THE MINIMAX RATE OF ONLINE ISOTONIC REGRESSION

Anonymous authors
Paper under double-blind review

ABSTRACT

Kotłowski, Koolen and Malek posed as an open problem, at COLT 2016 and again at NeurIPS 2017, the design of efficient algorithms for online isotonic regression on arbitrary partial orders. Wang (NeurIPS 2026) settled the product orders, conjectured the minimax rate for every finite partially ordered set, and posed the general case as an open problem. We resolve all of these open problems on every finite poset P, determining the squared-loss minimax regret over T rounds up to absolute constants:

$$R_T^*(P) = \Theta\Big(\min\Big\{T,\ \inf_{K\geq 1}\Big[H_P(K) + \frac{T}{K^2}\Big]\Big\}\Big),$$

where $H_P(K)$ is the logarithm of Stanley's order polynomial at $K+1$, which counts the monotone maps from P into $\{0,\dots,K\}$. The same formula holds for log loss and, up to loss-dependent constants, for bounded exp-concave smooth losses that are strongly proper, while pinball loss at quantile level τ has minimax regret $\Theta(\min\{\rho T, \sqrt{\rho T\, H_P(1)}\})$, $\rho = \tau(1-\tau)$. Exponential weights over the monotone maps attains the upper bounds, but its predictions are #P-hard to compute exactly on a general partial order. Our learner instead perturbs the scores with low-dimensional Gumbel noise and samples their maximizers, one minimum cut per sample, and an inequality that we call Gibbs domination shows that this costs only a constant factor in regret. It runs in polynomial time and attains each of these rates without knowledge of the horizon.

1 INTRODUCTION

Isotonic regression fits a function that respects an order. It is seventy years old (Ayer et al., 1955; Brunk, 1955; Barlow et al., 1972; Robertson et al., 1988), and on a partial order it is a convex program whose combinatorial core is a minimum cut (Picard, 1976; Hochbaum & Queyranne, 2003). Two uses drive the demand for it today. In probability calibration, the conversion of a model's scores into probabilities, isotonic regression is a standard method (Zadrozny & Elkan, 2002; Niculescu-Mizil & Caruana, 2005; Guo et al., 2017; Kumar et al., 2019), and the calibration map increasingly takes several signals as input (Pan et al., 2020; Huang et al., 2022; Dai et al., 2025), so that an isotonic one is monotone in each of them. Domain knowledge imposes monotonicity in several features of a learned model (Sill, 1998; Gupta et al., 2016; You et al., 2017; Liu et al., 2020; Runje & Shankaranarayana, 2023) and in the doses of a drug combination (Conaway et al., 2004; Wages et al., 2011). In both uses the constraint is monotonicity in several signals, and the predictions are made as the data arrive. The inputs form a finite set of vectors ordered coordinatewise, and every finite partial order has this form. This is online isotonic regression on a partial order.

A finite partially ordered set, or poset, is known in advance. In each round the learner is shown one of its points, predicts a value in $[0,1]$, then sees the label in $[0,1]$ and pays a loss. Its regret after T rounds is its cumulative loss minus that of the best isotonic function on the set, chosen in hindsight. Points may recur, as in a stream, and nothing is assumed about the labels. The minimax regret is the smallest regret that some learner can guarantee, in expectation over its own randomness, against every sequence of points and labels.

Kotłowski, Koolen and Malek (2016) determined the totally ordered case, a chain of T points, for squared loss. The minimax regret is $T^{1/3}$ up to logarithmic factors, and exponential weights over the isotonic functions with values in a grid gives the upper bound. Their conclusion states a belief and a problem: "We believe that $|\mathcal{F}_K|$ is the right quantity to measure the complexity of the problem",

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.

where $|\mathcal{F}_K|$ is the number of isotonic functions with values in a grid of $K+1$ levels, and "It thus remains an open problem to find an efficient algorithm for the partial order case." They posed it again the next year (Kotłowski et al., 2017): "The next major challenge is the design and analysis of efficient algorithms for online isotonic regression on arbitrary partial orders." Wang (2026) settled the product orders $[m]^d$, in which d variables take every combination of m values, by a formula in $|\mathcal{F}_K|$, conjectured the formula for every finite poset, and posed the general case as an open problem, to establish the conjecture and to match it in polynomial time on every finite poset.

Exponential weights over the isotonic functions with values in a grid has been known since 2016, and Section 3 shows that it attains the minimax rate on every finite poset. Kotłowski et al. (2016) saw what stands in the way: on a partial order it "suffers from the same problems that plague inference in graphical models on general graphs." Even its first prediction is #P-hard to compute exactly, since on a suitable poset it determines the number of order ideals, the downward closed subsets, of a poset of height two, and counting those is #P-complete (Provan & Ball, 1983; Appendix D). Optimization over the ideals is easy, since a maximum-weight ideal is found by one minimum cut, but an optimization oracle alone does not make online learning efficient (Hazan & Koren, 2016).

> *For a finite partially ordered set, what is the smallest regret a learner can guarantee against every sequence of points and labels, relative to the best isotonic function in hindsight, and can a polynomial-time learner attain it?*

We answer both parts, which resolves the open problem of Wang (2026). For squared loss and log loss, the order polynomial of the poset, which counts its isotonic functions into a grid, determines the minimax regret up to absolute constants. It is thus the complexity measure of the problem, as Kotłowski, Koolen and Malek believed. A polynomial-time learner attains this regret on every finite poset. It computes one minimum cut per sampled prediction and matches the regret bound of the algorithm they could not run, up to a constant factor.

1.1 Our Results

For squared loss the answer is the conjectured formula: for every finite poset P and horizon $T \geq 1$,

$$R_T^*(P) = \Theta(r_P(T)), \qquad r_P(T) = \min\Big\{T,\ \inf_{K \geq 1}\Big[2H_P(K) + \frac{T}{4K^2}\Big]\Big\},$$

where $H_P(K) = \log \Omega(P, K+1)$ and $\Omega(P, K+1)$, the order polynomial of P evaluated at $K+1$ (Stanley, 1972; 2012), is the number of isotonic functions from P into $\{0, \ldots, K\}$. The constants are absolute (Theorem 1). Exponential weights over the isotonic functions with values in a grid of resolution K, the route of Kotłowski et al. (2016), pays $2H_P(K)$ once and $T/(4K^2)$ over the whole run, and the rate is the best such balance, capped by the horizon. This upper bound uses nothing about P beyond the count, and the content of the theorem is the lower bound. The isotonic functions with values in a grid correspond to the ideals of a poset built from P and the grid, and the lower bound is therefore a noisy search for an ideal hidden by the adversary. A central element (Linial & Saks, 1985) lies in a fixed positive fraction of the ideals still possible and outside another such fraction. Over a block of order K^2 rounds the adversary asks whether one belongs to the hidden ideal, and the learner pays, in expectation, at least a constant times the entropy revealed (Section 3). The same formula holds for log loss with absolute constants and, with constants depending only on the loss, for every bounded exp-concave smooth loss that is strongly proper (Theorem 2).

The perturbed-ideal learner of this paper attains these rates in polynomial time. It keeps a score for every point and level and partitions the poset built from P and the grid into chains. It encodes each ideal by the lengths of its intersections with the chains and attaches Gumbel noise to this encoding, one variable for each chain and each possible length in place of one for each ideal. One minimum cut computes an ideal of maximum score under the noise, and the law of the value this maximizer takes at a point is the learner's level distribution there. The learner predicts by averaging sampled predictions, the values of independent maximizers at the queried point, and then lowers the scores of the levels at that point by their losses, scaled by an update rate.

The analysis rests on one inequality, Gibbs domination (Theorem 3). The perturbed potential of Section 4.1 is the expected maximum of the perturbed scores. When the scores at a point change, the potential increases by at most a log-partition increment, the logarithm of the mean of the exponentials of the changes, at a fixed scale, under the learner's own level distribution at that point before

the change. The potential of exponential weights satisfies this relation with equality, with its own level distribution and at scale one, so its regret analysis carries over at the cost of a constant factor in the regret bound (Section 4.2). Aggregating one such learner per resolution $K = 1, 2, 4, \ldots$, each following a baseline that learns every point separately before round K, removes the need to know the horizon or the order polynomial. For log loss and for every bounded exp-concave smooth loss, the regret is $O(r_P(T))$ with loss-dependent constants, in expectation at each horizon and with high probability simultaneously for all horizons. For bounded losses the learner takes polylogarithmically many minimum cuts per round (Theorem 4).

Pinball loss (Koenker & Bassett, 1978; Steinwart & Christmann, 2011), the loss of quantile estimation at level $\tau \in (0, 1)$, obeys a second law. Its minimax regret is $\Theta(\min\{\rho T, \sqrt{\rho T H_P(1)}\})$ with absolute constants, where $\rho = \tau(1 - \tau)$ (Theorem 10 in Appendix H). There is no resolution to choose. The regret of a convex loss is at most that of its linearization, and a linear functional on the isotonic class attains its minimum at a binary isotonic function. So the binary isotonic functions determine the rate, and their Littlestone dimension (Littlestone, 1988; Ben-David et al., 2009), the optimal mistake bound of online classification when the labels come from the class, is $H_P(1)$ up to absolute constants. The perturbed-ideal learner at resolution one, with its prediction corrected for the two slopes of the loss, attains the rate in polynomial time. The rate $\min\{T, \sqrt{T H_P(1)}\}$ holds up to loss-dependent constants for every bounded convex Lipschitz loss under which some label law has two Bayes-optimal predictions whose losses differ with positive probability (Corollary 3).

The rate $r_P(T)$ has three phases on a chain and, as Wang (2026) showed, on a product order, and two on posets of bounded height (Table 1 in Appendix C). Beyond isotonic classes, for every class defined by nonnegative upper bounds on the differences of a function's values over ordered pairs of points, a maximum-weight member of its rounded class, in which the values lie in a grid and the bounds are rounded up to it, is again computed by one minimum cut. For squared loss such a class obeys the rate formula, with the count of the rounded class in place of the order polynomial, from a horizon determined by the bounds, and a polynomial-time learner attains it for a given horizon (Theorem 9). The appendix contains all proofs and the learners' pseudocode.

1.2 Related Work

The balanced query of the lower bound comes from the search literature. Linial & Saks (1985) proved that every nonempty finite poset has a central element in order to show that membership questions identify an unknown ideal within a constant factor of the information bound, the logarithm of the number of ideals. When the answers are noisy, the cost is still the number of queries (Rivest et al., 1980; Feige et al., 1994; Karp & Kleinberg, 2007; Ben-Or & Hassidim, 2008; Nowak, 2011; Emamjomeh-Zadeh et al., 2016; Dereniowski et al., 2025). With regret as the cost, the two scales of the search, the $H_P(K)$ nats of the hidden ideal and the blocks of order K^2 rounds that reveal them, become the two terms of the rate. On product orders, Wang (2026) proved the formula with lower bounds built on the rank levels of the grid and with learners that slice the grid into planes. Here both the search and the maximization run on the ideals of any finite poset.

Up to a factor of two in resolution, the logarithm of the order polynomial is the metric entropy of the isotonic class, the logarithm of the number of functions needed to approximate every member uniformly to a given accuracy (Remark 2). The theory of online nonparametric regression controls the regret of a class by its metric entropy or its sequential version (Vovk, 2006; Rakhlin & Sridharan, 2014; 2015; Gaillard & Gerchinovitz, 2015; Rakhlin et al., 2015; Block et al., 2021), giving the order of the upper bound. Its lower bounds for rates below $\sqrt{T}$ hold for modified classes (Rakhlin & Sridharan, 2014), and its efficient learners are for smoothness classes on Euclidean domains (Gaillard & Gerchinovitz, 2015; Kuzborskij & Cesa-Bianchi, 2020; Liautaud et al., 2025a;b). Here the lower bound is for the isotonic class itself, and the learner runs in polynomial time.

Follow-the-perturbed-leader (Hannan, 1957; Kalai & Vempala, 2005) and its smoothing view (Abernethy et al., 2014; 2015; 2016) obtain regret of order $\sqrt{T}$ from a bound on the curvature of the smoothed potential, the expected perturbed maximum as a function of the scores. Oracle-efficient online learning (Dudík et al., 2020; Syrgkanis et al., 2016; Wang et al., 2022), which asks for regret guarantees with polynomially many calls to an optimization oracle per round, does so with low-dimensional perturbations, and Gumbel noise on every action samples the exponential-weights distribution (Abernethy et al., 2016). Here each sampled prediction takes one minimum cut, and

the regret matches that of exponential weights up to a constant for exp-concave losses. Theorem 3 compares the increments of the perturbed potential with those of a log-partition function under the learner's own level distribution, the form of bound from which exp-concavity yields regret independent of the horizon against a fixed finite class. Its proof is a comparison of curvatures, as in differential consistency (Abernethy et al., 2015), for independent noise on every action, and in generalized mixability (Reid et al., 2015; Mhammedi & Williamson, 2018), for prediction with expert advice.

2 PRELIMINARIES

Sections 3 and 4 share the objects defined here. The learner's own are defined in Section 4.1.

Protocol and losses. Throughout, P is a finite partially ordered set, known to the learner, with $n \geq 1$ points, order $\preceq$, width w and height h, the sizes of a largest antichain and of a largest chain. A function $f\colon P \to [0,1]$ is isotonic if $f(x) \leq f(x')$ whenever $x \preceq x'$, and $\mathcal{F}_P$ is the class of isotonic functions. In round $t = 1, 2, \ldots$, the adversary reveals a point $x_t \in P$, the learner predicts $\widehat{y}_t \in [0,1]$, possibly using private randomness, the adversary reveals a label $y_t \in [0,1]$, and the learner pays $\ell(\widehat{y}_t, y_t)$. The loss $\ell\colon [0,1]^2 \to [0,\infty]$ is finite at every prediction in $(0,1)$. The regret after T rounds is

$$R_T = \sum_{t \leq T} \ell(\widehat{y}_t, y_t) - \min_{f \in \mathcal{F}_P} \sum_{t \leq T} \ell(f(x_t), y_t),$$

where the minimum is read as an infimum when it is not attained. The minimax regret $R_T^*(P, \ell)$ is the infimum over randomized learners of the supremum over sequences $(x_t, y_t)_{t \leq T}$ of $\mathbb{E}R_T$. For squared loss we write $R_T^*(P)$. Points may recur. Every upper bound in the paper holds for every sequence, hence also in the single-pass protocol of Kotłowski et al. (2016), where each point of a given set is queried once. It also holds against a non-anticipating adaptive adversary, who may choose (x_t, y_t) from the past realized predictions but not from the learner's fresh randomness of round t. The lower bounds use repeated queries and oblivious random sequences, drawn before play.

A design horizon is a horizon T given to the learner in advance, and a learner given none is horizon-free. For a fixed loss with fixed regularity bounds, a learner runs in polynomial time if its computation per round is polynomial in n, the round t, the design horizon if any, the input bit lengths and $\log(1/\delta)$, with $\delta \in (0,1)$ the confidence parameter of its guarantee. Computation is measured in the finite-precision model of Appendix I, which preserves the order of every regret bound (Theorem 12).

Logarithms are natural. $\mathrm{Ber}(p)$ is the Bernoulli law with mean p, Δ^K the probability simplex on $\{0, \ldots, K\}$, and $[K] = \{1, \ldots, K\}$; Θ and O hide absolute constants unless stated otherwise.

Definition 1 (Loss Conditions)**.** Let $\ell\colon [0,1]^2 \to [0,\infty]$ be a loss, finite at every prediction in $(0,1)$, and let $\eta, \beta, \Lambda, \alpha > 0$. (i) ℓ is η-exp-concave if $p \mapsto e^{-\eta \ell(p,y)}$ is concave on $[0,1]$ for every $y \in [0,1]$ (Kivinen & Warmuth, 1999; Cesa-Bianchi & Lugosi, 2006). (ii) ℓ is β-smooth if for every y the function $\ell(\cdot, y)$ is differentiable on $[0,1]$, with one-sided derivatives at the endpoints, and its derivative is β-Lipschitz. (iii) ℓ has range Λ if $0 \leq \ell \leq \Lambda$ on $[0,1]^2$. (iv) ℓ is α-strongly proper on $[1/4, 3/4]$ if $\mathbb{E}_{Y \sim \mathrm{Ber}(p)}[\ell(p', Y) - \ell(p, Y)] \geq \alpha(p' - p)^2$ for every $p \in [1/4, 3/4]$ and $p' \in [0,1]$, the interior form of strong properness (Agarwal, 2014) with parameter 2α.

Squared loss $\ell(p,y) = (p-y)^2$ is $\frac{1}{2}$-exp-concave, 2-smooth and 1-strongly proper, with range 1. Log loss $\ell_{\log}(p,y) = -y \log p - (1-y)\log(1-p)$, with $0 \log 0 = 0$ and limits at the predictions 0 and 1, is 1-exp-concave, 2-strongly proper by Pinsker's inequality, and neither smooth nor bounded.

Isotonic functions and order ideals. A grid is a finite set $G = \{z_0 < \cdots < z_K\} \subseteq [0,1]$ of $K+1$ levels, with resolution $K \geq 1$, and $G_K = \{0, 1/K, \ldots, 1\}$ is the uniform grid. We write $\mathcal{F}_G = \mathcal{F}_P \cap G^P$ for the isotonic functions with values in G and $\mathcal{F}_K = \mathcal{F}_{G_K}$. The integer form of $f \in \mathcal{F}_G$ is $F\colon P \to \{0, \ldots, K\}$ with $F(x) = j$ when $f(x) = z_j$, so that $F = Kf$ on G_K.

Stanley's order polynomial, evaluated at $K+1$, counts the isotonic functions from P into a chain of $K+1$ elements (Stanley, 1972; 2012). Thus $|\mathcal{F}_K| = \Omega(P, K+1)$, and for $K \geq 1$ we put

$$H_P(K) = \log \Omega(P, K+1), \quad \Psi_P(T) = \inf_{K \geq 1} \Big\{ 2H_P(K) + \frac{T}{4K^2} \Big\}, \quad r_P(T) = \min\{T, \Psi_P(T)\}.$$

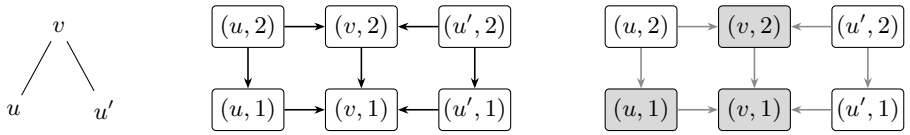


Figure 1: Left: the poset P with $u \preceq v$ and $u' \preceq v$. Middle: Q_2, with an arrow from each cell to each cell it covers. Right: the ideal I_F for $(F(u), F(u'), F(v)) = (1, 0, 2)$, shaded.

We call $r_P(T)$ the rate. Two facts recur (Lemma 1 in Appendix A): among posets with n points, $H_P(K)$ is smallest for a chain, where $\Omega(P, K+1) = \binom{n+K}{n}$, and largest for an antichain, where $\Omega(P, K+1) = (K+1)^n$; and H_P is subadditive, $H_P(K+K') \le H_P(K) + H_P(K')$ for $K, K' \ge 1$.

An ideal of a finite poset Q is a downward closed subset and a filter an upward closed one. $\mathcal{J}(Q)$ is the set of its ideals, a lattice under inclusion, and q covers q' if $q' < q$ with no element between.

Definition 2 (The Poset Q_K and Its Ideals)**.** Let $K \ge 1$. The poset Q_K is the set $P \times [K]$ ordered by $(x, j) \le (x', j')$ if and only if $x' \preceq x$ and $j \le j'$, with strict order $<$. An element $q = (x, j)$ of Q_K is called a cell; its point is $x(q) = x$ and its level is $j(q) = j$. The column of a point x is the set of its K cells. The ideal of an isotonic $F\colon P \to \{0, \dots, K\}$ is $I_F = \{(x, j) : j \le F(x)\}$.

The map $F \mapsto I_F$ is a bijection from the isotonic functions $P \to \{0, \dots, K\}$ onto $\mathcal{J}(Q_K)$, with inverse $F_I(x) = |I \cap (\{x\} \times [K])|$ (Lemma 6 in Appendix B), so $\Omega(P, K+1) = |\mathcal{J}(Q_K)|$ (Figure 1).

The theorem of Linial & Saks (1985) enters through one consequence. Every nonempty finite poset has a central element, one lying in a fraction of its ideals within $[\delta_{\mathrm{LS}}, 1 - \delta_{\mathrm{LS}}]$, where $\delta_{\mathrm{LS}} = (3 - \log_2 5)/4 = 0.1695\dots$; we put $\lambda = -\log(1 - \delta_{\mathrm{LS}}) = 0.1857\dots$. Whether it lies in a hidden ideal is a balanced binary question, each answer keeping at least a δ_{LS} fraction of the candidates.

3 Minimax Rates

This section states the rate theorems, proves the upper bounds and sketches the lower bound.

Theorem 1 (Minimax Rate under Squared Loss)**.** *For every finite poset P and every integer $T \ge 1$,*

$$\frac{\lambda}{2048}\, r_P(T) \le R^*_T(P) \le r_P(T),$$

where $\lambda = -\log(1 - \delta_{\mathrm{LS}}) = 0.1857\dots$, so that $\lambda/2048 > 9.06 \cdot 10^{-5}$. The upper bound is attained by a deterministic learner.

Thus the order polynomial of P determines the function $T \mapsto R^*_T(P)$ up to absolute constants.

Theorem 2 (Minimax Rate beyond Squared Loss)**.** *Let P be a finite poset and $T \ge 1$.*

(i) If ℓ is α-strongly proper on $[1/4, 3/4]$, then $R^*_T(P, \ell) \ge (\alpha\lambda/2048)\, r_P(T)$.

(ii) If ℓ is η-exp-concave, β-smooth and $0 \le \ell \le \Lambda$, then

$$R^*_T(P, \ell) \le \min\{\Lambda T,\ \max\{1/(2\eta), \beta/2\}\, \Psi_P(T)\}.$$

(iii) For log loss,

$$\frac{2\lambda}{2048}\, r_P(T) \le R^*_T(P, \ell_{\log}) \le \min\Big\{T \log 2,\ \inf_{K \ge 1}\Big[H_P(K) + \frac{6T}{K^2}\Big]\Big\} \le 24\, r_P(T).$$

Consequently $R^*_T(P, \ell) = \Theta(r_P(T))$ with absolute constants for log loss, and with constants depending only on $(\alpha, \eta, \beta, \Lambda)$ for every loss satisfying (i) and (ii).

A phase is a range of horizons on which the rate has one form, and Table 1 in Appendix C gives the linear, intermediate and logarithmic phases on three families of posets. On a chain at $n = T \ge 2$ the rate is of order $T^{1/3} \log^{2/3} T$, while Kotłowski et al. (2016) prove an upper bound of this order and a

lower bound of order $T^{1/3}$ in their single-pass protocol. On every poset, $\Psi_P(\varrho T) \le \lceil \varrho^{1/3} \rceil \Psi_P(T)$ for every $\varrho \ge 1$ by the subadditivity of H_P (Lemma 1), so once $\Psi_P(T) \le T$ the rate grows at most like $T^{1/3}$. Every exponent $\omega \in (0, 1/3)$ occurs, on families of posets on which the rate grows like T^ω over ranges of horizons $[T_1, T_2]$ with $T_2/T_1 \to \infty$ along the family, and such growth at horizons of larger order than n is possible only when the width is of smaller order than n (Appendix C). The logarithmic phase is reached by $T = \max\{2, n^4\}$ on every poset. From there on the rate is $n \log T$ up to absolute constants (Remark 3), the rate of n independent scalar problems, which the scalar baseline of Section 4.3 attains point by point.

The upper bounds are attained by exponential weights with the uniform prior over the isotonic functions into a grid of $K + 1$ levels. It predicts the mean of the values $f(x_t)$ under weights proportional to $\exp(-\eta \sum_{s<t} \ell(f(x_s), y_s))$, and for a finite η-exp-concave loss its cumulative loss exceeds that of every such function by at most $H_P(K)/\eta$ on every sequence (Vovk, 1998; Cesa-Bianchi & Lugosi, 2006; Lemma 3). On the uniform grid, common-shift rounding, $\lfloor f \rceil_\upsilon(x) = \lfloor Kf(x) + \upsilon \rfloor / K$ with one shift $\upsilon \in [0, 1)$ for all points, maps $\mathcal{F}_P$ into $\mathcal{F}_K$, and for a β-smooth loss it raises the cumulative loss of f by at most $\beta T/(8K^2)$ in expectation over a uniform shift, hence for some shift (Lemma 2). Squared loss has $\eta = 1/2$ and $\beta = 2$, the two costs add to $2H_P(K) + T/(4K^2)$, and the infimum over K is $\Psi_P(T)$. Since the regret is at most T, the upper bound of Theorem 1 is $r_P(T)$, and Theorem 2(ii) is the same with the cap ΛT (Theorem 5). For log loss, which is neither smooth nor bounded, the hindsight optimum is rounded along its level sets. On each queried level set its value is the mean of the labels there, so rounding it costs one Kullback–Leibler divergence between Bernoulli laws per query, below $6/K^2$ on a grid of interior levels (Lemmas 9 and 10).

For the lower bound, the adversary hides a uniformly random ideal of Q_K, that is, a uniformly random isotonic $F\colon P \to \{0, \dots, K\}$, and plays in blocks of $32K^2$ rounds, each block at one point with Bernoulli labels whose mean $1/4 + F/(2K)$ at that point encodes the hidden level there. This mean is an isotonic function of the point, so its cumulative loss is at least the hindsight minimum.

Sketch of the lower bound. Fix K and $T \ge 32K^2$. At the start of block ι, while more than one ideal is compatible with the levels revealed so far, the residual poset is the induced poset on the cells that lie in some but not all of these ideals. The adversary takes a central element ζ_ι of it, queries the point $x(\zeta_\iota)$ throughout the block, and then reveals the level $V_\iota = F(x(\zeta_\iota))$ to the learner, which only helps. Conditionally on the revealed levels, all earlier labels and the learner's private randomness, the unrevealed part of the hidden ideal is uniform on the ideals of the residual poset. The earlier labels have the same likelihood for every compatible ideal, because their means are fixed by the revealed levels, and Lemma 6(b) and (c) identify the compatible ideals with the ideals of the residual poset. So the theorem of Linial and Saks applies again, and the probability, given the past, that V_ι is at least the level of ζ_ι lies in $[\delta_{\mathrm{LS}}, 1 - \delta_{\mathrm{LS}}]$. In Figure 1, the cell $(u, 1)$ lies in 8 of the 14 ideals, a balanced query. Each revealed level keeps at most a $1 - \delta_{\mathrm{LS}}$ fraction of the compatible ideals, so over $M = \lfloor T/(32K^2) \rfloor$ blocks either the levels identify the ideal, revealing all $H_P(K)$ nats, or the logarithm of the number of compatible ideals falls by at least λM: the revealed Shannon entropy $\sum_\iota \mathsf{H}(V_\iota \mid V_{<\iota})$ is at least $\min\{H_P(K), \lambda M\}$. Each block costs the learner, in expectation, at least $3\alpha/32$ times the entropy it reveals (Lemma 7). By strong properness, the expected excess loss of a round is at least α times the expected variance, given the past, of the mean of its label, and hence at least $3\alpha/16$ times the information that the label adds about the level, because the means lie in $[1/4, 3/4]$. The $32K^2$ labels of a block, whose means at adjacent levels differ by $1/(2K)$, leave less than $1/25$ of a nat of the level's conditional entropy, which was at least $2/5$ of a nat at the start of the block because the query is balanced, so the labels alone disclose at least half of it. The blocks after the search has ended query a fixed point at nonnegative expected excess loss. Since $M \ge T/(64K^2)$, this gives the bound $(3\alpha\lambda/2048) \min\{H_P(K), T/K^2\}$, which for $T < 32K^2$ follows from a single round with a label mean drawn uniformly from $\{1/4, 3/4\}$ (Proposition 1). The best K gives the theorem, since $r_P(T) \le 3 \sup_{K \ge 1} \min\{H_P(K), T/K^2\}$ (Lemma 8). Both F and the label coins are drawn before play and fix the sequence, so the adversary is oblivious.

4 A Polynomial-Time Learner

Exponential weights over the isotonic functions with values in a grid attains the rates of Section 3, but on a general poset even its first prediction is #P-hard to compute exactly (Appendix D). The

perturbed-ideal learner predicts instead from the level distribution of a perturbed potential, drawing each level by one minimum cut. Section 4.1 defines its objects, Section 4.2 gives the inequality behind it and its guarantee at a fixed resolution, and Section 4.3 makes it horizon-free.

4.1 Perturbed Maximum-Weight Ideals

A chain partition of Q_K divides its cells into chains. A minimum one has $w(Q_K)$ chains, the width of Q_K, by Dilworth's theorem (Dilworth, 1950), and is found by maximum bipartite matching, as in Fulkerson's proof (Fulkerson, 1956). An ideal meets each chain in a prefix, an initial segment of it. A standard Gumbel variable is the negative logarithm of a standard exponential one. Its mean is Euler's constant $\gamma = 0.5772\ldots$, and the variable minus γ is a centered standard Gumbel variable.

Definition 3 (Perturbed Potential and Level Distribution)**.** Let $K \geq 1$ and fix a minimum chain partition $\mathcal{C}_1, \ldots, \mathcal{C}_{w_K}$ of Q_K, $w_K = w(Q_K)$. The prefix length of an ideal I in chain i is $k_i(I) = |I \cap \mathcal{C}_i|$. The set of noise coordinates is $\Xi = \{(i,k) : 1 \leq i \leq w_K,\ 0 \leq k \leq |\mathcal{C}_i|\}$. The noise $\xi = (\xi_{i,k})_{(i,k)\in\Xi}$ consists of independent centered standard Gumbel variables, and $\xi(I) = \sum_i \xi_{i,k_i(I)}$. A score array is an element $\theta \in \mathbb{R}^{P\times\{0,\ldots,K\}}$, and the score of an ideal is $\theta(I) = \sum_{x\in P} \theta_{x,F_I(x)}$. The perturbed potential is

$$\Phi(\theta) = \mathbb{E} \max_{I\in\mathcal{J}(Q_K)} \{\theta(I) + \xi(I)\}.$$

The perturbed maximizer $I^* = I^*(\theta, \xi)$ is the ideal attaining the maximum, unique with probability one (Appendix D). The level distribution of a point x at θ is the law $\mu_x(\theta) \in \Delta^K$ of $F_{I^*}(x)$.

The tuple $(k_i(I))_{i\leq w_K}$ of prefix lengths determines I and is the chain encoding of I. There are $|\Xi| = nK + w_K \leq 2nK$ noise coordinates. For $x \in P$ and $a \in \mathbb{R}^{K+1}$, $\theta +_x a$ is θ with a_j added to the score $\theta_{x,j}$ for every j. The level distribution is the gradient of the potential, $\mu_{x,j}(\theta) = \partial\Phi(\theta)/\partial\theta_{x,j}$ (Lemma 17). A learner with scores θ_t has $\mu_t = \mu_{x_t}(\theta_t)$, and $\mathrm{Var}_\mu(a) = \sum_j \mu_j a_j^2 - (\sum_j \mu_j a_j)^2$. In Figure 1, Q_2 has width three, one minimum chain partition is $(v,1) < (u,1) < (u,2)$, $(v,2) < (u',2)$ and $(u',1)$, and 14 of the 24 tuples of prefix lengths are chain encodings of ideals.

For every score array θ and every value of the noise ξ, $\theta(I) + \xi(I)$ is a constant plus a sum of node weights over the cells of I. The cell (x,j) that is the k-th element of $\mathcal{C}_i$ carries the weight $\theta_{x,j} - \theta_{x,j-1} + \xi_{i,k} - \xi_{i,k-1}$. The reduction of Picard (1976) therefore computes a maximum-weight ideal by one minimum cut, on a network with $nK + 2$ vertices and $O(n^2K)$ arcs (Appendix D).

4.2 Gibbs Domination

The potential dominates the log-partition function of exponential weights, and at zero scores the two agree up to a constant factor: for every score array θ,

$$\log \sum_{I\in\mathcal{J}(Q_K)} e^{\theta(I)} \leq \Phi(\theta) \qquad \text{and} \qquad H_P(K) \leq \Phi(0) \leq c_\Phi\, H_P(K),$$

where $c_\Phi = 2 + (\log\pi - \gamma)/\log 2 < 2.82$ (Lemma 14). The first inequality, which puts $\Phi(\theta)$ above the score of every ideal, is the bound of Hazan & Jaakkola (2012) for low-dimensional perturbations. The last holds because the noise adds at most $\log\pi - \gamma$ per chain to $2H_P(K)$, and a minimum chain partition has at most $H_P(K)/\log 2$ chains. As a by-product, an estimate of $\Phi(0)$ from $O(nK\log(2/\delta))$ minimum cuts approximates $H_P(K)$, the logarithm of the order polynomial, within the factor 2.82 with probability at least $1-\delta$ (Remark 4).

Theorem 3 (Gibbs Domination)**.** *Let P be a finite poset, $K \geq 1$, and let Φ be the perturbed potential of a minimum chain partition of Q_K. For every score array θ, every point $x \in P$ and every $a \in \mathbb{R}^{K+1}$,*

$$\Phi(\theta +_x a) - \Phi(\theta) \leq \frac{1}{\kappa} \log \sum_{j=0}^{K} \mu_{x,j}(\theta)\, e^{\kappa a_j}, \qquad \kappa = 4;$$

if $K = 1$, the inequality holds with $\kappa = 1$.

Theorem 3 turns the update identity of exponential weights into an inequality. Adding a to the scores at x changes $\log\sum_I e^{\theta(I)}$ by exactly the right side with $\kappa = 1$, once $\mu_x(\theta)$ is replaced by the law of $F_I(x)$ under the weights proportional to $e^{\theta(I)}$, and it changes Φ by at most the right side as it stands.

Sketch of Theorem 3. The proof bounds the curvature of $a \mapsto \Phi(\theta +_x a)$ by the variance of the level distribution and converts that bound into the inequality by a change of variables. First, if membership in a set of ideals depends only on the prefix lengths in L fixed chains, then raising by $\varsigma \geq 0$ the scores of its members multiplies the probability that the perturbed maximizer lies in it by at most $e^{L\varsigma}$. Indeed, adding ς to a member's noise coordinates in the L chains gives it at least as much extra noise as any other member and at least ς more than any ideal outside the set, which differs from it in one of these prefix lengths. Hence this translation maps the noise values at which the member attains the maximum after the raise to those at which it does before, and it multiplies the noise density by at least $e^{-L\varsigma}$ (Lemma 16). Second, the level fiber $\mathcal{J}_j(x) = \{I : F_I(x) = j\}$ consists of the ideals that contain (x, j) if $j \geq 1$ and not $(x, j+1)$ if $j < K$, so membership in it depends only on the prefix lengths in the at most two chains through these cells, and raising a_j raises exactly the scores of its members. By the first step with $L = 2$, the level distribution $\mu_x(\theta +_x a)$, which is the gradient of $a \mapsto \Phi(\theta +_x a)$, is Lipschitz in a, and its Jacobian, defined almost everywhere, has diagonal entries at most $2\mu_{x,j}$. Its off-diagonal entries are nonpositive, because raising a_j can only move the maximizer into $\mathcal{J}_j(x)$, and its rows and columns sum to zero, because a constant added to a changes no maximizer and the level probabilities sum to one. Its quadratic form is therefore that of a weighted Laplacian, and the diagonal bound gives, at almost every a,

$$\sum_{j,j'} a'_j a'_{j'} \frac{\partial \mu_{x,j'}(\theta +_x a)}{\partial a_j} \leq 4 \operatorname{Var}_{\mu_x(\theta +_x a)}(a') \qquad \text{for every } a' \in \mathbb{R}^{K+1}$$

(Lemma 17). Third, since this gradient is Lipschitz with values in Δ^K, the bound makes the function $\varpi \mapsto \exp(\kappa\Phi(\theta +_x \kappa^{-1} \log \varpi))$, with $\kappa = 4$ and the logarithm taken entrywise, concave on the positive orthant, and its tangent-plane inequality at the all-ones vector, evaluated at $(e^{\kappa a_j})_j$, is the inequality of the theorem (Theorem 6). At $K = 1$ the prefix length in the chain through $(x, 1)$ determines the level at x. Given the noise of the other chains, the expected increment is exactly a log-partition increment over these prefix lengths, and Jensen's inequality gives $\kappa = 1$ (Lemma 18).

At a fixed resolution the learner works on a grid $G = \{z_0 < \cdots < z_K\} \subseteq [0, 1]$, with an η-exp-concave loss that is finite on $[z_0, z_K]$. It starts from $\theta_1 = 0$, predicts $\bar{y}_t = \sum_j \mu_{t,j} z_j$, the mean prediction, and then lowers the scores at x_t by the losses of the grid values, $\theta_{t+1} = \theta_t +_{x_t} \big(-(\eta/\kappa)\, \ell(z_j, y_t)\big)_{j=0}^{K}$, with κ as in Theorem 3. By exp-concavity the loss of the mean prediction is at most the mix loss (de Rooij et al., 2014) of the losses $\ell(z_j, y_t)$ under μ_t, which Theorem 3, applied to this update, bounds by the potential's decrease divided by the update rate η/κ:

$$\ell(\bar{y}_t, y_t) \leq -\frac{1}{\eta} \log \sum_{j=0}^{K} \mu_{t,j}\, e^{-\eta \ell(z_j, y_t)} \leq \frac{\kappa}{\eta} \big(\Phi(\theta_t) - \Phi(\theta_{t+1})\big).$$

Summing over the rounds telescopes the right side, and $\Phi(\theta_{T+1})$ is at least the score of every ideal, which is $-\eta/\kappa$ times the cumulative loss of the corresponding member of $\mathcal{F}_G$. On every sequence the regret against $\mathcal{F}_G$ is therefore at most $\kappa\Phi(0)/\eta \leq \kappa c_\Phi H_P(K)/\eta$, which is the bound $H_P(K)/\eta$ of exponential weights over $\mathcal{F}_G$ (Lemma 3) up to the factor $\kappa c_\Phi < 11.3$ (Proposition 5(a)).

The mean prediction is an expectation over the noise. For a loss with values in $[0, 1]$ the learner replaces it by the average of an even number S of sampled predictions, the grid values $z_{F_{I^*}(x_t)}$ at the levels of independent perturbed maximizers I^* at θ_t, one minimum cut each, and updates at rate $\min\{\eta, 1\}/(16\kappa)$. Given the past and (x_t, y_t), exp-concavity puts the expected loss of the average of two sampled predictions below that of one by a multiple of the variance of that loss, and at this update rate that gain absorbs the variance term of the exponential-moment bound. A supermartingale and Ville's inequality (Howard et al., 2020) bound the regret against $\mathcal{F}_G$ by $O(H_P(K)/\min\{\eta, 1\})$ in expectation and, with $S = 2\max\{1, \lceil \log(1/\delta) \rceil\}$, with probability at least $1 - \delta$ simultaneously for all $T \geq 1$ (Proposition 5(b)).

4.3 The Horizon-Free Learner

The horizon-free learner runs the fixed-resolution learner at the dyadic resolutions $K = 1, 2, 4, \ldots$ and aggregates the runs by exponential weights. Algorithms 1 to 3 of Appendix E state its rules in full. A loss with values in $[0, \Lambda]$ is run as ℓ/Λ on the uniform grid G_K, with averages of sampled predictions. Log loss is run on a grid of interior levels, on which it is finite and Lipschitz

(Lemma 11), with a sampled estimate of the mean prediction whose cost in loss is summable over the rounds. For each dyadic K, the expert E_K follows the scalar baseline B up to round $K-1$ and then runs a fresh learner at resolution K. The baseline ignores the order. At each point separately it runs exponential weights over grids of growing resolution, and its regret against every function $P \to [0,1]$ over t rounds is at most $C_\ell\, n \log(1+t/n)$, with C_ℓ depending only on the loss (Lemma 20). A master aggregates the constant expert $1/2$, with prior $1/2$, and the experts E_K, with priors $\nu_K = 1/(2(\log_2 K+1)(\log_2 K+2))$, which sum to $1/2$. The experts not yet activated have all followed B, so a single weight represents them, and only the $1+\lfloor \log_2 t \rfloor$ experts with $K \le t$ run at round t. The baseline pays for this delayed activation, at most $C_\ell\, n \log(1+K/n)$ over the first K rounds, and this cost is at most $C_\ell\, H_P(K)$ (Lemma 1):

$$n \log(1+K/n) \le \log \binom{n+K}{n} \le H_P(K).$$

Theorem 4 (Horizon-Free Polynomial-Time Learning)**.** *Let ℓ be η-exp-concave and β-smooth with $0 \le \ell \le \Lambda$, $\Lambda > 0$, or let ℓ be log loss, and let $\delta \in (0,1)$. The horizon-free perturbed-ideal learner of Section 4.3, run on ℓ/Λ in the first case, satisfies, for every finite poset P and every non-anticipating adaptive adversary,*

$$\mathbb{E} R_T \le C_\ell\, r_P(T) \text{ for every } T \ge 1, \qquad \mathbb{P}\{R_T \le C'_\ell\, r_P(T) \text{ for all } T \ge 1\} \ge 1-\delta,$$

where C_ℓ and C'_ℓ depend only on (η, β, Λ) and are absolute constants for log loss; they do not depend on P, T or δ. For each fixed loss and fixed regularity bounds its computation per round is polynomial in n, t, the input bit lengths and $\log(1/\delta)$ in the finite-precision model; for bounded losses it uses $O\big((1+\log t)(1+\log(1/\delta)+\log\log(t+3))\big)$ minimum cuts at round t, each on a network with $O(nt)$ vertices and $O(n^2 t)$ arcs, and for log loss polynomially many.

For log loss and the strongly proper losses among these, the rate is minimax (Theorem 2).

At resolution K, the fixed-resolution learner and the rounding of its grid cost $C_\ell(H_P(K) + T/K^2)$ against $\mathcal{F}_P$ at every prefix, and the master pays a constant multiple of $H_P(K)$ beyond that: the delay, as above, and the prior ν_K, whose cost, a constant multiple of $\log(1/\nu_K)$, is absorbed by $H_P(K) \ge \log(K+1)$. Restricting the resolutions to powers of two costs a constant factor, since doubling a resolution at most doubles H_P (Lemma 21). Each expert E_K runs at the confidence $\delta_K = \delta\nu_K$, and its guarantee has constants free of δ_K. Since the confidences sum to $\delta/2$, the constants of Theorem 4 are free of δ as well. The cap T of $r_P(T)$ comes from the range of a bounded loss, and for log loss from the constant expert, whose loss is $\log 2$ per round.

A cheaper alternative drops the constraints between the chains of a minimum chain partition of P and runs exponential weights per chain. It pays the logarithm of the size of the enlarged grid class, at most $\min\{w,h\}\, H_P(K)$, instead of $H_P(K)$. The factor is sharp up to constants on a family of width-w posets, where at $T = w^5$ its regret bound is of order $w^{2/3}$ times the rate (Remark 5).

5 Discussion

The theorems of this paper determine the minimax regret of online isotonic regression on every finite poset, and a polynomial-time learner attains it. They all rest on one object, the lattice of ideals of Q_K, whose elements encode the isotonic functions at resolution K. Its size sets the rate of Theorems 1 and 2. Exponential weights on the grid pays the logarithm of the order polynomial $\Omega(P, K+1)$ plus the rounding cost T/K^2, and the best resolution balances the two. Searching the lattice shows that no learner has regret of smaller order. An adversary hides an ideal and answers, in blocks of noisy Bernoulli labels, whether central elements belong to it. Without the noise and at resolution one, the same questions build a tree of queries whose depth governs the rate under pinball loss (Appendix H). The learner maximizes over the lattice. A maximum-weight ideal is computed by one minimum cut, and Gibbs domination lets maximum-weight ideals under Gumbel noise take the place of exponential weights, at the cost of a constant factor in the regret bound. These arguments use the isotonic class only through closure under pointwise maxima and minima and composition with nondecreasing maps. By Remark 1 the nonempty classes with these closures are the isotonic classes of preorders, reflexive and transitive relations, and identifying the points of a preorder that are related in both directions leaves a poset with the same isotonic class. For every such class, presented by its preorder, the minimax regret is determined by counting ideals and attained in polynomial time by maximizing over them.

AI USE STATEMENT

The authors used OpenAI's ChatGPT and Anthropic's Claude to help organize and polish the exposition of this paper, based on the authors' own drafts. The same tools assisted in surveying the related literature. The authors formulated and proved all mathematical claims in this paper, and checked, revised, and finalized every statement, proof, and reference. They take full responsibility for the content of the submission.

REPRODUCIBILITY STATEMENT

All results of this paper are mathematical. Section 2 states the model and the loss conditions, and the appendix contains a complete proof of every claim: the minimax rates (Appendices A and B), the phases (Appendix C), the guarantees of the learners (Appendices D to F), the extension to difference constraints (Appendix G) and the results for pinball loss (Appendix H). Algorithms 1 to 4 specify the learners in full, and Appendix I proves the polynomial-time claims in a finite-precision model.

REFERENCES

Jacob Abernethy, Chansoo Lee, Abhinav Sinha, and Ambuj Tewari. Online linear optimization via smoothing. In *Proceedings of the 27th Conference on Learning Theory*, volume 35 of *Proceedings of Machine Learning Research*, pp. 807–823. PMLR, 2014.

Jacob Abernethy, Chansoo Lee, and Ambuj Tewari. Fighting bandits with a new kind of smoothness. In *Advances in Neural Information Processing Systems 28*, 2015.

Jacob Abernethy, Chansoo Lee, and Ambuj Tewari. Perturbation techniques in online learning and optimization. In Tamir Hazan, George Papandreou, and Daniel Tarlow (eds.), *Perturbations, Optimization, and Statistics*, pp. 233–264. MIT Press, 2016. doi: 10.7551/mitpress/10761.003.0009.

Shivani Agarwal. Surrogate regret bounds for bipartite ranking via strongly proper losses. *Journal of Machine Learning Research*, 15(49):1653–1674, 2014.

Noga Alon, Omri Ben-Eliezer, Yuval Dagan, Shay Moran, Moni Naor, and Eylon Yogev. Adversarial laws of large numbers and optimal regret in online classification. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC 2021)*, pp. 447–455, 2021. doi: 10.1145/3406325.3451041.

Miriam Ayer, H. D. Brunk, G. M. Ewing, W. T. Reid, and Edward Silverman. An empirical distribution function for sampling with incomplete information. *The Annals of Mathematical Statistics*, 26(4):641–647, 1955. doi: 10.1214/aoms/1177728423.

R. E. Barlow, D. J. Bartholomew, J. M. Bremner, and H. D. Brunk. *Statistical Inference under Order Restrictions: The Theory and Application of Isotonic Regression*. Wiley Series in Probability and Mathematical Statistics. Wiley, 1972.

Alexander Barvinok and Alex Samorodnitsky. Random weighting, asymptotic counting, and inverse isoperimetry. *Israel Journal of Mathematics*, 158(1):159–191, 2007. doi: 10.1007/s11856-007-0008-8.

Shai Ben-David, Dávid Pál, and Shai Shalev-Shwartz. Agnostic online learning. In *Proceedings of the 22nd Annual Conference on Learning Theory (COLT 2009)*, 2009.

Michael Ben-Or and Avinatan Hassidim. The Bayesian learner is optimal for noisy binary search (and pretty good for quantum as well). In *49th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2008)*, pp. 221–230, 2008. doi: 10.1109/FOCS.2008.58.

Garrett Birkhoff. Rings of sets. *Duke Mathematical Journal*, 3(3):443–454, 1937. doi: 10.1215/S0012-7094-37-00334-X.

Adam Block, Yuval Dagan, and Alexander Rakhlin. Majorizing measures, sequential complexities, and online learning. In *Proceedings of the 34th Conference on Learning Theory*, volume 134 of *Proceedings of Machine Learning Research*, pp. 587–590. PMLR, 2021.

Richard P. Brent and Paul Zimmermann. *Modern Computer Arithmetic*, volume 18 of *Cambridge Monographs on Applied and Computational Mathematics*. Cambridge University Press, 2010. doi: 10.1017/CBO9780511921698.

H. D. Brunk. Maximum likelihood estimates of monotone parameters. *The Annals of Mathematical Statistics*, 26(4):607–616, 1955. doi: 10.1214/aoms/1177728420.

Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006. doi: 10.1017/CBO9780511546921.

Mark R. Conaway, Stephanie Dunbar, and Shyamal D. Peddada. Designs for single- or multiple-agent phase I trials. *Biometrics*, 60(3):661–669, 2004. doi: 10.1111/j.0006-341X.2004.00215.x.

Quanyu Dai, Jiaren Xiao, Zhaocheng Du, Jieming Zhu, Chengxiao Luo, Xiao-Ming Wu, and Zhenhua Dong. MCNet: monotonic calibration networks for expressive uncertainty calibration in online advertising. In *Proceedings of the ACM on Web Conference 2025 (WWW '25)*, pp. 4408–4419, 2025. doi: 10.1145/3696410.3714802.

Steven de Rooij, Tim van Erven, Peter D. Grünwald, and Wouter M. Koolen. Follow the leader if you can, hedge if you must. *Journal of Machine Learning Research*, 15(37):1281–1316, 2014.

Dariusz Dereniowski, Aleksander Łukasiewicz, and Przemysław Uznański. Noisy (binary) searching: simple, fast and correct. In *42nd International Symposium on Theoretical Aspects of Computer Science (STACS 2025)*, volume 327 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pp. 29:1–29:18, 2025. doi: 10.4230/LIPIcs.STACS.2025.29.

R. P. Dilworth. A decomposition theorem for partially ordered sets. *Annals of Mathematics*, 51(1): 161–166, 1950. doi: 10.2307/1969503.

Miroslav Dudík, Nika Haghtalab, Haipeng Luo, Robert E. Schapire, Vasilis Syrgkanis, and Jennifer Wortman Vaughan. Oracle-efficient online learning and auction design. *Journal of the ACM*, 67(5), 2020. doi: 10.1145/3402203.

Jack Edmonds and Richard M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *Journal of the ACM*, 19(2):248–264, 1972. doi: 10.1145/321694.321699.

Ehsan Emamjomeh-Zadeh, David Kempe, and Vikrant Singhal. Deterministic and probabilistic binary search in graphs. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2016)*, pp. 519–532, 2016. doi: 10.1145/2897518.2897656.

Uriel Feige, Prabhakar Raghavan, David Peleg, and Eli Upfal. Computing with noisy information. *SIAM Journal on Computing*, 23(5):1001–1018, 1994. doi: 10.1137/S0097539791195877.

D. R. Fulkerson. Note on Dilworth's decomposition theorem for partially ordered sets. *Proceedings of the American Mathematical Society*, 7(4):701–702, 1956. doi: 10.1090/S0002-9939-1956-0078334-6.

Pierre Gaillard and Sébastien Gerchinovitz. A chaining algorithm for online nonparametric regression. In *Proceedings of the 28th Conference on Learning Theory*, volume 40 of *Proceedings of Machine Learning Research*, pp. 764–796. PMLR, 2015.

Alfred Galichon and Bernard Salanié. Cupid's invisible hand: social surplus and identification in matching models. *The Review of Economic Studies*, 89(5):2600–2629, 2022. doi: 10.1093/restud/rdab090.

Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1321–1330. PMLR, 2017.

Maya Gupta, Andrew Cotter, Jan Pfeifer, Konstantin Voevodski, Kevin Canini, Alexander Mangylov, Wojciech Moczydlowski, and Alexander van Esbroeck. Monotonic calibrated interpolated look-up tables. *Journal of Machine Learning Research*, 17(109):1–47, 2016.

James Hannan. Approximation to Bayes risk in repeated play. In Melvin Dresher, Albert W. Tucker, and Philip Wolfe (eds.), *Contributions to the Theory of Games, Volume III*, volume 39 of *Annals of Mathematics Studies*, pp. 97–139. Princeton University Press, 1957. doi: 10.1515/9781400882151-006.

Elad Hazan and Tomer Koren. The computational power of optimization in online learning. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2016)*, pp. 128–141, 2016. doi: 10.1145/2897518.2897536.

Tamir Hazan and Tommi Jaakkola. On the partition function and random maximum a-posteriori perturbations. In *Proceedings of the 29th International Conference on Machine Learning (ICML 2012)*, pp. 991–998. Omnipress, 2012.

Alexander Henzi, Johanna F. Ziegel, and Tilmann Gneiting. Isotonic distributional regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 83(5):963–993, 2021. doi: 10.1111/rssb.12450.

Dorit S. Hochbaum and Maurice Queyranne. Minimizing a convex cost closure set. *SIAM Journal on Discrete Mathematics*, 16(2):192–207, 2003. doi: 10.1137/S0895480100369584.

Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. doi: 10.1080/01621459.1963.10500830.

Josef Hofbauer and William H. Sandholm. On the global convergence of stochastic fictitious play. *Econometrica*, 70(6):2265–2294, 2002. doi: 10.1111/1468-0262.00376.

Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform Chernoff bounds via nonnegative supermartingales. *Probability Surveys*, 17:257–317, 2020. doi: 10.1214/18-PS321.

Siguang Huang, Yunli Wang, Lili Mou, Huayue Zhang, Han Zhu, Chuan Yu, and Bo Zheng. MBCT: tree-based feature-aware binning for individual uncertainty calibration. In *Proceedings of the ACM Web Conference 2022 (WWW '22)*, pp. 2236–2246, 2022. doi: 10.1145/3485447.3512096.

Alexander I. Jordan, Anja Mühlemann, and Johanna F. Ziegel. Characterizing the optimal solutions to the isotonic regression problem for identifiable functionals. *Annals of the Institute of Statistical Mathematics*, 74(3):489–514, 2022. doi: 10.1007/s10463-021-00808-0.

Adam Kalai and Santosh Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005. doi: 10.1016/j.jcss.2004.10.016.

Richard M. Karp and Robert Kleinberg. Noisy binary search and its applications. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2007)*, pp. 881–890, 2007.

Jyrki Kivinen and Manfred K. Warmuth. Averaging expert predictions. In *Computational Learning Theory: 4th European Conference (EuroCOLT '99)*, volume 1572 of *Lecture Notes in Computer Science*, pp. 153–167. Springer, 1999. doi: 10.1007/3-540-49097-3_13.

Roger Koenker and Gilbert Bassett. Regression quantiles. *Econometrica*, 46(1):33–50, 1978. doi: 10.2307/1913643.

Wojciech Kotłowski, Wouter M. Koolen, and Alan Malek. Online isotonic regression. In *Proceedings of the 29th Annual Conference on Learning Theory*, volume 49 of *Proceedings of Machine Learning Research*, pp. 1165–1189. PMLR, 2016.

Wojciech Kotłowski, Wouter M. Koolen, and Alan Malek. Random permutation online isotonic regression. In *Advances in Neural Information Processing Systems 30*, 2017.

Ananya Kumar, Percy Liang, and Tengyu Ma. Verified uncertainty calibration. In *Advances in Neural Information Processing Systems 32*, 2019.

Ilja Kuzborskij and Nicolò Cesa-Bianchi. Locally-adaptive nonparametric online learning. In *Advances in Neural Information Processing Systems 33*, pp. 1679–1689, 2020.

Paul Liautaud, Pierre Gaillard, and Olivier Wintenberger. Minimax-optimal and locally-adaptive online nonparametric regression. In *Proceedings of the 36th International Conference on Algorithmic Learning Theory*, volume 272 of *Proceedings of Machine Learning Research*, pp. 702–735. PMLR, 2025a.

Paul Liautaud, Pierre Gaillard, and Olivier Wintenberger. Minimax adaptive online nonparametric regression over Besov spaces. In *Advances in Neural Information Processing Systems 38*, pp. 193125–193164, 2025b. doi: 10.52202/085713-5799.

Nathan Linial and Michael Saks. Every poset has a central element. *Journal of Combinatorial Theory, Series A*, 40(2):195–210, 1985. doi: 10.1016/0097-3165(85)90087-1.

Nick Littlestone. Learning quickly when irrelevant attributes abound: a new linear-threshold algorithm. *Machine Learning*, 2(4):285–318, 1988. doi: 10.1007/BF00116827.

Xingchao Liu, Xing Han, Na Zhang, and Qiang Liu. Certified monotonic neural networks. In *Advances in Neural Information Processing Systems 33*, pp. 15427–15438, 2020.

Zakaria Mhammedi and Robert C. Williamson. Constant regret, generalized mixability, and mirror descent. In *Advances in Neural Information Processing Systems 31*, 2018.

Jaouad Mourtada and Odalric-Ambrym Maillard. Efficient tracking of a growing number of experts. In *Proceedings of the 28th International Conference on Algorithmic Learning Theory*, volume 76 of *Proceedings of Machine Learning Research*, pp. 517–539. PMLR, 2017.

Alexandru Niculescu-Mizil and Rich Caruana. Predicting good probabilities with supervised learning. In *Proceedings of the 22nd International Conference on Machine Learning (ICML 2005)*, pp. 625–632, 2005. doi: 10.1145/1102351.1102430.

Robert D. Nowak. The geometry of generalized binary search. *IEEE Transactions on Information Theory*, 57(12):7893–7906, 2011. doi: 10.1109/TIT.2011.2169298.

Feiyang Pan, Xiang Ao, Pingzhong Tang, Min Lu, Dapeng Liu, Lei Xiao, and Qing He. Field-aware calibration: a simple and empirically strong method for reliable probabilistic predictions. In *Proceedings of The Web Conference 2020 (WWW '20)*, pp. 729–739, 2020. doi: 10.1145/3366423.3380154.

Jean-Claude Picard. Maximal closure of a graph and applications to combinatorial problems. *Management Science*, 22(11):1268–1272, 1976. doi: 10.1287/mnsc.22.11.1268.

J. Scott Provan and Michael O. Ball. The complexity of counting cuts and of computing the probability that a graph is connected. *SIAM Journal on Computing*, 12(4):777–788, 1983. doi: 10.1137/0212053.

Alexander Rakhlin and Karthik Sridharan. Online non-parametric regression. In *Proceedings of the 27th Conference on Learning Theory*, volume 35 of *Proceedings of Machine Learning Research*, pp. 1232–1264. PMLR, 2014.

Alexander Rakhlin and Karthik Sridharan. Online nonparametric regression with general loss functions. *arXiv preprint arXiv:1501.06598*, 2015.

Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. Online learning via sequential complexities. *Journal of Machine Learning Research*, 16(6):155–186, 2015.

Mark D. Reid, Rafael M. Frongillo, Robert C. Williamson, and Nishant Mehta. Generalized mixability via entropic duality. In *Proceedings of the 28th Conference on Learning Theory*, volume 40 of *Proceedings of Machine Learning Research*, pp. 1501–1522. PMLR, 2015.

Ronald L. Rivest, Albert R. Meyer, Daniel J. Kleitman, Karl Winklmann, and Joel Spencer. Coping with errors in binary search procedures. *Journal of Computer and System Sciences*, 20(3):396–404, 1980. doi: 10.1016/0022-0000(80)90014-8.

Tim Robertson, F. T. Wright, and R. L. Dykstra. *Order Restricted Statistical Inference*. Wiley Series in Probability and Mathematical Statistics. Wiley, 1988.

Davor Runje and Sharath M. Shankaranarayana. Constrained monotonic neural networks. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 29338–29353. PMLR, 2023.

Joseph Sill. Monotonic networks. In *Advances in Neural Information Processing Systems 10*, pp. 661–667, 1998.

Richard P. Stanley. Ordered structures and partitions. *Memoirs of the American Mathematical Society*, 119, 1972. doi: 10.1090/memo/0119.

Richard P. Stanley. *Enumerative Combinatorics, Volume 1*, volume 49 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, second edition, 2012. doi: 10.1017/CBO9781139058520.

J. Michael Steele. An Efron–Stein inequality for nonsymmetric statistics. *The Annals of Statistics*, 14(2):753–758, 1986. doi: 10.1214/aos/1176349952.

Ingo Steinwart and Andreas Christmann. Estimating conditional quantiles with the help of the pinball loss. *Bernoulli*, 17(1):211–225, 2011. doi: 10.3150/10-BEJ267.

Vasilis Syrgkanis, Akshay Krishnamurthy, and Robert E. Schapire. Efficient algorithms for adversarial contextual learning. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 2159–2168. PMLR, 2016.

Vladimir Vovk. A game of prediction with expert advice. *Journal of Computer and System Sciences*, 56(2):153–173, 1998. doi: 10.1006/jcss.1997.1556.

Vladimir Vovk. Metric entropy in competitive on-line prediction. *arXiv preprint arXiv:cs/0609045*, 2006.

Nolan A. Wages, Mark R. Conaway, and John O'Quigley. Continual reassessment method for partial ordering. *Biometrics*, 67(4):1555–1563, 2011. doi: 10.1111/j.1541-0420.2011.01560.x.

Guanghui Wang, Zihao Hu, Vidya Muthukumar, and Jacob Abernethy. Adaptive oracle-efficient online learning. In *Advances in Neural Information Processing Systems 35*, pp. 23398–23411, 2022. doi: 10.52202/068431-1700.

Sichen Wang. The minimax rate of online isotonic regression on product orders. In *Advances in Neural Information Processing Systems 39*, 2026.

Seungil You, David Ding, Kevin Canini, Jan Pfeifer, and Maya Gupta. Deep lattice networks and partial monotonic functions. In *Advances in Neural Information Processing Systems 30*, 2017.

Bianca Zadrozny and Charles Elkan. Transforming classifier scores into accurate multiclass probability estimates. In *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2002)*, pp. 694–699, 2002. doi: 10.1145/775047.775151.

A TOOLS

Each theorem of the main text is restated in the appendix, with its number, before its proof. This part collects the tools used throughout: bounds on the order polynomial and on the rate, rounding, exponential weights, concentration and level-set stationarity. It also states and proves Remarks 1 to 3.

In the appendix, H denotes Shannon entropy in nats, $\mathbf{1}\{\cdot\}$ the indicator, $(\cdot)_+$ the positive part, and $\mathrm{kl}(p\,\|\,p') = p\log(p/p') + (1-p)\log((1-p)/(1-p'))$ the Bernoulli divergence. The minimum in the definition of the regret (Section 2) is finite because the constant $1/2$ is isotonic. It is attained for every loss named in this paper, each being lower semicontinuous in the prediction, and is read as an infimum otherwise. For $f\colon P \to [0,1]$ and $z \in [0,1]$, the set $\{x : f(x) = z\}$ is the level set of f at z and $\{x : f(x) \geq z\}$ its superlevel set. Superlevel sets of isotonic functions are filters of P, and for an ideal I of Q_K and $1 \leq j \leq K$, $F_I(x) \geq j$ if and only if $(x,j) \in I$. A preorder is a reflexive and transitive relation, and its isotonic class is defined as for a poset.

Lemma 1 (Order-Polynomial Bounds)**.** *Let P be a finite poset with n points and width w, and let $K, K' \geq 1$.*

(a) Every isotonic function from an induced subposet of P into $\{0,\dots,K\}$ extends to an isotonic function on P.

(b) The following bounds hold:

$$\max\Big\{w\log(K+1),\ \log\tbinom{n+K}{n}\Big\} \leq H_P(K) \leq n\log(K+1),$$
$$n\log(1+K/n) \leq \log\tbinom{n+K}{n}, \qquad H_P(K+K') \leq H_P(K) + H_P(K').$$

(c) Let $\Psi_P(T)$ be defined for every real $T \geq 0$ by the formula of Section 2. The infimum defining $\Psi_P(T)$ is attained, $\Psi_P(T) \geq \frac{1}{4}\log(1+T)$, and $\Psi_P(\varrho T) \leq \lceil \varrho^{1/3} \rceil \Psi_P(T)$ for every real $\varrho \geq 1$. If $T \geq 1$, then every minimizer is at most $\lceil\sqrt{T}\rceil$ and $r_P(T) \geq 1$.

Proof. Extensions by maxima give (a), counting arguments give (b), and a comparison of the two terms of the objective gives (c). (a) Let Q be an induced subposet of P and F' an isotonic function from Q into $\{0,\dots,K\}$. Put $F(x) = \max(\{F'(x') : x' \in Q,\ x' \preceq x\} \cup \{0\})$. If $x \preceq x''$, the set over which $F(x)$ is a maximum is contained in the set for $F(x'')$, so F is isotonic, and for $x \in Q$ the maximum is $F'(x)$, because F' is isotonic on Q.

(b) By (a), the $(K+1)^w$ functions on a largest antichain have distinct isotonic extensions, and every isotonic function of a linear extension of P, a total order refining $\preceq$, is isotonic on P. A chain of n points has $\binom{n+K}{n}$ isotonic functions into $\{0,\dots,K\}$, the nondecreasing sequences of length n in $K+1$ values. Counting all functions gives the upper bound. Next, each of the n factors $(K+1)/1,\dots,(K+n)/n$ of $\binom{n+K}{n}$ is at least $1+K/n$. Finally, $F \mapsto (\min\{F,K\}, (F-K)_+)$ maps the isotonic functions $P \to \{0,\dots,K+K'\}$ injectively to pairs of isotonic functions into $\{0,\dots,K\}$ and $\{0,\dots,K'\}$, which proves subadditivity.

(c) Since $H_P(K) \geq \log(K+1) \to \infty$, the objective exceeds its value at $K=1$ for all large K, so the infimum is attained. To locate the minimizers, let A_j be the number of isotonic surjections from P onto a chain with j elements. An isotonic function into $\{0,\dots,K\}$ with exactly j distinct values is determined by such a surjection together with the choice of its j values, so

$$\Omega(P,K+1) = \sum_{j\geq 1} A_j \binom{K+1}{j}, \qquad \frac{\Omega(P,K+2)}{\Omega(P,K+1)} \geq \min_{1\leq j\leq K+1} \frac{\binom{K+2}{j}}{\binom{K+1}{j}} = \frac{K+2}{K+1},$$

because $\binom{K+2}{j}/\binom{K+1}{j} = (K+2)/(K+2-j)$ and the term $j = K+2$ is nonnegative. Hence $H_P(K+1) - H_P(K) \geq \log(1+1/(K+1)) \geq 1/(K+2)$, while for $K \geq \sqrt{T}$ the approximation term decreases by

$$\frac{T}{4K^2} - \frac{T}{4(K+1)^2} = \frac{T(2K+1)}{4K^2(K+1)^2} \leq \frac{2K+1}{4(K+1)^2} < \frac{2}{K+2}.$$

So the objective of $\Psi_P(T)$ increases strictly from K to $K+1$ whenever $K \geq \sqrt{T}$, and every minimizer is at most $\lceil\sqrt{T}\rceil$. Also $\Psi_P(T) \geq 2H_P(1) \geq 2\log 2 > 1$, so $r_P(T) \geq 1$ when $T \geq 1$.

For the logarithmic bound, for every $K \geq 1$,

$$\log(1+T) \leq \log(1+K^2) + \log(1+T/K^2) \leq 2\log(K+1) + \frac{T}{K^2} \leq 4\Big(2H_P(K) + \frac{T}{4K^2}\Big),$$

using $(1+K^2)(1+T/K^2) \geq 1+T$ and $\log(K+1) \leq H_P(K)$. Take the infimum over K. For the scaling inequality put $\bar{\varrho} = \lceil \varrho^{1/3} \rceil$. Subadditivity gives $H_P(\bar{\varrho}K) \leq \bar{\varrho} H_P(K)$, and $\varrho \leq \bar{\varrho}^3$ gives $\varrho T/(4\bar{\varrho}^2 K^2) \leq \bar{\varrho} T/(4K^2)$, so

$$2H_P(\bar{\varrho}K) + \frac{\varrho T}{4(\bar{\varrho}K)^2} \leq \bar{\varrho}\Big(2H_P(K) + \frac{T}{4K^2}\Big)$$

for every K, and the infimum over K proves the claim. □

Lemma 2 (Common-Shift Rounding). *Let $K \geq 1$ and $\upsilon \in [0,1)$. Common-shift rounding $\lfloor p \rceil_\upsilon = \lfloor Kp + \upsilon \rfloor / K$ (Section 3), applied pointwise to functions, is nondecreasing on $[0,1]$, takes values in G_K and fixes 0 and 1, so $\lfloor f \rceil_\upsilon \in \mathcal{F}_K$ for every $f \in \mathcal{F}_P$. If υ is uniform on $[0,1)$, then for every $p \in [0,1]$*

$$\mathbb{E}\lfloor p \rceil_\upsilon = p, \qquad \mathbb{E}(\lfloor p \rceil_\upsilon - p)^2 \leq \frac{1}{4K^2},$$

and if ℓ is β-smooth, then for every sequence $(x_t, y_t)_{t \leq T}$ and every $f \in \mathcal{F}_P$

$$\mathbb{E}\sum_{t \leq T} \ell(\lfloor f \rceil_\upsilon(x_t), y_t) \leq \sum_{t \leq T} \ell(f(x_t), y_t) + \frac{\beta T}{8K^2},$$

so that some shift υ satisfies this inequality without the expectation.

Proof. Unbiased two-point rounding cancels the linear term of the smoothness bound in expectation. The formula gives monotonicity, the range and the fixed endpoints. Fix p and let $j = \lfloor Kp \rfloor$. Then $\lfloor Kp + \upsilon \rfloor$ is $j+1$ if $\upsilon \geq 1 - (Kp - j)$ and j otherwise, so $\lfloor p \rceil_\upsilon$ equals $(j+1)/K$ with probability $Kp - j$ and j/K with probability $1 - (Kp - j)$. Its mean is p, and its variance $(Kp - j)(1 - (Kp - j))/K^2$ is at most $1/(4K^2)$. For a β-smooth loss and every y, $\ell(p', y) - \ell(p, y)$ differs from $p' - p$ times the derivative of $\ell(\cdot, y)$ at p by at most $\beta(p' - p)^2/2$, because that derivative is β-Lipschitz. With $p' = \lfloor p \rceil_\upsilon$ the derivative term has mean zero and the quadratic term has mean at most $\beta/(8K^2)$. Summing over t with $p = f(x_t)$ proves the last inequality in expectation, and a shift no worse than the average satisfies it. □

Lemma 3 (Exponential Weights over a Countable Expert Set). *Let ℓ be η-exp-concave, let $\mathcal{E}$ be a finite or countable set of experts with prior $\nu_E > 0$, $\sum_{E \in \mathcal{E}} \nu_E = 1$, and let expert E predict $p_{E,t} \in [0,1]$ at round t, with $\ell(p_{E,t}, y_t) < \infty$, where $p_{E,t}$ may depend on the past and on the expert's own randomness. Write $\mathrm{Loss}_E(t) = \sum_{s \leq t} \ell(p_{E,s}, y_s)$, and let the master predict*

$$\widehat{y}_t = \sum_{E \in \mathcal{E}} \nu_{E,t}\, p_{E,t}, \quad \nu_{E,t} = \frac{\nu_E \exp(-\eta\, \mathrm{Loss}_E(t-1))}{W_t}, \quad W_t = \sum_{E \in \mathcal{E}} \nu_E \exp(-\eta\, \mathrm{Loss}_E(t-1)).$$

Then, on every realization, for every $T \geq 1$ and every $E \in \mathcal{E}$,

$$\sum_{t \leq T} \ell(\widehat{y}_t, y_t) \leq \mathrm{Loss}_E(T) + \frac{\log(1/\nu_E)}{\eta}.$$

Proof. This is the telescoping argument of Cesa-Bianchi & Lugosi (2006). By exp-concavity and Jensen's inequality, at every round

$$e^{-\eta \ell(\widehat{y}_t, y_t)} \geq \sum_{E \in \mathcal{E}} \nu_{E,t}\, e^{-\eta \ell(p_{E,t}, y_t)} = \frac{W_{t+1}}{W_t}.$$

Multiplying over $t \leq T$ and using $W_1 = 1$ and $W_{T+1} \geq \nu_E \exp(-\eta\, \mathrm{Loss}_E(T))$,

$$\exp\Big(-\eta \sum_{t \leq T} \ell(\widehat{y}_t, y_t)\Big) \geq W_{T+1} \geq \nu_E \exp(-\eta\, \mathrm{Loss}_E(T)),$$

and taking logarithms proves the claim. □

The learners' sampled estimates (Appendices E and H) use the following bound.

Lemma 4 (Conditional Concentration). *Let $X_1, \dots, X_S \in [0,1]$ be conditionally independent given a sigma-field $\mathcal{H}$, with a common conditional mean $\mathbb{E}[X_1 \mid \mathcal{H}]$. Then for every $\varepsilon > 0$,*

$$\mathbb{P}\Big(\Big|\frac{X_1 + \cdots + X_S}{S} - \mathbb{E}[X_1 \mid \mathcal{H}]\Big| \geq \varepsilon \;\Big|\; \mathcal{H}\Big) \leq 2e^{-2S\varepsilon^2}.$$

Proof. This is Hoeffding's inequality (Hoeffding, 1963), applied to the conditional law given $\mathcal{H}$. □

Lemma 2 and Lemma 3 combine into the upper bound of Section 3.

Theorem 5 (Universal Upper Bound). *Let ℓ be η-exp-concave and β-smooth. For every finite poset P and every $T \geq 1$ there is a deterministic learner whose regret on every sequence of length T satisfies*

$$R_T \leq \inf_{K \geq 1} \Big\{ \frac{H_P(K)}{\eta} + \frac{\beta T}{8K^2} \Big\}.$$

For squared loss the right side is $\Psi_P(T)$, and the same learner also has $R_T \leq T$.

Proof. The learner runs exponential weights over the grid class at the best resolution and is compared with the comparator rounded into that class. Since $H_P(K) \geq \log(K+1)$, the objective tends to infinity with K and its infimum is attained at some K, which the learner uses. It runs exponential weights (Lemma 3) over the finite expert set $\mathcal{F}_K$, expert f predicting $f(x_t)$, with the uniform prior $1/|\mathcal{F}_K|$. Its predictions are deterministic, and its cumulative loss exceeds that of every $f \in \mathcal{F}_K$ by at most $\log|\mathcal{F}_K|/\eta = H_P(K)/\eta$. For every $f \in \mathcal{F}_P$, Lemma 2 gives a shift υ with $\lfloor f \rceil_\upsilon \in \mathcal{F}_K$ and $\sum_t \ell(\lfloor f \rceil_\upsilon(x_t), y_t) \leq \sum_t \ell(f(x_t), y_t) + \beta T/(8K^2)$. Combining the two bounds and taking the infimum over $f \in \mathcal{F}_P$ gives the regret bound. Squared loss is $\frac{1}{2}$-exp-concave and 2-smooth, so the objective is $2H_P(K) + T/(4K^2)$. The loss lies in $[0,1]$ while the comparator's loss is nonnegative, so $R_T \leq T$. □

Log loss is not smooth, and its rounding cost is computed on the level sets of the hindsight optimum, unqueried points included (Appendix B). That computation rests on the following stationarity property.

Lemma 5 (Level-Set Stationarity). *Let $\ell(\cdot, y)$ be finite and differentiable on $(0,1)$ for every $y \in [0,1]$, and write $\ell'(z,y)$ for its derivative at z. Let $(x_t, y_t)_{t \leq T}$ be a sequence and let $f^* \in \mathcal{F}_P$ attain the minimum of $\sum_{t \leq T} \ell(f(x_t), y_t)$ over $\mathcal{F}_P$. Then for every value $z \in (0,1)$ of f^*,*

$$\sum_{t:\, f^*(x_t) = z} \ell'(z, y_t) = 0.$$

Proof. Shift the complete level set of f^* at z. Let $\varepsilon > 0$ be smaller than the distance from z to 0, to 1 and to every other value of f^*. Then $f^* + \varepsilon \mathbf{1}\{f^* = z\}$ and $f^* - \varepsilon \mathbf{1}\{f^* = z\}$ take values in $[0,1]$ and are isotonic. Two points at level z keep equal values, and a point at level z and a point at another level keep their order, because the shift is smaller than the gap between the levels. The cumulative loss of a shifted function is $\sum_{t: f^*(x_t) = z} \ell(z \pm \varepsilon, y_t)$ plus the terms at the other levels, which do not change, so the function $\varepsilon \mapsto \sum_{t: f^*(x_t) = z} \ell(z + \varepsilon, y_t)$, finite and differentiable near 0, attains its minimum at $\varepsilon = 0$, and its derivative there vanishes. □

The next remark, used in Section 5, characterizes isotonic classes by three closure properties.

Remark 1 (Which Classes Are Isotonic). A nonempty class of functions $\mathcal{X} \to [0,1]$ on a finite set $\mathcal{X}$ is closed under pointwise maxima and minima and under monotone rescaling (composition with every nondecreasing map $[0,1] \to [0,1]$) if and only if it is the isotonic class of a preorder on $\mathcal{X}$. The proofs of this paper use no other property of isotonic functions.

Proof. If $\preceq'$ is a preorder on $\mathcal{X}$, then its isotonic class has the three properties, since maxima, minima and monotone rescalings of functions that are nondecreasing along $\preceq'$ are nondecreasing along $\preceq'$, and the constants belong to the class. For the converse, the superlevel sets of the members form a family closed under unions and intersections that contains $\emptyset$ and $\mathcal{X}$, hence the family of

filters of a preorder (Birkhoff, 1937), and monotone rescaling and pointwise maxima recover every function whose superlevel sets lie in the family. Let $\mathcal{F}$ be a nonempty class with the three closure properties, and let $\mathcal{U}$ be the family of sets $A \subseteq \mathcal{X}$ with $\mathbf{1}_A \in \mathcal{F}$. For $f \in \mathcal{F}$ and $z \in (0, 1]$ the indicator of $\{f \geq z\}$ is the composition of f with the nondecreasing map $\mathbf{1}\{\cdot \geq z\}$, so every superlevel set of every member lies in $\mathcal{U}$. Composing a member with the constant maps 0 and 1 puts $\emptyset$ and $\mathcal{X}$ in $\mathcal{U}$, and $\mathbf{1}_{A \cup A'} = \max\{\mathbf{1}_A, \mathbf{1}_{A'}\}$ and $\mathbf{1}_{A \cap A'} = \min\{\mathbf{1}_A, \mathbf{1}_{A'}\}$ make $\mathcal{U}$ closed under unions and intersections. Define $x \preceq' x'$ if every member of $\mathcal{U}$ that contains x contains x'. This is a preorder. Every member of $\mathcal{U}$ is a filter of $\preceq'$ by the definition of $\preceq'$. Conversely, every filter of $\preceq'$ is a member. For each x in the filter, the intersection A_x of the members of $\mathcal{U}$ containing x is a member, $A_x = \{x' : x \preceq' x'\}$ lies in the filter, and the filter is the union of these A_x. Thus $\mathcal{U}$ is the family of filters of $\preceq'$.

Every $f \in \mathcal{F}$ is nondecreasing along $\preceq'$, because its superlevel sets are filters of $\preceq'$. If $x \preceq' x'$ and $f(x) > 0$, then x' lies in the superlevel set of f at $f(x)$. Conversely, let f be nondecreasing along $\preceq'$, let z range over the values of f, and put $A_z = \{f \geq z\}$, a filter of $\preceq'$ and hence a member of $\mathcal{U}$. Then $f = \max_z z\, \mathbf{1}_{A_z}$, each $z\, \mathbf{1}_{A_z}$ is the monotone rescaling of $\mathbf{1}_{A_z} \in \mathcal{F}$ by multiplication with z, and the maximum lies in $\mathcal{F}$. Hence $\mathcal{F}$ is the isotonic class of $\preceq'$. □

The sequential covering numbers (Rakhlin et al., 2015; Rakhlin & Sridharan, 2014) are defined on trees. A query tree is a complete binary tree whose internal nodes are labeled by points of P, its depth being the number of internal nodes on a root-to-leaf path. A prediction tree on it labels its internal nodes by real numbers. A sequential ε-cover for the query tree is a finite set of prediction trees such that, for every $f \in \mathcal{F}_P$ and every root-to-leaf path, some prediction tree in the set is within ε of f at the point of every node on the path. The worst-case sequential covering number at a depth is the largest, over query trees of that depth, of the least size of a sequential ε-cover.

Remark 2 (Metric Entropy). For $0 < \varepsilon < 1/2$ put $K_+ = \lceil 1/(2\varepsilon) \rceil$, and let $\mathcal{N}(\varepsilon, \mathcal{F}_P)$ be the sup-norm covering number of $\mathcal{F}_P$, the least number of functions $P \to \mathbb{R}$ such that every member of $\mathcal{F}_P$ is within ε of one of them at every point. Then $H_P(K_+ - 1) \leq \log \mathcal{N}(\varepsilon, \mathcal{F}_P) \leq H_P(K_+)$, and the worst-case sequential covering number equals $\mathcal{N}(\varepsilon, \mathcal{F}_P)$ at every depth of at least n. Both covers may be improper, that is, the covering functions need not belong to $\mathcal{F}_P$.

Proof. Rounding to G_{K_+} gives a cover and the separation of $\mathcal{F}_{K_+ - 1}$ bounds the size of every cover from below. A sup-norm cover evaluated on any query tree is a sequential cover, and a query tree that enumerates P turns a sequential cover into a sup-norm cover. Since $\varepsilon < 1/2$, $K_+ \geq 2$. Rounding every value to the nearest level of G_{K_+}, ties upward, maps $\mathcal{F}_P$ into $\mathcal{F}_{K_+}$ and moves no value by more than $1/(2K_+) \leq \varepsilon$, so $\mathcal{F}_{K_+}$ is a cover and $\log \mathcal{N}(\varepsilon, \mathcal{F}_P) \leq H_P(K_+)$. Two distinct members of $\mathcal{F}_{K_+ - 1}$ differ at some point by at least $1/(K_+ - 1)$, which exceeds 2ε because $K_+ < 1/(2\varepsilon) + 1$, so no function is within ε of both, every cover has at least $|\mathcal{F}_{K_+ - 1}|$ members, and $\log \mathcal{N}(\varepsilon, \mathcal{F}_P) \geq H_P(K_+ - 1)$.

Evaluating each function of a sup-norm cover at the point of every node turns it into a prediction tree, and the tree of a function within ε of f is within ε of f on every path, so every query tree has a sequential cover of size $\mathcal{N}(\varepsilon, \mathcal{F}_P)$. Conversely, at any depth of at least n, take the query tree whose internal nodes at distance $0, \ldots, n-1$ from the root are labeled by the n points of P in a fixed order, the same along every path. The predictions of a prediction tree at the first n nodes of one fixed root-to-leaf path define a function on P, and for every f the function of a tree within ε of f along that path is within ε of f at every point. So the functions of the trees of a sequential cover form a sup-norm cover, and the sequential cover has at least $\mathcal{N}(\varepsilon, \mathcal{F}_P)$ trees. □

The logarithmic phase of Section 3 follows from the bounds of Lemma 1. Beyond a threshold depending on n alone, $r_P(T)$ lies within absolute constant factors of $n \log T$.

Remark 3 (The Logarithmic Phase). For every finite poset P with n points and every $T \geq \max\{2, n^4\}$,

$$0.1\, n \log T \leq r_P(T) \leq 8\, n \log T.$$

Proof. Evaluate the objective of $\Psi_P(T)$ at $K = \lceil \sqrt{T} \rceil$ for the upper bound, which holds for every $T \geq 2$, and bound it below at every K through $H_P(K) \geq n \log(1 + K/n)$ for the lower bound. At

$K = \lceil \sqrt{T} \rceil$, where $K + 1 \le \sqrt{T} + 2 \le (1 + \sqrt{2})\sqrt{T}$, Lemma 1(b) gives

$$r_P(T) \le \Psi_P(T) \le 2n \log(K+1) + \frac{T}{4K^2} \le n \log T + 2n \log(1 + \sqrt{2}) + \frac{1}{4} \le 8n \log T,$$

because $2 \log(1 + \sqrt{2}) < 1.77$ and $\log T \ge \log 2$. For the lower bound, $n \le T^{1/4}$ and $\log T \le 10T^{3/4}$ give $T \ge 0.1n \log T$. Fix $K \ge 1$ and put $\varsigma = 4nK^2/T$. By Lemma 1(b), $H_P(K) \ge n \log(1 + K/n) \ge n \log(K/n)$, and

$$2n \log \frac{K}{n} + \frac{T}{4K^2} = n\Big[\log \frac{T}{4n^3} + \log \varsigma + \frac{1}{\varsigma}\Big] \ge n\Big[1 + \log \frac{T}{4n^3}\Big] \ge n\Big[\frac{1}{4} \log T + 1 - \log 4\Big],$$

using $\log \varsigma + 1/\varsigma \ge 1$ and $n^3 \le T^{3/4}$. The right side is at least $0.1n \log T$ when $T \ge 16$, because $0.15 \log 16 > \log 4 - 1$, and when $2 \le T < 16$ the hypothesis forces $n = 1$, so that $2H_P(K) \ge 2 \log 2 > 0.1 \log T$. Taking the infimum over K gives $\Psi_P(T) \ge 0.1n \log T$, and the lower bound follows. □

B Minimax Rates

This part proves Theorems 1 and 2. The upper bounds for smooth losses are those of Theorem 5. The lower bound combines the correspondence between isotonic functions and ideals (Lemma 6), the entropy bound for one block (Lemma 7), the search at a fixed resolution (Proposition 1) and the choice of the resolution (Lemma 8). The upper bound for log loss rounds the hindsight optimum along its level sets (Lemmas 9 and 10). Lemma 11 supplies the rational grid on which the learners of Appendix E run log loss.

For an ideal I_+ and a filter I_- of Q_K with $I_+ \cap I_- = \emptyset$, the residual poset of (I_+, I_-) is the induced poset on $Q_K \setminus (I_+ \cup I_-)$. The down-closure of a set of elements of a poset is the smallest ideal containing it, and the up-closure the smallest filter.

Lemma 6 (Ideals and Residual Posets)**.** *Let $K \ge 1$.*

(a) $F \mapsto I_F$ is a bijection from the isotonic functions $P \to \{0, \dots, K\}$ onto $\mathcal{J}(Q_K)$, with inverse $F_I(x) = |I \cap (\{x\} \times [K])|$.

(b) Let I_+ be an ideal and I_- a filter of Q_K with $I_+ \cap I_- = \emptyset$. The ideals of Q_K containing I_+ and disjoint from I_- are exactly the sets $I_+ \cup J$ with J an ideal of the residual poset of (I_+, I_-), and $J \mapsto I_+ \cup J$ is a bijection. In particular, a uniformly random ideal of Q_K conditioned on containing I_+ and avoiding I_- is $I_+ \cup J$ with J uniform on the ideals of the residual poset.

(c) Call the ideals of (b) compatible, and let $x \in P$ and $0 \le j \le K$ be such that some compatible ideal I has $F_I(x) = j$. Let I'_+ be I_+ together with the down-closure, in the residual poset, of the cells $(x, 1), \dots, (x, j)$ that lie in it, and let I'_- be I_- together with the up-closure, in the residual poset, of the cells $(x, j+1), \dots, (x, K)$ that lie in it. Then I'_+ is an ideal, I'_- is a filter, $I'_+ \cap I'_- = \emptyset$, and the compatible ideals with $F_I(x) = j$ are exactly the ideals containing I'_+ and disjoint from I'_-, so that (b) applies to (I'_+, I'_-).

Part (b) makes the search of Section 3 possible. Conditioning on the revealed memberships of cells leaves the unrevealed part of a uniformly random ideal uniform on the ideals of a residual poset, and part (c) says how a revealed level enlarges the pair (I_+, I_-).

Proof. Reading an ideal column by column gives (a), and splitting an ideal into a fixed part and an ideal of the residual poset gives (b) and (c). (a) If $(x, j) \in I_F$ and $(x', j') \le (x, j)$, then $x \preceq x'$ and $j' \le j \le F(x) \le F(x')$, so $(x', j') \in I_F$, and I_F is an ideal. It meets the column of x in the $F(x)$ cells $(x, 1), \dots, (x, F(x))$, so F is recovered from I_F as stated. Conversely, let I be an ideal and put $F(x) = |I \cap (\{x\} \times [K])|$. Since $(x, j') \le (x, j)$ for $j' \le j$, the ideal meets the column of x in the cells $(x, 1), \dots, (x, F(x))$. Since $(x', j) \le (x, j)$ when $x \preceq x'$, membership of (x, j) forces membership of (x', j), so $F(x) \le F(x')$ and F is isotonic, and $I = I_F$.

(b) Let Q be the residual poset. If I is an ideal containing I_+ and disjoint from I_-, then $J = I \cap Q$ is an ideal of Q, because a cell of Q below a cell of J lies in I and not in I_+, and $I = I_+ \cup J$.

Conversely, let J be an ideal of Q, and let $q \in I_+ \cup J$ and $q' \le q$. If $q \in I_+$, then $q' \in I_+$. If $q \in J$, then $q' \notin I_-$, since I_- is a filter that avoids q. So $q' \in I_+$, or $q' \in Q$ and then $q' \in J$. Thus $I_+ \cup J$ is an ideal, it contains I_+, and it is disjoint from I_-. The two maps are inverse to each other, and a bijection carries the uniform law to the uniform law.

(c) Let Q be the residual poset, J_+ the down-closure in Q of the cells $(x,1),\dots,(x,j)$ that lie in Q, and J_- the up-closure in Q of the cells $(x,j+1),\dots,(x,K)$ that lie in Q, so that $I'_+ = I_+ \cup J_+$ and $I'_- = I_- \cup J_-$. By (b), I'_+ is an ideal of Q_K, and the dual argument shows that I'_- is a filter. Fix a compatible ideal I° with $F_{I^\circ}(x) = j$. It contains $(x,1),\dots,(x,j)$ and, being an ideal, every cell below them, so $I'_+ \subseteq I^\circ$. It avoids $(x,j+1),\dots,(x,K)$ and every cell above them, so $I^\circ \cap I'_- = \emptyset$. In particular $I'_+ \cap I'_- = \emptyset$, and the same argument shows that every compatible ideal with level j at x contains I'_+ and avoids I'_-. Conversely, let I be an ideal containing I'_+ and disjoint from I'_-. It is compatible. For $1 \le j' \le j$ the cell (x,j') lies in I_+, in Q or in I_-, and the last is excluded because I° contains (x,j'). In the first two cases $(x,j') \in I'_+ \subseteq I$. For $j < j' \le K$ the cell (x,j') lies in I_-, in Q or in I_+, and the last is excluded because I° avoids (x,j'). In the first two cases $(x,j') \in I'_-$, so $(x,j') \notin I$. Hence $F_I(x) = j$. □

The adversary of Section 3 plays in blocks of $32K^2$ rounds, and the next lemma bounds the expected excess loss of a block from below by a constant times the conditional entropy of the level it encodes. Conditional entropy given a sigma-field is the expectation of the entropy of the conditional law.

Lemma 7 (Regret from Conditional Entropy). *Let $K \ge 1$ and let ℓ be α-strongly proper on $[1/4, 3/4]$. Conditionally on a sigma-field $\mathcal{H}$, let $V \in \{0,\dots,K\}$ satisfy $\mathbb{P}(V \ge j \mid \mathcal{H}) \in [\delta_{\mathrm{LS}}, 1-\delta_{\mathrm{LS}}]$ for an $\mathcal{H}$-measurable level j, and put $p_V = 1/4 + V/(2K)$. Given $(V,\mathcal{H})$, draw $32K^2$ independent labels $Y_1,\dots,Y_{32K^2}$ with law $\mathrm{Ber}(p_V)$, and let $\widehat{y}_s$ be measurable with respect to $\mathcal{H}$ and $Y_1,\dots,Y_{s-1}$, where $\mathcal{H}$ may include private randomness. Then*

$$\sum_{s=1}^{32K^2} \mathbb{E}\big[\ell(\widehat{y}_s, Y_s) - \ell(p_V, Y_s)\big] \ge \frac{3\alpha}{32}\,\mathsf{H}(V \mid \mathcal{H}).$$

Proof. The expected excess loss of the block is at least a constant times the entropy the block removes from V, through the conditional variances of p_V, and the block removes at least half of $\mathsf{H}(V \mid \mathcal{H})$. Write $\mathcal{H}_s$ for the sigma-field generated by $\mathcal{H}$ and $Y_1,\dots,Y_{s-1}$, so that $\mathcal{H}_1 = \mathcal{H}$ and $\widehat{y}_s$ is $\mathcal{H}_s$-measurable, write $Y = (Y_1,\dots,Y_{32K^2})$ for the block, and put $\bar{p}_s = \mathbb{E}[p_V \mid \mathcal{H}_s] \in [1/4, 3/4]$.

Conditionally on $\mathcal{H}_s$ and V, the label Y_s has law $\mathrm{Ber}(p_V)$ while $\widehat{y}_s$ is fixed, so strong properness gives $\mathbb{E}[\ell(\widehat{y}_s, Y_s) - \ell(p_V, Y_s) \mid \mathcal{H}_s, V] \ge \alpha(\widehat{y}_s - p_V)^2$, and $\mathbb{E}[(\widehat{y}_s - p_V)^2 \mid \mathcal{H}_s] \ge \mathrm{Var}(p_V \mid \mathcal{H}_s)$ because $\widehat{y}_s$ is $\mathcal{H}_s$-measurable. Summing over the block,

$$\sum_{s \le 32K^2} \mathbb{E}\big[\ell(\widehat{y}_s, Y_s) - \ell(p_V, Y_s)\big] \ge \alpha \sum_{s \le 32K^2} \mathbb{E}\,\mathrm{Var}(p_V \mid \mathcal{H}_s). \tag{1}$$

The entropy that Y_s removes from V is the expected divergence between the conditional laws of Y_s given $(\mathcal{H}_s, V)$ and given $\mathcal{H}_s$, which are $\mathrm{Ber}(p_V)$ and $\mathrm{Ber}(\bar{p}_s)$:

$$\mathsf{H}(V \mid \mathcal{H}_s) - \mathsf{H}(V \mid \mathcal{H}_{s+1}) = \mathbb{E}\,\mathrm{kl}(p_V \,\|\, \bar{p}_s) \le \frac{16}{3}\,\mathbb{E}\,\mathrm{Var}(p_V \mid \mathcal{H}_s),$$

using $\mathrm{kl}(p \,\|\, p') \le (p-p')^2/(p'(1-p'))$ and $\bar{p}_s(1-\bar{p}_s) \ge 3/16$. Summing over s telescopes to

$$\mathsf{H}(V \mid \mathcal{H}) - \mathsf{H}(V \mid \mathcal{H}, Y) \le \frac{16}{3} \sum_{s \le 32K^2} \mathbb{E}\,\mathrm{Var}(p_V \mid \mathcal{H}_s). \tag{2}$$

For the entropy left after the block, Jensen's inequality applied to the logarithm bounds the entropy of any probability vector by twice the logarithm of the sum of the square roots of its entries. For levels $j_1, j_2 \in \{0,\dots,K\}$ write $p_{j_1} = 1/4 + j_1/(2K)$, the value of p_V at $V = j_1$, and let $\mathrm{Aff}(j_1, j_2)$ be the affinity of the laws of the block given $V = j_1$ and given $V = j_2$, the sum, over all label sequences in $\{0,1\}^{32K^2}$, of the square root of the product of the probabilities of the sequence under

the two laws. Applying the entropy bound to the conditional law of V given $(\mathcal{H}, Y)$, then Jensen's inequality for the expectation given $\mathcal{H}$, and expanding the square with Bayes' formula gives

$$\begin{aligned}\mathsf{H}(V \mid \mathcal{H}, Y) &\le \mathbb{E}\log\mathbb{E}\Big[\Big(\sum_{j_1}\sqrt{\mathbb{P}(V=j_1 \mid \mathcal{H}, Y)}\Big)^2 \;\Big|\; \mathcal{H}\Big] \\ &= \mathbb{E}\log\sum_{j_1,j_2}\sqrt{\mathbb{P}(V=j_1 \mid \mathcal{H})\,\mathbb{P}(V=j_2 \mid \mathcal{H})}\;\mathrm{Aff}(j_1,j_2).\end{aligned}$$

The two block laws are the $32K^2$-fold products of $\mathrm{Ber}(p_{j_1})$ and of $\mathrm{Ber}(p_{j_2})$, and affinities multiply over independent observations, so $\mathrm{Aff}(j_1, j_2)$ is the $32K^2$-th power of the affinity $\sqrt{p_{j_1}p_{j_2}} + \sqrt{(1-p_{j_1})(1-p_{j_2})}$ of the two Bernoulli laws. For $p, p' \in [0,1]$,

$$1 - \big(\sqrt{pp'} + \sqrt{(1-p)(1-p')}\big)^2 = \big(\sqrt{p(1-p')} - \sqrt{p'(1-p)}\big)^2 \ge (p-p')^2,$$

the last step because $p - p' = (\sqrt{p(1-p')} - \sqrt{p'(1-p)})(\sqrt{p(1-p')} + \sqrt{p'(1-p)})$ and the second factor is at most one by the Cauchy–Schwarz inequality. So the affinity of $\mathrm{Ber}(p)$ and $\mathrm{Ber}(p')$ is at most $\sqrt{1-(p-p')^2} \le e^{-(p-p')^2/2}$. Since $p_{j_1} - p_{j_2} = (j_1 - j_2)/(2K)$, this gives $\mathrm{Aff}(j_1, j_2) \le \exp(-32K^2(j_1-j_2)^2/(8K^2)) = e^{-4(j_1-j_2)^2}$. The diagonal terms of the double sum add to one, and for each difference $\varsigma = j_2 - j_1 \ge 1$ the Cauchy–Schwarz inequality gives $\sum_{j_1}\sqrt{\mathbb{P}(V=j_1 \mid \mathcal{H})\,\mathbb{P}(V=j_1+\varsigma \mid \mathcal{H})} \le 1$. Therefore

$$\mathsf{H}(V \mid \mathcal{H}, Y) \le \log\Big(1 + 2\sum_{\varsigma\ge 1} e^{-4\varsigma^2}\Big) \le \log\Big(1 + \frac{2e^{-4}}{1-e^{-12}}\Big) < \frac{1}{25}, \tag{3}$$

using $\varsigma^2 \ge 1 + 3(\varsigma - 1)$. The threshold condition gives $\mathsf{H}(V \mid \mathcal{H}) \ge \mathsf{H}(\mathbf{1}\{V \ge j\} \mid \mathcal{H}) \ge \mathsf{H}(\mathrm{Ber}(\delta_{\mathrm{LS}})) > 2/5$, because $\mathbf{1}\{V \ge j\}$ is a function of V given $\mathcal{H}$ and the binary entropy is at least $\mathsf{H}(\mathrm{Ber}(\delta_{\mathrm{LS}}))$ on $[\delta_{\mathrm{LS}}, 1-\delta_{\mathrm{LS}}]$. Hence, by (3), the left side of (2) is at least $\mathsf{H}(V \mid \mathcal{H}) - 1/25 \ge \mathsf{H}(V \mid \mathcal{H})/2$, and (1) gives

$$\sum_{s \le 32K^2} \mathbb{E}\big[\ell(\widehat{y}_s, Y_s) - \ell(p_V, Y_s)\big] \ge \frac{3\alpha}{16}\big(\mathsf{H}(V \mid \mathcal{H}) - \mathsf{H}(V \mid \mathcal{H}, Y)\big) \ge \frac{3\alpha}{32}\,\mathsf{H}(V \mid \mathcal{H}). \qquad \square$$

Lemma 7, applied along the search of Section 3, gives the lower bound at each resolution.

Proposition 1 (Fixed-Resolution Lower Bound)**.** *Let ℓ be α-strongly proper on $[1/4, 3/4]$. For every finite poset P, every $K \ge 1$ and every $T \ge 1$,*

$$R_T^*(P, \ell) \ge \frac{3\alpha\lambda}{2048}\min\Big\{H_P(K), \frac{T}{K^2}\Big\}.$$

Proof. For every learner, the expected regret against a random sequence drawn before play, independently of the learner's private randomness, is at most the supremum over fixed sequences of the expected regret. So it suffices to give a random sequence against which every learner has expected regret at least the stated bound. Fix a point $x_0 \in P$, to be queried in the rounds outside the search. When $T \ge 32K^2$ the sequence encodes a random ideal in Bernoulli labels, and the learner is told, after each block, the level the block encoded. A learner may ignore this information, so a lower bound for the learners so informed is a lower bound for all learners.

The instance for $T \ge 32K^2$. Draw F uniformly among the isotonic functions $P \to \{0, \dots, K\}$, so that I_F is uniform on $\mathcal{J}(Q_K)$, and put $g_F = 1/4 + F/(2K)$, an isotonic function with values in $[1/4, 3/4]$. Put $M = \lfloor T/(32K^2) \rfloor \ge 1$. The sequence consists of M blocks of $32K^2$ rounds and $T - 32K^2M$ further rounds, which query x_0 with labels of law $\mathrm{Ber}(g_F(x_0))$. Let A_ι be the set of ideals I of Q_K with $F_I(x) = F(x)$ at every point x queried in the blocks before block ι, so that $A_1 = \mathcal{J}(Q_K)$. By Lemma 6(c), applied once per block, A_ι is the set of ideals containing an ideal I_+ and disjoint from a filter I_-, and by Lemma 6(b) it is in bijection with the ideals of the residual poset of (I_+, I_-), which is nonempty exactly when $|A_\iota| \ge 2$. If $|A_\iota| \ge 2$, block ι takes a central element ζ_ι of the residual poset, chosen by a fixed rule, queries $x(\zeta_\iota)$ in each of its rounds with labels drawn independently from $\mathrm{Ber}(g_F(x(\zeta_\iota)))$, and then reveals $V_\iota = F(x(\zeta_\iota))$ to the learner. If $|A_\iota| = 1$, block ι queries x_0 with labels of law $\mathrm{Ber}(g_F(x_0))$ and reveals $V_\iota = F(x_0)$. The query

points are functions of the revealed levels, hence of F, and the labels are independent coins given F, so the whole sequence is drawn before play.

Uniformity at the block boundaries. For $1 \le \iota \le M+1$ let $\mathcal{H}_\iota$ be the sigma-field generated by $V_{<\iota}$, the labels of the first $\iota - 1$ blocks and the learner's private randomness. Conditionally on $\mathcal{H}_\iota$, the ideal I_F is uniform on A_ι. Indeed, the private randomness is independent of F and of the label coins. The query point of each block $\iota' < \iota$ is a function of $V_{<\iota'}$, and the mean of every label in block ι' is $1/4 + V_{\iota'}/(2K)$, whether the queried point is $x(\zeta_{\iota'})$ or x_0. Hence, for given values of the levels revealed before block ι and of the labels of the first $\iota - 1$ blocks, the probability that $I_F = I$ and that these values occur is $\Omega(P, K+1)^{-1}$ times a function of the given values alone if F_I takes the given levels at the queried points, and zero otherwise. The conditional law of I_F given $\mathcal{H}_\iota$ is therefore uniform on A_ι. In particular, the conditional law of V_ι given $\mathcal{H}_\iota$ is a function of $V_{<\iota}$, so that $\mathsf{H}(V_\iota \mid \mathcal{H}_\iota) = \mathsf{H}(V_\iota \mid V_{<\iota})$.

The entropy revealed. If $|A_\iota| \ge 2$ and $\zeta_\iota = (x, j)$, then, by Lemma 6(b), the conditional probability given $\mathcal{H}_\iota$ that $\zeta_\iota \in I_F$ is the fraction of ideals of the residual poset that contain ζ_ι, which lies in $[\delta_{\mathrm{LS}}, 1 - \delta_{\mathrm{LS}}]$, and $\zeta_\iota \in I_F$ if and only if $V_\iota \ge j$. Every value of V_ι is therefore compatible with at most a $1 - \delta_{\mathrm{LS}}$ fraction of A_ι, so $|A_{\iota+1}| \le (1 - \delta_{\mathrm{LS}})|A_\iota|$. When $|A_\iota| = 1$, $A_{\iota+1} = A_\iota$. Hence $\log |A_{M+1}| \le (H_P(K) - \lambda M)_+$. If every block queries a central element, then $|A_{M+1}| \le (1 - \delta_{\mathrm{LS}})^M \Omega(P, K+1)$, and otherwise $|A_{M+1}| = 1$. Since $V_1, \dots, V_M$ are functions of F, and I_F is uniform on A_{M+1} given them by the uniformity at $\iota = M+1$, the chain rule gives

$$\begin{aligned}\sum_{\iota \le M} \mathsf{H}(V_\iota \mid V_{<\iota}) &= \mathsf{H}(V_1, \dots, V_M) = \mathsf{H}(F) - \mathsf{H}(F \mid V_1, \dots, V_M) \\ &= H_P(K) - \mathbb{E} \log |A_{M+1}| \ge \min\{H_P(K), \lambda M\}.\end{aligned}$$

The regret. Since $g_F \in \mathcal{F}_P$, $R_T \ge \sum_{t \le T} [\ell(\widehat{y}_t, y_t) - \ell(g_F(x_t), y_t)]$. Conditionally on F, on the learner's private randomness and on the observations before round t, the label y_t has law $\mathrm{Ber}(g_F(x_t))$ with $g_F(x_t) \in [1/4, 3/4]$ while $\widehat{y}_t$ is fixed, so by strong properness every term has nonnegative conditional expectation. Fix a block ι. On the event $\{|A_\iota| \ge 2\}$, which lies in $\mathcal{H}_\iota$, the hypotheses of Lemma 7 hold with $\mathcal{H} = \mathcal{H}_\iota$, $V = V_\iota$ and $j = j(\zeta_\iota)$. The labels of the block are independent with law $\mathrm{Ber}(1/4 + V_\iota/(2K))$ given $(\mathcal{H}_\iota, V_\iota)$, and each prediction in the block is a function of $\mathcal{H}_\iota$ and of the earlier labels of the block. Apply Lemma 7 under the conditional probability given this event. On the complementary event, V_ι is determined by $V_{<\iota}$, so its conditional law given $\mathcal{H}_\iota$ is a point mass, and the block's expected excess loss is nonnegative. Averaging over the two events shows that the expected excess loss of the block is at least $(3\alpha/32)\mathsf{H}(V_\iota \mid \mathcal{H}_\iota)$. Summing over the blocks,

$$\mathbb{E} R_T \ge \frac{3\alpha}{32} \sum_{\iota \le M} \mathsf{H}(V_\iota \mid V_{<\iota}) \ge \frac{3\alpha}{32} \min\{H_P(K), \lambda M\} \ge \frac{3\alpha\lambda}{2048} \min\Big\{H_P(K), \frac{T}{K^2}\Big\},$$

the last step because $M \ge T/(64K^2)$ and $\lambda < 1$.

The instance for $T < 32K^2$. Draw p uniformly from $\{1/4, 3/4\}$, query x_0 in every round, and draw the labels independently from $\mathrm{Ber}(p)$ given p. The constant p is isotonic. The first prediction is independent of p, so that its expected excess loss over p is at least $\alpha \mathbb{E}(\widehat{y}_1 - p)^2 \ge \alpha \operatorname{Var}(p) = \alpha/16$, and the later rounds have nonnegative conditional expected excess as above. Since $\min\{H_P(K), T/K^2\} < 32$,

$$\mathbb{E} R_T \ge \frac{\alpha}{16} \ge \frac{\alpha}{512} \min\Big\{H_P(K), \frac{T}{K^2}\Big\} \ge \frac{3\alpha\lambda}{2048} \min\Big\{H_P(K), \frac{T}{K^2}\Big\},$$

the last step because $3\lambda < 4$. □

The next lemma shows that the best of the bounds of Proposition 1 over K is within a constant factor of the rate.

Lemma 8 (Crossing). *For every finite poset P and every $T \ge 1$,*

$$r_P(T) \le 3 \sup_{K \ge 1} \min\Big\{H_P(K), \frac{T}{K^2}\Big\}.$$

Proof. Evaluate Ψ_P at the largest resolution K at which $H_P(K) \le T/K^2$. If $H_P(1) > T$, the supremum is at least $\min\{H_P(1), T\} = T \ge r_P(T)$. Otherwise let $K_\star$ be the largest K with $H_P(K) \le T/K^2$, which exists because $H_P(K) \to \infty$. The supremum is at least $H_P(K_\star)$, and at least $T/(K_\star+1)^2 \ge T/(4K_\star^2)$ because $H_P(K_\star+1) > T/(K_\star+1)^2$. Hence $r_P(T) \le \Psi_P(T) \le 2H_P(K_\star) + T/(4K_\star^2)$ is at most three times the supremum. □

Theorem 1 (Minimax Rate under Squared Loss). *For every finite poset P and every integer $T \ge 1$,*

$$\frac{\lambda}{2048}\, r_P(T) \le R_T^*(P) \le r_P(T),$$

where $\lambda = -\log(1-\delta_{\mathrm{LS}}) = 0.1857\ldots$, so that $\lambda/2048 > 9.06 \cdot 10^{-5}$. The upper bound is attained by a deterministic learner.

Proof. Squared loss is 1-strongly proper, so Proposition 1 for every K and Lemma 8 give

$$R_T^*(P) \ge \frac{3\lambda}{2048} \sup_{K \ge 1} \min\Big\{ H_P(K), \frac{T}{K^2} \Big\} \ge \frac{\lambda}{2048}\, r_P(T).$$

Squared loss is also $\frac{1}{2}$-exp-concave and 2-smooth, so the deterministic learner of Theorem 5 has $R_T \le \min\{T, \Psi_P(T)\} = r_P(T)$ on every sequence. □

A rounding map is a nondecreasing map $\psi\colon [0,1] \to [0,1]$, and for $f \in \mathcal{F}_P$ the function $\psi \circ f$ is isotonic, with values in the range of ψ. For log loss the cost of rounding the hindsight optimum is computed level set by level set. On each, Lemma 5 and the finiteness of the minimum identify the mean of the labels with the level, and the cost is the number of queries there times one Bernoulli divergence.

Lemma 9 (Log Loss on a Level Set). *Let $(x_t, y_t)_{t \le T}$ be a sequence and let $f^* \in \mathcal{F}_P$ attain the minimum of the cumulative log loss over $\mathcal{F}_P$. For each value z of f^* at a queried point put $\mathcal{T}_z = \{t \le T : f^*(x_t) = z\}$. Then for every $z' \in (0,1)$,*

$$\sum_{t \in \mathcal{T}_z} y_t = z\,|\mathcal{T}_z| \qquad \text{and} \qquad \sum_{t \in \mathcal{T}_z} \big[\ell_{\log}(z', y_t) - \ell_{\log}(z, y_t)\big] = |\mathcal{T}_z|\ \mathrm{kl}(z \,\|\, z').$$

Consequently, for every rounding map ψ with values in $(0,1)$,

$$\sum_{t \le T} \ell_{\log}(\psi(f^*(x_t)), y_t) \le \min_{f \in \mathcal{F}_P} \sum_{t \le T} \ell_{\log}(f(x_t), y_t) + T \sup_{z \in [0,1]} \mathrm{kl}(z \,\|\, \psi(z)).$$

Proof. Compute the mean of the labels on each level set and substitute it into the loss difference. For $z \in (0,1)$, log loss is finite and differentiable on $(0,1)$ with derivative $\ell'_{\log}(z,y) = (z-y)/(z(1-z))$, so Lemma 5 gives $\sum_{t \in \mathcal{T}_z}(z - y_t) = 0$. For $z \in \{0,1\}$, the loss $\ell_{\log}(z,y)$ is infinite unless $y = z$, while the minimum is finite, so $y_t = z$ for every $t \in \mathcal{T}_z$. In both cases $\sum_{t \in \mathcal{T}_z} y_t = z|\mathcal{T}_z|$. For $z' \in (0,1)$, substituting the label sum into the loss difference gives

$$\sum_{t \in \mathcal{T}_z} \big[\ell_{\log}(z', y_t) - \ell_{\log}(z, y_t)\big] = |\mathcal{T}_z| \Big[z \log \frac{z}{z'} + (1-z) \log \frac{1-z}{1-z'} \Big] = |\mathcal{T}_z|\ \mathrm{kl}(z \,\|\, z'),$$

with $0 \log 0 = 0$ when $z \in \{0,1\}$. The function $\psi \circ f^*$ is isotonic with values in $(0,1)$, and summing the second identity with $z' = \psi(z)$ over the values z of f^* at queried points, where $\sum_z |\mathcal{T}_z| = T$, gives the last claim. □

Both rounding maps below are analyzed through the bound $\mathrm{kl}(p \,\|\, p') \le (p-p')^2/(p'(1-p'))$, valid for $p \in [0,1]$ and $p' \in (0,1)$ because the logarithm lies below its tangent at 1. Its right side is the chi-square divergence of $\mathrm{Ber}(p)$ from $\mathrm{Ber}(p')$, and we call it the chi-square bound. The first map below gives the grid of Theorem 2(iii), the second the rational grid of the learners.

Lemma 10 (Arcsine Rounding Map). *Let $K \ge 1$, put $\varpi_K = \pi/(4(K+1))$ and $z_j = \sin^2((2j+1)\varpi_K)$ for $0 \le j \le K$. The arcsine rounding map $\psi_K^{\mathrm{arc}}\colon [0,1] \to \{z_0, \ldots, z_K\}$ sends $p = \sin^2\phi$ with $0 \le \phi < \pi/2$ to z_j for the index j with $2j\varpi_K \le \phi < 2(j+1)\varpi_K$, and sends 1 to z_K. It is nondecreasing, and*

$$\sup_{p \in [0,1]} \mathrm{kl}\big(p \,\|\, \psi_K^{\mathrm{arc}}(p)\big) \le \frac{9\pi^2}{16(K+1)^2} < \frac{6}{K^2}.$$

Proof. Rounding the angle controls the chi-square bound uniformly, endpoints included. The map is nondecreasing because $p \mapsto \phi$, $\phi \mapsto j$ and $j \mapsto z_j$ are. Fix $p = \sin^2\phi$ with $\phi \in [0, \pi/2]$, let j be the index with $\psi_K^{\mathrm{arc}}(p) = z_j$, put $\phi_j = (2j+1)\varpi_K$, the midpoint of the interval $[2j\varpi_K, 2(j+1)\varpi_K]$, and $z = \sin^2\phi_j$. Then $|\phi - \phi_j| \le \varpi_K$ and $\varpi_K \le \phi_j \le \pi/2 - \varpi_K$. The identities $\sin^2\phi - \sin^2\phi_j = \sin(\phi - \phi_j)\sin(\phi + \phi_j)$ and $z(1-z) = \sin^2(2\phi_j)/4$ give

$$\mathrm{kl}(p \,\|\, z) \le \frac{(p-z)^2}{z(1-z)} = 4\sin^2(\phi - \phi_j)\Big(\frac{\sin(\phi + \phi_j)}{\sin(2\phi_j)}\Big)^2.$$

Since $\phi + \phi_j = 2\phi_j + (\phi - \phi_j)$, the ratio equals $\cos(\phi - \phi_j) + \cot(2\phi_j)\sin(\phi - \phi_j)$. It is nonnegative, as $0 \le \phi + \phi_j < \pi$, and since $|\cot(2\phi_j)| \le \cot(2\varpi_K)$ for $2\phi_j \in [2\varpi_K, \pi - 2\varpi_K]$ and $|\sin(\phi - \phi_j)| \le \sin\varpi_K$, it is at most

$$1 + \cot(2\varpi_K)\sin\varpi_K = 1 + \frac{\cos(2\varpi_K)}{2\cos\varpi_K} \le \frac{3}{2}.$$

Hence $\mathrm{kl}(p \,\|\, z) \le 9\sin^2(\phi - \phi_j) \le 9\varpi_K^2$, and $9\pi^2/16 < 6$. □

Lemma 11 (Rational Rounding Map)**.** *Let $K \ge 1$ and*

$$z_j = \frac{(2j+1)^2}{(2j+1)^2 + (2K+1-2j)^2}, \qquad 0 \le j \le K.$$

The levels are strictly increasing rationals of bit length $O(\log(K+1))$, with $z_0 = 1/(1+(2K+1)^2)$ and $z_K = 1 - z_0$. There is a nondecreasing map $\psi_K^{\mathrm{rat}} \colon [0,1] \to \{z_0, \dots, z_K\}$, the rational rounding map, with

$$\sup_{p \in [0,1]} \mathrm{kl}\big(p \,\|\, \psi_K^{\mathrm{rat}}(p)\big) \le \frac{9}{(K+1)^2}.$$

For every $y \in [0,1]$, the function $p \mapsto \ell_{\log}(p, y)$ is Lipschitz on $[z_0, z_K]$ with constant $1 + (2K+1)^2$.

Proof. Round the transformed value $\sqrt{p}/(\sqrt{p} + \sqrt{1-p})$ to the midpoint of an interval of length $1/(K+1)$, where the chi-square bound is uniformly controlled. Put $\chi(\varsigma) = \varsigma^2/(\varsigma^2 + (1-\varsigma)^2)$ for $\varsigma \in [0,1]$. It is increasing, $\chi(\varsigma) = p$ exactly when $\varsigma = \sqrt{p}/(\sqrt{p} + \sqrt{1-p})$, and $z_j = \chi(\varsigma_j)$ for $\varsigma_j = (j + \frac{1}{2})/(K+1)$, so the levels increase with j, and their numerators and denominators are integers at most $2(2K+1)^2$. For $p \in [0,1]$ let $\varsigma = \sqrt{p}/(\sqrt{p} + \sqrt{1-p})$, let j be the index with $j/(K+1) \le \varsigma < (j+1)/(K+1)$, or $j = K$ if $p = 1$, and put $\psi_K^{\mathrm{rat}}(p) = z_j$. This map is nondecreasing, and $|\varsigma - \varsigma_j| \le 1/(2(K+1)) \le \min\{\varsigma_j, 1 - \varsigma_j\}$. Factoring $\chi(\varsigma) - \chi(\varsigma_j)$ and using $z_j(1-z_j) = \varsigma_j^2(1-\varsigma_j)^2/(\varsigma_j^2 + (1-\varsigma_j)^2)^2$ gives

$$\frac{p - z_j}{\sqrt{z_j(1-z_j)}} = \frac{\varsigma - \varsigma_j}{\varsigma^2 + (1-\varsigma)^2}\Big(\frac{\varsigma}{\varsigma_j} + \frac{1-\varsigma}{1-\varsigma_j}\Big).$$

The bracket equals $2 + (\varsigma - \varsigma_j)(1 - 2\varsigma_j)/(\varsigma_j(1-\varsigma_j))$ and lies in $[1,3]$, because $|\varsigma - \varsigma_j|\,|1 - 2\varsigma_j| \le \min\{\varsigma_j, 1-\varsigma_j\}\max\{\varsigma_j, 1-\varsigma_j\}$, and $\varsigma^2 + (1-\varsigma)^2 \ge 1/2$. So $|p - z_j|/\sqrt{z_j(1-z_j)} \le 6|\varsigma - \varsigma_j| \le 3/(K+1)$, and the chi-square bound gives the divergence bound. For $p \in [z_0, 1 - z_0]$ and $y \in [0,1]$, $|\ell'_{\log}(p, y)| = |-y/p + (1-y)/(1-p)| \le \max\{1/p, 1/(1-p)\} \le 1/z_0 = 1 + (2K+1)^2$. □

Theorem 2 (Minimax Rate beyond Squared Loss)**.** *Let P be a finite poset and $T \ge 1$.*

(i) If ℓ is α-strongly proper on $[1/4, 3/4]$, then $R_T^*(P, \ell) \ge (\alpha\lambda/2048)\, r_P(T)$.

(ii) If ℓ is η-exp-concave, β-smooth and $0 \le \ell \le \Lambda$, then

$$R_T^*(P, \ell) \le \min\{\Lambda T,\ \max\{1/(2\eta), \beta/2\}\, \Psi_P(T)\}.$$

(iii) For log loss,

$$\frac{2\lambda}{2048}\, r_P(T) \le R_T^*(P, \ell_{\log}) \le \min\Big\{T\log 2,\ \inf_{K \ge 1}\Big[H_P(K) + \frac{6T}{K^2}\Big]\Big\} \le 24\, r_P(T).$$

Consequently $R_T^*(P,\ell) = \Theta(r_P(T))$ with absolute constants for log loss, and with constants depending only on $(\alpha, \eta, \beta, \Lambda)$ for every loss satisfying (i) and (ii).

Proof. The lower bounds combine the fixed-resolution search with the crossing lemma, and the upper bounds come from exponential weights over a grid with the rounding suited to the loss.

(i) By Proposition 1 for every K and Lemma 8,

$$R_T^*(P,\ell) \geq \frac{3\alpha\lambda}{2048} \sup_{K\geq 1} \min\Big\{H_P(K), \frac{T}{K^2}\Big\} \geq \frac{\alpha\lambda}{2048}\, r_P(T).$$

(ii) The learner of Theorem 5 satisfies

$$R_T \leq \inf_{K\geq 1}\Big\{\frac{H_P(K)}{\eta} + \frac{\beta T}{8K^2}\Big\} \leq \max\Big\{\frac{1}{2\eta}, \frac{\beta}{2}\Big\}\Psi_P(T),$$

the second inequality by comparing the two terms with those of $2H_P(K) + T/(4K^2)$. With $R_T \leq \Lambda T$, which holds for every learner, this is (ii).

(iii) Log loss is 2-strongly proper on $[1/4, 3/4]$, because $\mathbb{E}_{Y\sim\mathrm{Ber}(p)}[\ell_{\log}(p', Y) - \ell_{\log}(p, Y)] = \mathrm{kl}(p \,\|\, p') \geq 2(p-p')^2$ by Pinsker's inequality, so (i) with $\alpha = 2$ gives the lower bound. For the upper bound fix $K \geq 1$ and let $G = \{z_0, \ldots, z_K\}$ be the grid of Lemma 10. Its levels lie in $(0,1)$, so log loss is finite on G, and the integer form identifies $\mathcal{F}_G$ with the isotonic functions into $\{0,\ldots,K\}$, so $\log|\mathcal{F}_G| = H_P(K)$. Log loss is 1-exp-concave, and the learner that runs exponential weights at rate 1 over the expert set $\mathcal{F}_G$ with the uniform prior, expert f predicting $f(x_t)$, has cumulative loss at most that of every $f \in \mathcal{F}_G$ plus $H_P(K)$ (Lemma 3). Let f^* attain the minimum of the cumulative log loss over $\mathcal{F}_P$. By Lemmas 9 and 10, $\psi_K^{\mathrm{arc}} \circ f^*$ belongs to $\mathcal{F}_G$ and its cumulative loss exceeds the minimum by less than $6T/K^2$. So for every K there is a deterministic learner with $R_T \leq H_P(K) + 6T/K^2$ on every sequence, and the learner that always predicts $1/2$ has $R_T \leq T\log 2$. The best of these bounds is the upper bound of (iii). It is at most $24\, r_P(T)$, because $T\log 2 \leq 24T$ and $H_P(K) + 6T/K^2 \leq 24\,[2H_P(K) + T/(4K^2)]$ for every K.

The last sentence follows from (iii) for log loss, and from (i) and (ii) otherwise, since the minimum in (ii) is at most $\max\{\Lambda, 1/(2\eta), \beta/2\}\, r_P(T)$. □

C Families and Phases

This part computes the rate from the growth of H_P. Lemma 12 turns two-sided bounds on H_P into two-sided bounds on r_P, Proposition 2 bounds the exponent of a power-law phase, Proposition 3 realizes every exponent in $(0, 1/3)$, and Lemma 13, Corollaries 1 and 2, and Proposition 4 give the rows of Table 1. Throughout, $K \geq 1$ is a real number in the role of a resolution, and a subscript on O, Θ, c or C names the parameters on which the constants may depend. The product order $[m]^d$ is the set $\{1,\ldots,m\}^d$ ordered coordinatewise, and its row of Table 1 is the three-phase law of Wang (2026).

Table 1: The phases of $r_P(T)$ on three families of posets, where n is the number of points and, in the product order, $d \geq 2$ and $m \geq 2$. The rate is of order T in the linear phase and of the orders shown in the other two, with constants that are absolute for the chain and depend only on d for the product order and only on h for posets of height at most h.

	Linear phase	Intermediate phase	Logarithmic phase	
Poset	$T \leq$	$r_P(T)$	$T \geq$	$r_P(T)$
chain of n points	$1 + \log n$	$T^{1/3}[1 + \log(n^3/T)]^{2/3}$	n^3	$n[1 + \log(T/n^3)]$
product order $[m]^d$	m^{d-1}	$m^{2(d-1)/3}\, T^{1/3}$	m^{d+2}	$m^d[1 + \log(T/m^{d+2})]$
height at most h	n	—	n	$n[1 + \log(T/n)]$

Lemma 12 (Rate from the Growth of the Order Polynomial)**.** *Let P be a finite poset and $T \geq 1$. In each part, the constants of the conclusion depend only on those of the hypotheses.*

(a) *Let* $\bar{H}\colon [1,\infty) \to (0,\infty)$ *be nondecreasing, with* $\bar{H}(2\bar{K}) = O(\bar{H}(\bar{K}))$ *uniformly over* $\bar{K} \geq 1$*, and let* $H_P(K) = \Theta(\bar{H}(K))$ *for all integers* $K \geq 1$*. If a real* $\bar{K} \geq 1$ *satisfies* $T/\bar{K}^2 = \Theta(\bar{H}(\bar{K}))$*, then* $\Psi_P(T) = \Theta(\bar{H}(\bar{K}))$.

(b) *If* $T = O(H_P(1))$*, then* $r_P(T) = \Theta(T)$.

(c) *Let* $A \geq 1$ *and let* $K' \geq 1$ *be a resolution with* $H_P(K) = \Theta(A[1 + \log(K/K')])$ *for all integers* $K \geq K'$*. Then* $r_P(T) = \Theta(A[1 + \log(T/(AK'^2))])$ *for every* $T \geq AK'^2$.

Proof. One term of the objective bounds it from below at every resolution, and the matching upper bound comes from one resolution in (a) and (c), and from the cap T in (b).

(a) For an integer $K \leq \bar{K}$ the second term of the objective of $\Psi_P(T)$ is at least $T/(4\bar{K}^2)$, and for an integer $K \geq \bar{K}$ the first term, $2H_P(K)$, is at least a constant times $\bar{H}(K) \geq \bar{H}(\bar{K})$, so the objective is at least a constant times $\bar{H}(\bar{K})$ at every K. At $K = \lceil \bar{K} \rceil \leq 2\bar{K}$ the first term is at most a constant times $\bar{H}(2\bar{K})$, hence times $\bar{H}(\bar{K})$, and the second term is at most $T/(4\bar{K}^2)$, so the infimum is at most a constant times $\bar{H}(\bar{K})$.

(b) $\Psi_P(T) \geq 2H_P(1)$, which is at least a constant times T, so $r_P(T) = \min\{T, \Psi_P(T)\}$ lies between a constant times T and T.

(c) Put $\varrho = T/(AK'^2) \geq 1$. At $K = \lceil K'\sqrt{\varrho} \rceil$, which lies between K' and $2K'\sqrt{\varrho}$, the first term of the objective is at most a constant times $A[1 + \log(2\sqrt{\varrho})] \leq 2A(1 + \log \varrho)$, and the second is at most $T/(4K'^2\varrho) = A/4$, so $\Psi_P(T)$ is at most a constant times $A(1 + \log \varrho)$. Conversely, for an integer $K \leq K'$ the second term is at least $T/(4K'^2) = A\varrho/4 \geq A(1 + \log \varrho)/4$, and for an integer $K \geq K'$ the objective is at least a constant times

$$A\Big[1 + \log \frac{K}{K'} + \varrho\Big(\frac{K'}{K}\Big)^2\Big] \geq \frac{A}{2}(1 + \log \varrho),$$

because $\log \varrho = 2\log(K/K') + \log(\varrho(K'/K)^2) \leq 2\log(K/K') + \varrho(K'/K)^2$. Hence $\Psi_P(T) = \Theta(A(1 + \log \varrho))$, and since $A(1 + \log \varrho) \leq A\varrho \leq T$, the minimum with T leaves this order of magnitude unchanged. □

Proposition 2 (Exponent Cap and Width Constraint)**.** (a) *For every finite poset* P *and every* $T \geq H_P(1)$*,* $\Psi_P(T) \leq 5H_P(1)^{2/3}T^{1/3}$.

(b) *Let* P_N*,* $N \geq 1$*, be finite posets with* n_N *points, let* $\omega > 0$ *and* $c > 0$*, and let* $[T_N, T'_N]$ *be ranges of horizons with* $T_N/n_N \to \infty$ *and* $T'_N/T_N \to \infty$ *as* $N \to \infty$*. Suppose that* $R^*_{T'}(P_N) \geq c\,(T'/T)^\omega\, R^*_T(P_N)$ *whenever* $T_N \leq T \leq T' \leq T'_N$*. Then* $\omega \leq 1/3$ *and* $w(P_N)/n_N \to 0$.

Proof. (a) Balance the two terms of the objective using $H_P(K) \leq KH_P(1)$, the subadditivity of Lemma 1(b). Put $\bar{K} = (T/H_P(1))^{1/3} \geq 1$ and $K = \lceil \bar{K} \rceil \leq 2\bar{K}$. Then

$$\Psi_P(T) \leq 2KH_P(1) + \frac{T}{4K^2} \leq 4\bar{K}H_P(1) + \frac{T}{4\bar{K}^2} = \frac{17}{4}H_P(1)^{2/3}T^{1/3}.$$

(b) The scaling inequality of Lemma 1(c) caps the exponent, and a width proportional to the size would make the regret grow logarithmically across the range. For all large N, every $T \in [T_N, T'_N]$ satisfies $T \geq 2n_N$, hence $\Psi_{P_N}(T) \leq 2H_{P_N}(1) + T/4 \leq 2n_N \log 2 + T/4 \leq T$ by Lemma 1(b), and $r_{P_N}(T) = \Psi_{P_N}(T)$. For horizons $T \leq T'$ in $[T_N, T'_N]$ put $\varrho = T'/T$. By Theorem 1 and Lemma 1(c),

$$c\varrho^\omega R^*_T(P_N) \leq R^*_{T'}(P_N) \leq \Psi_{P_N}(\varrho T) \leq \lceil \varrho^{1/3} \rceil \Psi_{P_N}(T) \leq \frac{4096}{\lambda}\varrho^{1/3}R^*_T(P_N),$$

so $\varrho^{\omega - 1/3} \leq 4096/(c\lambda)$. At $T = T_N$ and $T' = T'_N$ the ratio $\varrho = T'_N/T_N$ tends to infinity with N, so $\omega \leq 1/3$.

If $w(P_N)/n_N$ did not tend to zero, then along a subsequence the width would be at least a fixed fraction of the size, and Lemma 1(b) would put $H_{P_N}(K)$ between fixed multiples of $n_N \log(K+1)$ for all $K \geq 1$. Then Lemma 12(b) gives $r_{P_N}(T) = \Theta(T)$ for $T \leq n_N$, Lemma 12(c) with $A = n_N$

and $K' = 1$ gives $r_{P_N}(T) = \Theta(n_N[1+\log(T/n_N)])$ for $T \geq n_N$, both are $\Theta(n_N \log(1+T/n_N))$, and Theorem 1 turns this into $R^*_T(P_N) = \Theta(n_N \log(1 + T/n_N))$ for all $T \geq 1$, with constants depending only on the fraction. For horizons $T \leq T'$ in $[T_N, T'_N]$, with $\varrho = T'/T$ and $T \geq n_N$, and with C the ratio of the upper to the lower constant of that bound,

$$\frac{R^*_{T'}(P_N)}{R^*_T(P_N)} \leq C \frac{\log(1 + \varrho T/n_N)}{\log(1 + T/n_N)} \leq C\Big(1 + \frac{\log(2\varrho)}{\log 2}\Big),$$

which grows more slowly than the lower bound $c\varrho^\omega$ as $\varrho = T'_N/T_N \to \infty$. This contradiction proves $w(P_N)/n_N \to 0$. □

For a chain the order polynomial is a binomial coefficient, and for integers $n, K \geq 1$,

$$\frac{\log 2}{2} \leq \frac{\log \binom{n+K}{n}}{\min\{n,K\}\,[1 + \log(\max\{n,K\}/\min\{n,K\})]} \leq 2. \tag{4}$$

By the symmetry of $\binom{n+K}{n}$ in n and K, assume $n \leq K$. Each of the factors $(K+1)/1, \ldots, (K+n)/n$ of $\binom{n+K}{n}$ is at least $1+K/n$, so the logarithm is at least $n\log(1+K/n)$, and $\log(1+K/n) \geq \max\{\log 2, \log(K/n)\} \geq \frac{\log 2}{2}[1+\log(K/n)]$. One term of the binomial expansion of $(n+K)^{n+K}$ gives $\binom{n+K}{n} \leq (n+K)^{n+K}/(n^n K^K)$, which bounds the logarithm by $n\log(1+K/n)+K\log(1+n/K) \leq n[1+\log(2K/n)] \leq 2n[1+\log(K/n)]$.

A poset is connected if any two of its points are joined by a sequence of points each comparable to the next.

Proposition 3 (Connected Families with Every Exponent)**.** *Let $0 < \omega < 1/3$. There are connected posets P_N, $N = 2, 4, 8, \ldots$, with $n_N = \Theta_\omega(N)$ points and width $\Theta_\omega(N^{(1-3\omega)/(1-\omega)})$, such that $H_{P_N}(K) = \Theta_\omega\big(N^{(1-3\omega)/(1-\omega)} K^{2\omega/(1-\omega)}\big)$ for $1 \leq K \leq N$ and*

$$R^*_T(P_N) = \Theta_\omega\big(N^{1-3\omega} T^\omega\big) \qquad \textit{for } N^2 \leq T \leq N^{5/2}.$$

Proof. The poset is a union of chains of geometrically increasing lengths, in numbers chosen so that H_P grows like $K^{2\omega/(1-\omega)}$, with a minimum element adjoined for connectedness. Put $\varsigma = 2\omega/(1-\omega) \in (0,1)$, so that $\omega = \varsigma/(2+\varsigma)$, $(1-3\omega)/(1-\omega) = 1-\varsigma$ and $1-3\omega = 2(1-\varsigma)/(2+\varsigma)$. Let P°_N be the disjoint union, over $0 \leq i \leq \log_2 N$, of $A_i = \lfloor (N/2^i)^{1-\varsigma} \rfloor$ chains of 2^i points, and let P_N be P°_N with a new minimum element adjoined, which makes it connected. Since $1 \leq (N/2^i)^{1-\varsigma} \leq 2A_i$, the size $n_N = 1 + \sum_i A_i 2^i$ lies between $N/2$ and $1 + N2^\varsigma/(2^\varsigma - 1)$, by summing the geometric series $\sum_i 2^{\varsigma i}$. The width of P_N is the number of chains, $\sum_i A_i$, which lies between $\frac{1}{2}N^{1-\varsigma}$ and $N^{1-\varsigma}/(1-2^{-(1-\varsigma)})$.

The isotonic functions on a disjoint union are the tuples of isotonic functions on its parts, so for $1 \leq K \leq N$

$$H_{P^\circ_N}(K) = \sum_{i \leq \log_2 N} A_i \log \binom{2^i + K}{2^i}.$$

Put $i_K = \lfloor \log_2 K \rfloor$, so that $K/2 < 2^{i_K} \leq K$. For $i \leq i_K$, (4) with $2^i \leq K$ gives $\log \binom{2^i+K}{2^i} \leq 2 \cdot 2^i[1 + \log(K/2^i)] \leq 4 \cdot 2^i(1 + i_K - i)$, and $A_i \leq (N/2^i)^{1-\varsigma}$, so

$$A_i \log \binom{2^i + K}{2^i} \leq 4N^{1-\varsigma} 2^{\varsigma i}(1 + i_K - i) \leq 4N^{1-\varsigma} K^\varsigma (1 + i_K - i) 2^{-\varsigma(i_K - i)};$$

for $i > i_K$, (4) with $K < 2^i$ gives $\log \binom{2^i+K}{2^i} \leq 2K[1 + \log(2^i/K)] \leq 2K(1 + i - i_K)$, and $A_i \leq (N/2^{i_K})^{1-\varsigma} 2^{-(1-\varsigma)(i-i_K)} \leq (2N/K)^{1-\varsigma} 2^{-(1-\varsigma)(i-i_K)}$, so

$$A_i \log \binom{2^i + K}{2^i} \leq 4N^{1-\varsigma} K^\varsigma (1 + i - i_K) 2^{-(1-\varsigma)(i-i_K)}.$$

Since $\sum_{r \geq 0}(1+r)2^{-\varsigma r}$ and $\sum_{r \geq 0}(1+r)2^{-(1-\varsigma)r}$ are finite, summing over i gives $H_{P^\circ_N}(K) \leq C_\omega N^{1-\varsigma} K^\varsigma$. For the lower bound keep $i = i_K$ only. Then $A_{i_K} \geq \frac{1}{2}(N/2^{i_K})^{1-\varsigma} \geq \frac{1}{2}(N/K)^{1-\varsigma}$, and $\log \binom{2^{i_K}+K}{2^{i_K}} \geq 2^{i_K} \log(1 + K/2^{i_K}) \geq \frac{K \log 2}{2}$, so $H_{P^\circ_N}(K) \geq \frac{\log 2}{4} N^{1-\varsigma} K^\varsigma$.

An isotonic function on P_N is specified by its value j at the minimum and by the isotonic function $P_N^\circ \to \{0,\dots,K-j\}$ obtained on the rest by subtracting j. This bijection gives $\Omega(P_N, K+1) = \sum_{j=0}^{K} \Omega(P_N^\circ, K-j+1)$, which lies between $\Omega(P_N^\circ, K+1)$ and $(K+1)\Omega(P_N^\circ, K+1)$. Since $\log(K+1) \le K \le N^{1-\varsigma}K^\varsigma$ for $K \le N$, this gives $H_{P_N}(K) = \Theta_\omega(N^{1-\varsigma}K^\varsigma)$ for $1 \le K \le N$.

Let $N^2 \le T \le N^{5/2}$ and put $\bar{K} = (T/N^{1-\varsigma})^{1/(2+\varsigma)}$, so that $N^{1-\varsigma}\bar{K}^\varsigma = T/\bar{K}^2$. The bounds on T give $1 \le \bar{K} \le N$. As in the proof of Lemma 12(a), for an integer $K \le \bar{K}$ the objective of $\Psi_{P_N}(T)$ is at least $T/(4\bar{K}^2)$. For $\bar{K} \le K \le N$ it is at least $2H_{P_N}(K) \ge c_\omega N^{1-\varsigma}\bar{K}^\varsigma$, and for $K > N$ it is at least $2H_{P_N}(N) \ge c_\omega N \ge c_\omega N^{1-\varsigma}\bar{K}^\varsigma$. At $K = \lceil \bar{K} \rceil \le N$ it is at most $2H_{P_N}(\lceil \bar{K} \rceil) + T/(4\bar{K}^2) \le C_\omega N^{1-\varsigma}\bar{K}^\varsigma$. Hence $\Psi_{P_N}(T) = \Theta_\omega(N^{1-\varsigma}\bar{K}^\varsigma)$, the minimum with T leaves this order of magnitude unchanged because $N^{1-\varsigma}\bar{K}^\varsigma = T/\bar{K}^2 \le T$, and Theorem 1 gives $R_T^*(P_N) = \Theta_\omega(N^{1-\varsigma}\bar{K}^\varsigma)$. Substituting $\bar{K}$ and ς,

$$N^{1-\varsigma}\bar{K}^\varsigma = N^{(1-\varsigma)(1-\varsigma/(2+\varsigma))}\,T^{\varsigma/(2+\varsigma)} = N^{2(1-\varsigma)/(2+\varsigma)}\,T^{\varsigma/(2+\varsigma)} = N^{1-3\omega}\,T^{\omega}. \qquad \square$$

On the ranges $[N^2, N^{5/2}]$ the family meets the hypotheses of Proposition 2(b), and its width is of smaller order than its size, as that proposition requires.

On product orders, the growth of H_P changes from linear to logarithmic at the resolution m.

Lemma 13 (Order Polynomial of a Product Order)**.** *Let $d \ge 2$, $m \ge 2$ and $P = [m]^d$. Then $H_P(K) = \Theta_d(m^{d-1}K)$ for $1 \le K \le m$ and $H_P(K) = \Theta_d(m^d[1+\log(K/m)])$ for $K \ge m$.*

Proof. Upper bounds come from counting ideals and from encoding quotients and remainders, lower bounds from choosing values freely on sets of points of equal rank. Write $\mathrm{rk}(x) = \sum_{i \le d}(x^{(i)} - 1)$ for the rank of a point $x = (x^{(1)},\dots,x^{(d)})$. Points of equal rank are incomparable and form a rank level. A rank level has at most m^{d-1} points, because $d-1$ components determine the last, and there are at most dm rank levels, so the rank levels with fewer than $m^{d-1}/(2d)$ points contain fewer than $m^d/2$ points in all, and more than $m/2$ rank levels have at least $m^{d-1}/(2d)$ points each.

Binary isotonic functions are the indicators of filters, whose complements are the ideals, so $H_P(1) = \log|\mathcal{J}([m]^d)|$. An ideal of $[m]^d$ is determined by its intersections with the m copies of $[m]^{d-1}$ on which the last component is constant, each an ideal of $[m]^{d-1}$, and $[m]^2$ has $\binom{2m}{m} \le 4^m$ ideals, one for each nonincreasing sequence of m intersection sizes in $\{0,\dots,m\}$. By induction on d, $H_P(1) \le (2\log 2)m^{d-1}$, and subadditivity (Lemma 1(b)) gives $H_P(K) \le (2\log 2)m^{d-1}K$ for every K. For $K \le m/2$, choose K rank levels with at least $m^{d-1}/(2d)$ points each and list them as $\mathcal{A}_1,\dots,\mathcal{A}_K$ in increasing order of rank. For arbitrary subsets $\mathcal{S}_j \subseteq \mathcal{A}_j$, $1 \le j \le K$, let $F(x)$ be the number of chosen ranks below $\mathrm{rk}(x)$, plus one if $x \in \mathcal{S}_1 \cup \dots \cup \mathcal{S}_K$. Along $x \prec x'$ the rank increases, so the number of chosen ranks below it does not decrease, and if $x \in \mathcal{A}_j$ then $F(x) \le j \le F(x')$. Hence F is isotonic, with values in $\{0,\dots,K\}$. Since $\mathcal{S}_j = \{x \in \mathcal{A}_j : F(x) = j\}$, distinct choices of the subsets give distinct functions, so $H_P(K) \ge \log 2 \sum_j |\mathcal{A}_j| \ge \frac{\log 2}{2d} m^{d-1}K$. For $m/2 < K \le m$, the monotonicity of H_P and the bound at $\lfloor m/2 \rfloor \ge m/3$ give $H_P(K) \ge \frac{\log 2}{6d} m^{d-1}K$.

For $K \ge m$, an isotonic F into $\{0,\dots,K\}$ is determined by the isotonic function $\lfloor F/\lceil K/m \rceil \rfloor$, with values in $\{0,\dots,m\}$, and by the remainder $F - \lceil K/m \rceil \lfloor F/\lceil K/m \rceil \rfloor$, a function into $\{0,\dots,\lceil K/m \rceil - 1\}$. Hence $H_P(K) \le H_P(m) + m^d \log\lceil K/m \rceil \le 3m^d[1+\log(K/m)]$. For $K \ge 4dm$ put $\varphi = \lfloor (K+1)/(d(m-1)+1) \rfloor$, which is at least $K/(2dm) \ge 2$. For every function $\chi\colon P \to \{0,\dots,\varphi-1\}$ the function $\varphi\,\mathrm{rk} + \chi$ is isotonic, because the rank increases by at least one along $x \prec x'$ while χ changes by less than φ, and its values are at most $\varphi(d(m-1)+1) - 1 \le K$. Distinct χ give distinct functions. Hence

$$H_P(K) \ge m^d \log\varphi \ge m^d \log\frac{K}{2dm} \ge \frac{\log 2}{1+\log(4d)}\, m^d\Big[1 + \log\frac{K}{m}\Big],$$

the last step because $(\log(K/m) - \log(2d))/(1+\log(K/m))$ increases in K and equals $\log 2/(1+\log(4d))$ at $K = 4dm$. For $m \le K < 4dm$, the bound at $K = m$ gives $H_P(K) \ge H_P(m) \ge \frac{\log 2}{6d} m^d \ge \frac{\log 2}{6d(1+\log(4d))} m^d[1+\log(K/m)]$. $\square$

Corollary 1 (Three Phases on Product Orders)**.** *Let $d \ge 2$, $m \ge 2$ and $P = [m]^d$. Then $r_P(T) = \Theta_d(T)$ for $T \le m^{d-1}$, $r_P(T) = \Theta_d(m^{2(d-1)/3}T^{1/3})$ for $m^{d-1} \le T \le m^{d+2}$, and $r_P(T) = \Theta_d(m^d[1+\log(T/m^{d+2})])$ for $T \ge m^{d+2}$.*

Proof. Apply Lemma 12 to the bounds of Lemma 13. Let $\bar{H}(\bar{K}) = m^{d-1}\bar{K}$ for $1 \le \bar{K} \le m$ and $\bar{H}(\bar{K}) = m^d[1 + \log(\bar{K}/m)]$ for $\bar{K} \ge m$. It is continuous and nondecreasing, $H_P(K) = \Theta_d(\bar{H}(K))$ by Lemma 13, and $\bar{H}(\bar{K})/\bar{K}$ is nonincreasing, being m^{d-1} up to m and $m^d[1 + \log(\bar{K}/m)]/\bar{K}$ beyond, so $\bar{H}(2\bar{K}) \le 2\bar{H}(\bar{K})$. For $T \le m^{d-1}$, $H_P(1) \ge \frac{\log 2}{6d} m^{d-1}$ gives $T = O_d(H_P(1))$ and Lemma 12(b) applies. For $m^{d-1} \le T \le m^{d+2}$, the real resolution $\bar{K} = (T/m^{d-1})^{1/3}$ lies in $[1, m]$ and satisfies $\bar{H}(\bar{K}) = m^{d-1}\bar{K} = T/\bar{K}^2$, so Lemma 12(a) gives $\Psi_P(T) = \Theta_d(m^{d-1}\bar{K}) = \Theta_d(m^{2(d-1)/3}T^{1/3})$, and $m^{2(d-1)/3}T^{1/3} = T/\bar{K}^2 \le T$, so the minimum with T leaves this order of magnitude unchanged. For $T \ge m^{d+2}$, Lemma 12(c) with $A = m^d$ and $K' = m$ applies. □

Corollary 2 (The Chain)**.** *Let P be a chain of n points. With absolute constants, $r_P(T) = \Theta(T)$ for $T \le 1 + \log n$, $r_P(T) = \Theta\big(T^{1/3}[1 + \log(n^3/T)]^{2/3}\big)$ for $1 + \log n \le T \le n^3$, and $r_P(T) = \Theta(n[1 + \log(T/n^3)])$ for $T \ge n^3$.*

Proof. Apply Lemma 12 to the binomial estimate (4). Let $\bar{H}(\bar{K}) = \bar{K}[1 + \log(n/\bar{K})]$ for $1 \le \bar{K} \le n$ and $\bar{H}(\bar{K}) = n[1 + \log(\bar{K}/n)]$ for $\bar{K} \ge n$. By (4), $H_P(K) = \Theta(\bar{H}(K))$. Both $\bar{H}$ and $\bar{H}(\bar{K})/\bar{K}$ are continuous at n. Below n their derivatives are $\log(n/\bar{K}) \ge 0$ and $-1/\bar{K}$, and above n they are $n/\bar{K}$ and $-n\log(\bar{K}/n)/\bar{K}^2 \le 0$, so $\bar{H}$ is nondecreasing, $\bar{H}(\bar{K})/\bar{K}$ is nonincreasing, and $\bar{H}(2\bar{K}) \le 2\bar{H}(\bar{K})$. For $T \le 1+\log n$, $H_P(1) = \log(n+1) \ge \frac{\log 2}{2}(1+\log n)$ gives $T = O(H_P(1))$ and Lemma 12(b) applies. For $1+\log n \le T \le n^3$, the function $\bar{K} \mapsto \bar{K}^3[1+\log(n/\bar{K})]$ increases on $[1, n]$ from $1 + \log n$ to n^3, its derivative being $\bar{K}^2[2 + 3\log(n/\bar{K})]$. Let $\bar{K} \in [1, n]$ be the point where it equals T. Then $\bar{H}(\bar{K}) = \bar{K}[1 + \log(n/\bar{K})] = T/\bar{K}^2$, so Lemma 12(a) gives $\Psi_P(T) = \Theta(T/\bar{K}^2) = \Theta(T^{1/3}[1 + \log(n/\bar{K})]^{2/3})$. Since $\log(n^3/T) = 3\log(n/\bar{K}) - \log(1 + \log(n/\bar{K}))$ lies between $2\log(n/\bar{K})$ and $3\log(n/\bar{K})$, $1 + \log(n/\bar{K}) = \Theta(1 + \log(n^3/T))$, and $T/\bar{K}^2 \le T$, so the minimum with T leaves this order of magnitude unchanged. For $T \ge n^3$, Lemma 12(c) with $A = K' = n$ applies. □

Proposition 4 (Bounded Height)**.** *Let P have n points and height at most h. Then $\frac{n}{h}\log(K+1) \le H_P(K) \le n\log(K+1)$ for every $K \ge 1$, and $r_P(T) = \Theta_h(n\log(1 + T/n))$ for every $T \ge 1$. In particular $r_P(T) = \Theta_h(T)$ for $T \le n$ and $r_P(T) = \Theta_h(n[1 + \log(T/n)])$ for $T \ge n$.*

Proof. Assigning to each point the number of points of a longest chain ending at it partitions P into at most h antichains, so $w \ge n/h$, and Lemma 1(b) gives the bounds on H_P. Hence $H_P(K) = \Theta_h(n[1 + \log K])$ for $K \ge 1$, and the rate follows as in the proof of Proposition 2(b), by Lemma 12(b) for $T \le n$, using $T \le n \le hH_P(1)/\log 2$, and Lemma 12(c) with $A = n$ and $K' = 1$ for $T \ge n$. On each range the result is $\Theta_h(n\log(1 + T/n))$. □

D The Perturbed Potential

This part establishes the facts about the perturbed potential that Section 4 uses: the hardness of exact exponential weights, the minimum-cut computation of a perturbed maximizer, the range of the potential (Lemma 14) with the counting estimate it yields (Remark 4), Theorem 3, and the bound on the mix loss by the decrease of the potential (Lemma 19).

Throughout this part, Q_K, the minimum chain partition $\mathcal{C}_1, \dots, \mathcal{C}_{w_K}$, the noise ξ, the score arrays θ and the potential Φ are those of Definition 3, and within each chain the cells are indexed in increasing order. The potential is finite, because $|\max_I\{\theta(I)+\xi(I)\}| \le \max_I|\theta(I)|+\sum_{(i,k)\in\Xi}|\xi_{i,k}|$, and the right side is integrable. The maximizer is unique with probability one. Two distinct ideals $I \ne I'$ have different chain encodings, so $k_i(I) \ne k_i(I')$ for some chain i. Conditionally on all noise coordinates other than $\xi_{i,k_i(I)}$, the difference of the perturbed scores of I and I' is $\xi_{i,k_i(I)}$ plus a constant and vanishes with probability zero. There are finitely many pairs.

The hardness of exact exponential weights. On a general poset, even the first prediction of exponential weights over the isotonic functions with values in a grid is #P-hard to compute exactly (Section 4). Take a bipartite graph with m' right vertices, let Q be the poset on its vertices in which each left vertex lies below its neighbors on the right, and let P be Q with one new point added above every left vertex. The ideals of P that avoid the new point are the ideals of Q, and those that

contain it consist of the new point, all left vertices and an arbitrary set of right vertices, so under the uniform law on the ideals of P, the new point is absent with probability $|\mathcal{J}(Q)|/(|\mathcal{J}(Q)| + 2^{m'})$. At resolution one and with the uniform prior, the first prediction of exponential weights at the new point is the probability that a uniformly random binary isotonic function takes the value 1 there. Binary isotonic functions are the indicators of filters, and the complement of a filter is an ideal, so this prediction is the absence probability above and determines $|\mathcal{J}(Q)|$, and counting the ideals of such posets is #P-complete (Provan & Ball, 1983).

Node weights and the minimum cut. Fix θ and ξ, and give the cell $q = (x, j)$ that is the k-th cell of $\mathcal{C}_i$ the node weight $\mathrm{wt}(q) = \theta_{x,j} - \theta_{x,j-1} + \xi_{i,k} - \xi_{i,k-1}$. Since $(x, j) \in I$ exactly when $j \le F_I(x)$ and the k-th cell of $\mathcal{C}_i$ lies in I exactly when $k \le k_i(I)$, telescoping along columns and along chains gives, for every ideal I,

$$\theta(I) + \xi(I) = \sum_{x \in P} \theta_{x,0} + \sum_i \xi_{i,0} + \sum_{q \in I} \mathrm{wt}(q). \tag{5}$$

A perturbed maximizer is therefore an ideal of maximum weight $\sum_{q \in I} \mathrm{wt}(q)$, and the reduction of Picard (1976) finds one by a minimum cut. The network has a source, a sink and one vertex per cell, with an arc from the source to q of capacity $\mathrm{wt}(q)$ for every cell with $\mathrm{wt}(q) > 0$, an arc from q to the sink of capacity $-\mathrm{wt}(q)$ for every cell with $\mathrm{wt}(q) < 0$, and an arc $q \to q'$ of capacity $1 + \sum_q |\mathrm{wt}(q)|$ for every cell q and every cell q' that q covers. Let A be the set of cells on the source side of a cut. If A is not an ideal, some cell of A lies above a cell outside A, a chain of covers joins them, and its first arc leaving A crosses the cut, whose capacity then exceeds the capacity $\sum_{\mathrm{wt}(q)>0} \mathrm{wt}(q)$ of the cut with $A = \emptyset$. If A is an ideal, the cut consists of the arcs from the source to the cells $q \notin A$ with $\mathrm{wt}(q) > 0$ and the arcs to the sink from the cells $q \in A$ with $\mathrm{wt}(q) < 0$, so its capacity is $\sum_{\mathrm{wt}(q)>0} \mathrm{wt}(q) - \sum_{q \in A} \mathrm{wt}(q)$. A minimum cut therefore has an ideal of maximum weight on its source side. The network has $nK + 2$ vertices and at most $2nK + Kn^2/2$ arcs, because a cell (x, j) covers $(x, j-1)$ when $j \ge 2$ and (x', j) when x' covers x in P, and P has at most $\binom{n}{2}$ cover pairs.

The Gumbel maximum. Two facts about independent centered standard Gumbel variables are used below. Let φ_k be real numbers indexed by a nonempty set of prefix lengths k of a chain $\mathcal{C}_i$. Then $\max_k\{\varphi_k + \xi_{i,k}\} - \log \sum_k e^{\varphi_k}$ is a centered standard Gumbel variable, and the maximum is attained at k with probability $e^{\varphi_k} / \sum_{k'} e^{\varphi_{k'}}$. Indeed, each $e^{-\xi_{i,k} - \gamma}$ is a standard exponential variable, so the variables $e^{-\varphi_k - \xi_{i,k} - \gamma}$ are independent and exponential with rates e^{φ_k}. Their minimum is attained at k with probability $e^{\varphi_k} / \sum_{k'} e^{\varphi_{k'}}$ and equals a standard exponential variable divided by $\sum_k e^{\varphi_k}$. The maximum of the $\varphi_k + \xi_{i,k}$ is attained at the same k and equals $-\gamma$ minus the logarithm of this minimum, that is, $\log \sum_k e^{\varphi_k}$ plus a centered standard Gumbel variable.

Lemma 14 (Range of the Perturbed Potential)**.** *Let $K \ge 1$ and let Φ be the perturbed potential of a minimum chain partition of Q_K. For every score array θ,*

$$\log \sum_{I \in \mathcal{J}(Q_K)} e^{\theta(I)} \le \Phi(\theta),$$

so that $\Phi(\theta) \ge \theta(I)$ for every ideal I. At zero scores,

$$H_P(K) \le \Phi(0) \le c_\Phi \, H_P(K), \qquad c_\Phi = 2 + \frac{\log \pi - \gamma}{\log 2} < 2.82.$$

Proof. The lower bound on $\Phi(\theta)$ comes from a random ideal with the Gibbs distribution, the law ν on $\mathcal{J}(Q_K)$ with $\nu(I)$ proportional to $e^{\theta(I)}$, built by Gumbel maxima one chain at a time so that its expected noise equals the entropy of ν. The upper bound on $\Phi(0)$ comes from an exponential moment. Since $\log \nu(I) = \theta(I) - \log \sum_J e^{\theta(J)}$, the entropy of ν satisfies

$$\sum_I \nu(I)\, \theta(I) + \mathsf{H}(\nu) = \log \sum_I e^{\theta(I)}.$$

Select $\hat{k}_1, \hat{k}_2, \ldots, \hat{k}_{w_K}$ in turn. Given $\hat{k}_1, \ldots, \hat{k}_{i-1}$, let $\nu_i(k)$ be the conditional probability under ν that $k_i(I) = k$ given $k_{i'}(I) = \hat{k}_{i'}$ for all $i' < i$, and let $\hat{k}_i$ attain $\max_k\{\log \nu_i(k) + \xi_{i,k}\}$ over

the k with $\nu_i(k) > 0$. The noise of the earlier chains determines $\hat{k}_1, \dots, \hat{k}_{i-1}$ and is independent of the noise of chain i. By the Gumbel maximum, $\hat{k}_i$ therefore has conditional law ν_i, and the conditional expectation of $\log \nu_i(\hat{k}_i) + \xi_{i,\hat{k}_i}$ is $\log \sum_k \nu_i(k) = 0$. Hence the tuple $(\hat{k}_1, \dots, \hat{k}_{w_K})$ has positive probability under ν, so it is the chain encoding of an ideal $\hat{I}$ with law ν, and $\mathbb{E}\xi(\hat{I}) = -\mathbb{E}\sum_i \log \nu_i(\hat{k}_i) = -\mathbb{E}\log \nu(\hat{I}) = \mathsf{H}(\nu)$. Therefore

$$\Phi(\theta) \geq \mathbb{E}\big[\theta(\hat{I}) + \xi(\hat{I})\big] = \sum_I \nu(I)\,\theta(I) + \mathsf{H}(\nu) = \log \sum_I e^{\theta(I)},$$

and the right side is at least $\theta(I)$ for every ideal I, and at $\theta = 0$ it equals $H_P(K)$.

For the upper bound, $e^{-\xi_{i,k}-\gamma}$ is a standard exponential variable, whose moment of order $-1/2$ is $\Gamma(1/2) = \sqrt{\pi}$, so $\mathbb{E}e^{\xi_{i,k}/2} = \sqrt{\pi}\,e^{-\gamma/2}$. By Jensen's inequality, the bound of a maximum of nonnegative variables by their sum, and the independence of the noise of different chains,

$$\Phi(0) \leq 2\log \mathbb{E}\max_I e^{\xi(I)/2} \leq 2\log \sum_I \prod_i \mathbb{E}e^{\xi_{i,k_i(I)}/2} = 2H_P(K) + w_K(\log\pi - \gamma).$$

A maximum antichain of Q_K has w_K cells, and distinct subsets of it have distinct down-closures, which are ideals, so $w_K \log 2 \leq H_P(K)$, and since $\log\pi > \gamma$ the zero-score value $\Phi(0)$ is at most $c_\Phi H_P(K)$. □

The zero-score value $\Phi(0)$ is estimated by sampling. The learner for pinball loss sets its update rate from such an estimate (Appendix H), and Remark 4 turns it into an approximation of $H_P(K)$.

Lemma 15 (Reference-Subtracted Sampling)**.** *The zero-score sample*

$$\widehat{\Phi} = \max_{I \in \mathcal{J}(Q_K)} \sum_i \big(\xi_{i,k_i(I)} - \xi_{i,0}\big)$$

is the maximum weight of an ideal for the node weights $\xi_{i,k} - \xi_{i,k-1}$ on the k-th cell of $\mathcal{C}_i$, hence one minimum cut computes it, and it satisfies $\widehat{\Phi} \geq 0$, $\mathbb{E}\widehat{\Phi} = \Phi(0)$ and $\operatorname{Var}\widehat{\Phi} \leq 3|\Xi| \leq 6nK$.

Proof. The identity (5) gives the first claims, and the Efron–Stein inequality, applied to one noise coordinate at a time, gives the variance bound. With $\theta = 0$, (5) gives $\widehat{\Phi} = \max_I \sum_{q\in I} \mathrm{wt}(q)$ for the stated weights. The empty ideal has weight zero, so $\widehat{\Phi} \geq 0$. Since $\widehat{\Phi} = \max_I \xi(I) - \sum_i \xi_{i,0}$ and the noise is centered, $\mathbb{E}\widehat{\Phi} = \mathbb{E}\max_I \xi(I) = \Phi(0)$. The sample is square-integrable, since $0 \leq \widehat{\Phi} \leq \sum_{(i,k)\in\Xi} |\xi_{i,k}|$, so the Efron–Stein inequality (Steele, 1986) applies: if $\widehat{\Phi}^{(i,k)}$ denotes $\widehat{\Phi}$ with the coordinate $\xi_{i,k}$ replaced by an independent copy $\xi'_{i,k}$, then $\operatorname{Var}\widehat{\Phi} \leq \frac{1}{2}\sum_{(i,k)\in\Xi} \mathbb{E}(\widehat{\Phi} - \widehat{\Phi}^{(i,k)})^2$. Increasing $\xi_{i,k}$ with $k \geq 1$ by $\varsigma > 0$ raises the value $\sum_i(\xi_{i,k_i(I)} - \xi_{i,0})$ of the ideals with $k_i(I) = k$ by ς and leaves the others unchanged, so $\widehat{\Phi}$ rises by an amount in $[0, \varsigma]$. Increasing $\xi_{i,0}$ by ς lowers the value of the ideals with $k_i(I) \geq 1$ by ς and leaves the others unchanged, so $\widehat{\Phi}$ changes by an amount in $[-\varsigma, 0]$. In both cases $|\widehat{\Phi} - \widehat{\Phi}^{(i,k)}| \leq |\xi_{i,k} - \xi'_{i,k}|$, and a standard Gumbel variable has variance $\pi^2/6 < 3$. Hence $\operatorname{Var}\widehat{\Phi} \leq \frac{1}{2}\sum_{(i,k)\in\Xi} 2 \cdot \frac{\pi^2}{6} < 3|\Xi| \leq 6nK$. □

Remark 4 (Approximating the Logarithm of the Order Polynomial). Let P be a finite poset, $K \geq 1$, $\varepsilon \in (0, 1]$ and $\delta \in (0, 1)$, and let the zero-score estimate $\widetilde{\Phi}$ be ε plus a median of $2\lceil\log_4(1/\delta)\rceil$ batch means of $\lceil 96nK/\varepsilon^2\rceil$ zero-score samples each, all samples independent. It takes $O(nK\varepsilon^{-2}\log(2/\delta))$ minimum cuts and satisfies $\Phi(0) \leq \widetilde{\Phi} \leq \Phi(0) + 2\varepsilon$ with probability at least $1 - \delta$. At $\varepsilon = 1/2500$ these bounds give $H_P(K) \leq \widetilde{\Phi} \leq 2.82\,H_P(K)$.

Proof. Chebyshev's inequality controls each batch mean, and a union bound controls the median. By Lemma 15, a batch mean of $\lceil 96nK/\varepsilon^2\rceil$ independent zero-score samples has expectation $\Phi(0)$ and variance at most $6nK/\lceil 96nK/\varepsilon^2\rceil \leq \varepsilon^2/16$, so it differs from $\Phi(0)$ by more than ε with probability at most $1/16$. Put $A = \lceil\log_4(1/\delta)\rceil$. If the median of the $2A$ independent batch means differs from $\Phi(0)$ by more than ε, then at least A of them do. Independence and a union bound over the $\binom{2A}{A} \leq 4^A$ sets of A batches bound its probability by $4^A 16^{-A} = 4^{-A} \leq \delta$. Otherwise $\widetilde{\Phi}$, the

median plus ε, lies in $[\Phi(0), \Phi(0) + 2\varepsilon]$. The $2A\lceil 96nK/\varepsilon^2 \rceil$ samples take one minimum cut each, $O(nK\varepsilon^{-2}\log(2/\delta))$ in all. At $\varepsilon = 1/2500$, Lemma 14 gives, on the same event,

$$H_P(K) \le \Phi(0) \le \widetilde{\Phi} \le c_\Phi H_P(K) + \frac{1}{1250} \le 2.82\, H_P(K),$$

because $H_P(K) \ge \log 2$ and $(2.82 - c_\Phi)\log 2 > 1/1250$. □

This estimate has a precedent. The sets $\{(i, k_i(I)) : i \le w_K\}$ of noise coordinates, one per chain, form a family of $\Omega(P, K+1) \ge 2^{w_K}$ sets of size w_K, over which one minimum cut maximizes any sum of weights on the coordinates. For such families the random-weighting principle of Barvinok & Samorodnitsky (2007), with independent logistic weights, already approximates $H_P(K)$, the logarithm of the number of sets, within a constant factor in polynomial time. Remark 4 is the same principle with Gumbel weights, with the factor 2.82 and an explicit number of cuts.

Four results prove Theorem 3: a bound on the growth of the probability that the maximizer lies in a set of ideals when the scores of its members are raised (Lemma 16); the curvature bound it implies for the potential (Lemma 17); an exponential change of variables that turns such a curvature bound into the domination inequality (Theorem 6); and an exact computation for levels determined by one prefix length, which covers resolution one (Lemma 18).

A set of ideals is determined by the prefix lengths in a set of chains if whether an ideal belongs to it depends only on the ideal's prefix lengths in those chains. A cylinder on at most L chains is a set of ideals determined by the prefix lengths in at most L chains.

Lemma 16 (Cylinder Translation). *Let θ be a score array and $\mathcal{Z}$ a cylinder on at most L chains, and for $\vartheta \in \mathbb{R}$ let $\mu_{\mathcal{Z}}(\vartheta)$ be the probability that the maximizer of $\theta(I) + \vartheta\mathbf{1}\{I \in \mathcal{Z}\} + \xi(I)$ over the ideals lies in $\mathcal{Z}$. Let $\vartheta \in \mathbb{R}$ and $\varsigma \ge 0$. With probability one, the maximizers at ϑ and at $\vartheta + \varsigma$ coincide unless the first lies outside $\mathcal{Z}$ and the second in $\mathcal{Z}$, and*

$$\mu_{\mathcal{Z}}(\vartheta) \le \mu_{\mathcal{Z}}(\vartheta + \varsigma) \le e^{L\varsigma}\mu_{\mathcal{Z}}(\vartheta).$$

Proof. Under the same noise a larger boost can only move the maximizer into $\mathcal{Z}$, and translating the noise along the chains that determine $\mathcal{Z}$ preserves a maximizer in $\mathcal{Z}$ when the extra boost is removed, at a density at least $e^{-L\varsigma}$ times the original. The maximizers are unique with probability one, by the argument at the start of this part, which applies to any scores. From ϑ to $\vartheta + \varsigma$ every member of $\mathcal{Z}$ gains ς and every other ideal gains nothing, so a maximizer in $\mathcal{Z}$ remains the maximizer, and a maximizer outside $\mathcal{Z}$ is either kept or replaced by a member of $\mathcal{Z}$. This proves the first claim and the first inequality.

For the second inequality fix $I \in \mathcal{Z}$ and a set of $L' \le L$ chains whose prefix lengths determine $\mathcal{Z}$, and consider a value of the noise at which I is the unique maximizer at $\vartheta + \varsigma$. Add ς to the coordinates $\xi_{i,k_i(I)}$ of these L' chains, and compare the perturbed scores at ϑ and the translated noise with those at $\vartheta + \varsigma$ and the original noise. The perturbed score of I changes by exactly $(L'-1)\varsigma$, that of any other member of $\mathcal{Z}$ by at most $(L'-1)\varsigma$, and that of an ideal outside $\mathcal{Z}$ by at most $(L'-1)\varsigma$ as well, because such an ideal differs from I in its prefix length in at least one of the L' chains. So I is the unique maximizer at ϑ and the translated noise. As a function of its value, the noise $\xi_{i,k}$ at each coordinate has the density $\exp(-\xi_{i,k} - \gamma - e^{-\xi_{i,k}-\gamma})$, and adding ς to the value multiplies this density by $\exp(-\varsigma + e^{-\xi_{i,k}-\gamma}(1 - e^{-\varsigma})) \ge e^{-\varsigma}$, so the density of the noise at the translated value is at least $e^{-L\varsigma}$ times its density at the original value. Since translation preserves volume, the probability that I is the maximizer at $\vartheta + \varsigma$ is at most $e^{L\varsigma}$ times the probability that it is the maximizer at ϑ, and summing over $I \in \mathcal{Z}$ proves the second inequality. □

Lemma 17 (Curvature). *The perturbed potential Φ is convex and continuously differentiable, and $\partial\Phi(\theta)/\partial\theta_{x,j} = \mu_{x,j}(\theta)$ for every score array θ, point x and level j. For every θ and x, the map $a \mapsto \mu_x(\theta +_x a)$ is Lipschitz on $\mathbb{R}^{K+1}$, hence differentiable at almost every a by Rademacher's theorem, and at every such a, with ∂_j the partial derivative with respect to a_j,*

$$\sum_{j,j'=0}^{K} a'_j a'_{j'}\, \partial_j \mu_{x,j'}(\theta +_x a) \le 4\,\mathrm{Var}_{\mu_x(\theta+_x a)}(a') \qquad \text{for every } a' \in \mathbb{R}^{K+1}.$$

Proof. The gradient comes from differentiating the perturbed maximum, whose maximizer is almost surely unique, and Lemma 16, applied to the level fibers, bounds how the level distribution changes when one score is raised. For every value of the noise, $\theta \mapsto \max_I\{\theta(I) + \xi(I)\}$ is a maximum of finitely many affine functions of θ, hence convex, and a change of one entry of θ changes it by at most the absolute value of that change. When the maximizer I^* is unique, which happens with probability one, a small enough change of $\theta_{x,j}$ changes the maximum by exactly that change if $F_{I^*}(x) = j$ and leaves it unchanged otherwise. So Φ is convex, and dominated convergence gives $\partial\Phi(\theta)/\partial\theta_{x,j} = \mathbb{P}(F_{I^*}(x) = j) = \mu_{x,j}(\theta)$.

The level fiber $\mathcal{J}_j(x) = \{I : F_I(x) = j\}$ consists of the ideals that contain (x, j) if $j \geq 1$ and avoid $(x, j+1)$ if $j < K$, and an ideal contains the k-th cell of $\mathcal{C}_i$ exactly when its prefix length in $\mathcal{C}_i$ is at least k, so $\mathcal{J}_j(x)$ is a cylinder on at most two chains. Raising $\theta_{x,j}$ by $\varsigma \geq 0$ raises by ς exactly the scores of its members. By Lemma 16 with $L = 2$, the maximizer changes only by moving into $\mathcal{J}_j(x)$, and, with θ' the raised array, it changes with probability

$$\mu_{x,j}(\theta') - \mu_{x,j}(\theta) \leq (e^{2\varsigma} - 1)\, \mu_{x,j}(\theta). \tag{6}$$

Consequently, under the same noise, the level distribution at every point moves by at most twice this probability in ℓ_1 distance, hence by at most $2(e^2 - 1)\varsigma$ when $\varsigma \leq 1$, and a decrease of $\theta_{x,j}$ is the reverse change. So the level distributions are Lipschitz in each entry of θ, the gradient of Φ is continuous, and $a \mapsto \mu_x(\theta +_x a)$ is Lipschitz. Since the maximizer can only move into $\mathcal{J}_j(x)$, raising a_j does not raise $\mu_{x,j'}(\theta +_x a)$ for $j' \neq j$.

Now fix a at which $a \mapsto \mu_x(\theta +_x a)$ is differentiable, and write $\mu = \mu_x(\theta +_x a)$ and $\partial_j\mu_{j'} = \partial_j\mu_{x,j'}(\theta +_x a)$. The last observation gives $\partial_j\mu_{j'} \leq 0$ for $j' \neq j$, and dividing (6) by ς and letting $\varsigma \downarrow 0$ gives $\partial_j\mu_j \leq 2\mu_j$. Adding the same constant to every entry of a adds it to every ideal's score and changes no maximizer, so $\sum_j \partial_j\mu_{j'} = 0$ for every j'. Since $\sum_{j'} \mu_{j'} = 1$, also $\sum_{j'} \partial_j\mu_{j'} = 0$ for every j. With $\bar{a}_j = a'_j - \sum_{j'} \mu_{j'} a'_{j'}$, expanding the squares and using these zero row and column sums at each equality,

$$\begin{aligned}\sum_{j,j'} a'_j a'_{j'}\, \partial_j\mu_{j'} &= \sum_{j,j'} \bar{a}_j \bar{a}_{j'}\, \partial_j\mu_{j'} = \frac{1}{2}\sum_{j\neq j'} (-\partial_j\mu_{j'})(\bar{a}_j - \bar{a}_{j'})^2 \\ &\leq \sum_{j\neq j'} (-\partial_j\mu_{j'})(\bar{a}_j^2 + \bar{a}_{j'}^2) = 2\sum_j \partial_j\mu_j\, \bar{a}_j^2,\end{aligned}$$

and the right side is at most $4\sum_j \mu_j \bar{a}_j^2 = 4\,\mathrm{Var}_\mu(a')$. □

Theorem 6 (Gibbs Domination from a Curvature Bound)**.** *Let $\kappa > 0$, and let $\Upsilon\colon \mathbb{R}^{K+1} \to \mathbb{R}$ be differentiable, with a locally Lipschitz gradient $\nabla\Upsilon$ taking values in Δ^K. Write ∂_j for the partial derivative with respect to the j-th entry of the argument, and suppose that at almost every $a \in \mathbb{R}^{K+1}$ the gradient is differentiable and*

$$\sum_{j,j'=0}^{K} a'_j a'_{j'}\, \partial_j\partial_{j'}\Upsilon(a) \leq \kappa\, \mathrm{Var}_{\nabla\Upsilon(a)}(a') \qquad \textit{for every } a' \in \mathbb{R}^{K+1}.$$

Then for all $a, a' \in \mathbb{R}^{K+1}$,

$$\Upsilon(a + a') - \Upsilon(a) \leq \frac{1}{\kappa} \log \sum_{j=0}^{K} \partial_j\Upsilon(a)\, e^{\kappa a'_j}.$$

Proof. After an exponential change of variables, the curvature condition says that the transformed function is concave, and the claim is its tangent-plane inequality. Suppose first that Υ is twice continuously differentiable, so that, by continuity, the condition holds at every a. On the positive orthant put $\chi(\varpi) = \exp(\kappa\Upsilon(\kappa^{-1}\log\varpi))$, the logarithm taken entrywise. With $a = \kappa^{-1}\log\varpi$, the chain rule gives $\partial_j\chi(\varpi) = \chi(\varpi)\,\partial_j\Upsilon(a)/\varpi_j$ and

$$\partial_j\partial_{j'}\chi(\varpi) = \frac{\chi(\varpi)}{\kappa\varpi_j\varpi_{j'}}\Big[\partial_j\partial_{j'}\Upsilon(a) + \kappa\,\partial_j\Upsilon(a)\,\partial_{j'}\Upsilon(a) - \kappa\mathbf{1}\{j = j'\}\,\partial_j\Upsilon(a)\Big].$$

Every vector of $\mathbb{R}^{K+1}$ has the form $(\varpi_j a'_j)_j$, and since the entries of $\nabla\Upsilon(a)$ sum to one,

$$\sum_{j,j'} \varpi_j a'_j \, \varpi_{j'} a'_{j'} \, \partial_j \partial_{j'} \chi(\varpi) = \frac{\chi(\varpi)}{\kappa} \Big[\sum_{j,j'} a'_j a'_{j'} \, \partial_j \partial_{j'} \Upsilon(a) - \kappa \operatorname{Var}_{\nabla\Upsilon(a)}(a') \Big] \le 0.$$

So χ is concave, and its tangent-plane inequality at $(e^{\kappa a_j})_j$, evaluated at $(e^{\kappa(a_j + a'_j)})_j$, reads

$$e^{\kappa\Upsilon(a+a')} \le e^{\kappa\Upsilon(a)} \Big[1 + \sum_j \partial_j \Upsilon(a) \big(e^{\kappa a'_j} - 1 \big) \Big] = e^{\kappa\Upsilon(a)} \sum_j \partial_j \Upsilon(a) \, e^{\kappa a'_j};$$

taking logarithms proves the claim.

In general, let X be a random vector with a smooth density supported in the unit ball, and for $\varepsilon > 0$ put $\Upsilon_\varepsilon(a) = \mathbb{E}\Upsilon(a - \varepsilon X)$. This function, the convolution of Υ with a smooth density of compact support, is infinitely differentiable, its gradient $\mathbb{E}\nabla\Upsilon(a - \varepsilon X)$ lies in Δ^K, and, because $\nabla\Upsilon$ is locally Lipschitz and $a - \varepsilon X$ has a density, dominated convergence gives $\partial_j \partial_{j'} \Upsilon_\varepsilon(a) = \mathbb{E}\, \partial_j \partial_{j'} \Upsilon(a - \varepsilon X)$. The function $\mu \mapsto \operatorname{Var}_\mu(a')$ is concave on Δ^K, being linear minus the square of a linear function, so the condition and Jensen's inequality give

$$\sum_{j,j'} a'_j a'_{j'} \, \partial_j \partial_{j'} \Upsilon_\varepsilon(a) \le \kappa \, \mathbb{E} \operatorname{Var}_{\nabla\Upsilon(a - \varepsilon X)}(a') \le \kappa \operatorname{Var}_{\nabla\Upsilon_\varepsilon(a)}(a').$$

The first case therefore applies to Υ_ε. As $\varepsilon \downarrow 0$, Υ_ε and $\nabla\Upsilon_\varepsilon$ converge pointwise to Υ and $\nabla\Upsilon$, which are continuous, and the argument of the logarithm stays at least $\min_j e^{\kappa a'_j} > 0$, so the inequality passes to the limit. □

The Lipschitz condition on the gradient cannot be dropped. Let $K = 1$ and $\varepsilon > 0$, let Cantor be the Cantor function on $[0, 1]$, which is continuous and nondecreasing with derivative zero at almost every point, extended by 0 to the left and by 1 to the right, and put

$$\Upsilon(a) = a_1 + \int_0^{a_0 - a_1} \Big(\frac{1}{4} + \frac{1}{2} \operatorname{Cantor}(\varsigma/\varepsilon) \Big) d\varsigma.$$

This function is convex and continuously differentiable, with $\partial_0 \Upsilon(a) = \frac{1}{4} + \frac{1}{2} \operatorname{Cantor}((a_0 - a_1)/\varepsilon)$ and $\partial_1 \Upsilon = 1 - \partial_0 \Upsilon$, so its gradient takes values in Δ^1 and has derivative zero at almost every point. The curvature condition therefore holds for every $\kappa > 0$. Since $\operatorname{Cantor}(\varsigma) + \operatorname{Cantor}(1 - \varsigma) = 1$, the Cantor function integrates to $1/2$ over $[0, 1]$, and for $a' = (\varepsilon, 0)$

$$\Upsilon(a') - \Upsilon(0) = \frac{\varepsilon}{2} > \frac{e^{\kappa\varepsilon} - 1}{4\kappa} \ge \frac{1}{\kappa} \log\Big(\frac{1}{4} e^{\kappa\varepsilon} + \frac{3}{4} \Big) = \frac{1}{\kappa} \log \sum_j \partial_j \Upsilon(0) \, e^{\kappa a'_j}$$

whenever $\kappa\varepsilon \le 1$, because $e^{\kappa\varepsilon} - 1 - 2\kappa\varepsilon$ is a convex function of $\kappa\varepsilon$ that vanishes at 0 and is negative at 1.

Lemma 18 (Levels Determined by One Prefix Length)**.** *Let $x \in P$, and suppose that for some chain $\mathcal{C}_i$ the level $F_I(x)$ depends on the ideal I only through its prefix length $k_i(I)$. Then for every score array θ and every $a \in \mathbb{R}^{K+1}$,*

$$\Phi(\theta +_x a) - \Phi(\theta) \le \log \sum_{j=0}^{K} \mu_{x,j}(\theta) \, e^{a_j}.$$

Proof. Integrating out the noise of chain i turns the increment into an average of exact log-partition increments, and Jensen's inequality moves the average inside the logarithm. Let $\mathcal{H}$ be the sigma-field generated by the noise coordinates of the chains other than $\mathcal{C}_i$. For each prefix length k that some ideal has in $\mathcal{C}_i$, let j_k be the common level at x of the ideals with $k_i(I) = k$, and let φ_k be the maximum of $\theta(I) + \sum_{i' \ne i} \xi_{i', k_{i'}(I)}$ over these ideals, an $\mathcal{H}$-measurable variable. The noise of $\mathcal{C}_i$ is independent of $\mathcal{H}$, and the perturbed maximum is $\max_k \{\varphi_k + \xi_{i,k}\}$. By the Gumbel maximum, its conditional expectation given $\mathcal{H}$ is $\log \sum_k e^{\varphi_k}$, and the conditional level probabilities are

$$\bar{\mu}_j = \mathbb{P}\big(F_{I^*}(x) = j \;\big|\; \mathcal{H} \big) = \frac{\sum_{k:\, j_k = j} e^{\varphi_k}}{\sum_k e^{\varphi_k}},$$

with expectations $\mathbb{E}\bar{\mu}_j = \mu_{x,j}(\theta)$. Adding a to the scores at x adds a_{j_k} to φ_k, so by Jensen's inequality for the logarithm

$$\Phi(\theta +_x a) - \Phi(\theta) = \mathbb{E}\log\frac{\sum_k e^{\varphi_k + a_{j_k}}}{\sum_k e^{\varphi_k}} = \mathbb{E}\log\sum_j \bar{\mu}_j\, e^{a_j} \le \log\sum_j \mu_{x,j}(\theta)\, e^{a_j}. \qquad \square$$

Theorem 3 (Gibbs Domination). *Let P be a finite poset, $K \ge 1$, and let Φ be the perturbed potential of a minimum chain partition of Q_K. For every score array θ, every point $x \in P$ and every $a \in \mathbb{R}^{K+1}$,*

$$\Phi(\theta +_x a) - \Phi(\theta) \le \frac{1}{\kappa}\log\sum_{j=0}^{K} \mu_{x,j}(\theta)\, e^{\kappa a_j}, \qquad \kappa = 4;$$

if $K = 1$, the inequality holds with $\kappa = 1$.

Proof. Lemma 17 supplies the hypothesis of Theorem 6 with $\kappa = 4$, and Lemma 18 covers resolution one. Fix θ and x. By Lemma 17, the function $a \mapsto \Phi(\theta +_x a)$ is differentiable with gradient $a \mapsto \mu_x(\theta +_x a)$, which takes values in Δ^K and is Lipschitz, and the condition of Theorem 6 holds with $\kappa = 4$, its left side being the second derivative of $a \mapsto \Phi(\theta +_x a)$ in the direction a'. Theorem 6 at the point 0 with the increment a gives the inequality with $\kappa = 4$. If $K = 1$, then $F_I(x) = 1$ exactly when I contains $(x, 1)$, that is, when $k_i(I) \ge k$ for the chain $\mathcal{C}_i$ of which $(x, 1)$ is the k-th cell, so Lemma 18 gives the inequality with $\kappa = 1$. $\square$

Differential consistency (Abernethy et al., 2015) bounds the diagonal second derivatives of a potential smoothed by independent noise on each entry of its argument by constant multiples of its first derivatives, the probabilities that each entry attains the perturbed maximum, provided the hazard rate of the noise, its density divided by its tail probability, is bounded. Generalized mixability (Reid et al., 2015; Mhammedi & Williamson, 2018) obtains regret independent of the horizon in prediction with expert advice from a comparison of the curvature of a convex function on the simplex with that of negative Shannon entropy. For noise shared by the ideals through a chain encoding, Lemma 17 bounds the curvature of the perturbed potential in the scores at one point, as a quadratic form, by four times that of a log-partition function, and Theorem 3 is the resulting comparison of increments.

In discrete choice theory the perturbed potential, as a function of the ideal scores, is the surplus function of an additive random-utility model over the ideals, the expected maximum utility when the utility of each ideal is its score plus noise. For such surplus functions Hofbauer & Sandholm (2002) represent the choice probabilities through the convex conjugate of the surplus function under positive-density and smoothness conditions, and Galichon & Salanié (2022) name the negative of that conjugate the generalized entropy of choice, which in the logit case, independent Gumbel noise on every alternative at unit scale, is Shannon entropy up to an additive constant.

The choice of the chain partition. Definition 3 fixes a minimum chain partition to keep the zero-score value small. The proof of Lemma 14 adds $\log\pi - \gamma$ to the bound on $\Phi(0)$ for each chain, and a minimum chain partition has $w_K \le H_P(K)/\log 2$ chains. The lower bound of Lemma 14, Lemmas 16 to 18 and the inequality of Theorem 3 hold for every chain partition of Q_K, with the potential and the level distributions defined from it as in Definition 3. The partition into the n columns makes the level at x the prefix length in the column of x, so Lemma 18 gives $\kappa = 1$ at every resolution, but its zero-score bound $2H_P(K) + (\log\pi - \gamma)n$ is $O(H_P(K))$ precisely when $n = O(H_P(K))$, while a minimum chain partition gives $\Phi(0) \le c_\Phi H_P(K)$ with $\kappa = 4$, and $\kappa = 1$ at resolution one.

For $\eta > 0$, the mix loss of a random variable X with finitely many values is $\mathrm{mix}_\eta(X) = -\eta^{-1}\log\mathbb{E}e^{-\eta X}$ (de Rooij et al., 2014). For the learners of Appendix E, the following lemma links the decrease of the potential to the mean and the variance of the loss of a sampled prediction.

Lemma 19 (Potential Decrease and Mix Loss). *Let θ be a score array, $x \in P$, $y \in [0, 1]$ and $\eta > 0$, let $G = \{z_0 < \cdots < z_K\}$ be a grid on which $\ell(\cdot, y)$ is finite, let κ be as in Theorem 3, and let $X = \ell(z_{F_{I^*}(x)}, y)$ for the perturbed maximizer I^* at θ. Then*

$$\frac{\kappa}{\eta}\Big(\Phi(\theta) - \Phi\big(\theta +_x \big(-\tfrac{\eta}{\kappa}\,\ell(z_j, y)\big)_{j=0}^{K}\big)\Big) \ge \mathrm{mix}_\eta(X).$$

If the values $\ell(z_0, y), \dots, \ell(z_K, y)$ lie in an interval of length at most $1/\eta$, then

$$\operatorname{mix}_\eta(X) \geq \mathbb{E}X - \eta \operatorname{Var} X.$$

Proof. Apply Theorem 3 to the update by the losses, and bound the exponential moment of X by a quadratic. Since $X = \ell(z_j, y)$ on the event $F_{I^*}(x) = j$, which has probability $\mu_{x,j}(\theta)$, we have $\mathbb{E}e^{-\eta X} = \sum_j \mu_{x,j}(\theta) e^{-\eta \ell(z_j, y)}$. The right side of Theorem 3 for the increment $(-(\eta/\kappa)\ell(z_j, y))_{j=0}^K$ is therefore $\kappa^{-1} \log \mathbb{E}e^{-\eta X} = -(\eta/\kappa) \operatorname{mix}_\eta(X)$, which gives the first inequality. For the second, $\eta(\mathbb{E}X - X) \leq 1$, and $e^\varsigma \leq 1 + \varsigma + \varsigma^2$ for every real $\varsigma \leq 1$, so

$$\mathbb{E}e^{-\eta X} = e^{-\eta \mathbb{E}X}\, \mathbb{E}e^{\eta(\mathbb{E}X - X)} \leq e^{-\eta \mathbb{E}X}\big(1 + \eta^2 \operatorname{Var} X\big) \leq e^{-\eta \mathbb{E}X + \eta^2 \operatorname{Var} X};$$

taking logarithms and dividing by $-\eta$ gives the claim. □

E Learners

This part states and proves the guarantee of the learner at a fixed resolution (Proposition 5), proves Theorem 4, and states the perturbed-ideal learner in pseudocode. The learner for pinball loss is in Appendix H.

Proposition 5 (Learning at a Fixed Resolution)**.** *Fix $K \geq 1$, a grid $G = \{z_0 < \dots < z_K\} \subseteq [0, 1]$ and an η-exp-concave loss ℓ finite on $[z_0, z_K]$. Let Φ, μ be those of a minimum chain partition of Q_K, let κ be as in Theorem 3, and write $R_T(\mathcal{F}_G)$ for the regret against $\mathcal{F}_G$.*

(a) The learner that starts from $\theta_1 = 0$, predicts $\bar{y}_t = \sum_j \mu_{x_t,j}(\theta_t)\, z_j$ and updates $\theta_{t+1} = \theta_t +_{x_t} \big(-(\eta/\kappa)\, \ell(z_j, y_t)\big)_{j=0}^K$ satisfies, on every sequence and at every prefix, $R_T(\mathcal{F}_G) \leq \kappa\Phi(0)/\eta \leq \kappa c_\Phi H_P(K)/\eta$.

(b) If moreover $0 \leq \ell \leq 1$, the learner that updates as in (a) with the rate $\min\{\eta, 1\}/(16\kappa)$ in place of η/κ and predicts the average of $z_{F_{I^*}(x_t)}$ over $S \in \{2, 4, 6, \dots\}$ independent perturbed maximizers I^* at θ_t satisfies $\mathbb{E}R_T(\mathcal{F}_G) \leq 16\kappa\Phi(0)/\min\{\eta, 1\}$ for every $T \geq 1$, and for every $\delta \in (0, 1)$

$$\mathbb{P}\Big\{ R_T(\mathcal{F}_G) \leq \frac{16\kappa\Phi(0)}{\min\{\eta, 1\}} + \frac{64 \log(1/\delta)}{S \min\{\eta, 1\}} \text{ for all } T \geq 1 \Big\} \geq 1 - \delta.$$

Each perturbed maximizer is computed by one minimum cut on a network with $nK + 2$ vertices and $O(n^2 K)$ arcs.

Proof. Both parts compare the loss of each round with the decrease of the potential under that round's update, sum these decreases over the rounds by telescoping, and bound the final potential from below by Lemma 14. In (b) the comparison holds for exponential moments, and Ville's inequality turns it into a bound for all T at once. Both learners start from $\theta_1 = 0$ and lower $\theta_{x_t,j}$ at round t by the update rate times $\ell(z_j, y_t)$, for every j. So for every $f \in \mathcal{F}_G$ with integer form F, the score of I_F at θ_{T+1} is minus the update rate times $\sum_{t \leq T} \ell(f(x_t), y_t)$, and $\Phi(\theta_{T+1})$ is at least this score by Lemma 14.

(a) Since $\bar{y}_t = \sum_j \mu_{t,j} z_j$ and $p \mapsto e^{-\eta \ell(p, y_t)}$ is concave, Jensen's inequality gives $e^{-\eta \ell(\bar{y}_t, y_t)} \geq \sum_j \mu_{t,j} e^{-\eta \ell(z_j, y_t)}$, that is, $\ell(\bar{y}_t, y_t) \leq \operatorname{mix}_\eta(X)$ for the loss $X = \ell(z_{F_{I^*}(x_t)}, y_t)$ of a sampled prediction. Lemma 19 at θ_t, x_t and y_t bounds $\operatorname{mix}_\eta(X)$ by $\kappa(\Phi(\theta_t) - \Phi(\theta_{t+1}))/\eta$, and summing over $t \leq T$ gives, for every $f \in \mathcal{F}_G$,

$$\sum_{t \leq T} \ell(\bar{y}_t, y_t) \leq \frac{\kappa}{\eta}\big(\Phi(0) - \Phi(\theta_{T+1})\big) \leq \frac{\kappa\Phi(0)}{\eta} + \sum_{t \leq T} \ell(f(x_t), y_t).$$

This holds on every sequence and at every prefix, and Lemma 14 gives $\Phi(0) \leq c_\Phi H_P(K)$.

(b) An η-exp-concave loss is also $\min\{\eta, 1\}$-exp-concave, because a power with exponent in $(0, 1]$ of a nonnegative concave function is concave. We may therefore assume $\eta \leq 1$ and prove (b) with η

in place of $\min\{\eta, 1\}$, so that the rate is $\eta/(16\kappa)$. Let $\mathcal{H}_t$ be the sigma-field generated by the history through round $t-1$ together with the adversary's choices x_t and y_t for round t. The scores θ_t and θ_{t+1} are $\mathcal{H}_t$-measurable, and the noise of round t is independent of $\mathcal{H}_t$, because the adversary is non-anticipating. Given $\mathcal{H}_t$, the S sampled predictions of round t are therefore independent, each equal to z_j with probability $\mu_{t,j}$. Let X and X' be the losses at y_t of two of them. In the calculation for round t, $\mathbb{E}$ and Var are conditional on $\mathcal{H}_t$.

Pairing the samples produces a variance term that offsets the one in the exponential-moment bound. The loss is convex in the prediction, since $e^{-\eta\ell}$ is concave and positive, so if the S samples are grouped into $S/2$ pairs, $\ell(\widehat{y}_t, y_t)$ is at most the average of the losses of the $S/2$ pair averages, and by exp-concavity the average of a pair with losses X and X' has loss at most

$$\phi = -\frac{1}{\eta}\log\frac{e^{-\eta X}+e^{-\eta X'}}{2} = \frac{X+X'}{2} - \frac{1}{\eta}\log\cosh\frac{\eta(X-X')}{2}.$$

The bound $\log\cosh\varsigma \geq \varsigma^2/4$ for $|\varsigma| \leq 1/2$ follows by integrating $(\log\cosh)'' \geq 1/2$ twice from zero. Since $|\eta(X-X')/2| \leq 1/2$, the quantity subtracted from $(X+X')/2$ is at least $\eta(X-X')^2/16$, whose expectation is $\eta \operatorname{Var} X/8$. The function ϕ takes values in $[0,1]$ and is 1-Lipschitz in each argument, so the Efron–Stein inequality (Steele, 1986) bounds its variance by $\operatorname{Var} X + \operatorname{Var} X'$. Thus

$$\mathbb{E}\phi \leq \mathbb{E}X - \frac{\eta}{8}\operatorname{Var} X, \qquad \operatorname{Var}\phi \leq 2\operatorname{Var} X.$$

The update of (b) is that of Lemma 19 with $\eta/16$ in place of η. Its second bound applies because the losses lie in $[0,1]$, so

$$\frac{16\kappa}{\eta}\big(\Phi(\theta_t) - \Phi(\theta_{t+1})\big) \geq \mathbb{E}X - \frac{\eta}{16}\operatorname{Var} X.$$

Since $(\eta/32)|\phi - \mathbb{E}\phi| \leq 1$, the bound $e^{\varsigma} \leq 1 + \varsigma + \varsigma^2$ at $\varsigma = (\eta/32)(\phi - \mathbb{E}\phi)$ gives $\log \mathbb{E}e^{(\eta/32)(\phi-\mathbb{E}\phi)} \leq (\eta^2/1024)\operatorname{Var}\phi$, and the three bounds combine into

$$\begin{aligned}\log\mathbb{E}\exp\Big(\frac{\eta}{32}\phi - \frac{\kappa}{2}\big(\Phi(\theta_t) - \Phi(\theta_{t+1})\big)\Big) &\leq \frac{\eta}{32}\Big(\mathbb{E}\phi - \frac{16\kappa}{\eta}\big(\Phi(\theta_t) - \Phi(\theta_{t+1})\big)\Big) + \frac{\eta^2}{1024}\operatorname{Var}\phi \\ &\leq -\frac{\eta^2}{512}\operatorname{Var} X + \frac{\eta^2}{512}\operatorname{Var} X = 0.\end{aligned}$$

The pairs are independent, and $(S/2)\ell(\widehat{y}_t, y_t)$ is at most the sum of their values of ϕ. Exponentiating the last bound and multiplying over the pairs gives

$$\mathbb{E}\Big[\exp\Big(\frac{S}{4}\Big(\frac{\eta}{16}\ell(\widehat{y}_t, y_t) - \kappa\big(\Phi(\theta_t) - \Phi(\theta_{t+1})\big)\Big)\Big) \;\Big|\; \mathcal{H}_t\Big] \leq 1.$$

Hence the process

$$\exp\Big(\frac{S}{4}\Big[\frac{\eta}{16}\sum_{t\leq T}\ell(\widehat{y}_t, y_t) - \kappa\big(\Phi(0) - \Phi(\theta_{T+1})\big)\Big]\Big), \qquad T \geq 0,$$

equals one at $T = 0$ and, with respect to the sigma-fields $\mathcal{H}_{T+1}$, is a nonnegative supermartingale, a process whose expected next value, given the past, is at most its current value. So Ville's inequality (Howard et al., 2020) bounds by δ the probability that it ever reaches $1/\delta$. Outside that event, for every $T \geq 1$ and every $f \in \mathcal{F}_G$,

$$\begin{aligned}\sum_{t\leq T}\ell(\widehat{y}_t, y_t) &< \frac{16\kappa}{\eta}\big(\Phi(0) - \Phi(\theta_{T+1})\big) + \frac{64\log(1/\delta)}{S\eta} \\ &\leq \frac{16\kappa\Phi(0)}{\eta} + \frac{64\log(1/\delta)}{S\eta} + \sum_{t\leq T}\ell(f(x_t), y_t),\end{aligned}$$

which is the probability bound. Jensen's inequality applied to the one-round bound gives $\mathbb{E}[\ell(\widehat{y}_t, y_t) \mid \mathcal{H}_t] \leq (16\kappa/\eta)(\Phi(\theta_t) - \Phi(\theta_{t+1}))$. Summing over $t \leq T$, taking expectations and bounding $\Phi(\theta_{T+1})$ from below gives $\mathbb{E}R_T(\mathcal{F}_G) \leq 16\kappa\Phi(0)/\eta$.

The last claim is the minimum-cut computation of Appendix D. □

Algorithm 1 The perturbed-ideal learner at a fixed resolution

Input: a resolution K; a grid $G = \{z_0 < \cdots < z_K\}$; an η-exp-concave loss ℓ finite on $[z_0, z_K]$; $\kappa = 4$, or $\kappa = 1$ if $K = 1$; a mode: pair mode, which requires $0 \le \ell \le 1$, with an even number $S \ge 2$ of samples per round, or mean mode, with exact means or, for a loss that is Λ_1-Lipschitz on $[z_0, z_K]$, with $S_{K,t}$ samples at round t chosen for accuracy targets $\varepsilon_{K,t}$ and confidences $\delta_{K,t}$

compute a minimum chain partition $\mathcal{C}_1, \dots, \mathcal{C}_{w_K}$ of Q_K by maximum bipartite matching (Fulkerson, 1956), and set $\theta_1 = 0$
for $t = 1, 2, \dots$ **do**
  observe x_t
  if mean mode is exact **then**
    predict $\bar{y}_t = \sum_{j=0}^{K} \mu_{x_t,j}(\theta_t)\, z_j$
  else
    for each sample of round t **do**
      draw an independent copy of the noise ξ
      compute a maximum-weight ideal I^* for the node weights at θ_t and ξ by one minimum cut (Appendix D)
    end for
    predict $\widehat{y}_t$, the average of the sampled predictions $z_{F_{I^*}(x_t)}$; in mean mode, $S_{K,t}$ is large enough that $|\widehat{y}_t - \bar{y}_t| \le \varepsilon_{K,t}$ with probability at least $1 - \delta_{K,t}$ (Lemma 4)
  end if
  observe y_t
  if pair mode **then**
    $\theta_{t+1} \leftarrow \theta_t +_{x_t} \left(-\frac{\min\{\eta,1\}}{16\kappa}\, \ell(z_j, y_t)\right)_{j=0}^{K}$
  else
    $\theta_{t+1} \leftarrow \theta_t +_{x_t} \left(-\frac{\eta}{\kappa}\, \ell(z_j, y_t)\right)_{j=0}^{K}$
  end if
end for

The update of (a) does not use the prediction, so the bound of (a) holds for the mean predictions along the sequence that any learner with this update realizes. For a learner that predicts $\widehat{y}_t \in [z_0, z_K]$, adding the loss differences therefore gives, on every sequence and at every prefix,

$$R_T(\mathcal{F}_G) \le \frac{\kappa \Phi(0)}{\eta} + \sum_{t \le T} \big(\ell(\widehat{y}_t, y_t) - \ell(\bar{y}_t, y_t)\big), \tag{7}$$

where $R_T(\mathcal{F}_G)$ is the regret of its own predictions. When $\ell(\cdot, y)$ is Lipschitz on $[z_0, z_K]$ with constant Λ_1 for every y, each summand is at most $\Lambda_1 |\widehat{y}_t - \bar{y}_t|$, and if $\widehat{y}_t$ averages sampled predictions, these are independent given the past and (x_t, y_t), with conditional mean $\bar{y}_t$, so Lemma 4 bounds the deviation.

Pair mode is the learner of Proposition 5(b). Mean mode uses the update of (a) and predicts either the mean prediction, as the learner of (a) does, or an average of sampled predictions, whose regret (7) bounds. Algorithm 1 states both modes, each sampled prediction being one minimum cut (Appendix D). A loss with values in $[0, \Lambda]$, $\Lambda > 0$, is run in pair mode as ℓ/Λ, which is $\Lambda\eta$-exp-concave, and (β/Λ)-smooth when ℓ is β-smooth, so that the scores move by $\min\{\Lambda\eta, 1\}/(16\kappa\Lambda)$ times the original losses. The horizon-free learner of Section 4.3 runs Algorithm 1 on the uniform grid G_K for smooth losses and on the rational grid of Lemma 11 for log loss, on which log loss is Lipschitz with constant $\Lambda_1 = 1 + (2K+1)^2$.

The next two theorems bound the regret of Algorithm 1 at one resolution against the whole isotonic class, one theorem for each loss family of Theorem 4. Lemmas 20 to 22 then combine them into the horizon-free learner.

Theorem 7 (Smooth Losses at a Fixed Resolution)**.** *Let ℓ be η-exp-concave and β-smooth with $0 \le \ell \le \Lambda$, $\Lambda > 0$, and let $K \ge 1$ and $\delta \in (0,1)$. Run Algorithm 1 at resolution K in pair mode, on the uniform grid G_K and the loss ℓ/Λ, with $S = 2\max\{1, \lceil \log(1/\delta) \rceil\}$ samples per round. With probability at least $1 - \delta$, its regret against $\mathcal{F}_P$ satisfies at every prefix*

$$R_T \le \frac{\big(16\kappa\Phi(0) + 32\big)\Lambda}{\min\{\Lambda\eta, 1\}} + \frac{\beta T}{8K^2}.$$

For any even $S \geq 2$, $\mathbb{E}R_T$ obeys the same bound without the term 32, at every $T \geq 1$. Both bounds are at most $C_\ell(H_P(K) + T/K^2)$ with C_ℓ depending only on (η, β, Λ).

Proof. Apply Proposition 5(b) to the normalized loss, and add the cost of rounding the comparator into the grid. The loss ℓ/Λ has values in $[0,1]$ and is $\Lambda\eta$-exp-concave, so at resolution K its regret against $\mathcal{F}_K$ is, with probability at least $1-\delta$ and at every prefix, at most $[16\kappa\Phi(0) + 64\log(1/\delta)/S]/\min\{\Lambda\eta, 1\}$, and its expected regret is at most $16\kappa\Phi(0)/\min\{\Lambda\eta, 1\}$. Since $S \geq 2\log(1/\delta)$, the choice of S gives $64\log(1/\delta)/S \leq 32$, and multiplying by Λ turns both bounds into bounds for ℓ. For the comparator, Lemma 2 gives at each prefix and for each $f \in \mathcal{F}_P$ a shift υ with $\lfloor f \rceil_\upsilon \in \mathcal{F}_K$ whose cumulative loss exceeds that of f by at most $\beta T/(8K^2)$, so the regret against $\mathcal{F}_P$ exceeds the regret against $\mathcal{F}_K$ by at most that much. Finally $\Phi(0) \leq c_\Phi H_P(K)$ by Lemma 14, and $H_P(K) \geq \log 2$ absorbs the constant term. □

Theorem 8 (Log Loss at a Fixed Resolution)**.** *Let $K \geq 1$ and $\delta \in (0,1)$, and put $\Lambda_1 = 1 + (2K+1)^2$. Run Algorithm 1 at resolution K in mean mode, on the rational grid of Lemma 11 and log loss, drawing at the t-th round of its run*

$$S_{K,t} = \left\lceil 2\varepsilon_{K,t}^{-2} \log \frac{2}{\delta_{K,t}} \right\rceil \text{ samples}, \qquad \varepsilon_{K,t} = \frac{1}{\Lambda_1 (t+1)^2}, \qquad \delta_{K,t} = \frac{\delta}{t(t+1)}.$$

With probability at least $1-\delta$, its regret against $\mathcal{F}_P$ satisfies at every prefix

$$R_T \leq \kappa\Phi(0) + 1 + \frac{9T}{(K+1)^2},$$

and $\mathbb{E}R_T$ obeys this bound for every $T \geq 1$. Both bounds are at most $C(H_P(K) + T/K^2)$ with an absolute constant C.

Proof. Three costs add: the potential bound, the sampling error of the means and the rounding of the comparator. Log loss is 1-exp-concave and finite on the grid, whose levels are interior, so the learner updates at the rate $1/\kappa$ of Proposition 5(a), and (7) bounds its regret against $\mathcal{F}_G$, at every prefix, by $\kappa\Phi(0) + \Lambda_1 \sum_{t \leq T} |\widehat{y}_t - \bar{y}_t|$, since $\ell_{\log}(\cdot, y)$ is Λ_1-Lipschitz on the convex hull of the grid (Lemma 11).

Given the history and the round's x_t and y_t, the $S_{K,t}$ sampled predictions of round t are independent with values in $[0,1]$ and mean $\bar{y}_t$. So Lemma 4 bounds the conditional probability that $|\widehat{y}_t - \bar{y}_t|$ exceeds $\varepsilon_{K,t}$ by $2e^{-2S_{K,t}\varepsilon_{K,t}^2} \leq \delta_{K,t}$. These bounds sum to δ over the rounds of the run, so with probability at least $1-\delta$ every round of it satisfies $|\widehat{y}_t - \bar{y}_t| \leq \varepsilon_{K,t}$, and then $\Lambda_1 \sum_t |\widehat{y}_t - \bar{y}_t| \leq \sum_{t \geq 1} (t+1)^{-2} < 1$. The conditional variance of $\widehat{y}_t$ is at most $1/(4S_{K,t})$, so by the conditional Cauchy–Schwarz inequality the conditional mean of $\Lambda_1 |\widehat{y}_t - \bar{y}_t|$ is at most $\Lambda_1/(2\sqrt{S_{K,t}}) \leq \Lambda_1 \varepsilon_{K,t}/(2\sqrt{2})$. These means again sum to less than 1, and taking expectations in the bound of the first paragraph gives the bound on $\mathbb{E}R_T$.

It remains to round the hindsight optimum along its level sets. By Lemma 9 with the rounding map of Lemma 11, at every prefix some member of $\mathcal{F}_G$ has cumulative loss exceeding the minimum over $\mathcal{F}_P$ by at most $9T/(K+1)^2$. The last claim uses $\Phi(0) \leq c_\Phi H_P(K)$ and $1 \leq H_P(K)/\log 2$. □

Before its resolution is activated, an expert follows the scalar baseline, which ignores the order.

Lemma 20 (The Scalar Baseline)**.** *Let ℓ be η-exp-concave and β-smooth, or log loss. The learner B of Algorithm 2 is deterministic, computes in time polynomial in the round, and satisfies on every sequence, for every $t \geq 1$,*

$$\mathrm{Loss}_B(t) - \min_{f\colon P \to [0,1]} \sum_{s \leq t} \ell(f(x_s), y_s) \leq C_\ell \sum_{x \in P} \log\bigl(1 + N_t(x)\bigr) \leq C_\ell\, n \log(1 + t/n),$$

where $N_t(x)$ is the number of rounds $s \leq t$ with $x_s = x$, and C_ℓ depends only on (η, β) and is absolute for log loss.

Proof. Bound the regret at each point separately, then sum and apply Jensen's inequality. Since the counts $N_t(x)$ sum to t and $\log(1+\cdot)$ is concave, a bound $C_\ell \log(1 + N_t(x))$ at each point gives both

Algorithm 2 The scalar baseline B

Input: an η-exp-concave loss ℓ, β-smooth or log loss ($\eta = 1$); the uniform grid at each resolution for a smooth loss, the rational grid of Lemma 11 for log loss

1: design lengths $T_1 = 1$, $T_2 = 2$ and $T_{\iota+1} = T_\iota^2$ for $\iota \geq 2$, epoch ι using the resolution $K_\iota = \lceil \sqrt{T_\iota} \rceil$
2: **for** $t = 1, 2, \ldots$ **do**
3: observe x_t; if no epoch is in progress at x_t, begin its next epoch ι there, with the weights $\nu_j = 1/(K_\iota + 1)$ on the levels $z_0 < \cdots < z_{K_\iota}$ of the grid of resolution K_ι
4: predict $\sum_j \nu_j z_j / \sum_j \nu_j$ for the epoch in progress at x_t
5: observe y_t and multiply ν_j by $e^{-\eta \ell(z_j, y_t)}$ for every j; end the epoch after T_ι rounds at x_t
6: **end for**

inequalities. In the argument below, B plays the rounds at one point as a sequence of its own, and t counts them.

At a point, B plays in epochs, the runs between two restarts, of design lengths $T_1 = 1$, $T_2 = 2$ and $T_{\iota+1} = T_\iota^2$ for $\iota \geq 2$. In epoch ι it runs exponential weights afresh, at rate η, over the $K_\iota + 1$ levels of the grid of resolution $K_\iota = \lceil \sqrt{T_\iota} \rceil$, with the uniform prior, and predicts the weighted mean of the levels. By Lemma 3 its loss in the epoch exceeds that of each of these levels by at most $\log(K_\iota + 1)/\eta$, and on the epoch's rounds some level is nearly as good as every constant. For a β-smooth loss, Lemma 2 applied to those rounds and to a constant comparator gives a level worse by at most $\beta T_\iota/(8K_\iota^2) \leq \beta/8$. For log loss, Lemmas 9 and 11 round the best constant on those rounds into the grid at a cost of at most $9T_\iota/(K_\iota+1)^2 \leq 9$. The loss of a constant comparator splits across the epochs, so the regret over the first t rounds at the point is at most the sum over the epochs ι' begun by then of $\log(K_{\iota'} + 1)/\eta + \beta/8$. For log loss, which is 1-exp-concave, the summand is $\log(K_{\iota'} + 1) + 9$.

The squaring of the design lengths keeps that sum logarithmic. Let ι be the epoch in progress at time t. Then $T_\iota \leq t^2$, for $\iota = 1$ because $T_1 = 1$, for $\iota = 2$ because $t \geq 2$ and $T_2 = 2$, and for $\iota \geq 3$ because the preceding epoch alone has taken $T_{\iota-1} = \sqrt{T_\iota}$ rounds, so that $t > \sqrt{T_\iota}$. Since $T_{\iota'} = 2^{2^{\iota'-2}}$ for $\iota' \geq 2$, the logarithms of the design lengths double from one epoch to the next, and

$$\sum_{\iota' \leq \iota} \log T_{\iota'} \leq 2 \log T_\iota \leq 4 \log t, \qquad \iota \leq 3 + \log_2 \log_2 t \quad (t \geq 2).$$

With $\log(K_{\iota'} + 1) \leq \log 3 + \frac{1}{2} \log T_{\iota'}$, the sum of the epoch bounds is therefore at most a constant multiple of $1 + \log t$, hence of $\log(1 + t)$, the constant depending only on (η, β).

The predictions are weighted means, so B is deterministic, and the epoch in progress at a point keeps $K_\iota + 1 \leq t + 1$ weights. □

The experts run the dyadic resolutions only, which costs a constant factor.

Lemma 21 (Dyadic Resolution Selection)**.** *Let $T \geq 1$ and let $A, A' \geq 0$. Some power of two K satisfies*

$$A\, H_P(K) + \frac{A'T}{K^2} \leq \max\{A, 4A'\}\, \Psi_P(T).$$

Proof. Round a minimizer up to the next power of two. The infimum defining $\Psi_P(T)$ is attained (Lemma 1(c)). Let K_Ψ attain it and let K be the least power of two with $K \geq K_\Psi$, so that $K \leq 2K_\Psi$. Monotonicity of H_P, an isotonic function into $\{0, \ldots, K\}$ being one into $\{0, \ldots, K+1\}$, and subadditivity (Lemma 1(b)) give $H_P(K) \leq H_P(2K_\Psi) \leq 2H_P(K_\Psi)$, while $T/K^2 \leq T/K_\Psi^2$. Hence

$$\begin{aligned} A\, H_P(K) + \frac{A'T}{K^2} &\leq A \cdot 2H_P(K_\Psi) + 4A' \cdot \frac{T}{4K_\Psi^2} \\ &\leq \max\{A, 4A'\}\Big(2H_P(K_\Psi) + \frac{T}{4K_\Psi^2}\Big) = \max\{A, 4A'\}\, \Psi_P(T). \end{aligned}$$

□

Algorithm 3 The horizon-free perturbed-ideal learner

Input: a confidence $\delta \in (0,1)$; a β-smooth loss with values in $[0,\Lambda]$, run as ℓ/Λ in pair mode, or log loss, run in mean mode; the rate η, the exp-concavity parameter of the loss run, which is 1 for log loss; the priors $\nu_K = 1/(2(\log_2 K+1)(\log_2 K+2))$ of the dyadic resolutions K

start B (Algorithm 2); weight $1/2$ on the constant expert predicting $1/2$, and $\bar{W}_1 = 1/2$ on the inactive experts
for $t = 1, 2, \ldots$ **do**
  if t is a power of two **then**
    activate E_t: move the share $\bar{W}_t/(\log_2 t + 2)$ of $\bar{W}_t$ to it
    start a fresh copy of Algorithm 1 at resolution t, in pair mode with $S_t = 2\max\{1, \lceil \log(1/(\delta\nu_t))\rceil\}$ samples per round, or in mean mode with the sampling schedule of Theorem 8 at confidence $\delta\nu_t$
  end if
  observe x_t and collect the predictions of the constant expert, of B, and of each active E_K, which plays this round of its copy of Algorithm 1
  predict the weighted mean of these predictions, each active expert and the constant expert carrying its own weight and B carrying $\bar{W}_t$
  observe y_t and multiply the weight of the constant expert and of each active expert by the exponential of $-\eta$ times its own loss, and $\bar{W}_t$ by the exponential of $-\eta$ times the loss of B
  advance B and every active copy of Algorithm 1 with (x_t, y_t)
end for

Aggregation over an expert set that grows with the rounds is studied by Mourtada & Maillard (2017). Here the experts are fixed in advance and only the resolutions already reached are active. For a set of rounds, the regret of a learner over those rounds is its cumulative loss there minus the least cumulative loss there of a member of $\mathcal{F}_P$.

Lemma 22 (Aggregation with Delayed Activation)**.** *Let ℓ be η-exp-concave, with $0 \le \ell \le 1$ or $\ell = \ell_{\log}$, and consider the master of Algorithm 3. Suppose that its baseline satisfies the bound of Lemma 20 and that, for every dyadic K, the learner that E_K runs has regret at most $C_\ell(H_P(K) + T/K^2)$ over the rounds K through T for every $T \ge K$, both with the same constant C_ℓ. Then the master satisfies $R_T \le C'_\ell\, r_P(T)$ for every $T \ge 1$, with C'_ℓ depending only on C_ℓ and η. If the second hypothesis holds in expectation at each fixed T, so does the conclusion. At round t the master runs the learners of the experts with $K \le t$, $1 + \lfloor \log_2 t \rfloor$ of them.*

Proof. Aggregate the experts by Lemma 3, absorb the delay and the prior into $H_P(K)$, and optimize over the dyadic resolutions. The master of Algorithm 3 implements exponential weights over the countable expert set. Let K_t be the least dyadic resolution at least t. The experts E_K with $K \ge t$ have made the predictions of B through round $t-1$, so their unnormalized weights at round t are their priors times $e^{-\eta\,\mathrm{Loss}_B(t-1)}$, and since $\nu_K = \frac{1}{2}[(\log_2 K+1)^{-1} - (\log_2 K+2)^{-1}]$, those priors sum to $1/(2(\log_2 K_t + 1))$. Their combined weight is therefore the $\bar{W}_t$ of Algorithm 3, and at $t = K_t$ the share $\nu_t e^{-\eta\,\mathrm{Loss}_B(t-1)}$ of the newly activated expert is $\bar{W}_t/(\log_2 t + 2)$.

By Lemma 3, on every realization and for every expert E and every $T \ge 1$,

$$\sum_{t \le T} \ell(\widehat{y}_t, y_t) \le \mathrm{Loss}_E(T) + \frac{\log(1/\nu_E)}{\eta}.$$

Fix a dyadic K, put $s = \min\{T, K-1\}$ and let $f \in \mathcal{F}_P$. Up to round s the expert E_K follows B, whose cumulative loss there exceeds that of f by at most $C_\ell n \log(1+s/n) \le C_\ell H_P(K)$ (Lemma 20 and Lemma 1(b)). If $T \ge K$, its own learner runs over the rounds K through T, with cumulative loss there exceeding that of f by at most $C_\ell(H_P(K) + T/K^2)$. For the prior, $\log(1+\varsigma) \le \varsigma$ and $H_P(K) \ge \log(K+1) \ge \max\{\log 2, (\log 2)\log_2 K\}$ give

$$\log(1/\nu_K) = \log 2 + \log(1+\log_2 K) + \log(2+\log_2 K) \le 2\log 2 + \tfrac{3}{2}\log_2 K \le 5 H_P(K).$$

Adding the three bounds,

$$R_T \le \Big(2C_\ell + \frac{5}{\eta}\Big) H_P(K) + C_\ell \frac{T}{K^2}$$

for every dyadic K, and Lemma 21 bounds the least of these values by $\max\{2C_\ell+5/\eta, 4C_\ell\}\Psi_P(T)$.

For the cap, a loss with values in $[0,1]$ has $R_T \le T$. For log loss, the display above at the constant expert gives $\sum_{t\le T}\ell_{\log}(\widehat{y}_t, y_t) \le T\log 2 + (\log 2)/\eta \le (1+1/\eta)T\log 2$, and the comparator's cumulative loss is nonnegative. So $R_T \le C'_\ell \min\{T, \Psi_P(T)\} = C'_\ell r_P(T)$.

For the version in expectation, fix T and let K be the power of two that Lemma 21 supplies for the coefficients above. It depends only on P, on T and on those coefficients. Every step above except the hypothesis holds on each realization, so taking expectations in the displayed bound for that K proves the claim. Finally, the experts active at round t are those with $K \le t$. □

Theorem 4 (Horizon-Free Polynomial-Time Learning)**.** *Let ℓ be η-exp-concave and β-smooth with $0 \le \ell \le \Lambda$, $\Lambda > 0$, or let ℓ be log loss, and let $\delta \in (0,1)$. The horizon-free perturbed-ideal learner of Section 4.3, run on ℓ/Λ in the first case, satisfies, for every finite poset P and every non-anticipating adaptive adversary,*

$$\mathbb{E}R_T \le C_\ell\, r_P(T) \text{ for every } T \ge 1, \qquad \mathbb{P}\{R_T \le C'_\ell\, r_P(T) \text{ for all } T \ge 1\} \ge 1-\delta,$$

where C_ℓ and C'_ℓ depend only on (η, β, Λ) and are absolute constants for log loss; they do not depend on P, T or δ. For each fixed loss and fixed regularity bounds its computation per round is polynomial in n, t, the input bit lengths and $\log(1/\delta)$ in the finite-precision model; for bounded losses it uses $O\big((1+\log t)(1+\log(1/\delta)+\log\log(t+3))\big)$ minimum cuts at round t, each on a network with $O(nt)$ vertices and $O(n^2 t)$ arcs, and for log loss polynomially many.

Proof. The learner is the master of Algorithm 3, and what remains is to supply the hypothesis of Lemma 22 for each loss family and to bound the computation. A loss with values in $[0,\Lambda]$ is run as ℓ/Λ, which has values in $[0,1]$ and is $\Lambda\eta$-exp-concave. The regret of ℓ is Λ times that of ℓ/Λ, a factor the constants below absorb.

For a bounded loss, E_K runs Algorithm 1 in pair mode at resolution K, at the confidence $\delta_K = \delta\nu_K$ and with $S_K = 2\max\{1, \lceil\log(1/\delta_K)\rceil\}$ samples per round. By Theorem 7, its regret over the rounds K through T is at most $C_\ell(H_P(K) + T/K^2)$ for every $T \ge K$, on an event of probability at least $1-\delta_K$ and also in expectation at each fixed T. For log loss, E_K runs Algorithm 1 in mean mode on the rational grid with the sampling schedule of Theorem 8 at confidence δ_K, and Theorem 8 gives the same two conclusions. The exceptional events have total probability at most $\sum_K \delta_K = \delta/2$, and Lemma 22, with C_ℓ enlarged to cover the baseline of Lemma 20 as well, turns the resolution-wise guarantees into the two statements of the theorem.

Each activated learner is analyzed along the realized sequence, conditionally on the past randomness of all the experts together with x_t and y_t. The adversary is non-anticipating, so x_t and y_t are fixed before any expert's noise of round t, which is the conditioning of Proposition 5(b).

At round t the active experts are the $1+\lfloor\log_2 t\rfloor$ with $K \le t$, and each sample they draw takes one minimum cut on a network with $nK+2 \le nt+2$ vertices and $O(n^2K) = O(n^2t)$ arcs (Appendix D). For a bounded loss, $\log(1/\delta_K) = \log(1/\delta) + \log(1/\nu_K)$ with $\log(1/\nu_K) = O(1+\log\log(t+3))$ for $K \le t$, so each active expert draws $O(1+\log(1/\delta)+\log\log(t+3))$ samples, which gives the cut count of the theorem. For log loss, an expert's own round is at most t, so the counts $S_{K,t}$ of Theorem 8 are at most $O(K^4t^4\log(2/\delta_{K,t}))$, polynomial in t and $\log(1/\delta)$. Besides the cuts, the learner computes a minimum chain partition of Q_K once per activation, by maximum bipartite matching (Fulkerson, 1956), advances B, whose state at a point has at most $N_t(x)+1$ weights (Lemma 20), and keeps one weight for each active expert, one for the constant expert and one for the inactive group. In the finite-precision model all of this has polynomial bit complexity (Appendix I). □

F Chain Relaxation

The remark of this part bounds the regret of exponential weights run separately on the chains of a minimum chain partition of P. Its example is an ordinal sum of antichains, in which every element of one antichain lies below every element of the next.

Remark 5 (Chain Relaxation). For squared loss and a design horizon T, drop the constraints between the chains of a minimum chain partition of P with chain sizes $n_1, \dots, n_w$ and run exponential

weights on each chain, whose level marginals are computed from prefix and suffix sums in $O(n_i K)$ arithmetic operations per chain and round. The resulting regret is at most

$$r_P^{\mathrm{ch}}(T) = \min\Big\{T,\ \inf_{K\ge 1}\Big[2H_P^{\mathrm{ch}}(K) + \frac{T}{4K^2}\Big]\Big\}, \qquad H_P^{\mathrm{ch}}(K) = \sum_i \log\tbinom{n_i+K}{n_i}.$$

Since $H_P(K) \le H_P^{\mathrm{ch}}(K) \le \min\{w,h\} H_P(K)$, this is within the factor $\min\{w,h\}$ of the rate, and the count factor is sharp up to constants on the ordinal sum of $h = 2^w$ antichains of size w, where $H_P(1) = \Theta(w)$ while $H_P^{\mathrm{ch}}(1) = \Theta(w^2)$, and at $T = w^5$ the rate is of order $w^{7/3}$ while $r_P^{\mathrm{ch}}(T)$ is of order w^3. The perturbed-ideal learner avoids the factor $\min\{w,h\}$.

Proof. Dropping the constraints between the chains replaces P by another poset, on which the upper bound of Appendix A is the relaxed rate. Fix a minimum chain partition of P into w chains, the i-th of n_i points, and let Q be the poset on the points of P that keeps the relations inside the chains and drops the others. Every isotonic function on P is isotonic on Q, and an isotonic function on Q with values in a grid of resolution K is a tuple of nondecreasing sequences, one along each chain, so $H_Q(K) = H_P^{\mathrm{ch}}(K)$ and $r_Q(T) = r_P^{\mathrm{ch}}(T)$. By Theorem 5, applied to Q and squared loss, there is a deterministic learner whose regret against $\mathcal{F}_Q$, hence against the smaller class $\mathcal{F}_P$, is at most $r_P^{\mathrm{ch}}(T)$.

Given T, that learner takes a resolution $K \le \lceil\sqrt{T}\rceil$ that minimizes $2H_P^{\mathrm{ch}}(K) + T/(4K^2)$ (Lemma 1(c) for Q). It runs exponential weights with the uniform prior over the isotonic functions on Q with values in the uniform grid of that resolution, predicting the weighted mean of their values at x_t. Its weights factor over the chains. The cumulative loss of such a function is the sum over the chains of the losses it incurs at their points, so the weight of a function is the product over the chains of the weights of its restrictions. The law of the level at a point of the i-th chain therefore comes from the factor of that chain alone. Index the points of that chain by their positions $k = 1, \dots, n_i$ along it, and let $\nu_{k,j}$ be the exponential of $-\frac{1}{2}$ times the cumulative loss that the level z_j has incurred at the k-th point. A nondecreasing assignment of levels to the points then has as its weight the product over k of the $\nu_{k,j}$ at the levels it assigns. Put

$$\nu^{\le}_{1,j} = \nu_{1,j}, \quad \nu^{\le}_{k,j} = \nu_{k,j} \sum_{j' \le j} \nu^{\le}_{k-1,j'}, \qquad \nu^{\ge}_{n_i,j} = 1, \quad \nu^{\ge}_{k,j} = \sum_{j' \ge j} \nu_{k+1,j'}\, \nu^{\ge}_{k+1,j'}.$$

Then $\nu^{\le}_{k,j}$ is the total weight of the nondecreasing assignments to the first k points that end at the level j, and $\nu^{\ge}_{k,j}$ that of the assignments to the remaining points that start at j or above. The probability that the level at the k-th point is j is therefore $\nu^{\le}_{k,j}\nu^{\ge}_{k,j}$ divided by $\sum_{j'} \nu^{\le}_{n_i,j'}$. Prefix and suffix sums over the levels produce all these quantities in $O(n_i K)$ arithmetic operations, and a round changes the $K+1$ weights of the queried point.

For the counts, restriction to the chains is an injection from the isotonic functions on P into those on Q, so $H_P(K) \le H_P^{\mathrm{ch}}(K)$. In the other direction, every chain has at most h points, so each summand $\log\binom{n_i+K}{n_i}$ is at most $\log\binom{h+K}{h}$. The isotonic functions of a longest chain of P extend to distinct isotonic functions on P (Lemma 1(a)), so $\log\binom{h+K}{h} \le H_P(K)$ and $H_P^{\mathrm{ch}}(K) \le w H_P(K)$. Also $H_P^{\mathrm{ch}}(K) \le \sum_i n_i \log(K+1) = n\log(K+1) \le wh\log(K+1) \le h H_P(K)$, by Lemma 1(b). Since $\min\{w,h\} \ge 1$, at every resolution the objective of $r_P^{\mathrm{ch}}(T)$ is at most $\min\{w,h\}$ times that of $r_P(T)$, and the same holds for the cap T, so $r_P^{\mathrm{ch}}(T) \le \min\{w,h\} r_P(T)$.

In the ordinal sum of $h = 2^w$ antichains of size w, every ideal other than P itself consists of all the antichains before a uniquely determined antichain together with a proper subset of that antichain. The isotonic functions $P \to \{0,1\}$ are the indicators of the filters of P, whose complements are its ideals, so $\Omega(P,2) = 1 + h(2^w - 1)$ and $w\log 2 \le H_P(1) \le 2w\log 2$. The width of P is w and its height is h, and w chains, each meeting every antichain once, form a minimum chain partition. Hence $H_P^{\mathrm{ch}}(K) = w\log\binom{h+K}{K}$, and $H_P^{\mathrm{ch}}(1) = w\log(h+1)$ lies between $w^2\log 2$ and $w(w+1)\log 2$.

For $1 \le K \le 2^{w/2}$ the binomial estimate (4) gives $\log\binom{h+K}{K} = \Theta(K[1 + \log(h/K)]) = \Theta(wK)$, since $\log(h/K)$ lies between $(w/2)\log 2$ and $w\log 2$. Hence $H_P^{\mathrm{ch}}(K) = \Theta(w^2 K)$ on that range,

and $H_P(K) = \Theta(wK)$ there as well, from below because a longest chain of P has h points and from above because $H_P(K) \le KH_P(1)$ by subadditivity. Both orders follow by evaluating each objective at one resolution. For a nondecreasing count H and every resolution K',

$$\min\Big\{2H(K'), \frac{T}{4K'^2}\Big\} \le \inf_{K\ge 1}\Big\{2H(K) + \frac{T}{4K^2}\Big\} \le 2H(K') + \frac{T}{4K'^2},$$

because the first term of the objective is at least $2H(K')$ for $K \ge K'$ and the second is at least $T/(4K'^2)$ for $K \le K'$. At $T = w^5$ and for all large w, take $K' = \lceil w^{4/3} \rceil$ for H_P and $K' = w$ for H_P^{ch}. Both lie below $2^{w/2}$, and at them the two terms of the objective are of the same order, $w^{7/3}$ in the first case and w^3 in the second. Both orders are below the cap $T = w^5$, so $r_P(w^5) = \Theta(w^{7/3})$ while $r_P^{\mathrm{ch}}(w^5) = \Theta(w^3)$. □

G Lattice-Closed Classes and Difference Constraints

The lower bound of Section 3 and the learner of Section 4 use two properties of the isotonic class: closure under pointwise maxima and minima, and closure under monotone rescaling (Remark 1). The first alone suffices for the fixed-resolution learner of Section 4.2 and for Lemma 24, which Appendix H proves from the representation lemma of this part. The lower bound also uses the second, but only through the constants 0 and 1 and the contraction $f \mapsto 1/4 + f/2$. This part proves these facts about the learner and the lower bound and applies them to the classes cut out by difference constraints $f(x) - f(x') \le b_{xx'}$.

Throughout, $\mathcal{X}$ is a finite set of n points, and for a nonempty class $\mathcal{F}$ of functions $\mathcal{X} \to [0,1]$ the protocol of Section 2 with $\mathcal{F}$ in place of $\mathcal{F}_P$ defines the regret and the minimax regret $R_T^*(\mathcal{F}, \ell)$, written $R_T^*(\mathcal{F})$ for squared loss. A class is lattice-closed if it contains the pointwise maximum and the pointwise minimum of any two of its members. It is contraction-closed if it contains $1/4 + f/2$ whenever it contains f. Isotonic classes have both properties and contain the constants 0 and 1. For a class $\mathcal{G}$ of maps $\mathcal{X} \to \{0, \dots, K\}$ we keep the notation of Section 2. The pairs $(x, j) \in \mathcal{X} \times [K]$ are cells, the cells of a fixed x form its column, and $I_F = \{(x, j) : j \le F(x)\}$ for $F \in \mathcal{G}$.

The next lemma is a concrete form of Birkhoff's representation of finite distributive lattices by the ideals of a poset (Birkhoff, 1937).

Lemma 23 (Ideal Representation of a Lattice-Closed Class)**.** *Let $K \ge 1$ and let $\mathcal{G}$ be a nonempty lattice-closed class of maps $\mathcal{X} \to \{0, \dots, K\}$. Call a cell free if it lies in I_F for some but not for every $F \in \mathcal{G}$, and write $\Pi_{\mathcal{G}}$ for the set of free cells. For $q, q' \in \Pi_{\mathcal{G}}$ put $q \le q'$ when $q \in I_F$ for every $F \in \mathcal{G}$ with $q' \in I_F$. Then this relation is a preorder, and $F \mapsto I_F \cap \Pi_{\mathcal{G}}$ is a bijection from $\mathcal{G}$ onto the downward closed subsets of $\Pi_{\mathcal{G}}$. Each of them is a union of equivalence classes of the preorder, so it corresponds to an ideal of the quotient poset $Q_{\mathcal{G}}$ on these classes, and $|\mathcal{G}| = |\mathcal{J}(Q_{\mathcal{G}})|$. Write $F_{\min}$ and $F_{\max}$ for the pointwise minimum and the pointwise maximum of $\mathcal{G}$, both members of $\mathcal{G}$. The free cells of the column of x are $(x, F_{\min}(x) + 1), \dots, (x, F_{\max}(x))$. They are totally ordered by the preorder, every I_F meets them in a prefix, and*

$$F(x) = F_{\min}(x) + \big|I_F \cap \Pi_{\mathcal{G}} \cap (\{x\} \times [K])\big|.$$

Proof. The images $I_F \cap \Pi_{\mathcal{G}}$ form a family closed under unions and intersections, and every downward closed subset is the union, over its cells q, of the least member of that family containing q. Write $\mathcal{L}$ for the family. It is closed under finite unions and intersections because $I_{\max\{F,F'\}} = I_F \cup I_{F'}$ and $I_{\min\{F,F'\}} = I_F \cap I_{F'}$. The class is finite and lattice-closed, so $F_{\min}$ and $F_{\max}$ lie in it. The ideal of $F_{\min}$ is contained in every other, so it contributes $\emptyset$, and $F_{\max}$ contributes $\Pi_{\mathcal{G}}$.

The relation is reflexive and transitive, and its definition makes each member of $\mathcal{L}$ downward closed. Conversely let $J \subseteq \Pi_{\mathcal{G}}$ be downward closed. For a free cell q the intersection of the members of $\mathcal{L}$ that contain q is again a member, and it equals $\{q' \in \Pi_{\mathcal{G}} : q' \le q\}$. A cell $q' \le q$ lies in every I_F containing q, and a cell lying in every member that contains q satisfies $q' \le q$ because each F with $q \in I_F$ contributes such a member. That set is contained in J whenever $q \in J$, so J is the union over its cells of these sets, hence a member of $\mathcal{L}$.

A member's ideal meets the column of x in the cells $(x, 1), \dots, (x, F(x))$, so a cell (x, j) of that column is free exactly when $F_{\min}(x) < j \le F_{\max}(x)$, the displayed formula holds, and the map is

injective. Free cells of one column are comparable, since an ideal containing (x, j) contains (x, j') for $j' \le j$, so every I_F meets them in a prefix. Finally the downward closed subsets of a preorder are the unions of its equivalence classes that are downward closed, that is, the ideals of the quotient poset. □

The fixed-resolution learner of Section 4.2 uses the isotonic class only through the correspondence of Lemma 6, the count of ideals and the three facts listed in the proof of (b) below, and the lower bound of Section 3 uses in addition that the class contains the contracted comparators $1/4+F/(2K)$ and the constants $1/4$ and $3/4$. Both transfer to $Q_{\mathcal{G}}$.

Proposition 6 (Lattice-Closed Classes at a Fixed Resolution)**.** *Let $K \ge 1$ and let $\mathcal{G}$ be a nonempty lattice-closed class of maps $\mathcal{X} \to \{0, \dots, K\}$, with the poset $Q_{\mathcal{G}}$ of Lemma 23.*

(a) Let $\mathcal{F}$ be a class of functions $\mathcal{X} \to [0, 1]$ that contains $1/4 + F/(2K)$ for every $F \in \mathcal{G}$ and contains the constants $1/4$ and $3/4$. Then for every loss ℓ that is α-strongly proper on $[1/4, 3/4]$ and every $T \ge 1$,

$$R_T^*(\mathcal{F}, \ell) \ge \frac{3\alpha\lambda}{2048} \min\Big\{\log|\mathcal{G}|,\ \frac{T}{K^2}\Big\}.$$

(b) Let $G = \{z_0 < \cdots < z_K\} \subseteq [0, 1]$ be a grid and let ℓ be an η-exp-concave loss finite on $[z_0, z_K]$. Then Proposition 5 holds with $Q_{\mathcal{G}}$ in place of Q_K, the class of the functions $x \mapsto z_{F(x)}$ with $F \in \mathcal{G}$ in place of $\mathcal{F}_G$, and $\log|\mathcal{G}|$ in place of $H_P(K)$. Given a digraph generating the preorder of Lemma 23, a maximum-weight ideal of $Q_{\mathcal{G}}$, hence a perturbed maximizer, is computed by one minimum cut.

Proof. (a) The proof of Proposition 1 applies with $\mathcal{G}$ as the latent class. Draw F uniformly from $\mathcal{G}$, so that its free cells form a uniformly random ideal of $Q_{\mathcal{G}}$ (Lemma 23), and take $1/4 + F/(2K)$ as the latent comparator, which lies in $\mathcal{F}$ by hypothesis, so that the hindsight minimum is at most its cumulative loss. Parts (b) and (c) of Lemma 6 hold for $Q_{\mathcal{G}}$, because the proof of (b) uses only that Q_K is a finite poset, and that of (c) only that the cells of a column form a chain met by every ideal in a prefix, which Lemma 23 supplies. A central element of the residual poset is a class of free cells. Querying a point x with (x, j) in that class and revealing $F(x)$ answers whether the element lies in the hidden ideal, since $F(x) \ge j$ exactly then, so Lemma 7 bounds the expected excess loss of the block from below as it does there, with $V = F(x)$ and the threshold j. Over the $M = \lfloor T/(32K^2) \rfloor$ blocks the entropy revealed is at least $\min\{\log|\mathcal{G}|, \lambda M\}$, and the instance for $T < 32K^2$ uses the constants $1/4$ and $3/4$, which lie in $\mathcal{F}$.

(b) Definition 3 and Appendix D apply to any finite poset in place of Q_K. Fix a minimum chain partition of $Q_{\mathcal{G}}$, put independent centered Gumbel noise on the prefix lengths of its chains, and let a score array assign $\theta_{x,j}$ to the point x and the level j, the score of a member being $\sum_x \theta_{x,F(x)}$. The arguments transfer through three facts. First, the score of a member telescopes along the free cells of each column,

$$\sum_x \theta_{x,F(x)} = \sum_x \theta_{x,F_{\min}(x)} + \sum_{(x,j)\in I_F \cap \Pi_{\mathcal{G}}} \big(\theta_{x,j} - \theta_{x,j-1}\big),$$

and its noise along the prefix of each chain, as in (5). An element of $Q_{\mathcal{G}}$ therefore carries the sum of these differences over the cells it identifies, together with one noise increment for its position in its chain, and the perturbed score is a sum of these node weights over the ideal plus a term that does not depend on F. So a maximum-weight member is a maximum-weight ideal of $Q_{\mathcal{G}}$, which the reduction of Picard (1976) computes by one minimum cut. Second, distinct subsets of a maximum antichain of $Q_{\mathcal{G}}$ have distinct down-closures, which is all the proof of Lemma 14 uses about Q_K. Hence $\log|\mathcal{G}| \le \Phi(0) \le c_\Phi \log|\mathcal{G}|$. Third, write $\mathcal{A}_{x,j}$ for the set of members with $F(x) \ge j$, so that $\mathcal{A}_{x,0} = \mathcal{G}$ and $\mathcal{A}_{x,K+1} = \emptyset$. For $1 \le j \le K$ it is all of $\mathcal{G}$ when $j \le F_{\min}(x)$, is empty when $j > F_{\max}(x)$, and otherwise consists of the members whose ideal contains the free cell (x, j), a condition on the prefix length of one chain. The set of members with $F(x) = j$ is $\mathcal{A}_{x,j} \setminus \mathcal{A}_{x,j+1}$, hence a cylinder on at most two chains, and Lemma 17 and Theorem 3 hold with $\kappa = 4$. At $K = 1$ a level at x that is constant on $\mathcal{G}$ gives equality in Theorem 3 with $\kappa = 1$. Otherwise the level is decided by the free cell $(x, 1)$ alone, hence by one prefix length, and Lemma 18 gives $\kappa = 1$. With these three facts the proof of Proposition 5 applies word for word. □

If $\mathcal{F}$ is lattice-closed and contraction-closed and contains the constants 0 and 1, the class $\mathcal{G}$ of the integer forms Kf of its members with values in G_K is lattice-closed and satisfies the hypothesis of (a), because $1/4 + F/(2K)$ is the image of $F/K \in \mathcal{F}$ under the contraction and the constants $1/4$ and $3/4$ are the images of 0 and 1. So such a class obeys the lower bound of (a) with $\log|\mathcal{G}|$ at every resolution.

Difference constraints define the family of such classes in which the poset and the minimum cut are explicit.

Theorem 9 (Difference Constraints). *Let $b_{xx'} \in [0,\infty]$ for every ordered pair of points of $\mathcal{X}$, and let*

$$\mathcal{F}_b = \big\{ f\colon \mathcal{X} \to [0,1] : f(x) - f(x') \le b_{xx'} \text{ for all } x, x' \in \mathcal{X} \big\}.$$

For $K \ge 1$ let $\mathcal{G}_b(K)$ be the class of the maps $F\colon \mathcal{X} \to \{0,\dots,K\}$ with $F(x) - F(x') \le \lceil K b_{xx'} \rceil$ for all x, x', and put $H_b(K) = \log|\mathcal{G}_b(K)|$. As for $\mathcal{F}_P$, put $\Psi_b(T) = \inf_{K\ge1}\{2H_b(K) + T/(4K^2)\}$ and $r_b(T) = \min\{T, \Psi_b(T)\}$. Let $b_{\min}$ be the least finite positive value of b and $K_0 = \lceil 1/(2b_{\min}) \rceil$, with $K_0 = 1$ if b has no finite positive value. In the computational assertions the finite values of b are rational and ∞ is a symbol, so that $\lceil K b_{xx'} \rceil$ is obtained by integer division.

(a) $\mathcal{F}_b$ is lattice-closed and contraction-closed and contains the constants 0 and 1. $\mathcal{G}_b(K)$ is lattice-closed and contains the constants $0,\dots,K$, so $H_b(K) \ge \log(K+1)$, and H_b is nondecreasing. The sets I_F with $F \in \mathcal{G}_b(K)$ are exactly the closed sets of the digraph $\mathcal{D}_{b,K}$ on $\mathcal{X} \times [K]$, the sets of cells that contain the head of every arc leaving them, where $\mathcal{D}_{b,K}$ has an arc from (x,j) to $(x,j-1)$ for $2 \le j \le K$ and an arc from (x,j) to $(x', j - \lceil K b_{xx'} \rceil)$ for every pair with $b_{xx'} < \infty$ and every j with $j - \lceil K b_{xx'} \rceil \ge 1$. A maximum-weight member of $\mathcal{G}_b(K)$ under weights on the cells is computed by one minimum cut on a network with $nK + 2$ vertices and $O(n^2 K)$ arcs.

(b) For every $K \ge 1$, every $T \ge 1$ and every finite η-exp-concave β-smooth loss, $R_T^*(\mathcal{F}_b, \ell) \le H_b(K)/\eta + \beta T/(8K^2)$, and for squared loss $R_T^*(\mathcal{F}_b) \le r_b(T)$.

(c) For every $K \ge K_0$, every $T \ge 1$ and every loss ℓ that is α-strongly proper on $[1/4, 3/4]$,

$$R_T^*(\mathcal{F}_b, \ell) \ge \frac{3\alpha\lambda}{2048} \min\Big\{ H_b(K),\ \frac{T}{K^2} \Big\}.$$

(d) For squared loss and every $T \ge K_0^2 H_b(K_0)$,

$$\frac{\lambda}{2048}\, r_b(T) \le R_T^*(\mathcal{F}_b) \le r_b(T),$$

and for a loss that is α-strongly proper on $[1/4, 3/4]$ the lower constant is $\alpha\lambda/2048$.

(e) Let ℓ be η-exp-concave and β-smooth with values in $[0,1]$, let $T \ge 1$ and $\delta \in (0,1)$. For every $K \le \lceil \sqrt{T} \rceil$, run the learner of Proposition 6(b) on $\mathcal{G}_b(K)$ and the uniform grid G_K in pair mode with $S = 2\lceil 32 \log(1/\delta)/\min\{\eta, 1\} \rceil$ samples per round. Aggregate their predictions and the constant prediction $1/2$ by exponential weights at rate η, with the prior $1/(K(K+1))$ on the learner of resolution K and the remaining mass on the constant. The resulting learner satisfies

$$\mathbb{E} R_T \le C_\ell\, r_b(T), \qquad \mathbb{P}\{ R_t \le C_\ell\, r_b(T) \text{ for all } t \le T \} \ge 1 - \delta,$$

with C_ℓ depending only on (η, β). Its computation per round is polynomial in n, T, the input bit lengths and $\log(1/\delta)$, and it uses $O(\sqrt{T}(1 + \log(1/\delta))/\min\{\eta, 1\})$ minimum cuts per round.

The isotonic class of a poset is the case $b_{xx'} = 0$ for $x \preceq x'$ and $b_{xx'} = \infty$ otherwise. Then $K_0 = 1$, the class $\mathcal{G}_b(K)$ consists of the integer isotonic maps, and $H_b = H_P$.

Proof. The digraph represents the integer classes, common-shift rounding gives (b), Proposition 6 gives (c) and (e), and the largest resolution with $H_b(K) \le T/K^2$ gives (d).

(a) If $\max\{f, f'\}$ takes the value $f(x)$ at x, then $\max\{f, f'\}(x) - \max\{f, f'\}(x') \le f(x) - f(x') \le b_{xx'}$. For minima, the function attaining $\min\{f, f'\}$ at x' gives the same bound, and the same

computation applies to integer maps. Contraction holds because $(1/4+f(x)/2)-(1/4+f(x')/2)$ is half of $f(x)-f(x')$, hence at most $b_{xx'}$, and the constants belong to both classes because $b \geq 0$, which gives $|\mathcal{G}_b(K)| \geq K+1$. Monotonicity follows from $\lceil Kb_{xx'} \rceil \leq \lceil (K+1)b_{xx'} \rceil$, which embeds $\mathcal{G}_b(K)$ in $\mathcal{G}_b(K+1)$.

Write $\bar{b} = \lceil Kb_{xx'} \rceil$ for a pair with $b_{xx'} < \infty$. If $F \in \mathcal{G}_b(K)$ then I_F contains $(x, j-1)$ whenever it contains (x, j), and if $(x, j) \in I_F$ and $j - \bar{b} \geq 1$ then $F(x') \geq F(x) - \bar{b} \geq j - \bar{b}$, so $(x', j-\bar{b}) \in I_F$, and I_F is closed. Conversely a closed set I meets every column in a prefix, so $I = I_F$ with $F(x) = |I \cap (\{x\} \times [K])|$. If $F(x) - \bar{b} \geq 1$ then $(x, F(x)) \in I$ forces $(x', F(x) - \bar{b}) \in I$ and $F(x') \geq F(x) - \bar{b}$, while if $F(x) \leq \bar{b}$ the constraint holds because $F(x') \geq 0$. The closed sets of a digraph are the downward closed subsets of its reachability preorder, in which the head of every arc lies below its tail. Since a preorder is recovered from its downward closed subsets, that preorder is the one of Lemma 23, and every cell is free here because the constants 0 and K lie in $\mathcal{G}_b(K)$. A maximum-weight closed set is computed by one minimum cut (Picard, 1976) on the network of Appendix D, built from the arcs of $\mathcal{D}_{b,K}$. It has $nK+2$ vertices, and $\mathcal{D}_{b,K}$ has $O(n^2K)$ arcs.

(b) Repeat the proof of Theorem 5 with the expert set $\mathcal{G}_b(K)/K$, over which exponential weights costs $H_b(K)/\eta$ (Lemma 3). Common-shift rounding carries the comparator into that set. For $f \in \mathcal{F}_b$ and a shift υ, the difference $\lfloor Kf(x) + \upsilon \rfloor - \lfloor Kf(x') + \upsilon \rfloor$ is an integer smaller than $K(f(x) - f(x')) + 1$, hence at most $\lceil K(f(x) - f(x')) \rceil \leq \lceil Kb_{xx'} \rceil$, so $K\lfloor f \rfloor_\upsilon \in \mathcal{G}_b(K)$, and Lemma 2 supplies a shift that raises the cumulative loss of f by at most $\beta T/(8K^2)$. For squared loss, $\eta = 1/2$ and $\beta = 2$ give the objective $2H_b(K) + T/(4K^2)$, whose infimum is attained because $H_b(K) \geq \log(K+1)$, and the loss takes values in $[0,1]$, so the regret is also at most T.

(c) Apply Proposition 6(a) with $\mathcal{G} = \mathcal{G}_b(K)$, which is lattice-closed by (a), and with $\mathcal{F} = \mathcal{F}_b$, which contains $1/4$ and $3/4$. It also contains $1/4 + F/(2K)$ for every $F \in \mathcal{G}_b(K)$. A pair with $b_{xx'} = 0$ has $F(x) \leq F(x')$, a pair with $b_{xx'} = \infty$ imposes nothing, and for $0 < b_{xx'} < \infty$,

$$\frac{F(x) - F(x')}{2K} \leq \frac{\lceil Kb_{xx'} \rceil}{2K} \leq \frac{2Kb_{xx'}}{2K} = b_{xx'},$$

because $Kb_{xx'} \geq K_0 b_{\min} \geq 1/2$ and $\lceil \varsigma \rceil \leq 2\varsigma$ for $\varsigma \geq 1/2$.

(d) Let $K_\star$ be the largest resolution with $H_b(K_\star) \leq T/K_\star^2$. It exists because K_0 qualifies, by the hypothesis on T, and because $H_b(K) \geq \log(K+1)$ leaves only finitely many candidates, and $K_\star \geq K_0$. Apply (c) at $K_\star$ and at $K_\star + 1$, both at least K_0. At $K_\star$ the minimum in (c) is $H_b(K_\star)$, and at $K_\star+1$ it is $T/(K_\star+1)^2 \geq T/(4K_\star^2)$, because $H_b(K_\star+1) > T/(K_\star+1)^2$ by the maximality of $K_\star$. Hence

$$R_T^*(\mathcal{F}_b, \ell) \geq \frac{3\alpha\lambda}{2048} \max\Big\{ H_b(K_\star), \frac{T}{4K_\star^2} \Big\} \geq \frac{\alpha\lambda}{2048} \Big(2H_b(K_\star) + \frac{T}{4K_\star^2} \Big),$$

which is at least $(\alpha\lambda/2048)\Psi_b(T) \geq (\alpha\lambda/2048) r_b(T)$. Squared loss is 1-strongly proper, and the upper bound is (b).

(e) Fix $K \leq \lceil \sqrt{T} \rceil$. By (a) and Proposition 6(b), the learner of resolution K has regret against $\mathcal{G}_b(K)/K$ at most $64 c_\Phi H_b(K)/\min\{\eta, 1\}$ in expectation and, with probability at least $1-\delta$, at most $64 c_\Phi H_b(K)/\min\{\eta,1\} + 1$ at every prefix, the sample count S being even and making the confidence term $64\log(1/\delta)/(S\min\{\eta,1\})$ at most one. The master pays $\log(K(K+1))/\eta \leq 2H_b(K)/\eta$ against that learner on every realization and at every prefix (Lemma 3), and rounding the comparator costs at most $\beta t/(8K^2)$ at a prefix of length t by (b). So the regret against $\mathcal{F}_b$ at every prefix $t \leq T$ is at most $A_K + 1$ on the event of that learner, and its expectation at T is at most A_K, where

$$A_K = \Big(\frac{64 c_\Phi}{\min\{\eta, 1\}} + \frac{2}{\eta} \Big) H_b(K) + \frac{\beta T}{8K^2}.$$

Let K' minimize A_K over $K \leq \lceil \sqrt{T} \rceil$. It is determined by b, T, η and β, and the learner runs every resolution, so the two bounds are applied to the single resolution K' and no union bound is needed. Since H_b is nondecreasing and $T/(4\lceil \sqrt{T} \rceil^2) \leq 1/4$, restricting the infimum that defines $\Psi_b(T)$ to $K \leq \lceil \sqrt{T} \rceil$ raises it by at most $1/4$. As A_K is at most a constant depending only on (η, β) times $2H_b(K) + T/(4K^2)$, and $\Psi_b(T) \geq 2H_b(1) \geq 2\log 2$, this gives $A_{K'} + 1 \leq C_\ell \Psi_b(T)$. The loss has values in $[0,1]$, so $R_t \leq t \leq T$ as well, and both bounds hold with $r_b(T)$. Each of the $\lceil \sqrt{T} \rceil$ learners draws $S = O((1 + \log(1/\delta))/\min\{\eta, 1\})$ perturbed maximizers per round, one

minimum cut each, and $\mathcal{D}_{b,K}$, the preorder it generates, the quotient poset and a minimum chain partition of that poset are computed in polynomial time. Theorem 12 gives the implementation in the finite-precision model. □

The horizon-free construction of Section 4.3 rests on an inequality of isotonic classes. Its delayed activation is paid for by $n\log(1+K/n) \le H_P(K)$, which fails for the constant functions on two or more points, the case $b \equiv 0$, where $H_b(K) = \log(K+1)$.

H Pinball Loss

This part determines the minimax regret under pinball loss up to absolute constants and gives a polynomial-time learner with regret of the same order. It first fixes the notions it uses beyond those of the main text, then states the results with a sketch of the lower bound, and then proves Lemma 24, the rate (Lemma 25 and Proposition 7) and the guarantee of the learner (Lemma 26 and Theorem 11).

Pinball loss at quantile level $\tau \in (0,1)$ is $\ell_\tau(p,y) = \tau(y-p)_+ + (1-\tau)(p-y)_+$ (Koenker & Bassett, 1978; Steinwart & Christmann, 2011), and $\rho = \tau(1-\tau)$. A loss that is convex in its first argument has slope bounds $g_-, g_+ > 0$ if at every (p,y) it admits a subgradient in $[-g_-, g_+]$. The subgradient in this interval that a learner uses is its selected subgradient, and $\partial_p \ell(p,y)$ denotes the subdifferential of a convex loss in its first argument. Pinball loss has slope bounds τ and $1-\tau$. The class of binary isotonic functions is denoted by $\mathcal{F}_1$, the down-closure and the up-closure of an element q of a poset by $\downarrow q$ and $\uparrow q$, and $\lambda' = \log(1/\delta_{\mathrm{LS}}) = 1.774\ldots$.

For a nonempty class $\mathcal{G}$ of functions from a finite set to $\{0,1\}$, a mistake tree of depth D is a binary tree all of whose leaves have depth D, whose internal nodes are labeled by points and whose two edges out of each internal node are labeled 0 and 1, such that every root-to-leaf path is realized by a member of $\mathcal{G}$, one taking at each node's point the label of the edge the path follows. The Littlestone dimension $\mathrm{Ldim}(\mathcal{G})$ is the largest depth of a mistake tree (Littlestone, 1988; Ben-David et al., 2009). This part uses it for $\mathcal{F}_1$.

Pinball loss is the loss of quantile estimation. Under a label law, the predictions that minimize its expected loss are the τ-quantiles of the law. Offline, estimating them as an isotonic function of the point is a special case of isotonic distributional regression, the estimation of a conditional distribution under order constraints (Henzi et al., 2021), and Jordan et al. (2022) characterize the solutions of the quantile problem. The loss is piecewise linear, and its minimax rate involves the order polynomial only through its value at resolution one.

The reason is that the regret of a convex loss is at most that of its linearization. With g_t a selected subgradient at the learner's own prediction, convexity gives $\ell(\widehat{y}_t, y_t) - \ell(f(x_t), y_t) \le g_t(\widehat{y}_t - f(x_t))$ at every round, so the regret against f is at most the linearized regret $\sum_{t\le T} g_t(\widehat{y}_t - f(x_t))$. An isotonic function is the average of the indicators of its superlevel sets, $f = \int_0^1 \mathbf{1}\{f \ge z\}\,dz$, and those indicators are isotonic as well, so a linear functional on $\mathcal{F}_P$ attains its minimum on $\mathcal{F}_1$. The members of $\mathcal{F}_1$ are the indicators of the filters of P, which correspond to the ideals of Q_1, and there are $e^{H_P(1)}$ of them. Exponential weights over $\mathcal{F}_1$ at rate $\eta \le 1$, run on the linearized losses, has linearized regret at most $\eta\rho T + 2H_P(1)/\eta$ for pinball loss (proof of Lemma 25). The choice of η that balances the two terms, admissible once $\rho T \ge 2H_P(1)$, gives $2\sqrt{2\rho T H_P(1)}$, and the constant prediction τ, whose regret is at most ρ per round, gives the other branch of Theorem 10.

The lower bound measures $\mathcal{F}_1$ by its Littlestone dimension. Recall that a class of functions is lattice-closed if it contains the pointwise maximum and the pointwise minimum of any two of its members (Appendix G). Both $\mathcal{F}_P$ and $\mathcal{F}_1$ are lattice-closed.

Lemma 24 (Littlestone Dimension of Lattice-Closed Binary Classes)**.** *Let $\mathcal{X}$ be a finite set and $\mathcal{G}$ a nonempty lattice-closed class of functions $\mathcal{X} \to \{0,1\}$. Then*

$$\frac{\log|\mathcal{G}|}{\lambda'} \le \mathrm{Ldim}(\mathcal{G}) \le \frac{\log|\mathcal{G}|}{\log 2}.$$

In particular $H_P(1)/\lambda' \le \mathrm{Ldim}(\mathcal{F}_1) \le H_P(1)/\log 2$.

The Littlestone dimension determines the optimal regret of online classification (Littlestone, 1988; Ben-David et al., 2009; Alon et al., 2021), and for a general class it can lie far below $\log|\mathcal{G}|$. The

indicators of the singletons of a finite set $\mathcal{X}$ with $|\mathcal{X}| \geq 2$ form a class of $|\mathcal{X}|$ members whose Littlestone dimension is one. For a lattice-closed class the two agree up to the factor $\lambda'/\log 2$, because the search by central elements (Section 2) applies to it. Dropping the points at which all members agree and identifying points with the same membership pattern turns the class into the indicators of the ideals of a poset (Lemma 23), and a query at a central element of that poset keeps at least a δ_{LS} fraction of the members on either answer.

Exponential weights over $\mathcal{F}_1$ attains the upper bound, but its first prediction can already determine the number of ideals of a poset, a #P-complete count (Appendix D). The learner of Theorem 10 is the perturbed-ideal learner at resolution one, run on the linearized losses. It sets its update rate from one zero-score estimate in place of $H_P(1)$, and it corrects its prediction for the two slopes of the loss. Algorithm 4 states it in pseudocode. Given a design horizon T and a confidence $\delta \in (0,1)$, it proceeds as follows.

(1) If $\rho T \leq 4$, predict τ at every round.

(2) Otherwise fix a minimum chain partition of Q_1, compute the zero-score estimate $\widetilde{\Phi}$ at resolution one with $\varepsilon = 1$ and confidence $\min\{\delta, (T+1)^{-3}\}/2$ (Remark 4), and put $\Phi_+ = \max\{1, \widetilde{\Phi}\}$. If $\rho T \leq 4\Phi_+$, predict τ at every round.

(3) Otherwise put $\eta = 2\sqrt{\Phi_+/(\rho T)}$, which is less than one, and $\theta_1 = 0$. At round t, draw $S = \lceil 2T^2 \log(4T(T+1)^3/\delta) \rceil$ independent perturbed maximizers at θ_t, each computed by one minimum cut on a network with $n+2$ vertices. Let $\widetilde{\mu}_{t,1}$ be the fraction of them that contain the cell $(x_t, 1)$, and predict

$$\widehat{y}_t = \widetilde{\mu}_{t,1} + \eta(2\tau - 1)\,\widetilde{\mu}_{t,1}(1 - \widetilde{\mu}_{t,1}).$$

When y_t is revealed, lower the score $\theta_{x_t,1}$ by ηg_t, where g_t is the selected subgradient of $\ell_\tau(\cdot, y_t)$ at $\widehat{y}_t$.

We call $\widehat{y}_t$ the corrected prediction. The fraction $\widetilde{\mu}_{t,1}$ estimates $\mu_{t,1}$, the probability that the perturbed maximizer at θ_t contains $(x_t, 1)$. By Theorem 3 with $\kappa = 1$ and a quadratic bound on the exponential, the update lowers Φ by at least $\eta g_t \mu_{t,1} - \eta^2 g_t^2 \mu_{t,1}(1-\mu_{t,1})$. The linearized loss $g_t \mu_{t,1}$ of the uncorrected prediction therefore exceeds this decrease divided by η by at most the second-order term $\eta g_t^2 \mu_{t,1}(1-\mu_{t,1})$. That term can be as large as $\eta \max\{\tau, 1-\tau\}^2/4$, and balancing η against it would give a regret bound of order $\max\{\tau, 1-\tau\}\sqrt{T H_P(1)}$ in place of $\sqrt{\rho T H_P(1)}$. The shift by $\eta(2\tau-1)\mu_{t,1}(1-\mu_{t,1})$ turns the term into $\eta(g_t^2 + (2\tau-1)g_t)\mu_{t,1}(1-\mu_{t,1})$, which is at most $\eta\rho\mu_{t,1}(1-\mu_{t,1})$ for every $g_t \in [-\tau, 1-\tau]$, since $g_t^2 + (2\tau-1)g_t - \rho = (g_t + \tau)(g_t - 1 + \tau) \leq 0$ (Lemma 26).

With probability at least $1-\delta$, the estimate of step (2), when it is computed, satisfies $\Phi(0) \leq \Phi_+ < 6H_P(1)$ (Remark 4 and Lemma 14), and every $\widetilde{\mu}_{t,1}$ of step (3) is within $1/(2T)$ of $\mu_{t,1}$. On this event the potential of step (3) telescopes as in Proposition 5, and the linearized regret is at most $\Phi(0)/\eta + \eta\rho T/4 + 1 \leq \sqrt{\Phi_+ \rho T} + 1$, the term 1 paying for the errors of the estimates. Together with the constant predictions of steps (1) and (2), this bounds the regret by $5\min\{\rho T, \sqrt{\rho T H_P(1)}\}$, and the expected regret is at most $6\min\{\rho T, \sqrt{\rho T H_P(1)}\}$, because the failure probabilities are capped by $(T+1)^{-3}$ (Theorem 11). Restarting on epochs of lengths $1, 2, 4, \ldots$, with confidence $\delta 2^{-\iota}$ in the epoch of length $2^{\iota-1}$, removes the design horizon (Corollary 5).

Theorem 10 (Minimax Rate under Pinball Loss)**.** *Let $\tau \in (0,1)$, $\rho = \tau(1-\tau)$ and $\lambda' = \log(1/\delta_{\mathrm{LS}}) = 1.774\ldots$. For every finite poset P and every $T \geq 1$,*

$$\tfrac{1}{8}\min\bigl\{\rho T, \sqrt{\rho T\ \mathrm{Ldim}(\mathcal{F}_1)}\bigr\} \leq R_T^*(P, \ell_\tau) \leq \min\bigl\{\rho T, 2\sqrt{2\rho T H_P(1)}\bigr\};$$

with Lemma 24, $0.093\min\{\rho T, \sqrt{\rho T H_P(1)}\} \leq R_T^*(P, \ell_\tau) \leq 2\sqrt{2}\min\{\rho T, \sqrt{\rho T H_P(1)}\}$. On a chain of n points the rate is $\Theta(\min\{\rho T, \sqrt{\rho T \log(n+1)}\})$. Given a design horizon T and a confidence $\delta \in (0,1)$, the perturbed-ideal learner at resolution one with the corrected prediction runs in polynomial time, uses at most one estimate of $\Phi(0)$ at resolution one, and never computes $H_P(1)$. Its regret is $O(\min\{\rho T, \sqrt{\rho T H_P(1)}\})$ with probability at least $1-\delta$ and in expectation, with the constants 5 and 6, respectively, in exact arithmetic (Theorem 11). Restarted on epochs of lengths $1, 2, 4, \ldots$, it has regret of the same order at every T, in expectation and, with probability at least $1-\delta$, simultaneously for all T.

In words, the minimax regret is of order ρT while $T \le H_P(1)/\rho$ and of order $\sqrt{\rho T H_P(1)}$ beyond. The complexity measure is a single number, the logarithm of the number of filters of P, whereas Theorems 1 and 2 need the whole function $K \mapsto H_P(K)$. The lower bound is stated with the Littlestone dimension of $\mathcal{F}_1$, which Lemma 24 places between $H_P(1)/\lambda'$ and $H_P(1)/\log 2$.

Sketch of the lower bound. The search by central elements (Section 2), run on the filters, gives $\mathrm{Ldim}(\mathcal{F}_1) \ge H_P(1)/\lambda'$, and independent $\mathrm{Ber}(1-\tau)$ labels queried along a mistake tree of depth $\mathrm{Ldim}(\mathcal{F}_1)$ give the lower bound. A central element ζ of P lies in a fraction of the ideals within $[\delta_{\mathrm{LS}}, 1-\delta_{\mathrm{LS}}]$, hence in such a fraction of the filters, their complements. The filters containing ζ are obtained from those of $P \setminus \uparrow\zeta$ by adjoining $\uparrow\zeta$, and the filters avoiding ζ are exactly the filters of $P \setminus \downarrow\zeta$. Each answer therefore leaves the filters of a smaller poset, at least a δ_{LS} fraction of those still compatible, so every node at depth below $H_P(1)/\lambda'$ still carries more than one filter and the recursion builds a mistake tree of depth $\lceil H_P(1)/\lambda' \rceil$. In the other direction, the 2^D leaves of a tree of depth D are realized by distinct members. This is Lemma 24 for $\mathcal{F}_1$, proved below for every lattice-closed class.

Given a mistake tree of depth $D = \mathrm{Ldim}(\mathcal{F}_1)$, let every label be an independent $\mathrm{Ber}(1-\tau)$. Every prediction p then has expected loss $(1-\tau)\tau(1-p) + \tau(1-\tau)p = \rho$, so the expected regret is at least ρT minus the expected loss of any member of $\mathcal{F}_1$ that the labels select. If $T < D$, the adversary queries one node per round and follows the observed label. A member realizing the path taken reproduces every label and pays nothing, so $\mathbb{E}R_T \ge \rho T$. If $T \ge D$, it queries the point of each node $\varphi = \lfloor T/D \rfloor$ times, then follows the edge of the cheaper value, and spends the remaining rounds at a fixed point. With U_ι ones observed at node ι, a member taking the value 0 there pays τU_ι over those rounds and one taking the value 1 pays $(1-\tau)(\varphi - U_\iota)$, and the smaller of the two is

$$\min\{\tau U_\iota, (1-\tau)(\varphi - U_\iota)\} = \rho\varphi + \frac{(2\tau-1)\big(U_\iota - (1-\tau)\varphi\big)}{2} - \frac{\big|U_\iota - (1-\tau)\varphi\big|}{2}.$$

The middle term has mean zero, so at every node the member at the leaf gains half the expected absolute deviation of the count, at least $\frac{1}{4}\min\{\rho\varphi, \sqrt{\rho\varphi}\}$ by a fourth-moment estimate. Summing over the D nodes and using $\varphi \ge T/(2D)$ gives $\frac{1}{8}\min\{\rho T, \sqrt{\rho D T}\}$. The queried points and the labels are functions of the label coins alone, so the adversary is oblivious.

The lower bound uses only two properties of pinball loss under $\mathrm{Ber}(1-\tau)$ labels. The predictions 0 and 1 have the same expected loss, which no prediction beats, and the difference of their losses is random. Convex Lipschitz losses with two such predictions obey the same law. To state this, call a prediction Bayes-optimal under a loss and a label law if it minimizes the expected loss, and that minimum is the Bayes risk.

Corollary 3 (Lipschitz Losses with Two Bayes-Optimal Predictions)**.** *Let ℓ be a finite convex loss with slope bounds $g_-, g_+ > 0$, and suppose that under the law of a label Y two distinct predictions $p_0 < p_1$ are Bayes-optimal with a common finite Bayes risk, and that $Z = \ell(p_0, Y) - \ell(p_1, Y)$ satisfies $|Z| \le \bar{Z}$, $\mathbb{E}Z = 0$ and $\mathrm{Var}\, Z = \sigma_\ell^2 > 0$. Then for every finite poset and every $T \ge 1$,*

$$\frac{1}{8}\min\Big\{\frac{\sigma_\ell^2 T}{\bar{Z}},\ \frac{\sigma_\ell}{\sqrt{\lambda'}}\sqrt{T H_P(1)}\Big\} \le R_T^*(P,\ell) \le \min\Big\{\frac{g_- g_+ T}{g_- + g_+},\ 2\sqrt{2 g_- g_+ T H_P(1)}\Big\},$$

so $R_T^*(P,\ell) = \Theta_\ell(\min\{T, \sqrt{T H_P(1)}\})$ with constants depending only on $(g_-, g_+, \sigma_\ell, \bar{Z})$, and, given a polynomial-time subgradient oracle, a design horizon T and $\delta \in (0,1)$, a polynomial-time learner attains regret $O(\min\{g_- g_+ T/(g_- + g_+), \sqrt{g_- g_+ H_P(1) T}\})$ with probability at least $1-\delta$ and in expectation. Pinball loss is the case $g_- = \tau$, $g_+ = 1-\tau$, $\sigma_\ell^2 = \rho$, $\bar{Z} = \max\{\tau, 1-\tau\}$.

In its lower bound the two Bayes-optimal predictions take the place of the values 0 and 1 of a binary comparator. The branches the adversary follows determine a filter, and at each node that choice gains the comparator half the expected fluctuation of Z over the queries made there.

Running the search by central elements (Section 2) on the poset $Q_{\mathcal{G}}$ of Lemma 23 proves Lemma 24.

Proof of Lemma 24. The upper bound counts the leaves of a mistake tree. For the lower bound, fixing the value at a point leaves a lattice-closed class, so Lemma 23 represents the surviving members

by the ideals of a poset, and a central element of that poset gives a point at which they split in a balanced way.

Two root-to-leaf paths of a mistake tree part at some node, and the members realizing them differ at that node's point, so the 2^D leaves of a mistake tree of depth D need distinct members, and $2^{\mathrm{Ldim}(\mathcal{G})} \le |\mathcal{G}|$.

For the lower bound build a tree whose nodes carry subclasses of $\mathcal{G}$, the root carrying $\mathcal{G}$ itself. Let a node carry a subclass $\mathcal{G}'$ with at least two members. By Lemma 23 at $K = 1$, the members of $\mathcal{G}'$ correspond to the ideals of $Q_{\mathcal{G}'}$, which is nonempty. Let ζ be a central element of $Q_{\mathcal{G}'}$ and let x be a point whose cell $(x, 1)$ belongs to the class ζ. The node presents x, its label-1 child carries the members of $\mathcal{G}'$ with the value 1 at x and its label-0 child those with the value 0. Both subclasses are lattice-closed, and since a member has the value 1 at x exactly when its ideal contains ζ, each of them holds at least a δ_{LS} fraction of $\mathcal{G}'$. A node at depth ι therefore carries at least $\delta_{\mathrm{LS}}^{\iota}|\mathcal{G}|$ members, which exceeds one whenever $\iota < \log|\mathcal{G}|/\lambda'$, so every node at depth less than $\log|\mathcal{G}|/\lambda'$ can be split, and truncating at depth $\lceil \log|\mathcal{G}|/\lambda' \rceil$ leaves every leaf with a nonempty subclass. Any member of a leaf's subclass realizes the leaf's labeling, so the tree is a mistake tree and $\mathrm{Ldim}(\mathcal{G}) \ge \log|\mathcal{G}|/\lambda'$.

The class $\mathcal{F}_1$ is lattice-closed and has $e^{H_P(1)}$ members, which gives the last claim. □

Two facts give the bounds of Theorem 10 and Corollary 3. Against linear losses the best isotonic comparator can be taken in $\mathcal{F}_1$, which gives the upper bounds, and a label law that leaves two predictions equally good turns a mistake tree of $\mathcal{F}_1$ into a lower bound.

Lemma 25 (Convex Losses with Slope Bounds)**.** *Let ℓ be finite and convex in its first argument, with slope bounds $g_-, g_+ > 0$. For every finite poset P and every $T \ge 1$,*

$$R_T^*(P, \ell) \le \min\Big\{ \frac{g_- g_+ T}{g_- + g_+},\ 2\sqrt{2 g_- g_+ T H_P(1)} \Big\},$$

and the smaller of the two bounds is attained by a deterministic learner.

Proof. Convexity replaces the loss by its linearization, a linear functional on $\mathcal{F}_P$ is minimized at the indicator of a filter, and exponential weights over the filters pays the logarithm of their number. The constant prediction $g_-/(g_- + g_+)$ covers the horizons on which it is cheaper. Fix $f \in \mathcal{F}_P$ and let g_t be the selected subgradient at the learner's own prediction, so that at every round

$$\ell(\widehat{y}_t, y_t) - \ell(f(x_t), y_t) \le g_t\big(\widehat{y}_t - f(x_t)\big), \qquad g_t \in [-g_-, g_+]. \tag{8}$$

The constant prediction $\widehat{y}_t = g_-/(g_- + g_+)$ has $\widehat{y}_t - f(x_t) \in [-g_+/(g_- + g_+),\ g_-/(g_- + g_+)]$, so every term on the right of (8) is at most $g_- g_+/(g_- + g_+)$. Summing over the rounds and taking the infimum over f gives the first bound.

For the second bound fix $\eta \in (0, 1]$, and at round t give each member f' of $\mathcal{F}_1$ the shifted linearized loss

$$\chi_{t,f'} = \frac{g_t f'(x_t) - \min\{0, g_t\}}{g_- + g_+} \in [0, 1].$$

The learner runs exponential weights on these losses with the uniform prior. With W_t the average over $f' \in \mathcal{F}_1$ of $\exp(-\eta \sum_{s<t} \chi_{s,f'})$ and $\nu_{f',t} = \exp(-\eta \sum_{s<t} \chi_{s,f'})/(|\mathcal{F}_1|\, W_t)$ the normalized weight, it predicts $\widehat{y}_t = \sum_{f'} \nu_{f',t} f'(x_t)$. The subtracted term $\min\{0, g_t\}$ is common to all members, so it changes no normalized weight. Write $\widehat{\mathrm{Loss}} = \sum_{t \le T} \sum_{f'} \nu_{f',t} \chi_{t,f'}$ for the cumulative shifted loss of the mixture and $\mathrm{Loss}_\star = \min_{f' \in \mathcal{F}_1} \sum_{t \le T} \chi_{t,f'}$ for the least cumulative shifted loss of a member. Since $e^{-\eta\varsigma} \le 1 - (1 - e^{-\eta})\varsigma$ for $\varsigma \in [0, 1]$, by the convexity of $\varsigma \mapsto e^{-\eta\varsigma}$,

$$\frac{W_{t+1}}{W_t} = \sum_{f'} \nu_{f',t} e^{-\eta \chi_{t,f'}} \le 1 - (1 - e^{-\eta}) \sum_{f'} \nu_{f',t} \chi_{t,f'} \le \exp\Big(-(1 - e^{-\eta}) \sum_{f'} \nu_{f',t} \chi_{t,f'} \Big).$$

Multiplying over $t \le T$, with $W_1 = 1$ and $W_{T+1} \ge |\mathcal{F}_1|^{-1} e^{-\eta\, \mathrm{Loss}_\star}$, gives $(1 - e^{-\eta})\widehat{\mathrm{Loss}} \le \eta\, \mathrm{Loss}_\star + H_P(1)$, and the inequalities $\eta/(1 - e^{-\eta}) \le 1 + \eta$ and $1/(1 - e^{-\eta}) \le 2/\eta$ on $(0, 1]$ turn this into the small-loss bound (Cesa-Bianchi & Lugosi, 2006)

$$\widehat{\mathrm{Loss}} - \mathrm{Loss}_\star \le \eta\, \mathrm{Loss}_\star + \frac{2 H_P(1)}{\eta}.$$

It remains to compare the regret against f with $\widehat{\mathrm{Loss}} - \mathrm{Loss}_\star$. Pointwise $f = \int_0^1 \mathbf{1}\{f \ge z\}\,dz$, and every superlevel set of f is a filter, so $\sum_t g_t f(x_t)$ is the average over $z \in [0,1]$ of $\sum_t g_t \mathbf{1}\{f(x_t) \ge z\}$, and some member f° of $\mathcal{F}_1$ has $\sum_t g_t f^\circ(x_t) \le \sum_t g_t f(x_t)$. The mixture's shifted loss $\sum_{f'} \nu_{f',t}\chi_{t,f'}$ equals $[g_t\widehat{y}_t - \min\{0, g_t\}]/(g_- + g_+)$, so the shift cancels in

$$\sum_{t\le T} g_t\big(\widehat{y}_t - f^\circ(x_t)\big) = (g_- + g_+)\Big(\widehat{\mathrm{Loss}} - \sum_{t\le T}\chi_{t,f^\circ}\Big) \le (g_- + g_+)\big(\widehat{\mathrm{Loss}} - \mathrm{Loss}_\star\big),$$

and with (8) the regret against f is at most $(g_- + g_+)(\eta\,\mathrm{Loss}_\star + 2H_P(1)/\eta)$.

To bound $\mathrm{Loss}_\star$, use the constants 0 and 1, which lie in $\mathcal{F}_1$. At a round with $g_t \ge 0$ the constant 0 has shifted loss zero and the constant 1 at most $g_+/(g_- + g_+)$, and at a round with $g_t < 0$ the constant 1 has shifted loss zero and the constant 0 at most $g_-/(g_- + g_+)$. The average of their shifted losses with the weights $g_+/(g_- + g_+)$ on 0 and $g_-/(g_- + g_+)$ on 1 is therefore at most $g_-g_+/(g_- + g_+)^2$ at every round, and $\mathrm{Loss}_\star \le g_-g_+T/(g_- + g_+)^2$. Write X for that bound. If $X \ge 8H_P(1)$, the rate $\eta = \sqrt{2H_P(1)/X}$ is at most $1/2$ and

$$(g_- + g_+)\Big(\eta X + \frac{2H_P(1)}{\eta}\Big) = 2(g_- + g_+)\sqrt{2XH_P(1)} = 2\sqrt{2g_-g_+TH_P(1)} \le \frac{g_-g_+T}{g_- + g_+};$$

otherwise the first bound is the smaller of the two. Taking the infimum over f finishes the proof. □

Proposition 7 (Two Bayes-Optimal Predictions). *Let ℓ be a loss, and suppose that under the law of a label Y two distinct predictions $p_0 < p_1$ are Bayes-optimal with a common finite Bayes risk. Put $Z = \ell(p_0, Y) - \ell(p_1, Y)$ and suppose that $|Z| \le \bar{Z}$ and $\mathrm{Var}\,Z = \sigma_\ell^2 > 0$. Then for every finite poset P and every $T \ge 1$,*

$$R_T^*(P, \ell) \ge \frac{1}{8}\min\Big\{\frac{\sigma_\ell^2 T}{\bar{Z}},\ \sigma_\ell\sqrt{\mathrm{Ldim}(\mathcal{F}_1)\,T}\Big\}.$$

Proof. The adversary walks down a mistake tree of $\mathcal{F}_1$ with labels drawn from the law of Y, following at each node the cheaper of the two predictions. No prediction beats the Bayes risk, while the comparator read off at the leaf gains from the fluctuations of Z.

As in the proof of Proposition 1, it suffices to exhibit a random sequence drawn before play, independently of the learner's private randomness, against which every learner has the stated expected regret. Put $D = \mathrm{Ldim}(\mathcal{F}_1)$, and fix a mistake tree of depth D for $\mathcal{F}_1$ and, for each leaf, a member of $\mathcal{F}_1$ realizing its path. Both predictions have expected loss $\mathbb{E}\ell(p_0, Y)$, so $\mathbb{E}Z = 0$, and no prediction has smaller expected loss.

Suppose first that $T \ge D$, and put $\varphi = \lfloor T/D \rfloor \ge T/(2D)$. The adversary walks down from the root. At the node at depth $\iota - 1$ of its path it queries that node's point in φ consecutive rounds, with labels drawn independently from the law of Y, and then follows the edge labeled 1 if p_1 has the strictly smaller cumulative loss over those repetitions and the edge labeled 0 otherwise. The remaining $T - \varphi D$ rounds query a fixed point with fresh labels. Every queried point and every branch is a function of the labels, which are drawn before play.

Let $f \in \mathcal{F}_1$ be the member at the leaf reached. It takes at every node of the path the value of the edge the path follows, so the isotonic comparator $p_0 + (p_1 - p_0)f$ predicts the cheaper of p_0 and p_1 at every node. Writing $Z_s = \ell(p_0, y_s) - \ell(p_1, y_s)$ for the repetitions s at one node,

$$\min\Big\{\sum_s \ell(p_0, y_s),\ \sum_s \ell(p_1, y_s)\Big\} = \frac{1}{2}\sum_s\big[\ell(p_0, y_s) + \ell(p_1, y_s)\big] - \frac{1}{2}\Big|\sum_s Z_s\Big|,$$

whose expectation is $\varphi\,\mathbb{E}\ell(p_0, Y) - \frac{1}{2}\mathbb{E}|\sum_s Z_s|$. Each prediction of the learner has conditional expected loss at least $\mathbb{E}\ell(p_0, Y)$ given the past, because the label of its round is drawn independently of the past, so the learner's expected loss over the T rounds is at least $T\,\mathbb{E}\ell(p_0, Y)$. The comparator uses one of the two predictions on every round, at expected loss $\mathbb{E}\ell(p_0, Y)$ on each of the remaining rounds, so its expected loss is $T\,\mathbb{E}\ell(p_0, Y) - \frac{D}{2}\mathbb{E}|\sum_{s\le\varphi} Z_s|$, and $\mathbb{E}R_T \ge \frac{D}{2}\mathbb{E}|\sum_{s\le\varphi} Z_s|$.

The Z_s are independent copies of Z, so $\mathbb{E}(\sum_{s\le\varphi} Z_s)^2 = \varphi\sigma_\ell^2$ and $\mathbb{E}(\sum_{s\le\varphi} Z_s)^4 = \varphi\mathbb{E}Z^4 + 3\varphi(\varphi - 1)\sigma_\ell^4 \le \varphi\bar{Z}^2\sigma_\ell^2 + 3\varphi^2\sigma_\ell^4$. By Hölder's inequality the first absolute moment is at least the $\frac{3}{2}$

power of the second moment divided by the square root of the fourth, so

$$\mathbb{E}\Big|\sum_{s\le\varphi} Z_s\Big| \ge \frac{(\varphi\sigma_\ell^2)^{3/2}}{\sqrt{3\varphi^2\sigma_\ell^4+\varphi\bar{Z}^2\sigma_\ell^2}} = \frac{\varphi\sigma_\ell^2}{\sqrt{3\varphi\sigma_\ell^2+\bar{Z}^2}} \ge \frac{1}{2}\min\Big\{\frac{\varphi\sigma_\ell^2}{\bar{Z}},\ \sigma_\ell\sqrt{\varphi}\Big\},$$

the last step by bounding $3\varphi\sigma_\ell^2+\bar{Z}^2$ by $4\varphi\sigma_\ell^2$ when $\varphi\sigma_\ell^2 \ge \bar{Z}^2$ and by $4\bar{Z}^2$ otherwise. With $\varphi \ge T/(2D)$ and $4\sqrt{2} \le 8$ this gives

$$\mathbb{E}R_T \ge \frac{D}{4}\min\Big\{\frac{\varphi\sigma_\ell^2}{\bar{Z}},\ \sigma_\ell\sqrt{\varphi}\Big\} \ge \frac{1}{8}\min\Big\{\frac{\sigma_\ell^2 T}{\bar{Z}},\ \sigma_\ell\sqrt{DT}\Big\}.$$

If $T < D$, the adversary queries one node per round and follows the cheaper prediction after each label. With the member at any leaf below the path reached, the same identity with one repetition per node gives $\mathbb{E}R_T \ge \frac{T}{2}\mathbb{E}|Z| \ge T\sigma_\ell^2/(2\bar{Z})$, because $Z^2 \le \bar{Z}|Z|$, and this exceeds the stated bound. □

Proof of Theorem 10. The upper bound is Lemma 25 at the slope bounds of pinball loss, and the lower bound is Proposition 7 at the label law $\mathrm{Ber}(1-\tau)$, which leaves every prediction Bayes-optimal.

Pinball loss is convex in its first argument with slope bounds $g_- = \tau$ and $g_+ = 1-\tau$, and $g_-g_+/(g_-+g_+) = \rho$, so Lemma 25 gives the upper bound $\min\{\rho T, 2\sqrt{2\rho T H_P(1)}\}$.

For $Y \sim \mathrm{Ber}(1-\tau)$ and every $p \in [0,1]$,

$$\mathbb{E}\ell_\tau(p,Y) = (1-\tau)\,\tau(1-p) + \tau\,(1-\tau)p = \rho,$$

so every prediction is Bayes-optimal with Bayes risk ρ. Take $p_0 = 0$ and $p_1 = 1$. Then $Z = \tau Y - (1-\tau)(1-Y) = Y-(1-\tau)$, which has $|Z| \le \max\{\tau, 1-\tau\} \le 1$ and variance ρ. Proposition 7 therefore gives

$$R_T^*(P,\ell_\tau) \ge \frac{1}{8}\min\Big\{\frac{\rho T}{\max\{\tau,1-\tau\}},\ \sqrt{\rho\ \mathrm{Ldim}(\mathcal{F}_1)\,T}\Big\} \ge \frac{1}{8}\min\big\{\rho T, \sqrt{\rho T\ \mathrm{Ldim}(\mathcal{F}_1)}\big\}.$$

Since $\mathrm{Ldim}(\mathcal{F}_1) \ge H_P(1)/\lambda'$ by Lemma 24 and $1/\sqrt{\lambda'} > 0.75$, the lower bound is at least $\frac{1}{8}\min\{\rho T, 0.75\sqrt{\rho T H_P(1)}\} \ge 0.093\min\{\rho T, \sqrt{\rho T H_P(1)}\}$, while the upper bound is at most $2\sqrt{2}\min\{\rho T, \sqrt{\rho T H_P(1)}\}$. On a chain of n points $H_P(1) = \log(n+1)$. The algorithmic clause is Theorem 11 and Corollary 5, with the polynomial-time implementation of Theorem 12. □

Proof of Corollary 3. The upper bound is Lemma 25. The lower bound is Proposition 7 together with $\mathrm{Ldim}(\mathcal{F}_1) \ge H_P(1)/\lambda'$ (Lemma 24), which turns $\sigma_\ell\sqrt{\mathrm{Ldim}(\mathcal{F}_1)T}$ into $(\sigma_\ell/\sqrt{\lambda'})\sqrt{TH_P(1)}$. Both bounds are multiples of $\min\{T, \sqrt{TH_P(1)}\}$ with factors determined by $(g_-, g_+, \sigma_\ell, \bar{Z})$, which is the stated order. The learner of the last clause is that of Corollary 4, implemented in polynomial time by Theorem 12. Pinball loss has $g_- = \tau$ and $g_+ = 1-\tau$, and the label law $\mathrm{Ber}(1-\tau)$ gives $\sigma_\ell^2 = \rho$ and $\bar{Z} = \max\{\tau, 1-\tau\}$, as the proof of Theorem 10 computes. □

The learner of this part runs on the linearized losses at resolution one, where Theorem 3 holds with $\kappa = 1$. The next lemma bounds the linearized loss of one round by the decrease of the potential divided by the rate, up to a second-order term, and it holds for every pair of slope bounds.

Lemma 26 (Corrected Binary Prediction)**.** *Let Φ and μ be those of a minimum chain partition of Q_1, let $g_-, g_+ > 0$, and let $\eta > 0$ satisfy $\eta\max\{g_-, g_+\} \le 1$. For a score array θ and a point x, let $\bar{y} = \mu_{x,1}(\theta)$ be the mean prediction on the grid $\{0,1\}$, and put*

$$\widehat{y} = \bar{y} + \eta(g_- - g_+)\,\bar{y}(1-\bar{y}).$$

Then $\widehat{y} \in [0,1]$, and for every $g \in [-g_-, g_+]$

$$g\,\widehat{y} \le \frac{\Phi(\theta) - \Phi\big(\theta +_x (0, -\eta g)\big)}{\eta} + \eta g_- g_+\,\bar{y}(1-\bar{y}).$$

Algorithm 4 The perturbed-ideal learner at resolution one with the corrected prediction

Input: a quantile level $\tau \in (0,1)$; a design horizon T; a confidence $\delta \in (0,1)$

$\rho \leftarrow \tau(1-\tau)$
if $\rho T \le 4$ **then**
  predict τ at every round, and stop
end if
compute a minimum chain partition of Q_1 by maximum bipartite matching (Fulkerson, 1956)
compute the zero-score estimate $\widetilde{\Phi}$ at resolution one with $\varepsilon = 1$ and confidence $\min\{\delta, (T+1)^{-3}\}/2$ (Remark 4), and set $\Phi_+ \leftarrow \max\{1, \widetilde{\Phi}\}$
if $\rho T \le 4\Phi_+$ **then**
  predict τ at every round, and stop
end if
$\eta \leftarrow 2\sqrt{\Phi_+/(\rho T)}$; $S \leftarrow \lceil 2T^2 \log(4T(T+1)^3/\delta)\rceil$; $\theta_1 \leftarrow 0$
for $t = 1, \dots, T$ **do**
  observe x_t
  draw S independent copies of the noise ξ, and for each compute the perturbed maximizer at θ_t by one minimum cut (Appendix D)
  $\widetilde{\mu}_{t,1} \leftarrow$ the fraction of these maximizers that contain the cell $(x_t, 1)$
  predict $\widehat{y}_t = \widetilde{\mu}_{t,1} + \eta(2\tau - 1)\,\widetilde{\mu}_{t,1}(1 - \widetilde{\mu}_{t,1})$
  observe y_t, take the selected subgradient $g_t \in \partial_p \ell_\tau(\widehat{y}_t, y_t) \cap [-\tau, 1-\tau]$, and set $\theta_{t+1} \leftarrow \theta_t +_{x_t} (0, -\eta g_t)$
end for

Proof. Bound the change of the potential by Theorem 3 and the resulting exponential moment by a quadratic. The correction makes the quadratic term the same at both ends of the slope interval. Since $|\eta(g_- - g_+)| \le \eta \max\{g_-, g_+\} \le 1$,

$$\widehat{y} = \bar{y}\big[1 + \eta(g_- - g_+)(1-\bar{y})\big] \ge 0, \qquad 1 - \widehat{y} = (1-\bar{y})\big[1 - \eta(g_- - g_+)\bar{y}\big] \ge 0.$$

Fix $g \in [-g_-, g_+]$. Theorem 3 at $K = 1$, where $\kappa = 1$, with the increment $(0, -\eta g)$ gives

$$\Phi\big(\theta +_x (0, -\eta g)\big) - \Phi(\theta) \le \log\big(1 - \bar{y} + \bar{y}e^{-\eta g}\big) = -\eta g \bar{y} + \log\big((1-\bar{y})e^{\eta g \bar{y}} + \bar{y}e^{-\eta g(1-\bar{y})}\big).$$

Both exponents in the last logarithm have absolute value at most $\eta|g| \le 1$, and $e^{\varsigma} \le 1 + \varsigma + \varsigma^2$ for $\varsigma \le 1$. The linear terms cancel, so the argument of that logarithm is at most $1 + \eta^2 g^2 \bar{y}(1-\bar{y})$, and the logarithm is at most $\eta^2 g^2 \bar{y}(1-\bar{y})$. Dividing by η and adding $g(\widehat{y} - \bar{y}) = \eta(g_- - g_+) g\,\bar{y}(1-\bar{y})$,

$$g\,\widehat{y} \le \frac{\Phi(\theta) - \Phi\big(\theta +_x (0, -\eta g)\big)}{\eta} + \eta\big(g^2 + (g_- - g_+)g\big)\bar{y}(1-\bar{y}),$$

and $g^2 + (g_- - g_+)g = (g + g_-)(g - g_+) + g_- g_+ \le g_- g_+$, because $g + g_- \ge 0 \ge g - g_+$. □

Algorithm 4 states the learner in pseudocode, and the next theorem proves its guarantee for a design horizon.

Theorem 11 (Pinball Loss at Resolution One)**.** *Let $\tau \in (0,1)$, $\rho = \tau(1-\tau)$, $T \ge 1$ and $\delta \in (0,1)$, and run Algorithm 4 with design horizon T and confidence δ. Against every non-anticipating adaptive adversary, with probability at least $1 - \delta$,*

$$R_t \le 5\min\big\{\rho T, \sqrt{\rho T H_P(1)}\big\} \qquad \text{for every } t \le T,$$

and $\mathbb{E}R_t \le 6\min\{\rho T, \sqrt{\rho T H_P(1)}\}$ for every $t \le T$. The learner computes at most one zero-score estimate, with $O(n \log((T+1)/\delta))$ minimum cuts, and $O(T^2 \log((T+1)/\delta))$ perturbed maximizers per round, each one minimum cut on a network with $n+2$ vertices and $O(n^2)$ arcs.

Proof. The constant prediction τ covers the horizons at which ρT is at most a constant multiple of $H_P(1)$. Beyond them, Lemma 26 bounds the regret on the event that the zero-score estimate and every sampled fraction are accurate, and the complementary event costs at most T times its probability. Put $H = H_P(1)$ and $r = \min\{\rho T, \sqrt{\rho T H}\}$. By Lemma 1(b), $H \ge \log(n+1) \ge \log 2$.

At every label y the function $\ell_\tau(\cdot, y)$ is $\max\{\tau, 1-\tau\}$-Lipschitz, and $\ell_\tau(\tau, y) - \ell_\tau(p, y) \le \rho$ for every $p \in [0,1]$, since $\ell_\tau(\tau, y)$ is at most its value ρ at $y \in \{0,1\}$ and $\ell_\tau \ge 0$. So the regret is at most one per round, and at most ρ per round while the learner predicts τ.

If $\rho T \le 4$, the learner predicts τ, and $R_t \le \rho T$ for every $t \le T$. The bound ρT equals r if $\rho T \le H$, and otherwise $\rho T / r = \sqrt{\rho T / H} \le \sqrt{4/\log 2} < 5$.

Let $\rho T > 4$. With probability at least $1 - \min\{\delta, (T+1)^{-3}\}/2$ the zero-score estimate satisfies $\Phi(0) \le \widetilde{\Phi} \le \Phi(0) + 2$ (Remark 4), and then Lemma 14 gives

$$H \le \Phi(0) \le \Phi_+ \le \Phi(0) + 2 \le c_\Phi H + 2 < 6H,$$

because $2 \le (2/\log 2)H$ and $c_\Phi + 2/\log 2 < 6$. If on this event the learner predicts τ, then $\rho T \le 4\Phi_+ < 24H$, so that $R_t \le \rho T \le 5r$ for every $t \le T$, as in the first case with $\sqrt{24} < 5$.

Otherwise $\rho T > 4\Phi_+$, and $\eta = 2\sqrt{\Phi_+/(\rho T)} < 1$. At round t, given the history and x_t, the S indicators whose average is $\widetilde{\mu}_{t,1}$ are independent with mean $\mu_{t,1} = \mu_{x_t,1}(\theta_t)$, so by Lemma 4 and the choice of S the probability that $|\widetilde{\mu}_{t,1} - \mu_{t,1}| \ge 1/(2T)$ is at most $2e^{-S/(2T^2)} \le \min\{\delta, (T+1)^{-3}\}/(2T)$. By a union bound over the rounds and the zero-score estimate, that estimate is accurate and every $\widetilde{\mu}_{t,1}$ is within $1/(2T)$ of $\mu_{t,1}$, except with probability at most $\min\{\delta, (T+1)^{-3}\}$.

On this event, let $\widehat{y}'_s$ be the prediction of Lemma 26 at θ_s and x_s, with $g_- = \tau$ and $g_+ = 1-\tau$. The map $\varsigma \mapsto \varsigma + \eta(2\tau - 1)\varsigma(1-\varsigma)$ has derivative in $[0,2]$ on $[0,1]$, so $|\widehat{y}_s - \widehat{y}'_s| \le 1/T$. For $f \in \mathcal{F}_P$, apply convexity and Lemma 26 with $g = g_s$ and the learner's update $\theta_{s+1} = \theta_s +_{x_s} (0, -\eta g_s)$ to obtain, at every round s,

$$\ell_\tau(\widehat{y}_s, y_s) - \ell_\tau(f(x_s), y_s) \le g_s\big(\widehat{y}'_s - f(x_s)\big) + \frac{1}{T} \le \frac{\Phi(\theta_s) - \Phi(\theta_{s+1})}{\eta} + \frac{\eta\rho}{4} - g_s f(x_s) + \frac{1}{T}.$$

Sum over $s \le t$. The potential telescopes, and $\Phi(\theta_{t+1}) \ge \theta_{t+1}(I) = -\eta \sum_{s \le t} g_s F_I(x_s)$ for every ideal I of Q_1 (Lemma 14). As in the proof of Lemma 25, some $f^\circ \in \mathcal{F}_1$ has $\sum_{s \le t} g_s f^\circ(x_s) \le \sum_{s \le t} g_s f(x_s)$, and $f^\circ = F_I$ for an ideal I (Lemma 6(a)). Hence

$$R_t \le \frac{\Phi(0)}{\eta} + \frac{\eta\rho t}{4} + \frac{t}{T} \le \frac{\Phi_+}{\eta} + \frac{\eta\rho T}{4} + 1 = \sqrt{\Phi_+ \rho T} + 1.$$

Here $\rho T > 4\Phi_+ \ge 4H$, so $r = \sqrt{\rho T H} > 2H$ and $1 < r/(2H) \le r/(2\log 2)$, while $\sqrt{\Phi_+ \rho T} < \sqrt{6}\, r$. The bound is at most $(\sqrt{6} + 1/(2\log 2))\, r < 5r$.

For the expectation, $R_t \le 5r$ holds on every realization if $\rho T \le 4$. If $\rho T > 4$, then $r > \min\{4, \sqrt{4\log 2}\} > 1$. The event above fails with probability at most $(T+1)^{-3}$, and $R_t \le t \le T$ on every realization, so $\mathbb{E} R_t \le 5r + T(T+1)^{-3} \le 5r + 1/8 < 6r$.

The adversary is non-anticipating, so the sampling bounds hold conditionally on the history, and everything else holds on every realization. The zero-score estimate takes $2\lceil \log_4(2/\min\{\delta, (T+1)^{-3}\})\rceil \lceil 96n \rceil = O(n \log((T+1)/\delta))$ minimum cuts, and each round takes $S = O(T^2 \log((T+1)/\delta))$ minimum cuts. Every cut is on the network of Appendix D for Q_1, which has $n+2$ vertices and $O(n^2)$ arcs. □

Corollary 4 (Lipschitz Losses at Resolution One). *Let ℓ be finite and convex in its first argument, with slope bounds $g_-, g_+ > 0$, and let $T \ge 1$ and $\delta \in (0,1)$. Run Algorithm 4 with $\tau = g_-/(g_- + g_+)$, design horizon T and confidence δ on the loss $\ell/(g_- + g_+)$, using its selected subgradients in place of those of pinball loss. Then, with probability at least $1 - \delta$,*

$$R_t \le 5 \min\left\{ \frac{g_- g_+ T}{g_- + g_+},\ \sqrt{g_- g_+ T H_P(1)} \right\} \qquad \textit{for every } t \le T,$$

and $\mathbb{E} R_t \le 6 \min\{g_- g_+ T/(g_- + g_+), \sqrt{g_- g_+ T H_P(1)}\}$ for every $t \le T$.

Proof. With $\tau = g_-/(g_- + g_+)$ and $\rho = \tau(1-\tau) = g_- g_+/(g_- + g_+)^2$, the regret of ℓ is $g_- + g_+$ times that of $\ell/(g_- + g_+)$, and $g_- + g_+$ times the bounds of Theorem 11 are the stated ones. That proof uses pinball loss only through three properties, which $\ell/(g_- + g_+)$ shares. Its selected

subgradients lie in $[-\tau, 1-\tau]$, as Lemma 26 requires. Its value at τ exceeds its value at any $p \in [0,1]$ by at most ρ. With g the selected subgradient at τ, convexity bounds the difference by $g(\tau - p)$, which is at most $(1-\tau)\tau$ if $g \geq 0$ and at most $\tau(1-\tau)$ if $g < 0$. It is also $\max\{\tau, 1-\tau\}$-Lipschitz. For $p < p'$ in $[0,1]$, the subgradients in $[-\tau, 1-\tau]$ at p and at p' bound its increment from p to p' below by $-\tau(p'-p)$ and above by $(1-\tau)(p'-p)$. □

Corollary 5 (Horizon-Free Learning under Pinball Loss)**.** *Let $\tau \in (0,1)$, $\rho = \tau(1-\tau)$ and $\delta \in (0,1)$. Run Algorithm 4 afresh on the epochs $\iota = 1, 2, \ldots$ of lengths $T_\iota = 2^{\iota-1}$, epoch ι with design horizon T_ι and confidence $\delta 2^{-\iota}$. Against every non-anticipating adaptive adversary,*

$$\mathbb{P}\big\{R_T \leq 18 \min\big\{\rho T, \sqrt{\rho T H_P(1)}\big\} \text{ for all } T \geq 1\big\} \geq 1 - \delta,$$

and $\mathbb{E}R_T \leq 21 \min\{\rho T, \sqrt{\rho T H_P(1)}\}$ for every $T \geq 1$.

Proof. Split the regret over the epochs and add the bounds of Theorem 11, and both sums are geometric. Fix $T \geq 1$, put $H = H_P(1)$, and let $\iota_T = 1 + \lfloor \log_2 T \rfloor$ be the epoch that contains round T. For every $f \in \mathcal{F}_P$, the regret against f over the first T rounds is the sum, over the epochs $\iota \leq \iota_T$, of the regrets against f over their rounds up to round T. Each summand is at most the regret of that epoch's learner over a prefix of its run. With probability at least $1 - \sum_\iota \delta 2^{-\iota} \geq 1 - \delta$ every epoch satisfies the probability bound of Theorem 11, and then, for every T,

$$\begin{aligned} R_T &\leq 5 \sum_{\iota \leq \iota_T} \min\big\{\rho T_\iota, \sqrt{\rho H T_\iota}\big\} \leq 5 \min\Big\{\rho \sum_{\iota \leq \iota_T} T_\iota,\ \sqrt{\rho H} \sum_{\iota \leq \iota_T} \sqrt{T_\iota}\Big\} \\ &\leq 5 \min\big\{2\rho T, (2+\sqrt{2})\sqrt{\rho H T}\big\}, \end{aligned}$$

because $T_{\iota_T} \leq T$ gives $\sum_{\iota \leq \iota_T} T_\iota = 2T_{\iota_T} - 1 < 2T$ and $\sum_{\iota \leq \iota_T} \sqrt{T_\iota} < \sqrt{2T_{\iota_T}}/(\sqrt{2}-1) \leq (2+\sqrt{2})\sqrt{T}$. Since $5(2+\sqrt{2}) < 18$, this is the first claim. For the second, condition on the history before an epoch. The adversary's choices within the epoch form a non-anticipating adaptive adversary for that epoch's learner, so the expectation bound of Theorem 11 applies to it. The same sums with 6 in place of 5 give $\mathbb{E}R_T \leq 6(2+\sqrt{2}) \min\{\rho T, \sqrt{\rho H T}\} < 21 \min\{\rho T, \sqrt{\rho H T}\}$. □

I Finite Precision

This part specifies the finite-precision model of Section 2 and proves the polynomial-time claims of the paper in it, with the regret bounds of the exact learners up to constant factors (Theorem 12). In this model the order relation is given as a list of pairs. The labels and, for pinball loss, the quantile level are rational. The learner draws its random samples as rational numbers of polynomial bit length, and an oracle evaluates the loss, and for Lipschitz losses a subgradient, at rational arguments where the loss is finite, to any requested precision in polynomial time. The exact learners use real numbers in three places: the Gumbel noise and the scores, which determine the law of each perturbed maximizer; the weights of the masters and of the scalar baseline; and the subgradients of a Lipschitz loss. An implementation replaces them by rational approximations whose errors are summable over the rounds of each run. A real number is computed to accuracy ε if a rational number within ε of it is computed. The logarithm and the exponential function are computed to accuracy ε in time polynomial in $\log(1/\varepsilon)$ and in the bit length and the absolute value of the rational argument (Brent & Zimmermann, 2010). A minimum cut of a network with rational capacities is computed exactly in polynomial time. Multiplied by a common denominator, the capacities become integers, and the maximum-flow algorithm of Edmonds & Karp (1972) performs polynomially many arithmetic operations on integers bounded by the sum of the capacities.

Rational Gumbel samples. For $\varepsilon \in (0, 1/4]$, a rational Gumbel sample at accuracy ε is computed from $\lceil \log_2(4/\varepsilon^3) \rceil$ random bits. Read them as a dyadic number in $[0,1)$, clip it to $[\varepsilon, 1-\varepsilon]$, and compute $-\log(-\log(\cdot))$ at the clipped number to accuracy $3\varepsilon/4$. On $[\varepsilon, 1-\varepsilon]$ the function $-\log(-\log(\cdot))$ is increasing and has absolute value at most $\log(1/\varepsilon)$, because $\log(1/(1-\varepsilon)) \geq \varepsilon$ and $\log\log(1/\varepsilon) \leq \log(1/\varepsilon)$, so the sample has bit length $O(\log(1/\varepsilon))$. To compare it with the noise $\xi_{i,k}$ at a coordinate, let the bits be the leading binary digits of $\varpi = \exp(-e^{-\xi_{i,k}-\gamma})$, which is uniform on $(0,1)$, because $e^{-\xi_{i,k}-\gamma}$ is a standard exponential variable, and whose binary digits

are therefore independent fair bits. Then $\xi_{i,k} + \gamma = -\log(-\log \varpi)$. The sample $\tilde{\xi}_{i,k}$ so obtained satisfies

$$\mathbb{P}\big(|\tilde{\xi}_{i,k} - \xi_{i,k} - \gamma| > \varepsilon\big) \le 2\varepsilon, \qquad \mathbb{E}|\tilde{\xi}_{i,k} - \xi_{i,k} - \gamma| \le 4\varepsilon, \qquad |\tilde{\xi}_{i,k}| \le |\xi_{i,k} + \gamma| + \varepsilon. \tag{9}$$

To prove these bounds, let ϖ_ε be ϖ clipped to $[\varepsilon, 1-\varepsilon]$. On that interval the derivative of $-\log(-\log \varpi)$ is $1/(\varpi \log(1/\varpi)) \le \varepsilon^{-2}$, because $\varpi \log(1/\varpi)$ is concave and at least ε^2 at both endpoints. The bits determine ϖ to within $\varepsilon^3/4$ and clipping is 1-Lipschitz, so $\tilde{\xi}_{i,k}$ is within $\varepsilon/4 + 3\varepsilon/4 = \varepsilon$ of $-\log(-\log \varpi_\varepsilon)$. This gives the first bound, since $\varpi_\varepsilon = \varpi$ except with probability 2ε, and the third, since $-\log(-\log \varepsilon) < 0 < -\log(-\log(1-\varepsilon))$, so that clipping moves $\xi_{i,k}+\gamma$ toward 0. For the second, clipping from above lowers $\xi_{i,k}+\gamma$ by its excess over $-\log(-\log(1-\varepsilon))$, and the expected excess is at most $\log(1/(1-\varepsilon)) \le 2\varepsilon$ because $\mathbb{P}(\xi_{i,k} + \gamma > \varsigma) \le e^{-\varsigma}$. Clipping from below raises it by its shortfall under $-\log\log(1/\varepsilon)$, and the expected shortfall is

$$\int_{-\infty}^{-\log\log(1/\varepsilon)} \exp(-e^{-\varsigma})\, d\varsigma \le \frac{1}{\log(1/\varepsilon)} \int_{-\infty}^{-\log\log(1/\varepsilon)} e^{-\varsigma} \exp(-e^{-\varsigma})\, d\varsigma = \frac{\varepsilon}{\log(1/\varepsilon)} \le \varepsilon;$$

so $-\log(-\log \varpi_\varepsilon)$ differs from $\xi_{i,k} + \gamma$ by at most 3ε in expectation, and the sample by at most 4ε.

The total variation distance between two laws is the largest difference between the probabilities they give to an event. It is at most the probability that two random variables with these laws, defined on one probability space, differ, and the expectations of a function under the two laws differ by at most this distance times the length of an interval that contains the values of the function. When every perturbed score is approximated accurately, the approximation changes the perturbed maximizer only when the two best perturbed scores nearly tie, and independent noise makes near ties rare. The next lemma is stated for Q_K. It holds for the perturbed potential of every nonempty finite poset, in particular of the posets $Q_\mathcal{G}$ of Proposition 6(b) when $\mathcal{G}$ has at least two members.

Lemma 27 (Isolation under Independent Perturbations)**.** *Let θ be a score array and $\varepsilon \in (0, 1]$.*

- *(a) With probability at most $|\Xi|\varepsilon$, the largest and the second-largest of the perturbed scores $\theta(I) + \xi(I)$ over the ideals of Q_K differ by at most ε.*
- *(b) Let $\tilde{\theta}$ be a rational score array with $|\tilde{\theta}(I) - \theta(I)| \le \varepsilon/(8|\Xi|)$ for every ideal I. Let the $\tilde{\xi}_{i,k}$ be independent rational Gumbel samples at accuracy $\varepsilon/(8|\Xi|^2)$, one for each noise coordinate, and let $\tilde{I}$ be an ideal that maximizes $\tilde{\theta}(I) + \sum_i \tilde{\xi}_{i,k_i(I)}$, computed by one minimum cut with rational node weights. Then the law of $\tilde{I}$ is within total variation distance ε of the law of the perturbed maximizer at θ.*

Proof. (a) Two ideals with nearly equal perturbed scores differ in the prefix length of some chain, and the noise coordinate of that prefix length separates them. Fix a noise coordinate (i, k) and condition on all the others. The largest perturbed score over the ideals with $k_i(I) = k$ is $\xi_{i,k}$ plus a quantity that does not depend on $\xi_{i,k}$, and the largest over the other ideals does not depend on $\xi_{i,k}$. If both sets of ideals are nonempty, the two maxima are within ε of each other with conditional probability at most $2e^{-1}\varepsilon$, because the density of a centered standard Gumbel variable is at most e^{-1}. Now let I_1 and I_2 attain the largest and the second-largest perturbed score. Their chain encodings differ, say $k_i(I_1) = k \ne k_i(I_2)$. Then I_1 attains the largest perturbed score over the ideals with $k_i(I) = k$, and I_2 the largest over the others, because every ideal other than I_1 scores at most as much as I_2. So a near tie implies one of $|\Xi|$ events, each of probability at most $2e^{-1}\varepsilon \le \varepsilon$.

(b) Couple the samples with the noise as in the construction of (9). By (9) and $|\Xi| \ge 2$, outside an event of probability at most $|\Xi| \cdot 2\varepsilon/(8|\Xi|^2) \le \varepsilon/8$ every sample is within $\varepsilon/(8|\Xi|^2)$ of $\xi_{i,k} + \gamma$. Off that event, each computed perturbed score, which involves one noise coordinate in each of the $w_K \le |\Xi|$ chains, is within $w_K\varepsilon/(8|\Xi|^2) + \varepsilon/(8|\Xi|) \le \varepsilon/(4|\Xi|)$ of $\theta(I) + \xi(I) + w_K\gamma$, and the shift $w_K\gamma$ is common to all ideals. So $\tilde{I}$ is the perturbed maximizer at θ unless the two best perturbed scores are within $\varepsilon/(2|\Xi|)$ of each other, which by (a) has probability at most $\varepsilon/2$, and the two laws differ by at most $\varepsilon/8 + \varepsilon/2 < \varepsilon$ in total variation. □

Pair mode tolerates such errors in the laws of its sampled levels.

Lemma 28 (Numerical Stability of Pair Sampling)**.** *In the setting of Proposition 5(b), let θ_t be the scores that the update of Proposition 5(b) produces from the observed rounds, and suppose that at each round t the S sampled levels are independent given the history, x_t and y_t, each with a conditional law within total variation distance ε_t of μ_t, for a fixed sequence $(\varepsilon_t)_{t\ge1}$. Then both bounds of Proposition 5(b) hold with $5\sum_{t\le T}\varepsilon_t$ added to their right sides.*

Proof. In the proof of Proposition 5(b) the law of the samples enters only through the expectation of a bounded function of one pair. As in that proof, assume $\eta \le 1$, and let ϕ be the bound defined there on the loss of the average of a pair. The exponent $\frac{\eta}{32}\phi - \frac{\kappa}{2}(\Phi(\theta_t) - \Phi(\theta_{t+1}))$ lies in $[-\eta/32, \eta/32]$, because $\phi \in [0,1]$ and the update lowers the scores at x_t by at most $\eta/(16\kappa)$, which lowers Φ by at most as much. The law of a pair is within total variation distance $2\varepsilon_t$ of the exact one, under which the conditional expectation of the exponential of that exponent is at most one. Under the actual law this conditional expectation is therefore at most

$$1 + 2\varepsilon_t\big(e^{\eta/32} - e^{-\eta/32}\big) \le 1 + \frac{5\varepsilon_t\eta}{32} \le \exp\Big(\frac{5\varepsilon_t\eta}{32}\Big).$$

Multiplying over the $S/2$ independent pairs shows that the process of that proof, with $\ell(\widehat{y}_t, y_t) - 5\varepsilon_t$ in place of $\ell(\widehat{y}_t, y_t)$ at every round, is again a nonnegative supermartingale, and Ville's inequality and Jensen's inequality give the two bounds with the additional term. □

With these two lemmas, the implementations of the learners keep their guarantees.

Theorem 12 (Finite-Precision Implementations)**.** *Fix a loss with its regularity bounds. In the finite-precision model, the horizon-free learner of Theorem 4 (Algorithm 3, which runs Algorithms 1 and 2), the learners of Theorem 11 and Corollaries 4 and 5 (Algorithm 4) and the learner of Theorem 9(e) have implementations whose computation at round t is polynomial in n, t, the design horizon when there is one, the input bit lengths and $\log(1/\delta)$. These implementations satisfy the regret bounds of those results with larger constants, which depend on the same quantities as the original ones. The zero-score estimate of Remark 4, computed with rational Gumbel samples at accuracy $\varepsilon/(32|\Xi|)$ and with twice as many samples in each batch, satisfies the guarantee of Remark 4, with computation polynomial in n, K, $1/\varepsilon$ and $\log(1/\delta)$.*

Proof. An implementation stores the observations and adds errors, summable over the rounds of each run, in the laws of the sampled levels, in the predictions and, for Corollary 4, in the subgradients. Its analysis follows that of the exact learner along the realized sequence, with the exact scores of those observations. The combinatorial steps are exact. Maximum bipartite matching computes the minimum chain partitions, every minimum cut is computed from rational node weights, and, for Theorem 9(e), the digraphs $\mathcal{D}_{b,K}$, their reachability preorders and the quotient posets are computed from the rational bounds b. The regularity bounds may be taken rational. Replacing η by a rational number in $[\eta/2, \eta]$, and β, Λ and the slope bounds by rational numbers between them and their doubles, keeps every hypothesis, changes the constants by bounded factors, and makes the update rates of Algorithms 1 to 3, the normalization by Λ and the quantile level of Corollary 4 rational. The rate of Algorithm 4 is treated below. The adversary is non-anticipating, so the bounds on sampling errors hold conditionally on the history, as the exact analyses require.

At the t-th round of its run, a copy of Algorithm 1, or of the learner of Proposition 6(b), recomputes its scores from the stored observations, evaluating each of the at most t losses that enter the score of an ideal to accuracy $\varepsilon/(8|\Xi|t)$, for the tolerance ε chosen below. Its update rate is at most one, so every ideal's score is within $\varepsilon/(8|\Xi|)$ of the exact one, and by Lemma 27(b) every sampled ideal has a conditional law within total variation distance ε of that of the perturbed maximizer, and the law of its level at x_t is therefore within ε of μ_t. In pair mode take $\varepsilon = 1/(5(t+1)^2)$. Lemma 28 adds $\sum_t (t+1)^{-2} < 1$ to the bounds of Proposition 5(b), which the proofs of Theorem 7 and Theorem 9(e) absorb into their constants. In mean mode take $\varepsilon = \varepsilon_{K,t}/2$. The conditional mean of the average $\widehat{y}_t$ of the $S_{K,t}$ sampled grid values is then within $\varepsilon_{K,t}/2$ of $\bar{y}_t$, because the grid values lie in $[0,1]$. By Lemma 4, $\widehat{y}_t$ deviates from that mean by more than $\varepsilon_{K,t}/2$ with probability at most $2\exp(-S_{K,t}\varepsilon_{K,t}^2/2) \le \delta_{K,t}$, and the conditional mean of $|\widehat{y}_t - \bar{y}_t|$ is at most $1/(2\sqrt{S_{K,t}}) + \varepsilon_{K,t}/2 < \varepsilon_{K,t}$. The proof of Theorem 8 needs only these two bounds, so it applies unchanged.

The zero-score estimate uses the rational zero-score samples, the maximum weights of an ideal for the node weights $\tilde{\xi}_{i,k} - \tilde{\xi}_{i,k-1}$, in which the constant γ cancels. By the argument of Lemma 15,

changing one noise coordinate moves such a sample by at most the change of that coordinate. By (9), the expectation of a rational zero-score sample is therefore within $\sum_{(i,k)\in\Xi} \mathbb{E}|\tilde{\xi}_{i,k} - \xi_{i,k} - \gamma| \le \varepsilon/8$ of $\mathbb{E}\widehat{\Phi} = \Phi(0)$, and the Efron–Stein inequality bounds its variance by

$$\sum_{(i,k)\in\Xi} \mathbb{E}\tilde{\xi}_{i,k}^2 < 4|\Xi| \le 8nK,$$

since $\mathbb{E}\tilde{\xi}_{i,k}^2 \le 2\mathbb{E}(\xi_{i,k}+\gamma)^2 + 2(\varepsilon/(32|\Xi|))^2$ and $\mathbb{E}(\xi_{i,k}+\gamma)^2 = \gamma^2 + \pi^2/6 < 1.98$. By Chebyshev's inequality, a batch mean of $2\lceil 96nK/\varepsilon^2 \rceil$ such samples deviates from that expectation by more than $7\varepsilon/8$ with probability at most $(\varepsilon^2/24)/(7\varepsilon/8)^2 < 1/16$. The median argument of the proof of Remark 4 then puts the median within $7\varepsilon/8 + \varepsilon/8 = \varepsilon$ of $\Phi(0)$ with probability at least $1-\delta$, and adding ε gives $\Phi(0) \le \widetilde{\Phi} \le \Phi(0) + 2\varepsilon$.

The scalar baseline and the masters of Algorithm 3 and of Theorem 9(e) predict weighted means of predictions in $[0,1]$. A master weights an expert E in proportion to $\exp(\log \nu_E - \eta\,\mathrm{Loss}_E(t-1))$, with a rational prior ν_E. The scalar baseline keeps its epoch counts exactly and weights each level in proportion to the exponential of $-\eta$ times the loss of that level over the current epoch at the queried point. For $\varepsilon \in (0,1]$, compute each of these exponents to accuracy $\varepsilon/8$. Every normalized weight then changes by a factor in $[e^{-\varepsilon/4}, e^{\varepsilon/4}]$, hence the mean by at most $e^{\varepsilon/4} - 1 < \varepsilon/3$. Subtract the largest computed exponent, compute each exponential to accuracy $\varepsilon/16$ divided by the number of weights, clip it to $[0,1]$, and round the normalized mean to accuracy $\varepsilon/4$. Since the exponentials sum to at least one, a total change of at most $\varepsilon/16$ in the weights moves the normalized mean by at most $(\varepsilon/8)/(1-\varepsilon/16) < \varepsilon/7$, so these steps add less than $\varepsilon/2$. Projected onto the convex hull of the averaged predictions, which does not increase the error, the result is within ε of the exact weighted mean, whose cumulative loss obeys Lemma 3. Take $\varepsilon = 1/(\Lambda_1(t+1)^2)$ at round t, where Λ_1 is a Lipschitz constant of the loss on that hull, namely $1 + \beta/\Lambda$ for a β-smooth loss with values in $[0,\Lambda]$ run as ℓ/Λ, by the mean value theorem, and $1 + (2t+1)^2$ for log loss, whose predictions at round t lie between the extreme levels of rational grids of resolution at most t (Lemma 11). The cumulative loss then exceeds that of the exact weighted means by less than one, which the proofs of Lemma 20, Theorem 4 and Theorem 9(e) absorb into their constants.

With η as in Algorithm 4, the implementation of that algorithm keeps exact rational scores, uses the power of two $\hat{\eta} \in (\eta/2, \eta]$ in place of η, found by comparing powers of four with $4\Phi_+/(\rho T)$, and updates by rational subgradients. Since $\Phi_+/\hat{\eta} + \hat{\eta}\rho T/4 \le 2\Phi_+/\eta + \eta\rho T/4$, the bound below follows from the proof of Theorem 11 with $\hat{\eta}$ in place of η. The selected subgradients of pinball loss are exact. For Corollary 4, the oracle computes the selected subgradient of $\ell/(g_- + g_+)$ to accuracy $1/T$, and the implementation projects the result onto $[-\tau, 1-\tau]$. Lemma 26 holds for every value in that interval, and convexity charges the error at most $1/T$ per round, since predictions and comparator values lie in $[0,1]$. Lemma 27(b) with $\varepsilon = 1/(8T)$, together with four times as many samples, puts $\widetilde{\mu}_{t,1}$ within $1/(4T) + 1/(8T) < 1/(2T)$ of $\mu_{t,1}$ with the failure probability of the proof of Theorem 11, and the zero-score estimate above, at $K = 1$ and $\varepsilon = 1$, has the guarantee that proof uses. On the event that the zero-score estimate and every sampled fraction are accurate, the proof of Theorem 11 bounds the regret of the corrected predictions by

$$\frac{2\Phi_+}{\eta} + \frac{\eta\rho T}{4} + 2 = \frac{3}{2}\sqrt{\Phi_+\rho T} + 2,$$

and the remaining steps of that proof and of the proofs of Corollaries 4 and 5 hold with larger constants.

Every number that an implementation stores has polynomial bit length. A rational Gumbel sample at accuracy ε has $O(\log(1/\varepsilon))$ bits, and every tolerance above is the inverse of a polynomial in n, t, the resolution in use, the design horizon when there is one and, for the zero-score estimate, the inverse of its requested accuracy. Scores and cumulative losses are sums of at most t loss values computed to such tolerances, and at round t log loss is at most $\log(1 + (2t+1)^2)$ on the predictions in use. A sample count involving a logarithm uses an upper approximation of that logarithm to accuracy one, which preserves the count's order of magnitude. The numbers of samples, weights and cuts per round are those of the exact learners, with a factor of four for Algorithm 4. Every resolution in use at round t is at most t, or at most $\lceil\sqrt{T}\rceil$ in Theorem 9(e), and the zero-score estimate draws $4\lceil\log_4(1/\delta)\rceil\lceil 96nK/\varepsilon^2\rceil$ samples, so the computation is polynomial in the stated parameters. □