m/2 \rceil, \dots, m\},$$

so the tree can be bulk-built from sorted keys in $O(m)$ time. On round t, j_t is the minimum key in S_{t-1}; the endpoint invariant ensures that this set is nonempty. Only the bias at i_t or j_t changes, so one dictionary entry is updated and at most one key is inserted into or deleted from S_{t-1}. Both structures draw node records from fixed $O(m)$ pools and return deleted records to free lists, so repeated zero crossings do not grow the storage. These operations take $O(\log(m+1))$ worst-case time.

For exact arithmetic, store the integer $2mb_{t,j}$. The corresponding shifted-bias numerator is

$$2m\beta_{t,j} = 2mb_{t,j} + 2j - m.$$

An outcome-0 update at coordinate i_t adds $2i_t$, whereas an outcome-1 update at coordinate j_t adds $2(j_t - m)$. The invariant bounds the stored and shifted-bias numerators in magnitude by $2m$, so $O(\log(m+1))$ bits suffice. The announced forecast is the exact rational

$$p_t = \frac{2j_t - 1}{2m}.$$

All word operations therefore have constant cost in the stated model. □

B.2 Worst-case implementation of the anytime construction

It remains to ensure that restarting at a dyadic boundary does not create an occasional linear-time round. During an epoch of length L, prepare the initial search tree for the next epoch, whose parameter is m^*_{2L}. Its initial key set is sorted and known in advance. A linear-time bulk build can therefore be scheduled as $O(1)$ elementary node and pointer operations on each of the L current rounds. Since $m^*_{2L} = O(\sqrt{L})$, this budget suffices for every sufficiently large epoch; the finitely many initial cases are initialized directly. Allocate each epoch's search tree and sparse-bias dictionary in a single arena. At an epoch boundary, switch both current roots in $O(1)$ time (the new dictionary root is null) and retire the old arena without traversing it. The scheduled construction also initializes the two fixed-size node pools and their epoch-local free lists, so insertions reuse records rather than accumulating allocations. Retired arenas need not be reclaimed during play: up to any prefix T, their total size is

$$\sum_{k \le \lfloor \log_2 T \rfloor} O(\sqrt{2^k}) = O(\sqrt{T}).$$

The current sparse-bias dictionary is maintained by the balanced-tree method of Proposition 6.5. Thus each round performs $O(\log(T+1))$ worst-case current-epoch work plus $O(1)$ scheduled build work, while all current, future, and retired arenas together use $O(\sqrt{T})$ memory. This establishes the computational claim in Corollary 2.2.