DIRICHLET FOLLOW-THE-LEADER CLOSES THE GAP IN SIMULTANEOUS MULTICLASS U-CALIBRATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Can a single proper multiclass forecaster attain the minimax-optimal worst-case pseudo-U-calibration rate over all bounded proper losses? Can it do so while retaining the optimal logarithmic dependence on T for every bounded Euclidean-smooth proper loss? Recent work gave the first simultaneous guarantee, but its bounds are roughly $K^{5/4}\sqrt{T}$ in the worst case and include an additional $\beta\sqrt{K}\log K$ term for smooth losses. We show that both dimension losses are unnecessary. A one-line, horizon-free Dirichlet forecaster achieves the optimal rates. After observing class counts c_{t-1}, draw the next prediction from $\mathrm{Dir}(c_{t-1})$, on the face of classes seen so far. This is a fresh Bayesian bootstrap of the outcomes. The analysis rests on an exact identity: averaging any bounded proper loss under $\mathrm{Dir}(\alpha)$ equals a discrete derivative of its Dirichlet-averaged Bayes risk. The identity makes the be-the-perturbed-leader term telescope to a nonpositive Jensen gap. A one-count likelihood ratio then bounds stability by the inverse square root of that class's count. The resulting single, horizon-free algorithm satisfies

$$\sup_{\ell} \mathrm{Reg}_\ell \le 2S_T + 2\sum_{i:N_i>0}\sqrt{N_i} \le 4\sqrt{S_T T} \le 4\sqrt{\min\{K,T\}T},$$
$$\mathrm{Reg}_\ell \le \tfrac{5}{2}\beta(1+\log T) \quad \text{for every } \beta\text{-smooth proper } \ell.$$

Here N_i are terminal class counts and S_T is the number of observed classes. A class-aggregation reduction shows that the dimension-capped rate $\sqrt{\min\{K,T\}T}$ is minimax-optimal up to constants in every regime with $T \ge 24$, and known smooth-loss lower bounds match the logarithmic dependence on T of the second bound. The proof covers nondifferentiable losses and changes of the active simplex face.

1 INTRODUCTION

A probabilistic forecast is often consumed by an agent whose downstream utility is unknown to the forecaster. Proper losses capture this uncertainty. Each proper loss corresponds to a way of valuing probabilistic predictions, and truthful prediction minimizes expected loss (Gneiting & Raftery, 2007). U-calibration asks for one online sequence of forecasts with low regret under every bounded proper loss, and optimizing a single score such as the Brier score does not guarantee it (Kleinberg et al., 2023).

For K outcomes and $T \ge 12K$, Luo et al. (2024) established the optimal $\Theta(\sqrt{KT})$ worst-case pseudo-U-calibration rate. Their forecaster can nevertheless incur $\Omega(\sqrt{T})$ regret on squared loss, where Follow-the-Leader (FTL) has logarithmic regret (Frongillo et al., 2026). Frongillo et al. (2026) then gave a single forecaster with sublinear regret for every bounded proper loss and logarithmic regret for smooth losses. Their proper multiclass construction has worst-case rate $\widetilde{O}(K^{5/4}\sqrt{T})$ and an additional $O(\beta\sqrt{K}\log K)$ term in its smooth-loss guarantee. Their discussion explicitly identifies optimal dependence on K for general bounded losses as unresolved. We settle that question while retaining optimal logarithmic dependence on T for every Euclidean-smooth proper loss. Up to logarithmic factors, the leading worst-case dimension dependence improves by $K^{3/4}$, and the smooth guarantee loses its entire dimension-dependent additive term.

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.

We achieve both guarantees with a Dirichlet forecaster. Let $c_t = \sum_{s \le t} y_s$ be the vector of outcome counts. At time $t \ge 2$, independently sample

$$P_t \sim \mathrm{Dir}(c_{t-1}).$$

Zero count coordinates are fixed at zero, so the distribution lives on the face spanned by classes observed so far. Equivalently, assign independent Gamma weights to the observed outcomes and normalize. Thus the algorithm is exactly a fresh Bayesian bootstrap (Rubin, 1981). It requires no horizon, smoothness parameter, learning rate, or optimization oracle. The Bayesian-bootstrap draw itself is classical. To our knowledge, its loss-uniform count-deletion identity and its use to obtain optimal adversarial multiclass regret are new. They yield the first proper multiclass forecaster known to combine the minimax worst-case rate with logarithmic regret for every bounded Euclidean-smooth proper loss.

The main proof is not a generic posterior-sampling argument. If f is the concave Bayes risk of a proper loss and $F(\alpha) = \mathbb{E}_{P \sim \mathrm{Dir}(\alpha)} f(P)$, we prove

$$\mathbb{E}_{P \sim \mathrm{Dir}(\alpha)} \ell(P, e_j) = nF(\alpha) - (n-1)F(\alpha - e_j), \qquad n = \sum_i \alpha_i. \tag{1}$$

This identity applies to arbitrary bounded proper losses, including polyhedral and V-shaped losses. It turns the perturbed-leader contribution into $T(F(c_T) - f(c_T/T)) \le 0$. The remaining stability term compares two Dirichlet laws whose parameters differ by one count. Their likelihood ratio is linear in one coordinate, which gives stability $O(1/\sqrt{m})$ when that coordinate has appeared m times. Summing separately within each class produces $O(\sum_i \sqrt{N_i}) = O(\sqrt{KT})$.

The same distribution is automatically adapted to smooth losses. Its mean is empirical FTL and its squared radius is at most $1/t$. Centering removes the first-order smoothness term, leaving $O(\beta/t)$. A short semiconcavity argument also gives an explicit $2\beta H_T$ FTL bound under the one-sided smoothness definition used in prior work.

Our contributions are:

- We give the first proper multiclass forecaster known to attain both the minimax $\Theta(\sqrt{\min\{K, T\}T})$ worst-case rate and $O(\beta \log T)$ regret for every bounded β-smooth proper loss. The algorithm is horizon-free, parameter-free, and frequency-adaptive.
- We prove the loss-uniform count-deletion identity in (1), including nondifferentiable Bayes risks, zero count coordinates, and the singular case $\alpha_j = 1$.
- We show that one-count Dirichlet total variation controls nonsmooth loss while centered Dirichlet covariance controls smooth loss and yields an exact transfer bound relative to empirical FTL.
- The same analysis gives an all-regime minimax characterization and realized high-probability guarantees for finite or uniformly coverable loss families.

The main guarantee uses the same expected-regret, or pseudo-U-calibration, criterion as the closest multiclass results (Kleinberg et al., 2023; Luo et al., 2024; Frongillo et al., 2026). Corollary 5 additionally gives realized high-probability control for finite loss families and uniformly coverable classes.

2 Setting and Algorithm

Let $[K] = \{1, \dots, K\}$, and let $e_1, \dots, e_K$ denote the vertices of Δ_K. At round t, a forecaster chooses a possibly random $P_t \in \Delta_K$, then observes an outcome $y_t \in \{e_1, \dots, e_K\}$. We first state results for an oblivious outcome sequence. Section 4.1 gives the standard extension to a nonanticipating adaptive adversary when fresh randomness is used each round.

A loss $\ell : \Delta_K \times \{e_i\}_{i=1}^K \to [-1, 1]$ is *proper* if it is Borel measurable in its first argument and

$$p \in \arg\min_{q \in \Delta_K} \mathbb{E}_{Y \sim p} \ell(q, Y) \quad \text{for every } p \in \Delta_K.$$

Algorithm 1 Dirichlet Follow-the-Leader

1: Initialize $c_0 = 0$ and predict $P_1 = (1/K, \dots, 1/K)$.
2: **for** $t = 1, 2, \dots, T$ **do**
3: **if** $t \geq 2$ **then**
4: Draw $P_t \sim \mathrm{Dir}(c_{t-1})$ independently of past draws.
5: **end if**
6: Observe y_t and set $c_t = c_{t-1} + y_t$.
7: **end for**

For an outcome sequence, propriety makes its empirical distribution $q_T = T^{-1} \sum_t y_t$ a hindsight minimizer. Define

$$\mathrm{Reg}_\ell(T) = \mathbb{E} \sum_{t=1}^{T} \ell(P_t, y_t) - T f(q_T), \qquad f(p) = \sum_i p_i \ell(p, e_i).$$

The expectation is over the forecaster. Let $\mathcal{L}$ be all proper losses with range in $[-1, 1]$ and let $\mathrm{PUCal}_T = \sup_{\ell \in \mathcal{L}} \mathrm{Reg}_\ell(T)$.

Following Frongillo et al. (2026), a differentiable loss is β-smooth when

$$\ell(q, y) - \ell(p, y) \leq \langle \nabla_p \ell(p, y), q - p \rangle + \frac{\beta}{2} \|q - p\|_2^2 \tag{2}$$

for every $p, q \in \Delta_K$ and outcome y. Gradients and differentiability on the closed simplex are understood relative to its affine hull, including one-sided differentiability on boundary faces.

For a nonzero nonnegative parameter vector α, write $A = \{i : \alpha_i > 0\}$. We define $\mathrm{Dir}(\alpha)$ on the face Δ_A by drawing independent $G_i \sim \mathrm{Gamma}(\alpha_i, 1)$ for $i \in A$, setting $P_i = G_i / \sum_{k \in A} G_k$, and setting $P_i = 0$ outside A.

The first-round prediction only handles the empty count vector and may be replaced by any fixed point of the simplex. Its loss contributes to the first-appearance term in the proof. Thereafter, every draw is supported exactly on the empirical active face. A class that has not appeared therefore receives probability zero, but it joins the face immediately after its first occurrence. This support change is why the analysis treats first appearances separately from repeated outcomes.

Theorem 1 (Simultaneously optimal regret)**.** *For every $K, T \geq 1$, let $N_i = c_{T,i}$ and let S_T be the number of distinct outcomes observed. Algorithm 1 satisfies*

$$\mathrm{PUCal}_T \leq 2S_T + 2 \sum_{i : N_i > 0} \sqrt{N_i} \leq 4\sqrt{S_T T} \leq 4\sqrt{\min\{K, T\} T}.$$

Let $\mathrm{Reg}_\ell^{\mathrm{FTL}}(T)$ denote the pathwise regret of empirical FTL, initialized with $q_0 = P_1$ and predicting $q_{t-1} = c_{t-1}/(t-1)$ for $t \geq 2$. Simultaneously, for every β-smooth $\ell \in \mathcal{L}$,

$$\mathrm{Reg}_\ell(T) \leq \mathrm{Reg}_\ell^{\mathrm{FTL}}(T) + \frac{\beta}{2}(H_T - 1) \leq \frac{5}{2}\beta H_T \leq \frac{5}{2}\beta(1 + \log T),$$

where $H_T = \sum_{t=1}^{T} 1/t$.

For $T \geq 12K$, the $\sqrt{KT}$ worst-case rate is matched by the multiclass V-shaped lower bound of Luo et al. (2024). Corollary 4 below lifts that result to every dimension regime. Squared loss gives an $\Omega(\log T)$ lower bound for the smooth subclass (Luo et al., 2024). Thus the smooth bound has the optimal dependence on T for fixed smoothness, while the support-sensitive refinement is an additional instance-dependent guarantee. Appendix B reports exact worst-case regret computations for binary sequences.

3 A Dirichlet identity for every proper loss

Write $\ell_p = (\ell(p, e_1), \dots, \ell(p, e_K))$. Propriety gives

$$f(q) \leq \langle q, \ell_p \rangle, \qquad f(p) = \langle p, \ell_p \rangle. \tag{3}$$

Thus f is concave and ℓ_p is a supergradient representation. This is the loss-sign form of the categorical Savage representation (Gneiting & Raftery, 2007). Since $\|\ell_p\|_\infty \le 1$, applying (3) in both directions also gives

$$|f(p) - f(q)| \le \|p - q\|_1. \tag{4}$$

Lemma 2 (Count deletion)**.** *Let $\alpha \ge 0$, let $n = \sum_i \alpha_i$, and suppose $\alpha_j \ge 1$. For $F(\alpha) = \mathbb{E}_{P\sim\mathrm{Dir}(\alpha)} f(P)$,*

$$\mathbb{E}_{P\sim\mathrm{Dir}(\alpha)} \ell(P, e_j) = nF(\alpha) - (n-1)F(\alpha - e_j).$$

The statement uses the active-face convention above. When $n = 1$, the formula means $\mathbb{E}_{P\sim\mathrm{Dir}(\alpha)} \ell(P, e_j) = F(\alpha)$ and does not invoke $F(0)$.

Proof. If the active face consists only of e_j, then all relevant Dirichlet laws are point masses and the result follows from $\ell(e_j, e_j) = f(e_j)$. We henceforth assume that the active face has positive dimension.

First suppose $\alpha_j > 1$ and work relative to the affine hull of the active face. A concave function is differentiable almost everywhere on its relative interior. At each such p, (3) makes the tangent component of ℓ_p the unique supergradient of f. Since $e_j - p$ is tangent to the face,

$$D_{e_j - p} f(p) = \langle e_j - p, \ell_p \rangle = \ell(p, e_j) - f(p). \tag{5}$$

This is the directional form of the usual representation of a proper loss by its Bayes risk. It also shows that the derivative in (5) is bounded even when f is nonsmooth elsewhere.

Let ρ_α be the Dirichlet density and set $v(p) = e_j - p$. A direct calculation on the active face gives

$$\operatorname{div} v + D_v \log \rho_\alpha = \frac{\alpha_j - 1}{p_j} - (n-1). \tag{6}$$

Indeed, if the active set has size r, then $\operatorname{div} v = -(r-1)$. Differentiating $\log \rho_\alpha(p) = \text{const} + \sum_{i\in A}(\alpha_i - 1)\log p_i$ along v supplies the remaining terms in (6).

The boundary flux of $f\rho_\alpha v$ vanishes. To see the relevant orders, approach a facet $p_i = 0$. If $i \ne j$, the normal component is proportional to $v_i = -p_i$, so the flux is $O(p_i^{\alpha_i})$. At the facet $p_j = 0$, the normal component is bounded and the flux is $O(p_j^{\alpha_j - 1})$, which vanishes because $\alpha_j > 1$. Boundedness of f and integrability of the remaining Dirichlet factors justify these estimates also near intersections of facets. Appendix A supplies a truncation-and-mollification argument that makes this step valid for a merely concave Lipschitz f.

Integration by parts now gives

$$\begin{aligned}\mathbb{E}_\alpha D_v f(P) &= -\int f(p)\, \operatorname{div}(\rho_\alpha(p) v(p))\, dp \\ &= (n-1)F(\alpha) - (\alpha_j - 1)\mathbb{E}_\alpha \frac{f(P)}{P_j}.\end{aligned}$$

The ratio of Dirichlet normalizers is

$$\frac{B(\alpha - e_j)}{B(\alpha)} = \frac{n-1}{\alpha_j - 1},$$

and hence

$$(\alpha_j - 1)\mathbb{E}_\alpha \frac{f(P)}{P_j} = (n-1)F(\alpha - e_j),$$

where all integrals are over the same active face. Combining these displays with (5) proves the claim for $\alpha_j > 1$.

If $\alpha_j = 1$, apply the proved identity to $\alpha + \varepsilon e_j$ and let $\varepsilon \downarrow 0$. The first Dirichlet law converges in total variation, which is enough for the possibly discontinuous bounded function $\ell(\cdot, e_j)$. After deleting one count, the law with parameter ε in coordinate j converges weakly to the lower-dimensional face. Equation (4) makes f continuous, so its expectations converge. The case $n = 1$ is immediate from $\ell(e_j, e_j) = f(e_j)$. □

4 WORST-CASE ANALYSIS

For analysis, after observing round t, draw a ghost prediction $Q_{t+1} \sim \mathrm{Dir}(c_t)$. This variable is used only to compare the count vectors before and after the current outcome. For any fixed loss, adding and subtracting its loss gives

$$\begin{aligned}\mathrm{Reg}_\ell(T) = &\sum_{t=1}^{T} \mathbb{E}[\ell(P_t, y_t) - \ell(Q_{t+1}, y_t)] \\ &+ \sum_{t=1}^{T} \mathbb{E}\ell(Q_{t+1}, y_t) - Tf(q_T). \end{aligned} \tag{7}$$

If $y_t = e_j$, Lemma 2 gives

$$\mathbb{E}\ell(Q_{t+1}, y_t) = tF(c_t) - (t-1)F(c_{t-1}).$$

When $t = 1$, the second term is omitted and $F(c_0) = F(0)$ is not invoked. The second line of (7) therefore telescopes to

$$T\{F(c_T) - f(q_T)\} \le 0, \tag{8}$$

because $\mathbb{E}Q_{T+1} = q_T$ and f is concave.

It remains to control stability. The next elementary estimate is where the coordinate adaptivity enters.

Lemma 3 (One-count stability)**.** *Let $c \ge 0$, $n = \sum_i c_i$, and $c_j = m > 0$. Then*

$$\mathrm{TV}(\mathrm{Dir}(c), \mathrm{Dir}(c + e_j)) \le \frac{1}{2}\sqrt{\frac{n-m}{m(n+1)}} \le \frac{1}{2\sqrt{m}}.$$

Proof. On their common active face, the likelihood ratio is

$$\frac{d\,\mathrm{Dir}(c + e_j)}{d\,\mathrm{Dir}(c)}(p) = \frac{np_j}{m}.$$

Consequently,

$$\mathrm{TV}(\mathrm{Dir}(c), \mathrm{Dir}(c + e_j)) = \frac{1}{2}\mathbb{E}_{P\sim\mathrm{Dir}(c)}\left|\frac{nP_j}{m} - 1\right|.$$

Cauchy–Schwarz and $\mathrm{Var}(P_j) = m(n-m)/(n^2(n+1))$ prove both inequalities. When $b = n-m > 0$, direct integration at the likelihood-ratio crossing $p_j = m/n$ also gives the exact expression

$$\mathrm{TV} = \frac{(m/n)^m(1-m/n)^b}{mB(m,b)}.$$

For $b = 0$, both laws are the same point mass and the distance is zero. □

For any function taking values in $[-1, 1]$, the difference of its expectations under two laws is at most twice their total variation. Thus, if class j has appeared $m > 0$ times before round t, Lemma 3 bounds the corresponding stability term by $2\,\mathrm{TV} \le 1/\sqrt{m}$. If it is the first appearance, the two Dirichlet laws lie on different faces and the trivial range bound is 2. Let $N_i = c_{T,i}$ and $S = |\{i : N_i > 0\}|$. Grouping rounds by the class observed, Equations (7) and (8) give

$$\begin{aligned}\mathrm{Reg}_\ell(T) &\le 2S + \sum_{i=1}^{K}\sum_{m=1}^{N_i-1}\frac{1}{\sqrt{m}} \\ &\le 2S + 2\sum_{i:N_i>0}\sqrt{N_i} \\ &\le 2S + 2\sqrt{ST} \le 4\sqrt{ST},\end{aligned}$$

where $S \le T$ implies $S \le \sqrt{ST}$. The bound is independent of ℓ, so taking the supremum proves the first part of Theorem 1.

The upper bound remains sharp when the nominal alphabet is larger than the horizon. The reduction below is useful because the cited lower bound is stated only for $T \ge 12K$.

Corollary 4 (All-regime minimax rate)**.** *For $K \geq 2$ and $T \geq 24$,*

$$\inf_{\mathcal{A}} \sup_{y_{1:T} \in [K]^T} \sup_{\ell \in \mathcal{L}} \mathrm{Reg}_\ell^{\mathcal{A}}(T) = \Theta\left(\sqrt{\min\{K, T\}T}\right),$$

where the infimum is over randomized forecasting algorithms and regret includes their internal expectation.

Proof. The upper bound is Theorem 1. If $K \leq T/12$, the lower bound of Luo et al. (2024) applies directly. Otherwise let $r = \lfloor T/12 \rfloor$, so $2 \leq r \leq K$ and $r \geq \min\{K, T\}/24$.

It remains to transfer the r-outcome lower bound to K outcomes. Choose a surjection $\pi : [K] \to [r]$ and define $(A_\pi p)_j = \sum_{i:\pi(i)=j} p_i$. Given a proper loss $\widetilde{\ell}$ on Δ_r, its lift

$$\ell(p, e_i) = \widetilde{\ell}(A_\pi p, e_{\pi(i)})$$

is proper, because under any $q \in \Delta_K$ its expected value equals the expected value of $\widetilde{\ell}$ under $A_\pi q$ and is therefore minimized at $p = q$. Restrict the adversary to one representative outcome from each fiber of π and map every forecast through A_π. Incurred losses agree, and the hindsight minima agree because A_π is onto. Thus every r-outcome lower bound transfers unchanged. Applying the cited $\Omega(\sqrt{rT})$ bound completes the proof. □

4.1 Adaptive outcome sequences

The same expected bounds hold for a nonanticipating adaptive adversary that may depend on past predictions but not on the fresh draw P_t. For the analysis, let $\mathcal{F}_{t-1}$ contain all information through round $t-1$ and expose the adversary's choice through $\mathcal{G}_t = \sigma(\mathcal{F}_{t-1}, y_t)$ before the fresh algorithmic draw. Nonanticipation makes y_t conditionally independent of P_t given $\mathcal{F}_{t-1}$, so $P_t \mid \mathcal{G}_t \sim \mathrm{Dir}(c_{t-1})$. After adjoining P_t, write $\mathcal{F}_t = \sigma(\mathcal{G}_t, P_t)$. The one-step identity and stability bounds therefore hold conditionally. Summing conditional expectations gives the same count-potential telescoping along the realized outcome path. Writing $B_T = 2S_T + 2\sum_{i:N_i>0} \sqrt{N_i}$, this yields

$$\sup_{\ell \in \mathcal{L}} \mathbb{E}\left[\sum_{t=1}^{T} \ell(P_t, y_t) - T f(q_T)\right] \leq \mathbb{E} B_T \leq 4\sqrt{\min\{K, T\}T}.$$

This is the usual fresh-randomization reduction for perturbed leader (Hutter & Poland, 2005). The argument does not cover an adversary that observes the current randomized forecast before choosing the same-round outcome.

Fresh randomness also controls the stronger order of quantifiers for finite loss families. The argument is the one Luo et al. (2024, Section 3.3) give for their perturbed leader, using the high-probability bounds for fresh perturbations of Hutter & Poland (2005, Section 9), and we record it for completeness. For a realized forecast sequence, define

$$\widehat{\mathrm{Reg}}_\ell(T) = \sum_{t=1}^{T} \ell(P_t, y_t) - T f(q_T).$$

Corollary 5 (Finite loss families and covers)**.** *Let $\mathcal{L}_0 \subseteq \mathcal{L}$ contain $M < \infty$ losses. Against an oblivious or nonanticipating adaptive adversary, with probability at least $1 - \delta$,*

$$\max_{\ell \in \mathcal{L}_0} \widehat{\mathrm{Reg}}_\ell(T) \leq B_T + \sqrt{2T \log(M/\delta)}. \tag{9}$$

For a fixed oblivious outcome sequence,

$$\mathbb{E} \max_{\ell \in \mathcal{L}_0} \widehat{\mathrm{Reg}}_\ell(T) \leq B_T + \sqrt{2T \log M}. \tag{10}$$

If a class $\mathcal{C} \subseteq \mathcal{L}$ has an internal uniform ε-cover of size $M(\varepsilon)$, meaning that every $\ell \in \mathcal{C}$ lies within ε in sup norm on $\Delta_K \times \{e_i\}_{i=1}^K$ of a cover element in $\mathcal{C}$, then replacing the maxima above by $\sup_{\ell \in \mathcal{C}}$ and adding $2\varepsilon T$ to the right-hand sides gives the corresponding cover bounds.

Proof. Use the interleaved filtration $\mathcal{F}_{t-1} \subseteq \mathcal{G}_t \subseteq \mathcal{F}_t$ defined above and set

$$Z_{t,\ell} = \ell(P_t, y_t) - \mathbb{E}[\ell(P_t, y_t) \mid \mathcal{G}_t].$$

For every fixed loss, these are martingale differences with respect to the interleaved filtration $\mathcal{F}_{t-1} \subseteq \mathcal{G}_t \subseteq \mathcal{F}_t$, and they lie conditionally in intervals of length two. Conditional Hoeffding's lemma gives

$$\mathbb{E} \exp\left(\lambda \sum_t Z_{t,\ell}\right) \leq \exp(\lambda^2 T/2).$$

Because conditioning on $\mathcal{G}_t$ fixes y_t while preserving the Dirichlet law of P_t, the pathwise conditional version of the worst-case argument gives

$$\sum_{t=1}^{T} \mathbb{E}[\ell(P_t, y_t) \mid \mathcal{G}_t] - T f(q_T) \leq B_T.$$

Hence $\widehat{\mathrm{Reg}}_\ell(T) \leq B_T + \sum_t Z_{t,\ell}$. Chernoff's method and a union bound, applied to the preceding moment-generating-function estimate, prove (9). For a fixed oblivious outcome sequence,

$$\mathbb{E} \max_{\ell \in \mathcal{L}_0} \sum_t Z_{t,\ell} \leq \inf_{\lambda > 0} \left\{ \frac{\log M}{\lambda} + \frac{\lambda T}{2} \right\} = \sqrt{2T \log M},$$

which proves (10). Replacing a loss by a uniformly ε-close cover element changes its realized cumulative loss and its hindsight comparator by at most εT each, proving the cover bounds. □

5 Smooth-loss adaptation

For $t \geq 2$, Dirichlet moments give

$$\mathbb{E} P_t = q_{t-1}, \qquad \mathbb{E}\|P_t - q_{t-1}\|_2^2 = \frac{1 - \|q_{t-1}\|_2^2}{t} \leq \frac{1}{t}. \tag{11}$$

Set $q_0 = P_1$ for convenience. Applying (2) with $p = q_{t-1}$ and $q = P_t$, then taking expectations, makes the linear term vanish because $\mathbb{E}(P_t - q_{t-1}) = 0$. The covariance bound in (11) therefore yields

$$\mathbb{E}[\ell(P_t, y_t) - \ell(q_{t-1}, y_t)] \leq \frac{\beta}{2t}, \qquad t \geq 2. \tag{12}$$

For completeness, we prove the FTL estimate under precisely the one-sided definition (2). This avoids assuming that (2) implies full gradient Lipschitzness.

Lemma 6 (Smooth proper-loss FTL)**.** *For every differentiable proper loss satisfying (2), empirical FTL satisfies*

$$\sum_{t=1}^{T} \ell(q_{t-1}, y_t) - T f(q_T) \leq 2\beta H_T.$$

Proof. We first establish an endpoint condition from propriety and boundary differentiability. Fix $i \neq j$ and set $q_\varepsilon = (1-\varepsilon)e_j + \varepsilon e_i$. Properness at q_ε, comparing the truthful prediction with e_j, says

$$(1-\varepsilon)\ell(q_\varepsilon, e_j) + \varepsilon \ell(q_\varepsilon, e_i) \leq (1-\varepsilon)\ell(e_j, e_j) + \varepsilon \ell(e_j, e_i).$$

Expanding at zero gives $D_{e_i - e_j} \ell(e_j, e_j) \leq 0$. Propriety at e_j also says that e_j globally minimizes $\ell(\cdot, e_j)$, so the reverse inequality holds. Thus every tangent derivative of $\ell(\cdot, e_j)$ at e_j is zero.

Fix $p \in \Delta_K$, write $d = \|e_j - p\|_2$, and define $\phi(s) = \ell((1-s)p + s e_j, e_j)$. Equation (2) is equivalent to concavity of $h(s) = \phi(s) - (\beta d^2/2)s^2$. Since $\phi'(1) = 0$, monotonicity of the derivative of h gives, for $s < 1$,

$$\phi'(s) - \beta d^2 s = h'(s) \geq h'(1) = -\beta d^2.$$

Hence $\phi'(s) \geq -\beta d^2(1-s)$. The FTL update after outcome e_j is $q_t = (1 - 1/t)q_{t-1} + e_j/t$. Integrating from 0 to $1/t$ gives

$$\ell(q_{t-1}, e_j) - \ell(q_t, e_j) \leq \frac{\beta \|e_j - q_{t-1}\|_2^2}{t} \leq \frac{2\beta}{t}.$$

The standard be-the-leader inequality $\sum_t \ell(q_t, y_t) \leq T f(q_T)$ (Cesa-Bianchi & Lugosi, 2006, Lemma 3.1) and summation prove the result. □

Summing (12) and adding the common hindsight comparator first gives the transfer inequality

$$\mathrm{Reg}_\ell(T) \le \mathrm{Reg}_\ell^{\mathrm{FTL}}(T) + \frac{\beta}{2}(H_T - 1).$$

The comparison itself uses only smoothness and the centered Dirichlet moments. Propriety is needed to identify the common empirical comparator and to control FTL. Lemma 6 then gives at most $2\beta H_T + (\beta/2)(H_T - 1) \le (5\beta/2)H_T$, completing Theorem 1.

6 Related work and discussion

U-calibration. Kleinberg et al. (2023) introduced U-calibration for forecasts consumed by agents with unknown utilities. They obtained $O(\sqrt{T})$ expected U-calibration error for binary outcomes and an $O(K\sqrt{T})$ bound on the pseudo-U-calibration criterion for K outcomes. Luo et al. (2024) observed that the geometric-count perturbed leader of Daskalakis & Syrgkanis (2016) attains $O(\sqrt{KT})$ pseudo-U-calibration and proved a matching V-shaped lower bound for $T \ge 12K$. They also identified loss classes on which FTL is logarithmic, and bounded actual U-calibration over finite loss families and families with finite covering numbers. The self-concordant prediction-space perturbation of Frongillo et al. (2026) adapts to smooth losses, and more generally to losses smooth relative to the log barrier, whereas the worst-case-optimal perturbed leader does not adapt to smooth losses. For the classes considered here, its bounds are approximately $K^{5/4}\sqrt{T}$ and $\beta \log T + \beta\sqrt{K}\log K$. Our result removes these gaps.

Two recent calibeating papers establish nearby but distinct guarantees. Fichtl et al. (2026) give a simple FTRL procedure that simultaneously obtains low regret over classes of proper losses including Lipschitz and α-Tsallis losses, and use it to calibeat all losses in such a class at once. Foster & Hart (2026) define proper-calibration and proper-calibeating by requiring errors to vanish uniformly over all bounded proper scoring rules, and relate proper-calibration to universal no regret for agents who best reply to forecasts. These results use calibration or calibeating benchmarks, or apply only to restricted loss families, and do not give worst-case-optimal pseudo-U-calibration over every bounded proper loss. Calibeating and downstream swap-regret results more broadly study related forecasting guarantees under different benchmarks (Chen et al., 2026; Roth & Shi, 2024).

Perturbation and posterior sampling. Follow-the-Perturbed-Leader is classical (Kalai & Vempala, 2005; Cesa-Bianchi & Lugosi, 2006), and fresh perturbations yield the standard nonanticipating adaptive-adversary extension (Hutter & Poland, 2005). Instead of adding noise, van Erven et al. (2014) randomly set past losses to zero or one, which gives a tuning-free perturbed leader for prediction with expert advice. Algorithm 1 can be viewed as posterior sampling for a categorical model under the limiting zero-mass Dirichlet prior, or as Rubin's Bayesian bootstrap. Its closest adversarial sampling precedent is binary. Lewi et al. (2020) analyzed Thompson sampling with a uniform Beta prior for adversarial bit prediction and obtained $O(\sqrt{T})$ regret for weighted error losses. Their theorem is stated for a fixed weighted loss and a thresholded bit action, not for one multiclass probability forecast evaluated under all bounded proper losses, and it does not provide the count-deletion identity or smooth-loss adaptation established here. Other Dirichlet constructions appear in the loss-likelihood bootstrap (Lyddon et al., 2019), online Bayesian bagging (Lee & Clyde, 2004), Dirichlet Bayes mixtures for universal coding under log loss (Watanabe et al., 2013), and size-biased characterizations of Dirichlet laws (Last, 2020). None of these works gives adversarial regret guarantees over all proper losses.

Why Dirichlet geometry fits both regimes. The update $c \mapsto c + e_j$ has likelihood ratio np_j/c_j, so its total variation automatically scales with the number of previous appearances of the updated class. This is the right geometry for discontinuous proper losses. At the same time,

$$\mathrm{Cov}(P_t) = \frac{\mathrm{diag}(q_{t-1}) - q_{t-1}q_{t-1}^\top}{t},$$

which gives the centered $O(1/t)$ mean-squared radius needed for smooth losses. The exact identity in Lemma 2 is what lets these two local facts control global regret without a separate perturbed-leader path-length penalty.

7 LIMITATIONS AND CONCLUSION

Our theorem establishes minimax-optimal worst-case pseudo-U-calibration together with optimal logarithmic dependence on T for Euclidean-smooth proper losses, using one proper multiclass forecaster. It does not control actual U-calibration over the entire infinite class, where the loss can be selected after the forecaster's random trajectory is observed. Corollary 5 handles finite families and classes with finite uniform covers, but for an unrestricted class the resulting complexity term need not be finite. We also do not address the broader class of losses smooth relative to a log barrier, where Euclidean Dirichlet covariance need not be the right measure.

The result shows that the previous tradeoff was geometric rather than information-theoretic. A Bayesian-bootstrap sample of empirical FTL has exactly the count stability needed by nonsmooth scores and exactly the centering needed by smooth scores. The count-deletion identity makes those two properties compatible in one short analysis.

AI USE STATEMENT

AI tools were used to assist with editing the manuscript, TeX formatting, checking mathematical proofs, and writing the supplement code. All outputs were reviewed by the authors, who take full responsibility for the content. AI tools were not used for generating synthetic data sets, creating experiments, assisting with translation, cleaning or reformatting datasets, or data analysis.

REFERENCES

Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.

Yurong Chen, Zhiyi Huang, Michael I. Jordan, and Haipeng Luo. Calibeating made simple. In *Conference on Learning Theory*, pp. 1373–1398. PMLR, 2026.

Constantinos Daskalakis and Vasilis Syrgkanis. Learning in auctions: Regret is hard, envy is easy. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 219–228. IEEE, 2016. doi: 10.1109/FOCS.2016.31.

Maximilian Fichtl, Cristóbal Guzmán, and Nishant A. Mehta. Calibeating for general proper losses: A Bregman divergence approach, 2026. arXiv:2605.17269.

Dean P. Foster and Sergiu Hart. Proper calibeating, 2026. arXiv:2605.26703.

Rafael Frongillo, Haipeng Luo, Nishant A. Mehta, and Jon Schneider. Toward simultaneously optimal regret in U-calibration. In *Conference on Learning Theory*, pp. 2503–2534. PMLR, 2026. arXiv:2606.18527.

Tilmann Gneiting and Adrian E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007. doi: 10.1198/016214506000001437.

Marcus Hutter and Jan Poland. Adaptive online prediction by following the perturbed leader. *Journal of Machine Learning Research*, 6(22):639–660, 2005.

Adam Kalai and Santosh Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005.

Bobby Kleinberg, Renato Paes Leme, Jon Schneider, and Yifeng Teng. U-calibration: Forecasting for an unknown agent. In *Conference on Learning Theory*, pp. 5143–5145. PMLR, 2023. Extended abstract. Full version at arXiv:2307.00168.

Günter Last. An integral characterization of the Dirichlet process. *Journal of Theoretical Probability*, 33(2):918–930, 2020.

Herbert K. H. Lee and Merlise A. Clyde. Lossless online Bayesian bagging. *Journal of Machine Learning Research*, 5:143–151, 2004.

Yuval Lewi, Haim Kaplan, and Yishay Mansour. Thompson sampling for adversarial bit prediction. In Aryeh Kontorovich and Gergely Neu (eds.), *Proceedings of the 31st International Conference on Algorithmic Learning Theory*, volume 117 of *Proceedings of Machine Learning Research*, pp. 518–553. PMLR, 2020.

Haipeng Luo, Spandan Senapati, and Vatsal Sharan. Optimal multiclass U-calibration error and beyond. *Advances in Neural Information Processing Systems*, 37:7521–7551, 2024.

Simon P. Lyddon, Chris C. Holmes, and Stephen G. Walker. General Bayesian updating and the loss-likelihood bootstrap. *Biometrika*, 106(2):465–478, 2019.

Aaron Roth and Mirah Shi. Forecasting for swap regret for all downstream agents. In *ACM Conference on Economics and Computation*, pp. 466–488, 2024.

Donald B. Rubin. The Bayesian bootstrap. *The Annals of Statistics*, 9(1):130–134, 1981.

Tim van Erven, Wojciech Kotłowski, and Manfred K. Warmuth. Follow the leader with dropout perturbations. In *Conference on Learning Theory*, pp. 949–974. PMLR, 2014.

Kazuho Watanabe, Teemu Roos, and Petri Myllymäki. Achievability of asymptotic minimax regret in online and batch prediction. In *Asian Conference on Machine Learning*, pp. 181–196. PMLR, 2013.

A Boundary justification for Lemma 2

We give the omitted analytic details for the count-deletion identity. Let $A = \mathrm{supp}(\alpha)$, assume $|A| \geq 2$ and $\alpha_j > 1$, and use relative Lebesgue measure on Δ_A. By (4), the restriction of f to this face is Lipschitz. It therefore belongs to $W^{1,\infty}$ on the relative interior, is differentiable almost everywhere, and has weak directional derivatives bounded by the corresponding Lipschitz constant.

For $\delta > 0$, truncate the face to

$$\Delta_A^\delta = \{p \in \Delta_A : p_i \geq \delta \text{ for all } i \in A\}.$$

Choose $\delta < 1/|A|$ and mollify f in the affine hull on a neighborhood of Δ_A^δ. For fixed δ, the smooth approximants converge uniformly to f, while their gradients converge almost everywhere to the weak gradient and remain uniformly bounded. We may consequently apply the ordinary divergence theorem to each approximant and pass to the limit by dominated convergence. This yields integration by parts on Δ_A^δ with an explicit boundary-flux term.

It remains to let $\delta \downarrow 0$. On the boundary component $p_i = \delta$, the surface integral of the Dirichlet density has, up to a constant independent of δ, the same endpoint order as the beta marginal density

$$\delta^{\alpha_i - 1}(1-\delta)^{n-\alpha_i-1}.$$

If $i \neq j$, the normal component of $v = e_j - p$ is proportional to $p_i = \delta$, so the absolute flux is $O(\delta^{\alpha_i})$. If $i = j$, that normal component is bounded and the flux is $O(\delta^{\alpha_j - 1}) = o(1)$. These estimates use only $|f| \leq 1$ and remain valid when several coordinates approach zero because the conditional proportions of the other coordinates form an integrable Dirichlet law. Summing over the finitely many boundary components shows that the total flux tends to zero.

The interior terms also converge. The quantity $D_v f$ is bounded almost everywhere. One may use either Lipschitzness or the sharper identity $D_v f = \ell(p, e_j) - f(p)$. The only additional factor is $1/p_j$ from (6), and

$$\int_{\Delta_A} \frac{\rho_\alpha(p)}{p_j}\, dp < \infty \quad \Longleftrightarrow \quad \alpha_j > 1.$$

Dominated convergence therefore gives the boundary-free integration-by-parts formula used above.

Finally suppose $\alpha_j = 1$ and write $\alpha^{(\varepsilon)} = \alpha + \varepsilon e_j$. On the common active face, the densities of $\mathrm{Dir}(\alpha^{(\varepsilon)})$ converge pointwise to that of $\mathrm{Dir}(\alpha)$ and are probability densities. Scheffé's lemma gives L^1, hence total-variation, convergence. This is strong enough to pass to the expectation of the bounded, possibly discontinuous function $\ell(\cdot, e_j)$.

If $n > 1$, then after deleting e_j the jth coordinate under $\mathrm{Dir}(\alpha^{(\varepsilon)} - e_j)$ has distribution $\mathrm{Beta}(\varepsilon, n-1)$. It converges to zero in probability, while conditional on that coordinate the normalized remaining coordinates have law $\mathrm{Dir}(\alpha_{-j})$. The full vector thus converges weakly to $\mathrm{Dir}(\alpha - e_j)$ on the deleted face. The Bayes risk f is continuous by (4), so its expectations converge. If $n = 1$, the coefficient of the deleted term is zero and the identity reduces directly to $\ell(e_j, e_j) = f(e_j)$.

B NUMERICAL CHECKS

A supplementary script checks quantities used in the analysis for binary outcomes. No numerical result is used in any proof.

It first integrates $\mathrm{TV}(\mathrm{Beta}(a, b), \mathrm{Beta}(a + 1, b))$ numerically for six parameter pairs with $a + b \in \{2, 10, 100\}$ and confirms the bound of Lemma 3 in every case. It then compares Monte Carlo estimates of the Dirichlet mean and of $\mathbb{E}\|P - q\|_2^2$ with the closed forms behind (11) for three count vectors.

The main check computes worst-case expected regret exactly. With two outcomes, the expected loss of Algorithm 1 at a round depends only on the current counts. A dynamic program over count vectors therefore maximizes the cumulative expected loss over all 2^T outcome sequences, and subtracting the hindsight comparator gives the worst-case expected regret. We use $P_1 = (1/2, 1/2)$ and two losses. The first is the V-shaped loss with $\ell(p, e_1) = -\frac{1}{2}\,\mathrm{sign}(p_1 - \frac{1}{2})$ and $\ell(p, e_2) = \frac{1}{2}\,\mathrm{sign}(p_1 - \frac{1}{2})$, which is the $K = 2$ case of the lower-bound loss of Luo et al. (2024). The second is the squared loss $\frac{1}{2}\|p - y\|_2^2$, which is 1-smooth. Table 1 reports the results.

Table 1: Exact worst-case expected regret of Algorithm 1 over all binary outcome sequences of length T.

T	V-shaped	V-shaped/$\sqrt{T}$	squared	squared/$\log T$
16	2.799	0.700	1.772	0.639
64	5.951	0.744	2.465	0.593
256	12.272	0.767	3.158	0.569
1024	24.965	0.780	3.851	0.556

Both normalized columns change little over a 64-fold increase in T, consistent with $\sqrt{T}$ growth for the V-shaped loss and logarithmic growth for squared loss. At $T = 1024$ the entries are far below the corresponding bounds $4\sqrt{2T} \approx 181$ and $\frac{5}{2}(1 + \log T) \approx 19.8$ from Theorem 1.