

MULTISCALE REWARD HEDGING FROM CORRECT DEMONSTRATIONS

Pahan Dewasurendra
Johns Hopkins University

ABSTRACT

Learning from correct demonstrations is harder than supervised learning when many answers are correct: after predicting, the learner sees one valid answer but not whether its own answer was valid, nor any reward. Existing reward-hedging guarantees consequently assume a finite reward class. We give the first horizon-free guarantee for continuous classes. The key is to hedge in one shared vote over tolerant optimality tests at every accuracy scale. A target reward has one surviving proxy per scale, and a prediction with gap above that scale doubles the proxy. This yields the simultaneous tail bound

$$|\{t : \ell_t > 2^{-j}\}| \leq \log_2 \mathcal{N}(\mathcal{G}, 2^{-j-1}) + j,$$

where $\mathcal{G}$ is the class of optimality-gap functions. Integrating the tails gives cumulative hidden gap bounded by a metric-entropy integral, independently of the number of rounds. Polynomial entropy $(A/\epsilon)^d$ gives $O(d \log A)$ total gap and a fast $O(d/m)$ statistical rate. For bounded linear contextual recommendation, the result is $O(d)$ regret for arbitrary compact menus. This is the first polynomial finite bound without structural restrictions on the menus, at the price of improper prediction. Although the general vote can be expensive, it is exactly polynomial-time for one-dimensional Lipschitz parameter curves. Fixed-radius rank-two recommendation takes $O(KT^2)$ time for menus of size K . We also prove an $\Omega(d)$ lower bound, low-rank and bounded ReLU-network corollaries, and a robust theorem that adds only the demonstrator’s cumulative suboptimality. A reproducible adaptive stress test illustrates the predicted scale adaptation. After factorization, an exact MovieLens audit runs in 1.7 CPU seconds across ten users and improves mean latent gap over both a demonstrated-rating policy and a proper online baseline. The learner uses only action demonstrations and never observes a reward or a loss.

1 INTRODUCTION

A question can have many correct answers. A recommendation menu can have several optimal items. An expert demonstration reveals one of them, but it does not reveal the reward function, the full set of good answers, or whether a learner’s different answer was also good. This is the feedback pattern of supervised fine-tuning when utility, rather than cloning one demonstrator’s distribution, is the objective. It is also the one-step form of apprenticeship learning (Abbeel & Ng, 2004; Syed & Schapire, 2007).

Joshi et al. (2026) formalized this problem as learning to answer from correct demonstrations. For a finite reward class $\mathcal{R}$, their reward-hedging rule achieves $O(\log |\mathcal{R}|)$ cumulative gap against an optimal demonstrator and a corresponding $O(\log |\mathcal{R}|/m)$ statistical rate. Their main open question asks for continuous reward classes, naming generalized linear rewards, low-rank bilinear scores, and neural rewards as motivating examples. Applying a single fine discretization does not answer the question. An ϵ -cover gives the familiar tradeoff $T\epsilon + \log \mathcal{N}(\epsilon)$, which is $\Theta(d \log T)$ for a d -parameter class.

The same barrier appears independently in contextual recommendation. Proper algorithms progressed from $O(\sqrt{T})$ regret (Bärman et al., 2017) to polynomial-time $O(d^4 \log T)$ regret (Besbes

et al., 2025). More recent work sharpens the dimension dependence. Here a learner sees a menu $X_t \subset \mathbb{R}^d$, recommends $\hat{x}_t$, and then observes a user’s optimal choice x_t^* for an unknown linear utility. The best general bounds have been $O(d \log T)$ or $\exp(O(d \log d))$ (Gollapudi et al., 2021; Sakaue et al., 2025). A 2026 result obtains finite $O(d \log d)$ regret for M-convex menus and identifies a general polynomial finite bound as open (Oki & Sakaue, 2026). We resolve the menu-aware, improper version of this question with $O(d)$ regret for arbitrary bounded menus.

Our algorithm starts from a simple observation. At accuracy ϵ , every reward-gap model g defines a binary proxy that calls an action good when $g(x, a) \leq \epsilon$. A cover element within ϵ of the target always calls an optimal demonstration good. If it also calls the learner’s action good, that action has target gap at most 2ϵ . Thus every larger-gap round can multiply the proxy’s weight. Running this argument separately at each scale would produce incompatible actions. Instead, we put all proxies at all scales into one weighted vote, with scale prior 2^{-j} . One potential controls the whole hierarchy. The resulting mistake bound holds at every threshold simultaneously, and layer-cake integration converts it into a finite entropy sum.

Our contributions are:

- We introduce a deterministic multiscale reward-hedging algorithm for continuous reward classes. For the decision-relevant gap metric, its total unobserved loss is at most

$$\sum_{j \geq 1} 2^{-j} \log_2 \mathcal{N}(\mathcal{G}, 2^{-j-1}) + 2. \quad (1)$$

The guarantee is adversarial, horizon-free, and invariant to context-dependent reward offsets. A finite version uses only $\lceil \log_2 T \rceil$ scales and loses at most one more.

- If the normalized gap class has covering numbers at most $(A/\epsilon)^d$, the bound is $d(\log_2 A + 3) + 2$. Bounded linear contextual recommendation therefore has $O(BKd)$ regret when parameters have radius B and each menu has diameter K . Low-rank $m \times n$ bilinear rewards have $O(BKr(m+n))$ regret. For a p -parameter depth- L bounded ReLU reward network, we obtain an explicit $O(Gp[1 + \log(1 + BRS^{L-1}\sqrt{L}/G)])$ bound.
- The finite algorithm is exact and polynomial-time when the gap functions follow a one-dimensional Lipschitz curve. It uses $O(T)$ proxies, $O(KT^2)$ total time, and $O(T)$ memory for fixed geometric constants. This includes fixed-radius rank-two recommendation with arbitrary finite menus.
- We prove a matching $\Omega(d)$ lower bound for a continuous class with entropy exponent d . Under i.i.d. contexts, online-to-batch conversion gives population gap $O((d + \log(1/\delta))/m)$ with high probability.
- A soft multiscale update handles suboptimal demonstrations. If their cumulative target gap is Δ_T , the learner incurs $O(\Delta_T + d)$ for polynomial-entropy classes. No reward or correctness bit is observed in either algorithm. A finite stress test verifies the scale-adaptive behavior. A second deterministic audit uses actual MovieLens ratings as suboptimal demonstrations for ten rank-two preference circles.

The general-dimensional result is information-theoretic. Cover enumeration can be expensive and the action predictor is improper. Section 5.1 isolates a nontrivial case where the same exact algorithm is polynomial-time. Proper objective estimation remains distinct (Sakaue et al., 2025; Sakaue, 2026).

2 PROBLEM AND CLOSEST WORK

Let $\mathcal{X}$ be a context space and let $\mathcal{A}(x)$ be the nonempty action menu at x . A reward $r \in \mathcal{R}$ assigns values to context-action pairs. Its decision-relevant object is the *optimality-gap function*

$$g_r(x, a) = \sup_{b \in \mathcal{A}(x)} r(x, b) - r(x, a). \quad (2)$$

We first normalize so $g_r \in [0, 1]$, and restore scale in applications. Write $\mathcal{G} = \{g_r : r \in \mathcal{R}\}$ and equip it with

$$d_{\text{gap}}(r, s) = \|g_r - g_s\|_{\infty}.$$

This pseudometric ignores arbitrary context-dependent reward offsets. It is also no larger than $2\|r - s\|_\infty$, so ordinary reward covers remain valid.

The gap metric can be much sharper than the reward sup metric. For example, if $r_\theta(x, a) = u_\theta(x) + v(x, a)$ with an arbitrarily rich class of context offsets u_θ , then every member has the same gap function. Its decision-relevant covering number is one even when the raw reward class is not totally bounded. This quotient geometry is necessary for utility claims, since demonstrations cannot identify an offset that changes no decision.

On round t , the learner receives x_t , predicts $\hat{a}_t \in \mathcal{A}(x_t)$, then observes one demonstration a_t^* . There is a fixed unknown $r_* \in \mathcal{R}$. In the realizable case, $g_{r_*}(x_t, a_t^*) = 0$, but the learner does not observe this value or its own loss

$$\ell_t = g_{r_*}(x_t, \hat{a}_t). \quad (3)$$

Contexts and the choice among optimal demonstrations may be adaptive. Let $\mathcal{N}(\mathcal{G}, \epsilon)$ denote the minimum cardinality of a proper ϵ -cover under the sup metric, and assume these covering numbers are finite.

Closest feedback models. The finite-class rule of [Joshi et al. \(2026\)](#) is the direct starting point. A countable-class extension assigns a summable prior and obtains a target-specific $\log(1/w_0)$ bound ([Kleinberg et al., 2026](#)). It cannot assign positive mass to every member of an uncountable class. [Shao et al. \(2026\)](#) characterize 0–1 learning when only one member of a hidden acceptable set is shown, using partial-feedback Littlestone dimensions over collections of policy hypotheses. Their result does not control real-valued reward gap or connect reward-parameter entropy to regret. In the same 0–1 setting, [Pour et al. \(2026\)](#) give exact tree-dimension characterizations for mistake-unknown, mistake-known, and set-valued feedback, including infinite multi-label classes. Their tree dimensions do not grade reward gap or connect parameter entropy. [Raman et al. \(2024\)](#) reveal the entire acceptable set. [Cohen et al. \(2026\)](#) instead predict sets and receive sampled precision or recall membership bits, rather than one demonstrated action with hidden correctness.

Closest multiscale result. [Attias et al. \(2023\)](#) characterize minimax realizable online regression by a scaled-Littlestone dimension. Building on that characterization, [Doron-Arad et al. \(2026\)](#) bound the dimension by a Dudley-type entropy potential. Both protocols reveal the exact numeric label after every prediction. In our protocol, the numeric target gap and even the correctness of the prediction stay hidden. Our entropy shape is related, but the tolerant-proxy reduction and its one-vote potential are specific to action-only feedback.

Online inverse optimization. [Bärmann et al. \(2017\)](#) give a proper exponential-weights learner with $O(\sqrt{T})$ regret over arbitrary optimization-oracle menus. [Besbes et al. \(2025\)](#) introduce inverse exploration and a polynomial-time ellipsoidal-cone method with $O(d^4 \log T)$ general regret. The line culminating in $O(d \log T)$ and structured finite bounds is summarized in Table 1. Our general $O(d)$ result is improper and expensive, while Section 5.1 gives an exact polynomial-time slice when the intrinsic parameter dimension is one.

Two concurrent preprints clarify complementary frontiers. [Jimenez et al. \(2026\)](#) give polynomial-time mistake-bounded generators for conjunctions, parities, and monotone Boolean functions with polynomially many maxterms. Their discrete language-generation algorithms do not address real-valued gap entropy. Their positive results underscore the importance of finding efficient special cases of our vote. [Fatemi et al. \(2026\)](#) prove tight $O(d/T)$ mismatch and $O(d \log T)$ cumulative regret for noiseless inverse optimization under i.i.d. contexts, using proper consistency or incentive estimators. Our objective is different: total target reward gap under adversarial menus, without estimating a reward parameter. Their convex estimators and experiments provide a useful computational counterpart to our information-theoretic result.

3 ONE VOTE ACROSS ALL SCALES

For $j \geq 1$, set

$$\epsilon_j = 2^{-j-1}, \quad \theta_j = 2^{-j}, \quad \alpha_j = 2^{-j}.$$

Choose a proper ϵ_j -cover $C_j \subset \mathcal{G}$. Every $h \in C_j$ induces the tolerant binary proxy

$$q_{j,h}(x, a) = \mathbf{1}\{h(x, a) \leq \epsilon_j\}. \quad (4)$$

Setting	prior general guarantee	additional restriction	this work
Correct demonstrations	$O(\log \mathcal{R})$	finite $\mathcal{R}$	entropy integral
Inverse optimization	$O(d^4 \log T)$	proper, efficient	$O(d)$
Linear recommendation	$O(d \log T)$	efficient	$O(d)$
Linear recommendation	$\exp(O(d \log d))$	efficient	$O(d)$
M-convex recommendation	$O(d \log d)$	M-convex menus	$O(d)$

Table 1: Cumulative hidden reward gap. Our general bound permits arbitrary bounded menus but is improper and can be inefficient. Its rank-two fixed-radius slice is exact and polynomial-time.

Initialize $w_1(j, h) = \alpha_j / |C_j|$, so the total initial weight is one. At round t , predict an action maximizing the shared vote

$$\hat{a}_t \in \arg \max_{a \in \mathcal{A}(x_t)} S_t(a), \quad S_t(a) = \sum_{j \geq 1} \sum_{h \in C_j} w_t(j, h) q_{j,h}(x_t, a). \quad (5)$$

After observing a_t^* , update each proxy by

$$w_{t+1}(j, h) = \begin{cases} 0, & q_{j,h}(x_t, a_t^*) = 0, \\ 2w_t(j, h), & q_{j,h}(x_t, a_t^*) = 1, \quad q_{j,h}(x_t, \hat{a}_t) = 0, \\ w_t(j, h), & \text{otherwise.} \end{cases} \quad (6)$$

For finite menus, the maximum exists. Section A gives two alternatives: a finite hierarchy for a known horizon, and summably approximate maximization for arbitrary menus.

The update is “mistake unaware.” It never uses ℓ_t or checks whether $\hat{a}_t$ is good. It only compares each known proxy’s two binary judgments. Importantly, the vote is global. Running one learner per scale and choosing among their actions would not give the potential below.

Why one discretization stalls. At a single tolerance ϵ , the same argument gives at most $\log_2 \mathcal{N}(\mathcal{G}, \epsilon)$ rounds with gap above 2ϵ . Bounding every other round by 2ϵ yields

$$\sum_{t \leq T} \ell_t \leq \log_2 \mathcal{N}(\mathcal{G}, \epsilon) + 2T\epsilon.$$

For polynomial entropy, optimizing ϵ leaves a $\Theta(d \log T)$ term. The shared vote does not select one resolution. It pays each scale’s complexity only on the loss interval represented by that scale.

Lemma 1 (Shared potential). *The total proxy weight $W_t = \sum_{j,h} w_t(j, h)$ never increases.*

Proof. Partition current weight according to the pair $(q(x_t, a_t^*), q(x_t, \hat{a}_t)) \in \{0, 1\}^2$, and call the four parts $W_{00}, W_{01}, W_{10}, W_{11}$. Vote maximality against the demonstrated action gives $W_{01} + W_{11} \geq W_{10} + W_{11}$, hence $W_{10} \leq W_{01}$. The update changes total weight by $W_{10} - W_{01} - W_{00} \leq 0$. $\square$

Lemma 2 (One surviving witness per scale). *Fix j and choose $h_j \in C_j$ with $\|h_j - g_{r_*}\|_\infty \leq \epsilon_j$. This proxy is never deleted. Whenever $\ell_t > \theta_j$, its weight doubles.*

Proof. Optimality of the demonstration gives $g_{r_*}(x_t, a_t^*) = 0$, so $h_j(x_t, a_t^*) \leq \epsilon_j$ and its proxy calls the demonstration good. If the proxy also calls $\hat{a}_t$ good, then

$$\ell_t \leq h_j(x_t, \hat{a}_t) + \epsilon_j \leq 2\epsilon_j = \theta_j.$$

The contrapositive gives the doubling claim. $\square$

4 HORIZON-FREE ENTROPY BOUNDS

Let $M_T(u) = |\{t \leq T : \ell_t > u\}|$.

Theorem 3 (Multiscale demonstration bound). *For every realizable adaptive sequence and every $T, j \geq 1$, the algorithm satisfies*

$$M_T(2^{-j}) \leq \log_2 \mathcal{N}(\mathcal{G}, 2^{-j-1}) + j. \quad (7)$$

Consequently,

$$\sum_{t=1}^T \ell_t \leq \underbrace{\sum_{j \geq 1} 2^{-j} \log_2 \mathcal{N}(\mathcal{G}, 2^{-j-1})}_{\mathfrak{H}_{\text{dyad}}(\mathcal{G})} + 2. \quad (8)$$

Moreover,

$$\mathfrak{H}_{\text{dyad}}(\mathcal{G}) \leq 4 \int_0^{1/4} \log_2 \mathcal{N}(\mathcal{G}, \epsilon) d\epsilon. \quad (9)$$

Proof. The witness at scale j starts with weight $2^{-j}/|C_j|$. Lemmas 1 and 2 give

$$\frac{2^{-j}}{|C_j|} 2^{M_T(2^{-j})} \leq w_{T+1}(j, h_j) \leq W_{T+1} \leq 1,$$

which proves (7). Since $\ell_t \in [0, 1]$, layer cake gives $\sum_t \ell_t = \int_0^1 M_T(u) du$. On $u \in (2^{-j}, 2^{-j+1}]$, monotonicity gives $M_T(u) \leq M_T(2^{-j})$. The interval has width 2^{-j} . Summing (7) over these intervals proves (8), using $\sum_{j \geq 1} j 2^{-j} = 2$. Finally, monotonicity of covering numbers gives

$$2^{-j} \log_2 \mathcal{N}(\mathcal{G}, 2^{-j-1}) \leq 4 \int_{2^{-j-2}}^{2^{-j-1}} \log_2 \mathcal{N}(\mathcal{G}, \epsilon) d\epsilon.$$

The intervals partition $(0, 1/4]$, proving (9). $\square$

The dyadic bound is primary. Unlike a single discretization, no term proportional to $T\epsilon$ remains. The integral form makes its effective-dimension content transparent.

Corollary 4 (Polynomial entropy). *If $\mathcal{N}(\mathcal{G}, \epsilon) \leq (A/\epsilon)^d$ for all $\epsilon \leq 1/4$, with $A \geq 1$, then*

$$\sum_{t=1}^T \ell_t \leq d(\log_2 A + 3) + 2. \quad (10)$$

Proof. At scale j , the logarithm in (8) is at most $d(\log_2 A + j + 1)$. Use $\sum_j 2^{-j} = 1$ and $\sum_j j 2^{-j} = 2$. $\square$

This dependence cannot be improved in general.

Theorem 5 (Dimension lower bound). *For every d , there is a continuous reward class with $\log \mathcal{N}(\mathcal{G}, \epsilon) = \Theta(d \log(1/\epsilon))$ for which every randomized learner has worst-case expected cumulative gap at least $d/2$.*

Construction. Let $\mathcal{X} = [d]$, $\mathcal{A} = \{-1, +1\}$, $w \in [-1, 1]^d$, and $r_w(i, a) = (1 + aw_i)/2$. Draw w uniformly from $\{-1, +1\}^d$ and present contexts $1, \dots, d$. Before the first demonstration at context i , the transcript is independent of w_i . Every prediction is wrong with probability $1/2$, and a wrong action has unit gap. Yao's principle (Yao, 1977) proves the lower bound. On the positive subcube, the gap metric equals the ℓ_∞ parameter metric, which gives the entropy claim. Full details are in Appendix D. $\square$

5 CONTINUOUS REWARD CLASSES

We state results at their natural gap scale. If $0 \leq g_r \leq G$, apply the preceding theorem to g_r/G and multiply the bound by G .

Corollary 6 (Linear contextual recommendation). *Let $r_w(x, a) = \langle w, \phi(x, a) \rangle$, with $\|w\|_2 \leq B$, and suppose each feature menu has Euclidean diameter at most K . There is a deterministic menu-aware action predictor with*

$$\sum_{t=1}^T \left[\max_{a \in \mathcal{A}(x_t)} \langle w_*, \phi(x_t, a) \rangle - \langle w_*, \phi(x_t, \hat{a}_t) \rangle \right] \leq BK \{d(\log_2 3 + 3) + 3\}. \quad (11)$$

The bound holds for arbitrary compact feature menus. For finite menus, the final additive constant in (11) improves from three to two.

Proof. The gap range is at most BK . Rewriting the gap as a maximum of $\langle w, \phi(x, b) - \phi(x, a) \rangle$ gives

$$\|g_w - g_v\|_\infty \leq K\|w - v\|_2.$$

After normalizing by BK , an ϵB -cover of the parameter ball is enough. The volumetric bound is $(1 + 2/\epsilon)^d \leq (3/\epsilon)^d$ for $\epsilon \leq 1$. The finite compact-menu implementation in Appendix A adds one to Corollary 4. $\square$

In the standard normalization $B = 1$ and $\mathcal{A}(x) \subseteq \{a : \|a\|_2 \leq 1\}$, one has $K \leq 2$, giving $O(d)$ regret. This improves both general bounds in Table 1 and answers the polynomial finite-regret question for the original menu-aware recommendation protocol. The distinction from proper online inverse optimization is important. Our vote need not be maximized by any single w , while proper methods output an objective estimate and take its optimal action (Sakaue et al., 2025; Oki & Sakaue, 2026). We do not improve their computational guarantees.

5.1 AN EXACT POLYNOMIAL-TIME SLICE

The exponential dependence above is not unavoidable. It disappears when the gap class has one-dimensional intrinsic geometry.

Proposition 7 (Exact runtime on a Lipschitz curve). *Suppose every normalized gap function is g_s for some $s \in [0, Q]$ and $\|g_s - g_{s'}\|_\infty \leq L|s - s'|$. Assume finite menus of size at most k and reward evaluation cost C , with $T \geq 2$. The known-horizon hard algorithm is exact, uses at most*

$$P_T \leq \lceil \log_2 T \rceil + 8LQT \quad (12)$$

proxies, and takes $O(kCTP_T)$ total time and $O(P_T + k)$ memory. The robust algorithm has the same guarantees with 8 replaced by 12.

Proof. Partitioning $[0, Q]$ into equal subintervals and taking their midpoints gives a proper ϵ -cover of size at most $1 + LQ/\epsilon$. For the hard hierarchy, summing at $\epsilon_j = 2^{-j-1}$ over $j \leq J = \lceil \log_2 T \rceil$ gives $P_T \leq J + LQ(2^{J+2} - 4) \leq J + 8LQT$. At each round, evaluating all k rewards for one proxy gives its gaps and tolerant set in $O(kC)$ time. Streaming over proxies computes all votes and updates. The robust cover radius $2^{-j}/3$ gives $J + 12LQT$ proxies. $\square$

Corollary 8 (Efficient fixed-radius rank-two recommendation). *Let $r_\vartheta(x, a) = b(x, a) + \langle B(\cos \vartheta, \sin \vartheta), \phi(x, a) \rangle$, where $\vartheta \in [0, 2\pi]$, feature-menu diameter is at most D , menu cardinality is at most k , and all gaps are at most $G > 0$. The exact finite algorithm has cumulative gap at most*

$$G \left\{ \log_2 \left(1 + \frac{2\pi BD}{G} \right) + 6 \right\} \quad (13)$$

and total time $O(kC[1 + BD/G]T^2)$. The robust version is also exact and polynomial-time, with cumulative gap $O_\gamma(\Delta_T + G[1 + \log(1 + BD/G)])$.

Indeed, changing the angle by u changes the normalized gap by at most $(BD/G)|u|$. Thus Proposition 7 applies with $Q = 2\pi$, while the circle has normalized covering number at most $(1 + 2\pi BD/G)/\epsilon$. Corollary 4 and the one-unit finite-hierarchy truncation prove (13). The menus remain arbitrary. The vote is still potentially improper, but it is enumerated exactly rather than approximated by a surrogate.

Corollary 9 (Low-rank bilinear rewards). *Let $r_W(x, a) = \langle W, \Phi(x, a) \rangle_F$ for $m \times n$ matrices with $\text{rank}(W) \leq r$ and $\|W\|_F \leq B$. If every feature menu has Frobenius diameter at most K , then cumulative gap is $O(BKr(m+n))$.*

The proof uses the standard $(C/\epsilon)^{(m+n+1)r}$ cover of the rank- r Frobenius ball and appears in Appendix C. The same calculation gives $O(Gp \log(1+L))$ for a bounded p -parameter class whose normalized gap map has parameter Lipschitz constant L .

Corollary 10 (Generalized linear rewards). *Let $r_w(x, a) = \sigma(\langle w, \phi(x, a) \rangle) \in [0, G]$, where $\|w\|_2 \leq B$, $\|\phi(x, a)\|_2 \leq R$, and σ is L -Lipschitz. Then cumulative gap is*

$$O\left(Gd \left[1 + \log\left(1 + \frac{BLR}{G}\right)\right]\right).$$

Changing w to v changes the optimal value and the chosen action's value by at most $LR\|w - v\|_2$ each. Thus the gap map is $2LR$ -Lipschitz, and the generic parameter-cover calculation in Appendix C applies. This covers another continuous family posed by Joshi et al. (2026).

Corollary 11 (Bounded ReLU reward networks). *Let $z(x, a) \in \mathbb{R}^{d_0}$ satisfy $\|z(x, a)\|_2 \leq R$. Consider depth- L scalar ReLU networks with p weights, no biases, layer spectral norms at most S , concatenated parameter norm at most B , and output clipped to $[0, G]$. There is a demonstration learner whose cumulative target gap is at most*

$$G \left\{ p \left[\log_2 \left(1 + \frac{4BRS^{L-1}\sqrt{L}}{G} \right) + 3 \right] + 2 \right\}. \quad (14)$$

The same statement covers bounded biases after the standard constant-coordinate augmentation.

The proof in Appendix C.4 telescopes perturbations through the layers. This is an explicit finite-gap result for neural rewards, rather than an assumption that their entropy integral is finite. It is parameter-count dependent and information-theoretic. It does not assert that the exponentially large proxy vote is a practical neural-network training algorithm.

6 SUBOPTIMAL DEMONSTRATIONS

Suppose now that the observed action has target gap $\delta_t = g_{r_*}(x_t, a_t^*)$, and let $\Delta_T = \sum_t \delta_t$. At scale j , set $\rho_j = 2^{-j}/3$, take a ρ_j -cover, and use the more tolerant proxy $q_{j,h}(x, a) = \mathbf{1}\{h(x, a) \leq 2\rho_j\}$. Initialize with the same priors and predict by the shared vote. For fixed $\gamma \in (0, 1)$, replace (6) by

$$w_{t+1} = w_t(1 + \gamma)^{1-q(x_t, \hat{a}_t)}(1 - \gamma)^{1-q(x_t, a_t^*)}. \quad (15)$$

Theorem 12 (Robust multiscale hedging). *Let*

$$A_\gamma = \frac{\ln(1/(1-\gamma))}{\ln(1+\gamma)}, \quad \mathfrak{H}_\gamma(\mathcal{G}) = \sum_{j \geq 1} 2^{-j} \ln \mathcal{N}(\mathcal{G}, 2^{-j}/3).$$

For every sequence of demonstrations,

$$\sum_{t=1}^T \ell_t \leq 6A_\gamma \Delta_T + \frac{\mathfrak{H}_\gamma(\mathcal{G}) + 2 \ln 2}{\ln(1+\gamma)}. \quad (16)$$

Thus fixed γ , polynomial entropy, and bounded gap give $O(\Delta_T + d)$ cumulative gap.

Proof sketch. Using the same four weight categories as Lemma 1, the exact potential change is $\gamma W_{10} - \gamma W_{01} - \gamma^2 W_{00} \leq 0$. A target proxy within ρ_j calls every action of target gap at most ρ_j good. If it calls the prediction good, then $\ell_t \leq 3\rho_j = 2^{-j}$. If it calls the demonstration bad, then $\delta_t > \rho_j$. Taking logs of its final weight gives

$$M_T(2^{-j}) \leq A_\gamma M_T^{\text{dem}}(2^{-j}/3) + \frac{\ln \mathcal{N}(\mathcal{G}, 2^{-j}/3) + j \ln 2}{\ln(1+\gamma)}.$$

Layer cake and $\sum_j 2^{-j} M_T^{\text{dem}}(2^{-j}/3) \leq 6\Delta_T$ prove the result. Appendix B gives all details. $\square$

The coefficient on demonstrator gap can be made arbitrarily close to one. This requires a finer geometric scale grid and wider proxy tolerance, which increase only the entropy term.

Corollary 13 (Near-unit optimistic comparison). *For every $\eta \in (0, 1]$, there is a choice of scales, tolerance, and soft-update rate such that*

$$\sum_{t=1}^T \ell_t \leq (1 + \eta) \Delta_T + O\left(\frac{d \ln(A/\eta) + \ln(1/\eta)}{\eta}\right) \quad (17)$$

whenever $\mathcal{N}(\mathcal{G}, \epsilon) \leq (A/\epsilon)^d$ and $A \geq 1$.

Appendix B.1 proves a more general entropy form. In particular, the learner can compete nearly one-for-one with a suboptimal demonstrator while retaining a horizon-free complexity penalty.

This is optimistic robustness: exact demonstrations recover a finite bound, while errors cost their actual reward gap rather than their count. The theorem is incomparable computationally with the efficient $O(d \log T)$ robust linear methods of [Sakaue et al. \(2025\)](#); [Sakaue \(2026\)](#) and the structured M-convex result of [Oki & Sakaue \(2026\)](#).

7 FINITE AUDITS

7.1 ADAPTIVE STRESS TEST

We instantiate the finite algorithm on a two-action continuous class. Let $r_w(\varphi, a) = a \langle w, v_\varphi \rangle / 2$ for unit vectors $w, v_\varphi \in \mathbb{R}^2$ and $a \in \{-1, +1\}$. Uniform circular nets are proper gap covers. On every round, an adaptive adversary selects the largest-gap current error from a fixed set of probes placed geometrically around the target’s two decision boundaries. If no probe is an error, it supplies an uninformative singleton menu. Each method faces its own adversary, so this is a white-box stress test rather than a common-transcript performance comparison.

The left panel of Figure 1 reports 12 target rotations with $T = 512$ and $J = 9$. The multiscale learner uses 6,427 proxies. Its mean cumulative gap is 2.07 and its observed maximum is 2.54, below the executable finite-hierarchy certificate 7.67. Single-tolerance learners at $\epsilon \in \{1/4, 1/32\}$ use 13 and 101 proxies and have mean gaps 88.92 and 21.32. An oracle that selects, separately for each target and after seeing the final losses, the best of all nine hierarchy tolerances has mean gap 2.00. It selects $\epsilon = 1/512$ or $1/1024$ and uses 1,609 or 3,217 proxies. Thus multiscale is within 3.3% of this in-hindsight comparator without choosing a tolerance. The point is not that an enumerated vote is efficient. It is that one hierarchy adapts to losses spanning all displayed resolutions, in the precise finite setting used by the proof.

7.2 MOVIELENS RANK-TWO AUDIT

We next instantiate Corollary 8 on MovieLens 100K, which contains 100,000 integer ratings from 943 users on 1,682 movies ([Harper & Konstan, 2015](#)). A deterministic biased rank-two matrix factorization has training RMSE 0.8504. For each of the ten users with most ratings, we freeze the item biases and factors, fix the fitted user-vector norm, and hide only its direction on the circle. We normalize the entire resulting reward family to gap range one. Each of 256 rounds draws a ten-movie menu from that user’s rated movies. The demonstrated action is the movie with highest actual rating, with identifier tie-breaking independent of the hidden target. It need not maximize the fitted latent reward, so we run the robust hierarchy with $\gamma = 1/4$.

Each user has 4,810 proper proxies. After factorization, the full ten-user online audit, including a baseline, takes 1.62 CPU seconds. Mean cumulative latent gap is 3.53 for multiscale, 14.82 for the demonstrated-rating policy, and 7.74 for a proper pairwise-hinge learner with unit-circle retraction. The corresponding medians are 2.90, 13.78, and 5.49. Multiscale improves on the demonstrator for all ten users and on the proper learner for seven. Every potential stays at most one, and every run is below its executable robust certificate. Figure 1 shows both audits.

This audit is deliberately modest. The fitted model defines the latent target and is not evaluated for held-out prediction. Menus are resampled from historical ratings, whose exposure is shaped by the MovieLens interface and prior recommenders ([Harper & Konstan, 2015](#)). The experiment




Figure 1: **Left:** adaptive boundary stress test, with means and 10th–90th percentiles over target rotations. **Right:** MovieLens audit, with means and 10th–90th percentiles over ten users. All gaps are hidden from the learners.

tests the exact algorithm on non-synthetic menus and noisy demonstrations. It is not a production recommender benchmark.

8 FAST STATISTICAL RATES AND LIMITATIONS

Let contexts be i.i.d. from a distribution D , while the optimal demonstration may be chosen adaptively after seeing each context. Run the online learner for m rounds and return the uniform mixture of its round policies.

Corollary 14 (Online to batch). *Let $B_G = \mathfrak{H}_{\text{dyad}}(\mathcal{G}) + 2$. The mixture $\bar{\pi}_m$ satisfies $\mathbb{E}[\text{gap}_D(\bar{\pi}_m)] \leq B_G/m$. With probability at least $1 - \delta$,*

$$\text{gap}_D(\bar{\pi}_m) \leq \frac{B_G + \ln(1/\delta)}{(1 - e^{-1})m}. \quad (18)$$

Here gap_D is expected target suboptimality on a fresh context. The proof in Appendix E uses the conditional inequality $\mathbb{E}[e^{-\ell_t} \mid \mathcal{F}_{t-1}] \leq \exp(-(1 - e^{-1})\mathbb{E}[\ell_t \mid \mathcal{F}_{t-1}])$. Thus polynomial entropy yields the fast $O((d + \log(1/\delta))/m)$ rate asked for by the continuous-class problem.

Three limitations delimit the result. First, the class must be totally bounded in the uniform gap metric, which is stronger than distribution-dependent entropy. Second, the finite implementation may contain $T^{\Theta(d)}$ proxies and optimize a nonconvex vote. Finding a polynomial-time surrogate is open beyond tractable geometries such as Proposition 7. Third, an improper recommendation need not reveal a coherent estimate of the user’s objective. If parameter recovery or properness is required, our contextual-recommendation corollary does not apply. Within this scope, integrable gap entropy yields finite total hidden gap from action-only demonstrations.

REPRODUCIBILITY STATEMENT

Every formal claim is proved in the main text or appendix. The supplement contains fixed-seed scripts for both update invariants and both panels of Figure 1, together with exact reference JSON. The synthetic audit needs no downloaded data. The MovieLens script fetches the public 100K release from GroupLens and uses no specialized hardware.

REFERENCES

- Pieter Abbeel and Andrew Y. Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the Twenty-First International Conference on Machine Learning*, 2004.
- Idan Attias, Steve Hanneke, Alkis Kalavasis, Amin Karbasi, and Grigoris Velezgas. Optimal learners for realizable regression: PAC learning and online learning. In *Advances in Neural Information Processing Systems*, volume 36, 2023. URL

- https://proceedings.neurips.cc/paper_files/paper/2023/hash/8c22e5e918198702765ecff4b20d0a90-Abstract-Conference.html.
- Andreas Bärnmann, Sebastian Pokutta, and Oskar Schneider. Emulating the expert: Inverse optimization through online learning. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 400–410. PMLR, 2017. URL <https://proceedings.mlr.press/v70/barmann17a.html>.
- Omar Besbes, Yuri Fonseca, and Ilan Lobel. Contextual inverse optimization: Offline and online learning. *Operations Research*, 73(1):424–443, 2025. doi: 10.1287/opre.2021.0369. URL <https://arxiv.org/abs/2106.14015>.
- Lee Cohen, Yishay Mansour, Shay Moran, and Han Shao. Online set learning from precision and recall feedback. *arXiv preprint arXiv:2605.09565*, 2026. URL <https://arxiv.org/abs/2605.09565>.
- Ilan Doron-Arad, Idan Mehalel, and Elchanan Mossel. Online realizable regression and applications for ReLU networks. In *Proceedings of the Thirty-Ninth Conference on Learning Theory*, volume 336 of *Proceedings of Machine Learning Research*, pp. 2105–2106. PMLR, 2026. URL <https://proceedings.mlr.press/v336/doron-arad26a.html>.
- Pouria Fatemi, Hoomaan Maskan, Suvrit Sra, and Peyman Mohajerin Esfahani. Tight generalization bounds for noiseless inverse optimization. *arXiv preprint arXiv:2605.08866*, 2026. URL <https://arxiv.org/abs/2605.08866>.
- Sreenivas Gollapudi, Guru Guruganesh, Kostas Kollias, Pasin Manurangsi, Renato Paes Leme, and Jon Schneider. Contextual recommendations and low-regret cutting-plane algorithms. In *Advances in Neural Information Processing Systems*, volume 34, pp. 22498–22508, 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/bdc6c33585d0cf5d2a8cb83141cd037f-Abstract.html>.
- F. Maxwell Harper and Joseph A. Konstan. The MovieLens datasets: History and context. *ACM Transactions on Interactive Intelligent Systems*, 5(4):19:1–19:19, 2015. doi: 10.1145/2827872. URL <https://doi.org/10.1145/2827872>.
- Héctor Jimenez, Alexander Kozachinskiy, and Vicente Opazo. Polynomial-time mistake-bounded language generation. *arXiv preprint arXiv:2606.16077*, 2026. URL <https://arxiv.org/abs/2606.16077>.
- Nirmit Joshi, Gene Li, Siddharth Bhandari, Shiva Prasad Kasiviswanathan, Cong Ma, and Nathan Srebro. Learning to answer from correct demonstrations. In *International Conference on Learning Representations*, 2026. URL <https://openreview.net/forum?id=69fIHgLjyH>.
- Jon Kleinberg, Charlotte Peale, and Omer Reingold. Mistake-bounded language generation. *arXiv preprint arXiv:2605.10809*, 2026. URL <https://arxiv.org/abs/2605.10809>.
- Taihei Oki and Shinsaku Sakaue. Finite and corruption-robust regret bounds in online inverse linear optimization under M-convex action sets. *arXiv preprint arXiv:2602.01682*, 2026. URL <https://arxiv.org/abs/2602.01682>.
- Alireza F. Pour, Farnam Mansouri, and Shai Ben-David. Learning with multiple correct answers: Regret bounds under different feedback models. *arXiv preprint arXiv:2602.09402*, 2026. URL <https://arxiv.org/abs/2602.09402>.
- Vinod Raman, Unique Subedi, and Ambuj Tewari. Online learning with set-valued feedback. In *Proceedings of the Thirty-Seventh Conference on Learning Theory*, volume 247 of *Proceedings of Machine Learning Research*, pp. 4381–4412. PMLR, 2024. URL <https://proceedings.mlr.press/v247/raman24b.html>.
- Shinsaku Sakaue. Simple projection-free algorithm for contextual recommendation with logarithmic regret and robustness. *arXiv preprint arXiv:2603.20826*, 2026. URL <https://arxiv.org/abs/2603.20826>.

Shinsaku Sakaue, Taira Tsuchiya, Han Bao, and Taihei Oki. Online inverse linear optimization: Efficient logarithmic-regret algorithm, robustness to suboptimality, and lower bound. In *Advances in Neural Information Processing Systems*, 2025. URL <https://arxiv.org/abs/2501.14349>.

Shihao Shao, Cong Fang, Zhouchen Lin, and Dacheng Tao. Partial feedback online learning. *arXiv preprint arXiv:2601.21462*, 2026. URL <https://arxiv.org/abs/2601.21462>.

Umar Syed and Robert E. Schapire. A game-theoretic approach to apprenticeship learning. In *Advances in Neural Information Processing Systems*, volume 20, 2007.

Andrew Chi-Chih Yao. Probabilistic computations: Toward a unified measure of complexity. *Proceedings of the 18th Annual Symposium on Foundations of Computer Science*, pp. 222–227, 1977.

A TECHNICAL VARIANTS OF THE MAIN ALGORITHM

A.1 A FINITE HIERARCHY FOR A KNOWN HORIZON

For horizon $T \geq 2$, let $J = \lceil \log_2 T \rceil$ and retain only scales $1, \dots, J$. Normalize the priors as

$$\tilde{\alpha}_j = \frac{2^{-j}}{Z_J}, \quad Z_J = \sum_{k=1}^J 2^{-k} = 1 - 2^{-J}.$$

Since $\tilde{\alpha}_j \geq 2^{-j}$, the proof of (7) is unchanged for $j \leq J$. Layer cake above 2^{-J} is therefore bounded by the truncated sum in (8). The remaining interval contributes at most $T2^{-J} \leq 1$. Hence the fully finite algorithm satisfies

$$\sum_{t=1}^T \ell_t \leq \mathfrak{H}_{\text{dyad}}(\mathcal{G}) + 3. \quad (19)$$

It uses $\sum_{j=1}^J \mathcal{N}(\mathcal{G}, 2^{-j-1})$ proxies. For polynomial entropy this can be polynomial in T but exponential in d . The same truncation adds at most one to Theorem 12.

A.2 APPROXIMATE VOTE MAXIMIZATION

The infinite vote in (5) is absolutely convergent whenever $W_t < \infty$. Suppose the learner selects an action whose score is within η_t of the supremum, where $\eta_t \geq 0$ and $E = \sum_t \eta_t < \infty$. The proof of Lemma 1 becomes

$$W_{t+1} - W_t \leq W_{10} - W_{01} \leq \eta_t,$$

so $W_{T+1} \leq 1 + E$. The threshold bound incurs an additive $\log_2(1 + E)$, which integrates to the same additive constant in cumulative loss. Taking $E \leq 1$ gives $\mathfrak{H}_{\text{dyad}} + 3$. This also covers nonattained vote suprema. All sums can be justified by monotone convergence because the summands are nonnegative and total weight is finite.

For compact menus and the finite known-horizon algorithm, no approximation is needed when the gap functions are continuous. Each set $\{a : h(x, a) \leq \epsilon_j\}$ is closed, its indicator is upper semicontinuous, and a finite nonnegative weighted sum of these indicators attains a maximum.

B COMPLETE PROOF OF ROBUST HEDGING

Write $b_d = q(x_t, a_t^*)$ and $b_p = q(x_t, \hat{a}_t)$. Let $W_{b_d b_p}$ denote current weight in each category. Vote maximality gives $W_{10} \leq W_{01}$. Under (15), the four multiplicative factors are

$$(1 - \gamma^2), \quad (1 - \gamma), \quad (1 + \gamma), \quad 1$$

for categories 00, 01, 10, 11, respectively. Therefore

$$W_{t+1} - W_t = -\gamma^2 W_{00} - \gamma W_{01} + \gamma W_{10} \leq 0.$$

Fix scale j and a center h_j with $\|h_j - g_{r_*}\|_\infty \leq \rho_j$. If its proxy calls $\hat{a}_t$ good, then

$$\ell_t \leq h_j(x_t, \hat{a}_t) + \rho_j \leq 3\rho_j = 2^{-j}.$$

If it calls a_t^* bad, then $g_{r_*}(x_t, a_t^*) \geq h_j(x_t, a_t^*) - \rho_j > \rho_j$. Let P_j and D_j count rounds on which this target proxy calls the prediction and demonstration bad. Its final weight is exactly

$$\frac{2^{-j}}{|C_j|} (1 + \gamma)^{P_j} (1 - \gamma)^{D_j} \leq 1.$$

Consequently,

$$M_T(2^{-j}) \leq P_j \tag{20}$$

$$\leq \frac{\ln |C_j| + j \ln 2 + D_j \ln(1/(1 - \gamma))}{\ln(1 + \gamma)} \tag{21}$$

$$\leq A_\gamma M_T^{\text{dem}}(2^{-j}/3) + \frac{\ln \mathcal{N}(\mathcal{G}, 2^{-j}/3) + j \ln 2}{\ln(1 + \gamma)}. \tag{22}$$

For any $z \in [0, 1]$, if j_0 is the first integer with $2^{-j_0} < 3z$, then

$$\sum_{j \geq 1} 2^{-j} \mathbf{1}\{z > 2^{-j}/3\} \leq \sum_{j \geq j_0} 2^{-j} = 2^{-j_0+1} < 6z.$$

Summing this inequality over demonstration gaps proves

$$\sum_{j \geq 1} 2^{-j} M_T^{\text{dem}}(2^{-j}/3) \leq 6\Delta_T.$$

Integrating (22) over dyadic loss intervals and using $\sum_j j 2^{-j} = 2$ proves (16).

B.1 GENERAL GRID AND NEAR-UNIT COMPARISON

Fix a geometric ratio $\beta > 1$ and a proxy-width parameter $c > 1$. Set

$$\theta_j = \beta^{-j}, \quad \alpha_j = (\beta - 1)\beta^{-j}, \quad \rho_j = \frac{\theta_j}{c + 1}.$$

The priors sum to one. Use a ρ_j -cover and let a proxy call an action good when its center gap is at most $c\rho_j$. The soft update remains (15). A target proxy that calls the prediction good certifies target gap at most $(c + 1)\rho_j = \theta_j$. If it calls the demonstration bad, that demonstration has target gap greater than $(c - 1)\rho_j = \kappa\theta_j$, where $\kappa = (c - 1)/(c + 1)$. The target-weight argument therefore gives

$$M_T(\theta_j) \leq A_\gamma M_T^{\text{dem}}(\kappa\theta_j) + \frac{\ln \mathcal{N}(\mathcal{G}, \rho_j) + \ln(1/\alpha_j)}{\ln(1 + \gamma)}. \tag{23}$$

The width of the layer $(\theta_j, \theta_{j-1}]$ is exactly α_j . For any $z \in [0, 1]$, the geometric tail telescopes to

$$\sum_{j \geq 1} \alpha_j \mathbf{1}\{z > \kappa\theta_j\} < \frac{\beta}{\kappa} z.$$

Consequently,

$$\sum_{t=1}^T \ell_t \leq \frac{\beta A_\gamma}{\kappa} \Delta_T + \frac{\mathfrak{H}_{\beta,c}(\mathcal{G}) + P_\beta}{\ln(1 + \gamma)}, \tag{24}$$

where

$$\mathfrak{H}_{\beta,c}(\mathcal{G}) = \sum_{j \geq 1} \alpha_j \ln \mathcal{N}\left(\mathcal{G}, \frac{\beta^{-j}}{c + 1}\right),$$

$$P_\beta = \sum_{j \geq 1} \alpha_j \ln(1/\alpha_j) = \ln \frac{\beta}{\beta - 1} + \frac{\ln \beta}{\beta - 1}.$$

Theorem 12 is the special case $\beta = c = 2$.

To prove Corollary 13, choose

$$\gamma = \frac{\eta}{16}, \quad \beta = 1 + \frac{\eta}{16}, \quad c = 1 + \frac{16}{\eta}.$$

For $\gamma \leq 1/2$,

$$A_\gamma \leq \frac{1+\gamma}{1-\gamma} \leq 1 + \frac{\eta}{4}.$$

Also $1/\kappa = 1 + \eta/8$. Multiplying the three factors shows $\beta A_\gamma / \kappa \leq 1 + \eta$ for $\eta \in (0, 1]$. Under the polynomial entropy assumption,

$$\begin{aligned} \mathfrak{H}_{\beta,c}(\mathcal{G}) &\leq d \left[\ln(A(c+1)) + \frac{\beta \ln \beta}{\beta - 1} \right] \leq d[\ln(18A/\eta) + 2], \\ P_\beta &\leq \ln(17/\eta) + 1, \\ \frac{1}{\ln(1+\gamma)} &\leq \frac{17}{\eta}. \end{aligned}$$

Substitution into (24) proves (17).

C COVERING CALCULATIONS

C.1 LINEAR REWARDS AND GAP GEOMETRY

For completeness, fix an action a and write

$$g_w(x, a) = \max_{b \in \mathcal{A}(x)} \langle w, \phi(x, b) - \phi(x, a) \rangle.$$

For any w, v , the elementary inequality $|\max_b f_b - \max_b k_b| \leq \max_b |f_b - k_b|$ gives

$$\begin{aligned} |g_w(x, a) - g_v(x, a)| &\leq \max_{b \in \mathcal{A}(x)} |\langle w - v, \phi(x, b) - \phi(x, a) \rangle| \\ &\leq K \|w - v\|_2. \end{aligned}$$

The same expression gives $g_w(x, a) \leq BK$. The Euclidean ball of radius B has an ηB -net with cardinality at most $(1 + 2/\eta)^d$, by comparing volumes of disjoint $\eta B/2$ balls around a maximal separated set. This proves Corollary 6.

C.2 LOW-RANK BILINEAR REWARDS

The set

$$\{W \in \mathbb{R}^{m \times n} : \text{rank}(W) \leq r, \|W\|_F \leq B\}$$

has an ηB -net in Frobenius norm of cardinality at most $(9/\eta)^{(m+n+1)r}$. One way to obtain this standard bound is to cover the two Stiefel factors and the r singular values in a compact singular-value decomposition, then combine the three approximation errors. The same maximum-difference calculation as above gives

$$\|g_W - g_V\|_\infty \leq K \|W - V\|_F, \quad 0 \leq g_W \leq BK.$$

Corollary 4 with entropy exponent $(m + n + 1)r$ and $A = 9$ proves Corollary 9.

C.3 GENERIC PARAMETERIZATION

Let $\Theta \subset \mathbb{R}^p$ have an η -cover of size at most $(C/\eta)^p$. Suppose gaps lie in $[0, G]$ and $\|g_\theta - g_{\theta'}\|_\infty \leq L \|\theta - \theta'\|$. If the parameter set has radius B , then a normalized gap cover at accuracy ϵ is induced by a parameter cover at radius $G\epsilon/L$. Thus

$$\mathcal{N}(\mathcal{G}/G, \epsilon) \leq \left(\frac{CBL}{G\epsilon} \right)^p$$

whenever $BL/G \geq 1$, and the harmless replacement by $C(1 + BL/G)/\epsilon$ covers all cases. Corollary 4 gives $O(Gp[1 + \log(1 + BL/G)])$.

C.4 BOUNDED RELU NETWORKS

Write the unclipped depth- L network as

$$f_W(z) = W_L \sigma(W_{L-1} \sigma(\cdots \sigma(W_1 z))),$$

where W_L has one row and $\sigma(u) = \max\{u, 0\}$ coordinatewise. Clipping to $[0, G]$ is nonexpansive. For two parameter tuples W, V , insert one layer at a time. ReLU is 1-Lipschitz and fixes zero, while every layer has spectral norm at most S . Thus

$$\begin{aligned} |f_W(z) - f_V(z)| &\leq RS^{L-1} \sum_{l=1}^L \|W_l - V_l\|_F \\ &\leq RS^{L-1} \sqrt{L} \|W - V\|_2. \end{aligned}$$

The first line follows because the perturbation at any one layer is preceded and followed by $L - 1$ maps whose product Lipschitz constant is at most S^{L-1} . Maximization and subtraction then give

$$\|g_W - g_V\|_\infty \leq 2RS^{L-1} \sqrt{L} \|W - V\|_2.$$

The allowed parameter set lies in a Euclidean p -ball of radius B . Its radius- u covering number is at most $(1 + 2B/u)^p$. Taking $u = G\epsilon/(2RS^{L-1}\sqrt{L})$ yields

$$\mathcal{N}(\mathcal{G}/G, \epsilon) \leq \left(\frac{1 + 4BR S^{L-1} \sqrt{L}/G}{\epsilon} \right)^p, \quad \epsilon \leq 1.$$

Corollary 4 proves (14). Bias vectors can be represented by adding a constant coordinate at each layer. The same calculation then applies with the corresponding augmented input and operator-norm bounds.

D LOWER-BOUND DETAILS

For $w \in [-1, 1]^d$, define

$$r_w(i, a) = \frac{1 + aw_i}{2}, \quad g_w(i, a) = \frac{|w_i| - aw_i}{2}.$$

At a vertex, the two rewards are one and zero, hence a wrong action has gap one.

For $w, v \in [1/2, 1]^d$, action +1 is optimal at every context, and

$$\sup_{i,a} |g_w(i, a) - g_v(i, a)| = \|w - v\|_\infty.$$

Therefore the positive subcube requires at least $(c/\epsilon)^d$ balls and admits at most $(C/\epsilon)^d$ balls for numerical constants c, C . The full cube also admits such an upper bound because the displayed gap map is Lipschitz in ℓ_∞ . This proves the entropy order.

Under a uniform prior on the vertices, the observable history before round i contains only $w_1, \dots, w_{i-1}$. It is independent of w_i . Conditional on any learner randomness, either action is therefore wrong with probability $1/2$. The first d contexts contribute expected loss $d/2$. Yao's min-max principle converts this distributional lower bound for deterministic learners to a worst-case expected lower bound for randomized learners.

E ONLINE-TO-BATCH PROOF

Let π_t be the policy used on round t , based on the first $t - 1$ samples, and let $L_t \in [0, 1]$ be its realized target gap. Write

$$\mu_t = \mathbb{E}[L_t \mid \mathcal{F}_{t-1}] = \text{gap}_D(\pi_t).$$

Convexity of e^{-x} on $[0, 1]$ gives $e^{-x} \leq 1 - (1 - e^{-1})x$. Hence, with $c = 1 - e^{-1}$,

$$\mathbb{E}[e^{-L_t} \mid \mathcal{F}_{t-1}] \leq 1 - c\mu_t \leq e^{-c\mu_t}.$$

It follows that

$$Z_s = \exp \left(c \sum_{t=1}^s \mu_t - \sum_{t=1}^s L_t \right)$$

is a nonnegative supermartingale with $\mathbb{E}Z_0 = 1$. Markov's inequality gives, with probability at least $1 - \delta$,

$$c \sum_{t=1}^m \mu_t \leq \sum_{t=1}^m L_t + \ln(1/\delta) \leq B_G + \ln(1/\delta).$$

The population gap of the uniform mixture is $m^{-1} \sum_t \mu_t$, proving (18). Taking expectations and using the independence of each fresh x_t from π_t gives the expectation statement directly.

F EXECUTABLE INVARIANT CHECKS

The supplementary artifact contains four small NumPy programs. The first constructs finite linear reward covers and checks, round by round, that total weight never increases, every target proxy survives, every large target gap doubles it, each threshold count obeys its proof bound, and cumulative gap obeys the truncated entropy sum. The second injects suboptimal demonstrations and checks the soft potential, target-proxy log inequality, and robust cumulative bound. The third reproduces the left panel of Figure 1, its proxy counts, and its numeric values. The fourth downloads MovieLens 100K and reproduces the right panel, runtimes, per-user results, and robust certificates. All use fixed random seeds. These scripts are finite-precision consistency checks, not evidence in place of the proofs.