
When Lower-Order Terms Dominate: Adaptive Expert Algorithms for Heavy-Tailed Losses

Antoine Moulin*

Universitat Pompeu Fabra

antoine.moulin@upf.edu

Emmanuel Esposito*

Università degli Studi di Milano

emmanuel@emmanueleesposito.it

Dirk van der Hoeven

Leiden University

dirk@dirkvanderhoeven.com

Abstract

We consider the problem setting of prediction with expert advice with possibly heavy-tailed losses, i.e. the only assumption on the losses is an upper bound on their second moments, denoted by θ . We develop adaptive algorithms that do not require any prior knowledge about the range or the second moment of the losses. Existing adaptive algorithms have what is typically considered a lower-order term in their regret guarantees. We show that this lower-order term, which is often the maximum of the losses, can actually dominate the regret bound in our setting. Specifically, we show that even with small constant θ , this lower-order term can scale as \sqrt{KT} , where K is the number of experts and T is the time horizon. We propose adaptive algorithms with improved regret bounds that avoid the dependence on such a lower-order term and guarantee $\mathcal{O}(\sqrt{\theta T \log(K)})$ regret in the worst case, and $\mathcal{O}(\theta \log(KT)/\Delta_{\min})$ regret when the losses are sampled i.i.d. from some fixed distribution, where Δ_{\min} is the difference between the mean losses of the second best expert and the best expert. Additionally, when the loss function is the squared loss, our algorithm also guarantees improved regret bounds over prior results.

1 Introduction

We study the problem of prediction with expert advice [Vovk, 1990, Littlestone and Warmuth, 1994], a sequential decision-making setting over T rounds where each round t goes as follows. The learner selects a probability distribution $p_t \in \mathcal{P} = \{p \in \mathbb{R}^K : p(i) \geq 0, \langle p_t, \mathbb{1} \rangle = 1\}$ over K experts, suffers some loss $f_t(p_t) \in \mathbb{R}$, and subsequently observes the loss function $f_t : \mathcal{P} \rightarrow \mathbb{R}$. The goal in this problem, also known more briefly as the experts setting, is to control the (pseudo-)regret

$$R_T = \max_{i \in [K]} R_T(e_i) = \max_{i \in [K]} \mathbb{E} \left[\sum_{t=1}^T (f_t(p_t) - f_t(e_i)) \right],$$

where the expectation is with respect to the randomness in $(f_t)_{t \in [T]}$, and $e_i \in \mathcal{P}$ is the i -th standard basis vector. The pseudo-regret measures the expected difference between the learner's cumulative loss and that of the best expert in hindsight. Throughout the paper, we assume that either the losses or the *outcomes* are heavy tailed, and provide a concrete answer to the concluding remarks of Mhammedi et al. [2019] that highlight the challenge of dealing with infrequent large values in the full-information setting. Such conditions arise in various practical settings involving noisy data,

*Equal contribution.

outliers, or mechanisms like those in local differential privacy [Van der Hoeven, 2019]. Further motivations for our setting can be found in financial markets [Bradley and Taqqu, 2003], electricity forecasting [Li and Jones, 2019, Devaine et al., 2013], and predicting views on articles or videos (see, e.g., Cha et al. [2007]).

A primary goal of this paper is to develop algorithms that achieve sub-linear regret without prior knowledge of a bound on the second moment of the losses, outcomes, or range of the losses $\max_{t \in [T], p \in \mathcal{P}} |f_t(p)|$. We will refer to such algorithms as *loss-range adaptive algorithms*. Crucially, unlike the related multi-armed bandit problem, where adaptivity to an unknown second moment is generally impossible without further assumptions [Genalti et al., 2024], we demonstrate this is achievable in the experts setting. Our second goal is to aim for algorithms that exhibit strong guarantees both in worst-case scenarios and in more benign stochastic environments, often referred to as *best-of-both-worlds* guarantees.

While several loss-range adaptive algorithms exist for linear losses, $f_t(p) = \langle p, \ell_t \rangle = \sum_{i=1}^K p(i) \ell_t(i)$ for some $\ell_t \in \mathbb{R}^K$, they are not designed to handle heavy-tailed losses [Blackwell, 1956, Cesa-Bianchi et al., 2007, De Rooij et al., 2014, Orabona and Pál, 2015, Mhammedi et al., 2019]. Standard algorithms achieve $\mathcal{O}(M\sqrt{T \log K})$ regret for losses bounded as $\max_{t,i} |\ell_t(i)| \leq M$. However, if the losses are drawn from some distribution supported on \mathbb{R} and such that $\max_{t,i} \mathbb{E}[\ell_t(i)^2] \leq \theta$ for some $\theta > 0$, the regret guarantees of existing algorithms can degrade to $\mathcal{O}(\sqrt{KT})$, which is exponentially worse in K . The issue lies in what prior work consider a lower-order term, namely $\max_{t,i} |\ell_t(i)|$. Indeed, while this term is innocuous for bounded losses, we demonstrate that for heavy-tailed losses with second moment $\theta = \mathcal{O}(1)$, it can be as large as $\Omega(\sqrt{KT})$, thereby becoming the dominant factor in the regret bound.

This paper introduces adaptive algorithms that overcome this limitation. Specifically:

1. For linear losses, under Assumption 2.1, Algorithm 2 achieves $\mathcal{O}(\sqrt{\theta T \log(K)})$ worst-case regret, effectively removing the detrimental \sqrt{K} dependency.
2. Under Assumption 2.1, Algorithm 1 achieves $\mathcal{O}(\sqrt{\theta T \log(KT)})$ worst-case regret, while also providing a $\mathcal{O}(\theta \log(KT)/\Delta_{\min})$ regret in self-bounded environments (Assumption 2.3), where Δ_{\min} is the gap in expected loss between the two best experts. This contrasts with prior algorithms whose guarantees would still be hampered by the $\mathcal{O}(\sqrt{KT})$ term.
3. For the squared loss, $f_t(p) = (\langle p, \mathbf{z}_t \rangle - y_t)^2$, where expert predictions $\mathbf{z}_{t,i}$ are bounded by Y and the second moment of the outcomes y_t is bounded by σ , Algorithm 3 achieves a regret of $\mathcal{O}((Y^2 + \sigma^2) \log(K))$, avoiding a $\mathcal{O}(T)$ regret that can arise for heavy-tailed outcomes.

The remainder of the paper is structured as follows: Section 2 presents the problem settings and our main results. Section 3 discusses related work in more detail. In Section 4, we establish why terms previously considered “lower-order” become dominant under heavy-tailed losses for existing algorithms. Section 5 presents our algorithms and sketches their regret analyses, highlighting the techniques used to circumvent the challenges posed by the lack of control over the range of losses. We conclude with a discussion of future work in Section 6.

2 Preliminaries and Results

We consider two settings. The first, sometimes referred to as the “Hedge setting” [Littlestone and Warmuth, 1994], can be seen as a special case of the experts setting [Vovk, 1990] with linear losses. In the second, we consider quadratic losses. Throughout the paper, we assume $K, T \geq 2$ without loss of generality.

2.1 Hedge Setting

In this section, each round $t = 1, \dots, T$, goes as follows: the learner issues $p_t \in \mathcal{P}$, suffers loss $\langle p_t, \ell_t \rangle$, and then observes $\ell_t \in \mathbb{R}^K$. Let $\mathbb{E}_t[\cdot] = \mathbb{E}[\cdot | \mathcal{F}_t]$, where \mathcal{F}_t is the sigma-field generated by the history $(\ell_1, \dots, \ell_{t-1})$ up to round t . We will use the following assumption in the Hedge setting.

Assumption 2.1 (finite second moments). There exists a scalar $\theta \in \mathbb{R}$ such that $\mathbb{E}_t[\ell_t(i)^2] \leq \theta$ for any $t \in [T]$ and any $i \in [K]$.

This assumption is more general than the standard bounded losses assumption: if $|\ell_t(i)| \leq M$ for all $i \in [K]$, $t \in [T]$, then Assumption 2.1 is satisfied with $\theta = M^2$. We *do not assume* to know θ , which is the central algorithmic challenge in the Hedge setting. One of our main results is the following.

Theorem 2.2. *Consider the Hedge setting and suppose Assumption 2.1 holds. Then, there exists an algorithm that, without prior information on the losses, guarantees that $R_T = \mathcal{O}(\sqrt{\theta T \log(K)})$.*

This is achieved by Algorithm 2, which we discuss later in Section 5.2. The proof can be found in Appendix C. While such a result seems expected, we found that loss-range adaptive algorithms in the literature can only guarantee $\mathcal{O}(\sqrt{\theta T \log(K)} + \sqrt{KT})$ regret. We detail prior results in Section 3 and their issues in Section 4.

The following assumption allows us to obtain better regret bounds.

Assumption 2.3. There exist $\Delta_{\min}, C > 0$ and a unique $i^* = \arg \min_{i \in [K]} \mathbb{E} \left[\sum_{t=1}^T \ell_t(i) \right]$ such that $R_T \geq \mathbb{E} \left[\sum_{t=1}^T (1 - p_t(i^*)) \Delta_{\min} \right] - C$.

This is known as a self-bounded environment [Zimmert and Seldin, 2021] and has been studied in the Hedge setting by Amir et al. [2020]. An example setting in which Assumption 2.3 is satisfied is when the losses of the experts are sampled identically and independently from a fixed distribution, in which case $C = 0$. To see why, denote by $\Delta_i = \mathbb{E}[\ell_t(i)] - \mathbb{E}[\ell_t(i^*)]$ and by $\Delta_{\min} = \min_{i \in [K] \setminus i^*} \Delta_i$. We have that

$$R_T = \mathbb{E} \left[\sum_{t=1}^T \sum_{i \neq i^*} p_t(i) \Delta_i \right] \geq \mathbb{E} \left[\sum_{t=1}^T \sum_{i \neq i^*} p_t(i) \Delta_{\min} \right] = \mathbb{E} \left[\sum_{t=1}^T (1 - p_t(i^*)) \Delta_{\min} \right]$$

and so Assumption 2.3 is satisfied. We have the following result under Assumption 2.3.

Theorem 2.4. *Consider the Hedge setting and suppose Assumption 2.1 holds. Then, there exists an algorithm that, without prior information on the losses, guarantees $R_T = \mathcal{O}(\sqrt{\theta T \log(KT)})$. Furthermore, if Assumption 2.3 also holds, then the same algorithm simultaneously guarantees $R_T = \mathcal{O}\left(\frac{\theta \log(KT)}{\Delta_{\min}} + \sqrt{\frac{C \theta \log(KT)}{\Delta_{\min}}}\right)$.*

This is achieved by Algorithm 1, which we discuss in Section 5.1 together with a sketch of the proof. We provide a full proof in Appendix B.

2.2 Quadratic Losses

Here we introduce a second setting we consider. In each round $t \in [T]$, the learner observes the expert predictions $(\mathbf{z}_{t,i})_{i \in [K]}$, issues a prediction $\bar{\mathbf{z}}_t = \langle p_t, \mathbf{z}_t \rangle$, suffers the quadratic loss $f_t(p_t) = (\bar{\mathbf{z}}_t - y_t)^2$, and then observes y_t . We have the following result.

Theorem 2.5. *Consider quadratic losses and suppose that $\max_{t,i} |\mathbf{z}_{t,i}| \leq Y$ and $\max_t \mathbb{E}_t [y_t^2] \leq \sigma$. Then there exists an algorithm such that $R_T = \mathcal{O}((Y^2 + \sigma) \log(KT))$.*

We prove this result in Appendix D. A sketch of the proof is provided in Section 5.3. Our analysis could also be extended to strongly convex losses. We leave this extension for future work.

Notation. For any $t \in [T]$, $i \in [K]$, let $r_t(i) = \langle p_t, \ell_t \rangle - \ell_t(i)$ be the instantaneous regret, $v_t(i) = r_t(i)^2$, and $\bar{v}_t = \sum_{i=1}^K p_t(i) v_t(i)$ be the variance of $\ell_t(I_t)$ with $I_t \sim p_t$. Furthermore, let $\bar{V}_T = \sum_{t=1}^T \bar{v}_t$, $V_T(i) = \sum_{t=1}^T v_t(i)$, $M_T = \mathbb{E} [\max_{t \in [T], i \in [K]} |r_t(i)|]$, and $L_T = \mathbb{E} [\max_{t \in [T], i \in [K]} |\ell_t(i)|]$. Finally, we denote by $i^* \in \arg \min_{i \in [K]} \mathbb{E} \left[\sum_{t=1}^T \ell_t(i) \right]$ the best expert in hindsight.

3 Related Work

There are several loss-range adaptive algorithms in the Hedge setting. A summary can be found in Table 1. Perhaps the most well-known loss-range adaptive algorithm is the exponential weights

	Assumption 2.1	Assumptions 2.1 and 2.3 with $C = 0$
Cesa-Bianchi et al. [2007]; De Rooij et al. [2014]	$\sqrt{\theta T \log(K)} + M_T$	$\theta \log(K) \Delta_{\min}^{-1} + M_T$
Mhammedi et al. [2019]	$\sqrt{\theta T \log(K)} + M_T$	$\theta \log(K) \Delta_{\min}^{-1} + M_T$
Algorithm 1 (Ours)	$\sqrt{\theta T \log(KT)}$	$\theta \log(KT) \Delta_{\min}^{-1}$
Algorithm 2 (Ours)	$\sqrt{\theta T \log(K)}$	$\sqrt{\theta T \log(K)}$

Table 1: An overview of the most relevant loss-range adaptive algorithms in the literature, ignoring constants. Some of the results in the middle column follow from an application of Jensen’s inequality.

algorithm [Vovk, 1990, Littlestone and Warmuth, 1994, Cesa-Bianchi et al., 1997] combined with the doubling trick [Auer et al., 1995]. This combination leads to a $\mathcal{O}\left(L_T \sqrt{T \log(K)}\right)$ regret bound. Cesa-Bianchi et al. [2007] provide a refined version of the exponential weights algorithm combined with a refined doubling trick to obtain a $\mathcal{O}\left(\mathbb{E}\left[\sqrt{V_T \log(K)}\right] + M_T \log(K)\right)$ regret bound. De Rooij et al. [2014] also prove the same regret bound for an algorithm based on the exponential weights algorithm, with the added benefit that their algorithm does not use restarts to achieve this result. Orabona and Pál [2015] provide a generic analysis of follow the regularized leader and online mirror descent that, under mild assumptions on the regularizer, leads to loss-range adaptive algorithms which guarantee that the regret is at most $\mathcal{O}\left(\mathbb{E}\left[\sqrt{\sum_t \max_i \ell_t(i)^2}\right]\right)$, but this can be trivially seen to be at least M_T . Mhammedi et al. [2019] provide a loss-range adaptive version of Squint [Koolen and Van Erven, 2015], which guarantees a regret bound scaling as $\mathcal{O}\left(\mathbb{E}\left[\sqrt{V_T(i^*) \log(K)}\right] + M_T \log(K)\right)$. Flaspohler et al. [2021] show that regret matching [Blackwell, 1956, Hart and Mas-Colell, 2000] and regret matching+ [Tammelin et al., 2015] are both loss-range adaptive algorithm, but unfortunately with unsatisfactory regret bounds. Under Assumption 2.1, Wintenberger [2024] provides a loss-range adaptive algorithm with $\mathcal{O}\left(\log(K) \sqrt{\theta T} + \mathbb{E}\left[\sum_{i=1}^K \log\left(1 + \frac{\max_t |r_t(i)|}{m_{t,i}}\right)\right]\right)$ regret, where $m_{t,i}$ is the first non-null observation of $|r_t(i)|$. Unfortunately, it is not clear how to control $\frac{\max_t |r_t(i)|}{m_{t,i}}$ as $m_{t,i}$ can be made arbitrarily close to 0 by an adversary.

Orseau and Hutter [2024], Cutkosky [2019] provide generic templates to make most Hedge algorithms loss-range adaptive, at the cost of a M_T term in the regret bound. Orseau and Hutter [2024] make use of the following observation: if the algorithm makes too extreme of an update it might never recover. Therefore, it is sometimes better to ignore certain losses. Cutkosky [2019] observed that one can always guarantee that the algorithm does not make such an extreme update by feeding the algorithm slightly shrunken losses, a technique also employed by Mhammedi et al. [2019], Mhammedi [2022]. Specifically, they feed the algorithm the losses

$$\tilde{\ell}_t(i) = \ell_t(i) \min\left\{1, \frac{\max_{s \in [t-1], j \in [K]} |\ell_s(j)|}{\max_{j \in [K]} |\ell_t(j)|}\right\} = \ell_t(i) \cdot \frac{\max_{s \in [t-1], j \in [K]} |\ell_s(j)|}{\max_{s \in [t], j \in [K]} |\ell_s(j)|}.$$

This ensures the range of the losses we feed to the algorithm is known before choosing p_t . The cost for using $\tilde{\ell}_t$ rather than ℓ_t in the algorithm is minor at first sight:

$$\begin{aligned} \sum_{t=1}^T \sum_{i=1}^K p_t(i) (\ell_t(i) - \tilde{\ell}_t(i)) &= \sum_{t=1}^T \sum_{i=1}^K p_t(i) \ell_t(i) \left(1 - \frac{\max_{s \in [t-1], j \in [K]} |\ell_s(j)|}{\max_{s \in [t], j \in [K]} |\ell_s(j)|}\right) \\ &\leq \sum_{t=1}^T \max_{s' \in [t], i \in [K]} |\ell_{s'}(i)| \left(1 - \frac{\max_{s \in [t-1], j \in [K]} |\ell_s(j)|}{\max_{s \in [t], j \in [K]} |\ell_s(j)|}\right) = \max_{t \in [T], i \in [K]} |\ell_t(i)|. \end{aligned}$$

However, this term can be prohibitively large as we will show in Section 4. Instead, we adapt and combine the ideas of Cutkosky [2019], Orseau and Hutter [2024]. We develop a coordinate-wise version of the clipping technique of Cutkosky [2019]. Unfortunately, this is not sufficient for our needs and we need to combine it with a coordinate-wise version of the null-updates of Orseau and Hutter [2024] and the multi-scale entropic regularizer of Bubeck et al. [2019] to guarantee a satisfactory regret bound. Gökcesu and Kozat [2022] also claim to provide loss-range adaptive algorithms, but two of their results seem to contain mistakes. We provide details in Appendix E.

Scale-free algorithms. A related but different objective is obtaining scale-free or equivalently scale-invariant algorithms. An algorithm is said to be scale-free if the predictions of the algorithm do not change if the sequence of losses is multiplied by a positive constant. While scale-free algorithms are loss-range adaptive the converse is not necessarily true. [Mhammedi \[2022\]](#) provide a generic wrapper to make any algorithm scale-free, under some mild assumptions on the algorithm. However, this comes at the cost of an additive M_T . The algorithms of [De Rooij et al. \[2014\]](#), [Orabona and Pál \[2015\]](#) are known to be scale-free. Unfortunately, it is not clear whether our algorithms are also scale-free.

Adaptive algorithms for bounded losses. If we assume that losses are bounded, *e.g.*, $\ell_t(i) \in [0, 1]$, then there are several works that provide best-of-both-worlds results. [Gaillard et al. \[2014\]](#), [Koolen et al. \[2016\]](#) show that so-called second-order bounds simultaneously guarantee $\mathcal{O}(\sqrt{T \log(K)})$ regret without further assumptions on the loss and $\mathcal{O}(\log(K)/\Delta_{\min})$ regret under Assumption 2.3 with $C = 0$. This also implies that the results of [Cesa-Bianchi et al. \[2007\]](#), [De Rooij et al. \[2014\]](#), [Koolen and Van Erven \[2015\]](#), [Chen et al. \[2021\]](#) all lead to small regret under Assumption 2.3 with $C = 0$ while also being robust to more difficult environments. [Mourtada and Gaïffas \[2019\]](#) show that the exponential weights algorithm with a decreasing learning rate guarantees $\mathcal{O}(\log(K)/\Delta_{\min})$ regret under Assumption 2.3 with $C = 1$ while simultaneously guaranteeing $\mathcal{O}(\sqrt{T \log(K)})$ regret in the worst case. The only work that we are aware of that treats Assumption 2.3 with $C > 0$ is by [Amir et al. \[2020\]](#), who show that the exponential weights algorithm with a decreasing learning rate guarantees $\mathcal{O}(\log(K)/\Delta_{\min} + C)$ regret.

Online learning with heavy-tailed losses. To the best of our knowledge, we are the first to consider heavy tailed losses in the expert setting. A related setting is studied by [Bubeck et al. \[2013\]](#), who introduce the multi-armed bandits with heavy tails setting. The main difference with our setting is that the learner only gets to see the loss for the chosen action each round, which significantly reduces the amount of available feedback to the learner. Remarkably, without further assumptions, one cannot adapt to θ , as shown by [Genaldi et al. \[2024\]](#). However, under some additional assumptions on the loss, it is possible to adapt to θ with bandit feedback [[Lee et al., 2020](#), [Ashutosh et al., 2021](#), [Huang et al., 2022](#), [Genaldi et al., 2024](#), [Chen et al., 2025](#)].

4 Lower-Order Terms with Unbounded Losses

In this section, we show that what have been considered lower-order terms in the literature may actually dominate regret bounds.

4.1 Lower-Order Terms in the Hedge Setting

We provide three results that imply that if a regret bound contains M_T or L_T , then this “lower-order term” can dominate the regret bound. We rely on the following observation, proved in Appendix A.

Lemma 4.1. *Fix any $n \in \mathbb{N} \setminus \{0\}$. Let X_1, \dots, X_n be non-negative i.i.d. random variables such that*

$$X_j = \begin{cases} 0 & \text{w.p. } 1 - \frac{1}{n} \\ \sqrt{n} & \text{w.p. } \frac{1}{n} \end{cases}.$$

for each $j \in [n]$. Then, $\mathbb{E}[X_j^2] = 1$ for any $j \in [n]$ and $\frac{1}{2}\sqrt{n} \leq \mathbb{E}[\max\{X_1, \dots, X_n\}] \leq \sqrt{n}$.

While the distribution in Lemma 4.1 might appear unnatural, the Fréchet distribution, which is often used in economics, satisfies similar properties as the distribution of Lemma 4.1; see Lemma A.1 in Appendix A. Lemma 4.1 immediately leads to the following result.

Proposition 4.2. *There exists a distribution for ℓ_1, \dots, ℓ_T that satisfies Assumption 2.1 with $\theta = 4$ such that $L_T \geq \frac{1}{2}\sqrt{KT}$.*

Proof. Simply choose $\ell_t(i) = \mu_{t,i} + \xi_{t,i}\varepsilon_{t,i}$ where $\mu_{t,i} \in [0, 1]$ is chosen arbitrarily, $\xi_{t,i}$ is a Rademacher random variable, and $\varepsilon_{t,i}$ follow the distribution specified in Lemma 4.1 with $n = KT$. Then, Lemma 4.1 provides the result after using $\mathbb{E}_t[\ell_t(i)^2] \leq 2\mu_{t,i}^2 + 2\mathbb{E}_t[\varepsilon_{t,i}^2] \leq 4$. \square

In the worst case, Proposition 4.2 implies that any algorithm with a regret bound of the form $R_T \leq \sqrt{\theta T \log(K)} + L_T$ will be dominated by L_T for large enough K . Existing algorithms sometimes have a $M_T \leq 2L_T$ term in the regret instead. It is natural to question whether M_T can be small enough. The following proposition shows that M_T , like L_T , can be prohibitively large.

Proposition 4.3. *Suppose one can guarantee a bound $R_T \leq B_T^* + M_T$, with $B_T^* \geq 0$. There exists a distribution with $\theta = 3$ such that $B_T^* + M_T \geq B_T^* + \frac{1}{16}\sqrt{KT} - 1$.*

The proof of Proposition 4.3 can be found in Appendix A. While Proposition 4.3 shows that any regret bound with a term M_T is $\mathcal{O}(\sqrt{KT})$, Theorem 2.2 guarantees that Algorithm 2 achieves $R_T = \mathcal{O}(\sqrt{\theta T \log(K)})$. For loss sequences that satisfy Assumption 2.3, the situation is even more dire, as one then hopes to guarantee $\mathcal{O}(\theta \log(K)/\Delta_{\min})$ regret. Proposition 4.3 shows that if one has a regret bound of the form $\theta \log(K)/\Delta_{\min} + M_T$, in the worst case this bound is $\mathcal{O}(\sqrt{KT})$. In contrast, in Theorem 2.4 we provide a $\mathcal{O}(\theta \log(KT)/\Delta_{\min})$ regret bound, which is exponentially better in K and T . Since these results only affect upper bounds on the regret, one can ask whether existing algorithms could benefit from a more careful analysis. We do not think this is possible for existing algorithms and sketch an argument in Appendix A. Furthermore, to illustrate their failure modes, we compare empirically those algorithms with the ones we introduce in Section 5 in Appendix F.

4.2 Lower-Order Terms for the Squared Loss

Since the squared loss is exp-concave (or mixable), an appealing approach is to use exponential weights with a suitable learning rate to obtain a regret bound that is seemingly independent of T . Indeed, a simple calculation shows that the function $g_t(\cdot) = (y_t - \cdot)^2$ is $(2 \max_{t,i} (\mathbf{z}_{t,i} - y_t)^2)^{-1}$ -exp-concave. If one knows $\max_{t,i} (\mathbf{z}_{t,i} - y_t)^2$, this would lead to the following bound [Vovk, 1990]:

$$R_T \leq 2 \log(K) \cdot \mathbb{E} \left[\max_{t,i} (\mathbf{z}_{t,i} - y_t)^2 \right] \quad (1)$$

However, the $\max_{t,i} (\mathbf{z}_{t,i} - y_t)^2$ factor might lead to a poor guarantee. Suppose that $|\mathbf{z}_{t,i}| \leq Y$ and that $\mathbb{E}[y_t^2] \leq \theta < \infty$. If $y_t = \mu_t + \xi_t \varepsilon_t$, where $|\mu_t| \leq 1$, ξ_t is a Rademacher random variable, and ε_t follows the distribution specified in Lemma 4.1 with $n = T$, then

$$\begin{aligned} \mathbb{E} \left[\max_{t,i} (\mathbf{z}_{t,i} - y_t)^2 \right] &\geq \mathbb{E} \left[\max_t y_t^2 - 2Y|y_t| \right] \geq \mathbb{E} \left[\max_t |y_t| \right]^2 - 2Y \mathbb{E} \left[\max_t |y_t| \right] \\ &\geq \mathbb{E} \left[\max_t |\varepsilon_t| \right]^2 - 2(Y+1) \left(1 + \mathbb{E} \left[\max_t |\varepsilon_t| \right] \right) \geq \frac{1}{2}T - 2(Y+1)\sqrt{T}, \end{aligned}$$

which can be excessively large. Even more so, one would need to know $\max_{t,i} (\mathbf{z}_{t,i} - y_t)^2$ to obtain the regret bound in Equation (1). The results of Wintenberger [2017] (see also Van der Hoeven et al. [2022]) imply that the algorithm of Mhammedi et al. [2019], designed for the Hedge setting, applied to the losses $\ell_t(i) = 2\mathbf{z}_{t,i}(\bar{\mathbf{z}}_t - y_t)$ would lead to a bound of order

$$R_T = \mathcal{O} \left((Y^2 + \theta + \mathbb{E}[\max_t |\bar{\mathbf{z}}_t - y_t|]) \log(K) \right).$$

The $\mathbb{E}[\max_t |\bar{\mathbf{z}}_t - y_t|]$ term can also be problematic. Indeed, if $y_t = \mu_t + \xi_t \varepsilon_t$, where $\mu_t \in [-1, 1]$, ξ_t is a Rademacher random variable, and ε_t follows the distribution specified in Lemma 4.1 with $n = KT$, then $\mathbb{E}[\max_t |\bar{\mathbf{z}}_t - y_t|] \geq \mathbb{E}[\max_t |\bar{\mathbf{z}}_t|] - Y = \frac{1}{2}\sqrt{T} - Y$. Thus, if $Y \geq 2$ then $Y^2 + \theta + \mathbb{E}[\max_t |\bar{\mathbf{z}}_t - y_t|] \geq \frac{1}{2}(Y^2 + \sqrt{T})$. On the other hand, for the same y_t , Theorem 2.5 implies that Algorithm 3 guarantees a regret bound of order $(4 + Y^2) \log(KT)$.

5 Algorithms

In this section, we present two types of algorithms and provide a sketch for their regret analysis. The first algorithm, called LoOT-Free OMD (Lower-Order Term Free), is an instance of Online Mirror Descent (OMD), whereas the second, called LoOT-Free FTRL, is an instance of Follow The Regularized Leader (FTRL). Both algorithms are run on clipped versions of the instantaneous regrets,

Algorithm 1 LoOT-Free OMD

Inputs: Number of experts $K \geq 2$, truncation parameter $\alpha \in (0, 1]$, learning rate $\beta > 0$.
Initialize: $p_1(i) \leftarrow 1/K$ for all $i \in [K]$.
for $t = 1, \dots, T$ **do**
 Predict p_t , incur loss $\langle p_t, \ell_t \rangle$ and observe ℓ_t .
 Set $r_t(i) \leftarrow \langle p_t, \ell_t \rangle - \ell_t(i)$ for all $i \in [K]$.
 Set $v_t(i) \leftarrow r_t(i)^2$ for all $i \in [K]$, and $\bar{v}_t \leftarrow \sum_{i=1}^K p_t(i)v_t(i)$.
 if $\sum_{s \leq t} \bar{v}_s > 0$ **then**
 Set $\eta_{t,i} \leftarrow \beta \max\{\sum_{s \leq t} \bar{v}_s, \sum_{s \leq t} v_s(i)\}^{-1/2}$ for all $i \in [K]$.
 Set $\tilde{\ell}_t(i) \leftarrow -r_t(i) \mathbb{1}(|r_t(i)| \leq 1/\eta_{t,i})$ for all $i \in [K]$.
 Set $p_{t+1} \leftarrow \arg \min_{p \in \mathcal{P}_\alpha} \langle p, \tilde{\ell}_t \rangle + D_t(p \| p_t)$.
 end if
end for

but one could obtain similar worst-case guarantees by running the algorithms on clipped versions of the losses.

Before we describe the algorithms in more detail, we first need some definitions. For any $t \in [T]$, let $\eta_t \in \mathbb{R}_{>0}^K$ be a vector of learning rates and $\psi_t : x \in \mathbb{R}_{\geq 0}^K \mapsto \langle \text{Diag}(\eta_t^{-1})x, \log x \rangle$ be the regularizer, where y^{-1} is the coordinate-wise inverse, $\log y$ is the coordinate-wise logarithm, and $\text{Diag}(y)$ is the diagonal matrix with diagonal entries y_1, \dots, y_K for any $y \in \mathbb{R}^K$. Note that, for any $x \in \mathbb{R}_{\geq 0}^K$, the gradient of ψ_t is given by $\nabla \psi_t(x) = \text{Diag}(\eta_t^{-1}) \log x + \eta_t^{-1}$. Consider also its Bregman divergence, defined for any $x \in \mathbb{R}_{\geq 0}^K$ and $y \in \mathbb{R}_{>0}^K$ as $D_t(x \| y) = \psi_t(x) - \psi_t(y) - \langle \nabla \psi_t(y), x - y \rangle = \sum_{i=1}^K \frac{1}{\eta_{t,i}} \left[x_i \log \left(\frac{x_i}{y_i} \right) - x_i + y_i \right]$. Finally, denote by $\mathcal{P}_\alpha = \{p \in \mathcal{P} : \min_i p_i \geq \alpha/K\}$ the probability simplex truncated by α/K , for any $\alpha \in [0, 1]$.

5.1 OMD-based Algorithm

The main challenge in designing an algorithm for losses for which you do not know the range comes from proving the stability of the algorithm, which is to say that $p_t(i) \approx p_{t+1}(i)$. An analysis based on strong convexity circumvents this challenge. [Orabona and Pál \[2015\]](#) show that if the regularizer ϕ is strongly convex with respect to some norm $\|\cdot\|$, the regret can be $\mathcal{O}\left(\sqrt{\max_{p \in \mathcal{P}} \phi(p) \sum_{t=1}^T \|\ell_t\|_*^2}\right)$, where $\|\cdot\|_*$ is the dual norm. However, if one uses the (shifted) negative Shannon entropy as ϕ , this will lead to a $\mathcal{O}\left(\sqrt{\log(K) \sum_{t=1}^T \|\ell_t\|_\infty^2}\right)$ bound, which is $\mathcal{O}(\sqrt{KT})$ in the worst case. A more careful analysis that avoids strong convexity arguments can be found in [De Rooij et al. \[2014\]](#), but unfortunately this analysis does not avoid a problematic lower-order term. To avoid such issues, we make use of the multi-scale entropic regularizer of [Bubeck et al. \[2019\]](#). If the range of the losses is known a-priori, [Bubeck et al. \[2019\]](#) show that for bounded losses and known $\max_t |\ell_t(i)|$ the regret against expert i then scales as $\mathbb{E}[\max_t |\ell_t(i)|] \sqrt{T \log(K)}$. However, this alone is not sufficient, as it is not clear how to prove that the algorithm is stable without an a-priori uniform bound on the losses, nor is such a regret bound sufficiently small. Instead, we carefully clip the losses when they are excessively large. Combining these ideas leads to the following guarantee for Algorithm 1 with $\alpha = \frac{1}{T}$ and $\beta = \sqrt{\log(KT)}$:

$$R_T(e_{i^*}) = \sum_{t=1}^T r_t(i^*) = \mathcal{O}\left(\sqrt{\log(KT) (\bar{V}_T + V_T(i^*))}\right). \quad (2)$$

Notice that the above upper bound holds with probability one for any sequence of losses and, in turn, leads to the guarantees of Theorem 2.4. A full proof of the above result can be found in Appendix B. Here we provide some intuition.

Proof sketch of Theorem 2.4. First, observe that Algorithm 1 is an instance of OMD adopting the time-varying regularizer ψ_t with some additional tricks. We only update p_t if $\sum_{s \leq t} \bar{v}_s > 0$. The only

case where this is not true is if in all rounds up to and including round t , $r_t(i) = 0$ for all $i \in [K]$, in which case we can simply ignore these rounds in the analysis. From standard OMD analysis (see for example [Orabona \[2025\]](#)), we know that if we run OMD on losses $\tilde{\ell}_t$, for any fixed $j \in [K]$,

$$\sum_{t=1}^T (\langle p_t, \tilde{\ell}_t \rangle - \tilde{\ell}_t(j)) = \mathcal{O}\left(\frac{\log(K/\alpha)}{\eta_{T,j}} + \sum_{t=1}^T \sum_{i=1}^K \eta_{t,i} \tilde{p}_t(i) \tilde{\ell}_t(i)^2\right),$$

where $\tilde{p}_t(i) = p_t(i) \exp(-\eta_{t,i} \tilde{\ell}_t(i))$. At this point, we face our main challenge. We would like to show that $\tilde{p}_t(i) \approx p_t(i)$. To do so, we would need to show that $|\eta_{t,i} \tilde{\ell}_t(i)| \leq 1$. We force this to be true by simply setting $\tilde{\ell}_t(i) = 0$ if $|r_t(i)| > \frac{1}{\eta_{t,i}}$ and $\tilde{\ell}_t(i) = -r_t(i)$ otherwise. Of course, there is a price to pay for clipping the losses, but this price is negligible:

$$|\tilde{\ell}_t(i) - (-r_t(i))| = |r_t(i)| \mathbb{1}\left(|r_t(i)| > \frac{1}{\eta_{t,i}}\right) = \frac{r_t(i)^2}{|r_t(i)|} \mathbb{1}\left(|r_t(i)| > \frac{1}{\eta_{t,i}}\right) \leq \eta_{t,i} v_t(i). \quad (3)$$

Thus, after noting that ℓ_t and $-r_t$ only differ from a constant, we have that the regret against expert j is $\sum_{t=1}^T (\langle p_t, \ell_t \rangle - \ell_t(j)) = \sum_{t=1}^T (\langle p_t, -r_t \rangle - (-r_t(j)))$, and we can replace $-r_t$ by the truncated losses $\tilde{\ell}_t$ by paying the cost shown in Equation (3) and use the guarantee from OMD

$$\begin{aligned} \sum_{t=1}^T (\langle p_t, \ell_t \rangle - \ell_t(j)) &= \mathcal{O}\left(\sum_{t=1}^T (\langle p_t, \tilde{\ell}_t \rangle - \tilde{\ell}_t(j)) + \sum_{t=1}^T \left(\eta_{t,j} v_t(j) + \sum_{i=1}^K p_t(i) \eta_{t,i} v_t(i)\right)\right) \\ &= \mathcal{O}\left(\frac{\log(K/\alpha)}{\eta_{T,j}} + \sum_{t=1}^T \left(\eta_{t,j} v_t(j) + \sum_{i=1}^K p_t(i) \eta_{t,i} v_t(i)\right)\right). \end{aligned}$$

After controlling the sum by using the definition of $\eta_{t,i}$, the bound in Equation (2) follows as

$$\sum_{t=1}^T \left(\eta_{t,j} v_t(j) + \sum_{i=1}^K p_t(i) \eta_{t,i} v_t(i)\right) \leq \beta \sum_{t=1}^T \left(\frac{v_t(j)}{\sqrt{V_t(j)}} + \frac{\bar{v}_t}{\sqrt{\bar{V}_t}}\right) \leq 4\beta \sqrt{\bar{V}_T + V_T(j)},$$

where the first inequality follows from the definition of $\eta_{t,i}$ and the second follows from, for example, Lemma 4.13 in [Orabona \[2025\]](#). \square

Regarding computational aspects, note the optimization problem defining p_{t+1} is done on a truncated simplex, thus it does not have a closed-form solution but can be computed efficiently with a line-search as done by [Chen et al. \[2021\]](#).

5.2 FTRL-based Algorithm

Taking some ideas from the previous OMD-based algorithm, we derive an instantiation of FTRL described by Algorithm 2. For FTRL we use once more the multi-scale entropic regularizer ψ_t and a similar clipping of the losses to resolve the same issues that follow from the lack of any prior knowledge of the loss range. However, the FTRL framework has some fundamental differences compared to OMD, which in turn require further adjustments. To understand these differences, we will sketch the proof of Theorem 2.2 below, whose result is obtained by Algorithm 2 with $\beta = \sqrt{\log(K)}$. The detailed proof of Theorem 2.2 can be found in Appendix C. Here we provide the main ideas. We actually prove a stronger regret guarantee than provided in Theorem 2.2:

$$R_T(e_{i^*}) = \sum_{t=1}^T r_t(i^*) = \mathcal{O}\left(\sqrt{\log(K)(\bar{V}_T + V_T(i^*))} + \frac{1}{K} \sum_{i=1}^K \sqrt{V_T(i)}\right), \quad (4)$$

which then implies the statement of Theorem 2.2.

To prove this regret bound we make use of the standard FTRL analysis (e.g., see [Orabona \[2025\]](#)) with some additional tricks. Note that, differently from Algorithm 1, we replace the multi-scale entropic regularizer ψ_t with its Bregman divergence. This is to ensure the monotonicity of the regularizer, i.e., $\varphi_{t+1}(x) \geq \varphi_t(x)$ for all $x \in \mathcal{P}$. This monotonicity is necessary for the proof of the FTRL regret bound. The fact that we use the Bregman divergence generated by ψ_t rather than ψ_t directly

Algorithm 2 LoOT-Free FTRL

Inputs: Number of experts $K \geq 2$, learning rate coefficient $\beta > 0$.

Initialize: $p_1(i) \leftarrow 1/K$ and $b_0(i) \leftarrow 0$ for all $i \in [K]$.

for $t = 1, \dots, T$ **do**

 Predict p_t , incur loss $\langle p_t, \ell_t \rangle$ and observe ℓ_t .

 Set $r_t(i) \leftarrow \langle p_t, \ell_t \rangle - \ell_t(i)$ for all $i \in [K]$.

 Set $v_t(i) \leftarrow r_t(i)^2$ for all $i \in [K]$.

 Set $\bar{v}_t \leftarrow \sum_{i=1}^K p_t(i)v_t(i)$.

if $\sum_{s \leq t} \bar{v}_s > 0$ **then**

 Set $b_t(i) \leftarrow \max \left\{ \sum_{s \leq t} \bar{v}_s, \sum_{s \leq t} v_s(i) \right\}^{1/2}$ for all $i \in [K]$.

 Set $\eta_{t,i} \leftarrow \beta/b_t(i)$ for all $i \in [K]$.

 Set $\tilde{\ell}_t(i) \leftarrow -r_t(i) \frac{b_{t-1}(i)}{b_t(i)} \mathbb{1}(|r_t(i)| \leq 1/\eta_{t,i})$ for all $i \in [K]$.

 Set $p_{t+1} \leftarrow \arg \min_{p \in \mathcal{P}} \sum_{s \leq t} \langle p, \tilde{\ell}_s \rangle + D_t(p \| p_1)$.

end if

end for

as the regularizer also allows us to avoid the $\log(T)$ term of the OMD regret bound, at the cost of a $\frac{1}{K} \sum_{i=1}^K \sqrt{\sum_{t=1}^T v_t(i)}$ term in the regret. While this term is prohibitively big when we try to obtain improved regret bounds under Assumption 2.3, in expectation this term scales as $\mathcal{O}(\sqrt{\theta T})$ under Assumption 2.1, thus preserving the desired final bound.

Proof sketch of Theorem 2.2. Since the p_t in Algorithm 2 comes from an instance of FTRL over the losses $\tilde{\ell}_t$, we can apply a standard FTRL bound (for example, Lemma 7.16 of Orabona [2025]) to see that, for any fixed $j \in [K]$,

$$\sum_{t=1}^T (\langle p_t, \tilde{\ell}_t \rangle - \tilde{\ell}_t(j)) = \mathcal{O} \left(\frac{\log K}{\eta_{T,j}} + \frac{1}{K} \sum_{i=1}^K \frac{1}{\eta_{t,i}} + \sum_{t=1}^T \sum_{i=1}^K \eta_{t-1,i} \tilde{p}_t(i) \tilde{\ell}_t(i)^2 \right), \quad (5)$$

where $\tilde{p}_t(i) = p_t(i) \exp(-\eta_{t-1,i} \tilde{\ell}_t(i))$. Here we reach the main challenge of adapting FTRL to our setting. The third term illustrates crucial differences compared to OMD: the learning rate $\eta_{t-1,i}$ instead of $\eta_{t,i}$ appears as a multiplicative factor and in the definition of \tilde{p}_t . We resolve this issue by rescaling the loss $-r_t(i)$ by a factor $b_{t-1}(i)/b_t(i)$ in the definition of $\tilde{\ell}_t(i)$. Still, this rescaled loss can be too big, which is why we also clip the loss when necessary. Specifically, the clipping with a rescaling factor ensures, by the definitions of $b_t(i)$ and $\eta_{t,i}$, that $\eta_{t-1,i} \tilde{\ell}_t(i)^2 \leq \frac{b_{t-1}(i)}{b_t(i)} \eta_{t-1,i} r_t(i)^2 = \eta_{t,i} v_t(i)$ and, similarly, that $|\eta_{t-1,i} \tilde{\ell}_t(i)| \leq 1$. Thus, we can bound the third term on the right hand side of Equation (5) by $\mathcal{O}(\sum_{t=1}^T \sum_{i=1}^K \eta_{t,i} p_t(i) v_t(i))$, which we already know how to control from the analysis of Algorithm 1.

Another consequence of the new definition of $\tilde{\ell}_t$ is that it increases the cost of clipping the losses. In comparison to Equation (3), it now additionally presents terms

$$|r_t(i)| \left(1 - \frac{b_{t-1}(i)}{b_t(i)} \right) \mathbb{1}(|r_t(i)| \leq \frac{1}{\eta_{t,i}}) \leq \frac{1}{\eta_{t,i}} - \frac{1}{\eta_{t-1,i}}$$

but the sums involving them are also nicely behaved. Combining all these observations leads to Equation (4). \square

While the optimization problem suggests p_t is given by a softmax function, the coordinate-dependent learning rate prevents us from computing the normalization constant in closed-form (see Appendix C.1 for details). However, we can still compute it efficiently with a line-search as in Algorithm 1 from Bubeck et al. [2019].

5.3 Algorithm for the Squared Loss

The algorithm we use for the squared loss, Algorithm 3 can be found in Appendix D. It is exactly the same as Algorithm 1, but then run on the losses

$$\ell_t(i) = \mathbf{z}_{t,i}(\bar{\mathbf{z}}_t - y_t) + \frac{1}{2}(\mathbf{z}_{t,i} - y_t)^2,$$

where $\bar{\mathbf{z}}_t = \langle p_t, \mathbf{z} \rangle$. The inspiration from this surrogate loss comes from two inequalities for the squared loss:

$$(a - y)^2 - (b - y)^2 = 2(y - a)(b - a) - (a - b)^2,$$

$$\left(\frac{a+b}{2} - y\right)^2 \leq \frac{(a-y)^2}{2} + \frac{(b-y)^2}{2} - \frac{(a-b)^2}{8}.$$

By carefully applying these inequalities we find that, for any fixed $j \in [K]$,

$$\begin{aligned} & \sum_{t=1}^T ((\bar{\mathbf{z}}_t - y_t)^2 - (\mathbf{z}_{t,j} - y_t)^2) \\ & \leq \sum_{t=1}^T \left(\sum_{i=1}^K p_t(i) \ell_t(i) - \ell_t(j) - \sum_{i=1}^K p_t(i) \frac{(\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^2}{8} - \frac{(\mathbf{z}_{t,j} - \bar{\mathbf{z}}_t)^2}{2} \right). \end{aligned}$$

The negative quadratics on the right-hand-side of the equation above allow us prove the regret bound of Theorem 2.5. We show that, for any fixed $j \in [K]$,

$$\begin{aligned} R_T(e_j) &= \mathcal{O} \left(\mathbb{E} \left[\sqrt{\log(KT)(\bar{V}_T + V_T(j))} \right] - \mathbb{E} \left[\sum_{t=1}^T \left(\sum_{i=1}^K \frac{p_t(i)(\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^2}{8} + \frac{(\mathbf{z}_{t,j} - \bar{\mathbf{z}}_t)^2}{2} \right) \right] \right) \\ &= \mathcal{O}((Y^2 + \sigma) \log(KT)) . \end{aligned}$$

A detailed proof can be found in Appendix D.

6 Future Work

One direction to explore is whether the $\log(T)$ factor in Theorem 2.4 is necessary to obtain best-of-both-worlds guarantees. Potentially, an improved analysis of the FTRL algorithm will do the trick, as we know that OMD can be inferior to FTRL [Amir et al., 2020]. While we understand how to obtain best-of-both-worlds guarantees for FTRL with bounded losses [Mourtada and Gaiffas, 2019, Amir et al., 2020], it is unclear how to adapt these analyses to our setting. Another interesting direction is to see whether the ideas in Section 5.3 can be extended to strongly convex and exp-concave losses. We believe that the former is relatively straightforward, but the latter might be highly challenging. Another relatively straightforward extension could be to adapt to any moment of the loss, *i.e.*, adapt to $\mathbb{E}_t[|\ell_t(i)|^\alpha]$ for some $\alpha > 1$, without the prior knowledge of α or an upper bound for $\mathbb{E}_t[|\ell_t(i)|^\alpha]$.

Acknowledgements

EE acknowledges the financial support from the FAIR (Future Artificial Intelligence Research) project, funded by the NextGenerationEU program within the PNRR-PE-AI scheme (M4C2, investment 1.3, line on Artificial Intelligence), the EU Horizon CL4-2022-HUMAN-02 research and innovation action under grant agreement 101120237, project ELIAS (European Lighthouse of AI for Sustainability), and the One Health Action Hub, University Task Force for the resilience of territorial ecosystems, funded by Università degli Studi di Milano (PSR 2021-GSA-Linea 6). AM has received funding from the European Research Council (ERC), under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 950180). This work was partially done while DvdH was at the University of Amsterdam supported by Netherlands Organization for Scientific Research (NWO), grant number VI.Vidi.192.095.

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.
- Kumar Ashutosh, Jayakrishnan Nair, Anmol Kagrecha, and Krishna Jagannathan. Bandit algorithms: Letting go of logarithmic regret for statistical robustness. In *International Conference on Artificial Intelligence and Statistics*, pages 622–630. PMLR, 2021.
- Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E Schapire. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *Proceedings of IEEE 36th Annual Goundations of Computer Science*, pages 322–331. IEEE, 1995.
- David Blackwell. An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6(1):1–8, 1956.
- Brendan O Bradley and Murad S Taqqu. Financial risk and heavy tails. In *Handbook of heavy tailed distributions in finance*, pages 35–103. Elsevier, 2003.
- Sébastien Bubeck, Nicolò Cesa-Bianchi, and Gábor Lugosi. Bandits with heavy tail. *IEEE Transactions on Information Theory*, 59(11):7711–7717, 2013.
- Sébastien Bubeck, Nikhil R Devanur, Zhiyi Huang, and Rad Niazadeh. Multi-scale online learning: Theory and applications to online auctions and pricing. *Journal of Machine Learning Research*, 20(62):1–37, 2019.
- Nicolò Cesa-Bianchi, Yoav Freund, David Haussler, David P. Helmbold, Robert E. Schapire, and Manfred K. Warmuth. How to use expert advice. *J. ACM*, 44(3), 1997. doi: 10.1145/258128.258179.
- Nicolò Cesa-Bianchi, Yishay Mansour, and Gilles Stoltz. Improved second-order bounds for prediction with expert advice. *Machine Learning*, 66:321–352, 2007.
- Meeyoung Cha, Haewoon Kwak, Pablo Rodriguez, Yong-Yeol Ahn, and Sue Moon. I tube, you tube, everybody tubes: analyzing the world’s largest user generated content video system. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, pages 1–14, 2007.
- Gong Chen and Marc Teboulle. Convergence analysis of a proximal-like minimization algorithm using bregman functions. *SIAM Journal on Optimization*, 3(3):538–543, 1993.
- Liyu Chen, Haipeng Luo, and Chen-Yu Wei. Impossible tuning made possible: A new expert algorithm and its applications. In *Conference on Learning Theory*, pages 1216–1259, 2021.
- Yu Chen, Jiatai Huang, Yan Dai, and Longbo Huang. uniINF: Best-of-both-worlds algorithm for parameter-free heavy-tailed MABs. *arXiv preprint arXiv:2410.03284v2*, 2025.
- Ashok Cutkosky. Artificial constraints and hints for unbounded online learning. In *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99, pages 874–894, 25–28 Jun 2019.
- Marie Devaine, Pierre Gaillard, Yannig Goude, and Gilles Stoltz. Forecasting electricity consumption by aggregating specialized experts: A review of the sequential aggregation of specialized experts, with an application to slovakian and french country-wide one-day-ahead (half-) hourly predictions. *Machine Learning*, 90(2):231–260, 2013.
- Genevieve E. Flaspoehler, Francesco Orabona, Judah Cohen, Soukaina Mouatadid, Miruna Oprescu, Paulo Orenstein, and Lester Mackey. Online learning with optimism and delay. In *International Conference on Machine Learning*, pages 3363–3373. PMLR, 2021.
- Maurice Fréchet. Sur la loi de probabilité de l’écart maximum. *Ann. de la Soc. Polonaise de Math.*, 1927.
- Pierre Gaillard, Gilles Stoltz, and Tim van Erven. A second-order bound with excess losses. In *Conference on Learning Theory*, pages 176–196, 2014.

- Gianmarco Genalti, Lupo Marsigli, Nicola Gatti, and Alberto Maria Metelli. (ε, u) -adaptive regret minimization in heavy-tailed bandits. In *Conference on Learning Theory*, volume 247, pages 1882–1915, 2024.
- Hakan Gökcesu and Suleyman Serdar Kozat. Optimal tracking in prediction with expert advice. *arXiv preprint arXiv:2208.03708v1*, 2022.
- Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.
- Dirk van der Hoeven. User-specified local differential privacy in unconstrained adaptive online learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- Dirk van der Hoeven, Nikita Zhivotovskiy, and Nicolò Cesa-Bianchi. A regret-variance trade-off in online learning. In *Advances in Neural Information Processing Systems*, volume 35, pages 35188–35200, 2022.
- Jiatai Huang, Yan Dai, and Longbo Huang. Adaptive best-of-both-worlds algorithm for heavy-tailed multi-armed bandits. In *International Conference on Machine Learning*, pages 9173–9200. PMLR, 2022.
- Wouter M. Koolen and Tim van Erven. Second-order quantile methods for experts and combinatorial games. In *Conference on Learning Theory*, pages 1155–1175, 2015.
- Wouter M. Koolen, Peter Grünwald, and Tim van Erven. Combining adversarial guarantees and stochastic fast rates in online learning. *Advances in Neural Information Processing Systems*, 29, 2016.
- Kyungjae Lee, Hongjun Yang, Sungbin Lim, and Songhwai Oh. Optimal algorithms for stochastic multi-armed bandits with heavy tailed rewards. *Advances in Neural Information Processing Systems*, 33:8452–8462, 2020.
- Yun Li and Ben Jones. The use of extreme value theory for forecasting long-term substation maximum electricity demand. *IEEE Transactions on Power Systems*, 35(1):128–139, 2019.
- Nick Littlestone and Manfred K Warmuth. The weighted majority algorithm. *Information and Computation*, 108(2):212–261, 1994.
- Zakaria Mhammedi. Efficient projection-free online convex optimization with membership oracle. In *Conference on Learning Theory*, pages 5314–5390. PMLR, 2022.
- Zakaria Mhammedi, Wouter M. Koolen, and Tim van Erven. Lipschitz adaptivity with multiple learning rates in online learning. In *Conference on Learning Theory*, pages 2490–2511, 2019.
- Jaouad Mourtada and Stéphane Gaïffas. On the optimality of the Hedge algorithm in the stochastic regime. *Journal of Machine Learning Research*, 20:1–28, 2019.
- Gangadharan Muraleedharan, Carlos G. Soares, and Cláudia Lucas. Characteristic and moment generating functions of generalised extreme value distribution (gev). *Sea level rise, coastal engineering, shorelines and tides*, pages 269–276, 2011.
- Francesco Orabona. A modern introduction to online learning. *arXiv preprint arXiv:1912.13213v7*, 2025.
- Francesco Orabona and Dávid Pál. Scale-free algorithms for online linear optimization. In *International Conference on Algorithmic Learning Theory*, pages 287–301. Springer, 2015.
- Laurent Orseau and Marcus Hutter. Isotuning with applications to scale-free online learning. *arXiv preprint arXiv:2112.14586v3*, 2024.
- Steven de Rooij, Tim van Erven, Peter D. Grünwald, and Wouter M. Koolen. Follow the Leader if you can, Hedge if you must. *Journal of Machine Learning Research*, 15:1281–1316, 2014.
- Oskari Tammelin, Neil Burch, Michael Johanson, and Michael Bowling. Solving heads-up limit texas hold’em. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.

- Volodimir G. Vovk. Aggregating strategies. In *Proceedings of the 3rd Annual Workshop on Computational Learning Theory*, 1990.
- Olivier Wintenberger. Optimal learning with Bernstein online aggregation. *Machine Learning*, 106(1):119–141, 2017.
- Olivier Wintenberger. Stochastic online convex optimization. application to probabilistic time series forecasting. *Electronic Journal of Statistics*, 18(1):429–464, 2024.
- Julian Zimmert and Yevgeny Seldin. Tsallis-inf: An optimal algorithm for stochastic and adversarial bandits. *Journal of Machine Learning Research*, 22(28):1–49, 2021.

Appendix Contents

A Technical Results for Section 4	14
A.1 Proof of Lemma 4.1	14
A.2 Remarks on Fréchet Distributions	14
A.3 Proof of Proposition 4.3	15
A.4 Discussion on analyses of existing algorithms	16
B Regret Analysis for Online Mirror Descent	17
B.1 Adversarial Environments	17
B.2 Self-Bounded Environments	23
C Regret Analysis for Follow The Regularized Leader	25
C.1 Computing the Update in Algorithm 2	30
D Technical Results for Section 5.3	30
D.1 Proof of Theorem 2.5	30
E Comparison with Gökcesu and Kozat [2022]	34
F Details on the Experiments	35

A Technical Results for Section 4

In Section 4, we discussed lower bounds for the lower-order terms. These lower bounds are achieved by adopting specific binary random variables in the construction of the losses. This is shown in Lemma 4.1 and its proof is provided in what follows.

A.1 Proof of Lemma 4.1

Proof of Lemma 4.1. The fact that $\mathbb{E}[X_j^2] = 1$ follows by an immediate calculation after observing that each X_j is a binary random variable. By independence, we also have that $\mathbb{E}[\max_{j \in [n]} X_j] = \sqrt{n}(1 - (1 - \frac{1}{n})^n)$. The result now follows from observing that $1/2 < 1 - (1 - \frac{1}{n})^n \leq 1$, where we used the inequality $1 - x \leq \exp(-x)$. \square

A.2 Remarks on Fréchet Distributions

Here, we now discuss another family of random variables that guarantee similar properties. This family is that of Fréchet distributions [Fréchet, 1927]—see also Muraleedharan et al. [2011] for reference.

We denote by $\text{Fréchet}(\alpha, s, m)$ the Fréchet distribution whose parameters are the shape parameter $\alpha > 0$, the scale parameter $s > 0$, and the location (of the minimum) parameter $m \in \mathbb{R}$. Its cumulative distribution function (CDF) is

$$\mathbb{P}(X \leq x) = \exp\left(-\left(\frac{x-m}{s}\right)^{-\alpha}\right)$$

for $x > m$, where $X \sim \text{Fréchet}(\alpha, s, m)$. If $\alpha > 1$, its expected value corresponds to $\mathbb{E}[X] = m + s\Gamma(1 - 1/\alpha)$, where $\Gamma(z) = \int_0^\infty t^{z-1} \exp(-t) dt$ for $z > 0$ is the Gamma function.

Lemma A.1. *Fix any $n \in \mathbb{N} \setminus \{0\}$. Let $X_1, \dots, X_n \sim \text{Fréchet}(\alpha, s, 0)$ be i.i.d. random variables with $\alpha > 1$ and $s > 0$. Then, for any $1 \leq \beta < \alpha$ and any $j \in [n]$,*

$$\mathbb{E}[X_j^\beta] = s^\beta \Gamma\left(1 - \frac{\beta}{\alpha}\right).$$

Furthermore,

$$\mathbb{E}[\max\{X_1, \dots, X_n\}] = n^{1/\alpha} s \Gamma\left(1 - \frac{1}{\alpha}\right).$$

Proof. All results are standard. Here we provide explicit calculations for completeness. First, observe that for any $X \sim \text{Fr\'echet}(\alpha, s, 0)$ we have that

$$\mathbb{P}(X^\beta \leq x) = \mathbb{P}\left(X \leq x^{1/\beta}\right) = \exp\left(-\left(\frac{x^{1/\beta}}{s}\right)^{-\alpha}\right) = \exp\left(-\left(\frac{x}{s^\beta}\right)^{-\alpha/\beta}\right),$$

showing that $X^\beta \sim \text{Fr\'echet}(\alpha/\beta, s^\beta, 0)$. As a consequence, $\mathbb{E}[X^\beta] = s^\beta \Gamma(1 - \beta/\alpha)$ since $\alpha/\beta > 1$. Second, let $Y = \max\{X_1, \dots, X_n\}$ and $y > 0$. By independence, one has

$$\begin{aligned}\mathbb{P}(Y \leq y) &= \mathbb{P}\left(\bigcap_{j \in [n]} \{X_j \leq y\}\right) = \prod_{j \in [n]} \mathbb{P}(X_j \leq y) \\ &= \prod_{j \in [n]} \exp\left(-\left(y/s\right)^{-\alpha}\right) = \exp\left(-\left(\frac{y}{n^{1/\alpha}s}\right)^{-\alpha}\right),\end{aligned}$$

which means that $Y \sim \text{Fr\'echet}(\alpha, n^{1/\alpha}s, 0)$. Since $\alpha > 1$, its expectation is thus $\mathbb{E}[Y] = n^{1/\alpha}s\Gamma(1 - 1/\alpha)$. \square

Keeping the results from the lemma above in mind, we can consider a more specific case where the Fr\'echet random variables have scale $s = 1$ and $\alpha > 2$. Then, for $n = KT$ and $\beta = 2$, we clearly have that their second moment equals $\Gamma(1 - 2/\alpha)$, whereas the expectation of their maximum corresponds to $(KT)^{1/\alpha}\Gamma(1 - 1/\alpha)$. This shows that the i.i.d. Fr\'echet random variables as described above behave somewhat similarly to the random variables of Lemma 4.1 we employed in our lower bounds.

Moving back to the original construction from Lemma 4.1, we can adopt it to prove the main lower bound on any regret guarantee containing the M_T term (or a constant fraction of it). This result from Section 4 is stated in Proposition 4.3 and here we provide its proof.

A.3 Proof of Proposition 4.3

Proof. Let $\ell_t(i) = \mu_{t,i} + \xi_{t,i}\varepsilon_{t,i}$, where $\mu_{t,i} \in [-1, 1]$, $\xi_{t,i}$ is sampled from a Rademacher distribution, and $\varepsilon_{t,i}$ is sampled from the distribution specified in Lemma 4.1 with $n = KT$. Let $\sum_{t=1}^T \mu_{t,K} = 0$ and let $\sum_{t=1}^T \mu_{t,i} > 0$ for $i \neq K$. By assumption, we have that

$$R_T = \mathbb{E} \left[\sum_{t=1}^T \left(\sum_{i=1}^K p_t(i) \ell_t(i) - \ell_t(K) \right) \right] = \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^K p_t(i) \ell_t(i) \right] \leq B_T^\star + \mathbb{E} \left[\max_{t \in [T], i \in [K]} |r_t(i)| \right].$$

Observe that

$$\begin{aligned}
\mathbb{E} \left[\max_{t \in [T], i \in [K]} |r_t(i)| \right] &= \mathbb{E} \left[\max_{t \in [T], i \in [K]} |\ell_t(i) - \langle p_t, \ell_t \rangle| \right] \\
&\geq \mathbb{E} \left[\max_{t \in [T], i \in [K]} (|\ell_t(i) - \ell_t(K)| - |\ell_t(K) - \langle p_t, \ell_t \rangle|) \right] \\
&\quad \text{(reverse triangle inequality)} \\
&= \mathbb{E} \left[\max_{t \in [T], i \in [K]} (|\ell_t(i) - \ell_t(K)| - |r_t(K)|) \right] \\
&\geq \mathbb{E} \left[\max_{t \in [T], i \in [K]} |\ell_t(i) - \ell_t(K)| \right] - \mathbb{E} \left[\max_{t \in [T], i \in [K]} |r_t(i)| \right] \\
&\geq \mathbb{E} \left[\max_{t \in [T], i \in [K]} (|\ell_t(i)| - |\ell_t(K)|) \right] - \mathbb{E} \left[\max_{t \in [T], i \in [K]} |r_t(i)| \right] \\
&\quad \text{(reverse triangle inequality)} \\
&\geq \mathbb{E} \left[\max_{t \in [T], i \in [K]} |\ell_t(i)| \right] - \mathbb{E} \left[\max_{t \in [T]} |\ell_t(K)| \right] - \mathbb{E} \left[\max_{t \in [T], i \in [K]} |r_t(i)| \right] \\
&\geq \mathbb{E} \left[\max_{t \in [T], i \in [K]} |\varepsilon_{t,i}| \right] - \mathbb{E} \left[\max_{t \in [T]} |\varepsilon_{t,K}| \right] - \mathbb{E} \left[\max_{t \in [T], i \in [K]} |r_t(i)| \right] - 2 .
\end{aligned}$$

By carefully using the properties of the random variables $\varepsilon_{t,i}$ from Lemma 4.1, we have that

$$\mathbb{E} \left[\max_{t \in [T], i \in [K]} |\varepsilon_{t,i}| \right] \geq \left(1 - \frac{1}{e}\right) \sqrt{KT}$$

as well as

$$\mathbb{E} \left[\max_{t \in [T]} |\varepsilon_{t,j}| \right] = \sqrt{KT} \left(1 - \left(1 - \frac{1}{KT}\right)^T \right) \leq \sqrt{\frac{T}{K}}$$

for any fixed $j \in [K]$, where the last inequality follows from the fact that $(1 - x)^n \geq 1 - nx$ for all $x \leq 1$ and all $n \geq 1$. Therefore, by rearranging the terms in the previous inequality we obtain

$$\mathbb{E} \left[\max_{t \in [T], i \in [K]} |r_t(i)| \right] \geq \frac{\sqrt{T}}{2} \left(\left(1 - \frac{1}{e}\right) \sqrt{K} - \frac{1}{\sqrt{K}} \right) - 1 \geq \frac{\sqrt{KT}}{16} - 1 ,$$

where the last inequality holds for any $K \geq 2$. Hence, we immediately have that

$$B_T^* + \mathbb{E} \left[\max_{t \in [T], i \in [K]} |r_t(i)| \right] \geq B_T^* + \frac{1}{16} \sqrt{KT} - 1 .$$

Finally, we have

$$\mathbb{E}_t[\ell_t(i)^2] = \mu_{t,i}^2 + 2\mu_{t,i} \mathbb{E}_t[\varepsilon_{t,i}] + \mathbb{E}_t[\varepsilon_{t,i}^2] \leq 2 + 2\mathbb{E}_t[\varepsilon_{t,i}] = 2 + \frac{2}{\sqrt{KT}} \leq 3$$

since $K, T \geq 2$, which concludes the proof. \square

A.4 Discussion on analyses of existing algorithms

Our results from Section 4 show that the lower-order terms are a problematic component of the regret bounds of existing algorithms. In particular, we have shown that the lower-order terms can be larger than the ideal $\sqrt{\theta T \log K}$ in the case of heavy-tailed losses with bounded second moments. It is important to note that our general argument does not depend on the specific algorithm we consider, but rather on regret bounds involving terms such as L_T or M_T . Hence, generally speaking, one may believe that the presence of L_T or M_T in the regret guarantees could be an artefact of a loose analysis. However, here we show that this is not the case for most existing algorithms.

To illustrate, let us consider a simplified view that captures the essential mechanism. The algorithms in Table 1 are effectively sophisticated modifications of the exponential weights (EW) algorithm designed to adapt to $H_T = \max_{t,i} |\ell_t(i)|$. Their sampling distributions are of the form $p_t(i) \propto$

$\exp(-\eta_t \ell_t(i))$, where we assume a fixed learning $\eta > 0$ for simplicity. A crucial aspect of these adaptive algorithms is choosing a learning rate related to H_T , often resembling $\eta = \min\{\alpha, \frac{1}{H_T}\}$, where $\alpha > 0$ represents some cap and where we assume that H_T is known. Their choice ensures the control of terms like $\log(\mathbb{E}_{i \sim p_t} [\exp(-\eta \ell_t(i))])$, which is an essential component in any regret analysis of EW-based algorithms.

Now, we construct a sequence of losses for which the regret of such an algorithm is always at least $\mathbb{E}[\log(4/3)/\eta]$. Suppose $K \geq 5$ and let $\ell_t(i) = 1 + \varepsilon_{t,i}$ for each $i \in [K-1]$, where $\varepsilon_{t,i}$ follows the same distribution as in Lemma 4.1; for action K , we instead choose $\ell_t(K) = 0$. By Jensen's inequality and the definition of p_t , we have that

$$\begin{aligned} R_T &= \mathbb{E} \left[\sum_{t=1}^T \langle p_t, \ell_T \rangle \right] \\ &\geq \mathbb{E} \left[-\frac{1}{\eta} \sum_{t=1}^T \log(\mathbb{E}_{i \sim p_t} [\exp(-\eta \ell_t(i))]) \right] \\ &= \mathbb{E} \left[\frac{1}{\eta} \log(K) - \frac{1}{\eta} \log \left(\sum_{i=1}^K \exp \left(-\eta \sum_{t=1}^T \ell_t(i) \right) \right) \right] \\ &\geq \mathbb{E} \left[\frac{1}{\eta} \log(K) - \frac{1}{\eta} \log(1 + (K-1) \exp(-\eta T)) \right], \end{aligned}$$

where the last inequality is due to the definition of $\ell_t(i)$. It is safe to assume that $\eta T < 1$, as otherwise we would obtain a vacuous upper bound from the standard EW analysis. Therefore, we have that

$$\begin{aligned} R_T &\geq \mathbb{E} \left[\frac{1}{\eta} \log(K) - \frac{1}{\eta} \log(1 + (K-1)/2) \right] \\ &\geq \mathbb{E} \left[\frac{1}{\eta} \log(K) - \frac{1}{\eta} \log(3(K-1)/4) \right] \\ &\geq \mathbb{E} \left[\frac{1}{\eta} \right] \log(4/3) \\ &\geq L_T \log(4/3), \end{aligned}$$

where the last inequality follows from the definitions of η and $L_T = \mathbb{E}[H_T]$. This suggests the regret is lower bounded by a term proportional to the maximum loss range. For algorithms like Squint that run EW on the surrogate losses like $\ell_t(i) - \eta \ell_t(i)^2$, a similar but slightly more involved argument can be made. Consequently, the L_T (or M_T) term appears to be unavoidable in the regret of these algorithms.

Similarly, regarding Section 4.2, the existing algorithms for the squared loss setting are also fundamentally related to EW, adapted to quadratic losses. Their learning rates are typically dependent on the range $\max_t |\bar{z}_t - y_t|$ in a manner analogous to the dependence on H_T described above. Therefore, we believe a similar limitation applies, making it unlikely that a different analysis of previous algorithms could circumvent the dependence on lower-order terms in their regret.

B Regret Analysis for Online Mirror Descent

In Algorithm 1 we only update p_t if $\sum_{s \leq t} \bar{v}_t > 0$. The only case where this is not true is if in all rounds up to and including round t , $r_t(i) = 0$ for all $i \in [K]$, in which case we can simply ignore these rounds in the analysis, which is what we do: throughout this section we assume that $\bar{v}_1 > 0$ without loss of generality.

B.1 Adversarial Environments

We now analyze the regret incurred by Algorithm 1 for adversarial and self-bounded environments. We start by reminding a standard lemma.

Lemma B.1 (3-point identity, [Chen and Teboulle, 1993](#)). *Let $X \subseteq \mathbb{R}^d$ be a non-empty convex set, $\psi : X \rightarrow \mathbb{R}$ be strictly convex and differentiable on $\text{int}(X) \neq \emptyset$, and B_ψ be the associated Bregman divergence. Then for any $x \in X$ and any $y, z \in \text{int}(X)$, the following holds*

$$B_\psi(x, y) = B_\psi(x, z) + B_\psi(z, y) + \langle x - z, \nabla\psi(z) - \nabla\psi(y) \rangle.$$

We then have the following result about the regret of Algorithm 1.

Theorem B.2. *Consider Algorithm 1 run with parameters $\alpha \in (0, 1]$ and $\beta > 0$. Then, with probability one, for any $i^* \in [K]$ we have*

$$\begin{aligned} & \sum_{t=1}^T \langle p_t - e_{i^*}, \ell_t \rangle \\ & \leq \left(\sqrt{\alpha T} + 5\beta + \frac{4 + \log(K/\alpha)}{\beta} \right) \sqrt{\sum_{t=1}^T \bar{v}_t} + \left(\sqrt{\alpha T} + \frac{\log(K/\alpha)}{\beta} + 2\beta \right) \sqrt{\sum_{t=1}^T v_t(i^*)}. \end{aligned}$$

Furthermore, setting $\alpha = \frac{1}{T}$ and $\beta = \sqrt{\log(KT)}$ gives

$$\sum_{t=1}^T \langle p_t - e_{i^*}, \ell_t \rangle = \mathcal{O} \left(\sqrt{\log(KT) \left(\sum_{t=1}^T \bar{v}_t + \sum_{t=1}^T v_t(i^*) \right)} \right).$$

Proof. Let $i^* \in [K]$ be our point-mass comparator, let $\bar{\ell}_t = -r_t$ for any $t \in [T]$, and recall that $\mathcal{P}_\alpha = \{p \in \mathcal{P} : p \geq \frac{\alpha}{K} \mathbf{1}\}$ is the truncated probability simplex, for any $\alpha > 0$. Define $\text{trunc}_T(x) = \sum_{t=1}^T \langle x, \bar{\ell}_t - \tilde{\ell}_t \rangle$ to denote the cost of truncating the losses for any $x \in \mathbb{R}^K$ and, by convention, let $\text{trunc}_T(p_{1:T}) = \sum_{t=1}^T \langle p_t, \bar{\ell}_t - \tilde{\ell}_t \rangle$ where $p_{1:T} = (p_1, \dots, p_T)$. Moreover, define $\tilde{R}_T(q) = \sum_{t=1}^T \langle p_t - q, \tilde{\ell}_t \rangle$ for any $q \in \mathcal{P}$. We consider the following regret decomposition:

$$\begin{aligned} \sum_{t=1}^T \langle p_t - e_{i^*}, \ell_t \rangle &= \sum_{t=1}^T \langle p_t - e_{i^*}, \ell_t - \tilde{\ell}_t \rangle + \sum_{t=1}^T \langle p_t - e_{i^*}, \tilde{\ell}_t \rangle \\ &= \underbrace{\sum_{t=1}^T \langle p_t - e_{i^*}, \bar{\ell}_t - \tilde{\ell}_t \rangle}_{-\text{trunc}_T(p_{1:T}) + \text{trunc}_T(-e_{i^*})} + \underbrace{\sum_{t=1}^T \langle p_t - e_{i^*}, \tilde{\ell}_t \rangle}_{-\tilde{R}_T(e_{i^*})}, \end{aligned}$$

where we added a constant vector $\langle p_t, \ell_t \rangle \mathbf{1} = \ell_t - \bar{\ell}_t$ to ℓ_t at each term of the first sum in the second equality because, for any t , both p_t and e_{i^*} are probability distributions and thus $\langle p_t - e_{i^*}, c \cdot \mathbf{1} \rangle = c - c = 0$ for any $c \in \mathbb{R}$. The first term is just the cost of truncating the losses, and the second is just the regret of OMD on the truncated losses. We start with the latter.

Step 1: control the regret of OMD. Fix any $t \in [T]$. Note that by the first-order optimality condition of p_{t+1} we have, for any $u \in \mathcal{P}_\alpha$,

$$\langle \tilde{\ell}_t + \nabla\psi_t(p_{t+1}) - \nabla\psi_t(p_t), u - p_{t+1} \rangle \geq 0.$$

Following a simple sequence of derivations also referred to as the 3-point identity (see Lemma B.1), the inequality above can be rewritten as

$$\langle \tilde{\ell}_t, u - p_{t+1} \rangle - (D_t(u \| p_{t+1}) + D_t(p_{t+1} \| p_t) - D_t(u \| p_t)) \geq 0. \quad (6)$$

This inequality cannot be applied to $\tilde{R}_T(e_{i^*})$ since e_{i^*} does not belong to \mathcal{P}_α . Instead, we use a mixture of e_{i^*} and the uniform distribution, $u_{i^*} = (1 - \alpha) e_{i^*} + \frac{\alpha}{K} \mathbf{1} \in \mathcal{P}_\alpha$, and notice that by

definition of u_{i^*} and a triangle inequality

$$\begin{aligned}
\tilde{R}_T(e_{i^*}) - \tilde{R}_T(u_{i^*}) &= \sum_{t=1}^T \left\langle u_{i^*} - e_{i^*}, \tilde{\ell}_t \right\rangle \\
&= \alpha \left\langle \frac{1}{K} \mathbb{1} - e_{i^*}, \sum_{t=1}^T \tilde{\ell}_t \right\rangle \quad (\text{definition of } u_{i^*}) \\
&\leq \alpha \sum_{i=1}^K \frac{1}{K} \sum_{t=1}^T \left| \tilde{\ell}_t(i) \right| + \alpha \sum_{t=1}^T \left| \tilde{\ell}_t(i^*) \right|.
\end{aligned}$$

Next, following up with Cauchy-Schwarz inequality,

$$\begin{aligned}
\tilde{R}_T(e_{i^*}) - \tilde{R}_T(u_{i^*}) &\leq \alpha \sum_{i=1}^K \frac{1}{K} \sqrt{T \sum_{t=1}^T \tilde{\ell}_t(i)^2} + \alpha \sqrt{T \sum_{t=1}^T \tilde{\ell}_t(i^*)^2} \\
&\leq \alpha \sqrt{T \sum_{t=1}^T \sum_{i=1}^K \frac{1}{K} \tilde{\ell}_t(i)^2} + \alpha \sqrt{T \sum_{t=1}^T \tilde{\ell}_t(i^*)^2} \quad (\text{Jensen's inequality}) \\
&\leq \sqrt{\alpha T \sum_{t=1}^T \sum_{i=1}^K p_t(i) \tilde{\ell}_t(i)^2} + \alpha \sqrt{T \sum_{t=1}^T \tilde{\ell}_t(i^*)^2} \quad (p_t \in \mathcal{P}_\alpha) \\
&\leq \sqrt{\alpha T} \left(\sqrt{\sum_{t=1}^T \bar{v}_t} + \sqrt{\sum_{t=1}^T v_t(i^*)} \right). \quad (\text{Cauchy-Schwarz, } \tilde{\ell}_t^2 \leq \bar{\ell}_t^2)
\end{aligned}$$

Therefore, we can focus on bounding $\tilde{R}_T(u_{i^*})$. From Equation (6), we have

$$\tilde{R}_T(u_{i^*}) \leq \sum_{t=1}^T \left(\left\langle p_t - p_{t+1}, \tilde{\ell}_t \right\rangle - D_t(p_{t+1} \| p_t) + D_t(u_{i^*} \| p_t) - D_t(u_{i^*} \| p_{t+1}) \right).$$

The first difference within the sum can be analyzed with local norms arguments [Orabona, 2025, Section 6.5]. However, we do not have control over the range of the losses, thus we cannot use the standard arguments and instead resort to specific learning rates to account for this. On the other hand, the second difference is almost a telescoping sum and requires a more careful analysis for the same reason. We start with the first. Denote $\tilde{p}_{t+1} = \arg \min_{p \in \mathbb{R}_{\geq 0}^K} \left\{ \left\langle p, \tilde{\ell}_t \right\rangle + D_t(p \| p_t) \right\}$ the minimizer of the unconstrained optimization problem, and note that for any $i \in [K]$, $\tilde{p}_{t+1}(i) = p_t(i) \exp(-\eta_{t,i} \tilde{\ell}_t(i)) \leq 3p_t(i)$, where the inequality holds by definition of the losses which satisfy $|\tilde{\ell}_t(i)| \leq \frac{1}{\eta_{t,i}}$. Since the function ψ_t is twice differentiable on $\mathbb{R}_{>0}^K$, by Taylor's theorem there exists $z_t \in [p_t, \tilde{p}_{t+1}]$ such that $D_t(\tilde{p}_{t+1} \| p_t) = \frac{1}{2} \|\tilde{p}_{t+1} - p_t\|_{\nabla^2 \psi_t(z_t)}^2$. Therefore, the first difference in the inequality above becomes, for any $t \in [T]$,

$$\begin{aligned}
\left\langle p_t - p_{t+1}, \tilde{\ell}_t \right\rangle - D_t(p_{t+1} \| p_t) &\leq \left\langle p_t - \tilde{p}_{t+1}, \tilde{\ell}_t \right\rangle - D_t(\tilde{p}_{t+1} \| p_t) \quad (\text{optimality of } \tilde{p}_{t+1}) \\
&= \left\langle p_t - \tilde{p}_{t+1}, \tilde{\ell}_t \right\rangle - \frac{1}{2} \|\tilde{p}_{t+1} - p_t\|_{\nabla^2 \psi_t(z_t)}^2 \quad (\text{Taylor's theorem}) \\
&\leq \max_{x \in \mathbb{R}^K} \left\{ \left\langle x, \tilde{\ell}_t \right\rangle - \frac{1}{2} \|x\|_{\nabla^2 \psi_t(z_t)}^2 \right\} \\
&= \frac{1}{2} \|\tilde{\ell}_t\|_{(\nabla^2 \psi_t(z_t))^{-1}}^2,
\end{aligned}$$

where the last step follows from the Fenchel-Young inequality applied to the convex function $\frac{1}{2} \|\cdot\|_{\nabla^2 \psi_t(z_t)}$. Noting that $\nabla^2 \psi_t(z_t) = \text{Diag}(\eta_t \odot z_t)^{-1}$, we further have

$$\begin{aligned}
\langle p_t - p_{t+1}, \tilde{\ell}_t \rangle - D_t(p_{t+1} \| p_t) &\leq \frac{1}{2} \sum_{i=1}^K \eta_{t,i} z_t(i) \tilde{\ell}_t(i)^2 \\
&\leq \frac{3}{2} \sum_{i=1}^K \eta_{t,i} p_t(i) \tilde{\ell}_t(i)^2 \quad (z_t \leq 3p_t) \\
&= \frac{3\beta}{2} \sum_{i=1}^K \frac{p_t(i) \tilde{\ell}_t(i)^2}{\sqrt{\max \left\{ \sum_{s=1}^t \bar{v}_s, \sum_{s=1}^t v_s(i) \right\}}} \quad (\text{definition of } \eta_t) \\
&\leq \frac{3\beta}{2} \cdot \frac{\bar{v}_t}{\sqrt{\sum_{s=1}^t \bar{v}_s}} \quad (\tilde{\ell}_t^2 \leq \bar{\ell}_t^2).
\end{aligned}$$

Summing up for all $t \in [T]$, we obtain

$$\sum_{t=1}^T \langle p_t - p_{t+1}, \tilde{\ell}_t \rangle - D_t(p_{t+1} \| p_t) \leq \frac{3\beta}{2} \sum_{t=1}^T \frac{\bar{v}_t}{\sqrt{\sum_{s=1}^t \bar{v}_s}} \leq 3\beta \sqrt{\sum_{t=1}^T \bar{v}_t},$$

where the inequality follows from, *e.g.*, [Orabona \[2025, Lemma 4.13\]](#). Moving on to the second difference, we have

$$\begin{aligned}
&\sum_{t=1}^T (D_t(u_{i^*} \| p_t) - D_t(u_{i^*} \| p_{t+1})) \\
&= D_T(u_{i^*} \| p_T) - D_T(u_{i^*} \| p_{T+1}) + \sum_{t=1}^{T-1} (D_t(u_{i^*} \| p_t) - D_t(u_{i^*} \| p_{t+1})) \\
&= D_1(u_{i^*} \| p_1) - D_T(u_{i^*} \| p_{T+1}) + \sum_{t=2}^T (D_t(u_{i^*} \| p_t) - D_{t-1}(u_{i^*} \| p_t)) \\
&\leq D_1(u_{i^*} \| p_1) + \sum_{t=2}^T (D_t(u_{i^*} \| p_t) - D_{t-1}(u_{i^*} \| p_t)),
\end{aligned}$$

where the second equality is due to a telescopic sum, and the last inequality is because $D_T(u_{i^*} \| p_{T+1}) \geq 0$. The sum above is given by

$$\begin{aligned}
&\sum_{t=2}^T (D_t(u_{i^*} \| p_t) - D_{t-1}(u_{i^*} \| p_t)) \\
&= \sum_{t=2}^T \sum_{i=1}^K \left(\frac{1}{\eta_{t,i}} - \frac{1}{\eta_{t-1,i}} \right) \left(u_{i^*}(i) \log \left(\frac{u_{i^*}(i)}{p_t(i)} \right) - u_{i^*}(i) + p_t(i) \right) \\
&\leq \sum_{t=2}^T \sum_{i=1}^K \left(\frac{1}{\eta_{t,i}} - \frac{1}{\eta_{t-1,i}} \right) \left(u_{i^*}(i) \log \left(\frac{u_{i^*}(i)}{p_t(i)} \right) + p_t(i) \right),
\end{aligned}$$

where the inequality is due to having $\eta_{t,i} \leq \eta_{t-1,i}$ and $u_{i^*}(i) \geq 0$ for any i and any t . Since $p_t \in \mathcal{P}_\alpha$, we have for any $i \neq i^*$ that $\frac{u_{i^*}(i)}{p_t(i)} \leq 1$, i.e., $\log\left(\frac{u_{i^*}(i)}{p_t(i)}\right) \leq 0$. Thus,

$$\begin{aligned} & \sum_{t=2}^T (D_t(u_{i^*} \| p_t) - D_{t-1}(u_{i^*} \| p_t)) \\ & \leq \sum_{t=2}^T \left(\frac{1}{\eta_{t,i^*}} - \frac{1}{\eta_{t-1,i^*}} \right) u_{i^*}(i^*) \log\left(\frac{u_{i^*}(i^*)}{p_t(i^*)}\right) + \sum_{t=2}^T \sum_{i=1}^K \left(\frac{1}{\eta_{t,i}} - \frac{1}{\eta_{t-1,i}} \right) p_t(i) \\ & \leq \frac{\log(K/\alpha)}{\eta_{T,i^*}} + \sum_{t=1}^T \sum_{i=1}^K \left(\frac{1}{\eta_{t,i}} - \frac{1}{\eta_{t-1,i}} \right) p_t(i), \end{aligned}$$

where we used $u_{i^*}(i^*) \leq 1$, $p_t(i^*) \geq \frac{\alpha}{K}$, and $\eta_{1,i^*} > 0$ for the first sum and added the non-negative term for $t = 1$ in the second sum. For the remaining sum, notice that for any t, i , we have $\frac{1}{\eta_{t,i}} - \frac{1}{\eta_{t-1,i}} = \eta_{t,i} \left(\frac{1}{\eta_{t,i}^2} - \frac{1}{\eta_{t,i} \eta_{t-1,i}} \right) \leq \eta_{t,i} \left(\frac{1}{\eta_{t,i}^2} - \frac{1}{\eta_{t-1,i}^2} \right)$ due to $\eta_{t,i} \leq \eta_{t-1,i}$. Using the definition of the learning rates, we get

$$\begin{aligned} \frac{1}{\eta_{t,i}} - \frac{1}{\eta_{t-1,i}} & \leq \eta_{t,i} \left(\frac{1}{\eta_{t,i}^2} - \frac{1}{\eta_{t-1,i}^2} \right) \\ & = \frac{\max\left\{\sum_{s \leq t} \bar{v}_s, \sum_{s \leq t} v_s(i)\right\} - \max\left\{\sum_{s < t} \bar{v}_s, \sum_{s < t} v_s(i)\right\}}{\beta \sqrt{\max\left\{\sum_{s \leq t} \bar{v}_s, \sum_{s \leq t} v_s(i)\right\}}} \\ & \leq \frac{\max\{\bar{v}_t, v_t(i)\}}{\beta \sqrt{\max\left\{\sum_{s \leq t} \bar{v}_s, \sum_{s \leq t} v_s(i)\right\}}} \\ & \leq \frac{\bar{v}_t + v_t(i)}{\beta \sqrt{\sum_{s \leq t} \bar{v}_s}}. \end{aligned} \tag{7}$$

Therefore, adding everything up

$$\begin{aligned} \sum_{t=1}^T \sum_{i=1}^K \left(\frac{1}{\eta_{t,i}} - \frac{1}{\eta_{t-1,i}} \right) p_t(i) & \leq \sum_{t=1}^T \sum_{i=1}^K \frac{\bar{v}_t + v_t(i)}{\beta \sqrt{\sum_{s \leq t} \bar{v}_s}} p_t(i) \\ & = \frac{2}{\beta} \sum_{t=1}^T \frac{\bar{v}_t}{\sqrt{\sum_{s \leq t} \bar{v}_s}} \quad (p_t \in \mathcal{P}, \text{ definition } \bar{v}_t) \\ & \leq \frac{4}{\beta} \sqrt{\sum_{t=1}^T \bar{v}_t}, \end{aligned}$$

where the last inequality follows again from [Orabona \[2025, Lemma 4.13\]](#). Putting it back into the previous inequality, we get

$$\begin{aligned} \sum_{t=2}^T (D_t(u_{i^*} \| p_t) - D_{t-1}(u_{i^*} \| p_t)) & \leq \frac{\log(K/\alpha)}{\eta_{T,i^*}} + \frac{4}{\beta} \sqrt{\sum_{t=1}^T \bar{v}_t} \\ & \leq \frac{\log(K/\alpha)}{\beta} \left(\sqrt{\sum_{t=1}^T v_t(i^*)} + \sqrt{\sum_{t=1}^T \bar{v}_t} \right) + \frac{4}{\beta} \sqrt{\sum_{t=1}^T \bar{v}_t}. \end{aligned}$$

Finally, the regret incurred by OMD on the truncated losses is bounded by

$$\tilde{R}_T(e_{i^*}) \leq \left(\sqrt{\alpha T} + 3\beta + \frac{4 + \log(K/\alpha)}{\beta} \right) \sqrt{\sum_{t=1}^T \bar{v}_t} + \left(\sqrt{\alpha T} + \frac{\log(K/\alpha)}{\beta} \right) \sqrt{\sum_{t=1}^T v_t(i^*)}.$$

Step 2: control the cost of truncation. We have

$$\begin{aligned}
\text{trunc}_T(p_{1:T}) &= \sum_{t=1}^T \sum_{i=1}^K p_t(i) \left(\bar{\ell}_t(i) - \tilde{\ell}_t(i) \right) \\
&= \sum_{t=1}^T \sum_{i=1}^K p_t(i) \bar{\ell}_t(i) \mathbb{1}(|\bar{\ell}_t(i)| > \eta_{t,i}^{-1}) \quad (\text{definition of } \tilde{\ell}_t) \\
&\leq \sum_{t=1}^T \sum_{i=1}^K p_t(i) \frac{\bar{\ell}_t(i)^2}{|\bar{\ell}_t(i)|} \mathbb{1}(|\bar{\ell}_t(i)| > \eta_{t,i}^{-1}) \quad (p_t \geq 0) \\
&\leq \sum_{t=1}^T \sum_{i=1}^K p_t(i) \eta_{t,i} \bar{\ell}_t(i)^2.
\end{aligned}$$

Plugging the definition of the learning rates η_t , we obtain

$$\begin{aligned}
\text{trunc}_T(p_{1:T}) &\leq \beta \sum_{t=1}^T \sum_{i=1}^K p_t(i) \frac{\bar{\ell}_t(i)^2}{\sqrt{\max\left\{\sum_{s=1}^t \bar{v}_s, \sum_{s=1}^t v_s(i)\right\}}} \quad (\text{definition of } \eta_t) \\
&\leq \beta \sum_{t=1}^T \sum_{i=1}^K \frac{p_t(i) \bar{\ell}_t(i)^2}{\sqrt{\sum_{s=1}^t \bar{v}_s}} \\
&= \beta \sum_{t=1}^T \frac{\bar{v}_t}{\sqrt{\sum_{s=1}^t \bar{v}_s}} \\
&\leq 2\beta \sqrt{\sum_{t=1}^T \bar{v}_t},
\end{aligned}$$

where the last inequality follows from [Orabona \[2025, Lemma 4.13\]](#). Likewise,

$$\text{trunc}_T(-e_{i^\star}) \leq 2\beta \sqrt{\sum_{t=1}^T v_t(i^\star)}.$$

Overall, our regret is bounded by

$$\begin{aligned}
&\sum_{t=1}^T \langle p_t - e_{i^\star}, \ell_t \rangle \\
&\leq \left(\sqrt{\alpha T} + 5\beta + \frac{4 + \log(K/\alpha)}{\beta} \right) \sqrt{\sum_{t=1}^T \bar{v}_t} + \left(\sqrt{\alpha T} + \frac{\log(K/\alpha)}{\beta} + 2\beta \right) \sqrt{\sum_{t=1}^T v_t(i^\star)}. \quad \square
\end{aligned}$$

[Theorem B.2](#) is a more general result and it is indeed stronger than what we originally stated in [Theorem 2.4](#). In particular, we are able to show that the former result implies the latter under [Assumption 2.1](#). This is illustrated by the following corollary.

Corollary B.3. *Suppose [Assumption 2.1](#) holds. Then, [Algorithm 1](#) with $\alpha = \frac{1}{T}$ and $\beta = \sqrt{\log(KT)}$ guarantees*

$$R_T = \mathcal{O}\left(\sqrt{\theta T \log(KT)}\right).$$

Proof. Recall that, by [Theorem B.2](#), [Algorithm 1](#) with $\alpha = \frac{1}{T}$ and $\beta = \sqrt{\log(K)}$ already guarantees

$$\sum_{t=1}^T \langle p_t - e_{i^\star}, \ell_t \rangle = \mathcal{O}\left(\sqrt{\log(KT) \left(\sum_{t=1}^T \bar{v}_t + \sum_{t=1}^T v_t(i^\star) \right)}\right) \quad (9)$$

for any sequence of losses. Now, focus on the \bar{v}_t and $v_t(i)$ terms. First, for any $i \in [K]$, we can observe that $v_t(i)$ satisfies

$$\begin{aligned}
\mathbb{E}_t[v_t(i)] &= \mathbb{E}_t \left[\left(\ell_t(i) - \sum_{j=1}^K p_t(j) \ell_t(j) \right)^2 \right] && \text{(definition of } v_t(i) \text{)} \\
&= \mathbb{E}_t \left[\left(\sum_{j=1}^K p_t(j) (\ell_t(i) - \ell_t(j)) \right)^2 \right] \\
&\leq \mathbb{E}_t \left[\sum_{j=1}^K p_t(j) (\ell_t(i) - \ell_t(j))^2 \right] && \text{(Jensen's inequality)} \\
&\leq 2 \mathbb{E}_t[\ell_t(i)^2] + 2 \mathbb{E}_t \left[\sum_{j=1}^K p_t(j) \ell_t(j)^2 \right] && \text{(using } (a-b)^2 \leq 2a^2 + 2b^2 \text{ and } p_t \in \mathcal{P} \text{)} \\
&= 2 \mathbb{E}_t[\ell_t(i)^2] + 2 \sum_{j=1}^K p_t(j) \mathbb{E}_t [\ell_t(j)^2] \\
&\leq 4\theta ,
\end{aligned}$$

where the last inequality follows by Assumption 2.1 and the fact that $p_t \in \mathcal{P}$. Then, we can move to \bar{v}_t and notice that, by its definition, it satisfies

$$\begin{aligned}
\mathbb{E}_t[\bar{v}_t] &= \mathbb{E}_t \left[\sum_{i=1}^K p_t(i) v_t(i) \right] \\
&= \mathbb{E}_t \left[\sum_{i=1}^K p_t(i) \left(\ell_t(i) - \sum_{j=1}^K p_t(j) \ell_t(j) \right)^2 \right] && \text{(definition of } v_t(i) \text{)} \\
&\leq \mathbb{E}_t \left[\sum_{i=1}^K p_t(i) \ell_t(i)^2 \right] && \text{(using } \mathbb{V}[X] \leq \mathbb{E}[X^2] \text{)} \\
&= \sum_{i=1}^K p_t(i) \mathbb{E}_t [\ell_t(i)^2] \\
&\leq \theta ,
\end{aligned}$$

where the last inequality holds by both Assumption 2.1 and the fact that $p_t \in \mathcal{P}$.

At this point, we can observe that $R_T = \max_{i^* \in [K]} \mathbb{E}[\langle p_t - e_{i^*}, \ell_t \rangle]$, i.e., it corresponds to the expected value of the left-hand side of Equation (9). We can then focus on the expected value of its right-hand side (ignoring constant factors) and, by applying Jensen's inequality with respect to the square root and the tower rule of expectation, we infer that

$$\begin{aligned}
\mathbb{E} \left[\sqrt{\log(KT) \sum_{t=1}^T (\bar{v}_t + v_t(i^*))} \right] &\leq \sqrt{\log(KT) \mathbb{E} \left[\sum_{t=1}^T (\mathbb{E}_t[\bar{v}_t] + \mathbb{E}_t[v_t(i^*)]) \right]} \\
&\leq \sqrt{5\theta T \log(KT)} .
\end{aligned}$$

This concludes the proof. \square

B.2 Self-Bounded Environments

In this section, we provide a regret bound for self-bounded environments defined in [Zimmert and Seldin \[2021\]](#).

Theorem B.4. Let Assumption 2.1, 2.3 hold, and consider Algorithm 1 run with parameters $\alpha \in (0, 1]$, $\beta > 0$. We denote $C_3(K, T) = 2\sqrt{\alpha T} + 7\beta + \frac{4+2\log(K/\alpha)}{\beta}$. Then, we have

$$\mathbb{E}[R_T] \leq \begin{cases} \frac{16C_3(K, T)^2\theta}{\Delta_{\min}} + \frac{8C_3(K, T)\sqrt{\theta C}}{\sqrt{\Delta_{\min}}} & \text{if } C \leq \frac{16C_3(K, T)^2\theta}{\Delta_{\min}} \\ \frac{8C_3(K, T)\sqrt{\theta C}}{\sqrt{\Delta_{\min}}} & \text{otherwise} \end{cases}.$$

Furthermore, setting $\alpha = \frac{1}{T}$ and $\beta = \sqrt{\log(KT)}$ gives $C_3(K, T) = \mathcal{O}\left(\sqrt{\log(KT)}\right)$ and

$$\mathbb{E}[R_T] \leq \begin{cases} \frac{3600\log(KT)\theta}{\Delta_{\min}} + 120\sqrt{\frac{\log(KT)\theta C}{\Delta_{\min}}} & \text{if } C \leq \frac{16C_3(K, T)^2\theta}{\Delta_{\min}} \\ 120\sqrt{\frac{\log(KT)\theta C}{\Delta_{\min}}} & \text{otherwise} \end{cases}.$$

Proof. From Jensen's inequality applied to the regret bound proven for adversarial environments in Theorem B.2, we have

$$\begin{aligned} \mathbb{E}[R_T(e_{i^*})] &\leq C_1(K, T) \sqrt{\sum_{t=1}^T \sum_{i=1}^K \mathbb{E} \left[p_t(i) \left(\ell_t(i) - \sum_{j=1}^K p_t(j) \ell_t(j) \right)^2 \right]} \\ &\quad + C_2(K, T) \sqrt{\sum_{t=1}^T \mathbb{E} \left[\left(\ell_t(i^*) - \sum_{j=1}^K p_t(j) \ell_t(j) \right)^2 \right]}, \end{aligned}$$

where $C_1(K, T) = \sqrt{\alpha T} + 5\beta + \frac{4+\log(K/\alpha)}{\beta}$, and $C_2(K, T) = \sqrt{\alpha T} + \frac{\log(K/\alpha)}{\beta} + 2\beta$. Observe that for any $t \in [T]$, the expectation in the first term is a variance that can be bounded by the second moment ($\mathbb{V}[X] \leq \mathbb{E}[X^2]$ for any random variable X)

$$\begin{aligned} &\mathbb{E}_t \left[\sum_{i=1}^K p_t(i) \left(\ell_t(i) - \sum_{j=1}^K p_t(j) \ell_t(j) \right)^2 \right] \\ &= \mathbb{E}_t \left[\sum_{i=1}^K p_t(i) \left((\ell_t(i) - \ell_t(i^*)) - \sum_{j=1}^K p_t(j) (\ell_t(j) - \ell_t(i^*)) \right)^2 \right] \\ &\leq \mathbb{E}_t \left[\sum_{i \neq i^*} p_t(i) (\ell_t(i) - \ell_t(i^*))^2 \right]. \end{aligned}$$

Further, using the inequality $(a+b)^2 \leq 2(a^2 + b^2)$ for any $a, b \in \mathbb{R}$,

$$\begin{aligned} \mathbb{E}_t \left[\sum_{i=1}^K p_t(i) \left(\ell_t(i) - \sum_{j=1}^K p_t(j) \ell_t(j) \right)^2 \right] &\leq 2\mathbb{E}_t \left[\sum_{i \neq i^*} p_t(i) (\ell_t(i)^2 + \ell_t(i^*)^2) \right] \\ &\leq 4\theta(1 - p_t(i^*)), \end{aligned}$$

where the last inequality is by Assumption 2.1 and $\langle p_t, \mathbb{1} \rangle = 1$. Likewise, for the second term we have by Jensen's inequality

$$\begin{aligned} \mathbb{E}_t \left[\left(\ell_t(i^*) - \sum_{j=1}^K p_t(j) \ell_t(j) \right)^2 \right] &\leq \mathbb{E}_t \left[\sum_{j \neq i^*} p_t(j) (\ell_t(i^*) - \ell_t(j))^2 \right] \\ &\leq 2\mathbb{E}_t \left[\sum_{j \neq i^*} p_t(j) (\ell_t(i^*)^2 + \ell_t(j)^2) \right] \\ &\leq 4\theta(1 - p_t(i^*)), \end{aligned}$$

where the last inequality is again by Assumption 2.1 and $\langle p_t, \mathbf{1} \rangle = 1$. Denote $\Delta_{\min} = \min_{i \neq i^*} \Delta_i$, and $C_3(K, T) = C_1(K, T) + C_2(K, T)$. Combining the above and Assumption 2.3, we find that for any $\lambda \in (0, 1]$,

$$\begin{aligned}\mathbb{E}[R_T] &= (1 + \lambda) \mathbb{E}[R_T] - \lambda \mathbb{E}[R_T] \\ &\leq (1 + \lambda) C_3(K, T) \sqrt{4\theta \mathbb{E} \left[\sum_{t=1}^T (1 - p_t(i^*)) \right]} - \lambda \Delta_{\min} \mathbb{E} \left[\sum_{t=1}^T (1 - p_t(i^*)) \right] + \lambda C \\ &= \sqrt{4(1 + \lambda)^2 C_3(K, T)^2 \theta \mathbb{E} \left[\sum_{t=1}^T (1 - p_t(i^*)) \right]} - \lambda \Delta_{\min} \mathbb{E} \left[\sum_{t=1}^T (1 - p_t(i^*)) \right] + \lambda C.\end{aligned}$$

Using the inequality of arithmetic and geometric means, *i.e.* $|ab| = \inf_{\gamma > 0} \frac{a^2}{2\gamma} + \frac{\gamma b^2}{2}$,

$$\begin{aligned}\mathbb{E}[R_T] &= \inf_{\gamma > 0} \left\{ \frac{4(1 + \lambda)^2 C_3(K, T)^2 \theta}{\gamma} + \gamma \mathbb{E} \left[\sum_{t=1}^T (1 - p_t(i^*)) \right] \right\} \\ &\quad - \lambda \Delta_{\min} \mathbb{E} \left[\sum_{t=1}^T (1 - p_t(i^*)) \right] + \lambda C \\ &\leq \frac{4(1 + \lambda)^2 C_3(K, T)^2 \theta}{\lambda \Delta_{\min}} + \lambda C \\ &\leq \frac{16 C_3(K, T)^2 \theta}{\lambda \Delta_{\min}} + \lambda C,\end{aligned}$$

where the first inequality is by taking $\gamma = \lambda \Delta_{\min}$, and the second is by $1 + \lambda \leq 2$. In particular, setting $\lambda = \min \left\{ 1, \frac{4C_3(K, T)\sqrt{\theta}}{\sqrt{\Delta_{\min}C}} \right\}$ gives

$$\begin{aligned}\mathbb{E}[R_T] &\leq \frac{16 C_3(K, T)^2 \theta}{\Delta_{\min}} \max \left\{ 1, \frac{\sqrt{\Delta_{\min}C}}{4C_3(K, T)\sqrt{\theta}} \right\} + \min \left\{ 1, \frac{4C_3(K, T)\sqrt{\theta}}{\sqrt{\Delta_{\min}C}} \right\} C \\ &\leq \frac{16 C_3(K, T)^2 \theta}{\Delta_{\min}} + \frac{8 C_3(K, T)\sqrt{\theta C}}{\sqrt{\Delta_{\min}}}.\end{aligned}$$

If $C > \frac{16 C_3(K, T)^2 \theta}{\Delta_{\min}}$, then $\lambda = \frac{4C_3(K, T)\sqrt{\theta}}{\sqrt{\Delta_{\min}C}}$ and the previous inequality can be improved to

$$\mathbb{E}[R_T] \leq \frac{8 C_3(K, T)\sqrt{\theta C}}{\sqrt{\Delta_{\min}}}.$$

Recall that $C_3(K, T) = 2\sqrt{\alpha T} + 7\beta + \frac{4+2\log(K/\alpha)}{\beta}$. Setting $\alpha = \frac{1}{T}$ and $\beta = \sqrt{\log(KT)}$ gives $C_3(K, T) \leq 6 + 9\sqrt{\log(KT)} \leq 15\sqrt{\log(KT)}$ since $\log(KT) \geq 1$. Finally, we obtain

$$\mathbb{E}[R_T] \leq \begin{cases} \frac{3600 \log(KT)\theta}{\Delta_{\min}} + 120\sqrt{\frac{\log(KT)\theta C}{\Delta_{\min}}} & \text{if } C \leq \frac{16 C_3(K, T)^2 \theta}{\Delta_{\min}} \\ 120\sqrt{\frac{\log(KT)\theta C}{\Delta_{\min}}} & \text{otherwise} \end{cases}.$$

In the special case of a stochastic environment, $C = 0$ thus we set $\lambda = 1$ and obtain $\mathbb{E}[R_T] = \mathcal{O}\left(\frac{\theta \log(KT)}{\Delta_{\min}}\right)$. \square

C Regret Analysis for Follow The Regularized Leader

In Algorithm 2 we only update the prediction p_t if $\sum_{s \leq t} v_s > 0$. The only case where this is not true is if in all rounds up to and including round t , $r_t(i) = 0$ for all $i \in [K]$, meaning that the cumulative regret up to round t is null. In this case, we can simply ignore these rounds in the regret analysis, which is what we do: throughout this section we assume that $v_1 > 0$ without loss of generality.

Theorem C.1. Consider Algorithm 2 with parameter $\beta > 0$, providing predictions $p_1, \dots, p_T \in \mathcal{P}$ over a sequence of losses ℓ_1, \dots, ℓ_T . Then, with probability one, for any $i^* \in [K]$ we have

$$\begin{aligned} & \sum_{t=1}^T \langle p_t - e_{i^*}, \ell_t \rangle \\ & \leq \left(\frac{\log(K)}{\beta} + 2\beta \right) \sqrt{\sum_{t=1}^T v_t(i^*)} + \left(\frac{5 + \log(K)}{\beta} + 5\beta \right) \sqrt{\sum_{t=1}^T \bar{v}_t} + \frac{1}{\beta} \sum_{i=1}^K \frac{1}{K} \sqrt{\sum_{t=1}^T v_t(i)}. \end{aligned}$$

Furthermore, setting $\beta = \sqrt{\log(K)}$ gives

$$\sum_{t=1}^T \langle p_t - e_{i^*}, \ell_t \rangle = \mathcal{O} \left(\sqrt{\log(K) \left(\sum_{t=1}^T \bar{v}_t + \sum_{t=1}^T v_t(i^*) \right)} + \frac{1}{K} \sum_{i=1}^K \sqrt{\sum_{t=1}^T v_t(i)} \right).$$

Proof. Let $i^* \in [K]$ be our point-mass comparator and let $\bar{\ell}_t = -r_t$ for any $t \in [T]$. First, observe that the regret can be equivalently rewritten as

$$\sum_{t=1}^T \langle p_t - e_{i^*}, \ell_t \rangle = \sum_{t=1}^T \langle p_t - e_{i^*}, \bar{\ell}_t \rangle;$$

this follows from the fact that replacing ℓ_t with $\bar{\ell}_t$ only leads to a difference $\langle p_t - e_{i^*}, \ell_t - \bar{\ell}_t \rangle = \langle p_t - e_{i^*}, c \cdot \mathbf{1} \rangle = \langle p_t, c \cdot \mathbf{1} \rangle - \langle e_{i^*}, c \cdot \mathbf{1} \rangle = c - c = 0$ for the constant $c = \langle p_t, \ell_t \rangle$, by definition of $\bar{\ell}_t$. As in the proof of Theorem B.2, we take the same definitions of trunc_T and \tilde{R}_T , and consider the following regret decomposition:

$$\sum_{t=1}^T \langle p_t - e_{i^*}, \ell_t \rangle = \underbrace{\sum_{t=1}^T \langle p_t - e_{i^*}, \bar{\ell}_t - \tilde{\ell}_t \rangle}_{=\text{trunc}_T(p_{1:T}) + \text{trunc}_T(-e_{i^*})} + \underbrace{\sum_{t=1}^T \langle p_t - e_{i^*}, \tilde{\ell}_t \rangle}_{=\tilde{R}_T(e_{i^*})},$$

where the first term is the cost of truncating the losses, and the second is the regret of FTRL on the truncated losses $\tilde{\ell}_1, \dots, \tilde{\ell}_T$. Let us first focus on the former term.

Before we proceed, one final remark is in order. Recall that Algorithm 2 performs predictions defined such that

$$p_t = \arg \min_{p \in \mathcal{P}} \sum_{s=1}^{t-1} \langle p, \tilde{\ell}_s \rangle + D_{t-1}(p \| p_0)$$

for any $t > 1$, where $p_0 = \frac{1}{K} \mathbf{1} \in \mathcal{P}$ is the uniform distribution, while $p_1 = p_0$. Observe that we can define $\eta_{0,i} = \beta / \sqrt{\bar{v}_1}$ (never used nor set by Algorithm 2) for any $i \in [K]$ and, thus, we can equivalently denote p_1 as

$$p_1 \in \arg \min_{p \in \mathcal{P}} D_0(p \| p_0)$$

because $D_0(p \| p_0) = \frac{\sqrt{\bar{v}_1}}{\beta} \sum_{i=1}^K (p(i) \log(p(i)/p_0(i)) - p(i) + p_0(i))$ is minimized at p_0 . We remark that, while this step appears to require knowledge of \bar{v}_1 before computing the prediction p_1 , it is only part of the analysis and not algorithmically performed.

Step 1: control the cost of truncation. We can begin by focusing on $\text{trunc}_T(p_{1:T})$. Observe that

$$\begin{aligned} \text{trunc}_T(p_{1:T}) &= \sum_{t=1}^T \langle p_t, \bar{\ell}_t - \tilde{\ell}_t \rangle \\ &\leq \sum_{t=1}^T \sum_{i=1}^K p_t(i) |\bar{\ell}_t(i) - \tilde{\ell}_t(i)| \\ &= \sum_{t=1}^T \sum_{i=1}^K p_t(i) \left(|\bar{\ell}_t(i)| \left(1 - \frac{b_{t-1}(i)}{b_t(i)} \right) \mathbf{1} \left(|\bar{\ell}_t(i)| \leq \frac{1}{\eta_{t,i}} \right) + |\bar{\ell}_t(i)| \mathbf{1} \left(|\bar{\ell}_t(i)| > \frac{1}{\eta_{t,i}} \right) \right), \end{aligned}$$

where the last equality follows by definition of $\tilde{\ell}_t(i)$, after observing that $b_{t-1}(i) \leq b_t(i)$. Now, by using the definitions of $b_t(i)$ and $\eta_{t,i}$, observe that

$$\begin{aligned}
& \sum_{i=1}^K p_t(i) |\bar{\ell}_t(i)| \left(1 - \frac{b_{t-1}(i)}{b_t(i)}\right) \mathbb{1}(|\bar{\ell}_t(i)| \leq \frac{1}{\eta_{t,i}}) \\
&= \sum_{i=1}^K p_t(i) \frac{|\bar{\ell}_t(i)|}{b_t(i)} (b_t(i) - b_{t-1}(i)) \mathbb{1}(|\bar{\ell}_t(i)| \leq \frac{1}{\eta_{t,i}}) \\
&= \sum_{i=1}^K p_t(i) \eta_{t,i} |\bar{\ell}_t(i)| \left(\frac{1}{\eta_{t,i}} - \frac{1}{\eta_{t-1,i}}\right) \mathbb{1}(|\bar{\ell}_t(i)| \leq \frac{1}{\eta_{t,i}}) \\
&\leq \sum_{i=1}^K p_t(i) \left(\frac{1}{\eta_{t,i}} - \frac{1}{\eta_{t-1,i}}\right) \\
&\leq \frac{2\bar{v}_t}{\beta \sqrt{\sum_{s \leq t} \bar{v}_s}} ,
\end{aligned}$$

for $t > 1$, where the last inequality holds by Equation (8) given that the learning rates $\eta_{t,i}$ have the same definition; the same bound holds similarly for $t = 1$ by observing that $b_0(i) = 0$ and, hence, we have

$$\begin{aligned}
\sum_{i=1}^K p_1(i) |\bar{\ell}_1(i)| \mathbb{1}(|\bar{\ell}_1(i)| \leq \frac{1}{\eta_{1,i}}) &\leq \frac{1}{\beta} \sum_{i=1}^K p_1(i) \sqrt{\max\{\bar{v}_1, v_1(i)\}} \\
&\leq \frac{1}{\beta} \sqrt{\sum_{i=1}^K p_1(i) \max\{\bar{v}_1, v_1(i)\}} \leq \frac{\sqrt{2\bar{v}_1}}{\beta} \leq \frac{2\bar{v}_1}{\beta \sqrt{v_1}} ,
\end{aligned}$$

where the second step follows by Jensen's inequality. At the same time, we have that

$$\begin{aligned}
\sum_{i=1}^K p_t(i) |\bar{\ell}_t(i)| \mathbb{1}(|\bar{\ell}_t(i)| > \frac{1}{\eta_{t,i}}) &= \sum_{i=1}^K p_t(i) \frac{\bar{\ell}_t(i)^2}{|\bar{\ell}_t(i)|} \mathbb{1}(|\bar{\ell}_t(i)| > \frac{1}{\eta_{t,i}}) \\
&\leq \sum_{i=1}^K \eta_{t,i} p_t(i) \bar{\ell}_t(i)^2 \\
&\leq \frac{\beta \bar{v}_t}{\sqrt{\sum_{s \leq t} \bar{v}_s}} .
\end{aligned}$$

We can therefore combine the above inequalities and, together with [Orabona \[2025, Lemma 4.13\]](#), obtain that

$$\mathbf{trunc}_T(p_{1:T}) \leq \left(\beta + \frac{2}{\beta}\right) \sum_{t=1}^T \frac{\bar{v}_t}{\sqrt{\sum_{s \leq t} \bar{v}_s}} \leq 2 \left(\beta + \frac{2}{\beta}\right) \sqrt{\sum_{t=1}^T \bar{v}_t} .$$

Similarly, we can see that $\mathbf{trunc}_T(-e_{i^*})$ similarly satisfies

$$\begin{aligned}
\mathbf{trunc}_T(-e_{i^*}) &= \sum_{t=1}^T \left\langle -e_{i^*}, \bar{\ell}_t - \tilde{\ell}_t \right\rangle \\
&\leq \sum_{t=1}^T |\bar{\ell}_t(i^*) - \tilde{\ell}_t(i^*)| \\
&= \sum_{t=1}^T \left(|\bar{\ell}_t(i^*)| \left(1 - \frac{b_{t-1}(i^*)}{b_t(i^*)}\right) \mathbb{1}(|\bar{\ell}_t(i^*)| \leq \frac{1}{\eta_{t,i^*}}) + |\bar{\ell}_t(i^*)| \mathbb{1}(|\bar{\ell}_t(i^*)| > \frac{1}{\eta_{t,i^*}}) \right) .
\end{aligned}$$

Using similar calculations as before, it follows that

$$\begin{aligned} \sum_{t=1}^T |\bar{\ell}_t(i^*)| \left(1 - \frac{b_{t-1}(i^*)}{b_t(i^*)}\right) \mathbb{1}\left(|\bar{\ell}_t(i^*)| \leq \frac{1}{\eta_{t,i^*}}\right) &\leq \frac{1}{\beta} \sum_{t=1}^T (b_t(i^*) - b_{t-1}(i^*)) = \frac{1}{\eta_{T,i^*}} \\ &\leq \frac{1}{\beta} \sqrt{\sum_{t=1}^T \bar{v}_t} + \frac{1}{\beta} \sqrt{\sum_{t=1}^T v_t(i^*)}, \end{aligned}$$

where we used the subadditivity of the square root in the last inequality, and also that

$$\sum_{t=1}^T |\bar{\ell}_t(i^*)| \mathbb{1}\left(|\bar{\ell}_t(i^*)| > \frac{1}{\eta_{t,i^*}}\right) \leq \sum_{t=1}^T \eta_{t,i^*} \bar{\ell}_t(i^*)^2 \leq \beta \sum_{t=1}^T \frac{v_t(i^*)}{\sqrt{\sum_{s \leq t} v_s(i^*)}} \leq 2\beta \sqrt{\sum_{t=1}^T v_t(i^*)},$$

where the last inequality follows again by [Orabona \[2025, Lemma 4.13\]](#). Hence, we similarly conclude that

$$\text{trunc}_T(-e_{i^*}) \leq \frac{1}{\beta} \sqrt{\sum_{t=1}^T v_t} + \left(2\beta + \frac{1}{\beta}\right) \sqrt{\sum_{t=1}^T v_t(i^*)},$$

which implies that the total cost for truncating the losses is bounded from above as

$$\text{trunc}_T(p_{1:T}) + \text{trunc}_T(-e_{i^*}) \leq \left(2\beta + \frac{5}{\beta}\right) \sqrt{\sum_{t=1}^T \bar{v}_t} + \left(2\beta + \frac{1}{\beta}\right) \sqrt{\sum_{t=1}^T v_t(i^*)}.$$

Step 2: control the regret of FTRL. Let us now focus on the latter term in the regret decomposition, that is, the regret of FTRL on the losses $\tilde{\ell}_1, \dots, \tilde{\ell}_T$. Consider any $t \in [T]$ and define

$$\tilde{p}_{t+1} = \arg \min_{p \in \mathbb{R}^K} \left\{ \langle p, \tilde{\ell}_t \rangle + D_{t-1}(p \| p_t) \right\}.$$

Note that for any $i \in [K]$, $\tilde{p}_{t+1}(i) = p_t(i) \exp(-\eta_{t-1,i} \tilde{\ell}_t(i)) \leq 3p_t(i)$ by construction of $\tilde{\ell}_t$, which is such that

$$\left| \eta_{t-1,i} \tilde{\ell}_t(i) \right| = \eta_{t-1,i} |\bar{\ell}_t(i)| \frac{b_{t-1}(i)}{b_t(i)} \mathbb{1}\left(|\bar{\ell}_t(i)| \leq \frac{1}{\eta_{t,i}}\right) = \eta_{t,i} |\bar{\ell}_t(i)| \mathbb{1}\left(|\bar{\ell}_t(i)| \leq \frac{1}{\eta_{t,i}}\right) \leq 1$$

for any $i \in [K]$ and $t > 1$, while it immediately holds for $t = 1$ since $b_0(i) = 0$ for any i .

Let $\varphi_t = D_{t-1}(\cdot \| p_0)$ be the regularizer used in the FTRL update. Observe that φ_t is non-negative and twice-differentiable with Hessian having inverse $(\nabla^2 \varphi_t(x))^{-1} = \text{Diag}(\eta_{t-1} \odot x)$, and that $\varphi_{t+1}(x) \geq \varphi_t(x)$ for all $x \in \mathbb{R}_{\geq 0}^K$. Then, by standard results on FTRL with time-varying regularizers $\varphi_1, \dots, \varphi_{T+1}$ (e.g., see [Orabona \[2025, Lemma 7.16\]](#)), we obtain

$$\tilde{R}_T(e_{i^*}) \leq \varphi_{T+1}(e_{i^*}) + \frac{1}{2} \sum_{t=1}^T \left\| \tilde{\ell}_t \right\|_{(\nabla^2 \varphi_t(z_t))^{-1}}^2 \tag{10}$$

for some point z_t on the line segment between p_t and \tilde{p}_{t+1} ; this also follows by the monotonicity of φ_t , i.e., $\varphi_t(p_{t+1}) \leq \varphi_{t+1}(p_{t+1})$. The point z_t is such that $z_t(i) \leq \max\{p_t(i), \tilde{p}_{t+1}(i)\} \leq 3p_t(i)$ for all $i \in [K]$. Therefore, given the definition of the local norm, the second term of Equation (10)

satisfies

$$\begin{aligned}
\frac{1}{2} \sum_{t=1}^T \left\| \tilde{\ell}_t \right\|_{(\nabla^2 \varphi_t(z_t))^{-1}}^2 &= \frac{1}{2} \sum_{t=1}^T \sum_{i=1}^K \eta_{t-1,i} z_t(i) \tilde{\ell}_t(i)^2 \\
&\leq \frac{3}{2} \sum_{t=1}^T \sum_{i=1}^K \eta_{t-1,i} p_t(i) \tilde{\ell}_t(i)^2 && \text{(using } z_t(i) \leq 3p_t(i) \text{)} \\
&\leq \frac{3}{2} \sum_{t=1}^T \sum_{i=1}^K \eta_{t-1,i} p_t(i) \bar{\ell}_t(i)^2 \cdot \frac{b_{t-1}(i)^2}{b_t(i)^2} && \text{(using } |\tilde{\ell}_t(i)| \leq \frac{b_{t-1}(i)}{b_t(i)} |\bar{\ell}_t(i)| \text{)} \\
&= \frac{3\beta}{2} \sum_{t=1}^T \sum_{i=1}^K p_t(i) \bar{\ell}_t(i)^2 \cdot \frac{b_{t-1}(i)}{b_t(i)^2} && \text{(definition of } \eta_{t-1,i} \text{)} \\
&\leq \frac{3\beta}{2} \sum_{t=1}^T \sum_{i=1}^K \frac{p_t(i) \bar{\ell}_t(i)^2}{b_t(i)} && \text{(using } b_{t-1}(i) \leq b_t(i) \text{)} \\
&\leq \frac{3\beta}{2} \sum_{t=1}^T \frac{\bar{v}_t}{\sqrt{\sum_{s \leq t} \bar{v}_s}} && \text{(definition of } b_t(i) \text{ and } \bar{v}_t \text{)} \\
&\leq 3\beta \sqrt{\sum_{t=1}^T \bar{v}_t} ,
\end{aligned}$$

where the last inequality follows again by [Orabona \[2025, Lemma 4.13\]](#). On the other hand, the first term of Equation (10) is such that

$$\begin{aligned}
\varphi_{T+1}(e_{i^*}) &= \frac{\log(K) - 1}{\eta_{T,i^*}} + \frac{1}{K} \sum_{i=1}^K \frac{1}{\eta_{T,i}} \\
&\leq \frac{\log(K) - 1}{\beta} \left(\sqrt{\sum_{t=1}^T v_t(i^*)} + \sqrt{\sum_{t=1}^T \bar{v}_t} \right) + \frac{1}{\beta} \sqrt{\sum_{t=1}^T \bar{v}_t} + \frac{1}{\beta} \sum_{i=1}^K \frac{1}{K} \sqrt{\sum_{t=1}^T v_t(i)} \\
&\leq \frac{\log(K) - 1}{\beta} \sqrt{\sum_{t=1}^T v_t(i^*)} + \frac{\log(K)}{\beta} \sqrt{\sum_{t=1}^T \bar{v}_t} + \frac{1}{\beta} \sum_{i=1}^K \frac{1}{K} \sqrt{\sum_{t=1}^T v_t(i)} .
\end{aligned}$$

Combining all the above results together leads to

$$\begin{aligned}
&\sum_{t=1}^T \langle p_t - e_{i^*}, \ell_t \rangle \\
&\leq \left(\frac{\log(K)}{\beta} + 2\beta \right) \sqrt{\sum_{t=1}^T v_t(i^*)} + \left(\frac{5 + \log(K)}{\beta} + 5\beta \right) \sqrt{\sum_{t=1}^T \bar{v}_t} + \frac{1}{\beta} \sum_{i=1}^K \frac{1}{K} \sqrt{\sum_{t=1}^T v_t(i)} . \quad \square
\end{aligned}$$

We can now show that Theorem C.1 suffices to prove one of our main claims from Theorem 2.2. This is demonstrated by the following result.

Corollary C.2. *Suppose Assumption 2.1 holds. Then, Algorithm 2 with $\beta = \sqrt{\log(K)}$ guarantees*

$$R_T = \mathcal{O} \left(\sqrt{\theta T \log(K)} \right) .$$

Proof. Recall that, by Theorem C.1, Algorithm 2 with $\beta = \sqrt{\log(K)}$ guarantees

$$\sum_{t=1}^T \langle p_t - e_{i^*}, \ell_t \rangle = \mathcal{O} \left(\sqrt{\log(K) \left(\sum_{t=1}^T \bar{v}_t + \sum_{t=1}^T v_t(i^*) \right)} + \frac{1}{K} \sum_{i=1}^K \sqrt{\sum_{t=1}^T v_t(i)} \right) \quad (11)$$

for any sequence of losses. Additionally, in a similar way as in the proof of Corollary B.3, we have that $\mathbb{E}_t[\bar{v}_t] \leq \theta$ and that $\mathbb{E}_t[v_t(i)] \leq 4\theta$ for any $i \in [K]$, under Assumption 2.1.

We can analogously observe that $R_T = \max_{i^* \in [K]} \mathbb{E}[\langle p_t - e_{i^*}, \ell_t \rangle]$ or, in other words, that R_T essentially corresponds to the expectation of the left-hand side of Equation (11). Then, we consider the expectation of its right-hand side and, by applying Jensen's inequality with respect to the square root and the tower rule of expectation, we can finally show that

$$\begin{aligned} & \mathbb{E} \left[\sqrt{\log(K) \sum_{t=1}^T (\bar{v}_t + v_t(i^*))} + \frac{1}{K} \sum_{i=1}^K \sqrt{\sum_{t=1}^T v_t(i)} \right] \\ & \leq \sqrt{\log(K) \mathbb{E} \left[\sum_{t=1}^T (\mathbb{E}_t[\bar{v}_t] + \mathbb{E}_t[v_t(i^*)]) \right]} + \frac{1}{K} \sum_{i=1}^K \sqrt{\mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_t[v_t(i)] \right]} \\ & \leq C \sqrt{\theta T \log(K)} \end{aligned}$$

for some constant $C > 0$. This concludes the proof. \square

C.1 Computing the Update in Algorithm 2

We briefly discuss the update defining p_{t+1} in Algorithm 2. For any $\eta \in \mathbb{R}_{>0}^K$, consider an optimization problem of the form $\inf_{p \in \mathcal{P}} \langle p, L \rangle + D(p\|q)$, where $L \in \mathbb{R}^K$, $q \in \mathcal{P}$, and

$$D(p\|q) = \sum_{i=1}^K \frac{1}{\eta_i} \left[p(i) \log \left(\frac{p(i)}{q(i)} \right) + p(i) - q(i) \right].$$

As the probability simplex is compact and the mapping $p \mapsto \langle p, L \rangle + D(p\|q)$ is continuous, the infimum is attained at some $p^* \in \mathcal{P}$. The Lagrangian $\mathcal{L}(p, \lambda)$ of the optimization problem is defined for $p \in \mathbb{R}_{\geq 0}^K$ and $\lambda \in \mathbb{R}$ as

$$\mathcal{L}(p, \lambda) = \langle p, L \rangle + \sum_{i=1}^K \frac{1}{\eta_i} \left[p(i) \log \left(\frac{p(i)}{q(i)} \right) - p(i) + q(i) \right] + \lambda (\langle p, \mathbb{1} \rangle - 1).$$

For any $j \in [K]$, differentiate with respect to $p(j)$ to get

$$\frac{\partial \mathcal{L}}{\partial p(j)}(p, \lambda) = L(j) + \frac{1}{\eta_j} \left[\log \left(\frac{p(j)}{q(j)} \right) + 1 - 1 \right] + \lambda = L(j) + \lambda + \frac{1}{\eta_j} \log \left(\frac{p(j)}{q(j)} \right).$$

Setting it to zero, we get

$$p^*(j) = q(j) \exp(-\eta_j [L(j) + \lambda]).$$

One then wants to find the value of λ by enforcing the constraint on p^* , namely $\langle p^*, \mathbb{1} \rangle = 1$, which gives the following condition:

$$\sum_{i=1}^K q(i) \exp(-\eta_i [L(i) + \lambda]) = 1.$$

If the learning rate did not depend on the coordinate, we could take the term that depends on λ out of the sum to get a closed-form solution. Plugging it back into p^* would give a softmax distribution, but this is not possible here. Instead, one can efficiently compute the normalization constant λ with a line-search.

D Technical Results for Section 5.3

D.1 Proof of Theorem 2.5

Proof. We start with some useful inequalities. Let $a, b, y_t \in \mathbb{R}$ and let $\bar{\ell}_t = -r_t$ for any $t \in [T]$. By the 2-strong convexity of the function $x \mapsto (x - y_t)^2$, we have

$$\left(\frac{a+b}{2} - y_t \right)^2 \leq \frac{1}{2} (a - y_t)^2 + \frac{1}{2} (b - y_t)^2 - \frac{(a-b)^2}{8}. \quad (12)$$

Algorithm 3 LoOT-Free OMD for the Squared Loss

Inputs: Number of experts $K \geq 2$, minimum mass coefficient $\alpha \in (0, 1]$, learning rate coefficient $\beta > 0$.

Initialize: $p_1(i) \leftarrow 1/K$ for all $i \in [K]$.

for $t = 1, \dots, T$ **do**

- Receive $\mathbf{z}_{t,i}$ for all $i \in [K]$.
- Set $\bar{\mathbf{z}}_t \leftarrow \sum_{i=1}^K p_t(i) \mathbf{z}_{t,i}$.
- Observe y_t .
- Set $\ell_t(i) \leftarrow (\mathbf{z}_{t,i}(\bar{\mathbf{z}}_t - y_t) + \frac{1}{2}(\mathbf{z}_{t,i} - y_t)^2)$.
- Set $r_t(i) \leftarrow \langle \ell_t, p_t \rangle - \ell_t(i)$ for all $i \in [K]$.
- Set $v_t(i) \leftarrow r_t(i)^2$ for all $i \in [K]$.
- Set $\bar{v}_t \leftarrow \sum_{i=1}^K p_t(i) v_t(i)$.
- Set $\eta_{t,i} \leftarrow \beta \max \left\{ \sum_{s \leq t} \bar{v}_s, \sum_{s \leq t} v_s(i) \right\}^{-1/2}$ for all $i \in [K]$.
- Set $\tilde{\ell}_t(i) \leftarrow -r_t(i) \mathbb{1}(|r_t(i)| \leq 1/\eta_{t,i})$ for all $i \in [K]$.
- Set $p_{t+1} \leftarrow \arg \min_{p \in \mathcal{P}_\alpha} \langle p, \tilde{\ell}_t \rangle + D_t(p \| p_t)$.

end for

Furthermore, we also have that $(a - y_t)^2 - (b - y_t)^2 = 2(y_t - a)(b - a) - (a - b)^2$, sometimes referred to as a polarization identity (here, for the Euclidean norm on \mathbb{R}). Moving to the regret analysis, let us denote $i^* = \arg \min_{i \in [K]} \mathbb{E} \left[\sum_{t=1}^T (y_t - \mathbf{z}_{t,i})^2 \right]$. Splitting the sum in halves and rearranging the terms using that p_t is a probability distribution over $[K]$, the loss of the learner at time t can be bounded from above as

$$\begin{aligned} (\bar{\mathbf{z}}_t - y_t)^2 &= \left(\left\langle p_t, \frac{1}{2} \mathbf{z}_t \right\rangle + \frac{1}{2} \bar{\mathbf{z}}_t - y_t \right)^2 \\ &= \left(\sum_{i=1}^K p_t(i) \left[\frac{1}{2} \mathbf{z}_{t,i} + \frac{1}{2} \bar{\mathbf{z}}_t - y_t \right] \right)^2 \\ &\leq \sum_{i=1}^K p_t(i) \left[\frac{1}{2} \mathbf{z}_{t,i} + \frac{1}{2} \bar{\mathbf{z}}_t - y_t \right]^2, \end{aligned}$$

where the inequality follows from Jensen's inequality. Applying Equation (12) for any $i \in [K]$ with $a = \mathbf{z}_{t,i}$ and $b = \bar{\mathbf{z}}_t$, and again using $\langle p_t, \mathbb{1} \rangle = 1$, we further get

$$\begin{aligned} (\bar{\mathbf{z}}_t - y_t)^2 &\leq \frac{1}{2} \sum_{i=1}^K p_t(i) (\mathbf{z}_{t,i} - y_t)^2 + \frac{1}{2} (\bar{\mathbf{z}}_t - y_t)^2 - \frac{1}{8} \sum_{i=1}^K p_t(i) (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^2 \\ &= \frac{1}{2} \sum_{i=1}^K p_t(i) (\mathbf{z}_{t,i} - y_t)^2 + \underbrace{\frac{1}{2} (\bar{\mathbf{z}}_t - y_t)^2 - \frac{1}{2} (\mathbf{z}_{t,i^*} - y_t)^2}_{(\diamond)} \\ &\quad - \frac{1}{8} \sum_{i=1}^K p_t(i) (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^2 + (\mathbf{z}_{t,i^*} - y_t)^2 - \frac{1}{2} (\mathbf{z}_{t,i^*} - y_t)^2. \end{aligned}$$

Using the polarization identity on (\diamond) with $a = \bar{\mathbf{z}}_t$ and $b = \mathbf{z}_{t,i^*}$ leads to

$$\begin{aligned} (\bar{\mathbf{z}}_t - y_t)^2 &\leq \frac{1}{2} \sum_{i=1}^K p_t(i) (\mathbf{z}_{t,i} - y_t)^2 + (y_t - \bar{\mathbf{z}}_t) (\mathbf{z}_{t,i^*} - \bar{\mathbf{z}}_t) - \frac{1}{2} (\bar{\mathbf{z}}_t - \mathbf{z}_{t,i^*})^2 \\ &\quad - \frac{1}{8} \sum_{i=1}^K p_t(i) (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^2 + (\mathbf{z}_{t,i^*} - y_t)^2 - \frac{1}{2} (\mathbf{z}_{t,i^*} - y_t)^2. \end{aligned}$$

We recall that Algorithm 3 uses the loss $\ell_t(i) = \frac{1}{2}(\mathbf{z}_{t,i} - y_t)^2 + \mathbf{z}_{t,i}(\bar{\mathbf{z}}_t - y_t)$ for any $i \in [K]$. Rearranging the first two terms in the upper bound, we have

$$\begin{aligned}
(\bar{\mathbf{z}}_t - y_t)^2 &\leq \langle p_t, \ell_t \rangle - \underbrace{\mathbf{z}_{t,i^*}(\bar{\mathbf{z}}_t - y_t)}_{(\clubsuit)} - \frac{1}{2}(\bar{\mathbf{z}}_t - \mathbf{z}_{t,i^*})^2 \\
&\quad - \frac{1}{8} \sum_{i=1}^K p_t(i) (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^2 + (\mathbf{z}_{t,i^*} - y_t)^2 - \underbrace{\frac{1}{2}(\mathbf{z}_{t,i^*} - y_t)^2}_{(\spadesuit)} \\
&= \langle p_t, \ell_t \rangle - \ell_t(i^*) - \frac{1}{2}(\bar{\mathbf{z}}_t - \mathbf{z}_{t,i^*})^2 \quad (\text{using } (\clubsuit) + (\spadesuit) = \ell_t(i^*)) \\
&\quad - \frac{1}{8} \sum_{i=1}^K p_t(i) (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^2 + (\mathbf{z}_{t,i^*} - y_t)^2.
\end{aligned}$$

At this point, we can move the last term in the right-hand side to the left-hand side to obtain

$$(\bar{\mathbf{z}}_t - y_t)^2 - (\mathbf{z}_{t,i^*} - y_t)^2 \leq (\langle p_t, \ell_t \rangle - \ell_t(i^*)) - \frac{1}{2}(\bar{\mathbf{z}}_t - \mathbf{z}_{t,i^*})^2 - \frac{1}{8} \sum_{i=1}^K p_t(i) (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^2. \quad (13)$$

Now, from Theorem B.2 with $\alpha = \frac{1}{T}$ and $\beta = \sqrt{\log(KT)}$ we have

$$\sum_{t=1}^T (\langle p_t, \ell_t \rangle - \ell_t(i^*)) \leq 11\sqrt{\log(KT)} \left(\sqrt{\sum_{t=1}^T \bar{v}_t} + \sqrt{\sum_{t=1}^T v_t(i^*)} \right). \quad (14)$$

We continue by bounding $\bar{v}_t = \sum_{i=1}^K p_t(i) (\ell_t(i) - \langle p_t, \ell_t \rangle)^2$ from above. By definition of the loss ℓ_t , the square inside the sum can be bounded, for any $i \in [K]$, by

$$\begin{aligned}
(\ell_t(i) - \langle p_t, \ell_t \rangle)^2 &= \left(\frac{1}{2} \left[(\mathbf{z}_{t,i} - y_t)^2 - \langle p_t, (\mathbf{z}_t - y_t \cdot \mathbf{1})^2 \rangle \right] + (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)(\bar{\mathbf{z}}_t - y_t) \right)^2 \\
&\leq 2[(\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)(\bar{\mathbf{z}}_t - y_t)]^2 + \frac{1}{2} \left[(\mathbf{z}_{t,i} - y_t)^2 - \langle p_t, (\mathbf{z}_t - y_t \cdot \mathbf{1})^2 \rangle \right]^2,
\end{aligned} \quad (15)$$

where we used the inequality $(a + b)^2 \leq 2a^2 + 2b^2$ for any $a, b \in \mathbb{R}$. Plugging it into the definition of \bar{v}_t , this gives

$$\bar{v}_t \leq 2(\bar{\mathbf{z}}_t - y_t)^2 \sum_{i=1}^K p_t(i) [\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t]^2 + \frac{1}{2} \sum_{i=1}^K p_t(i) \left((\mathbf{z}_{t,i} - y_t)^2 - \sum_{j=1}^K p_t(j) (\mathbf{z}_{t,j} - y_t)^2 \right)^2.$$

We continue by bounding from above the second term on the right-hand side. The difference inside the square can be equivalently rewritten as

$$\begin{aligned}
&(\mathbf{z}_{t,i} - y_t)^2 - \sum_{j=1}^K p_t(j) (\mathbf{z}_{t,j} - y_t)^2 \\
&= [(\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t) - (y_t - \bar{\mathbf{z}}_t)]^2 - \sum_{j=1}^K p_t(j) [(\mathbf{z}_{t,j} - \bar{\mathbf{z}}_t) - (y_t - \bar{\mathbf{z}}_t)]^2 \\
&= (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^2 - 2(y_t - \bar{\mathbf{z}}_t)(\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t) + (y_t - \bar{\mathbf{z}}_t)^2 - \sum_{j=1}^K p_t(j) (\mathbf{z}_{t,j} - \bar{\mathbf{z}}_t)^2 \\
&\quad + 2(y_t - \bar{\mathbf{z}}_t) \sum_{j=1}^K p_t(j) (\mathbf{z}_{t,j} - \bar{\mathbf{z}}_t) - (y_t - \bar{\mathbf{z}}_t)^2 \quad (\langle p_t, \mathbf{1} \rangle = 1) \\
&= (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^2 - 2(y_t - \bar{\mathbf{z}}_t)(\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t) - \sum_{j=1}^K p_t(j) (\mathbf{z}_{t,j} - \bar{\mathbf{z}}_t)^2,
\end{aligned} \quad (16)$$

where the last equality is by definition of $\bar{\mathbf{z}}_t$. We plug it back in the sum above, and using again the inequality $(a + b)^2 \leq 2a^2 + 2b^2$, we get

$$\begin{aligned} & \sum_{i=1}^K p_t(i) \left((\mathbf{z}_{t,i} - y_t)^2 - \sum_{j=1}^K p_t(j) (\mathbf{z}_{t,j} - y_t)^2 \right)^2 \\ & \leq 2 \sum_{i=1}^K p_t(i) (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^4 + 2 \sum_{i=1}^K p_t(i) \left[2(y_t - \bar{\mathbf{z}}_t) (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t) + \sum_{j=1}^K p_t(j) (\mathbf{z}_{t,j} - \bar{\mathbf{z}}_t)^2 \right]^2. \end{aligned}$$

Expanding the square in the second term, we get that the cross-product is equal to zero by definition of $\bar{\mathbf{z}}_t$, thus

$$\begin{aligned} & \sum_{i=1}^K p_t(i) \left((\mathbf{z}_{t,i} - y_t)^2 - \sum_{j=1}^K p_t(j) (\mathbf{z}_{t,j} - y_t)^2 \right)^2 \\ & \leq 2 \sum_{i=1}^K p_t(i) (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^4 + 8(y_t - \bar{\mathbf{z}}_t)^2 \sum_{i=1}^K p_t(i) (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^2 + 2 \left[\sum_{j=1}^K p_t(j) (\mathbf{z}_{t,j} - \bar{\mathbf{z}}_t)^2 \right]^2, \end{aligned}$$

where we used $\langle p_t, \mathbf{1} \rangle = 1$. Using Jensen's inequality,

$$\begin{aligned} & \sum_{i=1}^K p_t(i) \left((\mathbf{z}_{t,i} - y_t)^2 - \sum_{j=1}^K p_t(j) (\mathbf{z}_{t,j} - y_t)^2 \right)^2 \\ & \leq 4 \sum_{i=1}^K p_t(i) (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^4 + 8(y_t - \bar{\mathbf{z}}_t)^2 \sum_{i=1}^K p_t(i) (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^2. \end{aligned}$$

Plugging it back into the upper-bound on \bar{v}_t , we obtain

$$\bar{v}_t \leq 6(\bar{\mathbf{z}}_t - y_t)^2 \sum_{i=1}^K p_t(i) [\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t]^2 + 2 \sum_{i=1}^K p_t(i) (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^4.$$

Using $\mathbb{E}_t[y_t^2] \leq \sigma$ and $|\mathbf{z}_{t,i}| \leq Y$, we find that

$$\mathbb{E}_t[\bar{v}_t] \leq (20Y^2 + 12\sigma) \sum_{i=1}^K p_t(i) (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^2. \quad (17)$$

We now bound $v_t(i)$ for any given $i \in [K]$. As with \bar{v}_t , Equation (15) gives

$$(\ell_t(i) - \langle p_t, \ell_t \rangle)^2 \leq 2(\bar{\mathbf{z}}_t - y_t)^2 (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^2 + \frac{1}{2} \left[(\mathbf{z}_{t,i} - y_t)^2 - \sum_{j=1}^K p_t(j) (\mathbf{z}_{t,j} - y_t)^2 \right]^2.$$

Reusing Equation (16), we can write the second term on the right-hand-side as

$$\begin{aligned} & \left((\mathbf{z}_{t,i} - y_t)^2 - \sum_{j=1}^K p_t(j) (\mathbf{z}_{t,j} - y_t)^2 \right)^2 \\ & = \left((\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^2 - 2(y_t - \bar{\mathbf{z}}_t) (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t) - \sum_{j=1}^K p_t(j) (\mathbf{z}_{t,j} - \bar{\mathbf{z}}_t)^2 \right)^2 \\ & \leq \left[3(\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^2 + 6(y_t - \bar{\mathbf{z}}_t)^2 \right] (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^2 + 3 \left(\sum_{j=1}^K p_t(j) (\mathbf{z}_{t,j} - \bar{\mathbf{z}}_t)^2 \right)^2 \\ & \leq \left[3(\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^2 + 6(y_t - \bar{\mathbf{z}}_t)^2 \right] (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^2 + 3 \left(\max_{k \in [K]} (\mathbf{z}_{t,k} - \bar{\mathbf{z}}_t)^2 \right) \sum_{j=1}^K p_t(j) (\mathbf{z}_{t,j} - \bar{\mathbf{z}}_t)^2, \end{aligned}$$

where the first inequality follows from $(a + b + c)^2 \leq 3a^2 + 3b^2 + 3c^2$ valid for any $a, b, c \in \mathbb{R}$ and the second inequality follows from Jensen's inequality. Thus, using $\mathbb{E}_t [y_t^2] \leq \sigma$ and $|\mathbf{z}_{t,i}| \leq Y$, we find that

$$\mathbb{E}_t [v_t(i)] \leq (16Y^2 + 10\sigma) (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^2 + 6Y^2 \sum_{i=1}^K p_t(i) (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^2. \quad (18)$$

Consider now any $\gamma > 0$. By using the bounds on $\mathbb{E}_t [v_t(i)]$ and $\mathbb{E}_t [\bar{v}_t]$, together with Equations (13) and (14), we finally find that

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=1}^T \left((\bar{\mathbf{z}}_t - y_t)^2 - (\mathbf{z}_{t,i^*} - y_t)^2 \right) \right] \\ & \leq \mathbb{E} \left[11\sqrt{\log(KT)} \left(\sqrt{\sum_{t=1}^T \bar{v}_t} + \sqrt{\sum_{t=1}^T v_t(i^*)} \right) \right] \\ & \quad - \mathbb{E} \left[\sum_{t=1}^T \left(\frac{1}{2} (\mathbf{z}_{t,i^*} - \bar{\mathbf{z}}_t)^2 + \frac{1}{8} \sum_{j=1}^K p_t(j) (\mathbf{z}_{t,j} - \bar{\mathbf{z}}_t)^2 \right) \right] \quad (\text{Equations (13) and (14)}) \\ & \leq \mathbb{E} \left[11\sqrt{2\log(KT)} \left(\sqrt{\sum_{t=1}^T \bar{v}_t} + \sum_{t=1}^T v_t(i^*) \right) \right] \\ & \quad - \mathbb{E} \left[\sum_{t=1}^T \left(\frac{1}{2} (\mathbf{z}_{t,i^*} - \bar{\mathbf{z}}_t)^2 + \frac{1}{8} \sum_{j=1}^K p_t(j) (\mathbf{z}_{t,j} - \bar{\mathbf{z}}_t)^2 \right) \right] \\ & \leq \frac{121\log(KT)}{\gamma} - \mathbb{E} \left[\sum_{t=1}^T \left(\frac{1}{2} (\mathbf{z}_{t,i^*} - \bar{\mathbf{z}}_t)^2 + \frac{1}{8} \sum_{j=1}^K p_t(j) (\mathbf{z}_{t,j} - \bar{\mathbf{z}}_t)^2 \right) \right] \\ & \quad + \gamma \mathbb{E} \left[(8Y^2 + 5\sigma) \sum_{t=1}^T (\mathbf{z}_{t,i^*} - \bar{\mathbf{z}}_t)^2 + (13Y^2 + 6\sigma) \sum_{t=1}^T \sum_{i=1}^K p_t(i) (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^2 \right] \\ & \leq \frac{121\log(KT)}{\gamma} - \mathbb{E} \left[\sum_{t=1}^T \left(\frac{1}{2} (\mathbf{z}_{t,i^*} - \bar{\mathbf{z}}_t)^2 + \frac{1}{8} \sum_{j=1}^K p_t(j) (\mathbf{z}_{t,j} - \bar{\mathbf{z}}_t)^2 \right) \right] \\ & \quad + \gamma (104Y^2 + 48\sigma) \mathbb{E} \left[\frac{1}{2} \sum_{t=1}^T (\mathbf{z}_{t,i^*} - \bar{\mathbf{z}}_t)^2 + \frac{1}{8} \sum_{t=1}^T \sum_{i=1}^K p_t(i) (\mathbf{z}_{t,i} - \bar{\mathbf{z}}_t)^2 \right] \\ & = C(Y^2 + \sigma) \log(KT) \end{aligned}$$

for a sufficiently large constant $C > 0$, where the third inequality follows from Equations (17) and (18) together with the fact that $\sqrt{ab} = \inf_{\gamma > 0} \frac{1}{2\gamma}a + \frac{\gamma}{2}b$ for any $a, b \geq 0$, whereas the last inequality follows by the choice of $\gamma = (104Y^2 + 48\sigma)^{-1}$. This completes the proof. \square

E Comparison with Gökcesu and Kozat [2022]

In this section, we provide further details about the comparison with Gökcesu and Kozat [2022]. The first regret guarantees we compare against are those provided by their Theorem IV.7 and Theorem V.2. One may immediately observe that those regret bounds present an additive term E_T (multiplied by a logarithmic factor) equivalent to our $M_T = \max_{t,i} |r_t(i)|$, for which we already prove that the \sqrt{KT} lower bound holds even with i.i.d. losses.

Hence, the main comparison is mainly with respect to Corollary IV.8 and Corollary V.3 in Gökcesu and Kozat [2022]. The proof of Corollary IV.8 and Corollary V.3 rely on their Lemma IV.1, which requires that, in the notation of Gökcesu and Kozat [2022], $\eta_t |l_{t,m} - \mu_t| \leq 1$ for all t and all m

Figure 1: Results of experiments with heavy-tailed losses.

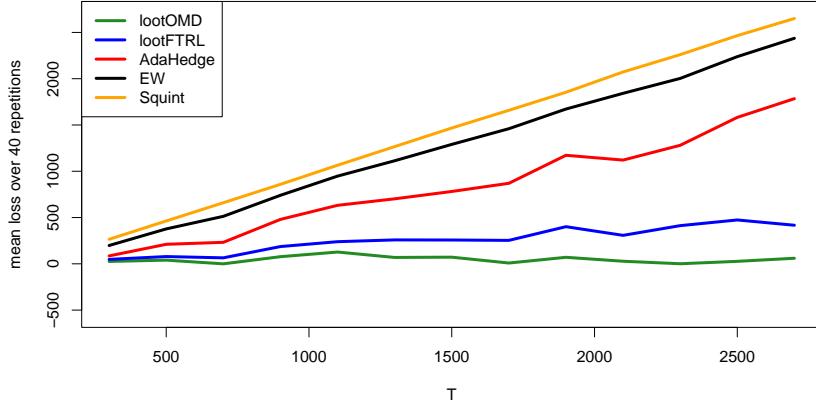
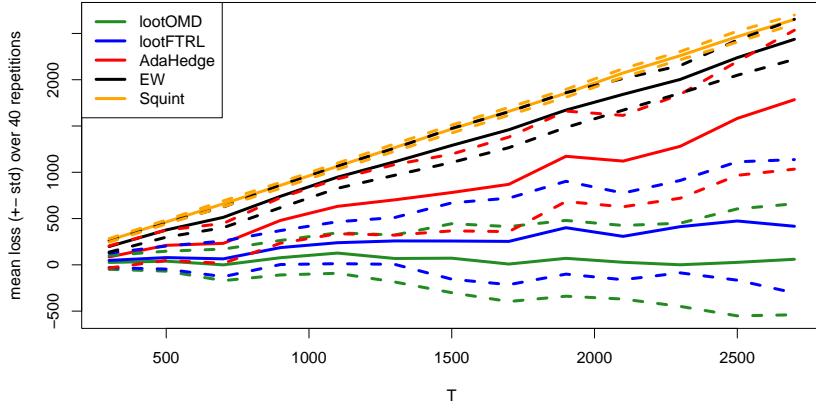


Figure 2: Results of experiments with heavy-tailed losses. Dotted lines represent mean \pm one standard deviation.



for Equation (b) in the proof of Lemma IV.1 to be true. However, with the learning rates given in Corollary IV.8 and Corollary V.3, $\eta_t |l_{t,m} - \mu_t| \leq 1$ does not hold. In Corollary IV.8 they choose $\eta_t \geq \sqrt{\frac{1}{\sum_{s=1}^t \mathbb{E}_{p_s}[(l_{s,m} - \mu_s)^2]}}$, which with $\mu_s = \min_m l_{s,m}$ for $t = 1$ with $l_{1,m} = 0$ if $m \neq K$ and $l_{1,K} = 1$, can be seen to lead to $\eta_t |l_{t,m} - \mu_t| \geq \sqrt{K} > 1$. With Corollary V.3, we run into a similar issue. We do not see a way to fix these issues.

F Details on the Experiments

In this section we provide details on the experiments. All experiments were run on a Macbook Air with 8GB of RAM and an Apple M2 processor. We ran two sets of experiments, and the experiments in each set were run for $K \in \{15, 25, \dots, 135\}$. For each instance we set $T = 20K$.

In the first set of experiments the losses mimic the construction we used in Section 4. The expert losses were equal to $\ell_t(i) = \mathbb{1}[i \neq 1] + \varepsilon_{t,i}$, where $\varepsilon_{t,i} = \zeta_{t,i} X_{t,i}$ with $\zeta_{t,i}$ a Rademacher random variable and

$$X_{t,i} = \begin{cases} 0 & \text{w.p. } 1 - 1/T \\ \sqrt{KT} & \text{w.p. } 1/T \end{cases}.$$

Figure 3: Results of experiments with non-heavy-tailed losses.

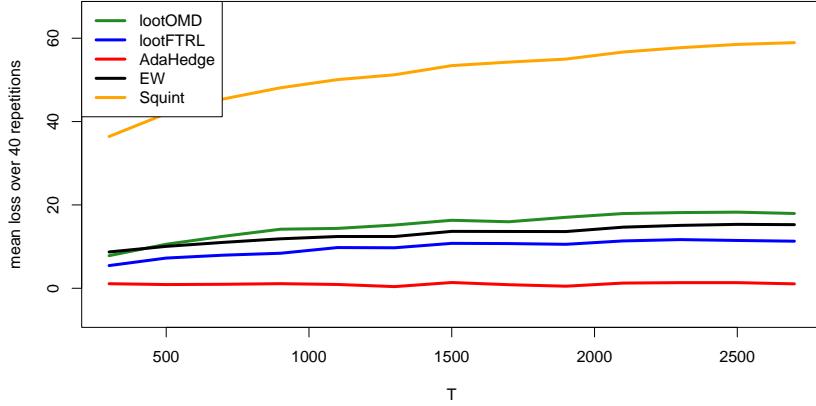
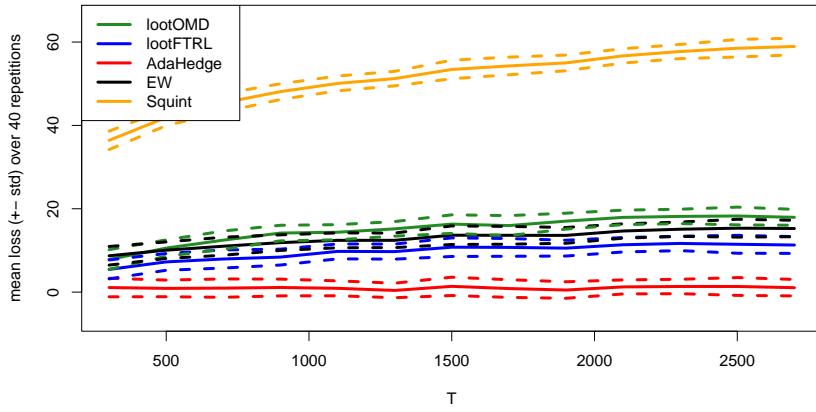


Figure 4: Results of experiments with non-heavy-tailed losses. Dotted lines represent mean \pm one standard deviation.



In the second set of experiments we use a similar construction, where the losses were equal to $\ell_t(i) = \mathbb{1}[i \neq 1] + \varepsilon_{t,i}$, where $\varepsilon_{t,i} = \zeta_{t,i} \tilde{X}_{t,i}$ with $\zeta_{t,i}$ a Rademacher random variable and

$$\tilde{X}_{t,i} = \begin{cases} 0 & \text{w.p. } 1 - 1/T \\ 2 & \text{w.p. } 1/T \end{cases}.$$

The algorithms we implemented were Squint with the improper prior [Koolen and Van Erven, 2015], AdaHedge [De Rooij et al., 2014], and an instance of Exponential Weights (EW) algorithm akin to the algorithm of [Cesa-Bianchi et al., 2007]. Specifically, we ran the FTRL version of exponential weights on the instantaneous regrets with learning rate $\eta_t = \min\left\{\frac{1}{\max_{t,i}|\ell_t(i)|}, \sqrt{\frac{\log(K)}{\sum_{s \leq t} \tilde{v}_t}}\right\}$. We gave this instance of EW the maximum loss, otherwise it would not provide any guarantees. We could have also opted for a doubling trick, but this is known to deteriorate performance. Likewise, we gave Squint the value of $\max_{t,i}|\ell_t(i)|$, as this is a required parameter for Squint. The algorithm of [Mhammedi et al., 2019] could have instead been used to learn $\max_{t,i}|\ell_t(i)|$ online, but seeing that a similar idea as the doubling trick is part of their algorithm, we suspect this would only deteriorate performance.

In the first set of experiments we expect Squint, AdaHedge, and EW to perform poorly due to the issues described in Section 4. In the second set of experiments we expect similar behaviour from all algorithms.

As can be seen from the results in Figure 1, algorithms not tailored to adapt to θ fare considerably worse in the heavy-tailed loss setting we consider. We therefore conclude that the lower-order terms in the regret bounds of these algorithms are not an artefact of the analysis, but rather represent the problematic behaviour of these algorithms in the face of heavy-tailed losses.

The results for the second set of experiments is similarly as expected with one exception: the performance of Squint. Squint fares considerably worse than the other algorithms. However, upon inspection, it seems that the performance of Squint is still below what is predicted by theory. The regret bound of Squint contains a $15 \log(1 + K(2 + \log(T + 1)))$ term, which for the smallest values of K, T is slightly larger than 71.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: See theorems in Section 2 and lower bounds in Section 4.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: See Section 6.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification: Results and assumptions are clearly stated in the main text. Proofs are provided in the appendix.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We clearly describe the experiments and how we chose the parameters of the relevant algorithms in Appendix F.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [\[Yes\]](#)

Justification: The data is simulated. The code to simulate the data is included in the supplementary material. The implementations of the algorithms is included in the supplementary material as well.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [\[Yes\]](#)

Justification: We clearly describe the experiments in Appendix F.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [\[No\]](#)

Justification: The guarantees in the paper are on the average performance, hence we only report the average behaviour.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer “Yes” if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: It is described in Appendix F.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [\[Yes\]](#)

Justification: This work being primarily theoretical, it does not raise any ethical concerns.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [\[NA\]](#)

Justification: The work is mainly theoretical and does not have any direct societal impact.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: This is mainly a theoretical paper.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: The paper does not use any existing assets.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: The experimental details are provided in the appendix.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.