

Tight Bounds for Online Learning with Adversarially Corrupted Experts and the Limits of Adaptive-rates EWA

Jin Sima

CSol & Purdue University

SIMAJ@PURDUE.EDU

Nikolaos Papagiannis

CSol & Purdue University

NPAPAGIA@PURDUE.EDU

Changlong Wu

Arizona State University

WUCHANGL@HAWAII.EDU

Wojciech Szpankowski

CSol & Purdue University

SZPAN@PURDUE.EDU

Editors: Steve Hanneke and Tor Lattimore

Abstract

We study learning from expert advice with corrupted feedback, where the learner observes only *noisy* losses but is evaluated on the underlying *true* losses. Unlike prior work on stochastic perturbations, we consider a fully adversarial model in which perturbations are arbitrary but bounded by a per-expert budget C . We establish the first minimax-optimal risk of order $\tilde{\Theta}(C \log K)$ for the corrupted-feedback setting, even when the corruption budget C is *unknown*. Our achievability result builds on a novel potential-based analysis of the Exponentially Weighted Average (EWA) algorithm together with an epoch-based scheduling scheme inspired by the doubling trick. By contrast, we prove that the self-confident EWA and, more generally, any adaptive learning-rate scheme suffer risk at least $\Omega(C^\beta)$ for any $1 < \beta < \frac{3}{2}$, highlight a fundamental separation between risk evaluations on noisy and true losses. Finally, we extend our guarantees beyond the realizable case by establishing sublinear regret bounds. Together, these results establish a complete minimax theory for the corrupted-expert model, capturing both the optimal rates and the algorithmic principles.

1. Introduction

Learning from expert advice (Cesa-Bianchi and Lugosi, 2006; Shalev-Shwartz, 2007; Rakhlin et al., 2010) is one of the most studied paradigms in online learning, forming the basis of algorithms for sequential prediction, game theory, and adaptive decision-making. In its classical form, the learner competes with the best expert in hindsight under adversarially chosen losses, and tight guarantees are well understood. Yet this clean picture changes dramatically when the learner’s observations are *corrupted*. In many realistic environments—from adversarial platforms and cybersecurity to noisy crowdsourcing—the feedback revealed to the learner may be an intentionally distorted version of the true losses, while performance is ultimately evaluated with respect to those hidden truths.

To further motivate our setting, consider online binary classification under *Massart noise* (Ben-David et al., 2009). Let $\mathcal{H} = \{h_1, \dots, h_K\} \subset \{0, 1\}^{\mathcal{X}}$ be a hypothesis class of size K , and let Nature fix an unknown target $h_{k^*} \in \mathcal{H}$. At each round $t = 1, \dots, T$: (i) a feature \mathbf{x}_t is revealed, (ii) the learner predicts $\hat{y}_t \in \{0, 1\}$, and (iii) a *noisy* label $\tilde{y}_t = \text{Bernoulli}(\gamma_t) \oplus h_{k^*}(\mathbf{x}_t)$ is observed,

where $\gamma_t \leq \gamma < \frac{1}{2}$ for some *unknown* γ_t (but the upper bound γ is known). The learner aims to minimize the cumulative loss with respect to the *true* labels:

$$\tilde{r}_T(\mathcal{H}, \gamma) := \sup_{\mathbf{x}^T, h_{k^*}, \gamma^T} \mathbb{E} \left[\sum_{t=1}^T \mathbf{1}\{h_{k^*}(\mathbf{x}_t) \neq \hat{y}_t\} \right].$$

It was shown in [Ben-David et al. \(2009\)](#) that

$$\tilde{r}_T(\mathcal{H}, \gamma) \leq \frac{\log |\mathcal{H}|}{1 - 2\sqrt{\gamma(1-\gamma)}}. \quad (1)$$

Note that a key nontrivial aspect of the result in (1) is that the noise parameters γ_t are *unknown*, which prevents the learner from directly estimating the actual noise levels. This phenomenon was further generalized by [Wu et al. \(2024\)](#) to more general stochastic noise mechanisms. However, a major limitation of [Ben-David et al. \(2009\)](#) and [Wu et al. \(2024\)](#) is that the noise upper bound γ remains assumed to be *known* to the learner, an assumption that is often unrealistic in practice.

To address these limitations and to capture scenarios where stochastic assumptions completely break down, [Papagiannis et al. \(2025\)](#) recently introduced a general framework (namely the so-called “nasty-noise” model) in which corruption can occur *arbitrarily*, without adhering to any stochastic mechanism, but subject only to a bounded corruption budget. Formally, at each time step t , Nature selects an arbitrary *true* loss vector $\ell_t \in [0, 1]^K$ (assigning losses to K experts) but reveals only a perturbed *noisy* version $\tilde{\ell}_t \in [0, 1]^K$. The choice of $(\ell_t, \tilde{\ell}_t)$ is *completely* arbitrary, except that: (i) for every expert, the discrepancy between its cumulative true and noisy losses is upper bounded by a parameter C (which may grow with the horizon T); and (ii) there exists at least one expert that incurs zero cumulative true loss (i.e., the well-specified case).

It was shown in [Papagiannis et al. \(2025\)](#) that if the corruption parameter C is *known*, then the minimax cumulative risk grows as $\Theta(C \log K)$, via an elimination-based algorithm. When C is *unknown*, however, [Papagiannis et al. \(2025\)](#) proved that the standard Exponentially Weighted Average (EWA) algorithm ([Cesa-Bianchi and Lugosi, 2006](#)) with a constant learning rate attains cumulative risk of order $O(C^2 + C \log K)$. This result was further employed in ([Papagiannis et al., 2025](#), Theorem 3) to derive a constant-risk bound for the Massart-noise model *without* assuming knowledge of the noise upper bound γ .¹ In addition, they showed an $\Omega(C^2 / \log K)$ lower bound for EWA with constant learning rates. These findings leave open the following fundamental questions: *Can one achieve the minimax risk of order $O(C \log K)$ without knowing the parameter C ? If so, can this optimal rate be obtained by the EWA algorithm with an adaptive schedule of learning rates based on the observed (noisy) losses?*

Summary of results. In this paper, we provide a complete resolution to the open questions above, which were also explicitly posed in [Papagiannis et al. \(2025\)](#). Specifically, we show (in Theorem 3) that for any learning rate η , the risk (on true losses) of the EWA algorithm is bounded by $O\left(\frac{\log K}{\eta} + \eta C^2 + C \log K\right)$. By choosing $\eta = 1/C$, this recovers the $O(C \log K)$ bound of [Papagiannis et al. \(2025\)](#) for the case of *known* C . Crucially, however, our result holds for *any* learning rate η , rather than only for given constants. The proof relies on a novel algebraic variant of the

1. In fact, any risk bound for the “nasty-noise” model *without* known C can be turned into such a guarantee via a black-box reduction.

“elliptic potential lemma” (Cesa-Bianchi and Lugosi, 2006), which allows us to tightly control the discrepancies between risks on noisy and true losses. This technique is of independent interest.

Note that Theorem 3 still requires tuning the parameter η , which is infeasible when C is unknown. To address this, we propose a novel epoch-based algorithm (see Algorithm 2 and Theorem 5) that employs a “guess-and-double” strategy for handling the unknown corruption budget. This achieves a risk upper bound of order $\tilde{O}(C \log K)$ in the case of *unknown* C , matching the general $\Omega(C \log K)$ lower bound up to only a $\log T$ factor, and substantially improving upon the $O(C^2 + C \log K)$ upper bound from Papagiannis et al. (2025) for large C .

Note that the “guess-and-double” strategy resembles the classical *doubling trick* in the online-learning literature, which *restarts* the algorithm whenever the current guess is violated. A natural question is whether a more *smoothed* adaptive schedule—such as the self-confident EWA of Auer et al. (2002b)—could attain the optimal $C \log K$ risk without restarts. Quite surprisingly, Theorem 4 shows that *any* adaptive selection of the learning rate (however it depends on the observed noisy losses) incurs minimax risk at least $\Omega(C^\beta)$ for any $1 < \beta < \frac{3}{2}$. This, to our knowledge, is the first evidence that demonstrates adaptive learning-rate schedules for EWA alone cannot match the guarantees of doubling-trick style algorithms, highlighting a fundamental separation between the two approaches. Finally, we extend our results to scenarios where the true losses are *misspecified*, i.e., the cumulative true loss of the best expert is nonzero. We introduce a notion of *regret* with respect to the true losses,² and establish in Theorem 6 a sublinear regret bound $O(\sqrt{C(C + L_{\min}^*) \ln K} + C \log K)$, where L_{\min}^* denotes the minimal cumulative true loss of the best expert. In Theorem 7 we show that Algorithm 2 also applies when C and L_{\min}^* are unknown.

Highlights. In summary, our results (i) settle the minimax rate for learning with adversarially corrupted experts by achieving $\tilde{O}(C \log K)$ risk *without* prior knowledge of C , (ii) reveal a fundamental separation between risks on noisy and true losses by showing the inherent suboptimality of EWA with adaptive learning-rate schedules, and (iii) extend the guarantees beyond the well-specified setting to obtain sublinear regret when the best expert incurs nonzero loss. Collectively, these contributions provide a complete minimax theory for corrupted-feedback expert advice.

1.1. Other related works

Our framework is mostly comparable to Amir et al. (2020), which considers an adversarial corruption model with a budget on the cumulative *supremum* corruption across experts at each time step. This assumption is more stringent, as it directly yields a trivial upper bound on the discrepancy between noisy and true losses. In contrast, our model imposes a *per-expert* corruption budget, which is more flexible and motivated by the semi-stochastic noise models studied in Ben-David et al. (2009); Wu et al. (2024) (see also Bshouty et al. (2002)). In fact, in our model the contribution of corruption to the risk can even dominate the base risk attainable under noisy losses, which is the main nontriviality of our setting. Beyond corruption models, our setting also bears resemblance to *prediction with limited feedback*, as treated in (Cesa-Bianchi and Lugosi, 2006, Section 6) (see also Kakade and Kalai (2005); Wu et al. (2023)). However, our “nasty expert” differs in that the noisy losses are selected *arbitrarily* (subject to certain constraints), rather than through a *time-independent* feedback function as in Cesa-Bianchi and Lugosi (2006). Prediction with *stochastic noise* was first investigated in Weissman et al. (2001), which studies sequential prediction with binary outcomes

2. Note that, both our notions of *risk* and *regret* are defined with respect to the true losses; the noisy losses serve only as observations and do not appear in the performance measure.

passing through a Binary Symmetric Channel (BSC). This setting was recently generalized in [Bhatt and Kostina \(2024, 2025\)](#) (see also references therein) under the assumption that the true losses are estimable. The regret analysis of *observable* losses using the Exponentially Weighted Average (EWA) algorithm has been widely studied; see [Auer et al. \(2002b\)](#); [Cesa-Bianchi and Lugosi \(2006\)](#); [Gaillard et al. \(2014\)](#); [Mourtada and Gaiffas \(2019\)](#).

2. Problem formulation and preliminaries

Let K be the number of experts. We consider the following online learning framework, which operates in T rounds. At each time step $t = 1, 2, \dots, T$, the following events occur sequentially: (i) Nature selects the *true* loss vector $\ell_t \in [0, 1]^K$ but keeps it secret from the learner; (ii) The learner selects *distribution* $\hat{p}_t \in \Delta([K])$ (where $\Delta([K])$ is a set of distributions over $[K]$) samples $\hat{k}_t \sim \hat{p}_t$, predicts \hat{k}_t , and incurs the loss $\ell_t[\hat{k}_t]$; (iii) Nature selects a *noisy* loss vector $\tilde{\ell}_t \in [0, 1]^K$ and reveals it to the learner.

Let $C \leq T$ be an arbitrary parameter that controls the noise level, which may grow with respect to T . The goal of the learner is to design a prediction strategy \hat{p}^T that minimizes:

$$\text{risk}_{T,K,C} := \sup_{\ell^T, \tilde{\ell}^T} \mathbb{E} \left[\sum_{t=1}^T \ell_t[\hat{k}_t] \right] = \sup_{\ell^T, \tilde{\ell}^T} \sum_{t=1}^T \langle \hat{p}_t, \ell_t \rangle, \quad (2)$$

where $\langle \cdot, \cdot \rangle$ is scalar product, and $\ell^T := \{\ell_1, \dots, \ell_T\}$, $\tilde{\ell}^T := \{\tilde{\ell}_1, \dots, \tilde{\ell}_T\}$ are selected *arbitrarily* subject to the following constraints: (A1) There *exists* $k^* \in [K]$ such that

$$\sum_{t=1}^T \ell_t[k^*] = 0; \quad (3)$$

(A2) For all $k \in [K]$, we have

$$\sum_{t=1}^T |\ell_t[k] - \tilde{\ell}_t[k]| \leq C. \quad (4)$$

Here, the first constraint ensures that there exists an expert incurring zero cumulative *true* loss (i.e., realizable), while the second constraint ensures that for any expert, the discrepancy between the true and noisy cumulative losses is upper bounded by C . Crucially, we assume that the parameter C is *unknown* to the learner. In Section 3.4, we will further relax the assumption (3) by considering a *regret* formulation. However, we should emphasize that even with the realizable assumption, the problem remains fundamentally challenging and leads to unexpected results.

Exponentially Weighted Average (EWA). Arguably the most established approach to learning from expert advice is the Exponentially Weighted Average (EWA) algorithm ([Cesa-Bianchi and Lugosi, 2006](#)), given in Algorithm 1 for the case of observable noisy losses. The following risk bound on *noisy* losses is classical ([Auer et al., 2002b](#)).

Lemma 1 (Risk of EWA on Noisy Losses) *Let $\tilde{\ell}_1, \dots, \tilde{\ell}_T \in \mathbb{R}^K$ be arbitrary noisy loss vectors, and let $\hat{p}_1, \dots, \hat{p}_T$ be the distributions produced by Algorithm 1. Then*

$$\sum_{t=1}^T \langle \hat{p}_t, \tilde{\ell}_t \rangle \leq O\left((\eta + 1)\tilde{L} + \frac{\log K}{\eta}\right) \leq O\left((\eta + 1)C + \frac{\log K}{\eta}\right),$$

where $\tilde{L} := \min_{k \in [K]} \sum_{t=1}^T \tilde{\ell}_t[k] \leq C$ denotes the minimal cumulative noisy loss.

Algorithm 1 EWA algorithm with Noisy Losses

- 1: **Input:** Learning rate $\eta > 0$, number of experts K
- 2: Initialize weights $w_k^1 \leftarrow 1$ for all $k \in [K]$
- 3: **for** $t = 1, 2, \dots, T$ **do**
- 4: Compute probability distribution over experts:

$$\hat{p}_t[k] \leftarrow \frac{w_k^t}{\sum_{j=1}^K w_j^t} \quad \text{for } k \in [K]$$

- 5: Sample \hat{k}_t from the distribution \hat{p}_t
 - 6: Predict using expert \hat{k}_t
 - 7: Observe the noisy loss vector $\tilde{\ell}_t \in [0, 1]^K$
 - 8: Update weights for all experts: $w_k^{t+1} \leftarrow w_k^t \cdot \exp(-\eta \tilde{\ell}_t[k])$ for $k \in [K]$
 - 9: **end for**
-

In order to convert such a risk bound to the *true* loss, we consider the following decomposition:

$$\sum_{t=1}^T \langle \hat{p}_t, \ell_t \rangle = \underbrace{\sum_{t=1}^T \langle \hat{p}_t, \tilde{\ell}_t \rangle}_{\text{Algorithmic (noisy) risk}} + \underbrace{\sum_{t=1}^T \langle \hat{p}_t, \ell_t - \tilde{\ell}_t \rangle}_{\text{Discrepancy between true and noisy risks}}. \quad (5)$$

Note that, in contrast to the cumulative *supremum* corruption considered in Amir et al. (2020)³, our condition (4) does *not* guarantee that the discrepancy term is bounded by C . In fact, the discrepancy may vary significantly across algorithms (i.e., the choices of \hat{p}_t), and controlling this behavior is the central technical challenge of our analysis.

To gain intuition about the behavior of the risk on *true* losses, we present the following lower bound (see also Papagiannis et al. (2025)) which – for completeness – we prove in Appendix A.

Lemma 2 (Lower Bound on True Risk) *For any prediction rule $\hat{p}_1, \dots, \hat{p}_T$ (with or without knowledge of C), there exist true and noisy loss sequences ℓ_1, \dots, ℓ_T and $\tilde{\ell}_1, \dots, \tilde{\ell}_T$ satisfying conditions (3) and (4) such that $\sum_{t=1}^T \langle \hat{p}_t, \ell_t \rangle \geq \Omega(C \cdot \log K)$.*

If we take $\eta = \sqrt{\log K / C}$ in Lemma 1, the risk of EWA on *noisy* losses scales as $\sqrt{C \cdot \log K} + C$. This is of lower asymptotic order than the lower bound in Lemma 2, implying that the discrepancy term can, in fact, be the *dominant* contribution to the true risk. Perhaps surprisingly, (Papagiannis et al., 2025, Lemma 4) further shows that when $\eta = 1$ (or any constant), the true risk of EWA is bounded below by $\Omega\left(\frac{C^2}{\log K}\right)$, even though the noisy risk is only $O(C + \log K)$.

3. Main results

In this section, we present our main results, which settle the minimax risk for the adversarially corrupted online learning model *without* prior knowledge of the corruption budget C .

3. They assume $\sum_{t=1}^T \|\ell_t - \tilde{\ell}_t\|_\infty \leq C$ for certain budget parameter C .

3.1. A tighter analysis of EWA algorithm

We start with the following key risk bound on true losses for EWA.

Theorem 3 *Let $\ell_1, \dots, \ell_T \in \mathbb{R}^K$ be arbitrary true loss vectors and $\tilde{\ell}_1, \dots, \tilde{\ell}_T \in \mathbb{R}^K$ be corresponding noisy loss vectors satisfying conditions (3) and (4). Let $\hat{p}_1, \dots, \hat{p}_T$ be the distributions produced by EWA Algorithm 1 with any learning rate $\eta > 0$. Then with $\tilde{L} = \min_k \sum_{t=1}^T \tilde{\ell}_t[k]$*

$$\sum_{t=1}^T \langle \hat{p}_t, \ell_t \rangle \leq O\left(\frac{\ln K}{\eta} + \eta C \tilde{L} + C \ln K\right),$$

Proof Using the decomposition in (5), the expected true risk can be written as

$$\begin{aligned} \sum_{t \in [T]} \frac{\sum_{i \in [K]} w_i^t \ell_t[i]}{\sum_{i \in [K]} w_i^t} &= \sum_{t \in [T]} \frac{\sum_{i \in [K]} w_i^t \tilde{\ell}_t[i]}{\sum_{i \in [K]} w_i^t} + \sum_{t \in [T]} \frac{\sum_{i \in [K]} w_i^t (\ell_t[i] - \tilde{\ell}_t[i])}{\sum_{i \in [K]} w_i^t} \\ &\stackrel{(a)}{\leq} O\left(\eta C + \frac{\ln K}{\eta}\right) + \sum_{t \in [T]} \frac{\sum_{i \in [K]} w_i^t (\ell_t[i] - \tilde{\ell}_t[i])}{\sum_{i \in [K]} w_i^t} \\ &\leq O\left(\eta C + \frac{\ln K}{\eta}\right) + \sum_{k \in [K]} \max_{t \in [T]} \frac{w_k^t}{\sum_{i \in [K]} w_i^t} \sum_{t \in [T]} |\ell_t[k] - \tilde{\ell}_t[k]| \\ &\stackrel{(b)}{\leq} O\left(\eta C + \frac{\ln K}{\eta}\right) + C \sum_{k \in [K]} \max_{t \in [T]} \frac{w_k^t}{\sum_{i \in [K]} w_i^t}, \end{aligned} \quad (6)$$

where (a) follows from Lemma 1 and (b) follows by condition (4).

Let $t_k = \arg \max_{t \in [T]} \frac{w_k^t}{\sum_{i \in [K]} w_i^t}$. Without loss of generality, assume that $t_1 \leq t_2 \leq \dots \leq t_{K-1}$. In addition, W.L.O.G. assume that expert K achieves the smallest accumulative noisy loss, i.e., $\sum_{t=1}^T \ell_t[K] = \min_k \sum_{t=1}^T \tilde{\ell}_t[k]$. Then, we have

$$\begin{aligned} \sum_{k \in [K-1]} \max_{t \in [T]} \frac{w_k^t}{\sum_{i \in [K]} w_i^t} &= \sum_{k \in [K-1]} \frac{w_k^{t_k}}{\sum_{i \in [K]} w_i^{t_k}} \stackrel{(a)}{\leq} \sum_{k \in [K-1]} \frac{w_k^{t_k}}{\sum_{i \in [K-1] \setminus [k-1]} w_i^{t_k} + w_K^{t_k}} \\ &\stackrel{(b)}{\leq} \sum_{k \in [K-1]} \frac{w_k^{t_k}}{\sum_{i \in [K-1] \setminus [k-1]} w_i^{t_k} + e^{-\eta \tilde{L}}} \\ &\stackrel{(c)}{\leq} \sum_{k \in [K-1]} \ln \frac{\sum_{i \in [K-1] \setminus [k-1]} w_i^{t_k} + e^{-\eta \tilde{L}}}{\sum_{i \in [K-1] \setminus [k]} w_i^{t_k} + e^{-\eta \tilde{L}}} \\ &\leq \ln \frac{e^{-\eta \tilde{L}} + \sum_{i \in [K-1]} w_i^{t_k}}{e^{-\eta \tilde{L}}} \leq \ln K + \eta \tilde{L}, \end{aligned}$$

where (a) follows from the fact that $w_k^t > w_k^{t'}$ for $t < t'$, (b) follows since $w_K^t \geq e^{-\eta \tilde{L}}$ for all $t \in [T]$, and (c) follows since $x \leq -\ln(1-x)$ for $x \in (0, 1)$. Combining with (6), we obtain

$$\sum_{t \in [T]} \frac{\sum_{i \in [K]} w_i^t \ell_t[i]}{\sum_{i \in [K]} w_i^t} \leq O\left(\eta C + \frac{\ln K}{\eta} + C \ln K + \eta C \tilde{L}\right),$$

as needed. ■

One may notice that the key difference between the risk bound in Theorem 3 and the classical bound in Lemma 1 (on noisy losses) is the additional term ηC^2 , which scales quadratically with C . In fact, as shown in (Papagiannis et al., 2025, Lemma 4), an $\Omega(\eta C^2 / \log K)$ lower bound holds for the true risk as well, implying that this quadratic dependence is inherent to EWA. For $\eta = 1$, Theorem 3 also recovers the risk bound of (Papagiannis et al., 2025, Lemma 3), which was previously proved only for *constant* learning rates. Note that our key technique in the proof of Theorem 3, which relates the discrepancy term to a log-potential function, resembles the classical “elliptic potential lemma” (Cesa-Bianchi and Lugosi, 2006) and is of independent interest.

3.2. Fundamental limits of EWA with adaptive learning rates

In this section, we study a class of EWA variants with an adaptive learning rate η_t that may depend on the past noisy losses (e.g., the *self-confident* EWA of Auer et al. (2002a)). We show that there exists a corruption level C such that, without knowing C , any such adaptive-rate EWA algorithm suffers risk at least $\Omega(C^\beta)$ for some $1 < \beta < 1.5$. This reveals a fundamental limitation of adapting EWA to the unknown- C setting. This result should be contrasted with the lower bound $\Omega(C \log K)$ in Lemma 2, which holds for *any* algorithm (not just EWA-type methods).

Theorem 4 *Let \mathcal{E} be any adaptive-rates EWA algorithm (Auer et al., 2002a), where the learning rate η_t is tuned at each step based only on the past noisy losses $\tilde{\ell}_1, \dots, \tilde{\ell}_{t-1}$. For any $1 < \beta < \frac{3}{2}$, there exists $C \geq \frac{K^{(3-2\beta)/\beta}}{\log^{2/\beta} K}$, and a choice of noisy and true loss sequences satisfying constraints (3) and (4) such that \mathcal{E} incurs risk at least $\Omega(C^\beta)$. In particular, this is strictly larger than $\tilde{O}(C \log K)$ (with polylog(C) factors) whenever $K = \text{poly}(C)$ or $C = \text{poly}(K)$.*

Proof Note that the adaptive EWA algorithm \mathcal{E} maintains, at each round t , the distribution

$$\hat{p}_t[k] = \frac{\exp\left(-\sum_{s=1}^{t-1} \eta_s \tilde{\ell}_s[k]\right)}{\sum_{u=1}^K \exp\left(-\sum_{s=1}^{t-1} \eta_s \tilde{\ell}_s[u]\right)}, \quad k \in [K],$$

where η_t is any learning-rate schedule based on $\tilde{\ell}_1, \dots, \tilde{\ell}_{t-1}$. For any given adaptive schedule, we construct specific hard noisy and true losses that attain the claimed lower bound.

Noisy losses $\tilde{\ell}$. Let us fix $1 < \beta < \frac{3}{2}$ and assume $K^{2-\beta}$ and $K^{\beta-1}$ are integers (otherwise round with floors or ceilings, which only affects constants). We construct the noisy losses $\tilde{\ell}_t[k] \in \{0, 1\}$ in two phases:

- (i) *Upper-triangle prefix:* for $1 \leq t \leq K$ and $1 \leq k \leq K$, set $\tilde{\ell}_t[k] = \mathbf{1}\{t \geq k\}$.
- (ii) *Main phase:* for $t \geq K + 1$,
 - for all non-special $k \notin \{iK^{2-\beta} : i = 1, \dots, K^{\beta-1}\}$, set $\tilde{\ell}_t[k] = \mathbf{1}\{t \geq k\}$;
 - for each *special* index $k = iK^{2-\beta}$:

$$\tilde{\ell}_t[k] = \begin{cases} 0, & K + 1 \leq t \leq 2K + (K^{\beta-1} - i)K, \\ 1, & t \geq 2K + (K^{\beta-1} - i)K + 1, \end{cases}$$

and additionally, set $\tilde{\ell}_t[K^{2-\beta}] = 0$ for $t \in \{2K + K^\beta + 1, \dots, T\}$.

Intuitively, in the first phase we place loss 1 in an “upper-triangle” pattern for all indices $k \in [K]$. In the second phase, we place loss 1 in a “lower-triangle” pattern for certain *special* indices from $\{iK^{2-\beta} : i = 1, \dots, K^{\beta-1}\}$. We also ensure that for the particular index $k = K^{2-\beta}$ we have $\sum_{t=1}^T \tilde{\ell}_t[K^{2-\beta}] = K + K^{2-\beta} \leq 2K$. Based on the noisy losses, at each time t , we track the cumulative expected noisy loss

$$S(t) = \sum_{s=1}^t \sum_{k=1}^K \hat{p}_s[k] \tilde{\ell}_s[k],$$

and compare it against the threshold

$$\Gamma(t) = \frac{\left(\max \left\{ \min_k \sum_{s=1}^t \tilde{\ell}_s[k], \frac{K^{(3-2\beta)/\beta}}{\log^{2/\beta} K} \right\} \right)^\beta}{2e}.$$

Since the algorithm’s behavior depends only on the noisy losses, we will assign the true losses adaptively according to its behavior in the following two cases.

Case A (early-stop occurs). If $S(t) \geq \Gamma(t)$ for some t , let $k^* \in \arg \min_k \sum_{s=1}^t \tilde{\ell}_s[k]$ and define the *true* losses

$$\ell_s[k] = \begin{cases} 0, & k = k^*, \\ \tilde{\ell}_s[k], & k \neq k^*. \end{cases}$$

In this case, the only corruption comes from coordinate k^* , with budget

$$C := \max \left\{ \frac{K^{(3-2\beta)/\beta}}{\log^{2/\beta} K}, \sum_{s=1}^t \tilde{\ell}_s[k^*] \right\}.$$

Since $S(t) \geq \Gamma(t)$ and $\Gamma(t) = (\max\{\min_k \sum_{s=1}^t \tilde{\ell}_s[k], \frac{K^{(3-2\beta)/\beta}}{\log^{2/\beta} K}\})^\beta / (2e)$, we have

$$\sum_{s=1}^T \langle \hat{p}_s, \ell_s \rangle \geq \sum_{s=1}^t \langle \hat{p}_s, \tilde{\ell}_s \rangle - \sum_{s=1}^t \tilde{\ell}_s[k^*] \geq \Gamma(t) - C = \Omega(C^\beta),$$

where the last step follows since $C = (2e\Gamma(t))^{1/\beta} = o(\Gamma(t))$ for $\beta > 1$.

Case B (no early-stop). If $S(t) < \Gamma(t)$ for all $t \leq T$. We show that, for all $i = 1, \dots, K^{\beta-1} - 1$,

$$\sum_{s=(i-1)K^{2-\beta}+1}^{iK^{2-\beta}} \eta_s \geq \log K, \tag{7}$$

i.e., the leaning rates cannot be too small. Indeed, if (7) failed for some i , then by partitioning that block into $\log K$ equal sub-blocks, there exists one sub-block

$$B = \{a+1, \dots, a+L\} \subset \{(i-1)K^{2-\beta}+1, \dots, iK^{2-\beta}\}, \quad L = \frac{K^{2-\beta}}{\log K},$$

that would satisfy $\sum_{s \in B} \eta_s < 1$. We show that this will contradict to our presumption $S(t) < \Gamma(t)$. Consider the upper-triangular region $\mathcal{R} = \{(s, k) : s \in B, k \in B, k \leq s\}$. By the prefix construction $\tilde{\ell}_s[k] = \mathbf{1}\{s \geq k\}$, so $\tilde{\ell}_s[k] = 1$ for all $(s, k) \in \mathcal{R}$. Moreover, for any $(s, k) \in \mathcal{R}$,

$$\sum_{u=1}^{s-1} \eta_u \tilde{\ell}_u[k] = \sum_{u=k}^{s-1} \eta_u \leq \sum_{u \in B} \eta_u < 1,$$

hence

$$\exp\left(-\sum_{u=1}^{s-1} \eta_u \tilde{\ell}_u[k]\right) \geq e^{-1}.$$

Since each term in the softmax denominator is at most 1, we find

$$\hat{p}_s[k] = \frac{\exp(-\sum_{u=1}^{s-1} \eta_u \tilde{\ell}_u[k])}{\sum_{j=1}^K \exp(-\sum_{u=1}^{s-1} \eta_u \tilde{\ell}_u[j])} \geq \frac{e^{-1}}{K}, \quad (s, k) \in \mathcal{R}.$$

Therefore,

$$S(K) = \sum_{s=1}^K \sum_{k=1}^K \hat{p}_s[k] \tilde{\ell}_s[k] \geq \sum_{(s,k) \in \mathcal{R}} \hat{p}_s[k] \geq \frac{e^{-1}}{K} \cdot \frac{L(L+1)}{2} \geq \frac{e^{-1}}{2 \log^2 K} K^{3-2\beta}.$$

For any $1 < \beta \leq \frac{3}{2}$, we conclude

$$S(K) \geq \frac{K^{(3-2\beta)}}{2e^{\frac{\log^2 K}{2}}} \geq \Gamma(K),$$

where the last equality holds because at $t = K$ the minimal noisy loss is $\min_k \sum_{s=1}^K \tilde{\ell}_s[k] = 1$, so

$$\Gamma(K) = \frac{(\max\{1, \frac{K^{(3-2\beta)/\beta}}{\log^{2/\beta} K}\})^\beta}{2e} = \frac{K^{(3-2\beta)}}{2e}.$$

This contradicts the no-early-stop assumption. Hence (7) holds.

From (7) and the definition of the special indices, we have, for each $i = 1, \dots, K^{\beta-1} - 1$,

$$\sum_{s=1}^K \eta_s \tilde{\ell}_s[iK^{2-\beta}] = \sum_{s=iK^{2-\beta}}^K \eta_s \geq \sum_{s=(i+1)K^{2-\beta}}^K \eta_s + \log K = \sum_{s=1}^K \eta_s \tilde{\ell}_s[(i+1)K^{2-\beta}] + \log K, \quad (8)$$

i.e., the weights (of EWA) of successive *special* indices differ by at least $\log K$ in the exponent.

We now define the *true* losses by placing a block of ones of length K on each special index (immediately before its noisy-loss block of ones), in descending order of i :

$$\ell_t[k] = \begin{cases} 1, & k = iK^{2-\beta}, t \in \{K + (K^{\beta-1} - i)K + 1, \dots, 2K + (K^{\beta-1} - i)K\}, \\ 0, & k = K^{2-\beta}, \\ \tilde{\ell}_i[k], & \text{otherwise.} \end{cases}$$

Fix $i \in \{1, \dots, K^{\beta-1} - 1\}$ and let $a := iK^{2-\beta}$ be the active index as defined above on the block

$$B_i = \{K + (K^{\beta-1} - i)K + 1, \dots, K + (K^{\beta-1} - i + 1)K\}.$$

Let $\mathcal{A}_i := \{jK^{2-\beta} : j \leq i - 1\}$ denote the lower special indices of a .

By our construction of the noisy and true losses, the only indices that incur zero true loss on B_i are those in \mathcal{A}_i . Moreover, by our noisy-loss schedule, the weights produced by the adaptive EWA algorithm on $\mathcal{A}_i \cup \{a\}$ are fully determined at time $t = K$ and remain fixed throughout the entire block B_i . Therefore, (8) implies that

$$\forall a' \in \mathcal{A}_i, \forall t \in B_i, \quad \exp\left(-\sum_{s=1}^{t-1} \eta_s(\tilde{\ell}_s[a'] - \tilde{\ell}_s[a])\right) \leq \exp(-\log K) \leq \frac{1}{K}. \quad (9)$$

Hence

$$\begin{aligned} \hat{p}_t[\mathcal{A}_i] &= \frac{\sum_{k \in \mathcal{A}_i} \exp(-\sum_{s=1}^{t-1} \eta_s \tilde{\ell}_s[k])}{\sum_{k \in \mathcal{A}_i} \exp(-\sum_{s=1}^{t-1} \eta_s \tilde{\ell}_s[k]) + \sum_{k \notin \mathcal{A}_i} \exp(-\sum_{s=1}^{t-1} \eta_s \tilde{\ell}_s[k])} \\ &= \frac{\sum_{k \in \mathcal{A}_i} \exp(-\sum_{s=1}^{t-1} \eta_s(\tilde{\ell}_s[k] - \tilde{\ell}_s[a]))}{\sum_{k \in \mathcal{A}_i} \exp(-\sum_{s=1}^{t-1} \eta_s(\tilde{\ell}_s[k] - \tilde{\ell}_s[a])) + \sum_{k \notin \mathcal{A}_i} \exp(-\sum_{s=1}^{t-1} \eta_s(\tilde{\ell}_s[k] - \tilde{\ell}_s[a]))} \\ &\leq \frac{\sum_{k \in \mathcal{A}_i} \exp(-\sum_{s=1}^{t-1} \eta_s(\tilde{\ell}_s[k] - \tilde{\ell}_s[a]))}{\sum_{k \in \mathcal{A}_i} \exp(-\sum_{s=1}^{t-1} \eta_s(\tilde{\ell}_s[k] - \tilde{\ell}_s[a])) + 1} \\ &\leq \frac{(K-1)/K}{(K-1)/K + 1} \leq \frac{1}{2}, \end{aligned}$$

where the first inequality follows since $\sum_{k \notin \mathcal{A}_i} \exp(-\sum_{s=1}^{t-1} \eta_s(\tilde{\ell}_s[k] - \tilde{\ell}_s[a])) \geq 1$ (taking $k = a$), and the second inequality follows from (9) together with the monotonicity of $x \mapsto x/(x+1)$.

The cumulative risk on the true losses over the block B_i is therefore bounded below by $K/2$. Summing over the $K^{\beta-1}$ such blocks gives $\sum_{t=1}^T \langle \hat{p}_t, \ell_t \rangle \geq \frac{K^\beta}{2}$. Since $C := \min_k \sum_{s=1}^T \tilde{\ell}_s[k] \leq 2K$, we obtain $\sum_{t=1}^T \langle \hat{p}_t, \ell_t \rangle \geq \Omega(C^\beta)$. Therefore, in both cases A and B, we obtain, for any $1 < \beta \leq \frac{3}{2}$ that

$$\sum_{t=1}^T \langle \hat{p}_t, \ell_t \rangle = \Omega\left(\max\left\{\frac{K^{(3-2\beta)/\beta}}{\log^{2/\beta} K}, \left(\min_k \sum_{s=1}^T \tilde{\ell}_s[k]\right)^\beta\right\}\right).$$

This yields the lower bound for some $C \geq \frac{K^{(3-2\beta)/\beta}}{\log^{2/\beta} K}$. In particular, this is strictly larger than $O(C \log K)$ whenever $K = \text{poly}(C)$ or $C = \text{poly}(K)$, which completes the proof. \blacksquare

3.3. Achieving minimax optimal risk for unknown C

In Theorem 4, we showed that the EWA algorithm with any *smoothed* adaptive learning rates cannot achieve $\tilde{O}(C \log K)$ regret when C is *unknown*. In contrast, Theorem 3 yields an $O(C \log K)$ risk if we set $\eta = 1/C$, matching the general lower bound in Lemma 2; however, this choice requires the learner to know C in advance. To address these issues, we introduce a “guess-and-double” strategy in Algorithm 2 to estimate C . The main idea is to monitor the minimal cumulative noisy

loss observed so far, which always underestimates C , and use this estimate to tune the learning rate. Once the observed noisy losses exceed our current estimate, we double the estimate and “restart” the EWA— this *restart* mechanism is the key distinction that bypasses the lower bound in Theorem 4.

In our next main result, we show that Algorithm 2 achieves risk $O(C \log C \log K)$, which nearly matches the general lower bound in Lemma 2.

Theorem 5 *Let $\ell_1, \dots, \ell_T \in \mathbb{R}^K$ be arbitrary true loss vectors and $\tilde{\ell}_1, \dots, \tilde{\ell}_T \in \mathbb{R}^K$ be corresponding noisy loss vectors satisfying conditions (3) and (4). Then Algorithm 2, which does not require prior knowledge of C , guarantees*

$$\sum_{t=1}^T \langle \hat{p}_t, \ell_t \rangle \leq O(C \cdot \log C \cdot \ln K).$$

Since $C \leq T$, this bound matches the lower bound in Lemma 2 up to only a $\log T$ factor.

Proof Since L_m doubles at each epoch and satisfies $L_m \leq C$ by (4), there are at most $\log C$ epochs. Let $I_m \subseteq [T]$ denote the set of time steps belonging to epoch m . During each epoch, the algorithm runs as a *refresh* EWA algorithm with learning rate $1/L_m$ and smallest accumulative noisy loss satisfies $\min_{k \in [K]} \tilde{L}_{t-1}(k) \leq 2L_m$. Invoking Theorem 3, we obtain

$$\sum_{t \in I_m} \langle \hat{p}_t, \ell_t \rangle \leq O\left(\log K \cdot L_m + \frac{2CL_m}{L_m} + C \log K\right) \leq O(\log K \cdot L_m + 2C + C \log K).$$

Summing over all (at most $\log C$) epochs and using fact $\sum_m L_m \leq 2C$, we obtain

$$\sum_{t=1}^T \langle \hat{p}_t, \ell_t \rangle \leq O(C \log K + C \log K \cdot \log C) = O(C \cdot \log C \cdot \log K),$$

as claimed. ■

Algorithm 2 EWA with Doubling for Unknown C

- 1: **Input:** Number of experts K
- 2: **Initialize:** epoch index $m \leftarrow 1$, threshold $L_1 \leftarrow 1$, learning rate $\eta_1 \leftarrow 1/L_1$, weights $w_k^1 \leftarrow 1$, cumulative noisy losses $\tilde{L}_0(k) \leftarrow 0$ for all $k \in [K]$
- 3: **for** $t = 1, 2, \dots, T$ **do**
- 4: **if** $\min_{k \in [K]} \tilde{L}_{t-1}(k) \geq 2L_m$ **then**
- 5: $m \leftarrow m + 1$
- 6: $L_m \leftarrow \min_{k \in [K]} \tilde{L}_{t-1}(k)$, $\eta_m \leftarrow 1/L_m$
- 7: Reset weights: $w_k^t \leftarrow 1$ for all $k \in [K]$
- 8: **end if**
- 9: Compute probability distribution over experts:

$$\hat{p}_t[k] \leftarrow \frac{w_k^t}{\sum_{j=1}^K w_j^t} \quad \text{for } k \in [K]$$

- 10: Sample \hat{k}_t from the distribution \hat{p}_t
- 11: Predict using expert \hat{k}_t
- 12: Observe the noisy loss vector $\tilde{\ell}_t \in [0, 1]^K$
- 13: Update weights for all experts:

$$w_k^{t+1} \leftarrow w_k^t \exp\left(-\eta_m \tilde{\ell}_t[k]\right) \quad \text{for } k \in [K]$$

- 14: Update cumulative noisy losses:

$$\tilde{L}_t(k) \leftarrow \tilde{L}_{t-1}(k) + \tilde{\ell}_t[k] \quad \text{for } k \in [K]$$

- 15: **end for**
-

Note that Algorithm 2 performs an explicit restart at the end of each epoch (Step 7). This design allows us to apply Theorem 3 within each epoch and then sum the resulting bounds. A natural question is whether these restarts are responsible for the extra $\log C$ factor in our final risk bound. In much of the online learning literature, one often tries to avoid restart overhead by *smoothing* the learning-rate tuning over time, most notably via the *self-confident* EWA algorithm of Auer et al. (2002b). However, Theorem 4 shows that such smoothing does not help in our setting: not only does it fail to eliminate the extra factor, it can lead to strictly worse performance, with risk growing superlinearly in C , regardless of the adaptive schedule (even when the schedule is allowed to depend on the noisy losses). In our analysis, the $\log C$ factor instead originates from the $C \log K$ term in the epoch-wise guarantee of Theorem 3. This term arises from the discrepancy bound used there, and it does not appear amenable to a simple replacement by something like $O(L_m \log K)$. We therefore suspect that the $\log C$ dependence may be an artifact of the proof technique, and we leave it as an open problem to determine whether it can be removed.

3.4. Extension to non-realizable true losses

In this section, we address the scenario where the true losses are not realizable, i.e., when condition (3) is violated. To handle this case, we introduce the notion of *regret*. For any learning strategy $\hat{p}_t = \Phi_t(\tilde{\ell}_1, \dots, \tilde{\ell}_{t-1})$, we define its worst-case *regret* as

$$\text{Regret}_{T,K,C} = \min_{\{\Phi_t\}_{t=1}^T} \sup_{\ell^T, \tilde{\ell}^T} \left(\mathbb{E} \left[\sum_{t=1}^T \ell_t[\hat{k}_t] \right] - \inf_{k \in [K]} \sum_{t=1}^T \ell_t[k] \right).$$

Note that our regret formulation is evaluated on the *true* losses, not on the observable noisy losses.

We now state our main result for this section as follows with proof delegated to Appendix C.

Theorem 6 *For any corruption budget C , the EWA algorithm (Algorithm 1) with any learning rate η achieves regret at most*

$$\text{Regret}_{T,K,C} \leq O\left(\frac{\ln K}{\eta} + \eta C(C + L_{\min}^*) + C \ln K\right),$$

where $L_{\min}^* = \min_{k \in [K]} \sum_{t=1}^T \ell_t[k]$. In particular, if C and L_{\min}^* are known, then with $\eta = \sqrt{\frac{\ln K}{C(C + L_{\min}^*)}}$, the regret of the EWA algorithm is at most

$$O\left(\sqrt{\ln K \cdot C(C + L_{\min}^*)} + C \ln K\right).$$

Note that the regret bound in Theorem 6 is non-trivial, as it grows sublinearly with L_{\min}^* . A naive extension of the realizable risk bound $O(C \log K)$ would instead yield an $O((C + L_{\min}^*) \log K)$ regret. Our sublinear guarantee benefits from the tighter analysis of EWA in Theorem 3. However, the bound still requires knowledge of both C and L_{\min}^* in order to tune the learning rate. To handle unknown C and L_{\min}^* , we leverage a similar doubling approach as in Algorithm 2 and prove below that it achieves slightly worse regret compared to Theorem 6. The proof is in Appendix D.

Theorem 7 *For arbitrary loss vectors ℓ_1, \dots, ℓ_T and noisy loss vectors $\tilde{\ell}_1, \dots, \tilde{\ell}_T$ satisfying only (4), Algorithm 2 with learning rates $\eta_m = 1/\sqrt{L_m}$ (instead of $\eta_m = 1/L_m$), without prior knowledge of C or L_{\min}^* , achieves*

$$\sum_{t=1}^T \langle \hat{p}_t, \ell_t \rangle \leq O\left(C \sqrt{L_{\min}^* + C} \log(L_{\min}^* + C) \ln K\right).$$

Acknowledgments

We thank a bunch of people and funding agency.

References

Idan Amir, Idan Attias, Tomer Koren, Yishay Mansour, and Roi Livni. Prediction with corrupted expert advice. *Advances in Neural Information Processing Systems*, 33:14315–14325, 2020.

- P. Auer, N. Cesa-Bianchi, and C. Gentile. Adaptive and self-confident on-line learning algorithms. *Journal of Computer and System Sciences*, 2002a.
- Peter Auer, Nicolo Cesa-Bianchi, and Claudio Gentile. Adaptive and self-confident on-line learning algorithms. *Journal of Computer and System Sciences*, 64(1):48–75, 2002b.
- Shai Ben-David, Dávid Pál, and Shai Shalev-Shwartz. Agnostic online learning. In *COLT*, volume 3, page 1, 2009.
- Alankrita Bhatt and Victoria Kostina. Prediction with noisy expert advice. In *2024 IEEE International Symposium on Information Theory (ISIT)*, pages 3546–3551. IEEE, 2024.
- Alankrita Bhatt and Victoria Kostina. Prediction with expert advice under additive noise. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL <https://openreview.net/forum?id=2rupj5mVpp>.
- Nader H Bshouty, Nadav Eiron, and Eyal Kushilevitz. Pac learning with nasty noise. *Theoretical Computer Science*, 288(2):255–275, 2002.
- N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning and Games*. Cambridge University Press, 2006.
- Pierre Gaillard, Gilles Stoltz, and Tim Van Erven. A second-order bound with excess losses. In *Conference on Learning Theory*, pages 176–196. PMLR, 2014.
- Sham Kakade and Adam T Kalai. From batch to transductive online learning. *Advances in Neural Information Processing Systems*, 18, 2005.
- Jaouad Mourtada and Stéphane Gaïffas. On the optimality of the hedge algorithm in the stochastic regime. *Journal of Machine Learning Research*, 20(83):1–28, 2019.
- Nikolaos Papagiannis, Wojciech Szpankowski, and Changlong Wu. Online learning with nasty experts. In *2025 IEEE International Symposium on Information Theory (ISIT)*, 2025.
- Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. Online learning: Random averages, combinatorial parameters, and learnability. *NIPS*, 2010.
- S. Shalev-Shwartz. *Online Learning: Theory, Algorithms, and Applications*. PhD thesis, The Hebrew University of Jerusalem, 2007.
- Tsachy Weissman, Neri Merhav, and Anelia Somekh-Baruch. Twofold universal prediction schemes for achieving the finite-state predictability of a noisy individual binary sequence. *IEEE Transactions on Information Theory*, 47(5):1849–1866, 2001.
- Changlong Wu, Ananth Grama, and Wojciech Szpankowski. Robust online classification: From estimation to denoising. *arXiv preprint arXiv:2309.01698*, 2023.
- Changlong Wu, Ananth Grama, and Wojciech Szpankowski. Information-theoretic limits of online classification with noisy labels. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=Ke3MSP8Nr6>.

Appendix A. Proof of Lemma 2

Our goal is to construct hard instances $\ell^T, \tilde{\ell}^T$ using a *probabilistic* argument. We partition the time horizon into $\log K$ epochs, each of size C (with the other $T - C \log K$ positions padded with 0 losses). Here, we assume $\log K$ is an integer; otherwise, we use the largest power of 2, K' , such that $\log K' \geq \log K - 1$. For each epoch $j \in [\log K]$, we maintain a *random* expert class $S^j \subset [K]$ with $|S^j| = \frac{K}{2^j}$. To construct this, we first select $S^0 := [K]$ and sample S^{j+1} as a random subset of S^j such that $|S^{j+1}| = |S^j|/2$ for all $j \geq 0$.

Having constructed the sets S^j , we define the true and noisy losses as follows. For any time step t in epoch j :

- $\ell_t[k] = 0$ if $k \in S^j$ and $\ell_t[k] = 1$ otherwise;
- $\tilde{\ell}_t[k] = 0$ if $k \in S^{j-1}$ and $\tilde{\ell}_t[k] = 1$ otherwise.

It is straightforward to verify that both constraints (3) and (4) are satisfied by this construction. This is because any expert accumulates discrepancies between the true and noisy losses only within a *single* epoch, and each epoch has size C .

Now, consider *any* prediction rule. At each epoch $j \leq \log K$, the predictor must choose $\hat{k}_t \in S^{j-1}$, otherwise the true loss is 1. Moreover, since the noisy losses are 0 for all $k \in S^{j-1}$, the predictor gains no information about S^j during epoch j . Therefore, the expected error at *every* step, taken over the randomness of both S^j and the predictor's internal randomness, is at least $\frac{1}{2}$ (since $|S^j| = |S^{j-1}|/2$).

This implies that the expected cumulative risk is at least $\frac{1}{2}C \cdot \log K$, since the predictor incurs an expected error of $1/2$ at *every* step, and by the linearity of expectation. The lemma then follows from the fact that there must *exist* realizations $\ell^T, \tilde{\ell}^T$ that achieve this expected cumulative risk bound.

Appendix B. Empirical Validation

We evaluate several EWA variants against the adversarial construction from Theorem 4 with parameter $\beta = \frac{4}{3}$ (see Figure 1). The number of experts is varied exponentially as $K \in \{16, 32, 64, \dots, 4096\}$, while the time horizon is fixed at $T = 20,000$. Performance is reported with respect to the corruption scale (as defined in the proof of Theorem 4)

$$C = \max \left\{ \frac{K^{(3-2\beta)/\beta}}{\log^{2/\beta} K}, \min_k \sum_{t=1}^T \tilde{\ell}_t[k] \right\}, \quad (C \leq 2K).$$

We compare the following learners:

- **EWA-oracle.** Uses $\eta = \sqrt{\log K}/C$, assuming parameter C is *known*.
- **EWA-doubling.** Employs the doubling schedule (Algorithm 2) with $\eta_t = \sqrt{\log K}/L_m$, where L_m is updated at each doubling and restarts are performed.
- **EWA-selfconf.** Uses the self-confident schedule of Auer et al. (2002b),

$$\eta_t = \sqrt{\log K / (\min_k \sum_{s \leq t} \tilde{\ell}_s[k] + 1)}.$$

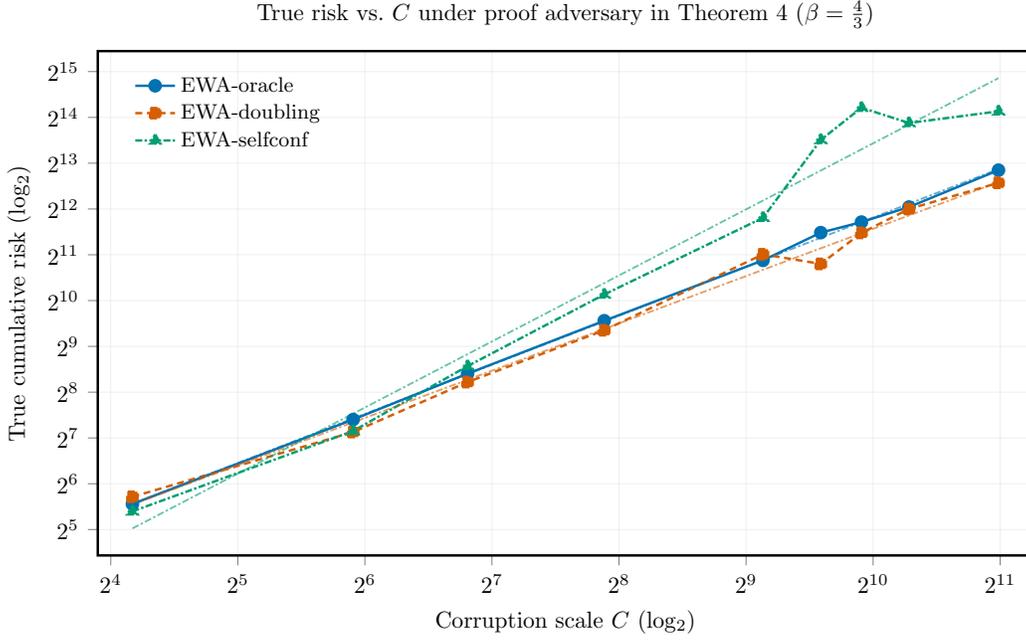


Figure 1: **True risk versus C (log–log scale)** under the adversary of Theorem 4 with $\beta = \frac{4}{3}$.

For a fair comparison, all schedules include the same $\sqrt{\log K}$ factor in the learning rate. This factor does not alter the asymptotic risk rates established in our analysis, but it makes the numerical performance more stable across different values of K .

Finding. Figure 1 shows that the self-confident EWA (Auer et al., 2002b) consistently incurs substantially higher risk than both the doubling-trick EWA (Algorithm 2) and the baseline oracle EWA tuned with knowledge of C . In contrast, the doubling-trick schedule nearly matches the performance of oracle EWA. The fitted slopes— $s_{\text{oracle}} \approx 1.07$, $s_{\text{doubling}} \approx 1.03$, and $s_{\text{selfconf}} \approx 1.44$ —are also consistent with the predictions of our theoretical results.

Appendix C. Proof of Theorem 6

Similar to (6), we decompose the total expected regret as

$$\begin{aligned} \text{regret}_{T,K,C} &= \sup_{\ell^T, \tilde{\ell}^T} \left(\mathbb{E} \left[\sum_{t=1}^T \tilde{\ell}_t[\hat{k}_t] \right] - \inf_{k \in [K]} \sum_{t=1}^T \ell_t[k] + \mathbb{E} \left[\sum_{t=1}^T (\ell_t[\hat{k}_t] - \tilde{\ell}_t[\hat{k}_t]) \right] \right) \\ &\leq \sup_{\ell^T, \tilde{\ell}^T} \left(\mathbb{E} \left[\sum_{t=1}^T \tilde{\ell}_t[\hat{k}_t] \right] - \inf_{k \in [K]} \sum_{t=1}^T \tilde{\ell}_t[k] \right) + \sup_{\ell^T, \tilde{\ell}^T} \left(\mathbb{E} \left[\sum_{t=1}^T (\ell_t[\hat{k}_t] - \tilde{\ell}_t[\hat{k}_t]) \right] \right) \end{aligned}$$

$$\begin{aligned}
 & + \sup_{\ell^T, \tilde{\ell}^T} \left(\inf_{k \in [K]} \sum_{t=1}^T \tilde{\ell}_t[k] - \inf_{k \in [K]} \sum_{t=1}^T \ell_t[k] \right) \\
 & \leq \sup_{\ell^T, \tilde{\ell}^T} \left(\mathbb{E} \left[\sum_{t=1}^T \tilde{\ell}_t[\hat{k}_t] \right] - \inf_{k \in [K]} \sum_{t=1}^T \tilde{\ell}_t[k] \right) + \sup_{\ell^T, \tilde{\ell}^T} \left(\mathbb{E} \left[\sum_{t=1}^T (\ell_t[\hat{k}_t] - \tilde{\ell}_t[\hat{k}_t]) \right] \right) + C \\
 & \stackrel{(a)}{\leq} \frac{\ln K}{\eta} + \eta L_{\min}^* + \sum_{t \in [T]} \frac{\sum_{i \in [K]} w_i^t (\ell_t[i] - \tilde{\ell}_t[i])}{\sum_{i \in [K]} w_i^t} + C \\
 & \leq \frac{\ln K}{\eta} + \eta L_{\min}^* + C \sum_{k \in [K]} \max_{t \in [T]} \frac{w_k^t}{\sum_{i \in [K]} w_i^t} + C,
 \end{aligned}$$

where (a) follows from the regret upper bound for EWA (Lemma 1). On the other hand, we have

$$\sum_{k \in [K-1]} \max_{t \in [T]} \frac{w_k^t}{\sum_{i \in [K]} w_i^t} \leq \ln K + \eta(C + L_{\min}^*)$$

by replacing C in (6) by $C + L_{\min}^*$. Therefore, we have

$$\text{regret}_{T,K,C} \leq \frac{\ln K}{\eta} + \eta L_{\min}^* + \eta C(C + L_{\min}^*) + C \ln K + C$$

$$\text{regret}_{T,K,C} \leq \frac{\ln K}{\eta} + \eta L_{\min}^* + \eta C \tilde{L}_{\min} + C \ln K + C$$

$$\eta_m = \sqrt{\frac{\log K}{C \tilde{L}_m}}$$

and this completes the proof.

Appendix D. Proof Of Theorem 7

Proof The proof is close to that of Theorem 5. Note that L_m doubles at every epoch and $L_m \leq L_{\min}^* + C$. Hence, there are at most $\log(L_{\min}^* + C)$ epochs. At each epoch m the accumulative noisy loss $\min_{k \in [K]} \tilde{L}_{t-1}(k) \leq 2L_m$. Then applying Theorem 3 with $\eta_m = \frac{1}{\sqrt{L_m}}$, we have

$$\sum_{t \in I_m} \langle \hat{p}_t, \ell_t \rangle \leq O \left(\log K \cdot \sqrt{L_m} + \frac{2CL_m}{\sqrt{L_m}} + C \log K \right) \leq O \left(\log K \cdot \sqrt{L_m} + 2C \sqrt{L_m} + C \log K \right).$$

Summing over at most $\log(L_{\min}^* + C)$ epochs m , we have

$$\begin{aligned}
 \sum_{t \in I_m} \langle \hat{p}_t, \ell_t \rangle & \leq O \left(\log K \cdot \sqrt{L_m} + \frac{2CL_m}{\sqrt{L_m}} + C \log K \right) \\
 & \leq O \left(\log K \log(L_{\min}^* + C) \sqrt{L_{\min}^*} + 2C \log(L_{\min}^* + C) \sqrt{L_{\min}^*} + C \log K \log(L_{\min}^* + C) \right)
 \end{aligned}$$

as needed. ■